Understanding Deep Learning Requires Rethinking Generalization

Chiyuan Zhang
CSAIL, CBMM, MIT

Poster: Wednesday Morning C23


Chiyuan Zhang

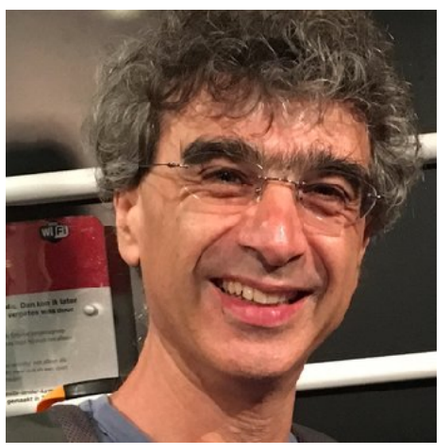
Samy Bengio

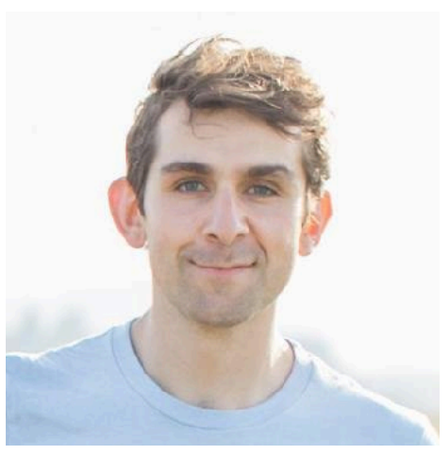
Moritz Hardt


Benjamin Recht


Oriol Vinyals

Model Selection

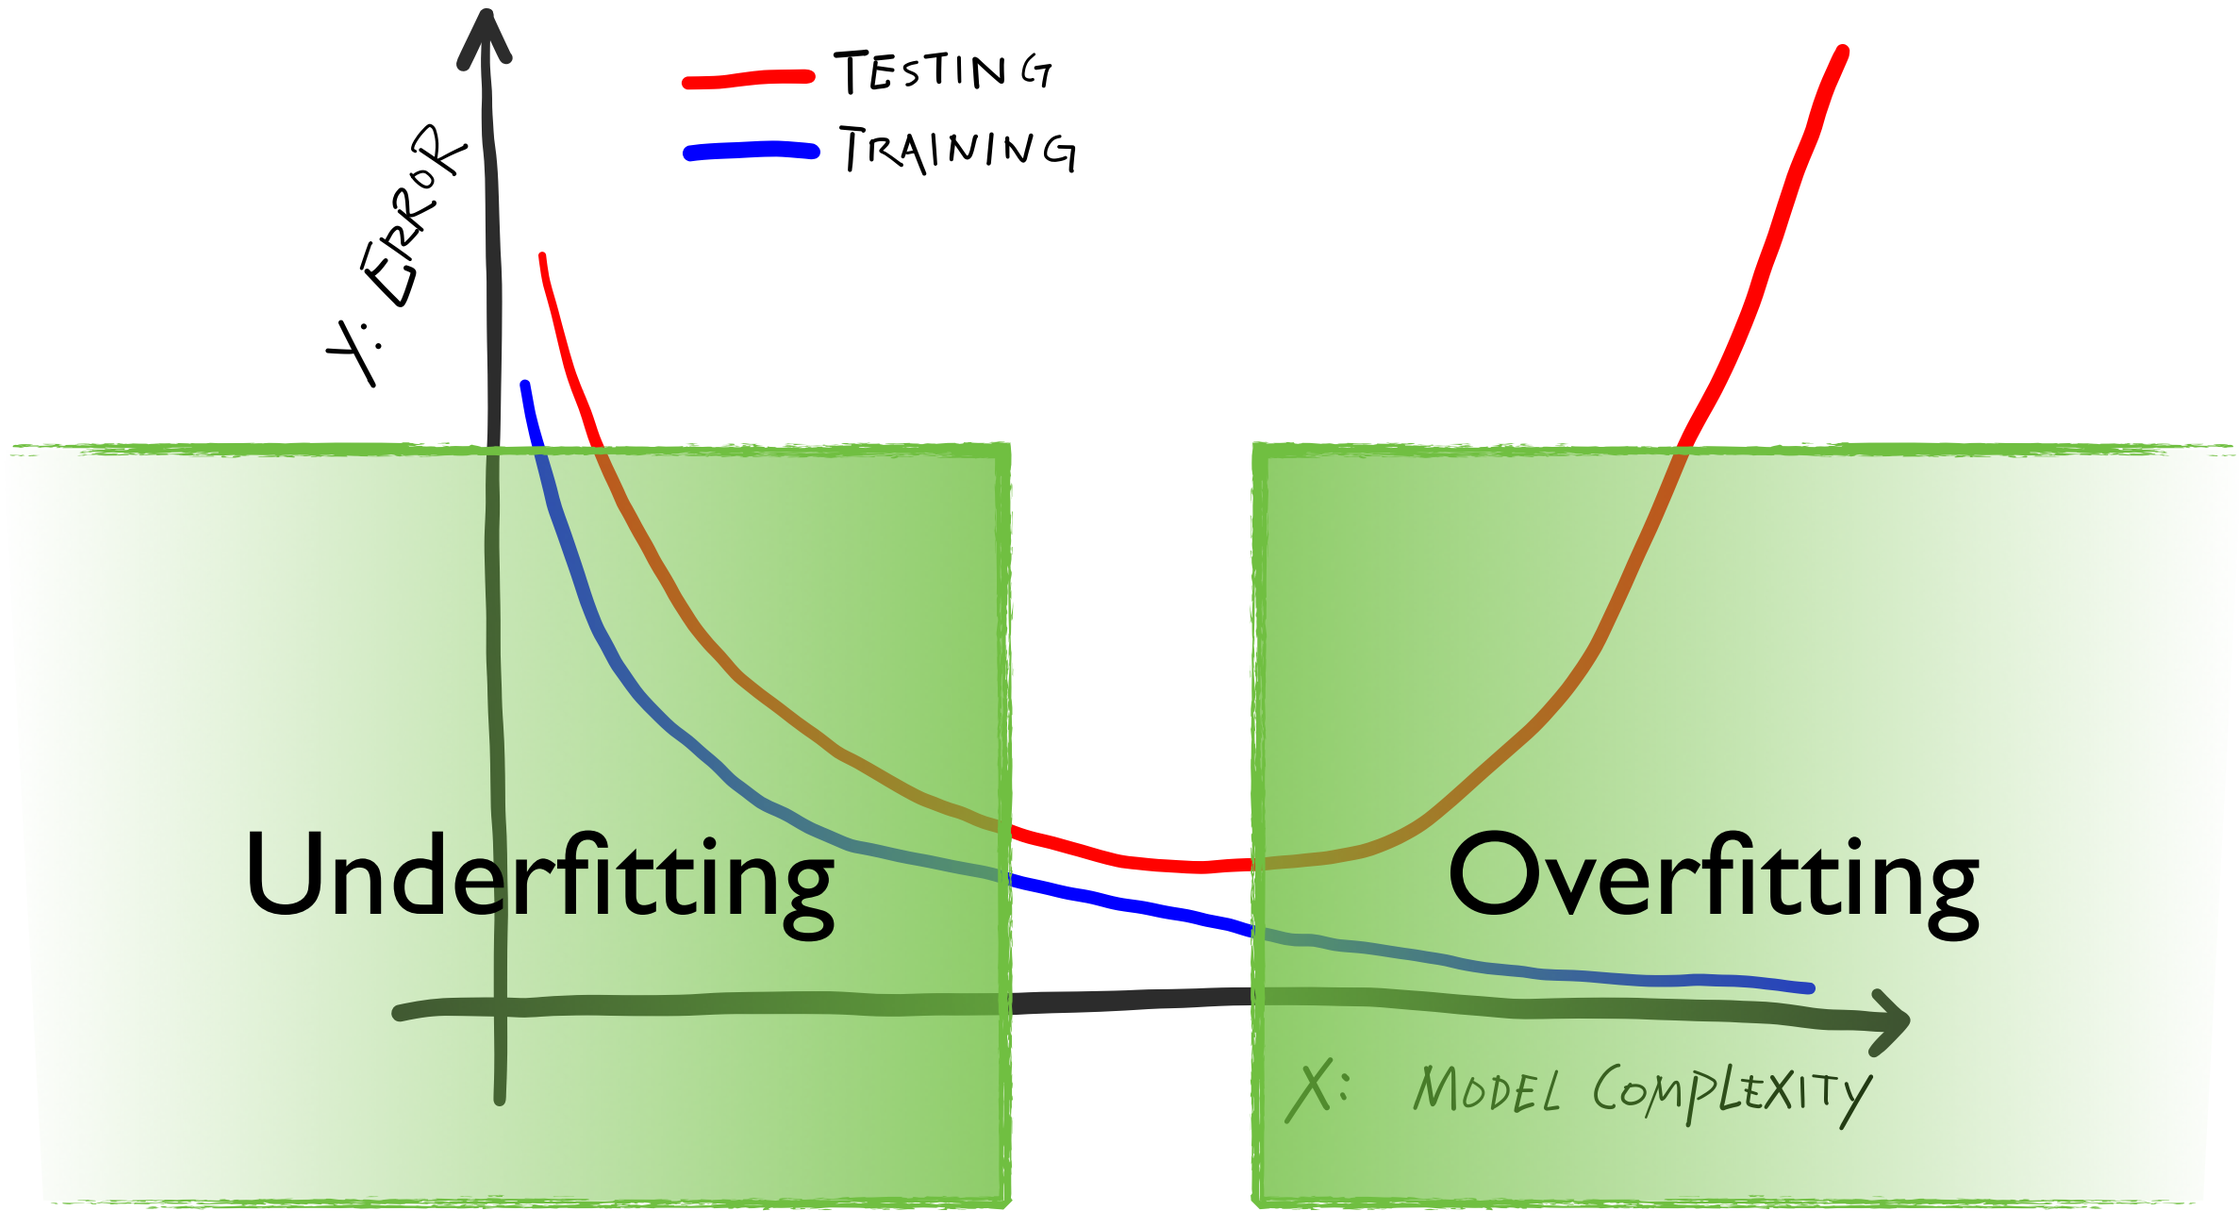

Over-parameterized Models

What are
the purple dots?

A Water Snake

The Constellation Hydra

https://www.wikiwand.com/en/Hydra_(constellation)
Screenshot from the "Star Walk 2" app.

