

Lesson 36 Summary: Overfitting Prevention

Data Science with Python – Key Concepts

Data Science Program

Overfitting Prevention in Deep Learning

Overfitting

Memorizes training
Poor generalization
Train \ll Test error

Dropout

Randomly drop neurons
Typical: 20-50%
Ensemble effect

Regularization

L1/L2 weight penalty
Smaller weights
Simpler model

Early Stopping

Monitor val loss
Stop when increases
Patience parameter

More Data

Data augmentation
Collect more samples
Synthetic generation

Batch Norm

Normalize activations
Regularizing effect
Faster training

Multiple techniques prevent neural networks from memorizing data

The Overfitting Problem

Why it happens:

- **Too many parameters:** Model capacity exceeds data
- **Training too long:** Memorizes noise
- **Symptom:** Low train error, high test error

Overfitting = poor generalization to new data

Random neuron deactivation:

- **Training:** Randomly zero neurons (20-50%)
- **Inference:** Use all neurons, scale outputs
- **Effect:** Prevents co-adaptation

Dropout creates implicit ensemble of networks

Weight penalty in loss:

- **L2 (Ridge):** $\lambda \sum w_i^2$ shrinks weights
- **L1 (Lasso):** $\lambda \sum |w_i|$ creates sparsity
- **Effect:** Simpler, smoother models

L2 is more common in deep learning (weight decay)

Stop before overfitting:

- **Monitor:** Validation loss each epoch
- **Stop when:** Val loss increases (patience)
- **Restore:** Best weights from training

Simple yet effective regularization technique

Normalize layer inputs:

- **Normalize:** Zero mean, unit variance per batch
- **Learn:** Scale and shift parameters
- **Bonus:** Regularizing effect, faster training

BatchNorm placed after linear, before activation

Increase effective dataset size:

- **Images:** Rotate, flip, crop, color shift
- **Time series:** Add noise, time warping
- **Finance:** Bootstrap, synthetic scenarios

More data is the best regularizer

Combine multiple techniques:

- **Start simple:** Small network first
- **Add capacity:** If underfitting
- **Add regularization:** If overfitting

Balance model capacity with regularization strength

Regularization Techniques:

Technique	How It Works
Dropout	Randomly zero neurons during training
L2/Weight decay	Penalize large weights
Early stopping	Stop when val loss increases
Batch norm	Normalize layer inputs
Data augmentation	Artificially expand dataset

Use multiple techniques together for best results