KYUNGPOOK NATIONAL UNIVERSITY

Recent Advances in Deep Anomaly Detection

Jaeil Kim

School of Computer Science and Engineering,
Kyungpook National University

jaeilkim@knu.ac.kr, threeyears@gmail.com

This tutorial ...

- A **short tutorial (50 min)** to introduce brief ideas and some promising approaches
  - Not an exhaustive survey
- Focus on anomaly detection in computer vision and medical image analysis
  - You can find much literature on novelty detection in many domains
- Review of recent anomaly detection methods in 2017 ~ 2020

Outline

- Medical imaging and anomaly detection
  - Clinical motivation
  - Anomaly detection system
  - Anomaly score
  - Challenges in anomaly detection
- Anomaly detection approaches in computer vision and medical imaging
  - Deep learning based approaches
    - Unsupervised and hybrid methods
    - One-class neural networks
  - Medical applications
  - Anomaly scores for medical images
- Summary
  - Limitations and Challenges

Anomaly Detection

Medical Imaging Applications

Medical Imaging

- “Process of creating visual representations of the interior of a body for clinical analysis and medical intervention”
  - Interpreted by medical doctors, radiologists or sonographers
  - For diagnostic purpose
- CT, MRI, ultrasound imaging, functional imaging, …


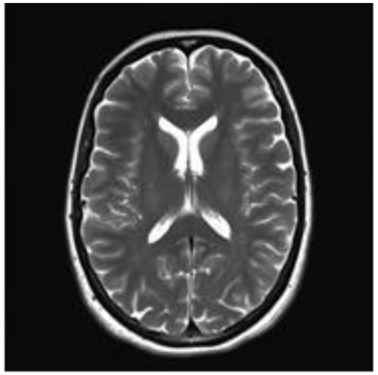

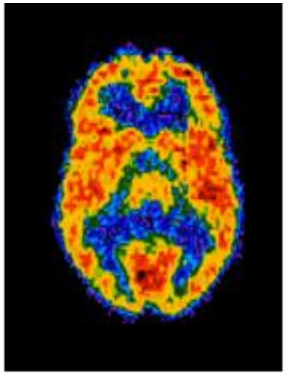
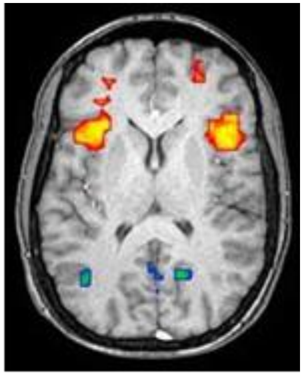

Brain Hemorrhage

- Diagnosis using medical images
  - **Find** any lesion
  - **Compare** and **measure**
- What can you see?
  - Same anatomical structure
    - Location and shape
  - Differences in pixel intensity and texture
  - Expansion of the structure?
  - **Compared to healthy subject's image**

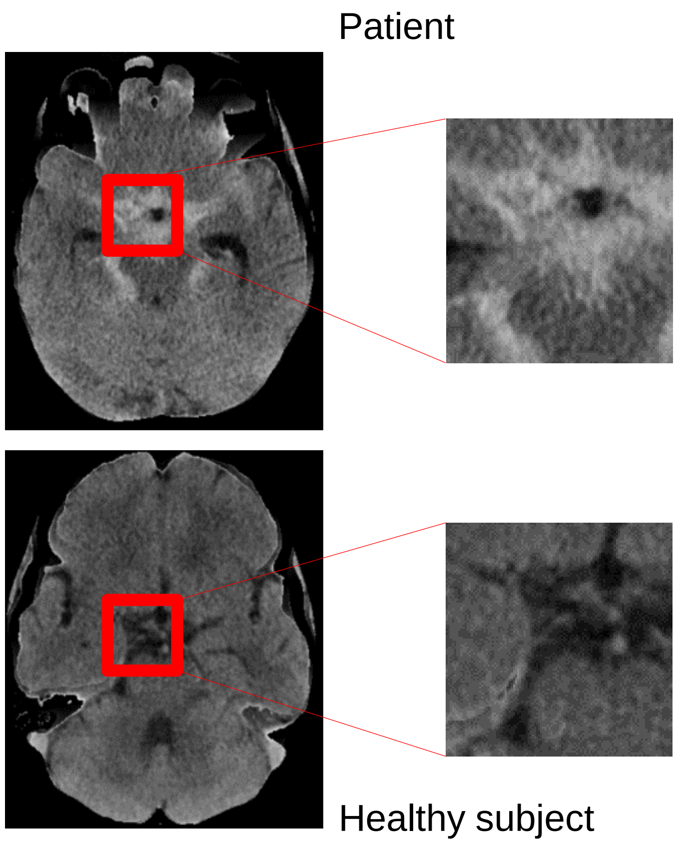

KYUNGPOOK NATIONAL UNIVERSITY

Anomalies in Medical Images

- Definition of **anomaly**
  - "A deviation from the common rule, type, arrangement, or form"
  - Also referred as abnormalities, outliers, or deviants
- Anomalies in diagnostic medical images
  - Uncommon patterns in images
  - **How to detect and quantify**

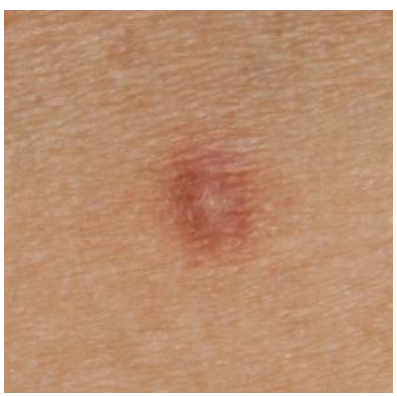

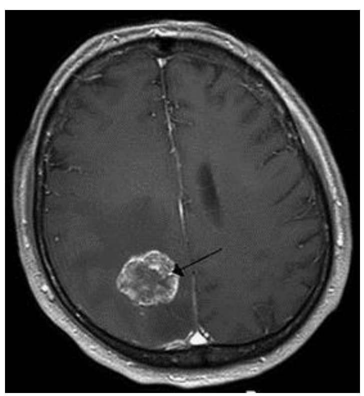

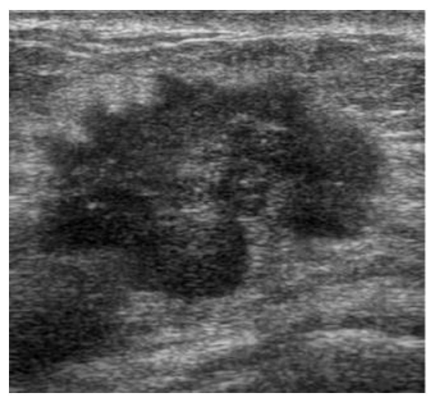

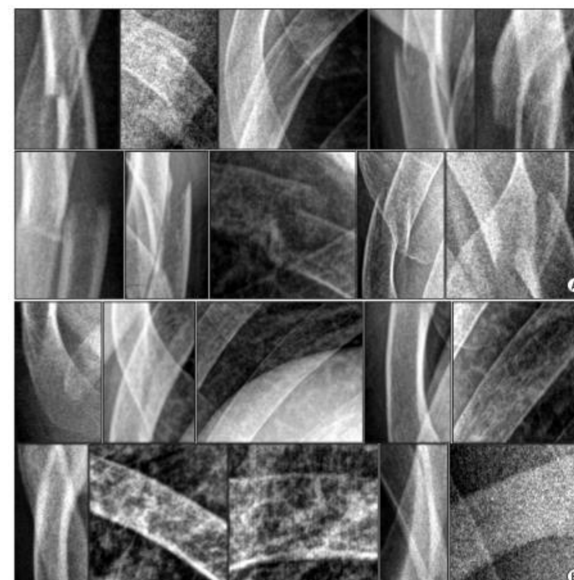

Anomaly Detection in Medical Imaging

- Early detection of lesions is critical
  - Better prognosis and able to prevent severe symptoms
- Problems in real world
  - Number of experienced doctors < number of patients and scans
  - Difficult to read all slices of 3D volumes → time-consuming
  - Dataset for possible anomalies → not available in most cases
- Clinical scenario

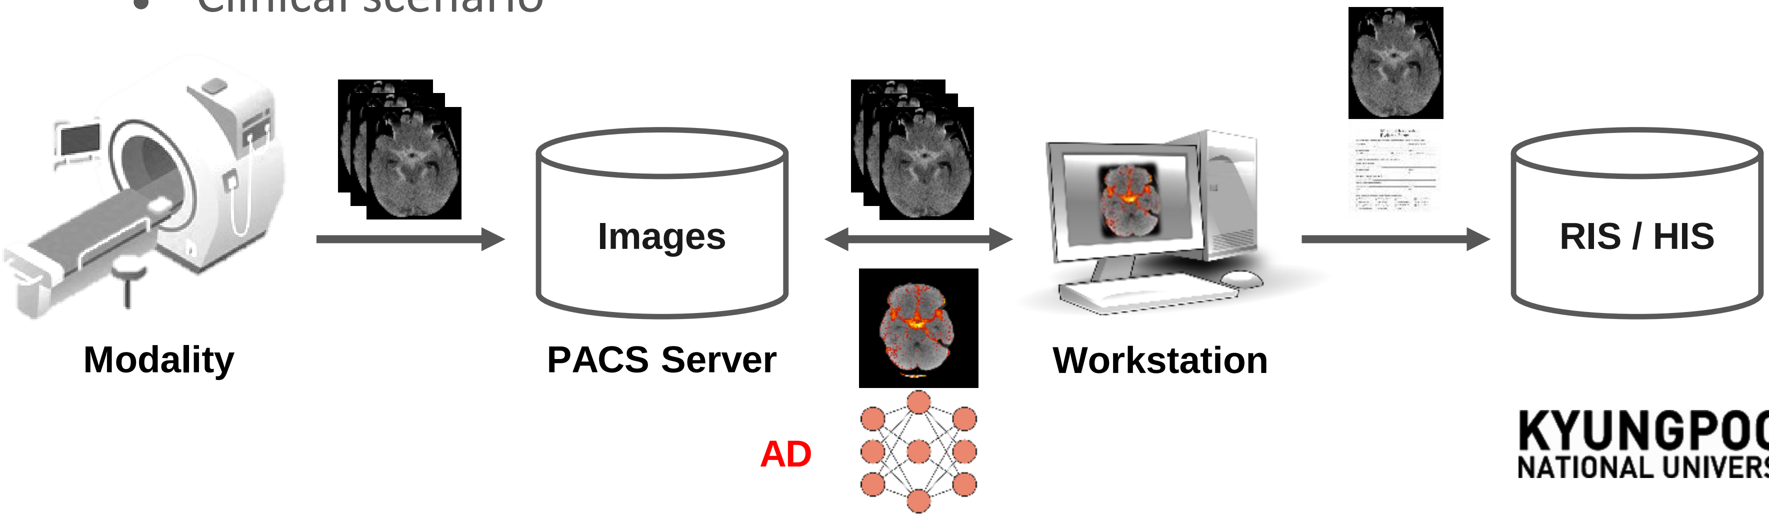

KYUNGPOOK NATIONAL UNIVERSITY

Characteristics of Anomalies

- Brain Hemorrhage
  - Hounsfield unit
  - Different grey levels to hemorrhage
  - Distinguishable from surrounding structures and normal cases
- **Anomalies**
  - Depend on organs and imaging modality
    - CT, MR, US, X-Ray, PET, ….
    - Brain, heart, lung, prostate, …
  - **Need information about normal cases**
  - **Understanding of abnormal changes**