Deep Learning

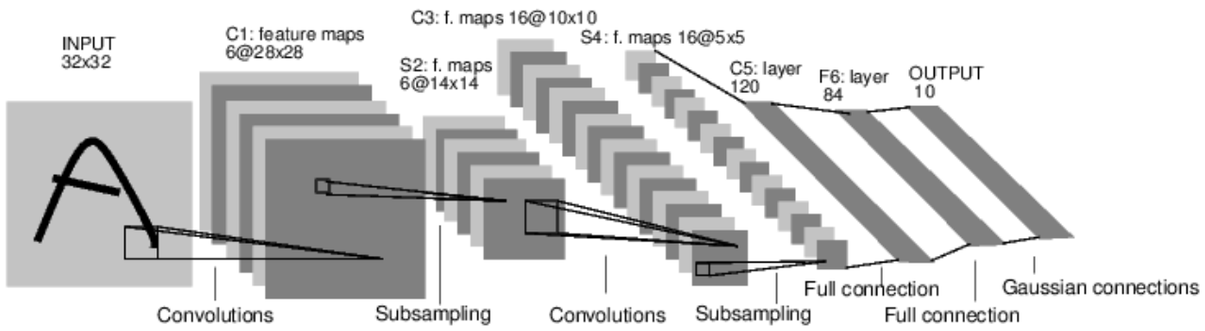

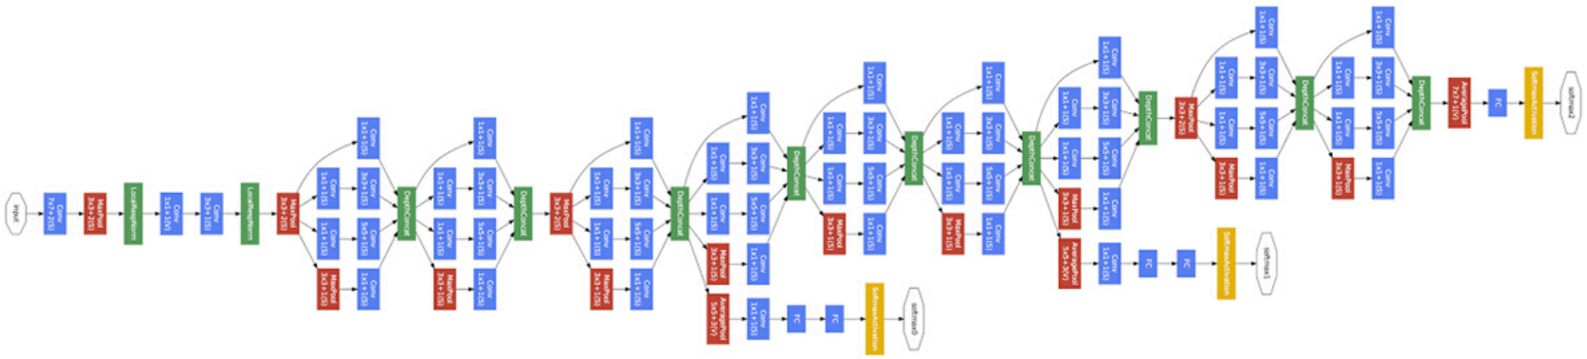

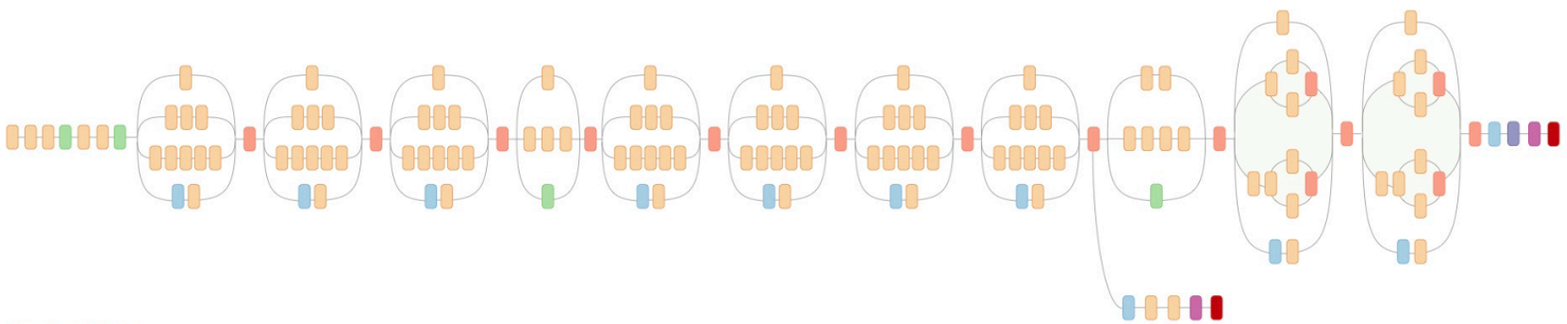

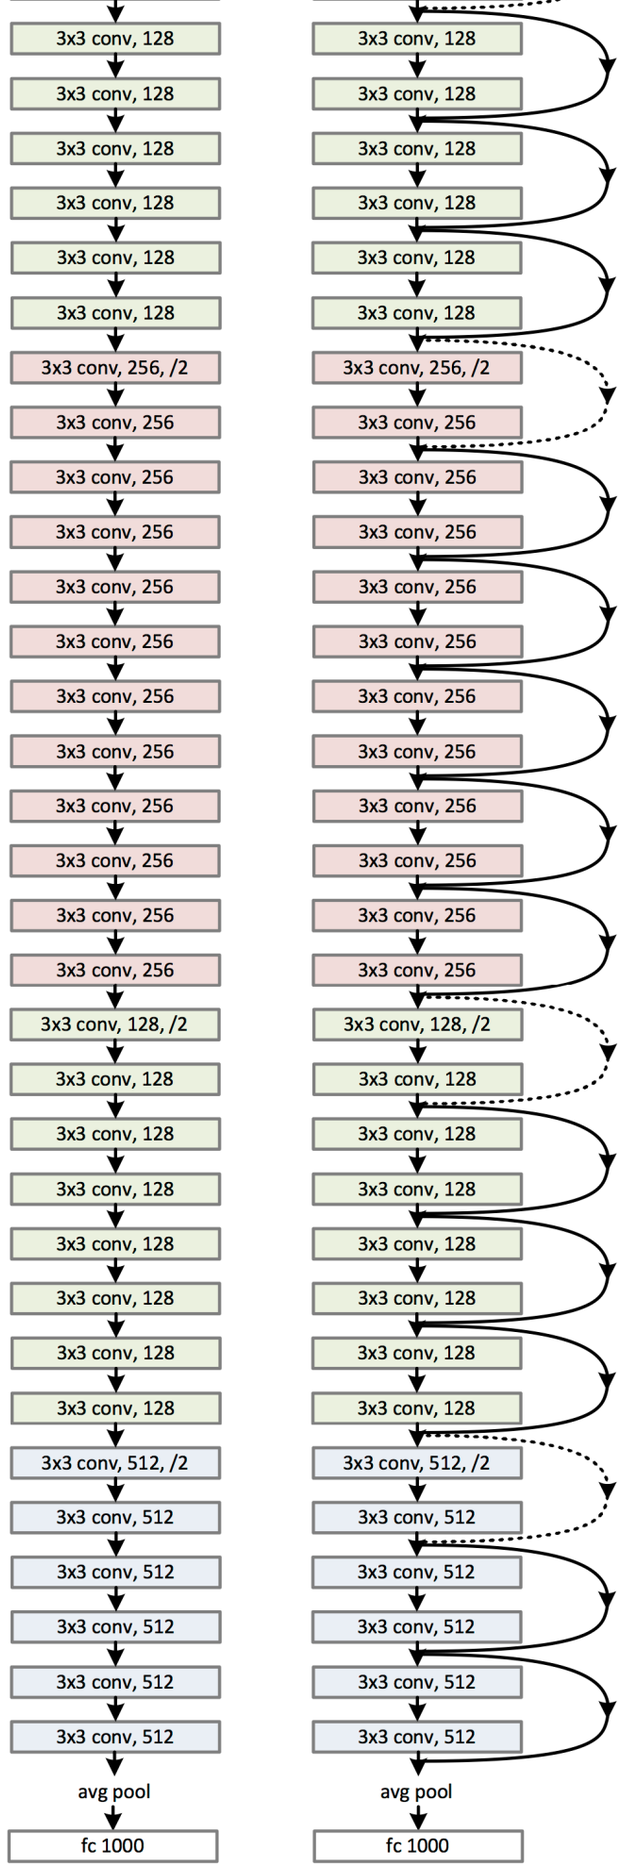

Bias — Variance

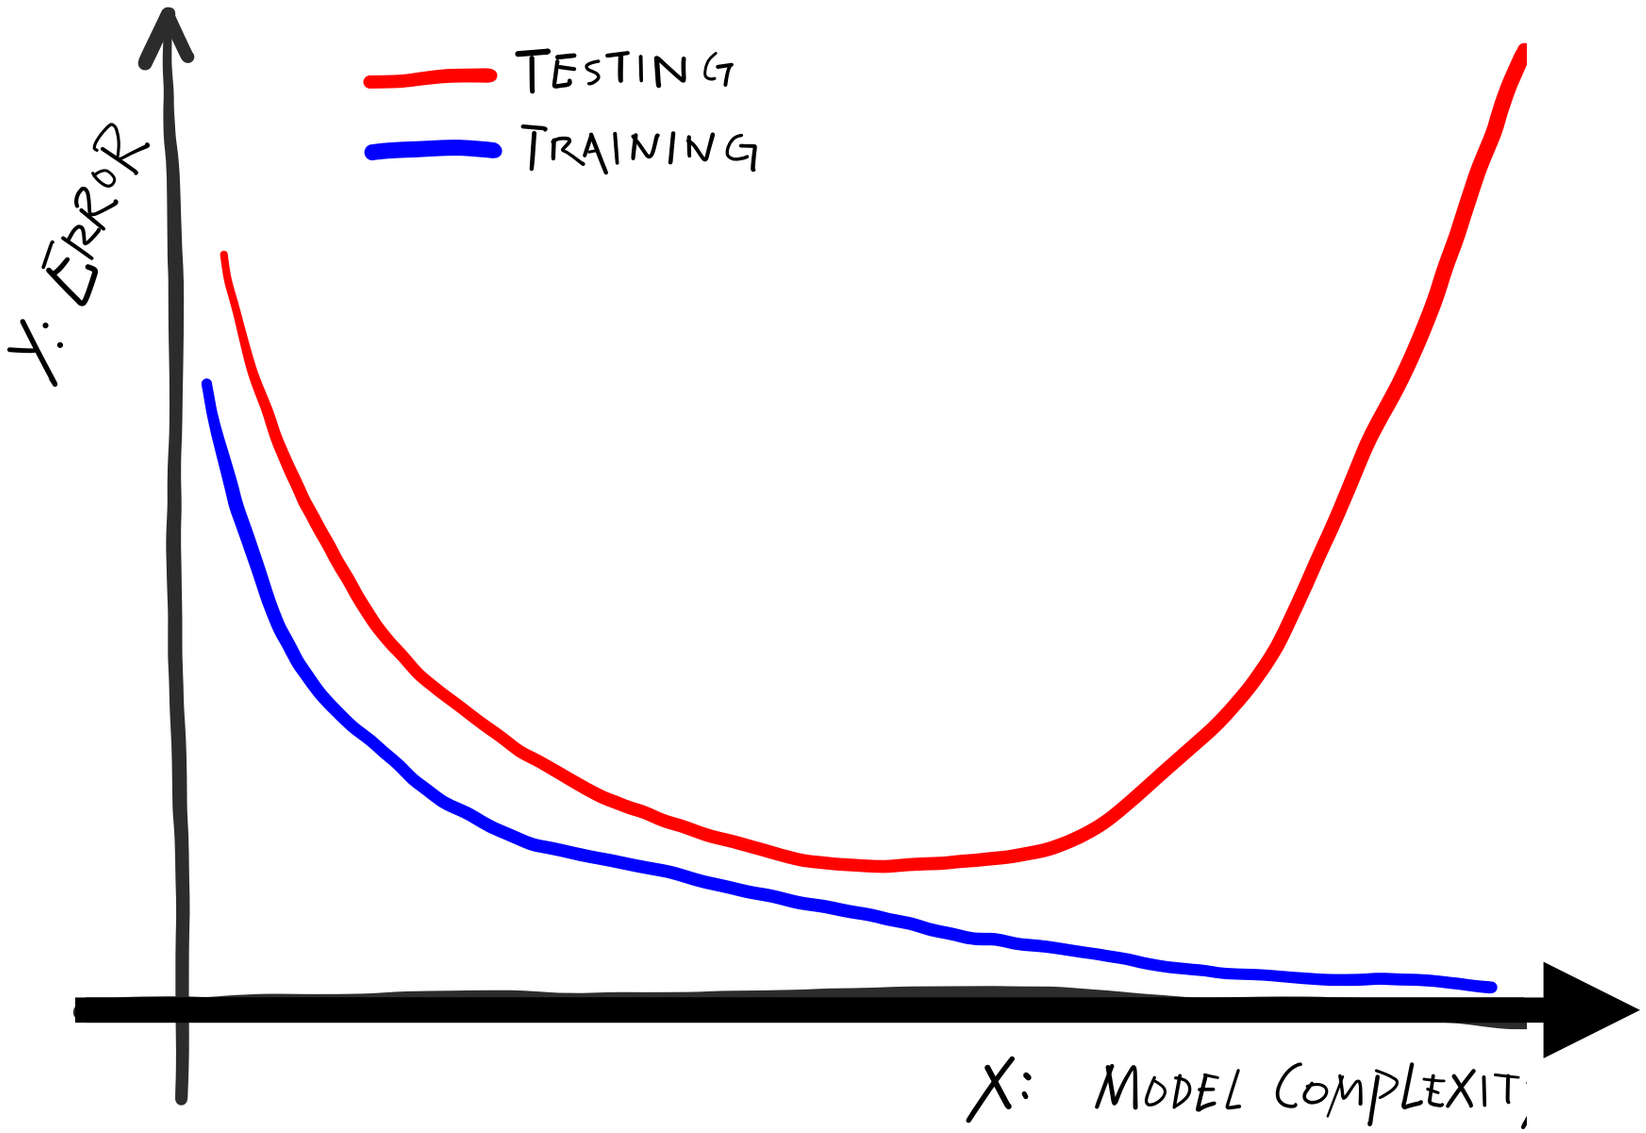

Deep Learning ★

Parameter Count

Num Training Samples

MLP 1x512
p/n: 24

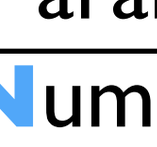

Alexnet
p/n: 28

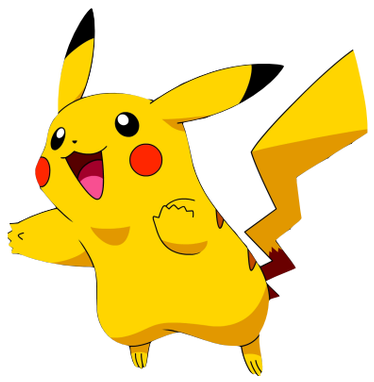

Inception
p/n: 33

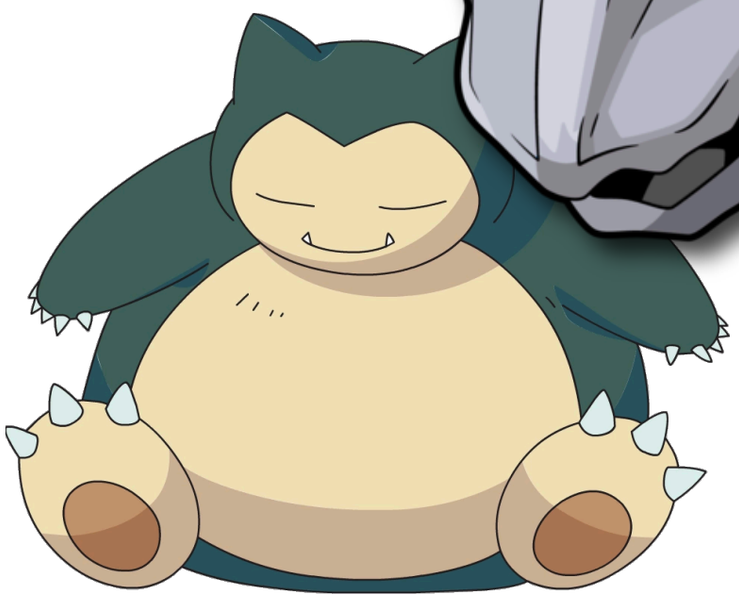

Wide Resnet
p/n: 179

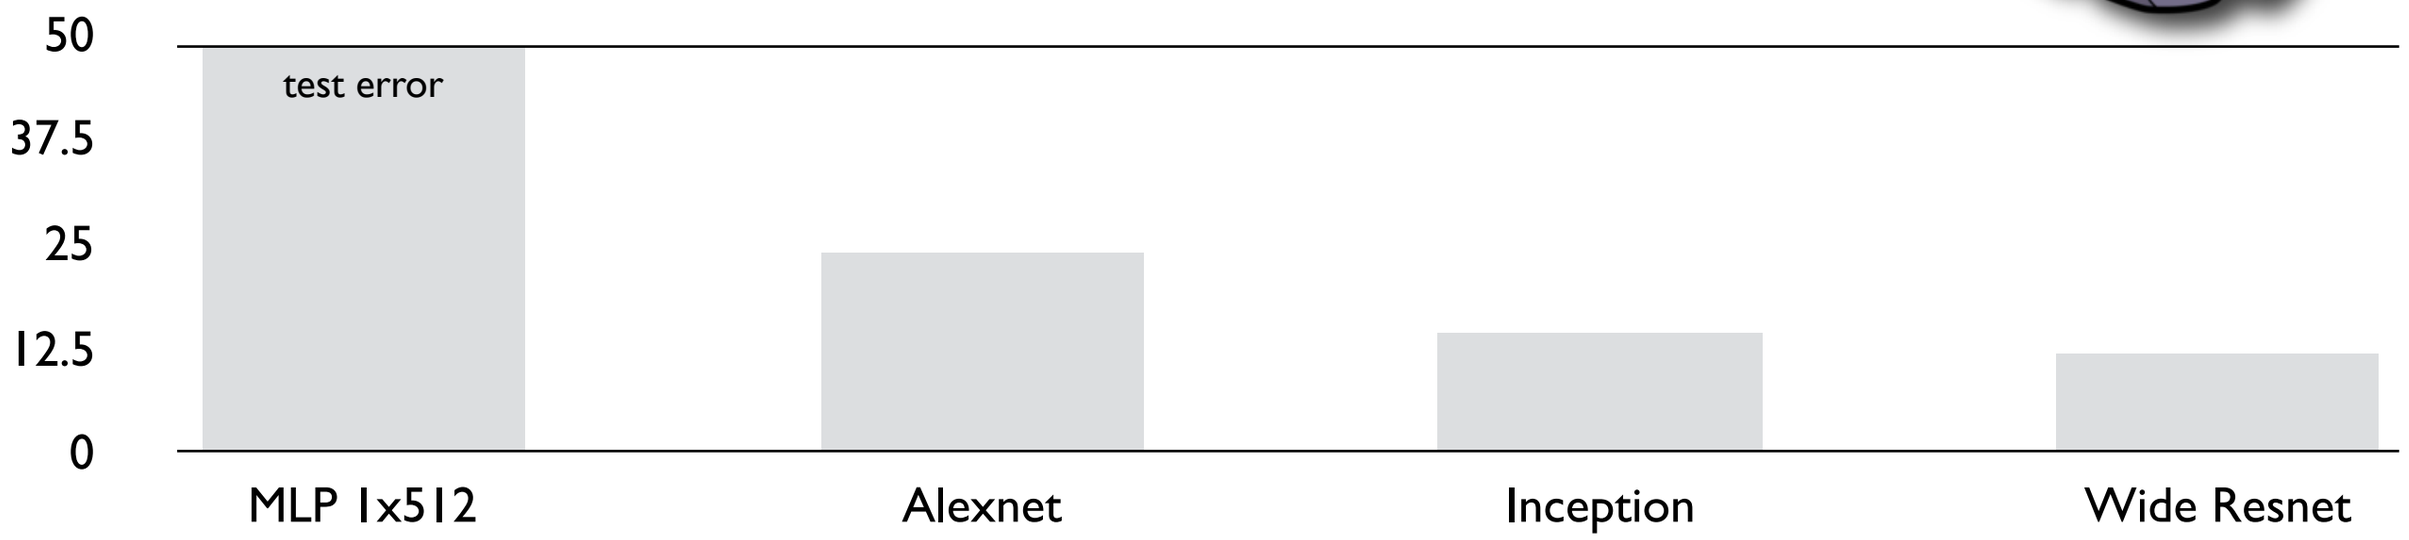

Randomization Test

Deep Neural Networks easily fit random labels.

