Batch Normalization Tells You Which Filter Is Important
– *Supplementary Material* –

Junghun Oh Heewon Kim Sungyong Baik Cheeun Hong Kyoung Mu Lee
Department of ECE, ASRI, Seoul National University
{dh6dh, ghimhw, dsybaik, cheeun914, kyoungmu}@snu.ac.kr

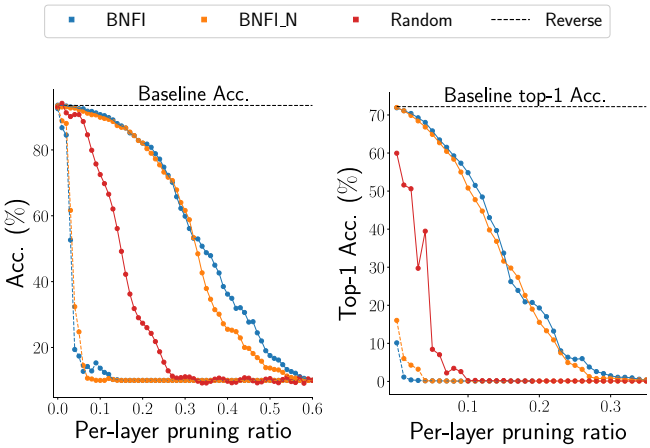

(b) Result on VGGNet-16 (c) Result on MobileNetV2

Figure A: Impact of sparse activation on the filter importance.

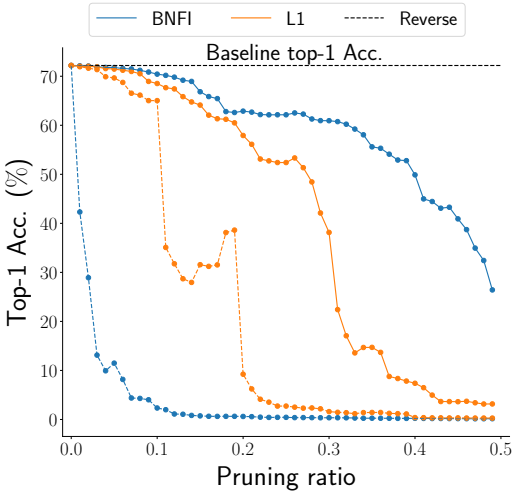

Figure B: Pruning results on the pre-trained MobileNetV2 without fine-tuning. For quantitative support on the results in Figure C, we prune the filters of the 3x3 convolutional layer in the 2nd residual block.

A. Filter importance on sparse activation

In this section, we provide more detailed explanations for Section 3.2. We claim that the proposed filter importance in Equation (7) should be slightly modified when the activation function is likely to induce sparse activation outputs (*e.g.* the activation function maps a certain range of input values to zero as in ReLU). The reason is that when there are many zero activation values, they can dominate the non-zero values, reducing the averaged value of the overall activation outputs. In Equation (8), we eliminate the impact of the zero activation values by dividing the probability of zero activation values. To support our claim, in Figure A, we measure and compare the without-fine-tuning performance of networks that are pruned according to filter importance that either considers the impact of zero activation values (BNFI) or not (BNFI_N). The results show that BNFI is slightly better than BNFI_N, demonstrating that eliminating the impact of zero activation values is helpful for more accurate filter importance when there is a chance for the high sparsity.

B. Feature map visualization

To further demonstrate the effectiveness of our method, we visualize activation output channels sorted by the filter importance of BNFI, which is shown in Figure C. The visualization results from BNFI seems to be sorted well in an order of the magnitude of activation compared to L1. For the input image of a dog, for example, the highlighted channel (by red boxes), which is considered to be the most important activation channel in BNFI, is considered relatively less important by L1. The results are quantitatively supported by the results shown in Figure B. These results seem to suggest that BNFI is a more accurate importance measure than L1, one of widely-used filter-weight-based methods.

