Automated Anomaly Detection Using Auto Encoders

Applied ML 2022

Carl-Johannes Johnsen
carl-johannes@di.ku.dk
Department of Computer Science
University of Copenhagen

What are we trying to solve?

What are we trying to solve?

What are we trying to solve?

Can we build a program that can do ***automatic anomaly detection*** of ***X-Ray images*** of food while keeping up with the ***high-throughput demands*** of food production?

Automated feature extraction

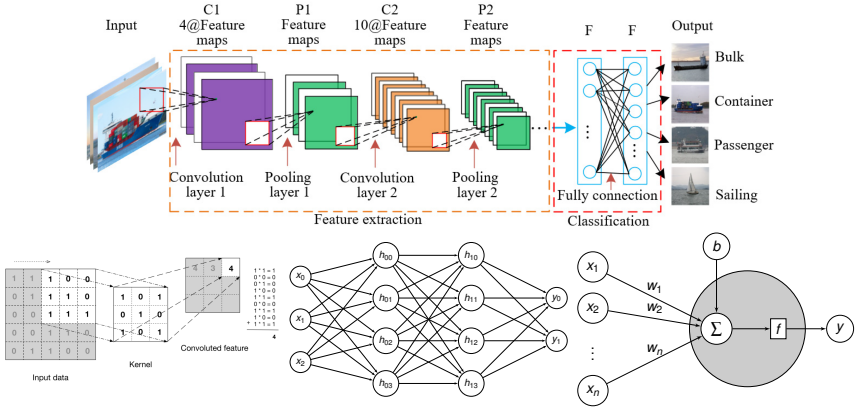

How are the CNN results?

How are the CNN results?

Good results (>90 %)

How are the CNN results?

Good results (>90 %) ... but could be better.

How are the CNN results?

Good results (>90 %) ... but could be better.

The approach is supervised.

How are the CNN results?

Good results (>90 %) ... but could be better.

The approach is supervised.

Anomalies are far in between.

How are the CNN results?

Good results (>90 %) ... but could be better.

The approach is supervised.

Anomalies are far in between.

There's a high probability that we cannot cover the entire problem domain.

How are the CNN results?

Good results (>90 %) ... but could be better.

The approach is supervised.

Anomalies are far in between.

There's a high probability that we cannot cover the entire problem domain.

Which gave me an idea ...

Can we construct a model purely from "good" data samples?

Can we construct a model purely from "good" data samples?

Alina Sodes Master's thesis investigated whether this is possible.

Can we construct a model purely from "good" data samples?

Alina Sodes Master's thesis investigated whether this is possible.

She found that Auto Encoders could be a viable approach.

Auto Encoder (AE)

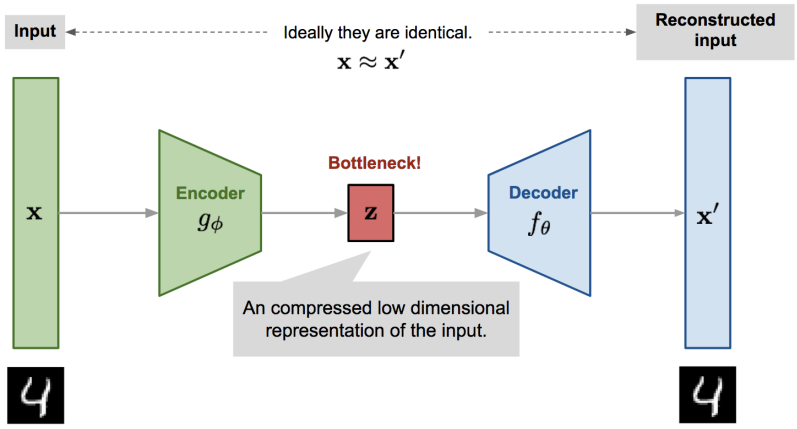

How do we predict with AEs?

1. Preprocess image to counter X-ray effects.

How do we predict with AEs?

1. Preprocess image to counter X-ray effects.
2. Apply the AE to the image.

How do we predict with AEs?

1. Preprocess image to counter X-ray effects.
2. Apply the AE to the image.
3. Compare the input to the output.

How do we predict with AEs?

1. Preprocess image to counter X-ray effects.
2. Apply the AE to the image.
3. Compare the input to the output.
4. If they match, the image was a "good" sample.