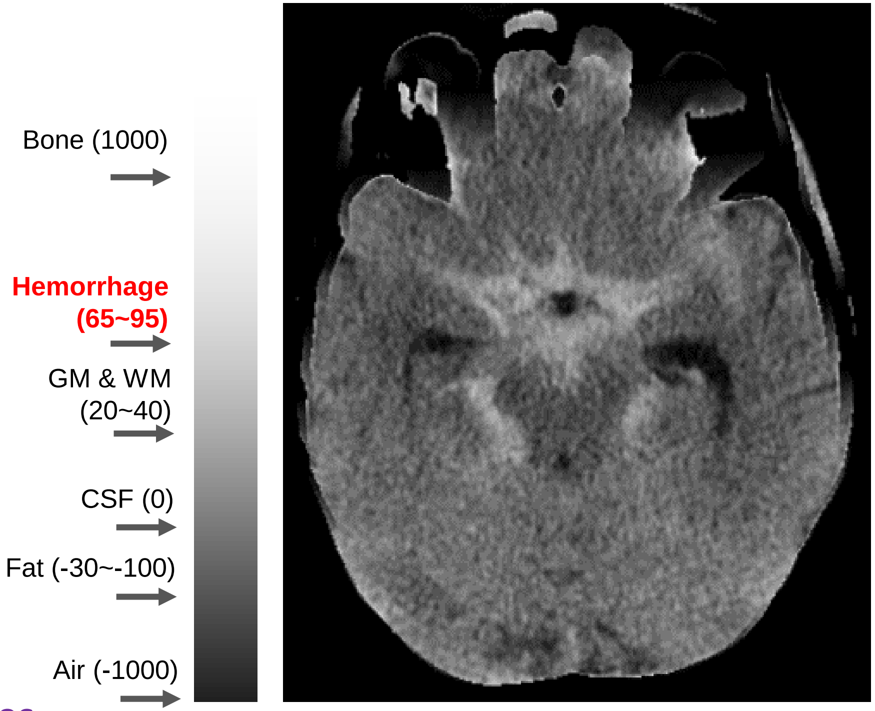

KYUNGPOOK NATIONAL UNIVERSITY

Various Representations of Brain in MRI

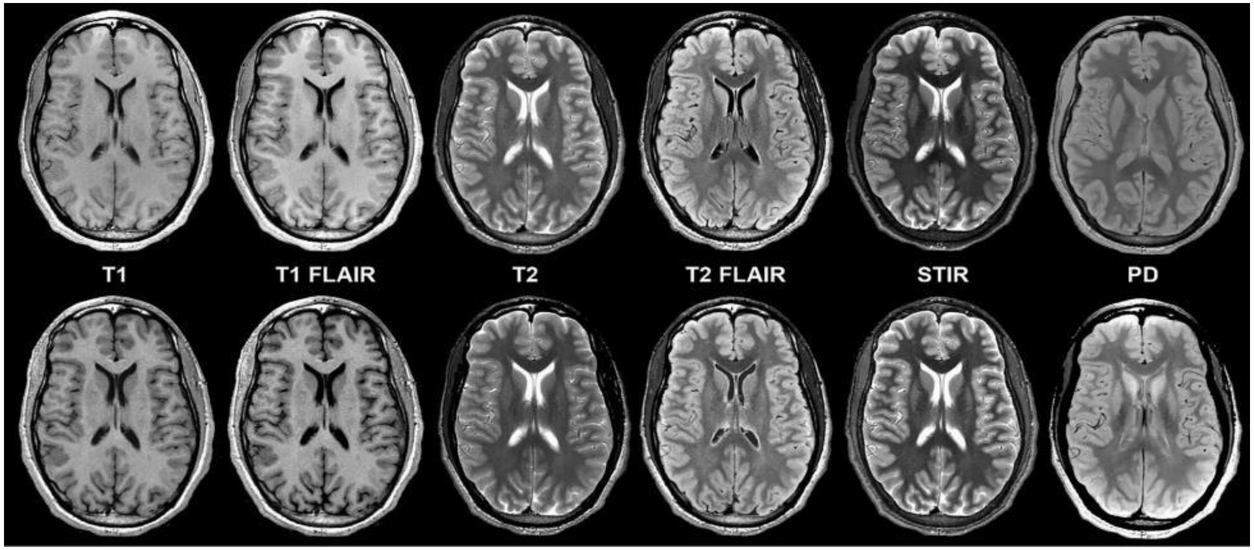

- Large variations across patients, imaging device, and hospital
- Difficult to define anomalies and image dissimilarity
    - **Data-driven approaches may be promising**

Tanenbaum, L. N., Tsiouris, A. J., Johnson, A. N., Naidich, T. P., DeLano, M. C., Melhem, E. R., ... & Field, A. S. (2017). Synthetic MRI for clinical neuroimaging: results of the Magnetic Resonance Image Compilation (MAGiC) prospective, multicenter, multireader trial. *American Journal of Neuroradiology*, *38*(6), 1103-1110.

KYUNGPOOK NATIONAL UNIVERSITY

Anomaly Detection System

- Aims to automatically **localize anomalies** and **quantify their differences** from **other observations** (normal cases)
  - Computer-aided anomaly detection
- Three fundamental approaches to anomaly detection
  - Using prior knowledge of anomalies and normality or not

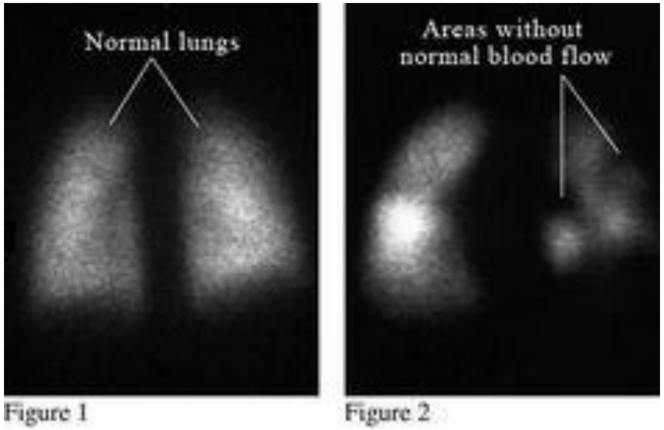

https://www.pinterest.com/pin/108790147231801412/

For medical imaging
- Strong or weak annotation by experts
- Regions of interest, landmarks, disease type labels

1. Supervised methods
2. Unsupervised methods
3. Semi-supervised methods

Hodge, V., & Austin, J. (2004). A survey of outlier detection methodologies. *Artificial intelligence review*, *22*(2), 85-126.

KYUNGPOOK NATIONAL UNIVERSITY

Anomaly Detection System

- Approach #1: **with prior knowledge of anomalies and normality**
  - Modeling normality and anomalies together
    - Using pre-labelled data
  - Limitations
    - Require a good spread of normal and abnormal data for generalization → not easy in most clinical cases
    - Classification is limited to the known distribution

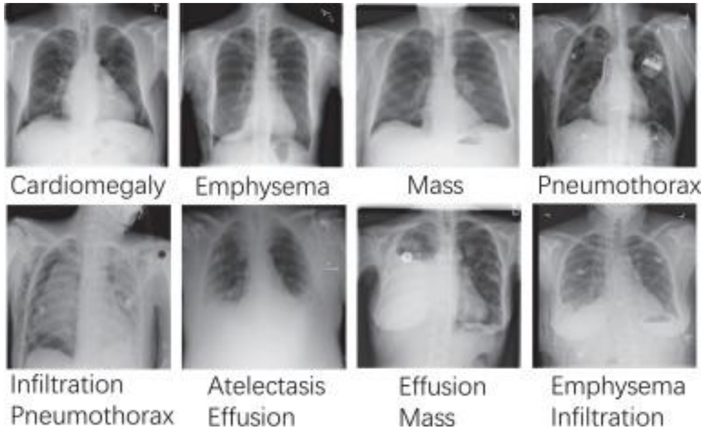

Multi-label chest X-Ray image classification

- Using 112,120 chest X-ray images (Chest X-ray14)
- Labelled for **14 disease pathologies** + **no-finding**

* https://www.kaggle.com/nih-chest-xrays/data

Guan, Qingji, and Yaping Huang. "Multi-label chest X-ray image classification via category-wise residual attention learning." *Pattern Recognition Letters* (2018).

KYUNGPOOK NATIONAL UNIVERSITY

Anomaly Detection System

- Approach #2: **without prior knowledge of training data**
  - Assume that **abnormal cases are well separated from the normal data**
  - Unsupervised clustering methods
    - Remove outliers in training data and fit models to the remaining
    - Use the outliers to build a robust classifier
  - Limitations
    - Difficult to recognize without expertise → machine can do it?
    - Require to define dissimilarity measures between images

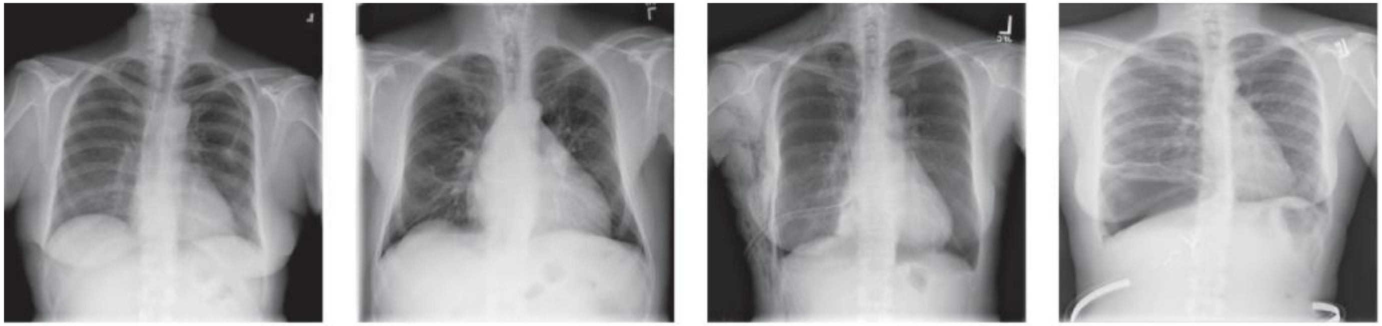

(a) “No Finding” (b) “Cardiomegaly” (c) “Pneumothorax” (d) “Pneumothorax”

Baltruschat, Ivo M., Hannes Nickisch, Michael Grass, Tobias Knopp, and Axel Saalbach. "Comparison of deep learning approaches for multi-label chest X-ray classification." *Scientific reports* 9, no. 1 (2019): 6381.

Anomaly Detection System

- Approach #3: **modeling normality only**
  - Learns to abnormality with supervision about normal class
    - Draw a boundary of normality
    - Suitable when the data set of possible anomalies is limited
    - A way medical doctors recognize lesions in medical images
  - Limitations
    - Lack of prior information on anomalies → high false alarms
    - Require enough samples for normality → how many?