Random Label Dataset

Dog

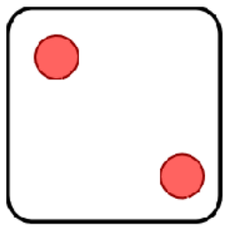

Cat

Flower

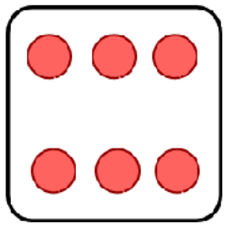

Dog

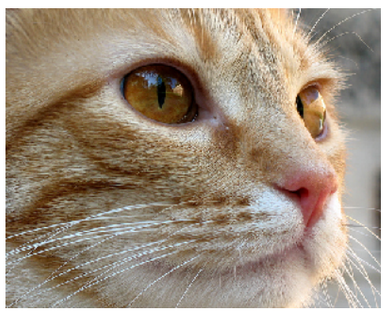

Cat

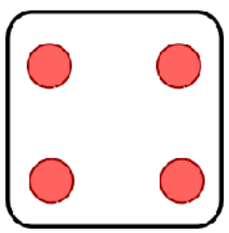

Bus

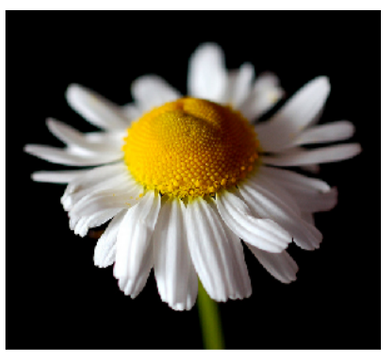

Flower

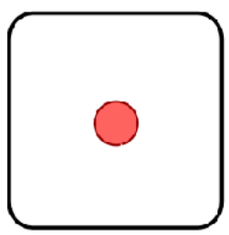

Bird

Randomization Test

arch: MLP

arch: Alexnet

arch: Inception
optim: SGD
batch: 128
lr: 0.1
momentum: 0.9
init: …

Recipes of Successful Models

Data with original labels


dog | flower

Data with **random** labels

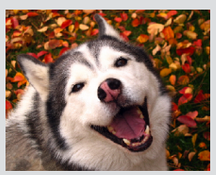
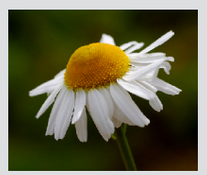

