

# DeepOBS

A Deep Learning Optimizer Benchmark Suite

Frank Schneider


30 Sep 2019

EBERHARD KARLS  
UNIVERSITÄT  
TÜBINGEN






Max Planck Institute for  
**Intelligent Systems**

**imprs-is**








## Everyone creates their own benchmark



-  Repeated work (including bugs)
-  Benchmarks are not comparable (not even the metric)
-  Potential for cherry-picking

## Doing proper benchmarks requires time

-  Use of rather small test problems (MNIST)
-  Use only a few test problems
-  Own optimizer gets more attention than the competition

# The Goal of DeepOBS

...we are not quite there yet!



1. Run the optimizer on a test problem
2. Compare to the state of the art
3. Plot the results

```
1 import tensorflow as tf
2 from deepobs import tensorflow as tfobs
3
4 optimizer_class = tf.train.RMSPropOptimizer
5 hyperparams = {"learning_rate": {"type": float},
6               "decay": {"type": float, "default": 0.9}}
7
8 runner = tfobs.runners.StandardRunner(optimizer_class, hyperparams)
9 runner.run()
```

Figure: run\_rmsprop.py

```
python run_rmsprop.py mnist_mlp --learning_rate=1e-3
```

```
*****
Evaluating after 0 of 100 epochs...
TRAIN: loss 353.934
VALID: loss 353.813
TEST: loss 353.447
*****
*****
Evaluating after 1 of 100 epochs...
TRAIN: loss 338.48
```

Data set	Model	Description	Conv	RNN	Drop	BN	WD
● 2D	Noisy Beale	Noisy version of the Beale function					
	Noisy Branin	Noisy version of the Branin function					
	Noisy Rosenbrock	Noisy version of the Rosenbrock function					
● Quadratic	N-Dimensional	100-dimensional ill-conditioned noisy quadratic					
● MNIST	Log. Regr.	Logistic regression					
	MLP	Four layer fully-connected network					
	2c2d	Two conv. and two fully-connected layers	✓				
	VAE	Variational Autoencoder	✓		✓		
● FASHION MNIST	Log. Regr.	Logistic regression					
	MLP	Four layer fully-connected network					
	2c2d	Two conv. and two fully-connected layers	✓				
● MNIST	VAE	Variational Autoencoder	✓		✓		
● CIFAR-10	3c3d	Three conv. and three fully-connected layers	✓				✓
	VGG 16	Adapted version of VGG 16	✓		✓		✓
	VGG 19	Adapted version of VGG 19	✓		✓		✓
● CIFAR-100	3c3d	Three conv. and three fully-connected layers	✓				✓
	VGG 16	Adapted version of VGG 16	✓		✓		✓
	VGG 19	Adapted version of VGG 19	✓		✓		✓
	All-CNN-C	The all convolutional net	✓		✓		✓
	Wide ResNet-40-4	Wide Residual Network	✓			✓	✓
● SVHN	3c3d	Three conv. and three fully-connected layers	✓				✓
	Wide ResNet-16-4	Wide Residual Network	✓			✓	✓
● IMAGENET	VGG 16	VGG 16	✓		✓		✓
	VGG 19	VGG 19	✓		✓		✓
	Inception-v3	Inception-v3 network	✓		✓	✓	✓
● Tolstoi	CharRNN	Recurrent Neural Network for character-level language modeling		✓	✓		



```
results
├── quadratic_deep
│   └── RMSPropOptimizer
│       ├── num_epochs__100__batch_size__128__decay__9.e-01__learning_rate__1.e-02
│       │   ├── random_seed__42__2019-09-27-15-15-59.json
│       │   └── random_seed__43__2019-09-27-15-19-33.json
│       └── num_epochs__100__batch_size__128__decay__9.e-01__learning_rate__1.e-03
│           └── random_seed__42__2019-09-27-15-01-49.json
```

# Tuning your Hyperparameters


...a semi-automated process



```
1 import numpy as np
2 from torch.optim import SGD
3 from deepobs.pytorch.runners import StandardRunner
4 from deepobs.tuner import GridSearch
5
6 optimizer_class = SGD
7 hyperparams = {"lr": {"type": float}}
8
9 grid = {"lr": np.logspace(-5, 2, 6)}
10
11 tuner = GridSearch(optimizer_class, hyperparams, grid, runner=StandardRunner)
12
13 tuner.tune('quadratic_deep', rerun_best_setting=True)
```



```
deepobs_plot_results results/ --full
```



