

# **Another BRIXEL in the Wall: Towards Cheaper Dense Features**

Supplementary Material

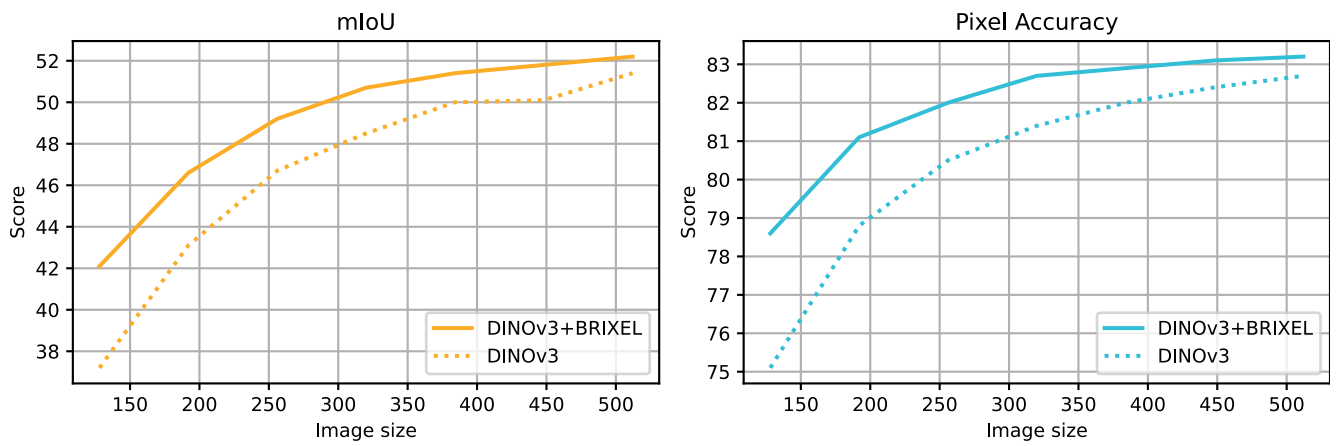


Figure 8. We evaluate the fine-tuned ViT-B BRIXEL model on semantic segmentation on ADE20k at a variety of input image sizes. BRIXEL outperforms the DINOv3 baseline at all image sizes, showing that the method generalizes well beyond the training image size.

