Self-Supervised Segmentation Loss: Learning to Estimate Loss Quality form Mask

Klára Janoušková

First experiments

First experiment: See what we can do when learning to estimate the IoU Loss on the training set (same as where the segmentation was trained), no augmentation or anything like that. The loss function is MSE: x is predicted mask

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N}(\mathrm{IoU}_{\mathrm{GT},x} - f(x))^2$$

The minimum validation loss is $l = 0.009631$; $\sqrt{l} = 0.0981$

We can note it converges very fast.

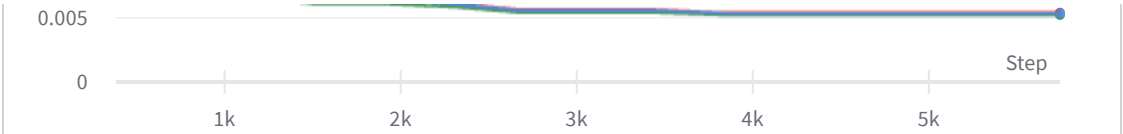
0.005
0
Step
1k
2k
3k
4k
5k


train/loss
train_loss_test_IoU_loss_lr_0.0001
train_loss_test_IoU_loss_lr_0.0005
train_loss_test_IoU_loss_lr_1e-05
0.1
0.08
0.06
0.04
0.02
0
1k
2k
3k
4k
5k

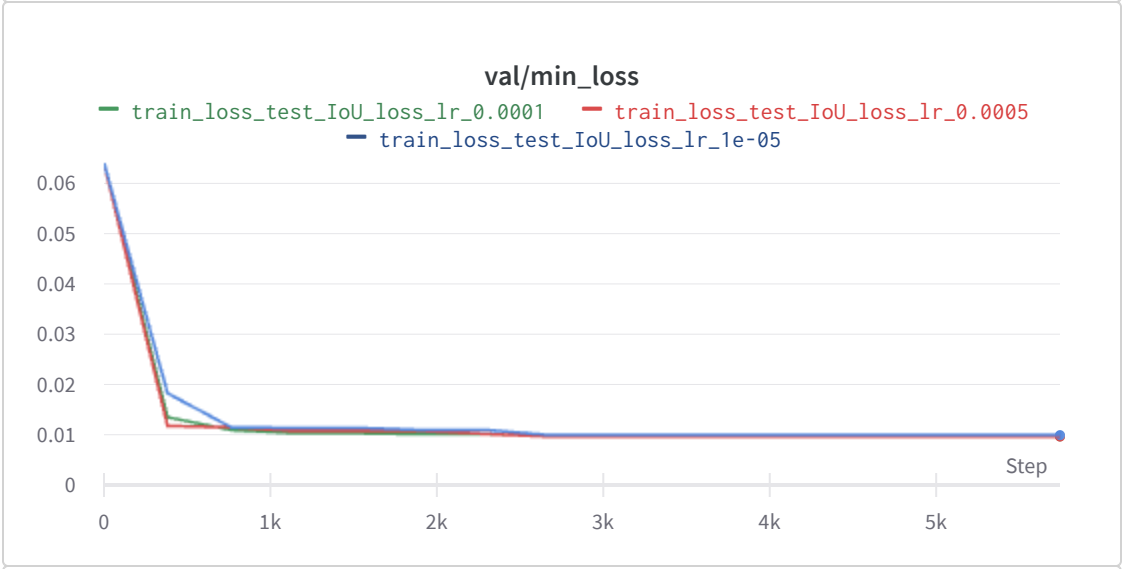
val/min_loss
train_loss_test_IoU_loss_lr_0.0001
train_loss_test_IoU_loss_lr_0.0005
train_loss_test_IoU_loss_lr_1e-05
0.06
0.05
0.04
0.03
0.02
0.01
0
Step
0
1k
2k
3k
4k
5k

