

Loss Surfaces, Mode Connectivity, and Fast Ensembling of DNNs

Timur Garipov^{1,2} Pavel Izmailov³ Dmitrii Podoprikin⁴
Dmitry Vetrov⁵ Andrew Gordon Wilson³


¹Samsung AI Center in Moscow, ²Skolkovo Institute of Science and Technology, ³Cornell University,
⁴Samsung-HSE Laboratory, ⁵National Research University Higher School of Economics

Neural Information Processing Systems
Montreal, Canada

December 4, 2018



Loss Surfaces

ResNet-164, CIFAR-100



Loss Surfaces

ResNet-164, CIFAR-100



Finding Paths between Modes

- Weights of pretrained networks: $\hat{w}_1, \hat{w}_2 \in \mathbb{R}^{|net|}$

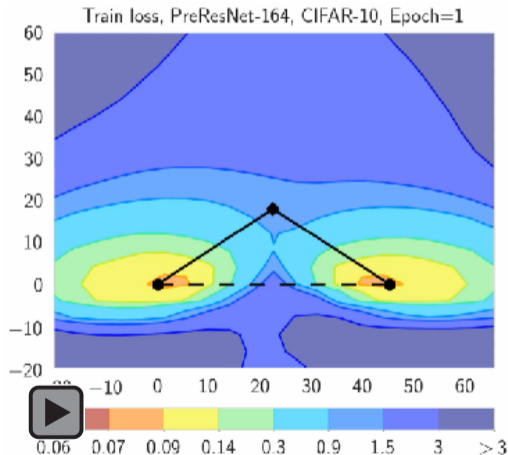
- Define parametric curve: $\phi_\theta(\cdot) [0, 1] \rightarrow \mathbb{R}^{|net|}$

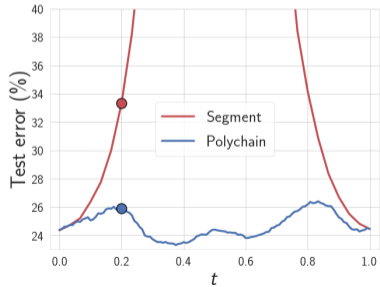
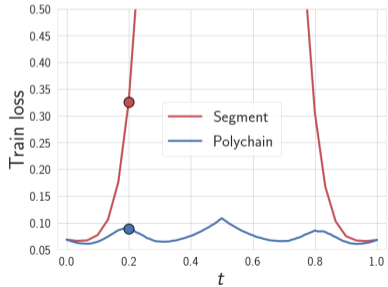

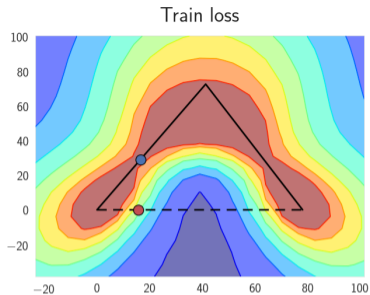
$$\phi_\theta(0) = \hat{w}_1, \quad \phi_\theta(1) = \hat{w}_2$$

- DNN loss function: $\mathcal{L}(w)$

- Minimize averaged loss w.r.t. θ

$$\underset{\theta}{\text{minimize}} \ell(\theta) = \int_0^1 \mathcal{L}(\phi_\theta(t)) dt = \mathbb{E}_{t \sim U(0,1)} \mathcal{L}(\phi_\theta(t))$$











Loss Surfaces

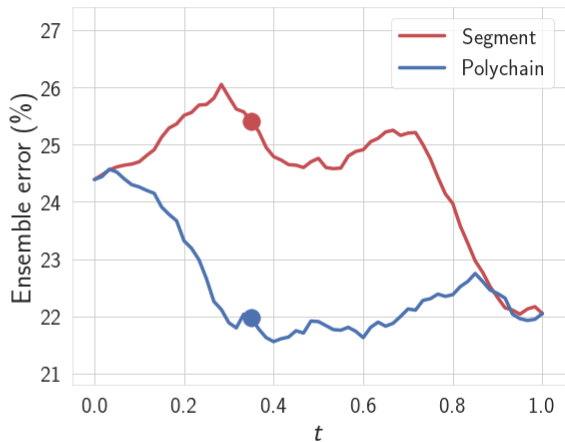

VGG-16, CIFAR-10

Train loss

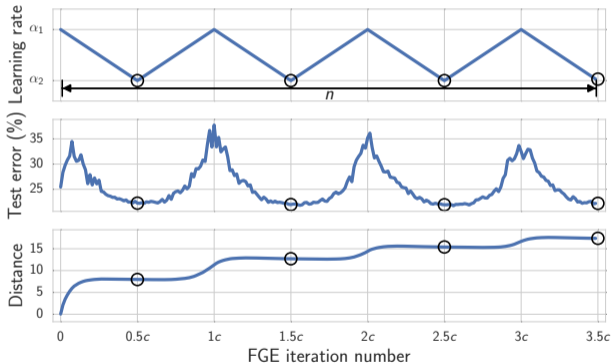



Test error (%)

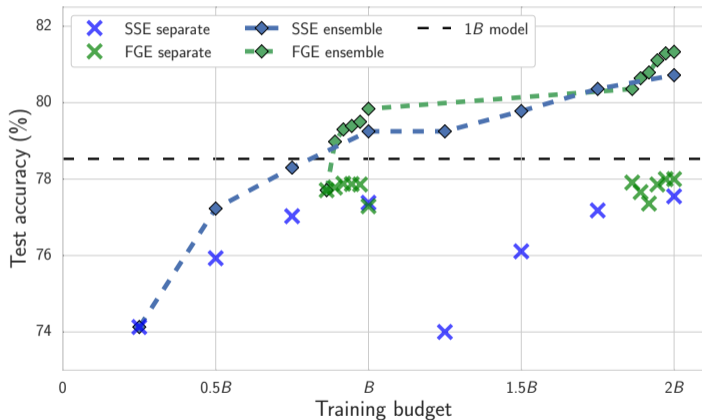




Fast Geometric Ensembles (FGE)



Ensembling Results



SSE = Huang et al., ("Snapshot ensembles: Train 1, get m for free"), ICLR 2017

Summary

- Local optima are connected by simple curves.
- To find these curves we minimize loss uniformly in expectation over a path from one mode to another.
- We are inspired by these insights to propose a fast ensembling algorithm.

PyTorch code released for both mode connectivity and FGE

Come to our poster #162!