Paper ID	1570744319
Title	Pre-trained vs. Random Weights for Calculating Fréchet Inception Distance in Medical Imaging
Author	Jamie A O'Reilly and Fawad Asadi (Rangsit University, Thailand)
Email	jamie.o@rsu.ac.th

Abstract

Fréchet Inception Distance (FID) is an evaluation metric for assessing the quality of images synthesized by generative models. Conventionally this involves using an Inception v3 convolutional neural network that has been pre-trained to classify everyday color images with the ImageNet dataset. The final classification section of this network is omitted, leaving an efficient feature extractor that outputs an encoded representation of each input image in the form of a 2048 element vector. Difference or similarity between samples of images can then be compared by measuring the distance between the distributions of their corresponding feature representations. Researchers have raised concerns about the utility of FID for evaluating unorthodox images (e.g. medical images) that are unlike those used for model training; suggesting that randomly initialized convolutional neural networks may be more appropriate. The aim of this study was to compare pre-trained and random approaches for evaluating medical images. Robustness to synthetic image distortions (Gaussian noise, blurring, swirl, and impulse noise) and different image types (x-ray, CT, fundus, and everyday images) was addressed. Feature representations were converted into two-dimensional space and visualized using t-distributed stochastic neighbor embedding (tSNE) and principal component analysis (PCA). Normalized FID between image classes was substantially larger and more consistent for the pre-trained model. Overall, this suggests that the pre-trained model is preferable to the randomly initialized model for evaluating medical images.