

Neural Networks

- RELU unscaled results
- Linear Unscaled results
- RNN scaled results
- RNN unscaled results

RELU unscaled results

Joy train on 9 folds, test on 1

```
2020-05-13 12:15:50,734 - INFO - allennlp.common.util - Metrics: {  
  "best_epoch": 12,  
  "peak_cpu_memory_MB": 2626.072,  
  "peak_gpu_0_memory_MB": 8085,  
  "peak_gpu_1_memory_MB": 21478,  
  "training_duration": "0:23:41.867688",  
  "training_start_epoch": 0,  
  "training_epochs": 21,  
  "epoch": 21,  
  "training_pearson": 0.9939782728494088,  
  "training_mae": 0.14345509548909208,  
  "training_loss": 0.03389307007519076,  
  "training_cpu_memory_MB": 2626.072,  
  "training_gpu_0_memory_MB": 7409,  
  "training_gpu_1_memory_MB": 18680,  
  "validation_pearson": 0.8524911266953036,  
  "validation_mae": 0.7244842611554498,  
  "validation_loss": 0.8840915312369665,  
  "best_validation_pearson": 0.8559774687113151,  
  "best_validation_mae": 0.6610433984638224,  
  "best_validation_loss": 0.7440112556020418  
}
```

Sadness train on 9 folds, test on 1

These are longer and they're causing clipping. I don't know how many are causing clipping though because allennlp only reports the first case of clipping.

```
2020-05-13 12:22:49,449 - INFO - allennlp.common.util - Metrics: {
  "best_epoch": 1,
  "peak_cpu_memory_MB": 2759.188,
  "peak_gpu_0_memory_MB": 7409,
  "peak_gpu_1_memory_MB": 18680,
  "training_duration": "0:17:02.438299",
  "training_start_epoch": 0,
  "training_epochs": 10,
  "epoch": 10,
  "training_pearson": -0.01633020200422353,
  "training_mae": 1.3776687749682646,
  "training_loss": 2.781699788029837,
  "training_cpu_memory_MB": 2759.188,
  "training_gpu_0_memory_MB": 7409,
  "training_gpu_1_memory_MB": 11,
  "validation_pearson": 0,
  "validation_mae": 1.355329878786777,
  "validation_loss": 2.7624370823515223,
  "best_validation_pearson": 0.23847302176509885,  "best_validation_mae":
1.3548242284896526,
  "best_validation_loss": 2.7604094781774156
}
```

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  
}
```

RNN scaled results

3 hid

lr 0.00001

```
2020-05-14 22:57:02,171 - INFO - allennlp.common.util - Metrics: {
  "best_epoch": 28,
  "peak_cpu_memory_MB": 1997.92,
  "peak_gpu_0_memory_MB": 2697,
  "peak_gpu_1_memory_MB": 11,
  "training_duration": "0:06:15.555519",
  "training_start_epoch": 0,
  "training_epochs": 37,
  "epoch": 37,
  "training_pearson": 0.8835897766720373,
  "training_mae": 0.6050379267542354,
  "training_loss": 0.012446582688072931,
  "training_cpu_memory_MB": 1997.892,
  "training_gpu_0_memory_MB": 2697,
  "training_gpu_1_memory_MB": 11,
  "validation_pearson": 0.7651979545423095,
  "validation_mae": 0.854339599609375,
  "validation_loss": 0.02388349245302379,
  "best_validation_pearson": 0.763310168046131,
  "best_validation_mae": 0.8508703021026365,
  "best_validation_loss": 0.023334396770223975
}
```

lr 0.0001

```
2020-05-14 23:00:13,281 - INFO - allennlp.common.util - Metrics: {  
  "best_epoch": 3,  
  "peak_cpu_memory_MB": 1984.856,  
  "peak_gpu_0_memory_MB": 2437,  
  "peak_gpu_1_memory_MB": 11,  
  "training_duration": "0:02:09.181165",  
  "training_start_epoch": 0,  
  "training_epochs": 12,  
  "epoch": 12,  
  "training_pearson": 0.9309341933147448,  
  "training_mae": 0.4662433107257326,  
  "training_loss": 0.007578380315483195,  
  "training_cpu_memory_MB": 1984.816,  
  "training_gpu_0_memory_MB": 2437,  
  "training_gpu_1_memory_MB": 11,  
  "validation_pearson": 0.7313824688523022,  
  "validation_mae": 0.8880522283261034,  
  "validation_loss": 0.028224910454203684,  
  "best_validation_pearson": 0.7701748503103649,  
  "best_validation_mae": 0.8408898201914168,  
  "best_validation_loss": 0.022758180275559425  
}
```

2 hid

RNN unscaled results

Glove embeddings

64 dimensional RNN

1 hidden layer

2 hidden layers

0.0001 lr

```
2020-05-19 23:19:25,110 - INFO - allennlp.common.util - Metrics: {  
  "best_epoch": 35,  
  "peak_cpu_memory_MB": 1715.32,  
  "peak_gpu_0_memory_MB": 1,  
  "peak_gpu_1_memory_MB": 636,  
  "training_duration": "0:01:44.572271",  
  "training_start_epoch": 0,  
  "training_epochs": 44,  
  "epoch": 44,  
  "training_pearson": 0.8755669240062097,  
  "training_mae": 0.6300580370557177,  
  "training_loss": 0.6342626507793154,  
  "training_cpu_memory_MB": 1715.3,  
  "training_gpu_0_memory_MB": 1,
```



```
"training_gpu_1_memory_MB": 636,  
"validation_pearson": 0.7133927439166854,  
"validation_mae": 0.913505204604321,  
"validation_loss": 1.405275821685791,  
"best_validation_pearson": 0.7112064102524507,  
"best_validation_mae": 0.9231068901617251,  
"best_validation_loss": 1.3673358609278996  
}
```

3 hidden layers