Supplementary material

It's DONE: Direct ONE-shot learning with Hebbian weight imprinting

This file includes

- Figures S1, S2, and S3

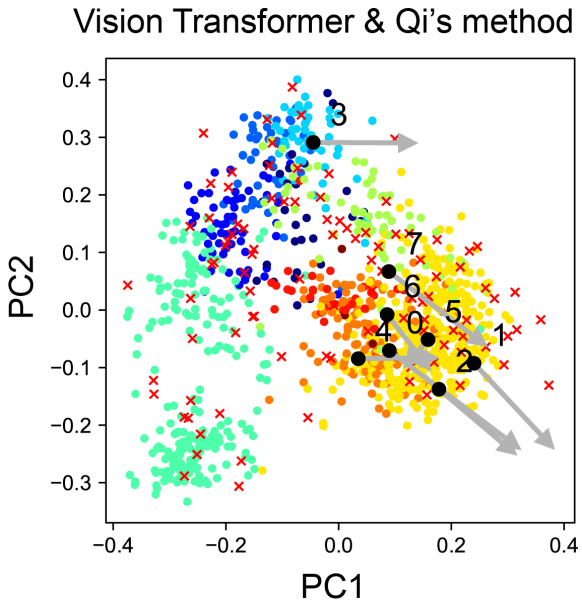

Figure S1. Principal component analysis of weight vectors (Figure 4 in the manuscript) in the case of Qi's method with ViT.

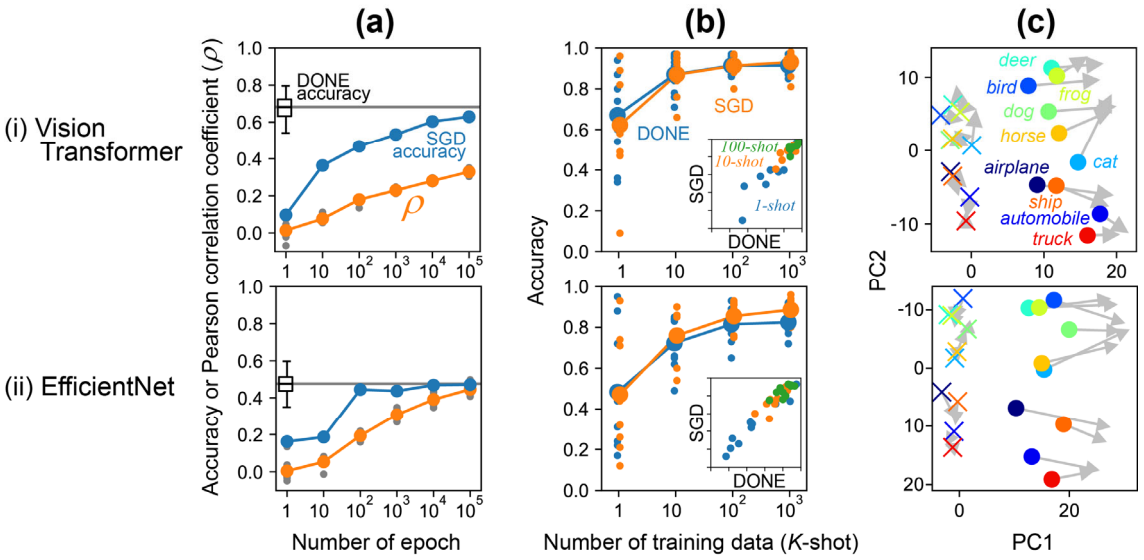

Figure S2. Transfer learning tasks of DONE. (a) Results of one-shot learning with DONE or SGD. (b) The results of *K*-shot learning. The small circle is the accuracy for each class, and the large circle is the average value. The inset graph is the relationship between the accuracy of each class obtained by DONE and SGD. (c) PCA of w_j vectors. Circles and crosses show the results obtained by one-shot learning with DONE and SGD, respectively. The color indicates the class. The initial and terminal points of the grey arrows indicate the results of one-shot and 1000-shot learning, respectively.

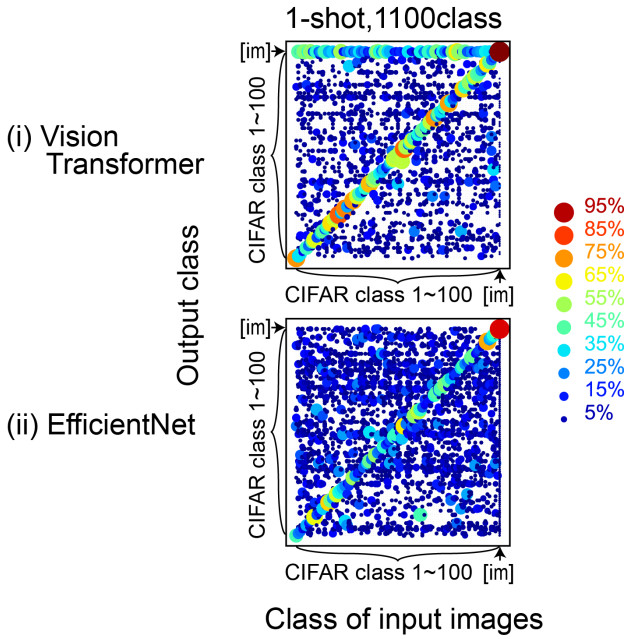

Figure S3. 100-class addition by 1-shot learning for investigating a scalability. We used images that showed median accuracy at one-class addition task such as shown in Figure 2(b) as a standard training image of each class. The horizontal and vertical axes show the class of the input images, and the output class, respectively. The color and size of the circles indicate the percentage of the output class in the input-class images. The class ID of CIFAR-100 was used. The class [im] contains 1000 classes of ImageNet. The accuracy of the 1100-class models for the classification of all 10,000 test images of CIFAR-100 was 37.2% and 26.3% (note that the chance level was not 1/100 but 1/1100), with decreasing the accuracy for the classification of all 50,000 ImageNet validation images by 0.1% and 2.9% from the original 1000-class model (65% and 69%), in ViT and EfficientNet, respectively. It is not clear whether these performance are at a practical level, but in any case, this task (1-shot 100-class addition to 1000 class) is expected to be difficult for humans as well.