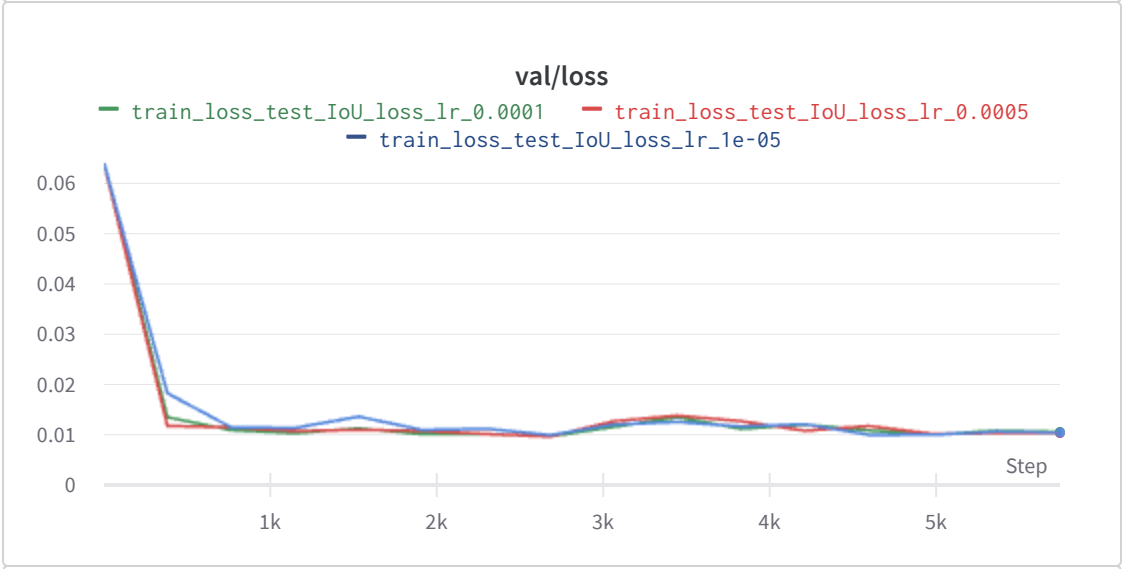
val/loss
train_loss_test_IoU_loss_lr_0.0001
train_loss_test_IoU_loss_lr_0.0005
train_loss_test_IoU_loss_lr_1e-05
0.06
0.05
0.04
0.03
0.02
0.01
0
Step
1k
2k
3k
4k
5k


lr
0.00050
val/min_loss
0.00992

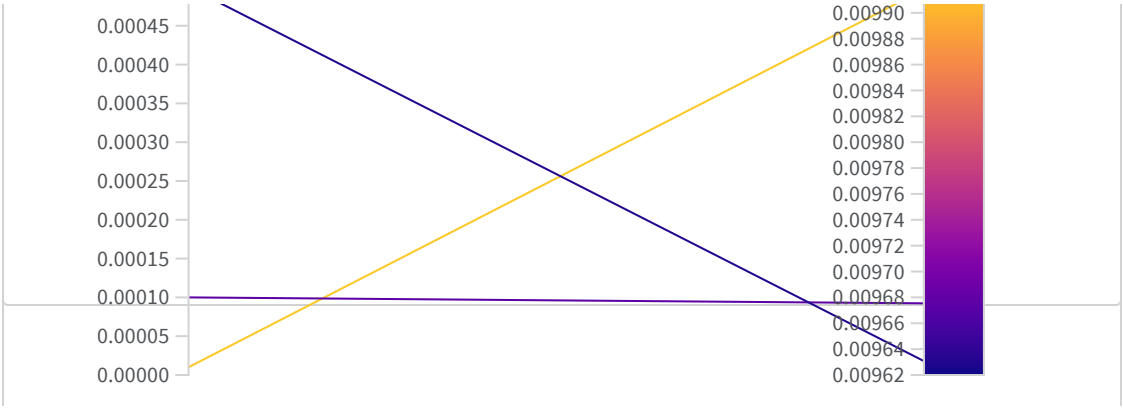

L1 vs L2 with distortion

L1 seems to be doing better.