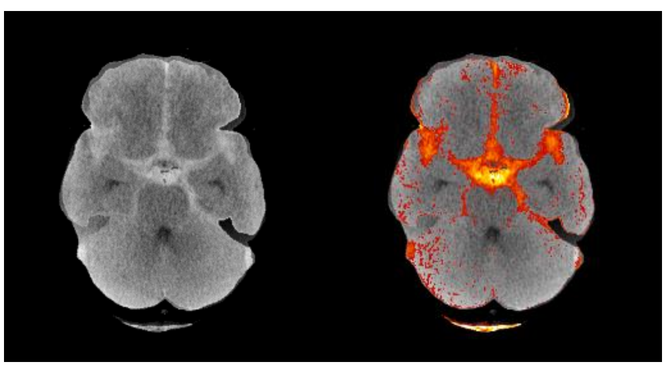

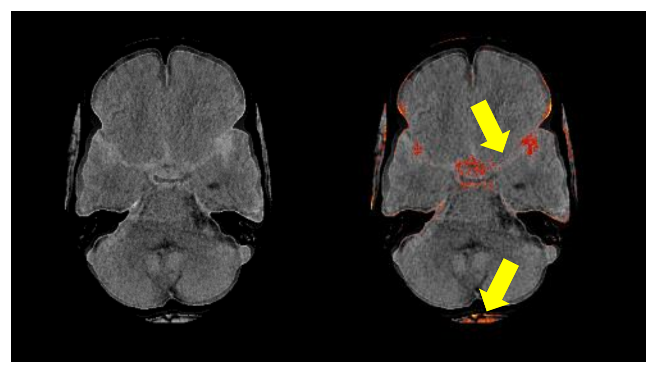

Voxelwise detection using z-score threshold

KYUNGPOOK NATIONAL UNIVERSITY

Challenges in Anomaly Detection

- **Modeling methods of normal cases**
  - Representation learning of medical images
    - Curse of dimensionality: 3D/4D volume (CT, MRI), large 2D image (X-Ray)
  - Lack of well-defined representative boundary of normal cases
- **Data collection from healthy subjects**
  - Good spread of normal cases: represent the diversity of target population
  - Factors to consider
    - Age, gender, imaging modality, hospital, and imaging protocol
- **Generalization of models**
  - How to build more generalized models of normality with limited sample
  - Data augmentation, hyperparameter settings, adding noise, …
- **Anomaly scores**
  - Margin from the boundaries of normal cases
  - Evaluation: pixel-wise, patch-wise, region-wise, and image-wise

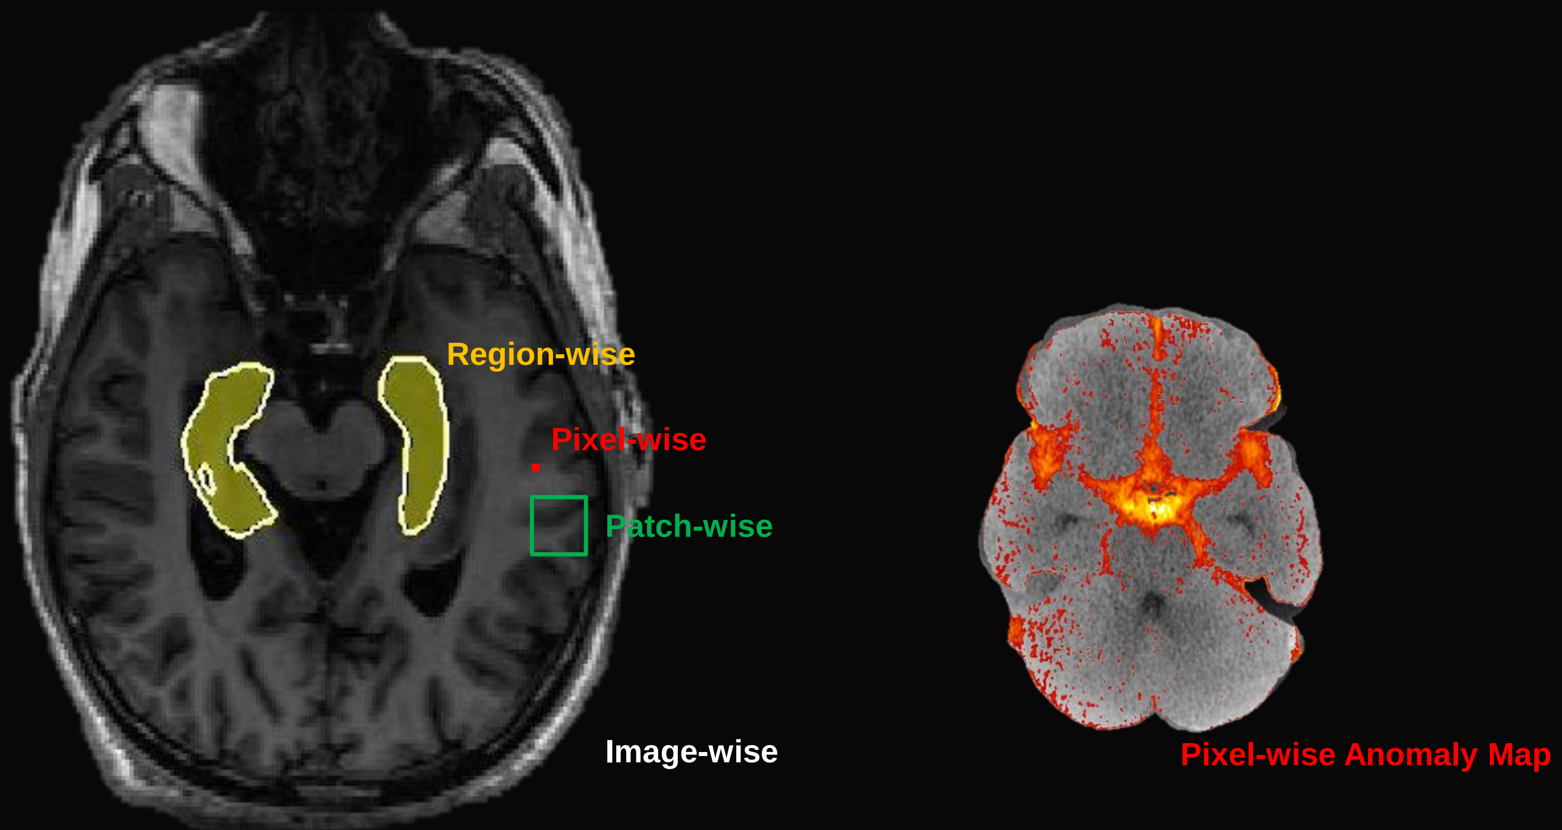
Region-wise
Pixel-wise
Patch-wise
Image-wise
Pixel-wise Anomaly Map

Deep Anomaly Detection

Anomaly Detection using Deep Learning

- Models
  - Unsupervised AD
    - Reconstruction-based
  - Hybrid method
    - Classifier using latent representations
  - Classification (one-class)
- Anomaly scores
  - Reconstruction loss
  - Distance from the center of data distribution
  - Depend on models and modalities

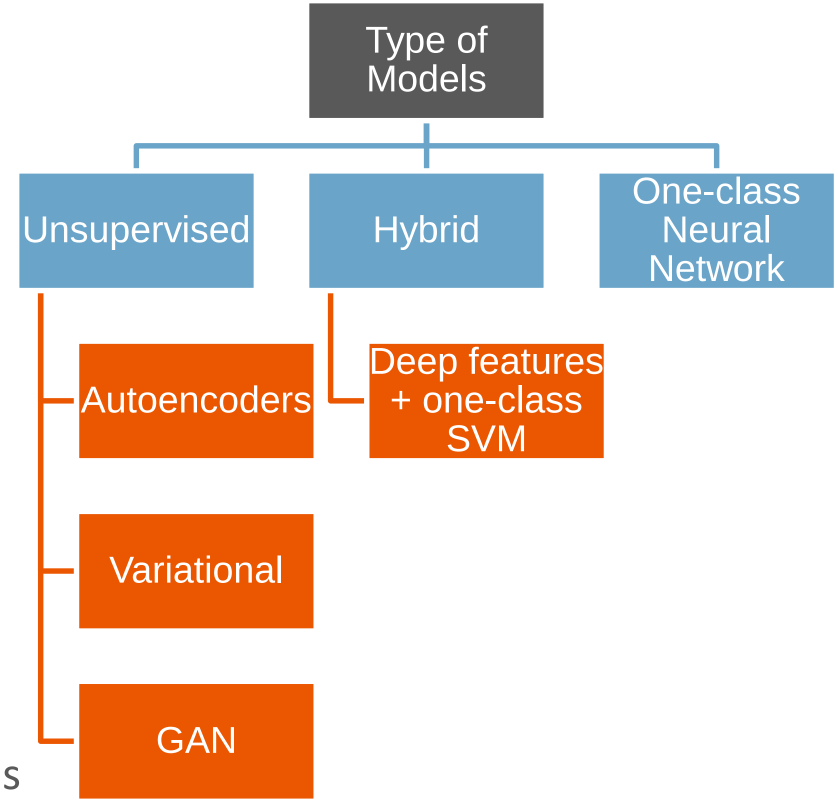

Chalapathy, Raghavendra, and Sanjay Chawla. "Deep learning for anomaly detection: A survey." *arXiv preprint arXiv:1901.03407* (2019).

Reconstruction-based Anomaly Detection

- Aims to localize anomalies in images
- Assume that anomalies yield **higher residuals** to the reconstruction of input data
- Deep autoencoders with/without adversarial framework

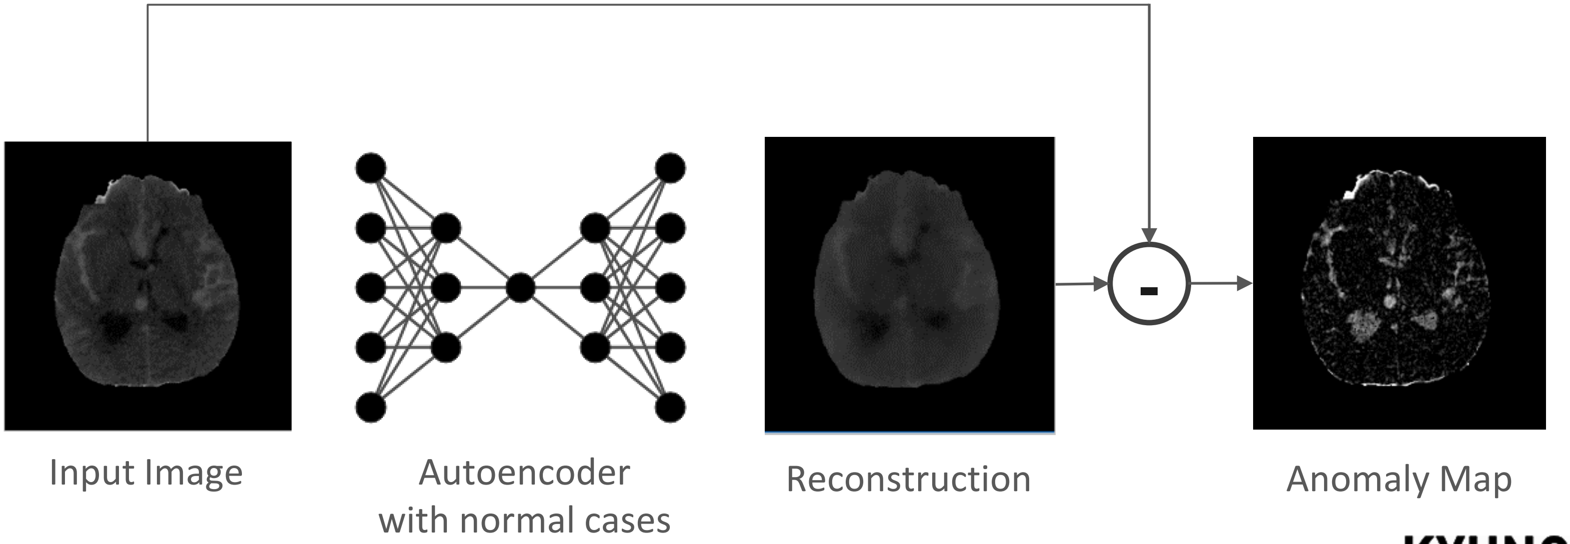

Autoencoder Approaches

- Robust deep autoencoder
- Deep autoencoding Gaussian mixture model
- Adversarial autoencoder

