Model Selection and Estimation of Generalization Cost

Volker Tresp
Summer 2019

Generalization Cost and Training Cost

Cost Functions

- Let's consider the prediction of a model $\mathcal{M}$ for input $\mathbf{x}$ and with parameters $\mathbf{w}$

$$f(\mathbf{x}, \mathbf{w}, \mathcal{M})$$

- Example: $\mathcal{M}$ can be a neural network with a specific architecture. Another example: $\mathcal{M}$ is a linear regression model

- We define a cost function for a data point $\mathbf{x}, y$

$$\text{cost}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}]$$

- We will use the terms **cost**, **loss** and **error** exchangeably

- Example (quadratic cost):

$$\text{cost}^q_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = (y - f(\mathbf{x}, \mathbf{w}, \mathcal{M}))^2$$

Cost Functions (cont'd)

- Misclassification cost ($y \in \{-1, 1\}$):

$$\text{cost}^{m}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = \frac{1}{2}|y - \text{sign}(f(\mathbf{x}, \mathbf{w}, \mathcal{M}))|$$

- Absolute deviation (AD) ($y \in \{-1, 1\}$):

$$\text{cost}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = \frac{1}{2}|y - f(\mathbf{x}, \mathbf{w}, \mathcal{M})|$$

- Cross-entropy cost function for logistic regression ($y \in \{-1, 1\}$):

$$\text{cost}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = \sum_{i=1}^{N} \log\left(1 + \exp\left[-y_i f(\mathbf{x}, \mathbf{w}, \mathcal{M})\right]\right)$$

Cost Functions (cont'd)

- Perceptron ($y \in \{-1, 1\}$):

$$\text{cost}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = |-yf(\mathbf{x}, \mathbf{w}, \mathcal{M})|_+$$

- Vapnik's optimal hyperplanes ($y \in \{-1, 1\}$):

$$\text{cost}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = |1 - yf(\mathbf{x}, \mathbf{w}, \mathcal{M})|_+$$

- Cross-entropy cost (negative log-likelihood cost)

$$\text{cost}^l_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = -\log P(y|f(\mathbf{x}, \mathbf{w}, \mathcal{M})$$

For binary classification, identical to the logistic regression cost function

Model Selection Based on Generalization Cost

- In statistics one is often interested in the estimation of the value and the uncertainty of particular parameters. Example: is parameter w_1 significantly nonzero?

- In machine learning one is often interested in the **generalization cost** which is the average expected cost over all possible data, **for any** fixed $\mathbf{w}$ and model $\mathcal{M}$

$$\text{cost}_{P(\mathbf{x}, y)}[\mathbf{w}, \mathcal{M}] = \int \text{cost}_{\mathbf{x}, y}[\mathbf{w}, \mathcal{M}] P(\mathbf{x}, y) \ d\mathbf{x} dy$$

- A typical assumption is that $P(\mathbf{x}, y) = P(\mathbf{x})P(y|\mathbf{x})$ is fixed but unknown

Average Test Set Cost Estimates the Generalisation Cost

- An estimator of the generalization cost is the **average test set cost**

$$\text{cost}_{\text{test}}[\mathbf{w}, \mathcal{M}] = \frac{1}{T} \sum_{i=1}^{T} \text{cost}_{\mathbf{x}_i, y_i}[\mathbf{w}, \mathcal{M}]$$

  which is the average cost on the T test data points with $(\mathbf{x}_i, y_i \in \text{test})$

- The test data are data points not used in training

- This is an unbiased estimator of the generalization cost, for any fixed $\mathbf{w}$ and any model $\mathcal{M}$; we can compare different models based on their average test set performance

- The variance of this estimator approaches zero for $T \to \infty$; typically one does not want to reserve a large set of available data as test data; a better alternative is a cross validation approach, described later

Average Training Set Cost

- To obtain a better understanding of model performance, one is interested in the relationship between average training set cost and generalization cost

- We define the **average training set cost** of the best parameter vector, trained and evaluated **on the training data** as

$$\mathsf{cost}_{\mathsf{train}}[\widehat{\mathbf{w}}(\mathsf{train}), \mathcal{M}] = \frac{1}{N} \sum_{i=1}^{N} \mathsf{cost}_{\mathbf{x}_i, y_i}[\widehat{\mathbf{w}}(\mathsf{train}), \mathcal{M}]$$

Here, $(\mathbf{x}_i, y_i) \in$ train; we use train and D interchangeably

Analysis of Different Quantities

- So far we were interested in is the generalization cost of a particular model $\mathcal{M}$ with particular best-fit parameters $\widehat{\mathbf{w}}(\text{train})$.

