Teaching Your MAML

an adaptive teacher for meta-learning

Gleb Shevchuk
glebs@stanford.edu

Motivation

We want general AI! But how do we train it? Let's look to humans!

Approach

Introduce a **teacher model** that changes problem difficulty **in response to student progress** during meta training.

Experiments

Three meta learning experiments:

1) MNIST classification[1]
2) Omniglot 5 way 1 shot classification[2]
3) Omniglot 20 way 1 shot classification[2]

Compare teacher-aided model to regular MAML[3].

Discussion

1) MAML already really hard to train.[4]
2) Hard to find best teacher hyperparameters.
3) Still need to know "hardest" problem.
4) Unclear how much teacher helps.
5) Experiments were pretty similar.

The Nitty Gritty

1) Task distribution has difficulty Ψ with current parameters ω
2) Perturb ω with meta train step
3) Approximate gradient ∇ ω
4) Teacher does gradient descent over task space parameters!

Models

Three layer conv net with batch norm and max pooling used for all 3 experiments.

Results

Adaptive teacher seems to help generalize!

Future!

Extend to:

1) Continuous task parameter space
2) Regression and reinforcement learning problems
3) Lifelong settings

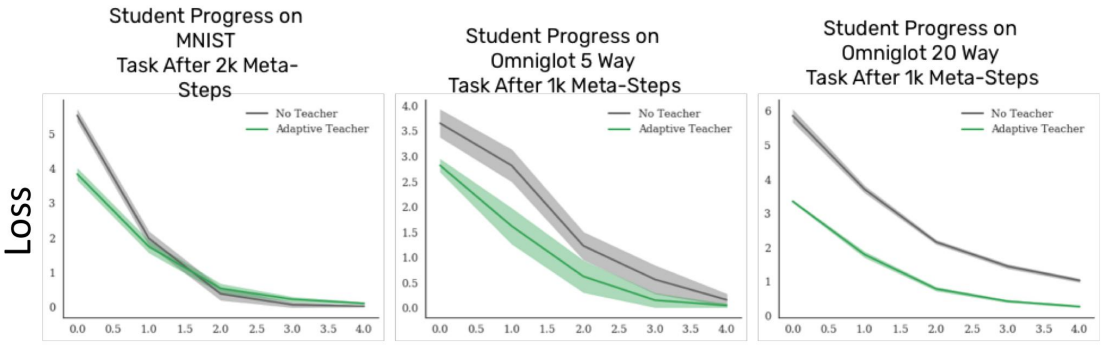

[1] Yann LeCun, Corinna Cortes, and CJ Burges. "MNIST handwritten digit database"
[2] Brenden Lake et al. "One shot learning of simple visual concepts."
[3] Chelsea Finn, Pieter Abbeel, and Sergey Levine. "Model-agnostic meta-learning for fast adaptation of deep networks"
[4] Antreas Antoniou, Harrison Edwards, and Amos Storkey. "How to train your MAML"