Robust Deep Autoencoders

- An extension of deep autoencoder
  - To discover high quality, non-linear features while eliminating noise
  - Inspired by Robust PCA
  - Optimization problem

    $\min_{D,E} \|X - D(E(X))\| \rightarrow$

    $\min_{\theta,S} \|L_D - D_\theta(E_\theta(L_D))\|_2 + \lambda \|S\|_1$

    (Denoising)

    $\min_{\theta,S} \|L_D - D_\theta(E_\theta(L_D))\|_2 + \lambda \|S\|_{2,1}$

    (Anomalies in group)

    subject to $X - L_D - S = 0$

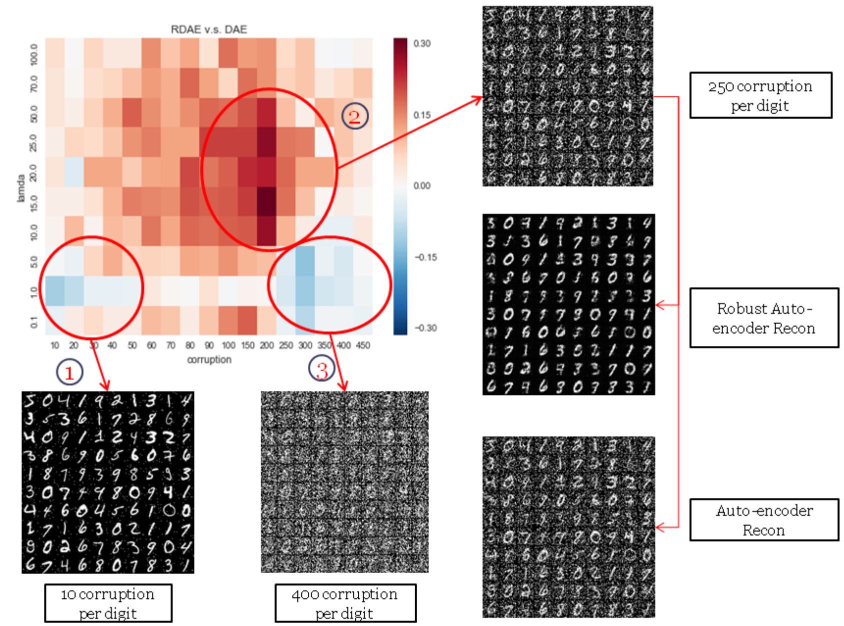

- Learning common patterns in data against image noise

Zhou, Chong, and Randy C. Paffenroth. "Anomaly detection with robust deep autoencoders." In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 665-674. ACM, 2017.

KYUNGPOOK NATIONAL UNIVERSITY

Robust Deep Autoencoders

- Anomaly detection
  - S: reconstruction error in $\min_{\theta,S} \|L_D - D_\theta(E_\theta(L_D))\|_2 + \lambda\|S\|_{2,1}$
  - Larger λ → sparser S
  - Grouped $l_{2,1}$ norm to find abnormal sample
- Working on abnormal sample with significant differences
  - Anomaly score → pixel-wise error

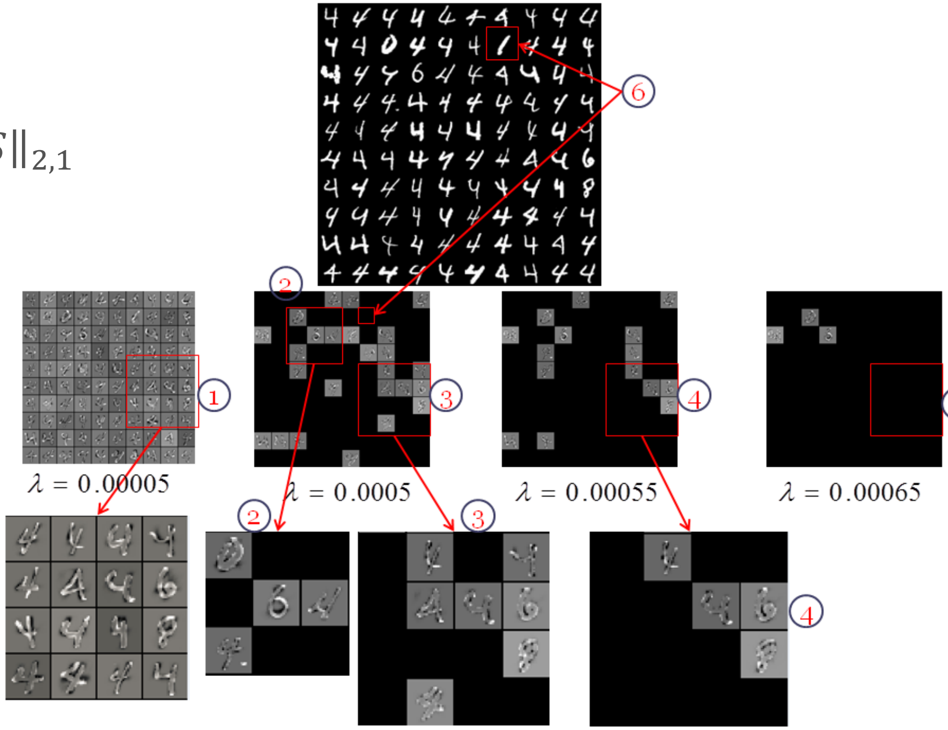

Zhou, Chong, and Randy C. Paffenroth. "Anomaly detection with robust deep autoencoders." In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 665-674. ACM, 2017.

Deep Autoencoding Gaussian Mixture Model

- Curse of dimensionality
  - Any sample could be a rare event with low probability to observe
  - Not easy for density estimation
  - Dimensionality reduction, but key features could be removed
- DAGMM
  - Compression: deep autoencoder
    - Latent features and reconstruction error
  - Estimation: predict the mixture membership for each sample
    - *K* mixture components

$$E(\mathbf{z}) = -\log\left(\sum_{k=1}^{K} \hat{\phi}_k \frac{\exp\left(-\frac{1}{2}(\mathbf{z}-\hat{\mu}_k)^T \hat{\Sigma}_k^{-1}(\mathbf{z}-\hat{\mu}_k)\right)}{\sqrt{|2\pi\hat{\Sigma}_k|}}\right)$$

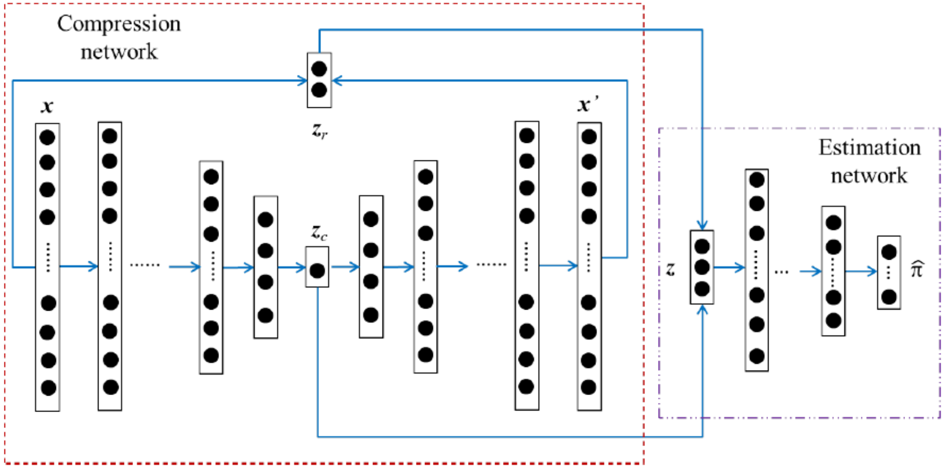

Zong, Bo, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. "Deep autoencoding gaussian mixture model for unsupervised anomaly detection." (2018).

Deep Autoencoding Gaussian Mixture Model

- Anomaly detection using DAGMM
  - Sample with high energy ($E(z)$)
    - Low likelihood in mixture-component distribution
  - Estimation network considers two features
    - Not only latent representation
    - Also consider reconstruction error

- Anomaly localization
  - How to use the membership information
  - Patch-wise or region-wise evaluation

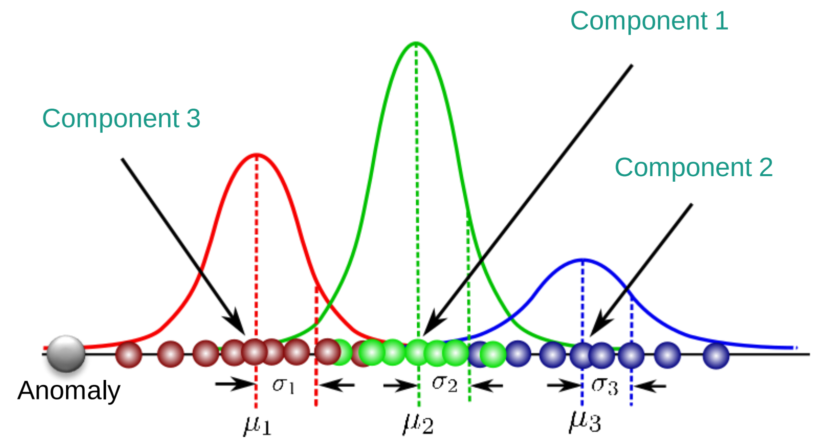

https://towardsdatascience.com/gaussian-mixture-models-explained-6986aaf5a95

Zong, Bo, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. "Deep autoencoding gaussian mixture model for unsupervised anomaly detection." (2018).

KYUNGPOOK NATIONAL UNIVERSITY

Adversarial Autoencoder

- Learning the probability distribution of normal sample
  - Controlling the latent distribution and image generation
  - To ensure good generative reconstruction of normal sample
- Autoencoder + two discriminators
  - D_z: distinguish btw latent representation and $\mathcal{N}(0,1)$
  - D_x: distinguish btw the reconstructed image (x') from $\mathcal{N}(0,1)$ and real image

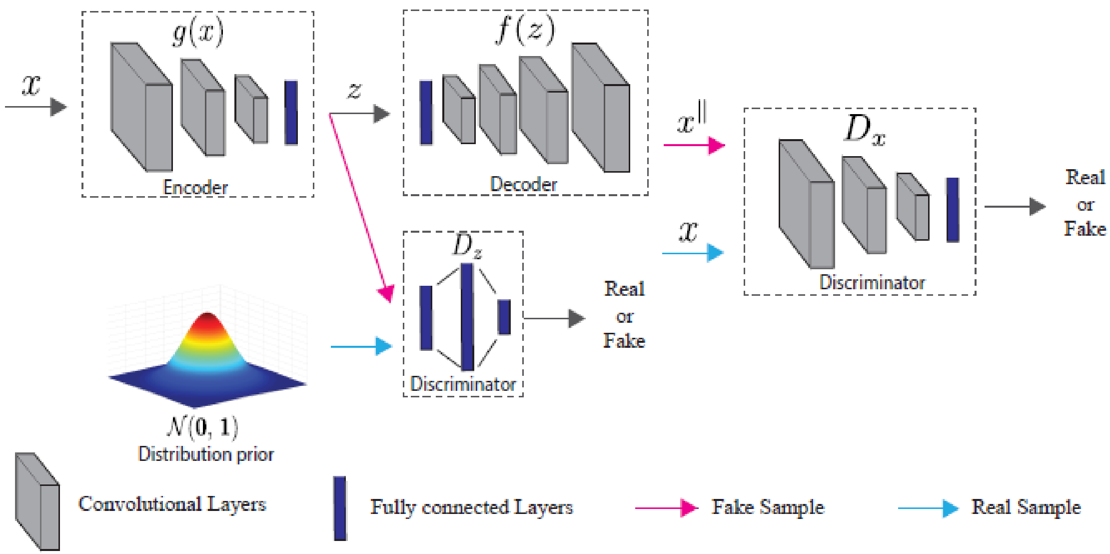

Pidhorskyi, Stanislav, Ranya Almohsen, and Gianfranco Doretto. "Generative probabilistic novelty detection with adversarial autoencoders." In *Advances in Neural Information Processing Systems*, pp. 6822-6833. 2018.

