Deep Generative Models for Human Activity Recognition

Lakshmi Mukkawar and Katayoun Farrahi

Electronics and Computer Science, University of Southampton

Obtaining labelled human activities from wearable data is costly and cumbersome. Moreover, in real-world scenarios, we have class imbalance and various sources of noise often resulting in missing data. In this work, we generate synthetic human activity data using deep generative models as a way to handle class imbalance and missing data. We consider Variational Autoencoders (VAE) and Generative Adversarial Networks (GAN) and explore how to optimize these models for sensor data generation. We explore various methodologies for generating data to boost a supervised activity recognition classifier and compare the VAE versus GAN architectures for a class imbalance problem.

We use data from the UCI repository and consider various sensor features like acceleration and angular velocity. First, we apply a supervised classifier to classify data into different activities and use this as our baseline. We then build a VAE model, consisting of an encoder and decoder, to generate data. We experiment with different optimizers (rmsprop, adam, etc.) and loss functions (MSE, binary cross entropy) and explore how to optimize the VAE for data generation.

To evaluate the quality of generated data, we pass new data to the supervised classifier that has been built to predict labels for data. We observe that newly generated data is classified correctly with the same supervised model. Another interesting insight is to check the score of the same model when we augment it with new data. We start with 7352 records of data and generated new data resulting in 13000 records. When we merge new and original records, score on supervised model varies from 0.929 to 0.9345. Finally, it has been observed that even if we generate a significantly large amount of data the accuracy score does not go down significantly. Originally, the f1-score is 0.9338 and the highest score we have is 0.9345 on 9000 and 20000 overall records. We also experiment with various architectures, including convolutional layers, and the results obtained are also good.

Next, we explore the use of GANs for the same task by first creating a generator and discriminator with three simple fully connected layers of 'LeakyReLU' activation. The generator generates data by taking noise as input and the discriminator predicts whether generated data is from the training set or not. We experiment with different optimizers (SGD, Adam and rmsprop), learning rate, batch normalization and dropout in both the discriminator and generator. We observe that we can generate data with the GAN but by applying lots of optimization techniques. We pass newly generated data to the same supervised model and data is classified correctly. After following the same methodology of generating 13000 new records, we receive the highest accuracy score of 0.9385 for 8000 records. When we change the architecture of discriminator and generator to include some convolutional layers, the results are not so good.

In conclusion, we experiment with VAEs and GANs for the problem of human activity data and are able to optimize both approaches to increase the performance of a HAR classifier.