- Another quantity of interest is the generalization cost **averaged over all possible training sets of size** N (where the training data set is generated from $P(\mathbf{x}, y)$)

Training Set Cost and Generalization Cost

- It turns out that if we calculate the expected average over all training sets of the same size, then

$$\mathbb{E}_{\text{train}} \left\{ \text{cost}_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] - \text{cost}_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \right\} \geq 0$$

- Thus in expectation, the average training cost underestimates the generalization cost for the estimated $\widehat{\mathbf{w}}$, optimized on the training data. Thus the performance of a trained model should not be evaluated on the training set but on the test set, which is an unbiased estimator of the generalization cost

- This expression is the focus in a frequentist analysis!

Preview: Theoretical Analysis

- Consider the special case of models with fixed basis functions and least squares cost and let's assume the model is free of bias (no regularization, no structural bias)

- Then the generalization cost of the true model (the best possible model)

$$\text{cost}^q_{P(\mathbf{x},y)}[f(\cdot), \mathcal{M}] = \text{Residual} = \sigma^2$$

  If our training procedure would identify the true parameters, this would be the generalization cost

- The expected generalization cost of the fitted model is

$$\mathbb{E}_{\text{train}}\left\{\text{cost}^q_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}]\right\} = \text{Residual} + \text{Var}$$

  This term is estimated by the average test set cost; in expectation, it is **larger** than the generalization costs of the best model by a term called the variance Var

Preview: Theoretical Analysis (cont'd)

- The expected average training set cost of the fitted model in some cases is

$$\mathbb{E}_{\text{train}} \left\{ \text{cost}_{train}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \right\} = \text{Residual} - \text{Var}$$

  This term is estimated by the average training costs; this term is **smaller** than the generalization costs of the best model by the variance Var

- Then we obtain

$$\mathbb{E}_{\text{train}} \left\{ \text{cost}_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] - \text{cost}_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \right\} = 2\text{Var}$$
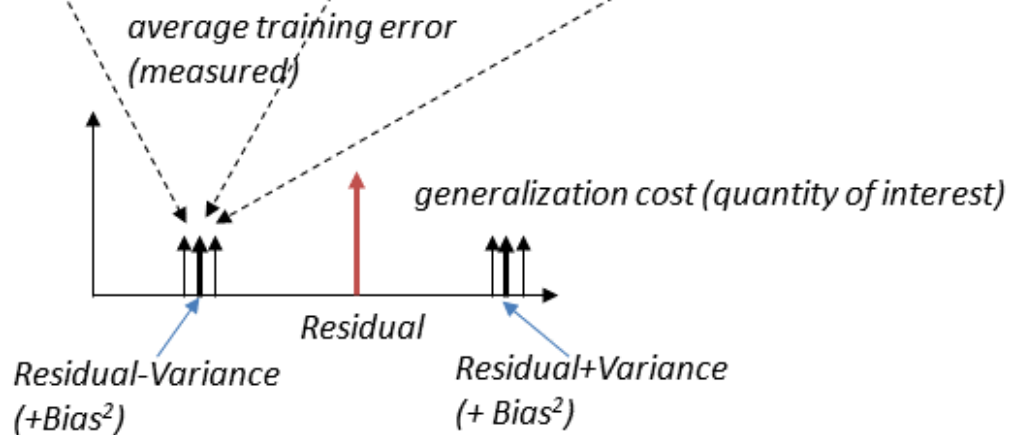
- For certain models, we can estimate Var

$x \sim P(x), y \sim f(x) + epsilon$

Fixed Basis functions with additive noise

$train_1$ ... ... $train_L$

w_1 ... ... w_L

average training error (measured)

generalization cost (quantity of interest)

Residual

Residual-Variance (+Bias^2)