KYUNGPOOK NATIONAL UNIVERSITY

Adversarial Autoencoder

- Anomaly detection
  - By evaluating the inlier probability distribution on test sample
    - Density estimation in latent space
  - Image-wise anomaly evaluation → patch-wise approach

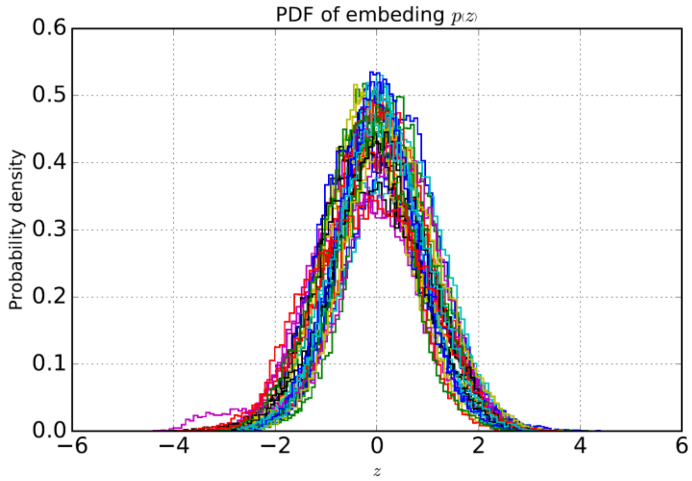

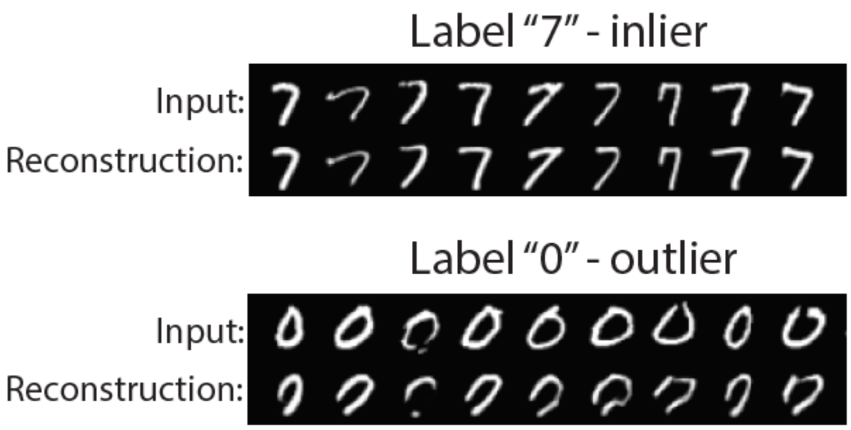

Reconstruction of inliers (7) and outliers (0)

Pidhorskyi, Stanislav, Ranya Almohsen, and Gianfranco Doretto. "Generative probabilistic novelty detection with adversarial autoencoders." In *Advances in Neural Information Processing Systems*, pp. 6822-6833. 2018.

Generative Adversarial Networks

- Deep convolutional GAN (AnoGAN)
- Ganomaly: a GAN for anomaly detection

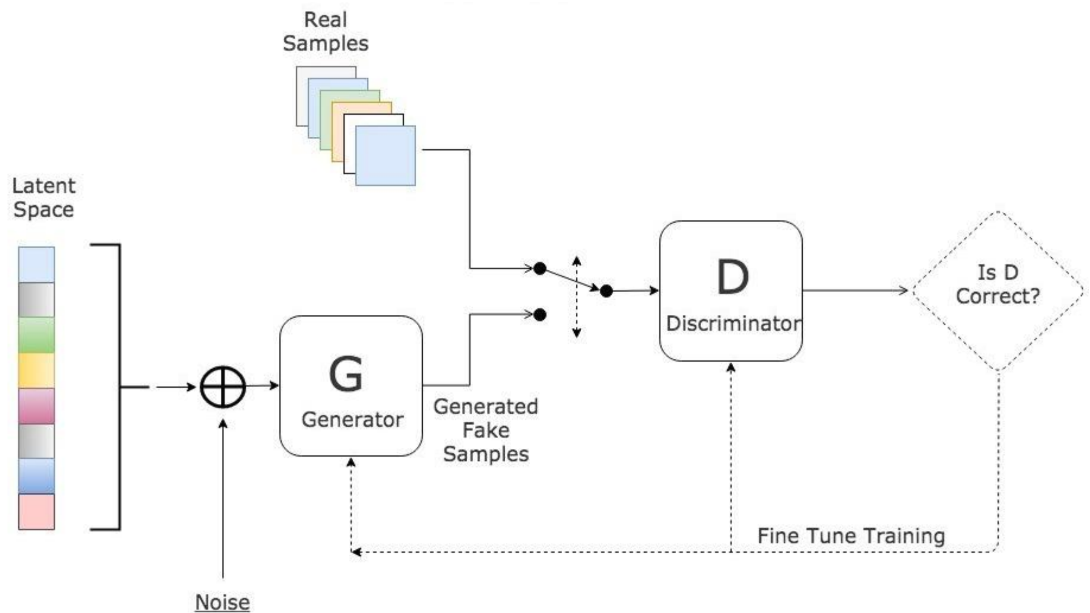

https://www.kdnuggets.com/2017/01/generative-adversarial-networks-hot-topic-machine-learning.html

Deep convolutional GAN (AnoGAN)

- Unsupervised learning
  - To create a rich generative model of normal cases
  - Manifold learning of normal anatomical variability

- Mapping images to latent space (z)
  - Latent space search
  - To find the best z for normal images
  - Using residual and discrimination losses

$$\mathcal{L}_R(\mathbf{z}_\gamma) = \sum |\mathbf{x} - G(\mathbf{z}_\gamma)|$$
$$\mathcal{L}_D(\mathbf{z}_\gamma) = \sum |\mathbf{f}(\mathbf{x}) - \mathbf{f}(G(\mathbf{z}_\gamma))|$$

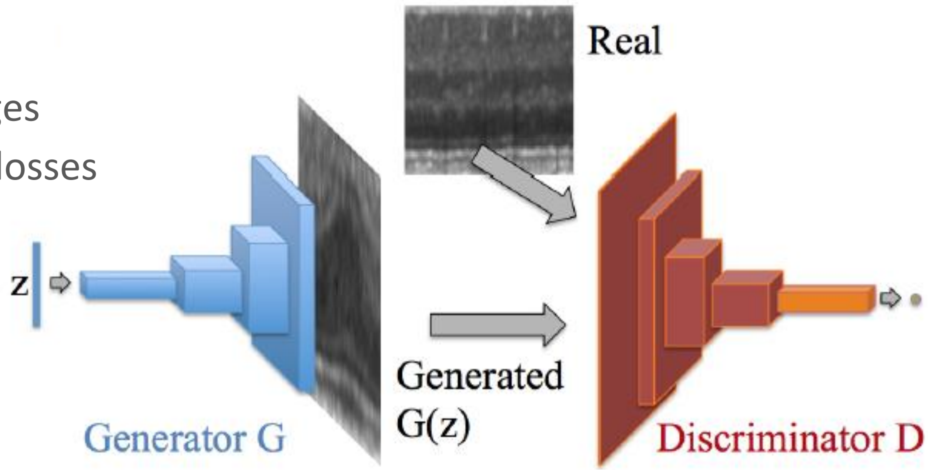

Schlegl, Thomas, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery." In *International Conference on Information Processing in Medical Imaging*, pp. 146-157. Springer, Cham, 2017.

Deep convolutional GAN (AnoGAN)

- Anomaly detection
  - Anomaly score: $A(\mathbf{x}) = (1 - \lambda) \cdot R(\mathbf{x}) + \lambda \cdot D(\mathbf{x})$
  - Pixel-wise residual loss to localize anomalies

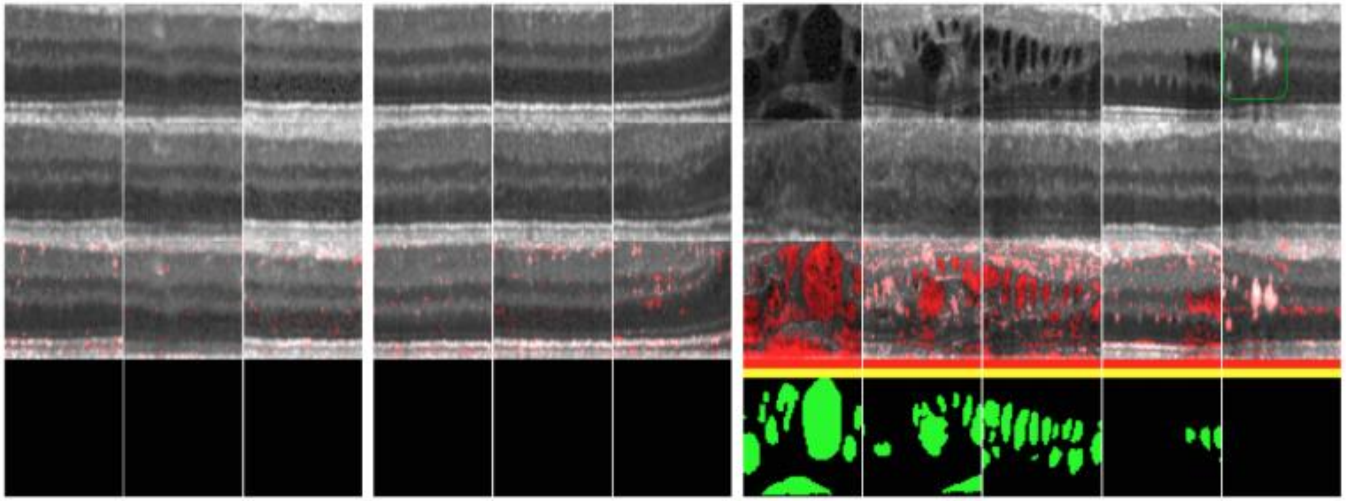

Schlegl, Thomas, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery." In *International Conference on Information Processing in Medical Imaging*, pp. 146-157. Springer, Cham, 2017.