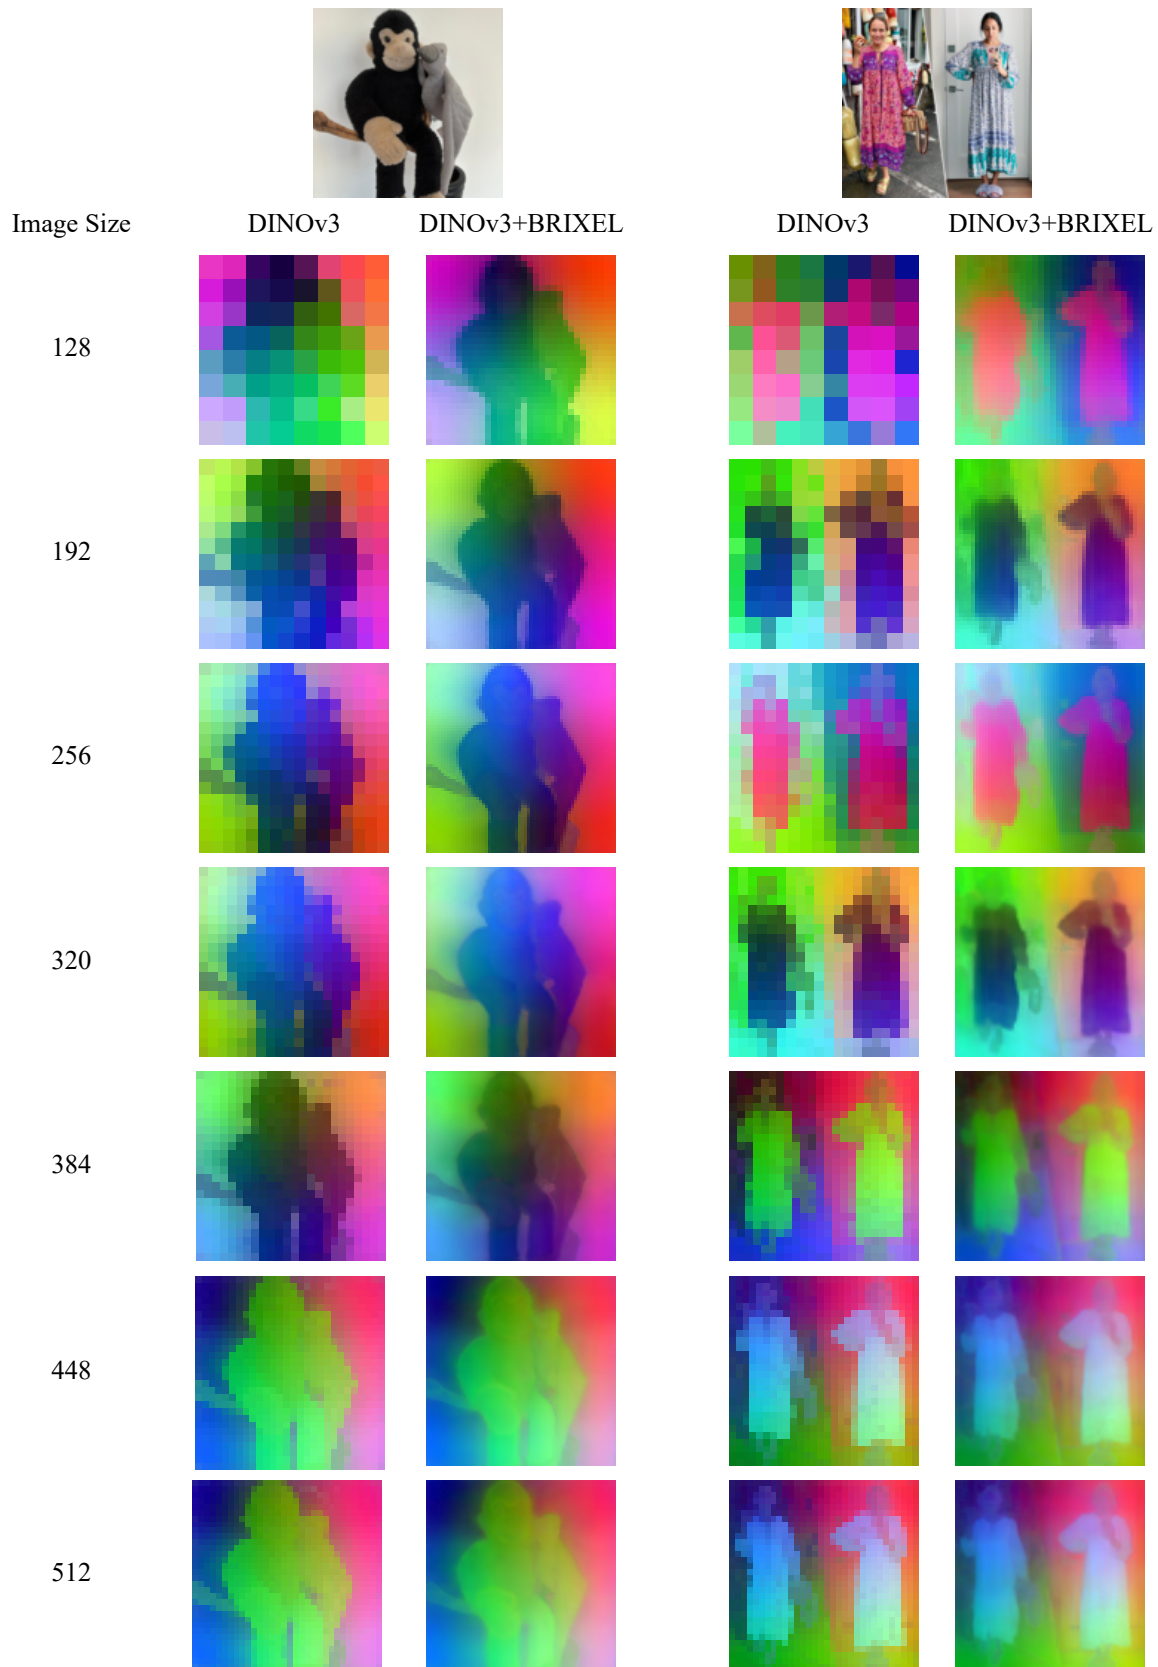


Figure 9. Feature maps of the fine-tuned ViT-B model for different input sizes. Different RGB maps across image sizes are due to the PCA being recomputed for each image size.