AE performance (kitkat)

AE performance (kitkat)

AE performance (kitkat)

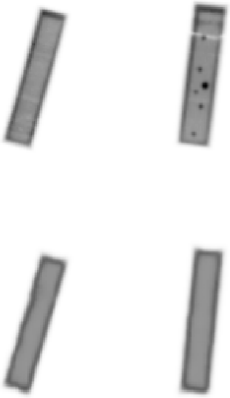

AE performance (kitkat)

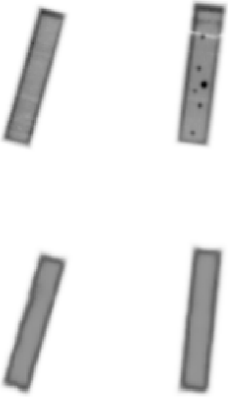

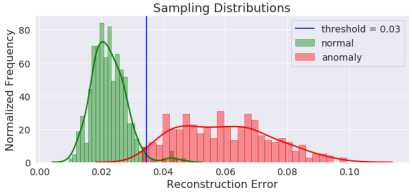

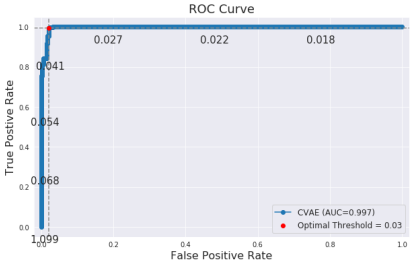

Accuracy of 98 %

AE performance (new potatoes)

AE performance (new potatoes)

AE performance (new potatoes)

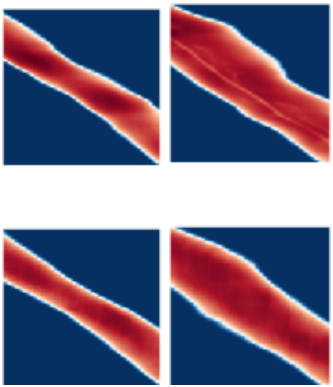

AE performance (new potatoes)

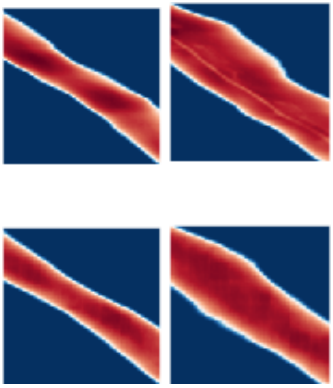

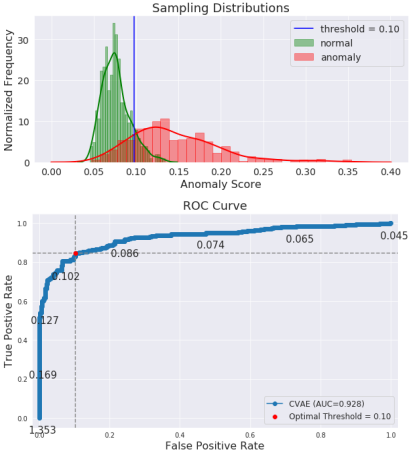

Accuracy of 86 %

How do we improve the performance?

How do we improve the performance?

After Alina Sode's master thesis, I further investigated using AEs.

How do we improve the performance?

After Alina Sode's master thesis, I further investigated using AEs.

Alina Sode trained on the entire image, which means that the model has to learn the shape.

The data

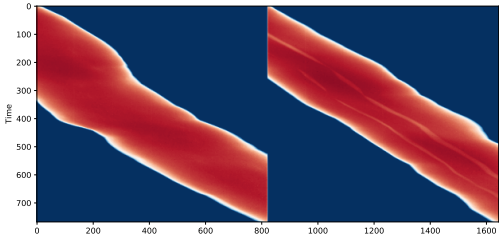
Time
0
100
200
300
400
500
600
700
0
200
400
600
800
1000
1200
1400
1600

The data

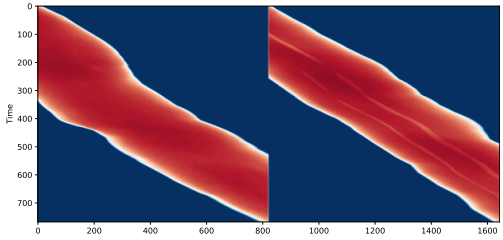
Time
0
100
200
300
400
500
600
700
0
200
400
600
800
1000
1200
1400
1600

