

CS 342: Section 5

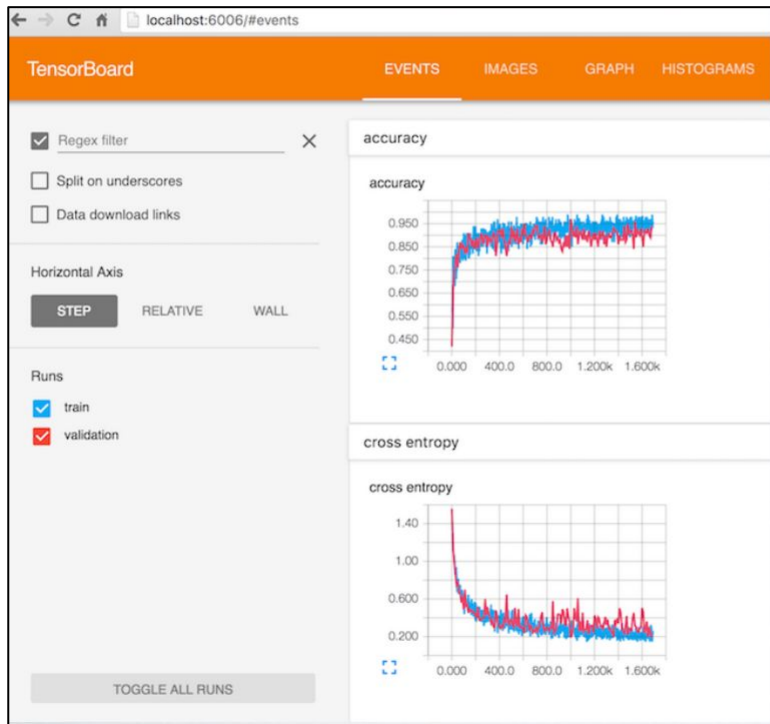
- Homework 4
 - Due 9/27, Can still submit by 9/30 (with penalties)
- Homework 5 released
 - Check course webpage
- Homework 3
 - Grades and reports posted on Canvas

Questions?

Preventing Overfitting

- Early Stopping
- Ensembles
- Data Augmentation
- Dropout
- Weight Regularization
- PyTorch walk through

Early Stopping



Pick the Model with best validation accuracy

Early Stopping

- Save models every T iterations
- Or watch it :P

```
While ...  
    Try:  
        ...  
    Except KeyboardInterrupt  
        Save model ...
```

Ensemble

- Train several models and combine them
 - Bias-Variance tradeoff (regression as e.g.)
 - $MSE(\hat{y}) = \text{bias}(\hat{y}, y) + \text{var}(\hat{y})$, where $\text{bias}(\hat{y}, y) := y - E(\hat{y})$
 - Bias: error from wrong model assumption (underfitting)
 - Variance: error from sensitivity to fluctuations (overfitting)
 - Intuition
 - Combining multiple models to reduce variance

Data Augmentation

- Intuition
 - Increase the number of data
 - Force the model not to “memorize” data, learn robust features that are invariant to these augmentation changes (e.g. random cropping, flipping, rotation, hue changes etc.)

Dropout

- Randomly drop weights in network with prob p
 - Keep (expected) magnitude the same by multiplying $1/p$
- Intuition
 - Implicit model ensembling, 2^H models
 - Learning more robust features
- Refer to lecture slides and

<https://www.cs.toronto.edu/~hinton/absps/JMLRdropout.pdf> for more details

Weight regularization

- Add a penalty (typically L2) on the weights to force them to be small
 - Intuition: penalize peaky weights
 - Interpretation from a probabilistic/bayesian view (e.g. linear regression)
 - Suppose a gaussian prior on weights $\theta \sim N(0, \Sigma)$
 - MAP: $\operatorname{argmax} \log P(\theta|y) = \operatorname{argmax} \log P(y|\theta)P(\theta) = \operatorname{argmax} \log P(y|\theta) + \log P(\theta) = \operatorname{argmax} \log P(y|\theta) + \|\theta\|^2$

Let's translate these into code now!