Residual+Variance (+ Bias^2)

$L \rightarrow infinity$

Empirical Model Comparison

Model Selection via Training and Test Data Performance

- This procedure can be applied with huge amounts of available data, $N >> M_p$, where M_p is the number of model parameters

- This procedure is typically used in deep learning with large data sets, where a cross validation approach would be too costly

- Divide the data set randomly into a training data set and a test data set

- Train all models only on the training data: find the best parameters for each model under consideration

- Evaluate the generalization performance based on the average test set performance and get $\text{cost}_{\text{test}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}]$ for the different models, as an estimate of the generalization costs $\text{cost}_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}]$

Cross Validation

- Cross validation uses all data in turn for testing

- Consider K- fold cross validation; typical: $K = 5$ oder $K = 10$

- The data is partitioned into K sets of approximately the same size

- For $k = 1, \ldots, K$: The $k-$th fold (test_k) is used for testing and the remaining data (train_k) is used for training (finding the best parameters)

Evaluating Performance with Cross Validation

- For each model one gets K test costs

$$\text{cost}_{\text{test}_k}[\widehat{\mathbf{w}}(\text{train}_k), \mathcal{M}], \quad k = 1, \ldots, K$$

- Now we now consider the generalization costs averaged over the parameter estimates obtained from different training data sets of size N

- We can estimate this expectation as

$$\mathbb{E}_{\text{train}}\text{cost}_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \approx \frac{1}{K} \sum_{k=1}^{K} \text{cost}_{\text{test}_k}[\widehat{\mathbf{w}}(\text{train}_k), \mathcal{M}] = m(\mathcal{M})$$

Variance Estimate

- We can estimate the uncertainty of $m(\mathcal{M})$ as the variance

$$\widehat{Var}(\mathcal{M}) = \frac{1}{K(K-1)} \sum_{k=1}^{K} (\text{cost}_{\text{test}}[\widehat{\mathbf{w}}(\text{train}_k), \mathcal{M}] - m(\mathcal{M}))^2$$

- Mean and mean-variance estimates can be used to decide if two models significantly differ in generalization performance: typically, one accepts that model $\mathcal{M}_i$ has smaller generalization cost than $\mathcal{M}_j$, if

$$\mathcal{M}_i + \sqrt{\widehat{Var}(\mathcal{M}_i)} < \mathcal{M}_j - \sqrt{\widehat{Var}(\mathcal{M}_j)}$$

5-times cross validation:
Blue: Trainings Data
Red: Test Data

Paired Tests

- With few data one can use a paired test

- Basic idea: let's assume that $K = 10$; if $\mathcal{M}_i$ in all test sets is better than $\mathcal{M}_j$, then this is a strong indication that $\mathcal{M}_i$ performs better, even if the variation in test set performance masks this behavior (error bars of the estimate are too large)

- Calculate the average difference between both model costs

$$\text{MeanDiff}_{i,j} = \frac{1}{K} \sum_{k=1}^{K} \text{cost}_{\text{test}_k}[\widehat{\mathbf{w}}(\text{train}_k), \mathcal{M}_j - \text{cost}_{\text{test}_k}[\widehat{\mathbf{w}}(\text{train}_k), \mathcal{M}_i]]$$

and analyse if this difference is significantly larger than zero

Paired Tests (cont'd)

- In the case of high correlation, the variance of $\mathsf{MeanDiff}_{i,j}$ can be much smaller than the variances of $m(\mathcal{M}_i)$ and $m(\mathcal{M}_j)$, due to the rule $var(X - Y) = var(X) + var(Y) - 2cov(X,Y)$

- For a statistical analysis, one employs the test statistics for the paired t-test

- Alternative approach: Wilcoxon signed-rank test

Test error

- Based on the test error on the different folds alone it is not possible to decide that model 1 (red) is significantly bettwer than model 2 (blue)

Difference in test error

- If we look at the difference in performance, it is clear that model 2 (red) performs better

Empirical Tuning of Hyperparameters

Hyperparameter

- In addition to the normal parameters, often one or several hyperparameters need to be tuned as well. Example: regularization weight λ

- The tuning should be done on the training fold. Part of the training fold becomes another fold on which the hyperparameters are tuned