bus | dog

https://www.tensorflow.org/
http://mxnet.io/

TensorFlow™ MXNet

Randomization Test

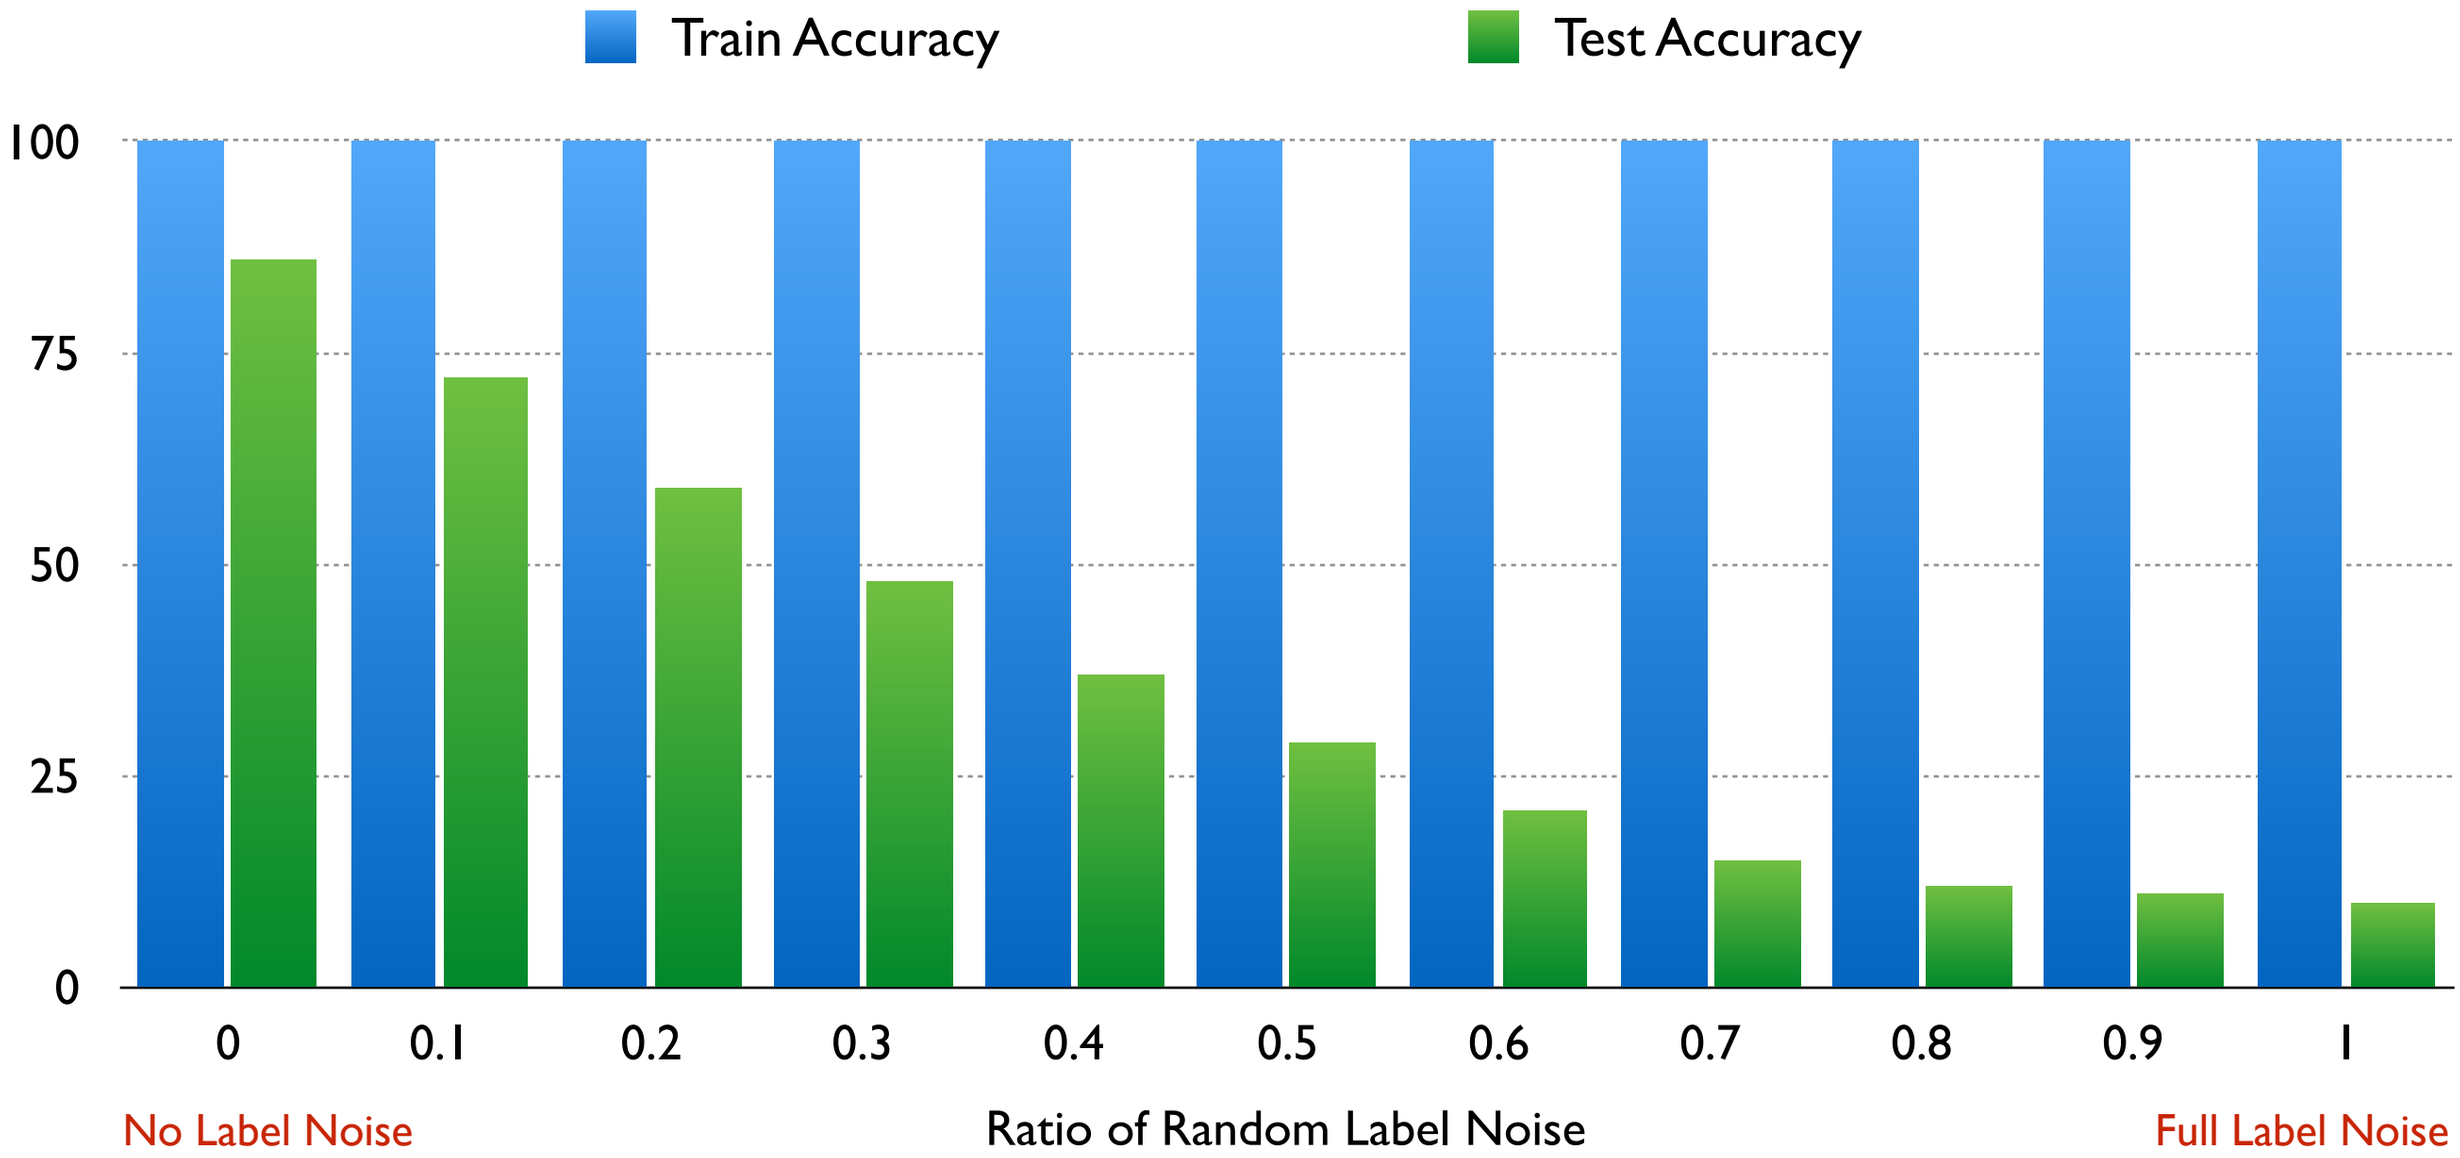

Randomization Test

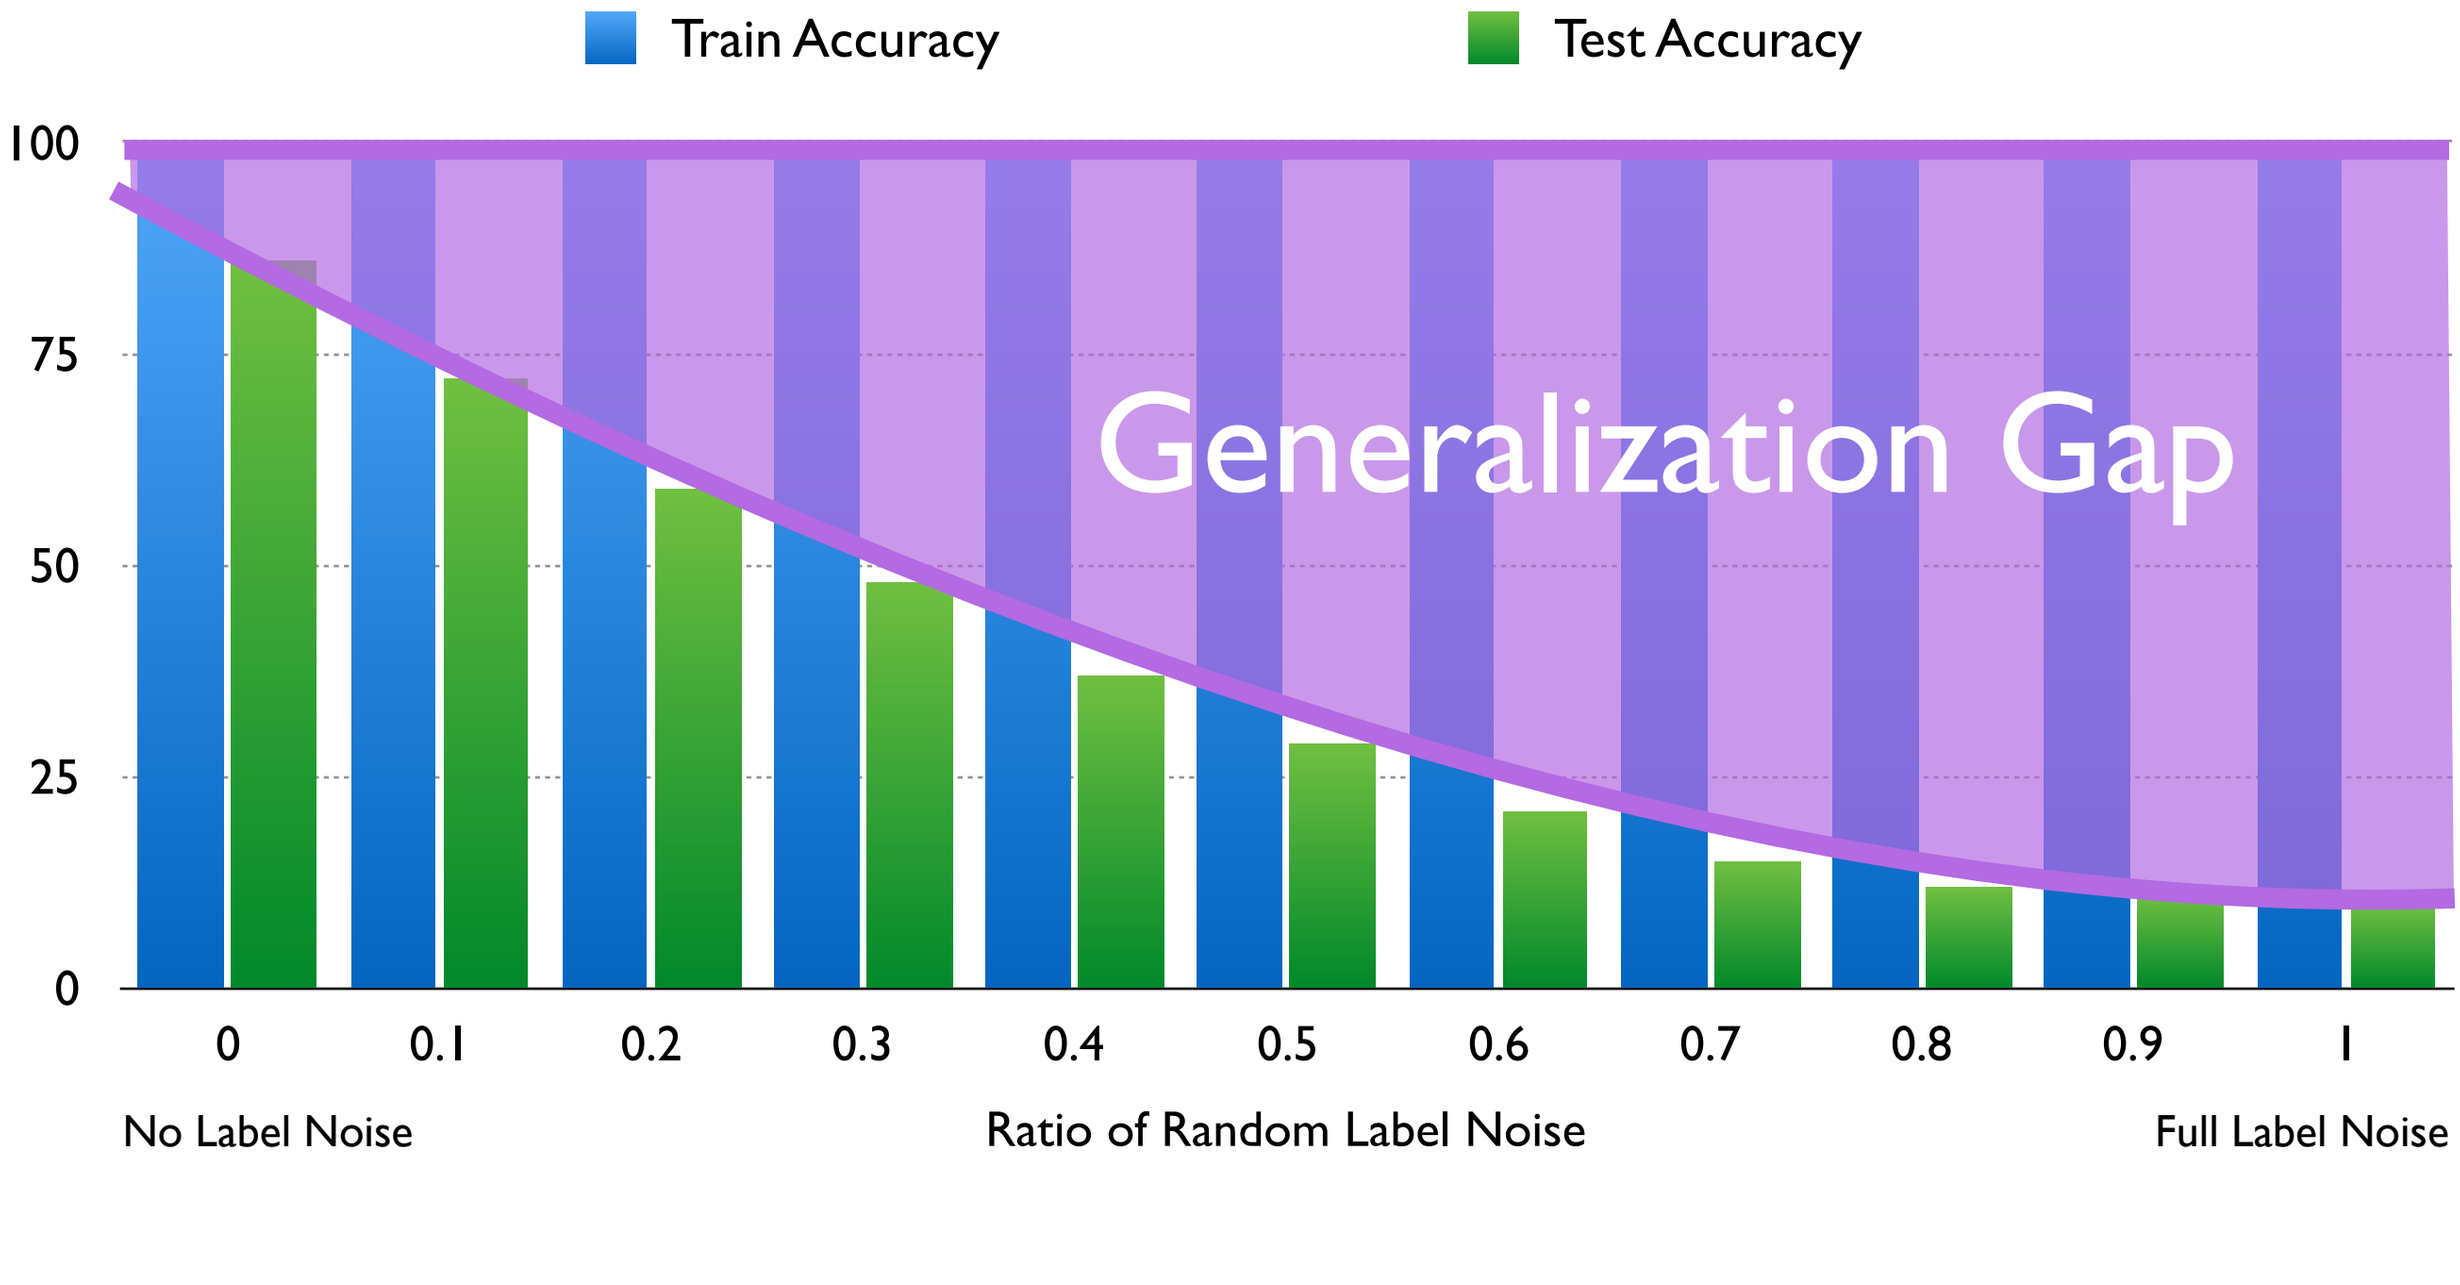

Randomization Test

Deep Neural Networks easily fit random labels.

Regularizers

⇐ Big Hypothesis Space

⇓ Regularized Models

Regularizers in Deep Learning

- Data augmentation: domain-specific transformations

- Weight decay: l2-regularizer on weights

- Dropout*: randomly mask out responses

* Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. JMLR, 15(1):1929–1958, 2014.