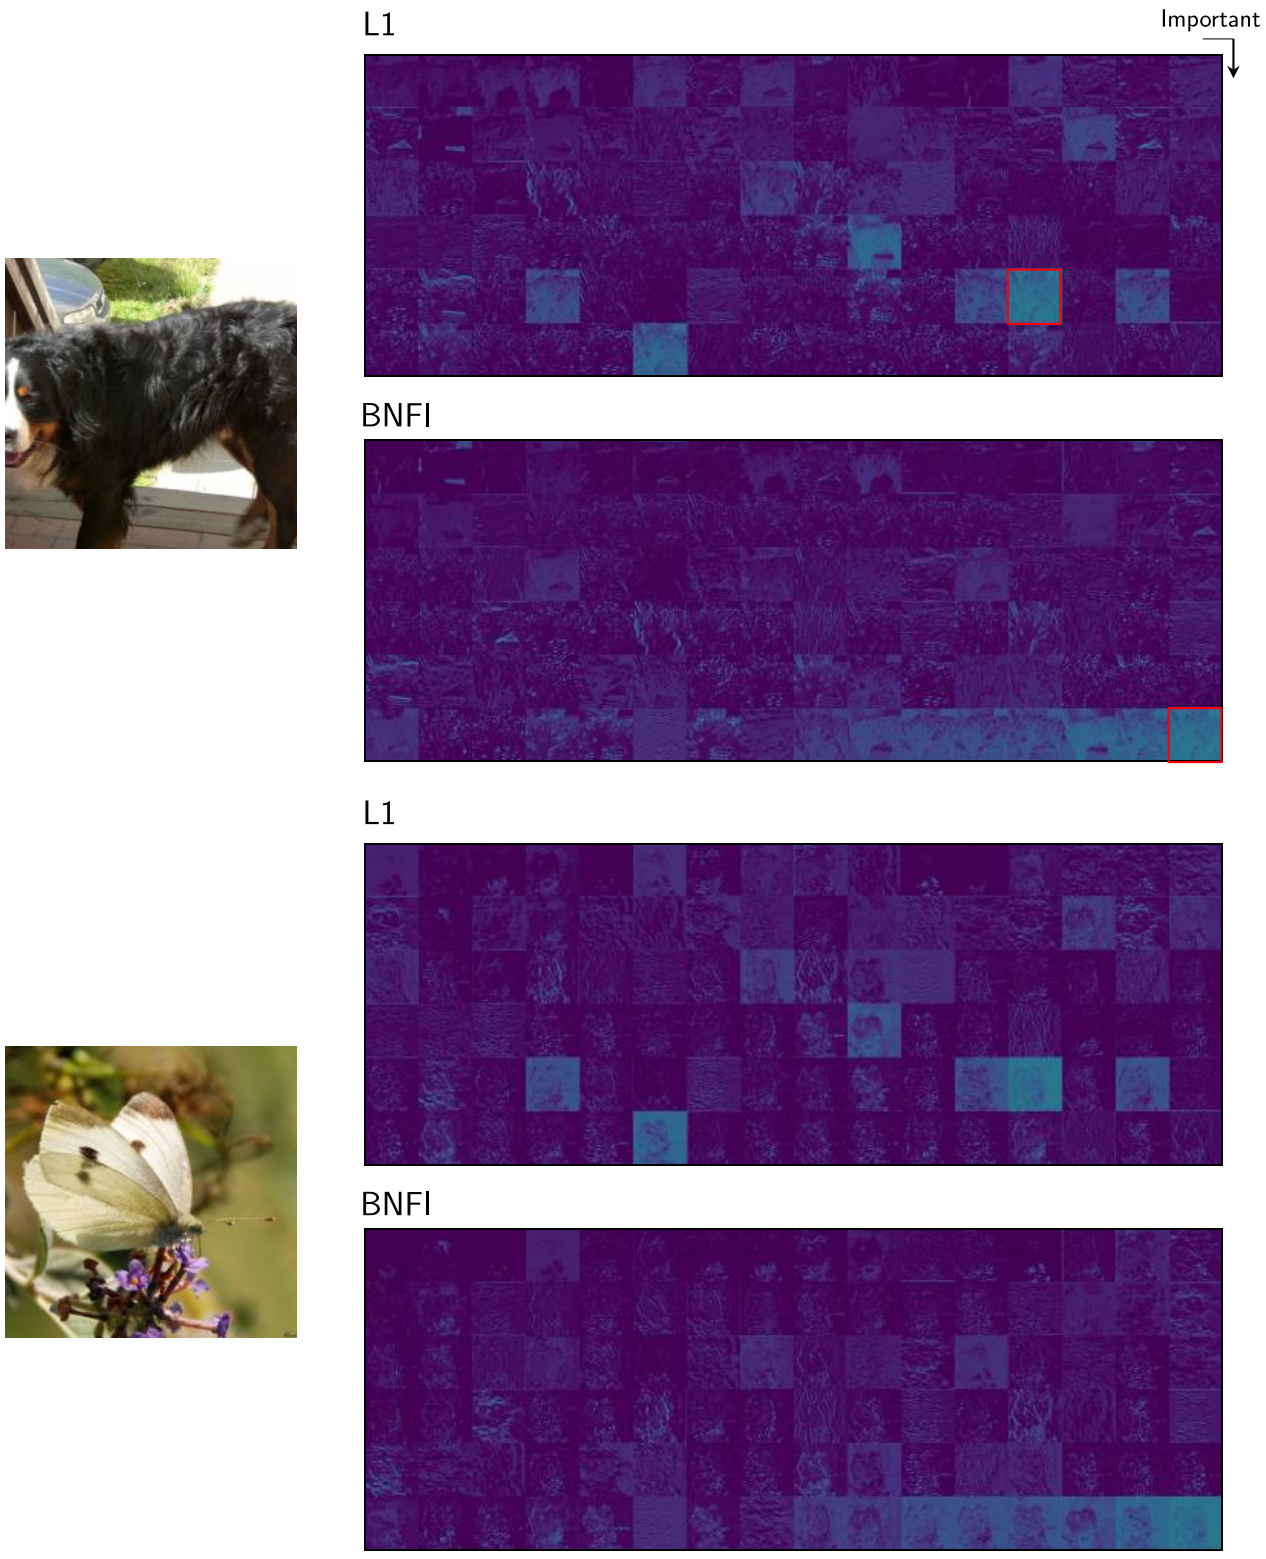

Figure C: Visualization results on the pre-trained MobileNetV2. We present the activation outputs after the 3x3 convolutional layer in the 2nd residual block. The input images are shown in left side and the corresponding activation channels are shown in right side, which are sorted by BNFI or L1. More important channels are located on the lower right side.