

Linear Unscaled results

Joy

$lr = 0.001$

Really bad, mean absolute error of 1.3ish at the end.

$lr = 0.0001$

Seemed to overfit, learning rate continued to decrease. for training but not validation.

```
2020-05-14 18:53:12,495 - INFO - allennlp.common.util - Metrics: {
  "best_epoch": 11,
  "peak_cpu_memory_MB": 2630.66,
  "peak_gpu_0_memory_MB": 1,
  "peak_gpu_1_memory_MB": 21478,
  "training_duration": "0:23:01.688213",
  "training_start_epoch": 0,
  "training_epochs": 20,
  "epoch": 20,
  "training_pearson": 0.9952644629823881,
  "training_mae": 0.12786512176923553,
  "training_loss": 0.02625432804009868,
  "training_cpu_memory_MB": 2630.66,
  "training_gpu_0_memory_MB": 1,
  "training_gpu_1_memory_MB": 18680,
  "validation_pearson": 0.8548239304162555,
  "validation_mae": 0.6628524492371757,
  "validation_loss": 0.7741623421510061,
```

```
"best_validation_pearson": 0.8563759958887212,  
"best_validation_mae": 0.6515413680166569,  
"best_validation_loss": 0.7303319076697031  
}
```

lr = 0.00001

```
2020-05-14 21:39:30,893 - INFO - allennlp.common.util - Metrics: {  
  "best_epoch": 2,  
  "peak_cpu_memory_MB": 2627.62,  
  "peak_gpu_0_memory_MB": 1,  
  "peak_gpu_1_memory_MB": 21478,  
  "training_duration": "0:13:09.258866",  
  "training_start_epoch": 0,  
  "training_epochs": 11,  
  "epoch": 11,  
  "training_pearson": 0.9709087494155519,  
  "training_mae": 0.31317798209277703,  
  "training_loss": 0.15988704243909965,  
  "training_cpu_memory_MB": 2627.62,  
  "training_gpu_0_memory_MB": 1,  
  "training_gpu_1_memory_MB": 18094,  
  "validation_pearson": 0.8498408919748213,  
  "validation_mae": 0.721538697934215,  
  "validation_loss": 0.8732728213071823,  
  "best_validation_pearson": 0.8493931180602585,  
  "best_validation_mae": 0.6772722928029187,  
  "best_validation_loss": 0.7859473476807276  
}
```

Revision #6

Created Thu, May 14, 2020 10:26 PM by [kenneth](#)

Updated Fri, May 15, 2020 1:48 AM by [kenneth](#)