Supplementary Information for:

Machine Learning and image analysis towards improved energy management in Industry 4.0: a practical case study on quality control

Mattia Casini[a], Paolo De Angelis[a], Marco Porrati[b], Paolo Vigo[a], Matteo Fasano[a], Eliodoro Chiavazzo[a], Luca Bergamasco[a,*]

[a] *Department of Energy, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy*
[b] *R&D Department, GEFIT S.p.a., Via Giovanni de Negri 9, 15121 Alessandria, Italy*

Choice of the model

In choosing the neural network models to be used for our case study, the first objective was to identify the appropriate complexity of the model to be used. The first comparison was made between a traditional CNN model (VGG-16), and a residual CNN model (ResNet50). Using the same training dataset and number of epochs, a direct comparison between the two models was made. Figure S1 and S2 show the evolution of the learning curves for the models during training. The performance of the VGG-16 model was found to be insufficient for our purposes (due to under-fitting and accuracy close to a dummy classifier). On the other hand, the ResNet50 showed an accuracy approaching 80-85%.

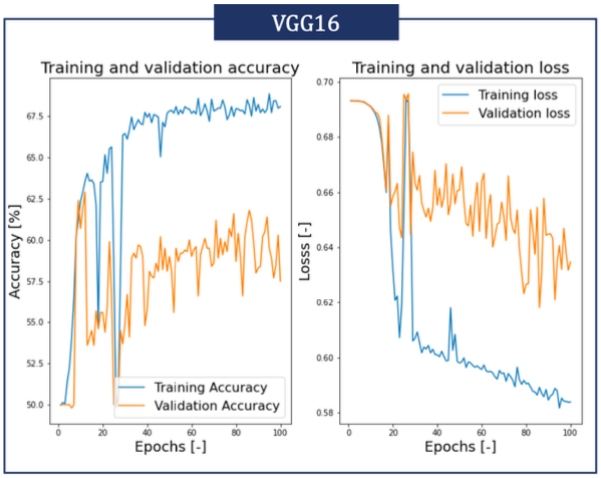

Figure S1: Training of the VGG-16 model

After the first test, the ResNet model was identified as a possible better candidate with respect to the VGG-16 model. We then proceeded to additionally test different depths of the ResNet model (50, 101 and 152) and other models with similar complexity but slightly different working principle: the Inception V1 and V4 model.

*Corresponding author: luca.bergamasco@polito.it

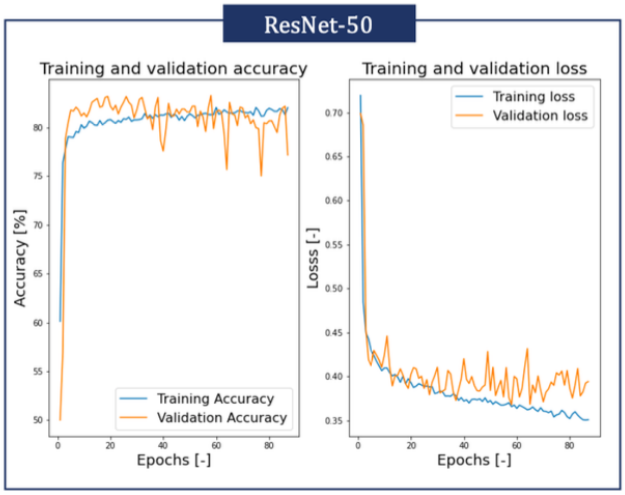

Figure S2: Training of the ResNet50 model

New training tests were then carried at extended number of epochs for the ResNet 50, ResNet 101, ResNet 152, Inception V1 and Inception V4 models. All models initially showed over-fitting problems; thus, regularization was adopted to improve the performance. The two models with the best performance per architecture (ResNet 101 and Inception V4) were chosen as suitable candidates and their performance additionally assessed. The two models were found to perform very similarly (see Fig. 6 in the main manuscript and Fig. S3), therefore we finally decided to adopt the ResNet 101 as it has a lower number of parameters to fit, and thus a faster response time.

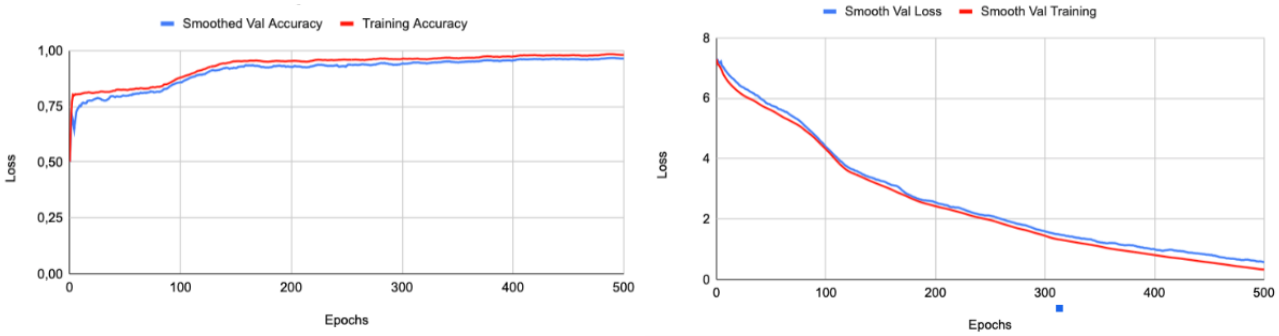

Figure S3: Final training of the InceptionV4 model: accuracy (left); loss (right)