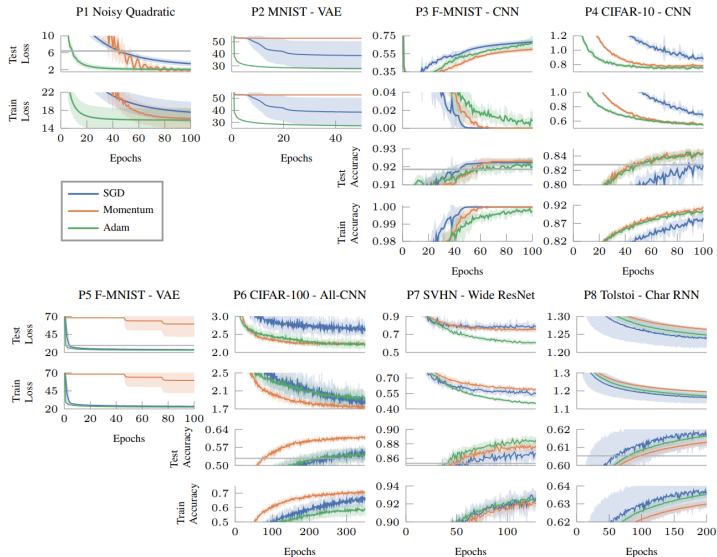


# Plotting the Results

...getting page 7 of your paper



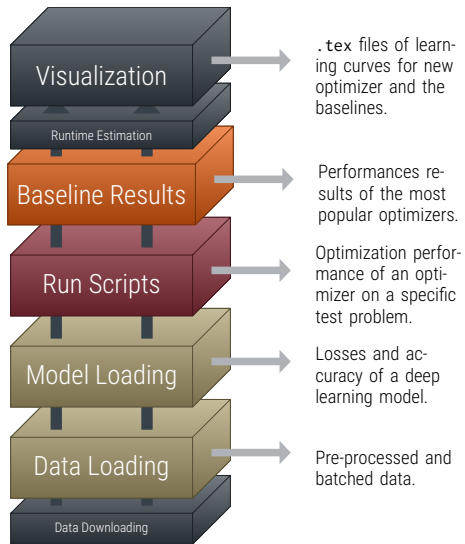
```
deepobs_plot_results results/ --full
```



```
deepobs_plot_results results/ --full
```

Table 2: DEEPOBS benchmark for the baseline optimizers.

Test Problem		SGD	Momentum	Adam	Test Problem		SGD	Momentum	Adam
P1 Noisy Quadratic	Performance	3.42	1.92	2.08	P5 F-MNIST VAE	Performance	23.94	59.44	23.11
	Speed	45.60	36.90	9.00		Speed	3.90	93.30	1.50
P2 MNIST VAE	Tuneability	$\alpha: 3.98e-03$	$\alpha: 3.98e-04$ $\mu: 0.99$	$\alpha: 1.00e-01$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$	P6 CIFAR-100 All CNN C	Tuneability	$\alpha: 3.98e-03$	$\alpha: 2.51e-04$ $\mu: 0.99$	$\alpha: 1.58e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$
	Performance	38.34	52.97	27.86		Performance	55.39 %	60.79 %	54.34 %
P3 F-MNIST CNN	Speed	1.00	1.00	1.00	P7 SVHN Wide ResNet	Speed	167.40	72.20	194.00
	Tuneability	$\alpha: 3.98e-03$	$\alpha: 1.58e-05$ $\mu: 0.99$	$\alpha: 1.58e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$		Tuneability	$\alpha: 1.58e-01$	$\alpha: 3.98e-03$ $\mu: 0.99$	$\alpha: 1.00e-03$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$
P4 CIFAR-10 CNN	Performance	92.25 %	92.32 %	92.03 %	P8 TOLSTOI Char RNN	Performance	86.69 %	87.74 %	88.53 %
	Speed	38.70	51.10	39.20		Speed	40.00	42.10	34.60
P4 CIFAR-10 CNN	Tuneability	$\alpha: 1.58e-01$	$\alpha: 1.00e-03$ $\mu: 0.99$	$\alpha: 2.51e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$	P8 TOLSTOI Char RNN	Tuneability	$\alpha: 2.51e-01$	$\alpha: 3.98e-03$ $\mu: 0.99$	$\alpha: 6.31e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$
	Performance	82.76 %	84.53 %	84.30 %		Performance	61.71 %	61.29 %	61.68 %
P4 CIFAR-10 CNN	Speed	77.10	41.00	44.60	P8 TOLSTOI Char RNN	Speed	57.90	96.30	79.20
	Tuneability	$\alpha: 1.58e-02$	$\alpha: 3.98e-04$ $\mu: 0.99$	$\alpha: 2.51e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$		Tuneability	$\alpha: 2.51e+00$	$\alpha: 3.98e-02$ $\mu: 0.99$	$\alpha: 6.31e-04$ $\epsilon: 1e-08$ $\beta_1: 0.9$ $\beta_2: 0.999$





