

# ASSESSMENT OF THE PERFORMANCE OF DEEP LEARNING TECHNIQUES ON HUMAN ACTIVITY CLASSIFICATION BASED ON RADAR SPECTRAL IMAGES

**Devis DISHA, Linda SENIGAGLIESI and Ennio GAMBÌ**  
Università Politecnica delle Marche, Ancona, Italy

---

## ABSTRACT

Classification and recognition of human activities has gained a huge relevance in many fields, ranging from the well-being of the elderly to video surveillance. Many methodologies to improve the prediction of human activities, such as falls or unexpected behaviors, have been proposed over the years, exploiting different technologies. In the present investigation, a dataset related to different walking activities is considered, which includes both velocity and hands position of the subjects. Human activity classification is performed by extracting features from the spectral images obtained by an automotive radar. Different deep learning architectures and different optimizers are compared to assess performance and obtain reliable results. Results show that an appropriate selection of the network parameters provide an accuracy of more than 97%.

**Keywords:** deep neural networks, human activity classification, radar

## 1. INTRODUCTION

Automatic human activity classification is an active research field with application in various contexts, such as smart monitoring, video retrieval, ambient assisted living, surveillance, etc. Mainstream techniques rely on multiple sensors placed in different parts of the human body (Wu *et al.*, 2012), on sensors of smartphones and wearable devices (such as wristbands and watches) (Zhang *et al.*, 2015), while contactless technologies are assuming a growing relevance, due to their ability to provide a constant monitoring without the need of placing sensors in precise positions. Among these methodologies, radar technology has been recently applied to identify people on the basis of their gait characteristics (Vandersmissen *et al.*, 2018; Senigagliesi *et al.*, 2020) and gestures recognition (Ryu *et al.*, 2018; Wang *et al.*, 2020).

Activity classification is usually performed via machine learning and deep learning techniques (Wang *et al.*, 2019; Xing *et al.*, 2019), in most cases using wearable sensors (Ramanujam *et al.*, 2021). The rise of convolutional neural networks (CNN) has resulted in the development of categorization networks such as AlexNet (Krizhevsky *et al.*, 2012) and VGG (Simonyan and Zisserman 2014). AlexNet that is the first DNN (deep neural network) published in 2012 significantly improved identification accuracy (10% higher) in comparison to traditional approaches on ImageNet's 1000 class. Since then, literature has focused on both, creating networks accurate and designing efficient in terms of computational cost. However, there is a lot of literature that discuss new architectures in terms of layer composition performance. There are few papers (Bianco *et al.*, 2018) which evaluate factors related to computational cost like execution time, memory usage, etc., and more relevantly, how computational efficiency impacts precision.

In the present paper, different experiments are made to assess the benefits and drawbacks of the application of several existing deep neural networks (DNNs) on human activity classification. In particular, we consider the selection of 6 DNNs originally developed for image recognition and we evaluate their performance on images derived from the raw signal of an automotive FMCW radar and then processed to obtain range- Doppler images as final output. Classification results are obtained considering the dataset in (Gambi *et al.*, 2020), made up of 171 images associated to three different activities, i.e., fast walk, slow walk and slow walk with hands in pockets. Two of these activities are very similar, so they are sometimes merged in order to achieve better results in terms of accuracy and loss functions. The dataset has been created with participants with different characteristics, such as age, gender, height and weight. Subjects walked back and forth the radar used in the experiments within a distance of 15 meters. Similar comparative studies have been developed in (Maeda-Gutierrez *et al.*, 2020; Chellapandi *et al.*, 2021) on a very large dataset of plant diseases, while we test DNNs for human activity classification, which involves the resolution of different issues, starting from the small dimension of our dataset.

The rest of the paper is organized as follows. Section 2 informs about the hardware and software used in experiments. In addition, the DNN architectures have been introduced. The performance metrics and the described results are finally presented in Section 3, while the final considerations and remarks are provided in Section 4.