GANomaly: a GAN for anomaly detection

- Problem of prior GAN based approaches
  - **Require latent space search** to remap the latent vector to new images
  - Expensive process and high complexity
- **GANomaly** = encoder-decoder-encoder pipeline
  - To capture the normal data distribution within image and latent space
  - Reconstruction network
    - **Adversarial auto-encoder (AAE)**
  - Latent vector generation
    - **Encoder** to learn the distributions from a latent space

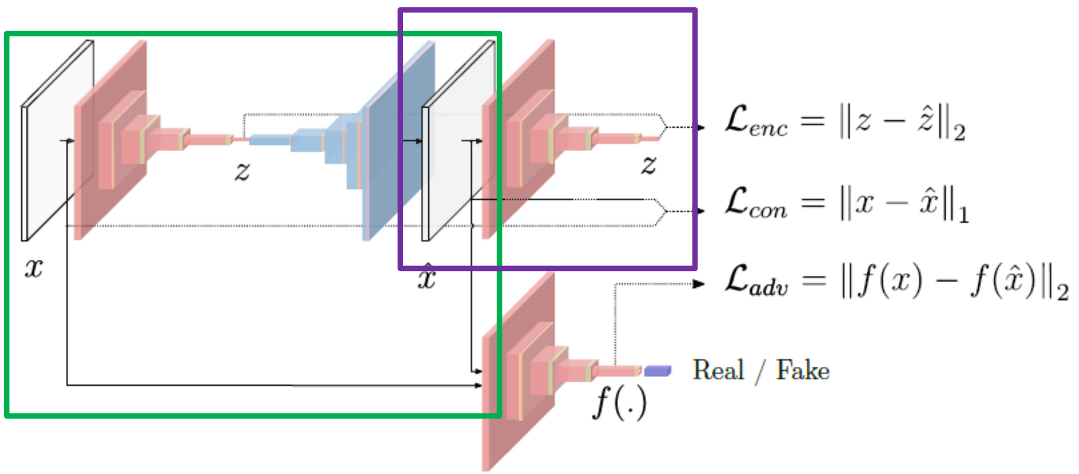

Akcay, Samet, Amir Atapour-Abarghouei, and Toby P. Breckon. "Ganomaly: Semi-supervised anomaly detection via adversarial training." In *Asian Conference on Computer Vision*, pp. 622-637. Springer, Cham, 2018.

GANomaly: a GAN for anomaly detection

- Anomaly detection
  - Anomaly score = distance between the latent vector of AAE and the inferred vector from the reconstructed image

$$\mathcal{A}(\hat{x}) = \|G_E(\hat{x}) - E(G(\hat{x}))\|_1$$

    - $\hat{x}$: “normal”-like image drawn from the normal data distribution

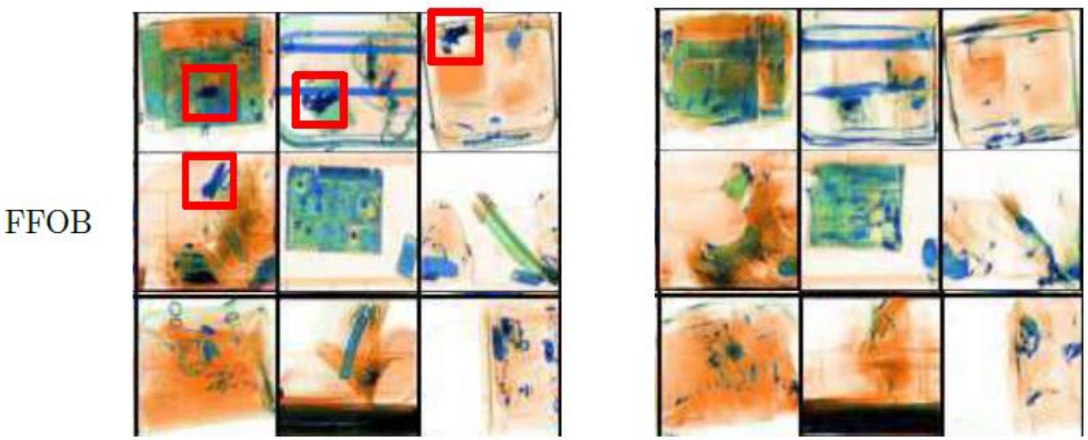

AAE fails to reconstruct abnormal sample → dissimilarity btw the latent vectors

Akcay, Samet, Amir Atapour-Abarghouei, and Toby P. Breckon. "Ganomaly: Semi-supervised anomaly detection via adversarial training." In *Asian Conference on Computer Vision*, pp. 622-637. Springer, Cham, 2018.

One-class Classification Methods

- Deep Support Vector Data Description (Deep SVDD)
- One-class Convolutional Neural Network (OC-CNN)
- Adversarially Learned One-Class Classifier for Novelty Detection

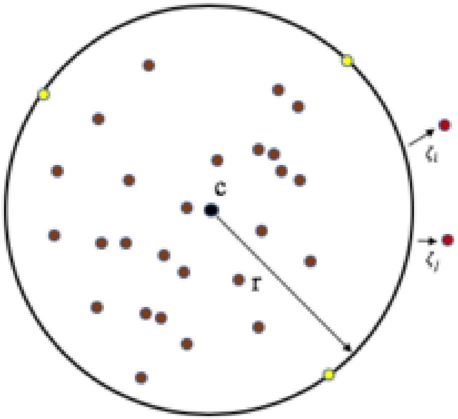

https://en.wikipedia.org/wiki/One-class_classification

Deep Support Vector Data Description

- Previous approaches in one-class classification
  - One-class SVM or kernel density estimation
  - Fails in high-dimensional, data-rich scenarios ← curse of dimensionality
- **Deep support vector data description**
  - A neural network minimizing the volume of a **hypersphere** that encloses the latent representations of normality data
  - Learning common features across normal instances

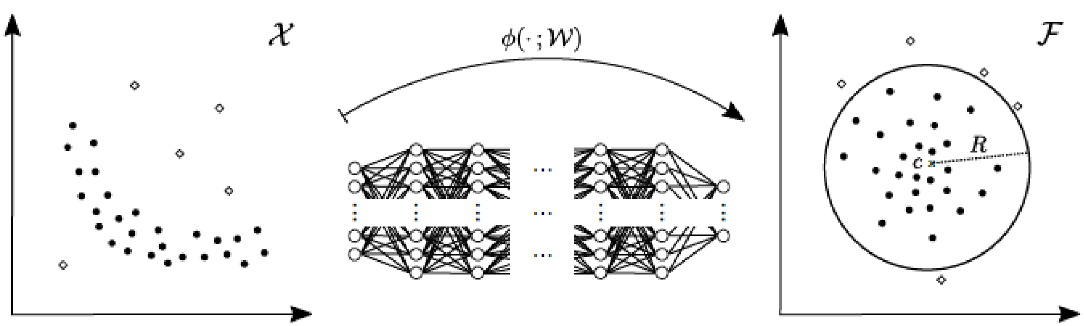

$$\min_{R,\mathcal{W}} \boxed{R^2} + \boxed{\frac{1}{\nu n}\sum_{i=1}^{n}\max\{0, \|\phi(\boldsymbol{x}_i;\mathcal{W}) - \boldsymbol{c}\|^2 - R^2\}} + \frac{\lambda}{2}\sum_{\ell=1}^{L}\|\boldsymbol{W}^{\ell}\|_F^2.$$

Ruff, Lukas, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. "Deep one-class classification." In *International conference on machine learning*, pp. 4393-4402. 2018.

Deep Support Vector Data Description

- Anomaly detection
  - Anomaly score = distance of the latent representation of new image to the center of the hypersphere

$$s(\boldsymbol{x}) = \|\phi(\boldsymbol{x}; \mathcal{W}^*) - \boldsymbol{c}\|^2$$

- Image-, patch-, and region-wise evaluation → rough localization

Most normal | Most anomalous | Most normal | Most anomalous

Ruff, Lukas, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. "Deep one-class classification." In *International conference on machine learning*, pp. 4393-4402. 2018.

One-class Convolutional Neural Network

- One-class classifier
  - Mapping image features distant from pseudo-negative data
  - Pseudo-negative data drawn from a Gaussian distribution
- Anomaly detection
  - Classification with softmax layer: membership of normality

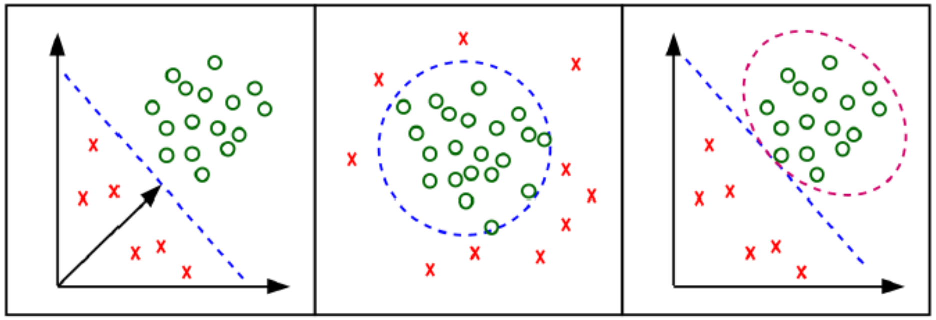

(a) OC-SVM (b) SVDD (c) Single-MPM

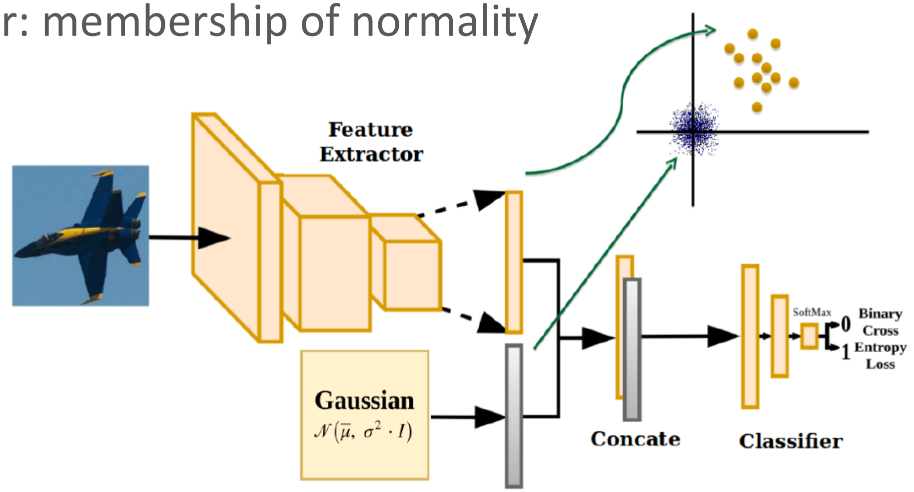

Oza, Poojan, and Vishal M. Patel. "One-class convolutional neural network." *IEEE Signal Processing Letters* 26, no. 2 (2018): 277-281.

Adversarially Learned One-Class Classifier

- Two networks with adversarial learning
  - Generator ($\mathcal{R}$): reconstruction from the normal data distribution
  - Discriminator ($\mathcal{D}$): classification for inliers and outliers
- Using discriminator for anomaly detection
- Input with random noise
  - Robust to image noise and distortions
  - Generalization

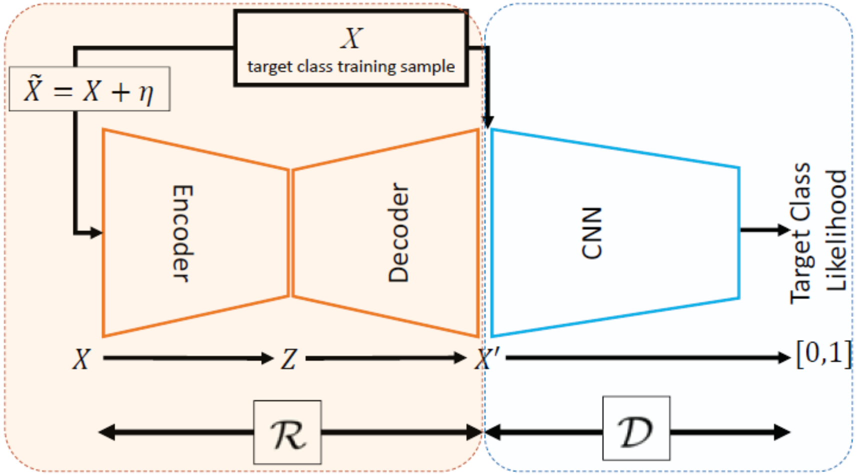

Sabokrou, Mohammad, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. "Adversarially learned one-class classifier for novelty detection." In Proceedings of the CVPR, pp. 3379-3388. 2018.