## deepobs.github.io

### Leaderboard

Overview over the current optimizer leaderboard on the DeepOBS test problems. Click on See Full Results to see the plots and tables of the full benchmarking results.

#### Quadratic Deep

A 100-dimensional noisy quadratic problem with an eigenspectrum similar to the one reported for deep neural networks.

-	Optimizer	Test Loss	Speed
#1	Momentum	87.05	70.5
#2	Adam	87.11	39.9
#3	SGD	87.40	51.1

[See Full Results](#)

#### MNIST - VAE

A basic variational autoencoder for the MNIST data set with three convolutional and three deconvolutional layers.

-	Optimizer	Test Loss	Speed
#1	Adam	27.83	1.0
#2	SGD	38.46	1.0
#3	Momentum	52.93	1.0

[See Full Results](#)

#### F-MNIST - CNN

A simple convolutional network for the Fashion-MNIST data set, consisting of two conv and two fully-connected layers.

-	Optimizer	Test Accuracy	Speed
#1	Adam	92.34 %	40.1
#2	SGD	92.27 %	40.6
#3	Momentum	92.14 %	59.1

[See Full Results](#)

#### CIFAR-10 - CNN

A slightly larger convolutional network for the Cifar-10 data set, with three conv and three fully-connected layers.

-	Optimizer	Test Accuracy	Speed
#1	Adam	84.75 %	36.0
#2	Momentum	84.41 %	40.7
#3	SGD	83.71 %	42.5

[See Full Results](#)

#### F-MNIST - VAE

A basic variational autoencoder for the Fashion-MNIST data set with three convolutional and three deconvolutional layers.

-	Optimizer	Test Loss	Speed
#1	Adam	23.07	1.0
#2	SGD	23.80	1.0
#3	Momentum	59.23	1.0

[See Full Results](#)

#### CIFAR-100 - All CNN C

Variante C of the All Convolutional Network from *Striving for Simplicity* for the CIFAR-100 data set consisting solely of convolutional layers.

-	Optimizer	Test Accuracy	Speed
#1	Momentum	60.33 %	72.8
#2	SGD	57.06 %	128.7
#3	Adam	56.15 %	152.6

[See Full Results](#)

#### SVHN - Wide ResNet 16-4

The **Wide ResNet 16-4** for the Street View House Numbers data set using the variant with 16 conv layers and a widening factor of 4.

-	Optimizer	Test Accuracy	Speed
#1	Momentum	95.53 %	10.8
#2	SGD	95.37 %	28.3
#3	Adam	95.25 %	12.1

[See Full Results](#)

#### Tolstoj - Char RNN

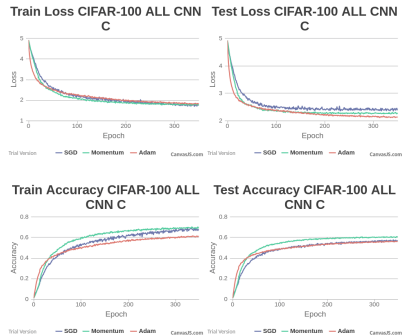
A recurrent neural network for character-level language modeling on the novel *War and Peace* by Leo Tolstoj using two LSTM layers.

