FAST ADAPTATION OF SEGMENTATION MODELS FOR TIRE SIDEWALL SYMBOLS

Michal Balogh
Supervisor: **doc. Ing. Michal Španěl, Ph.D**

In collaboration with Micro-Epsilon Inspection, this work addresses the current industry solution, which requires approximately 72 hours to retrain when new symbols are introduced. My approach combines state-of-the-art segmentation architecture (OneFormer) with efficient fine-tuning techniques including Low-Rank Adaptation (LoRA) and mixed precision training. The results show comparable segmentation quality to the current solution while reducing adaptation time from 72 hours to less than one hour — in extreme cases, fine-tuning a single symbol in just 5 minutes.

introduction & problem statement

Tire sidewalls of tires contain important information through various symbols and text that must be accurately segmented for quality control. The current industrial solution uses Mask R-CNN, but requires extensive retraining when new symbols are introduced:

current state

- 72-hour retraining time for new symbols
- Industrial requirement: Fast adaptation to new symbols while maintaining segmentation quality

dataset

- Depth images of tire sidewalls (65,536 × 1,300 pixels)

used hardware

- GPU with 16GB memory for training base model, 8GB for fine-tuning on new symbols

model architecture

OneFormer with Swin Transformer backbone

- Unified framework for semantic, instance, and panoptic segmentation
- Task-conditional joint training capability

optimization techniques

- Mixed precision training
- Low-Rank Adaptation (LoRA)
- Selective freezing of network components
- Kaiming uniform weight initialization

fine-tuning strategy

- Using 10% of model parameters
- Mixed dataset approach (blending new and original data)
- Preventing catastrophic forgetting through balanced sampling

conclusions

- OneFormer architecture with instance segmentation delivers high-quality masks for tire sidewall symbols
- Mixed precision training accelerates the process by ~27%
- Fine-tuning with 1:6 ratio of new:original data prevents catastrophic forgetting
- Selective parameter updating provides the best balance between adaptation speed and performance
- LoRA can achieve comparable results while training fewer parameters

new:old dataset ratio comparison

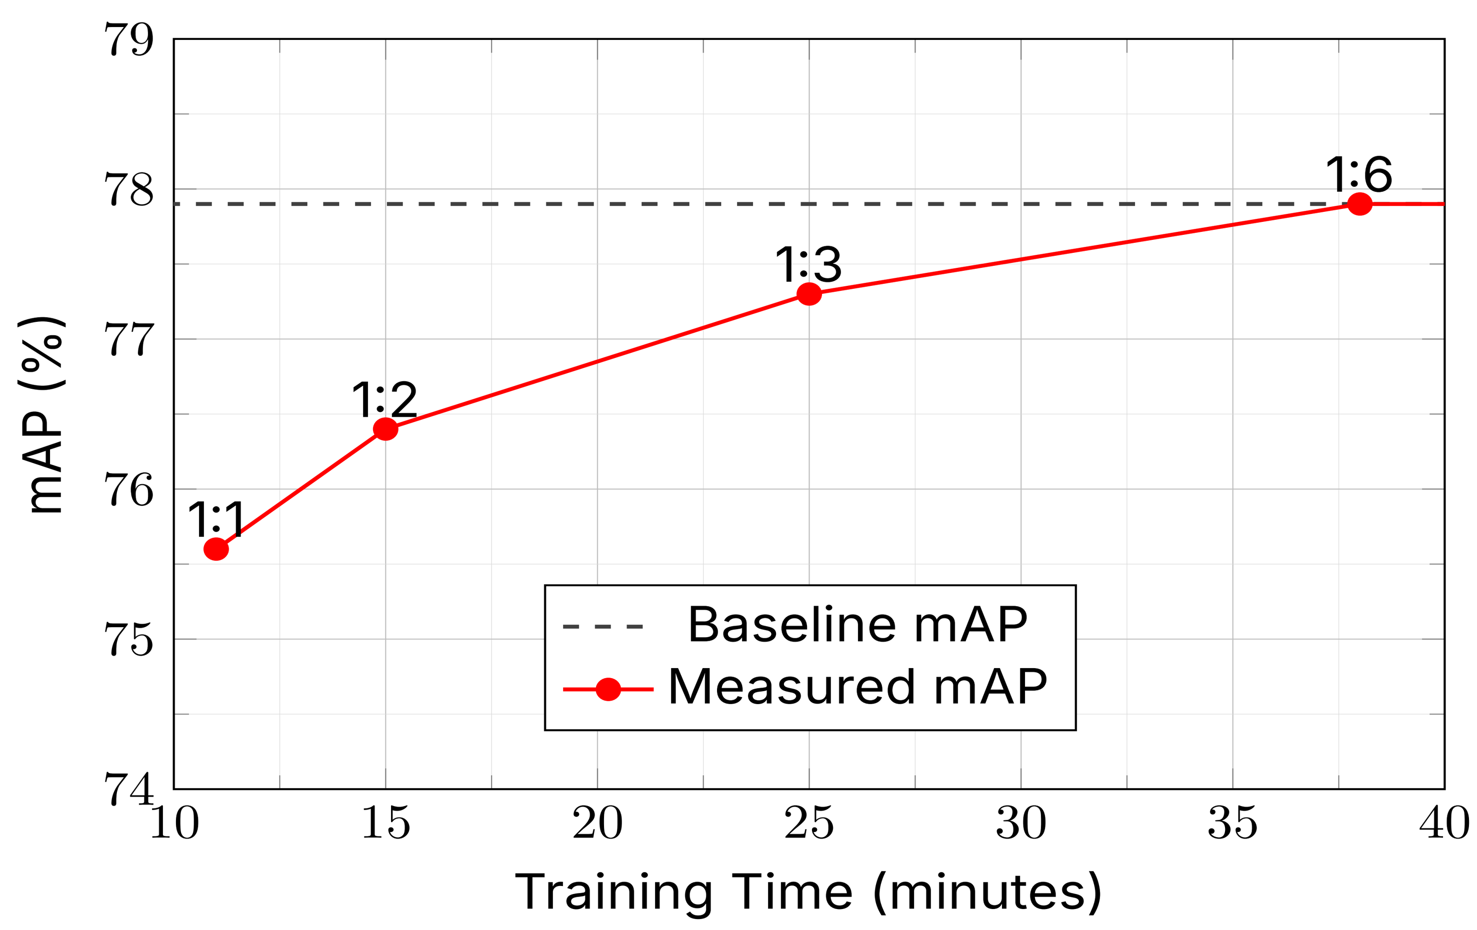

figure 1: preserving performance by leveraging old data.

segmentation type comparison

Task	F1 (%)	PQ (%)	AP (%)
Instance	94.4	90.9	77.9
Panoptic	−6.1	−8.3	−9.1
Semantic	−18.3	−23.0	−24.9

table 1: Impact of task type on F1, PQ, and mAP metrics.

fine-tuning approaches

- Full model fine-tuning - severe catastrophic forgetting
- Backbone freezing - moderate preservation of knowledge
- Selective parameter updating - best balance
- Low-Rank Adaptation - comparable to freezing but with less parameters