Adversarially Learned One-Class Classifier

- Anomaly detection
  - Using target class likelihood of reconstructed images
    - Simple threshold
    - Low probability → outlier
  - Generator usually fails to reconstruct outlier images

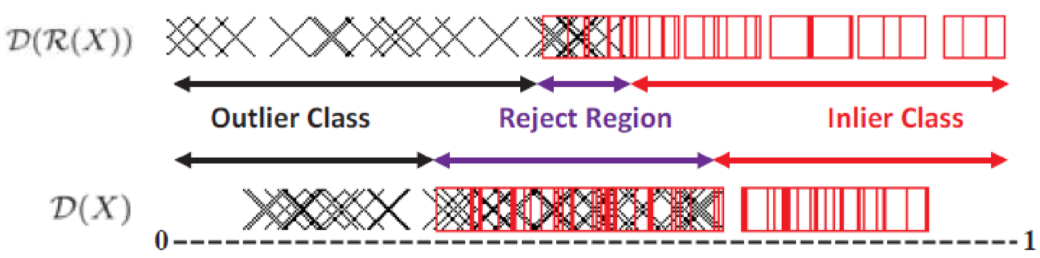

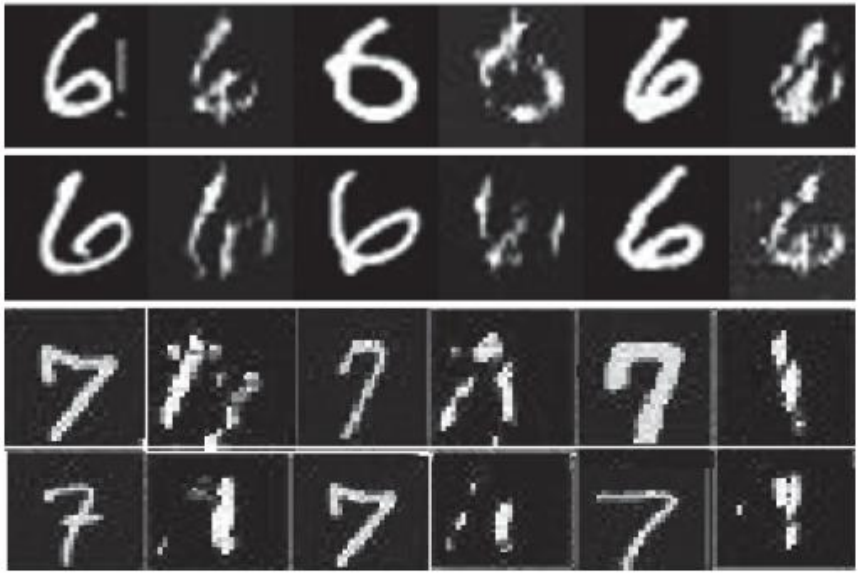

Sabokrou, Mohammad, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. "Adversarially learned one-class classifier for novelty detection." In Proceedings of the CVPR, pp. 3379-3388. 2018.

Deep Anomaly Detection for Medical Image Analysis

Anomaly Detection in Medical Images

- Autoencoder approaches
  - Aims to localize anomalies in medical images
  - Reconstruction errors → pixel-wise anomaly score
  - Using variational autoencoder and its variants for manifold learning
- Imaging domains
  - Optical coherence tomography (OCT): retinal disease
  - MRI, CT: brain tumor, brain hemorrhage
  - Diffusion MRI: multiple sclerosis
  - Neural networks designed for specific imaging modalities
- Anomaly scores

Anomaly Score using Variational Autoencoder Gradients

- Prior approaches
  - Reconstruction error as anomaly score
  - Assumption: models fail to reproduce anomalies not seen during training
- Anomaly score with gradients
  - Score = directions towards the normal data samples
  - Magnitude of the score = magnitude of abnormality
  - Score is the derivative of the error function w.r.t the input data

$$\frac{\partial \log p(x)}{\partial x} \approx \frac{\partial(-D_{KL}(q(z|x)||p(z)) + \mathbb{E}_{q(z|x)}[\log p(x|z)])}{\partial x}$$

Zimmerer, David, Jens Petersen, Simon AA Kohl, and Klaus H. Maier-Hein. "A Case for the Score: Identifying Image Anomalies using Variational Autoencoder Gradients." *arXiv preprint arXiv:1912.00003* (2019).

Anomaly Score using Variational Autoencoder Gradients

- Application to brain MR

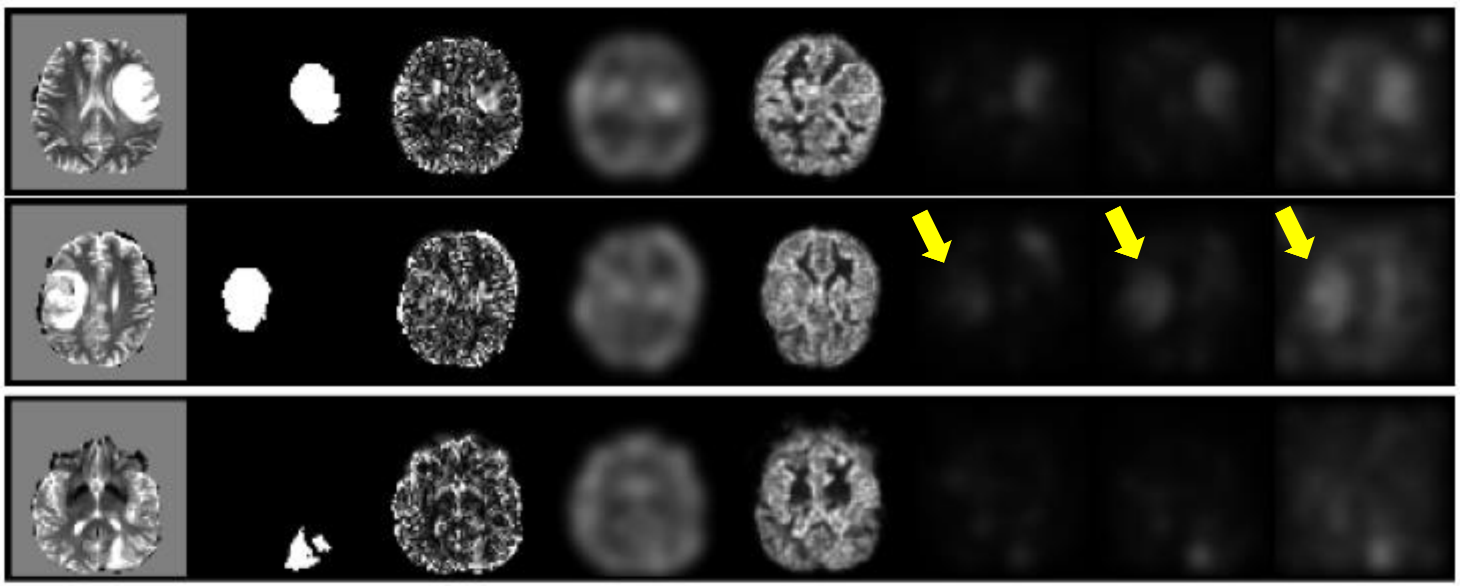

Input Lesion Reconstruction Error Variaional Autoencoder Gradients

Zimmerer, David, Jens Petersen, Simon AA Kohl, and Klaus H. Maier-Hein. "A Case for the Score: Identifying Image Anomalies using Variational Autoencoder Gradients." *arXiv preprint arXiv:1912.00003* (2019).

KYUNGPOOK NATIONAL UNIVERSITY

q-Space Novelty Detection with Variational Autoencoders

- Diffusion MRI
  - Mapping the diffusion process of water molecules
  - Multiple volumes with different b-values and gradient directions

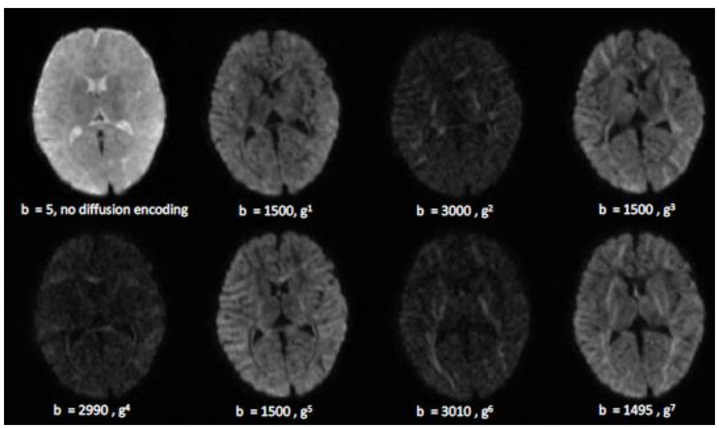

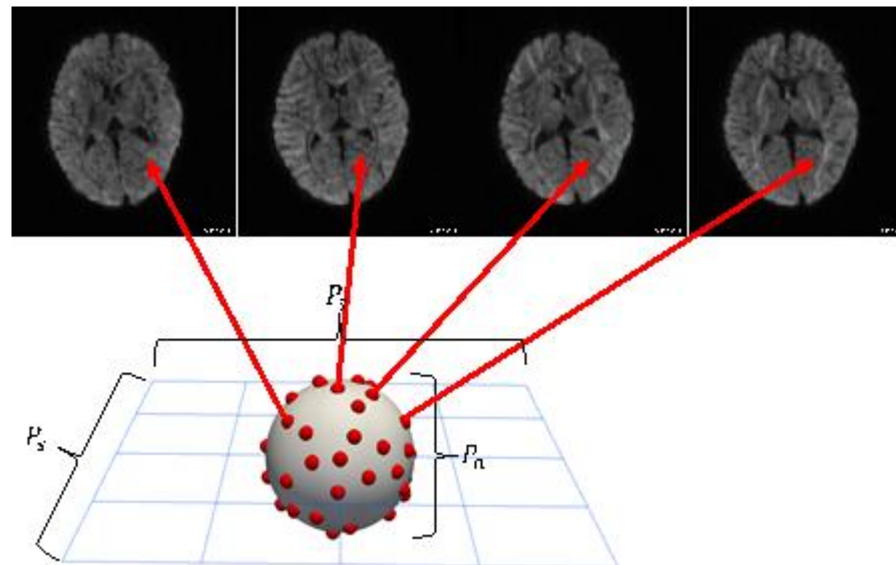

Diffusion Signals
of different gradient directions
at the same location

q-Space Novelty Detection with Variational Autoencoders

- Input feature
  - *d*-dimensional vector of diffusion directions
  - Voxel-wise modeling
- Normality modeling: Variational autoencoder
- Various anomaly scores
  - Deterministic error
  - Stochastic error
    - Multiple inference
    - Average and min error
  - Distance btw latent feature

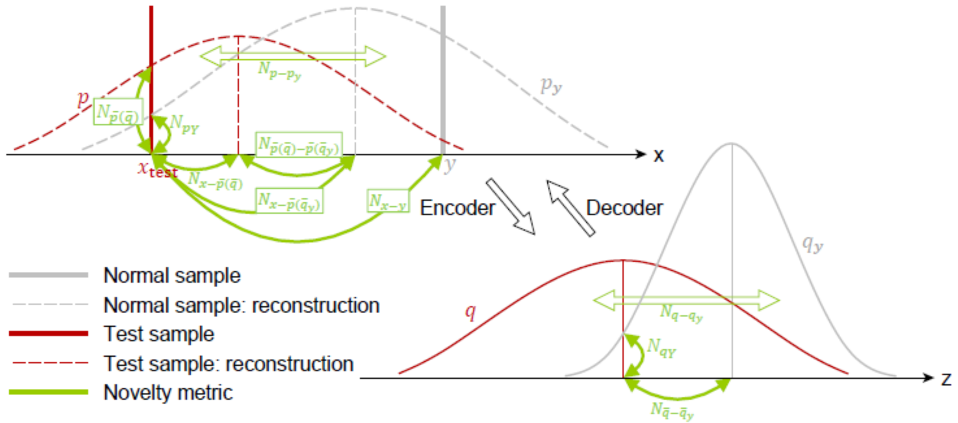

Vasilev, Aleksei, Vladimir Golkov, Marc Meissner, Ilona Lipp, Eleonora Sgarlata, Valentina Tomassini, Derek K. Jones, and Daniel Cremers. "q-Space Novelty Detection with Variational Autoencoders." arXiv preprint arXiv:1806.02997 (2018).