Fitting Natural Label with Regularizers

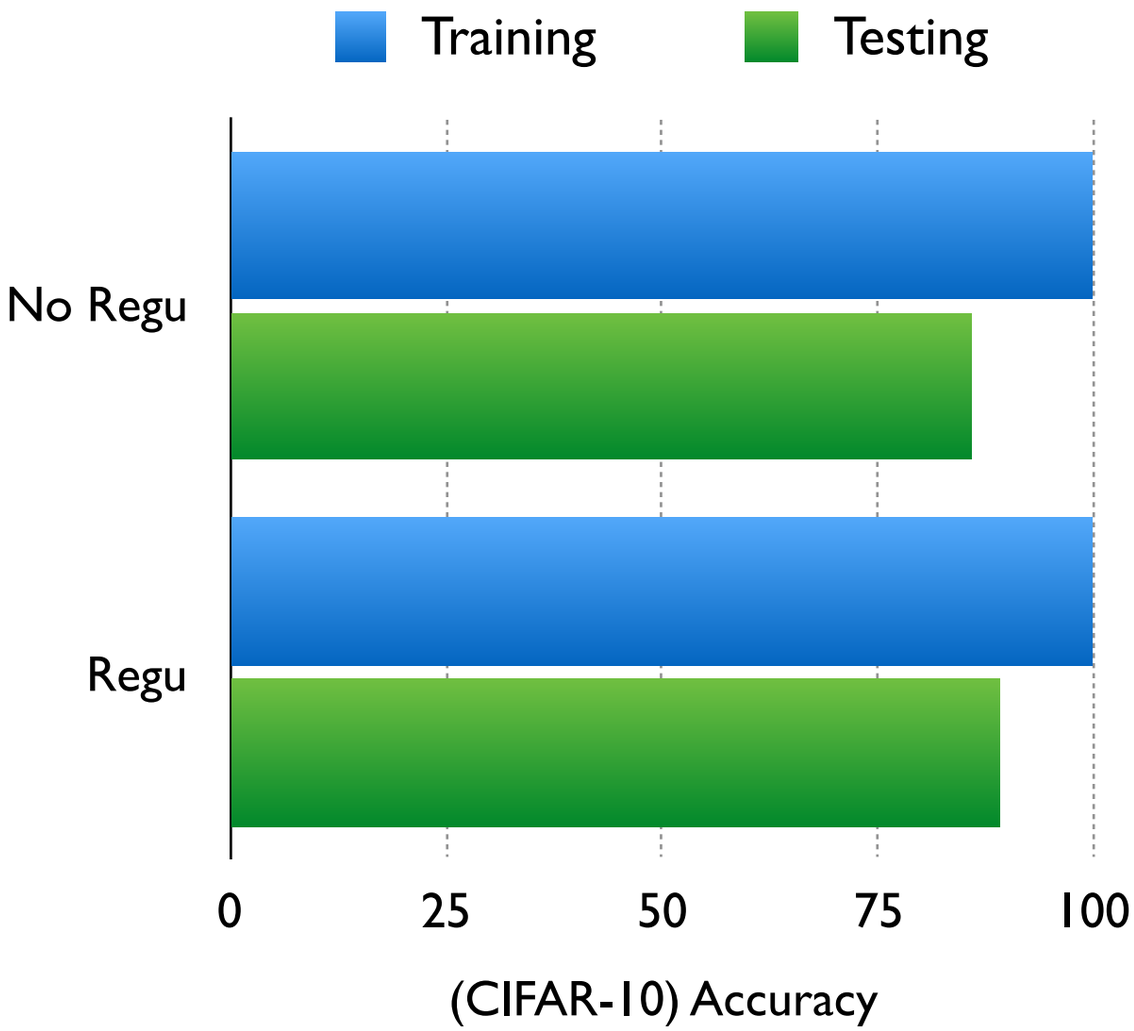

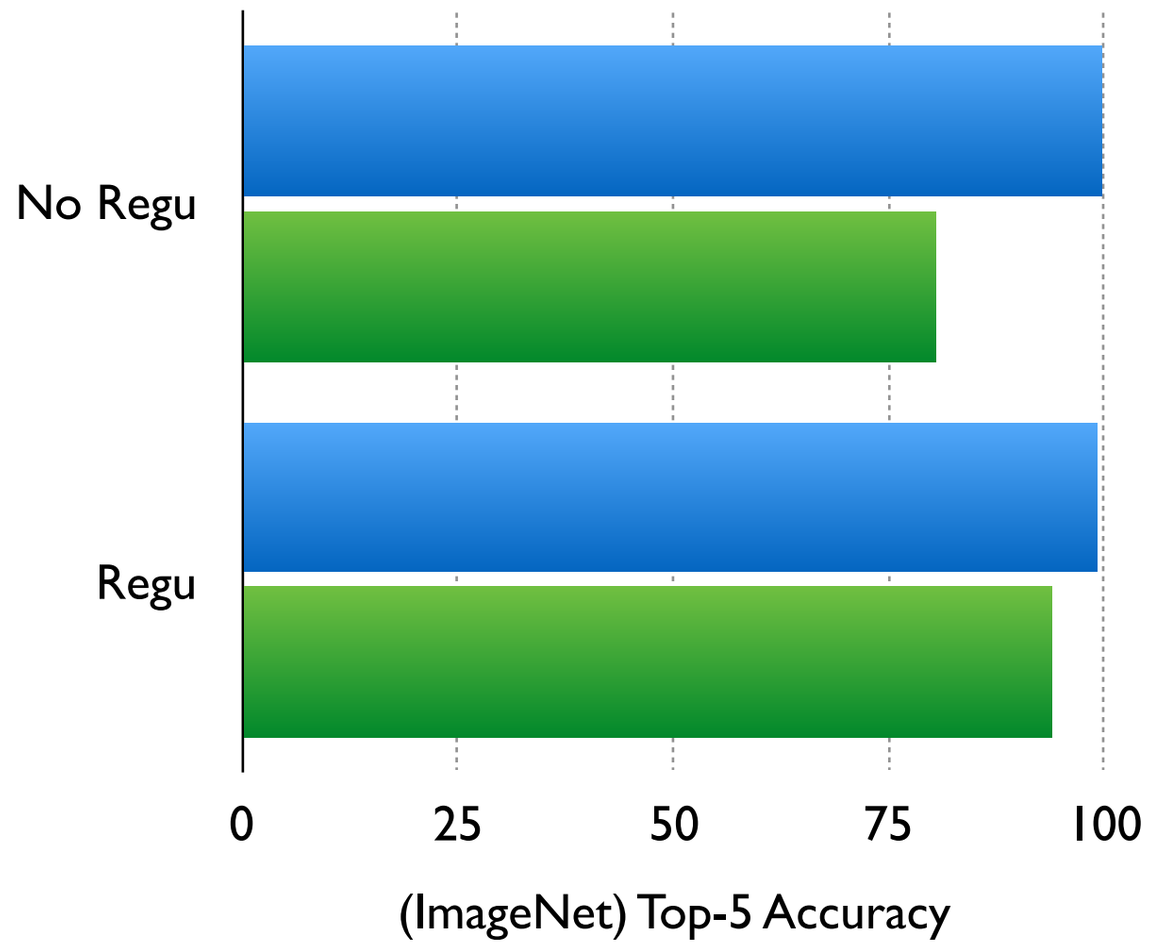

Fitting Random Label with Regularizers

Regularizer	Model	Training Accuracy
Weight decay	Inception	100%
	Alexnet	Failed to converge
	MLP 1x512	99.21%
Crop Augmentation*	Inception	99.93%

IMAGENET

Regularizer	Model	Training top-5
Dropout	Inception V3	96.15%
Dropout + Weight decay		97.95%

*We need to tune the hyperparams a bit and run for more epochs for this to converge, see paper for details.