Hyperparameters(cont'd)

- Let's call the folds parameter training fold, hyperparameter fold, and test fold

- In the outer loop we generate training data and test data (as part of K-fold cross validation)

- In the inner loop we divide the training data into parameter training fold and hyperparameter fold. We train the parameters using the parameter training fold with different values of the hyperparameters. We then select the hyperparameter values which give best performance on the hyper-parameter fold

- We use these hyperparameter values to optimize the model on all training data, and evaluate this model on the test set

Figure 7.1: *Behavior of test sample and training sample error as the model complexity is varied.*

Learning Theories

Overview: Statistical Theories and Learning Theories

VC-Theory (Statistical Learning theory)
- No assumption on function
- Worst-case analysis
- *Vapnik–Chervonenkis*

PAC Learning (probably approximate correct)
- Similar to VC-Theory
- Also considers computational complexity
- *Valient*

Regularization theory
- Regularization: -> increases stability of solution; ill-poses problems become well-posed
- *Hadamard, Tikhonov*

Probability
- Example: Bester linear Estimator
- Not really statistics but uses simple quantities (correlations) that can be estimated from data

(Subjective) Bayesian Statistics:
- Subjective knowledge can be formulated as probabilities and can be integrated into statistical modeling

Robust Statistics
- Non-Gaussian likelihoods
- *Huber*

Stein estimation
- Biased estimators can beat ML
- *Stein* estimator

Empirical Risc Minimization
- Vapnik

Least squares principle
- Gauss
- Gauss Likelihood

MDL – Theorie
  (minim. description length)
- Information Theory
- Rissanen, Wallace, Boulton

Function Approximation Theory

Empirical Bayes (technicality)
- Type II likelihood
- Evidence Framework

Objective Bayesian Statistic
- Non-informative Priors (Jeffrey)
- Maximum Entropy Priors

Frequentist Statistics
- Rejection of prior probability
- Dominant in the last century
- *Fisher*
  - *p-value*
- *Pearson und Neyman*
  - *Conf .interv., hypoth. test*

Algorithmic Statistics
- Focus on predictions (not parameter estimation)
- *Breiman, Huber, Friedman*

- **Green: Frequent.**
- **Blue: Bayes**
- **Gold:**
     Learn. Theory
- **Red: Related**

Learning Theories

- A: Classical Frequentist Approaches

  - C_p Statistics

  - Akaikes Information Criterion (AIC)

- B: Bayesian approaches

  - Strict Bayes: model averaging instead of model selection

  - Bayesian model selection and Bayesian Information Criterion (BIC)

- C: Modern Frequentist Approaches

  - Minimum Description Length (MDL) Principle

  - Statistical Learning Theory (Vapnik-Chervonenkis (VC) Theory)

A:Classical Frequentist Approaches

Frequentist Approaches

- We are again interested in the generalization cost averaged over the parameter estimates from different training data sets of size N

$$\mathbb{E}_{\mathsf{train}} \mathsf{cost}_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\mathsf{train}), \mathcal{M}]$$

- Thus we evaluate the quality of a particular model $\mathcal{M}$ and not a *particular* parameter vector

Bias-Variance Decomposition

- We assume a fixed $P(x)$ and than $y = f(x) + \epsilon$, where ϵ is uncorrelated uniform noise

- We use a quadratic cost function. Then one can decompose for the *squared cost*

$$\mathbb{E}_{\text{train}}\text{cost}^q_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] = \text{Bias}^2 + \text{Var} + \text{Residual}$$

Residual

- The residual cost is simply the cost of the true model

$$\text{Residual} = \int (f_{true}(\mathbf{x}) - y)^2 P(\mathbf{x}, y) d\mathbf{x} dy = \text{cost}_{P(\mathbf{x},y)}^q [f_{true}]$$

- In regression, this is simply the noise variance σ_ϵ^2

Bias

- The bias is the mean square of the difference between the true model and the average prediction of all models trained with different training sets of size N. A regularized model with $\lambda > 0$ would typically be biased. A linear model is biased if the true dependency is quadratic. With $m(\mathbf{x}) = \mathbb{E}_{\text{train}}(f(\mathbf{x}, \widehat{\mathbf{w}}(\text{train})))$

$$\text{Bias}^2 = \int \left[m(\mathbf{x}) - f_{true}(\mathbf{x}) \right]^2 P(\mathbf{x}) d\mathbf{x}$$

