

Happiness narratives

- Data Discrepancies
 - The stats
- Results
 - SVR Joy
 - SVM Joy
 - Lasso Joy
 - Lasso Sadness
 - SVR Sadness
 - SVM Sadness
- Correlation
- Neural Networks
 - RELU unscaled results
 - Linear Unscaled results
 - RNN scaled results
 - RNN unscaled results
- Meeting Notes
 - 5-13

Data Discrepancies

There's more data used for the hedometer evaluation than there is for the ml training and evaluation.

Data Discrepancies

The stats

The Hedometer data has 3656 entries The ML data has 3343 enttries.

The preprocessing done on the hedometer data makes it difficult to compare the two files.

Lowercasing, all punctuation removed, spaced injeted.

Results

Results

SVR Joy

MSE: 0.9964795758847595

Mean Absolute Error 0.8209001542496819

R2: -0.0969678363700448

Pearson correlation: (0.1630455970822213, 3.789529650766423e-23)0 values were squeezed into the score range

```
0 values were squeezed into the score range          precision    recall  f1-score
support
      1  0.00000000  0.00000000  0.00000000          288          2  0.32646048  0.21517554
0.25938567          883
      3  0.32398754  0.64148034  0.43053040          1297          4  0.40041494  0.19108911
0.25871314          1010
      5  0.00000000  0.00000000  0.00000000          169
accuracy          0.33315053          3647  macro avg  0.21017259  0.20954900
0.18972584          3647
weighted avg  0.30515370  0.33315053  0.28756121          3647
[[ 0  36  234  18  0]
 [ 5 190  640  48  0]
 [ 2 251  832 208  4]
 [ 0  93  720 193  4]
 [ 0  12  142  15  0]]
ROC - AUC : 0.5053219775610537
```


Results

Lasso Joy

```
MSE: 1.0281566480572382
Mean Absolute Error 0.8358763189438444
R2: -0.1318393281341672
Pearson correlation: (0.14760188286173698, 3.252804550962394e-19)3 values were squeezed into
the score range
0 values were squeezed into the score range
precision recall f1-score
support
      1 0.10000000 0.00347222 0.00671141      288      2 0.30782030 0.20951302
0.24932615      883
      3 0.32254244 0.62991519 0.42663185     1297      4 0.42738589 0.20396040
0.27613941     1010
      5 0.04761905 0.00591716 0.01052632      169
accuracy          0.33177954    3647 macro avg 0.24107354 0.21055560
0.19386703     3647
weighted avg 0.31769954 0.33177954 0.28958298    3647
[[ 1 26 247 14  0]
 [ 4 185 641 52  1]
 [ 5 270 817 196  9]
 [ 0 105 689 206 10]
 [ 0 15 139 14  1]]
ROC - AUC : 0.5058964520950757
```

Results

Lasso Sadness

SE: 0.9307384965022848

Mean Absolute Error 0.7960851914177507

R2: -0.04084814009093862

Pearson correlation: (0.23080343775556006, 3.2202046350021302e-46)20 values were squeezed into the score range

0 values were squeezed into the score range
support

precision recall f1-score

1	0.00000000	0.00000000	0.00000000	146	2	0.06217617	0.01826484
0.02823529	657						

3	0.33868243	0.64677419	0.44456763	1240	4	0.46069470	0.38858905
0.42158093	1297						

5	0.00000000	0.00000000	0.00000000	384			
---	------------	------------	------------	-----	--	--	--

accuracy		0.35392052	3724	macro avg	0.17231066	0.21072562
0.17887677	3724					

weighted avg	0.28419360	0.35392052	0.29984020	3724			
--------------	------------	------------	------------	------	--	--	--

[[0 2 125 19 0]

[1 12 551 92 1]

[3 136 802 281 18]

[1 41 706 504 45]

[0 2 184 198 0]]

ROC - AUC : 0.5100520397496775

.....

Results

SVM Sadness

MSE: 0.6640974977269949

Mean Absolute Error 0.6386105570026086

R2: 0.2573374283476926

Pearson correlation: (0.5643611440598052, 2.75759376237e-312)0 values were squeezed into the score range

```
0 values were squeezed into the score range          precision    recall  f1-score
support
      1  0.14285714  0.00684932  0.01307190          146          2  0.42553191  0.30441400
0.35492458          657
      3  0.43119266  0.53064516  0.47577730          1240          4  0.48706625  0.59521974
0.53573907          1297
      5  0.41176471  0.14583333  0.21538462          384
accuracy          0.45300752          3724  macro avg  0.37968253  0.31659231
0.31897949          3724
weighted avg  0.43634615  0.45300752  0.43034883          3724
[[ 1  56  73  16   0]
 [ 5 200 365  86   1]
 [ 1 162 658 414   5]
 [ 0  45 406 772  74]
 [ 0   7  24 297  56]]
ROC - AUC : 0.578488077944395
```

Correlation

Binyan was doing pearson's correlation. Thinks maybe we should do spearman

Neural Networks

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

Meeting Notes

Meeting Notes

5-13

Variety of different conferences and journals we could submit to.

Digital Scholarship in the Humanities is a DH journal we could submit to.