## **2. OVERVIEW OF CNNs**

In this section we provide a brief description of the CNNs architectures considered for our experiments. Because of the small dimension of the

dataset in (Gambi *et al.*, 2020), transfer learning is applied (Shu, 2019). TensorFlow package 2.3.0 is used to process neural networks with CUDA-V11.0 and cuDNN-v7.1 as back-end.

A CNN architecture is made up of different layers, namely, convolutional layer, pooling layer, reLU layer, fully connected layer, and loss layer. The convolutional layer is the main component of a CNN, comprising of filters (or *kernels*) that identify and forward various sorts of characteristics from the input. A pooling layer is included between consecutive convolutional layers to minimize the parameters in the network. The ReLu layer is a function that converts the negative numbers to zero.

We choose several types of networks, some of which are meant to be even more effective in terms of performance related to our small dataset. In certain situations, the number that appears with the name of architecture represents the number of layers containing parameters to be learnt (e.g., convolutional, or fully connected layers). The considered architectures are the following:

1. ResNet50
2. Inception ResNetV2
3. InceptionV3
4. VGG16
5. VGG19
6. Xception

A general overview of the different architectures is provided in the following, along with a brief description of the optimizers considered in our simulations.

The same sampling policies are applied with all the selected networks in order to have a direct fair and precisely comparison. InceptionResNet-v2, Inception-v3 and Xception models require normalized images with 229 pixels, while for all the other models 224 pixels are used. Our study is focused on accuracy rate, model complexity and inference time.

#### *VGG*

VGG16 (Simonyan and Zisserman, 2014) consists of 16 levels divided into convolutional layers, max pooling layers and fully linked layers. It has 5 blocks and a maximum pooling layer in each block. Similarly, VGG19 contains 19 layers, with three additional convolution layers in the last three blocks. Both VGG16 and VGG19 excel in the image recognition thanks to their deep layers.

#### *InceptionV3*

Deep convolutional architecture Inception V3 is frequently used for classification problems. Szegedy *et al.*, (2015) first suggested the model notion in the Google Net architecture where this model is proposed by upgrading the inception module. Each block of the Inception V3 network

contains many branches of convolutions, average pooling, max pooling, concatenated, dropouts, and fully connected layers, with each block having several symmetric and asymmetric construction blocks. The network has 42 layers and 29.3 million parameters; thus, its computing cost is only around 2.5 times that of GoogleNet. The combination of a reduced parameter count and additional regularization with batch-normalized auxiliary classifiers and label-smoothing allows for training high quality networks on small training sets.

#### *Residual Network (ResNet)*

He *et al.*, (2016) developed the ResNet models, which are based on deep architectures that have demonstrated strong convergence tendencies and convincing correctness. ResNet was created using many stacked residual units and a variety of layers. However, depending on the design, the number of operations can be changed. Convolutional, pooling, and layers make up the residual units for all the above. ResNet is comparable to VGG net but it is eight times deeper. The ResNet 50 network has 49 convolutional layers and a fully linked layer at the end.

#### *Xception*

Xception (Chollet, 2017) is based on depth-wise separable convolutional layers. This neural network architecture has 36 convolutional layers. So, the feature extraction base of the network is made up. Except for the first and the last modules, all the 36 layers are divided into 14 modules, all of which have linear residual connections surrounding them. In summary, the Xception design is a depth-wise separable convolution layer stack with residual connections, allowing the architecture to be easily modifiable and define simply using high-level libraries.