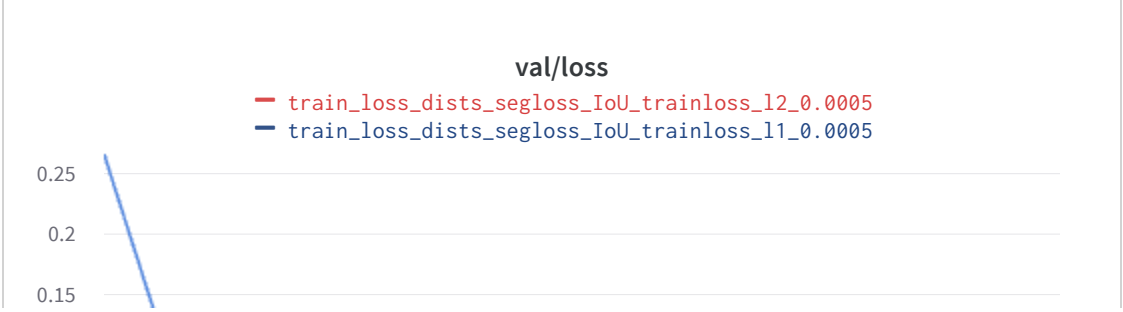

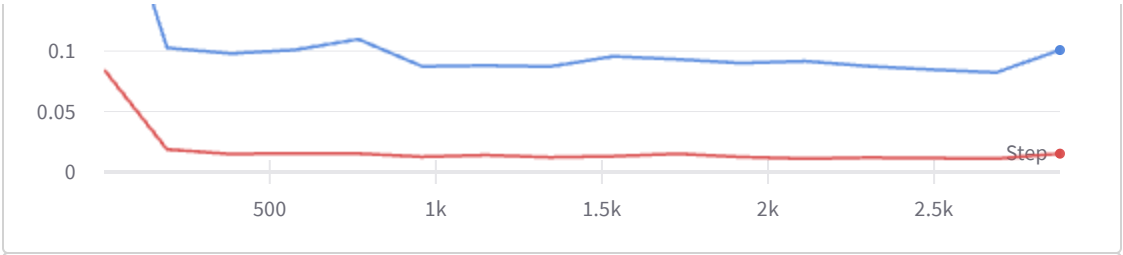
0.1
0.05
0
Step
500
1k
1.5k
2k
2.5k


train/min_loss
train_loss_dists_segloss_IoU_trainloss_l2_0.0005
train_loss_dists_segloss_IoU_trainloss_l1_0.0005
0.12
0.1
0.08
0.06
0.04
0.02
0
Step
500
1k
1.5k
2k
2.5k

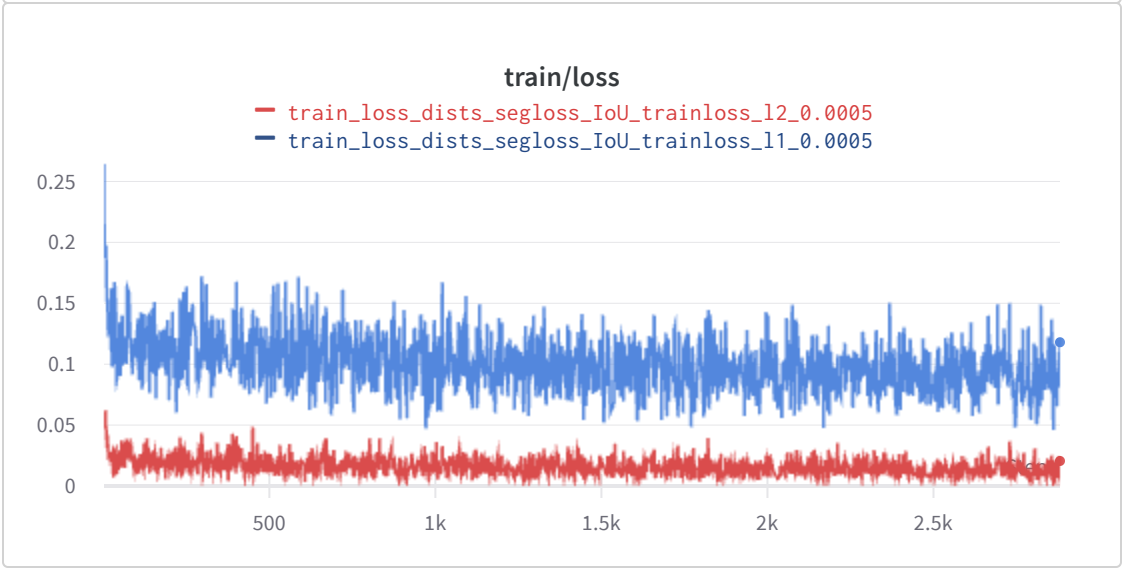
train/loss
train_loss_dists_segloss_IoU_trainloss_l2_0.0005
train_loss_dists_segloss_IoU_trainloss_l1_0.0005
0.25
0.2
0.15
0.1
0.05
0
500
1k
1.5k
2k
2.5k

Run set 2

Created with ❤️ on Weights & Biases.

https://wandb.ai/klara/TTA-finetune/reports/Self-Supervised-Segmentation-Loss-Learning-to-Estimate-Loss-Quality-form-Mask--VmlldzozNzMxMzAz