Variance

- The variance is the mean square of the difference between trained models and the average prediction of all models trained with different training sets of size N

$$\mathsf{Var} = \int \mathbb{E}_{\text{train}}[f(\mathbf{x}, \widehat{\mathbf{w}}(\text{train})) - m(\mathbf{x})]^2 P(\mathbf{x}) d\mathbf{x}$$

Background: Some Rules for Variances and Traces

- Let $\mathbf{y}$ be a random vector with covariance $\mathrm{Cov}(\mathbf{y})$ and let A be a fixed matrix

  If $\mathbf{z} = A\mathbf{y}$, then: $\mathrm{Cov}(\mathbf{z}) = A\,\mathrm{Cov}(\mathbf{y})A^T$

- The trace is the sum over the diagonal elements of a matrix. One can show that

$$\mathbf{trace}[\Phi(\Phi^T\Phi)^{-1}\Phi^T] = M$$

  where M is the number of columns of the matrix Φ. Special case: when Φ is a square matrix and has an inverse, then $\Phi(\Phi^T\Phi)^{-1}\Phi^T = I$ and the trace of I is obviously M

Example: Linear Models

- We assume that the data has been generated with

$$y_i = \phi(\mathbf{x}_i)\mathbf{w} + \epsilon_i$$

  where: ϵ_i is independent noise with variance σ^2

- We take the ML estimator which is known to be unbiased and is

$$\widehat{\mathbf{w}} = (\Phi^T\Phi)^{-1}\Phi^T\mathbf{y}$$

  Thus we now know that $Bias = 0$.

- With the rule we just learned we can calculate the parameter covariance

$$\mathsf{Cov}(\widehat{\mathbf{w}}) = (\Phi^T\Phi)^{-1}\Phi^T\mathsf{Cov}(\mathbf{y})\Phi(\Phi^T\Phi)^{-1}$$

$$= \sigma^2(\Phi^T\Phi)^{-1}\Phi^T\Phi(\Phi^T\Phi)^{-1} = \sigma^2(\Phi^T\Phi)^{-1}$$

- Great, now we know how certain the parameters are. We can now evaluate Var by taking a large sample of $P(\mathbf{x})$. Unfortunately such a large sample is not available and *we simply approximate it with the training data inputs*.

Variance Estimate

- The mean predictions of our model at the training data inputs is then $\mathbf{f} = \Phi\mathbf{w}$

- Applying the covariance formula again as before, we get

$$\mathrm{Cov}(\mathbf{f}) = \Phi\mathrm{Cov}(\mathbf{w})\Phi^T$$

- For the variance we really only need the mean over the diagonal terms

$$\mathrm{Var}(\mathbf{f}) = \frac{1}{N}trace(\Phi\mathrm{Cov}(\mathbf{w})\Phi^T)$$

Example: Linear Models (cont'D)

- Substituting, we get

$$\widehat{\mathsf{Var}} = \frac{1}{N} trace(\Phi \mathsf{Cov}(\mathbf{w}) \Phi^T) = \frac{\sigma^2}{N} trace(\Phi (\Phi^T \Phi)^{-1} \Phi^T)$$

- Now we apply our trace-rule and get

$$\widehat{\mathsf{Var}} = \frac{M_p}{N} \sigma^2$$

  where M_p is the number of parameters

- The solution is surprisingly simple, but makes sense: The predictive variance increases with more noise on the data and with more free parameters and decreases with more data!