Implicit Regularization

A REGULARIZER is ~~a mechanism~~ anything that

~~Constrain the model or empower the data.~~ hurt the training Process.

SGD fits Random Labels

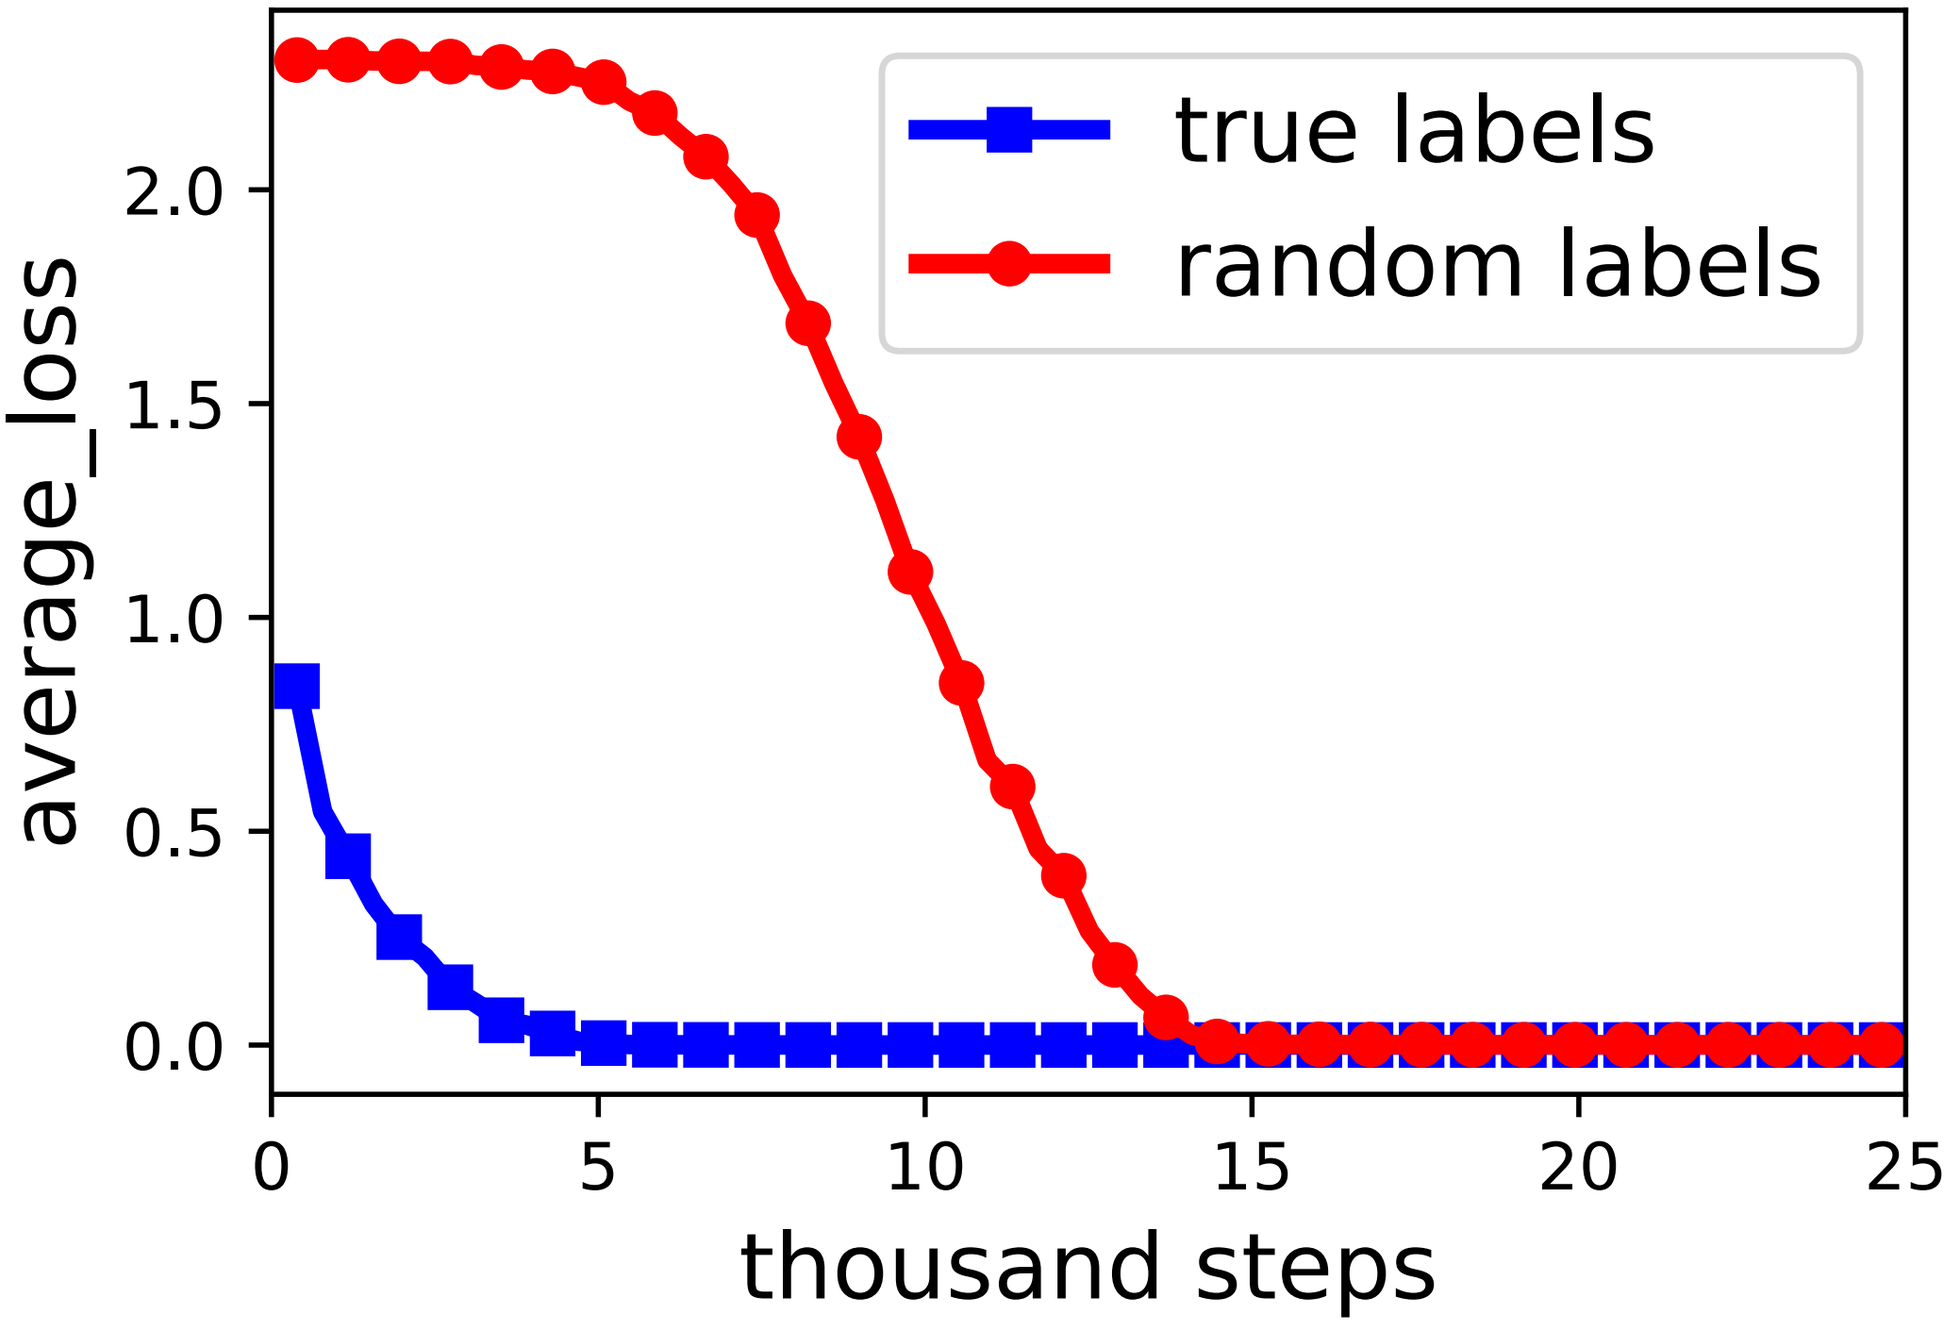

SGD fits Random Labels

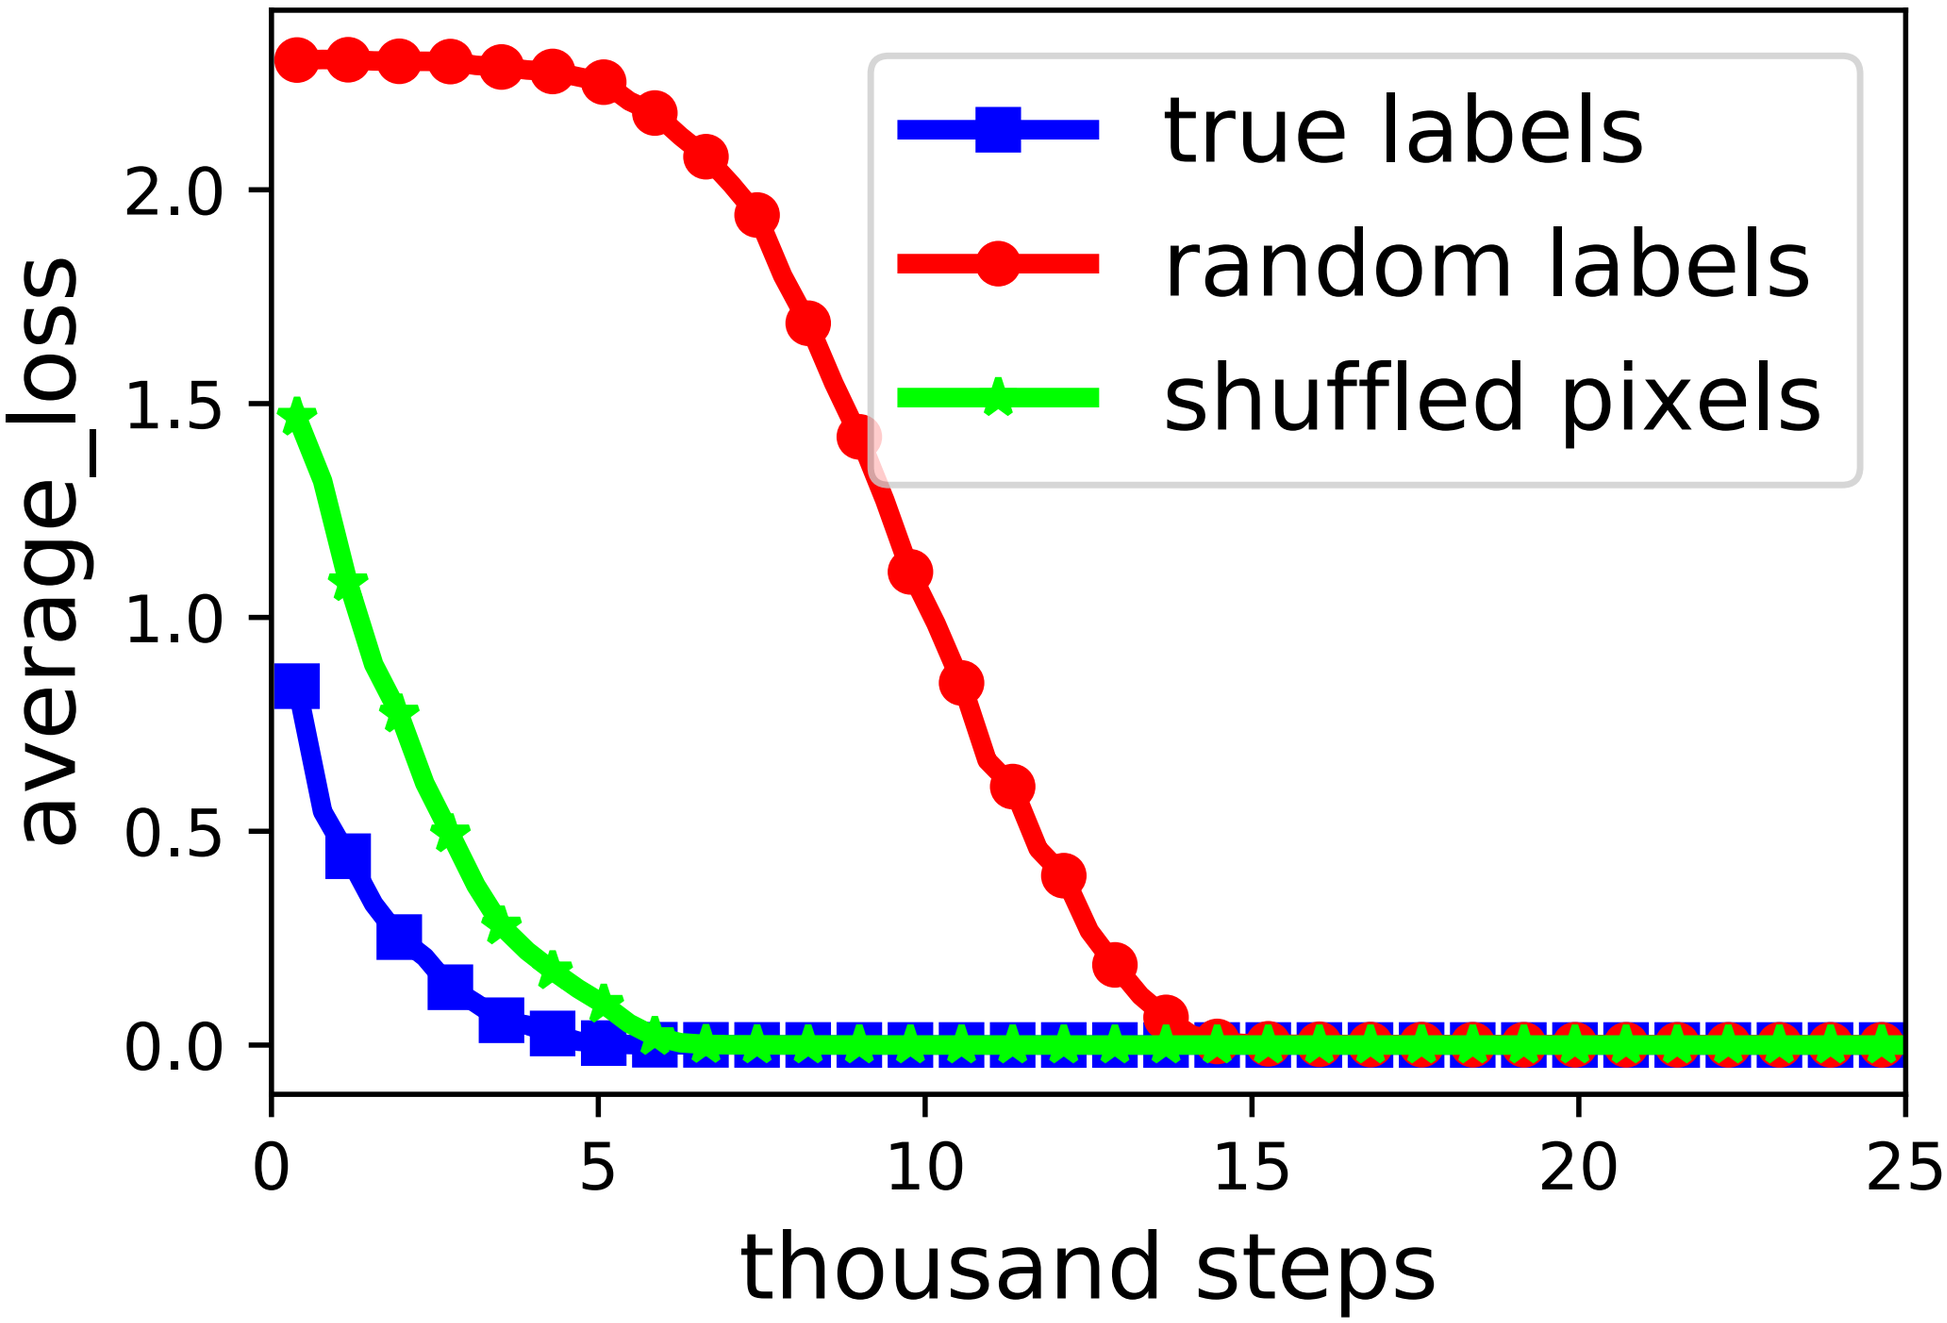

SGD fits Random Labels

