Mathematical Tools for Neural and Cognitive Science

Fall semester, 2023

Section 6

Model fitting:
comparison, selection and regularization


How do we avoid overfitting (i.e., concluding that $\mathrm{M}=7$ is "best")?

## Taxonomy of model-fitting errors

- Unexplainable variability (e.g., due to noisy measurements)
- Overfitting (too many params, not enough data)
- Optimization failures (e.g., local minima)
- Model failures (what you'd really like to know)


## Optimization...



## Model Comparison

- If models are optimized according to some objective, it is natural to compare them based on the value of that objective...
- for least squares regression, compare the residual squared error of two models (with different regressors).
- for ML estimates, compute the likelihood (or log likelihood) ratio, and compare to 1 (or zero).
- for MAP estimates, common to compute the posterior ratio
- Problem: evaluating the objective with the same data used to optimize the model leads to over-fitting! We really want to predict error on non-training data...


## Bayesian Model Comparison

- Eg: Is the coin fair? Compared to what?
- Consider two models: $M_{1}: p=0.5 \quad M_{2}: p=0.6$

$$
p\left(M_{k} \mid D\right)=\frac{p\left(D \mid M_{k}\right) P\left(M_{k}\right)}{p(D)}
$$

Compare their posterior ratio (the product of the Bayes factor, i.e., their likelihood ratio, and the prior odds):

$$
\frac{p\left(M_{1} \mid D\right)}{p\left(M_{2} \mid D\right)}=\frac{p\left(D \mid M_{1}\right) P\left(M_{1}\right)}{p\left(D \mid M_{2}\right) P\left(M_{2}\right)}
$$

## Comparing models' predictive performance

Option 1: Include a penalty for number of parameters:
For an ML estimate: $\quad \hat{\theta}=\arg \min _{\theta}[-\ln p(\vec{d} \mid \theta)]$
a. Akaike information criterion (AIC) [Akaike, 1974]

$$
E_{\mathrm{AIC}}(\vec{d}, \hat{\theta})=2 \operatorname{dim}(\hat{\theta})-2 \ln p(\vec{d} \mid \hat{\theta})
$$

b. Bayesian information criterion (BIC) [Schwartz, 1978]
$E_{\mathrm{BIC}}(\vec{d}, \hat{\theta})=\operatorname{dim}(\hat{\theta}) \ln [\operatorname{dim}(\vec{d})]-2 \ln p(\vec{d} \mid \hat{\theta})$
valid when $\operatorname{dim}(\vec{d}) \gg \operatorname{dim}(\hat{\theta})$
Option 2: Cross-validation
(simulated test of generalization to additional data)

## Cross-validation

A resampling method for estimating predictive error of a model. Widely used to identify/avoid over-fitting, and to provide a fair comparison of models.

Using cross-validation to select the degree of a polynomial model:
(1) Randomly partition data into a "training" set, and a "test" set.
(2) Fit model to training set. Measure error on test set.
(3) Repeat (many times).
(4) Choose model that minimizes the average crossvalidated ("test") error


## Ridge regression

(a.k.a. $L_{2}$ regularization)

Ordinary least squares regression:

$$
\arg \min _{\vec{\beta}}\|\vec{y}-X \vec{\beta}\|^{2}
$$

"Regularized" least squares regression:

$$
\arg \min _{\vec{\beta}}\|\vec{y}-X \vec{\beta}\|^{2}+\lambda\|\vec{\beta}\|^{2}
$$

Equivalent formulation: MAP estimate, assuming Gaussian likelihood \& prior!

$$
\hat{\beta}_{\text {ridge }}=\left(X^{T} X+\lambda I\right)^{-1} X^{T} \vec{y}
$$

Choose lambda by cross-validation:


Ridge Regression trades off bias and variance:

 solution is a "soft-threshold" on $\vec{y}^{T} \vec{x}$

MAP interpretation:
Gaussian noise, with
"Laplacian" prior


## Ridge regression vs. LASSO (2-dimensional) solutions



## LASSO vs. ridge regression



## The "Relaxed LASSO"

To reduce bias, re-solve the regression fit, after eliminating unused regressors (those with zero LASSO coefficients)

## Clustering

-K-Means (Lloyd, 1957)

- "Soft-assignment" version of K-means
(a form of Expectation-Maximization - EM)
- In general, alternate between:

1) Estimating cluster assignments (classification)
2) Estimating cluster parameters

- Coordinate descent: converges to (possibly local) minimum
- Need to choose K (number of clusters) - cross-validation!


## K-Means clustering algorithm

Alternate between two steps:

1. Estimate cluster assignments: given class centers, assign each point to closest one:
 estimate the centroid of each cluster.


## K-means clustering (Bishop Fig. 9.1)



## K-means optimization failures

Initialization matters (due to local minima) ...
Three solutions obtained with different random starting points:


## K-means systematic failures

Non-convex/non-round-shaped clusters


Clusters with different densities


Picture courtesy: Christof Monz (Queen Mary, Univ. of London)

ML for discrete mixture of Gaussians: soft K-means

$$
\begin{aligned}
& p\left(\vec{x}_{n} \mid a_{n k}, \vec{\mu}_{k}, \Lambda_{k}\right) \propto \sum_{k} \frac{a_{n k}}{\sqrt{\left|\Lambda_{k}\right|}} e^{-\left(\vec{x}_{n}-\vec{\mu}_{k}\right)^{T} \Lambda_{k}^{-1}\left(\vec{x}_{n}-\vec{\mu}_{k}\right) / 2} \\
& a_{n k} \quad=\text { assignment probability } \\
&\left\{\vec{\mu}_{k}, \Lambda_{k}\right\}=\text { mean/covariance of class } k
\end{aligned}
$$

Intuition: alternate between maximizing these two sets of variables ("coordinate descent")

Essentially, a version of K-means with "soft" (i.e., continuous, as opposed to binary) assignments!


EM: Gaussian mixture model (Bishop Fig. 9.5)


EM: Gaussian mixture model (Bishop Fig. 9.8)


## Application to neural "spike sorting"



Standard solution:

1. Threshold to find segments containing spikes
2. Reduce dimensionality of segments using PCA
3. Identify spikes using clustering (e.g., K-means)

Note: Fails for overlapping spikes!

Failures of clustering for near-synchronous spikes
us spiking


