IMPROVING OSTEOARTHRITIS DIAGNOSIS WITH GAN-BASED PRETRAINING

Teun W. den Boer | 5743877
9/5/2025

TU Delft | CSE3000
Supervisor : Gijs van Tulder
Responsible Professor: Jesse Krijthe

BACKGROUND

- Osteoarthritis is a progressive disease diagnosed via X-rays.
- Manual diagnosis is time-consuming and subjective, especially in early stages.
- Automated methods could improve diagnosis speed and accuracy.
- Labeling medical data is expensive
- Hence self-supervised methods
- Can GANs be used

RESEARCH OBJECTIVE

Can GAN-based pretraining improve osteoarthritis classification from X-rays compared with supervised methods?

- Test GAN-discriminator pretraining
- Compare to standard supervised learning.
- Explore fine-tuning strategies.

STEP 1: GAN PRE-TRAINING

- A GAN, consisting of a Generator and a Discriminator, is trained on thousands of unlabeled X-rays.
- The Generator's job: Create realistic fake X-ray images.
- The Discriminator's job: Learn to distinguish real vs. fake images. By doing this, it implicitly learns the defining features of a real X-ray.
- Optimized using WGAN-GP

STEP 2: SUPERVISED FINE-TUNING

- We discard the Generator and repurpose the pre-trained Discriminator as a classifier.
- This model is then trained on a small, labeled dataset to classify images as "Healthy" or "Osteoarthritis".
- Baseline for Comparison: We compare its performance against a control model, the same architecture trained from scratch with random weights.
- Does replacing the classification head improve the model

RESULTS

The GAN's generator successfully produced realistic X-rays, indicating the model learned the underlying structure and texture of the training data. The GAN pre-trained model was better at identifying osteoarthritis than the standard model trained from scratch. Replacing the classification layer provided an additional, smaller performance increase.

GENERATED IMAGES

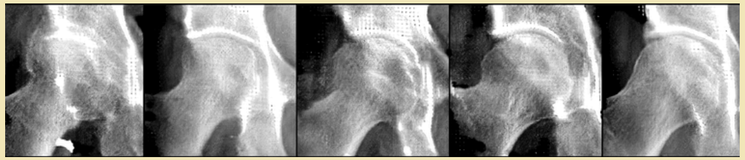

REAL IMAGES

- GAN-pretrained models achieved higher AUC than baseline
- Base Discriminator: 0.72 vs. 0.68
- Best Architecture: 0.734 vs. 0.727
- Discriminator learns useful features but correlation with diagnosis is weak.
- Training GANs turned out to be computationally expensive

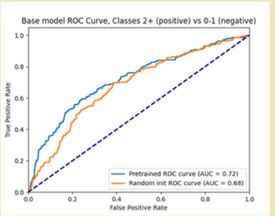

DISCUSSION

- Modest improvement suggests limited practical benefit considering computational cost.
- Standardized X-rays may limit GAN learning diversity.
- Future work: Larger datasets, varied architectures, hyperparameter tuning.

KEY TAKEAWAYS

- GAN-based pre-training improves osteoarthritis classification.
- But Gains are Modest
- The improvements do not justify the computational cost and complexity