Generalization Cost of the Best Fit

- Thus the generalization cost for the parameters that minimize the training set costs is on average

$$\mathbb{E}_{\text{train}}\text{cost}^q_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \approx \sigma^2 + \frac{M_p}{N}\sigma^2 = \sigma^2\frac{M_p + N}{N}$$

- Thus on average the generalization cost for the parameters optimized on the training set is larger by $\frac{M_p}{N}\sigma^2$, if compared to the generalization cost of the best possible model

Training Cost of the Best Fit

- We estimate σ^2 as

$$\widehat{\sigma}^2 = \frac{N}{N - M_p} \text{cost}^q_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}]$$

C_P-statistics

- By substitution we now get

$$\mathbb{E}_{\mathsf{train}}\mathsf{cost}^q_{P(\mathbf{x},y)}[\widehat{\mathbf{w}}(\mathsf{train}), \mathcal{M}] \approx \frac{N + M_p}{N - M_p}\mathsf{cost}^q_{\mathsf{train}}[\widehat{\mathbf{w}}(\mathsf{train}), \mathcal{M}]$$

- This is called Mallot's C_P-statistics

- Thus in model selection on would chose the model where Mallot's C_P is smallest

Expected Average Training Cost

- By taking the expectation of both sides,

$$\sigma^2 = \frac{N}{N - M_p} E_{\text{train}} \text{cost}_{\text{train}}^q [\widehat{\mathbf{w}}(\text{train}), \mathcal{M}]$$

and thus the expected training error is equal to the Residual minus the Variance,

$$E_{\text{train}} \text{cost}_{\text{train}}^q [\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] = \sigma^2 - \frac{M_p}{N} \sigma^2$$

which, as proclaimed, is the Residual minus the Variance

- We can refer this to the earlier discussion and the difference between generalization cost and expected average training cost is 2 times the Variance,

$$\mathbb{E}_{\text{train}} \left\{ \text{cost}_{P(\mathbf{x},y)}^q [\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] - \text{cost}_{\text{train}}^q [\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] \right\} \approx 2 \frac{M_p}{N} \sigma^2$$

Considering Bias

- Discussion: if two models are unbiased, the estimate of the noise level should be more or less the same and the C_P-statistics will choose the smaller one. If the model becomes too simple (we get bias), the average training cost will contain the bias as well and σ^2 will be estimated to be too large

- Thus one can argue that the smallest unbiased model will be optimal under the C_P-statistics

Estimates for Variance and Bias

- Let $\widehat{\sigma}_\infty^2$ be the estimate for a sufficiently large model (∞), which we can assume to be unbiased; from this model, we can obtain an unbiased estimate of the noise variance

- The estimates are then

$$\widehat{\text{Residual}} = \widehat{\sigma}_\infty^2 \qquad \widehat{\text{Var}} = \frac{M_p}{N}\widehat{\sigma}_\infty^2$$

$$\widehat{\text{Bias}}^2 = cost_{train}(\widehat{w}(train, \mathcal{M})) - \widehat{\sigma}_\infty^2 + \frac{M_p}{N}\widehat{\sigma}_\infty^2$$

$$= cost_{train}(\widehat{w}(train, \mathcal{M})) - \frac{N - M_p}{N}\widehat{\sigma}_\infty^2$$

- Note of caution: in the derivations of the formulas, we assumed that the input distribution $P(x)$ is the empirical distribution defined by the training data; in practice, we want to generalize to new x such that, with considerable structural bias, the Variance estimate, and thus also the Bias estimate, become inaccurate

Conceptual Plots

- The next figures show the behavior of bias, variance and residual and average training and average test costs

- The complexity is controlled by the number of parameters M_p, or the number of epochs (stopped training), of the inverse of the regularization parameter

- Note that the best models have a Bias > 0

Costs as a function of model complexity

Costs as a function of training epochs

Costs as a function of inverse regluarization parameter

Gen

Var

Train

Bias^2

Residual

Training epochs

$1/\lambda$

Akaikes Information Criterion (AIC)

- The analysis so far was only valid for models that minimized the squared cost. Consider the cross entropy cost function which minimizes the negative log-likelihood

$$\text{cost}^l_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = -\log P(y|\mathbf{x}, \mathbf{w}, \mathcal{M})$$