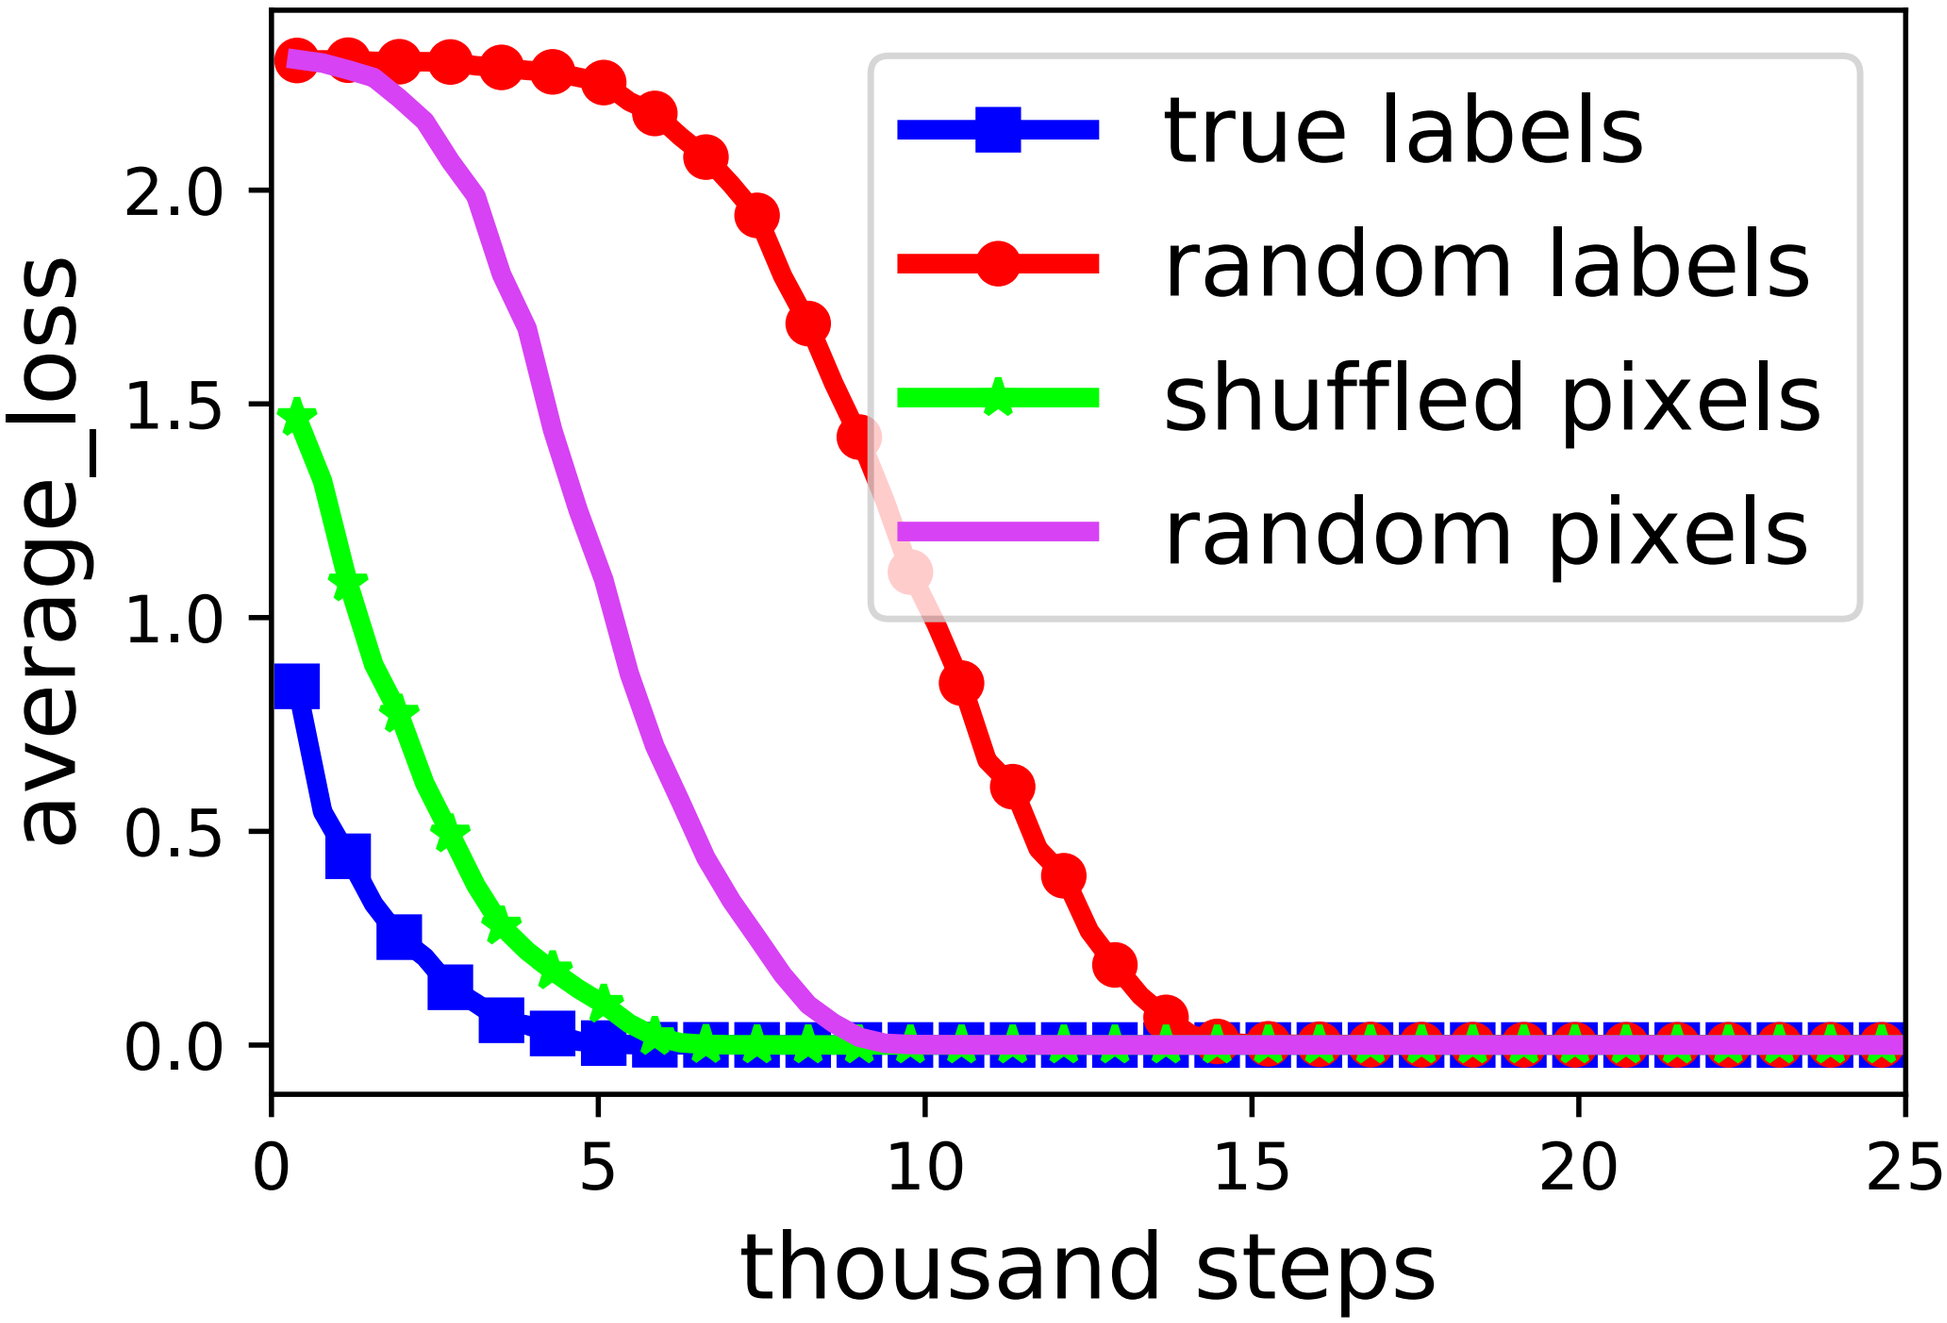

SGD fits Random Labels

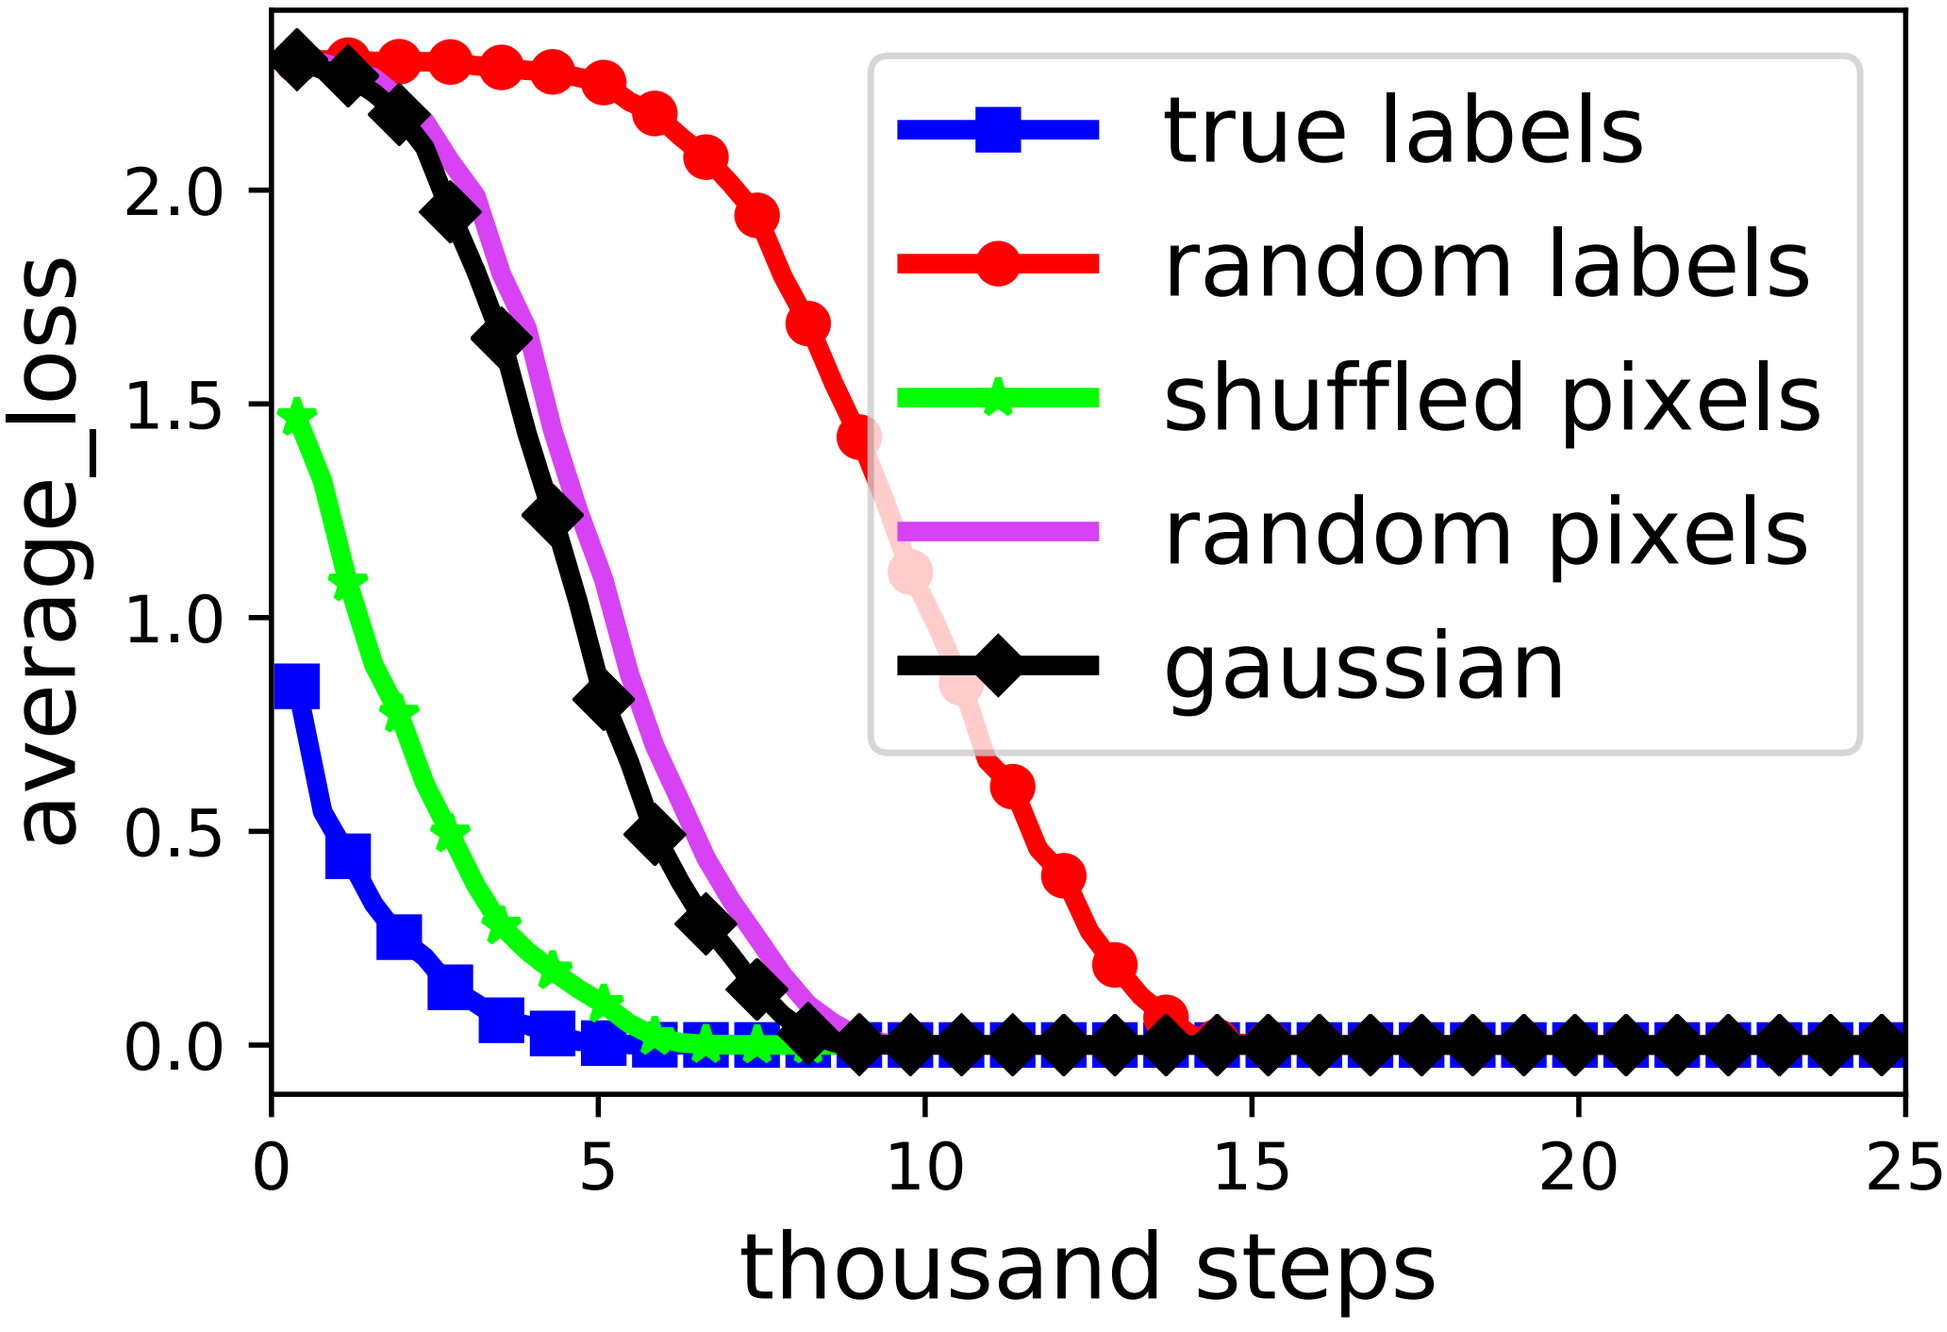

SGD fits Random Labels

Optimization is easy for deep learning.

Conclusion

Simple experimental framework for understanding the effective capacity of deep learning models