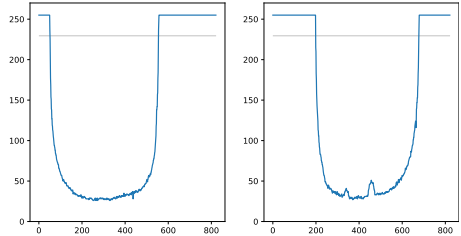
250
200
150
100
50
0
0
200
400
600
800

Virtually peel the potatoes

Virtually peel the potatoes

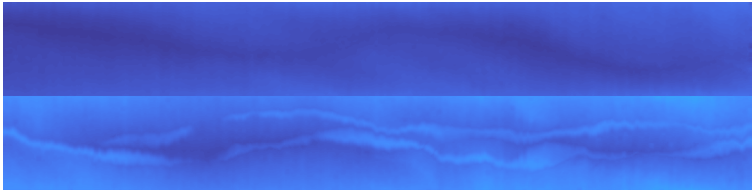

AE - good sample

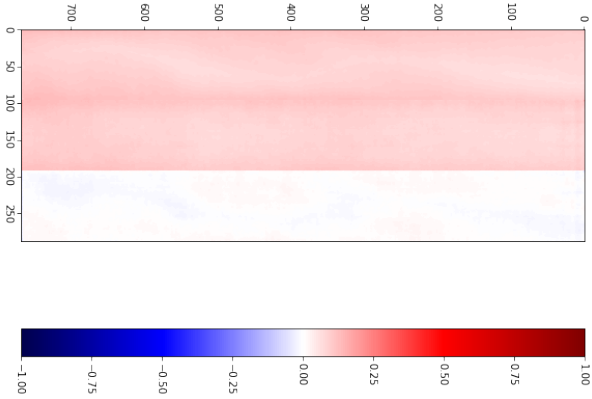

Input

Output

Difference

AE - bad sample

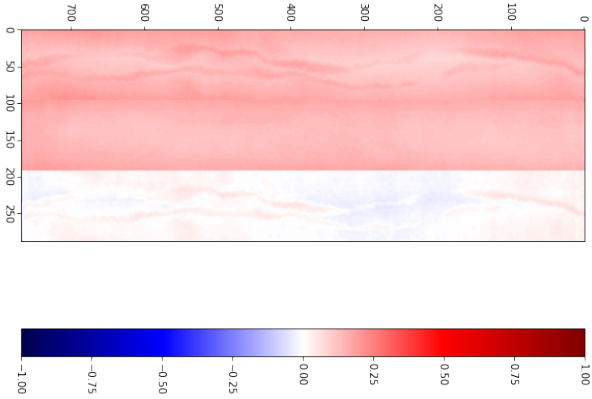

Input

Output

Difference

Image similarities

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Alina Sode looked at SE, IMED, SSIM, ZNCC, GMSD, with SSIM working best.

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Alina Sode looked at SE, IMED, SSIM, ZNCC, GMSD, with SSIM working best... in some cases.

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Alina Sode looked at SE, IMED, SSIM, ZNCC, GMSD, with SSIM working best... in some cases.

I further looked at RMSE, PSNR, FSIM, ISSM, SRE, SAM, UIQ and SH.

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Alina Sode looked at SE, IMED, SSIM, ZNCC, GMSD, with SSIM working best... in some cases.

I further looked at RMSE, PSNR, FSIM, ISSM, SRE, SAM, UIQ and SH.

From these, FSIM worked the best on these images.

Image similarities

We can clearly see the difference, so it should be possible for the machine as well.

Alina Sode looked at SE, IMED, SSIM, ZNCC, GMSD, with SSIM working best... in some cases.

I further looked at RMSE, PSNR, FSIM, ISSM, SRE, SAM, UIQ and SH.

From these, FSIM worked the best on these images.

FSIM works better than SSIM, due to the errors we see.

AE + FSIM - distribution

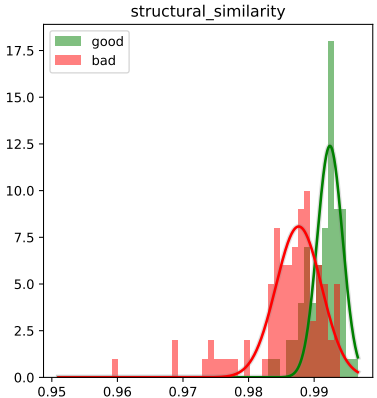

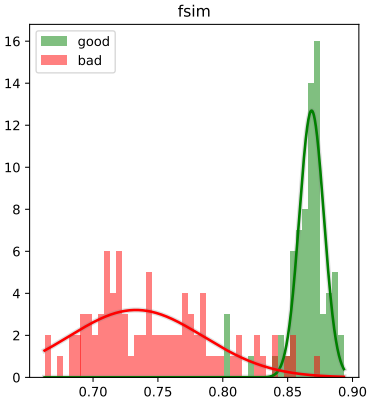

AE + FSIM - ROC

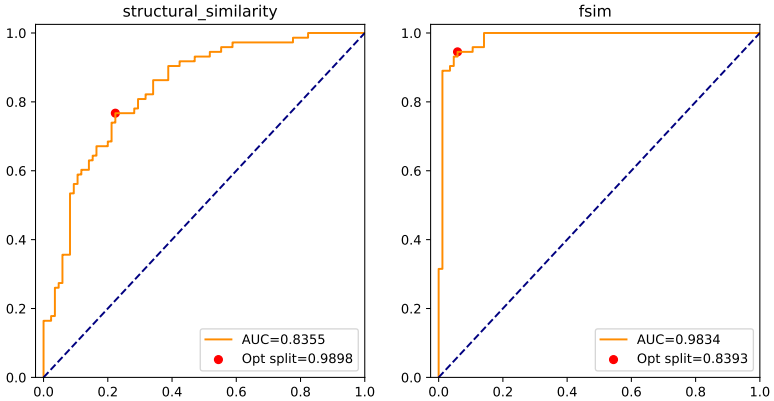
structural_similarity
AUC=0.8355
Opt split=0.9898
fsim
AUC=0.9834
Opt split=0.8393

Conclusion

- CNNs work great - but are supervised.
- AEs can work great - and are semi-supervised.
- Choosing the right similarity measure is important.