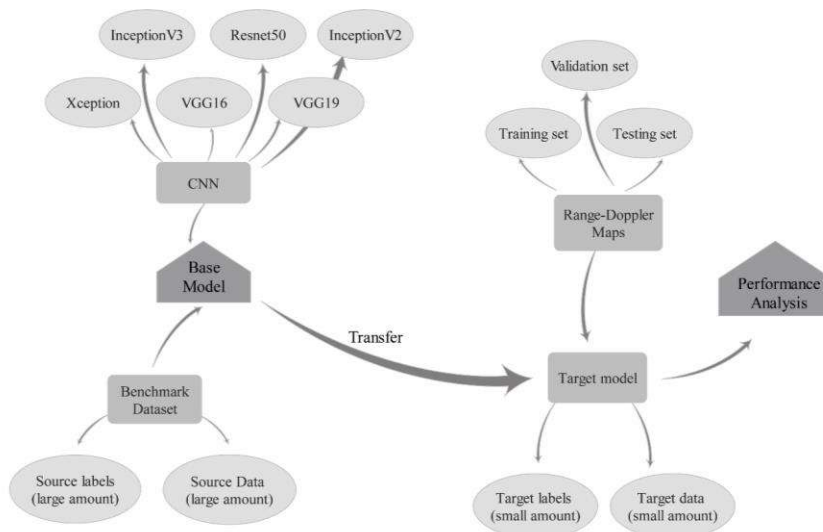
#### *Optimizers*

Gradient descent is the method most used to improve the learning of deep neural networks. It basically updates each model parameter, verifying how it affects the objective function, and iterates until this function converges to the minimum. Stochastic Gradient Descent (SGD) is a well-known variation of gradient descent. The Adaptive Gradient Algorithm (Adagrad) optimizer (Duchi *et al.*, 2011) is designed to deal with sparse data. It adapts the learning rate to the parameters, performing smaller updates (i.e., low learning rates) for parameters associated with frequently occurring features, and larger updates (i.e., high learning rates) for parameters associated with infrequent features. Adadelat is instead an extension of Adagrad developed in (Zeiler 2012), simultaneously to Root Mean Square Propagation (RMSprop), with the aim of solving Adagrad's diminishing learning rates. Adaptive Moment Estimation (Adam) (Kingma 2014) is another method that computes adaptive learning rates for each parameter.

## PROPOSED METHODOLOGY

As aforementioned said, the dataset in (Gambi *et al.*, 2020) is considered. Dataset was built at Università Politecnica delle Marche, Ancona, Italy. Different subjects of age, gender, height and weight participate in the tests, repeating each activity (slow walk, fast walk and walk with the hands in pockets) for three times. A total of 171 images of 224 x 224 pixels is built. The raw data obtained from the radar are processed to obtain range-doppler maps and feature detection is applied to smooth the signals.

The proposed methodology is proposed in the Figure 1. All the images are colored (RGB). Classification is performed by exploiting the range-doppler maps. Data augmentation is used to pre-process the range-doppler images to remove noise or redundancy. Also, it increases accuracy and rush the execution. For this purpose, Keras deep learning library in Python is used. When a large amount of data is available, a model is built from scratch to solve a specific problem. Differently, when the dataset dimension is small, transfer learning is usually applied, and a model built by training on a larger dataset can be used as a starting point to train the n`etwork on our specific data.



**Fig. 1:** Conceptual model of the proposed methodology.

Numerous pre-trained models are available under certain conditions. For our implementation we used six pre-trained models, namely, VGG16,

VGG19, InceptionResNetV2, InceptionV3, ResNet50, Xception. We fine-tuned all the models on the last layer. All the 171 images are randomly generated using ImageDataGenerator. Training and testing image dataset used the same data generators. These ImageDataGenerators apply the pre-processing while generating the images. We also use the label encoding to get the categorical output. When the validation accuracy stops improving, an earlyStopping is used to terminate the training process. The underlying model is then combined by a GlobalAveragePooling layer, which is used to reduce the data and prepare the model for the final classification. After that, a BatchNormalization layer is introduced for model stability and to have a faster execution, resulting in fewer training epochs. The last layer exploits a dense layer and a SoftMax activation function for extracting and categorizing the output. The weights are updated using different optimizers with an activation function of softmax in the final layer for classifying output. Optimizers are used to update the weights. The performance of the network is measured using a categorical loss function. A dropout is added to avoid overfitting.

After verifying that a number greater than 100 epochs leads to overfitting, the models are trained using a size batch of 50 and 100 epochs on the training set and then tested on the test set. At the end, we evaluate the performance of the chosen models as a function of different parameters.

