

6 Supplementary materials

This section contains the supplementary materials.

6.1 Source code

We share our implementation code to make it easy to reproduce our results. The source-code is attached to the supplementary materials in a ‘code’ directory. We also provide detailed instructions for training and evaluating our models in ‘README.md’ files.

6.2 Additional visualizations

Fig. 6 provides a visualization of object detection results of Table 3. We observe from this figure that in low data regimes model composition performs considerably better than supervised training with partial data. Fig. 7 shows an extended visualization on the cloud embodiment introduced in the paper. In this figure, we provide an easier comparison between before & after incorporating the model composition as a service. Moreover, Fig. 8 demonstrates an example of pseudo-label aggregation procedure of Algorithm 2. In addition, Fig. 9 visualizes the data splits of object detection scenario 4, where we combined 10 models trained on different COCO subsets.

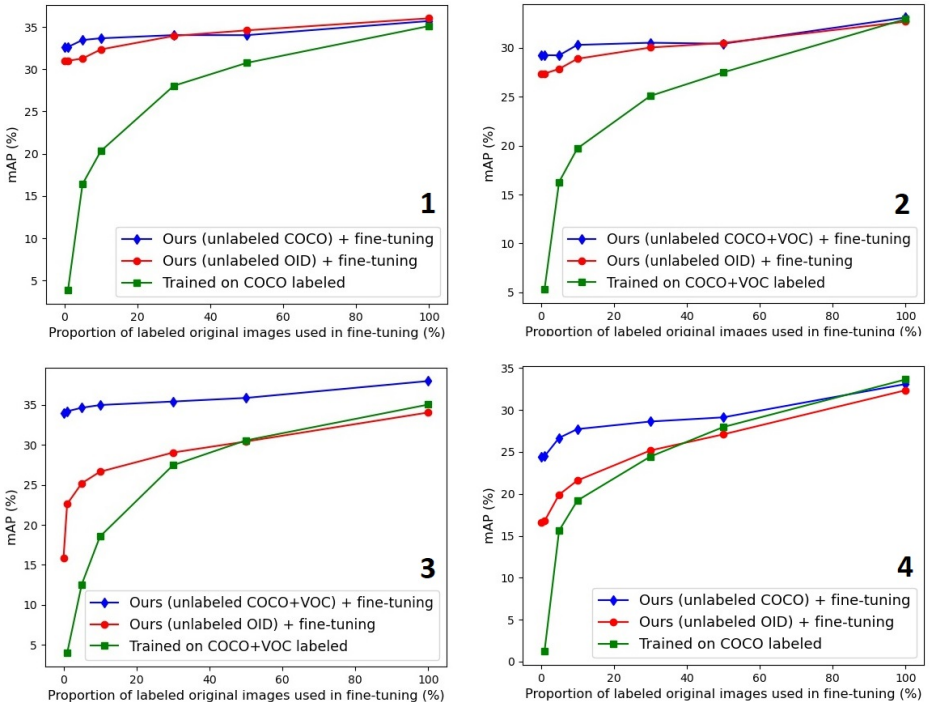


Figure 6: Results of different scenarios for fine-tuning. Each scenario is tagged with its number.

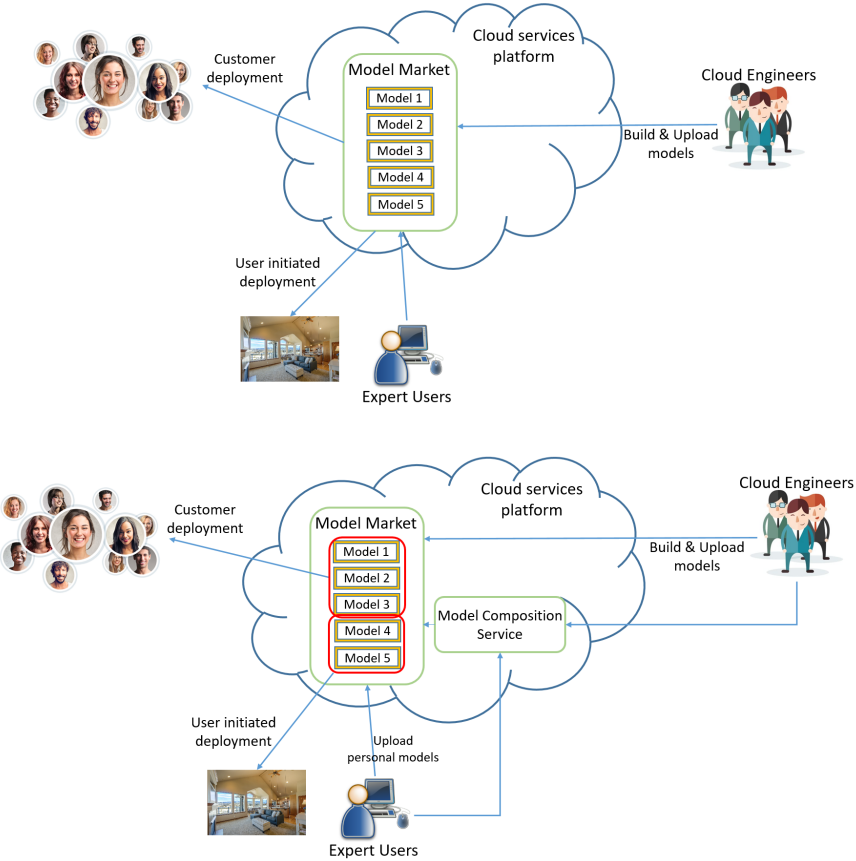


Figure 7: An embodiment of how our model composition strategy would be implemented in a cloud services platform. Top: existing system. Bottom: after applying model composition.

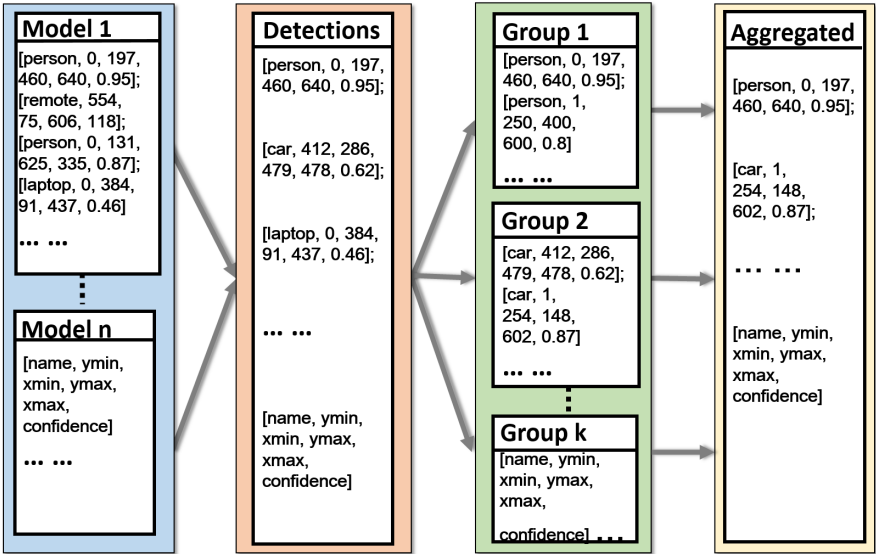


Figure 8: Pseudo-label aggregation flow diagram, an example.




Figure 9: Object categories and data size for models used in object detection scenario 4. Larger dots correspond to higher number of examples. 10 rows correspond to 10 different models.