Successful DeepNets are able to shatter the training set

Other formal measures of complexity for the models / algorithms / data distributions are needed to precisely explain the over-parameterized regime

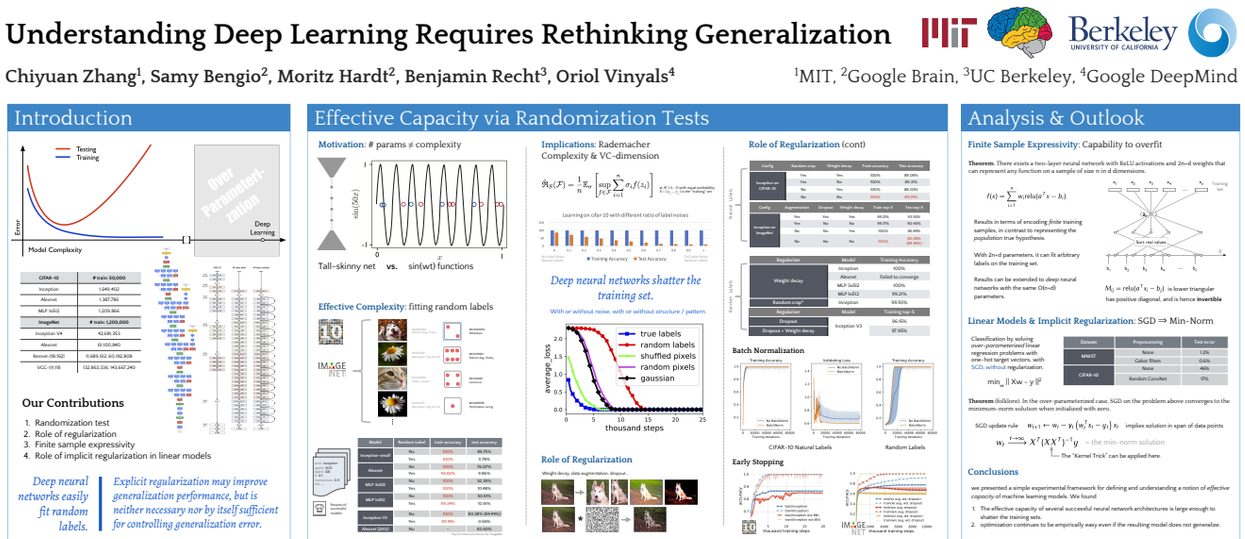

Poster: Wednesday Morning (April 26th, 10:30am to 12:30pm)

C23