### 3. EXPERIMENTS AND RESULTS

We here present a comparison of the results obtained by applying different neural networks and different optimizers to the dataset in (Gambi *et al.*, 2020). For the sake of reproducibility, we select a random seed equal to 1337. This choice does not influence the results, since they are averaged over multiple random selections. The dataset then is randomly split into 80% for training, 10% for validation and 10% for test. This is done to prevent results from repeating each time the algorithm is rerun.

We first define the metrics which will be used to evaluate the performance of the proposed methodologies. They are based on the so-called *confusion matrix*, which columns represent the predicted values for each class, while the rows represent the real values.

The most common metric is the accuracy, which is defined as

$$Acc = \frac{TP+TN}{TP+TN+FP+FN} \quad (1)$$

where *TP* and *TN* stand for true positives and true negatives, while *FP* and *FN* are the false positives and the false negatives.

A loss function measures the probabilities or uncertainty of a prediction based on how much the prediction varies from the true value. Differently from accuracy, the loss is computed as summation of the errors made for each sample in training or validation sets, and is not evaluated as a percentage. Loss is often used in the training process to find the “best” parameter values for the model (e.g., weights in neural network). During the training process the goal is to minimize this value. In the following we consider the categorical cross-entropy loss.

Another useful metric is the precision (also called positive predictive value), which corresponds to the ratio of correctly predicted positive observations to the total predicted positive observations, or

$$P = \frac{TP}{TP+FP} \quad (2)$$

Recall (also known as sensitivity) is the ratio of correctly predicted positive observations to the allobservations in actual class:

$$R = \frac{TP}{TP+FN} \quad (3)$$

The F1 score is the harmonic mean of the precision and recall, where an F1 score reaches its best value at 1 (perfect precision and recall). Therefore, this score takes both false positives and false negatives into account as follows:

$$F1 = \frac{2RP}{R+P} \quad (4)$$

The accuracy and loss for two and three activities are in Table 1 and 2 respectively reported. Different optimizers are selected for each type of CNN. In the two-activities test we have an imbalance between the folders because the slow walk and slow walk with hands in pockets classes are merged. This gives us some fluctuation in the results of the training and validation loss. It is possible to observe that the Inception family of algorithms leads to the best results for the two activities case, especially when selecting the RMSprop optimizer. The small dimension of the dataset deeply affects the results when three different activities are considered, as evident from the high values of the loss function. This reflects the fact that there are not enough images to perform a correct classification, and the model easily tends to overfitting. An accuracy of more than 86% is however achieved by InceptionV3 network is SGD optimizer, with an overall contained loss with respect to the other cases.

Table 3 reports the precision and recalls values and the F1 score for three CNNs, i.e., VGG16, ResNet50 and Inception ResNetV2. We select the optimizers which give the best results and correspond to Adagrad, RMSprop and Adam, respectively. As there is no CNN or optimizers that might result

the best in all cases, a suitable configuration must be selected depending on the scenario and the parameters that want to maximize.

**Table 1.** Accuracy and loss for two activities. Accuracy values are reported as percentages.

Model	VGG16		ResNet50		Inception ResNetV2		InceptionV3		Xception		VGG19	
	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS
Adam	93.8	0.17	85.7	0.32	95.9	0.63	98	0.11	94	0.13	93.8	0.29
RMSprop	93.8	0.17	81.6	0.36	98	0.06	98	0.05	98	0.07	93.8	0.25
SGD	87.5	0.27	67.4	0.54	89.8	0.22	95.9	0.16	91.8	0.20	81.3	0.37
Adadelta	81.3	0.48	67.4	0.51	89.8	0.27	87.8	0.25	89.8	0.25	81.3	0.50
Adagrad	93.8	0.19	89.8	0.34	98	0.04	95.9	0.06	98	0.06	93.8	0.27
Adamax	93.8	0.7	89.8	0.40	98	0.07	98	0.05	95.9	0.14	93.8	0.29