Uncertainty Autoencoder for Anomaly Detection

- Limitations of variational autoencoder
  - To regularize the latent space to follow normal distribution → difficult to model normal data distribution
- Spatial Uncertainty Auto-encoder
  - Add isotropic Gaussian noise to model the data distribution as Gaussian with fixed scalar variance
  - Spatial autoencoder to preserve anatomical features of brain images
  - Reconstruct "normal-like" brain image for disease cases

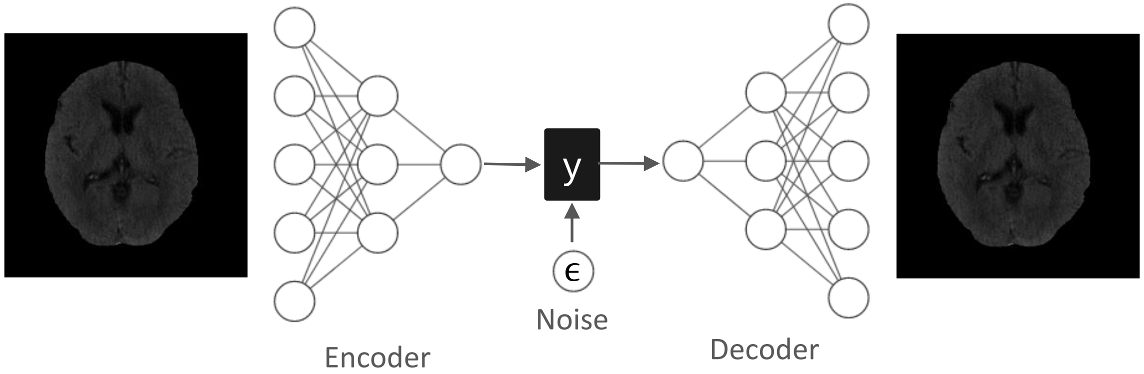

KYUNGPOOK NATIONAL UNIVERSITY

#4. Reconstruction-based Anomaly Detection

- Application to brain hemorrhage

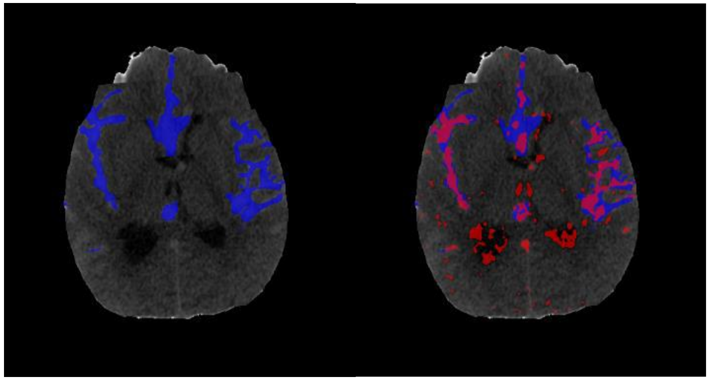

  - Better sensitivity compared to VAE, VAE-GAN

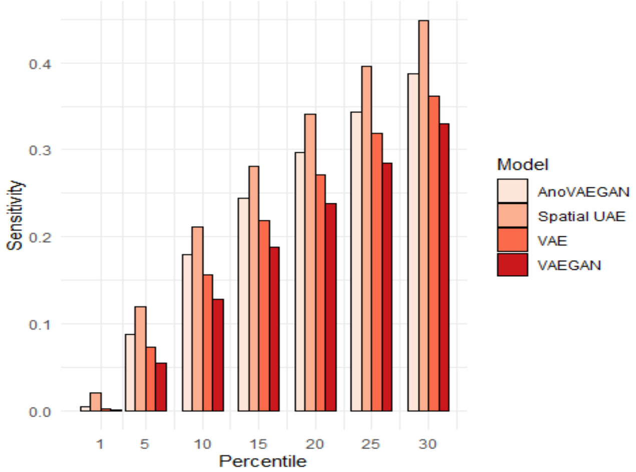

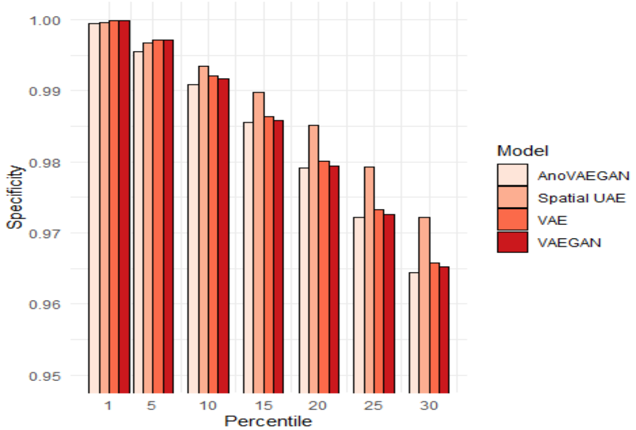

Grover, A., & Ermon, S. (2018). Uncertainty autoencoders: Learning compressed representations via variational information maximization. *arXiv preprint arXiv:1812.10539*.

Jaeil Kim, et al. Unsupervised anomaly detection in brain CT, in preparation

KYUNGPOOK NATIONAL UNIVERSITY

Summary

Summary

- Deep anomaly detection methods
  - Autoencoder approaches
  - Generative adversarial networks
  - One-class classification
- Open questions
  - How to reduce false positives and false negatives
    - Large variations across subjects and image quality
  - "Normal"-like reconstruction process
    - Reconstruction of abnormal region is unpredictable
  - Anomaly evaluation: image-wise, patch-wise, pixel-wise
  - Limited number of sample from normal cases
    - Generalization matters
  - Clinical interpretation of anomaly scores
    - how images are separated from normal cases

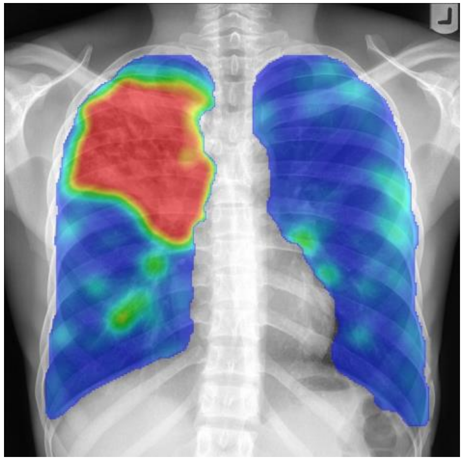

https://medium.com/analytics-vidhya/detecting-anomalies-in-x-ray-using-cnn-1e4c2e49f23a

KYUNGPOOK NATIONAL UNIVERSITY

Discussion

Jaeil Kim, Ph.D.

School of Computer Science and Engineering,
Kyungpook National University

jaeilkim@knu.ac.kr, threeyears@gmail.com

KYUNGPOOK NATIONAL UNIVERSITY

References

- Taboada-Crispi, A., Sahli, H., Hernandez-Pacheco, D., & Falcon-Ruiz, A. (2009). Anomaly detection in medical image analysis. In *Handbook of Research on Advanced Techniques in Diagnostic Imaging and Biomedical Applications* (pp. 426-446). IGI Global.
- Pimentel, M. A., Clifton, D. A., Clifton, L., & Tarassenko, L. (2014). A review of novelty detection. *Signal Processing*, *99*, 215-249.
- Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*.
- Tanenbaum, L. N., Tsiouris, A. J., Johnson, A. N., Naidich, T. P., DeLano, M. C., Melhem, E. R., ... & Field, A. S. (2017). Synthetic MRI for clinical neuroimaging: results of the Magnetic Resonance Image Compilation (MAGiC) prospective, multicenter, multireader trial. *American Journal of Neuroradiology*, *38*(6), 1103-1110.
- Hodge, V., & Austin, J. (2004). A survey of outlier detection methodologies. Artificial intelligence review, 22(2), 85-126.
- Guan, Qingji, and Yaping Huang. "Multi-label chest X-ray image classification via category-wise residual attention learning." Pattern Recognition Letters (2018).
- Baltruschat, Ivo M., Hannes Nickisch, Michael Grass, Tobias Knopp, and Axel Saalbach. "Comparison of deep learning approaches for multi-label chest X-ray classification." Scientific reports 9, no. 1 (2019): 6381.
- Zhou, Chong, and Randy C. Paffenroth. "Anomaly detection with robust deep autoencoders." In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 665-674. ACM, 2017.
- Zong, Bo, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. "Deep autoencoding gaussian mixture model for unsupervised anomaly detection." (2018).

References

- Schlegl, Thomas, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. "Unsupervised anomaly detection with generative adversarial networks to guide marker discovery." In International Conference on Information Processing in Medical Imaging, pp. 146-157. Springer, Cham, 2017.
- Akcay, Samet, Amir Atapour-Abarghouei, and Toby P. Breckon. "Ganomaly: Semi-supervised anomaly detection via adversarial training." In Asian Conference on Computer Vision, pp. 622-637. Springer, Cham, 2018.
- Ruff, Lukas, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. "Deep one-class classification." In International conference on machine learning, pp. 4393-4402. 2018.
- Oza, Poojan, and Vishal M. Patel. "One-class convolutional neural network." IEEE Signal Processing Letters 26, no. 2 (2018): 277-281.
- Sabokrou, Mohammad, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. "Adversarially learned one-class classifier for novelty detection." In Proceedings of the CVPR, pp. 3379-3388. 2018.
- Zimmerer, David, Jens Petersen, Simon AA Kohl, and Klaus H. Maier-Hein. "A Case for the Score: Identifying Image Anomalies using Variational Autoencoder Gradients." arXiv preprint arXiv:1912.00003 (2019).
- Vasilev, Aleksei, Vladimir Golkov, Marc Meissner, Ilona Lipp, Eleonora Sgarlata, Valentina Tomassini, Derek K. Jones, and Daniel Cremers. "q-Space Novelty Detection with Variational Autoencoders." arXiv preprint arXiv:1806.02997 (2018).

KYUNGPOOK NATIONAL UNIVERSITY