-	Optimizer	Test Accuracy	Speed
#1	SGD	62.07 %	47.7
#2	Momentum	61.30 %	88.0
#3	Adam	61.23 %	62.8

[See Full Results](#)

deepobs.github.io

## Leaderboard CIFAR-100 - All CNN C



### Overview Table

Rank	Optimizer	Final Test Accuracy	Best Test Accuracy	Final Train Accuracy	Best Train Accuracy	Final Test Loss	Best Test Loss	Final Train Loss	Best Train Loss	Speed
1	Momentum	60.33%	60.83%	69.25%	70.25%	2.29	2.24	1.79	1.75	72.8
2	SGD	57.06%	57.77%	68.48%	69.31%	2.39	2.36	1.74	1.70	128.7
3	Adam	56.15%	56.41%	60.89%	61.52%	2.13	2.12	1.83	1.81	152.6

# The DeepOBS ecosystem

...lots of ways to participate



>\_ Install DeepOBS

```
pip install deepobs
```



PyTorch  
Coming soon

Check out Github for the Beta of DeepOBS  
1.2.0

[fsschneider.github.io/DeepOBS](https://fsschneider.github.io/DeepOBS)

Visit [deepobs.github.io](https://deepobs.github.io)

## Leaderboard

Overview over the current optimizer leaderboard on the DeepOBS test problems. Click on See Full Results to see the plots and tables of the full benchmarking results.

Quadratic Deep			MNIST - VAE			F-MNIST - CNN			CIFAR-10 - CNN		
A 100-dimensional noisy quadratic problem with an eigenvalue spectrum similar to the one reported for deep neural networks.			A basic variational autoencoder for the MNIST data set with three convolutional and three deconvolutional layers.			A simple convolutional network for the Fashion-MNIST data set, consisting of two conv and two fully-connected layers.			A slightly larger convolutional network for the CIFAR-10 data set, with three conv and three fully-connected layers.		
Optimizer	Total Loss	Speed	Optimizer	Total Loss	Speed	Optimizer	Total Accuracy	Speed	Optimizer	Total Accuracy	Speed
H1 Momentum	47.05	51.5	H1 Adam	27.66	1.0	H1 Adam	92.84 %	45.1	H1 Adam	84.76 %	26.4
H1 Adadelta	87.15	28.9	H1 RMS	28.96	1.5	H1 RMS	92.27 %	40.9	H1 Momentum	89.45 %	46.7
H1 SGD	47.48	56.1	H1 Momentum	28.96	1.8	H1 Momentum	92.84 %	56.1	H1 SGD	89.73 %	45.8
<a href="#">See Full Results</a>			<a href="#">See Full Results</a>			<a href="#">See Full Results</a>			<a href="#">See Full Results</a>		
F-MNIST - VAE			CIFAR-100 - All CNN C			SVHN - Wide ResNet 16-4			Toitotl - Char RNN		
A basic variational autoencoder for the Fashion-MNIST data set with three convolutional and three deconvolutional layers.			Variant C of the All Convolutional Network from <a href="#">Rotating for Simplicity</a> for the CIFAR-100 data set consisting solely of convolutional layers.			The <a href="#">Wide ResNet 16-4</a> for the Street View House Numbers data set using the variant with 16 conv layers and a widening factor of 4.			A recurrent neural network for character-level language modeling on the novel <i>War and Peace</i> by Leo Tolstoy using two LSTM layers.		
Optimizer	Total Loss	Speed	Optimizer	Total Accuracy	Speed	Optimizer	Total Accuracy	Speed	Optimizer	Total Accuracy	Speed
H1 Adam	23.07	1.0	H1 Momentum	69.92 %	15.8	H1 Momentum	98.92 %	22.9	H1 RMS	82.07 %	47.7
H1 SGD	28.88	1.8	H1 SGD	67.96 %	16.2	H1 SGD	98.92 %	26.5	H1 Momentum	84.93 %	46.4
H1 Momentum	38.23	1.8	H1 Adam	68.53 %	100.9	H1 Adam	98.92 %	12.1	H1 Adam	83.29 %	42.4
<a href="#">See Full Results</a>			<a href="#">See Full Results</a>			<a href="#">See Full Results</a>			<a href="#">See Full Results</a>		