**Table 2.** Accuracy and loss for three activities. Accuracy values are reported as percentages

Model	VGG16		ResNet50		Inception ResNetV2		InceptionV3		Xception		VGG19	
	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS	ACC	LOSS
Adam	60	0.66	53.3	0.74	66.7	2.2	66.7	0.51	66.67	3.44	66.67	0.75
RMSprop	66.7	0.74	53.3	0.91	80	1.1	86.7	1.49	80	1.70	73.3	0.68
SGD	60	0.94	33.3	1.08	73.3	0.66	86.7	0.56	60	0.85	53.3	0.89
Adadelta	60	1.01	53.3	0.99	53.3	0.77	60	0.65	60	0.77	60	0.97
Adagrad	53.3	0.68	60	0.82	80	0.77	66.7	0.66	73.3	0.87	53.3	0.74
Adamax	80	0.67	46.6	0.82	80	0.90	86.7	2.50	66.7	1.30	60	0.77

**Table 3.** Precision, recall and F1 score for the FastWalk and SlowWalk activities

Model	VGG16			ResNet50			Inception ResNetV2		
	P	R	F1	P	R	F1	P	R	F1
FastWalk	1	0.43	0.6	1	0.57	0.73	0.6	0.43	0.5
SlowWalk	0.75	1	0.86	0.8	1	0.89	0.71	0.83	0.77

Finally, in Table 4 and 5 we compare the execution times required by each algorithm. Algorithms have been tested using a Windows 10 Pro 64-bit



(10.0, Build19042.789), Version 20H2, Intel Core I7-5500U, CPU @ 2.40Ghz, RAM 16 GB, Display NVIDIA GeForce 920M 4GB. Here, there are no evident differences between the cases with two or three activities. In both scenarios SGD takes less time than other optimizers, while VGG19 results the fastest network. SGD in fact is a variation of the classical gradient descent, but it only computes a tiny subset or random selection of data instances, rather than the entire dataset, which is redundant and wasteful. In this way SGD needs less iterations until the objective function converges to the minimum, thus resulting faster than other optimizers, and when the learning rate is modest, it is able to achieve the same results than traditional gradient descent. Adam is instead a gradient-based optimization technique for stochastic objective functions. It computes specific adaptive learning rates for distinct parameters by combining the benefits of two SGD extensions, RMSProp and AdaGrad. Although it has a great popularity, Adam has recently appeared to fail to converge to an optimal solution in some situations, thus requiring in some cases larger execution times.

**Table 4.** Execution times required by different algorithms for the three activities case. Time is expressed in seconds

Model	VGG16	ResNet50	Inception ResNetV2	InceptionV3	Xception	VGG19
Adam	30	100	38	55	107	21.04
RMSprop	52	195	75	93	150	42.4
SGD	19	68	24	30	47	17

**Tab. 5** Execution times required by different algorithms for the two activities case. Time is expressed in seconds.

Model	VGG16	ResNet50	Inception ResNetV2	InceptionV3	Xception	VGG19
Adam	31.2	196	45.7	57.2	114	24
RMSprop	56.1	104	89	117	215	60
SGD	19	86	39	42.1	73	24

#### 4. CONCLUSIONS

Deep learning is a branch of artificial intelligence that is widely applied in different fields. We have considered human activity classification performed by exploiting images derived from an automotive radar, which represents an emerging contactless technology. We have shown how different pre-trained

models built using transfer learning can achieve good results to discriminate different gait velocities. We have provided a comparison between different neural networks, in terms of accuracy, loss, classification parameters and execution times. The presented results might further improve by expanding the dataset dimension and including the tracking of the subjects and trajectory.

## REFERENCES

**Bianco S, Cadene R, Celona L, Napoletano P. 2018.** Benchmark analysis of representative deep neural network architectures. *IEEE Access*, **6**: 64270–64277. doi: 10.1109/ACCESS.2018.2877890.

**Chellapandi B, Vijayalakshmi M, Chopra S. 2021.** Comparison of pre-trained models using transfer learning for detecting plant disease. *2021 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)*, pp. 383-387, doi: 10.1109/ICCCIS51004.2021.9397098.

**Chollet F. 2017.** Xception: Deep learning with depth-wise separable convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, doi: 10.1109/CVPR.2017.195.

**Duchi J, Hazan E, Singer Y. 2011.** Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* **12(7)**: 2121-2159.

**Gambi E, Ciattaglia G, De Santis A, Senigagliesi L. 2020.** Millimeter wave radar data of people walking. Data in Brief 31, 105,996. doi: 10.1016/j.dib.2020.105996. eCollection 2020 Aug.

**He K, Zhang X, Ren S, Sun J. 2016.** Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778 <https://doi.org/10.48550/arXiv.1512.03385>.

**Kingma DP, Ba J. 2014.** A method for stochastic optimization. arXiv preprint, arXiv:1412.6980.

**Krizhevsky A, Sutskever I, Hinton GE. 2012.** ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, **25**: 1097–1105.

**Maeda-Gutierrez V, Galvan-Tejada CE, Zanella-Calzada LA, Celaya-Padilla JM, Galván-Tejada JI, Gamboa-Rosales H, Luna-García H, Magallanes-Quintanar R, Guerrero Mendez CA, Olvera-Olvera, CA. 2020.** Comparison of convolutional neural network architectures for classification of tomato plant diseases. *Applied Sciences*, **10(4)**: 1245.

**Ramanujam E, Perumal T, Padmavathi S. 2021.** Human activity recognition with smartphone and wearable sensors using deep learning

techniques: A Review. *IEEE Sensors Journal*, **21(12)**:13029-13040, doi: 10.1109/JSEN.2021.3069927.

**Ryu SJ, Suh JS, Baek SH, Hong S, Kim JH. 2018.** Feature-based hand gesture recognition using an fmcw radar and its temporal feature analysis. *IEEE Sensors Journal*, **18(18)**: 7593–7602.

**Senigagliesi L, Ciattaglia G, De Santis A, Gambi E. 2020.** People walking classification using automotive radar. *Electronics*, **9(4)**: 588.

**Shu M. 2019.** Deep learning for image classification on very small datasets using transfer learning. *Creat. Components, Jan.* [content \(iastate.edu\)](https://content.iastate.edu) Corpus ID: 204089661.

**Simonyan K, Zisserman A. 2014.** Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556. [1409.1556.pdf \(arxiv.org\)](https://arxiv.org/abs/1409.1556).

**Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A. 2015.** Going deeper with convolutions. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1–9, doi: 10.1109/CVPR.2015.7298594.

**Vandersmissen B, Knudde N, Jalalvand A, Couckuyt I, Bourdoux A, De Neve W, Dhaene T. 2018.** Indoor person identification using a low-power FMCW radar. *IEEE Transactions on Geoscience and Remote Sensing*, **56(7)**: 3941–3952.

**Wang J, Chen Y, Hao S, Peng X, Hu L. 2019.** Deep learning for sensor-based activity recognition: A survey. *Pattern Recognition Letters*, **119**: 3-11.

**Wang Y, Ren A, Zhou M, Wang W, Yang X. 2020.** A novel detection and recognition method for continuous hand gesture using fmcw radar. *IEEE Access*, **8**: 167264–167275, 10.1109/ACCESS.2020.3023187.

**Wu W, Dasgupta S, Ramirez EE, Peterson C, Norman GJ. 2012.** Classification accuracies of physical activities using smartphone motion sensors. *Journal of medical Internet research*, **14(5)**: e2208.

**Xing Y, Lv C, Wang H, Cao D, Velenis E, Wang FY. 2019.** Driver activity recognition for intelligent vehicles: A deep learning approach. *IEEE transactions on Vehicular Technology*, **68(6)**: 5379-5390.

**Zeiler MD. 2012.** Adadelta: an adaptive learning rate method. arXiv preprint arXiv:1212.5701.

**Zhang L, Wu X, Luo D. 2015.** Recognizing human activities from raw accelerometer data using deep neural networks. In: 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA), pp. 865–870.