- The log likelihood of the ML-solution is

$$\log L = \sum_{i=1}^{N} \log P(y_i|\mathbf{x}_i, \mathbf{w}_{ML})$$

- Here, one can apply Akaike's *Information Criterion (AIC)* (as defined in Wikipedia)

$$AIC = -2\log L + 2M_p$$

- A model with a smaller *AIC* is preferred

Comments on AIC

- The expression

$$\frac{AIC}{2N} = \left(\text{cost}^l_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + \frac{M_p}{N} \right)$$

estimates the generalization log-likelihood cost

C_p and AIC

- The minimizers for *AIC* and C_p are identical for Gaussian noise with known noise variance σ^2; here,

$$\log L = -\frac{N}{2}\log(2\pi\sigma^2) - \frac{1}{2\sigma^2}\sum_{i=1}^{N}(y_i - f(\mathbf{x}_i, \mathbf{w}))^2$$

Here, $const = -\frac{N}{2}\log(2\pi\sigma^2)$ is independent of N and M_p

$$\frac{AIC}{2N} - \frac{1}{2}\log(2\pi\sigma^2) = \frac{1}{N}\frac{1}{2\sigma^2}\sum_{i=1}^{N}(y_i - f(\mathbf{x}_i, \mathbf{w}))^2 + \frac{M_p}{N}$$

$$= \frac{1}{2\sigma^2}\left(\frac{1}{N}\sum_{i=1}^{N}(y_i - f(\mathbf{x}_i, \mathbf{w}))^2 + 2\frac{M_p}{N}\sigma^2\right)$$

$$= \frac{1}{2\sigma^2}\left(\text{cost}_{\text{train}}^{q}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + 2\frac{M_p}{N}\sigma^2\right)$$

AIC for Likelihood Cost Function and for 1/0 Cost Function

Proof: Bias-Variance Decomposition

- One can reduce the problem to estimating the decomposition for one parameter. μ is the parameter and x are the data. We add and subtract $\mathbb{E}_{\text{train}}(\widehat{\mu}))$ and we add and subtract μ (true parameter). Then,

$$\mathbb{E}_{\text{train}}\mathbb{E}_x(\widehat{\mu}-x)^2 = \mathbb{E}_{\text{train}}\mathbb{E}_x \left[(\widehat{\mu} - \mathbb{E}_{\text{train}}(\widehat{\mu})) + (\mathbb{E}_{\text{train}}(\widehat{\mu}) - \mu) + (\mu - x)\right]^2$$

- One gets

$$\mathbb{E}_{\text{train}}\mathbb{E}_x(\widehat{\mu} - x)^2 = \text{Bias}^2 + \text{Var} + \text{Rest}$$

$$\text{Rest} = \mathbb{E}_x(x - \mu)^2$$

$$\text{Bias} = \mathbb{E}_{\text{train}}(\widehat{\mu}) - \mu$$

$$\text{Var} = \mathbb{E}_{\text{train}}[\widehat{\mu} - \mathbb{E}_{\text{train}}(\widehat{\mu})]^2$$

Proof: Bias-Variance Decomposition (cont'd)

$$\mathbb{E}_{\mathsf{train}}\mathbb{E}_x(\widehat{\mu} - x)^2 = \mathbb{E}_{\mathsf{train}}\mathbb{E}_x \left[(\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu})) + (\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu) + (\mu - x) \right]^2$$

- We get 6 terms. Three are: Bias^2, Var, Rest. We need to show that the three cross terms become zero.

$$\mathbb{E}_{\mathsf{train}}\mathbb{E}_x[(\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu}))(\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu)] = \mathbb{E}_{\mathsf{train}}[(\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu}))(\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu)]$$

$$= (\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu)\mathbb{E}_{\mathsf{train}}[\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu})] = \mathsf{Bias} \times 0 = 0$$

$$\mathbb{E}_{\mathsf{train}}\mathbb{E}_x[(\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu}))(\mu - x)] = \mathbb{E}_{\mathsf{train}}[\widehat{\mu} - \mathbb{E}_{\mathsf{train}}(\widehat{\mu})]\mathbb{E}_x[\mu - x] = 0 \times 0 = 0$$

$$\mathbb{E}_{\mathsf{train}}\mathbb{E}_x[(\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu)(\mu - x)] = \mathbb{E}_{\mathsf{train}}[\mathbb{E}_{\mathsf{train}}(\widehat{\mu}) - \mu]\mathbb{E}_x[\mu - x] = \mathsf{Bias} \times 0 = 0$$

Bayesian Approaches

The Bayesian Perspective

- The Bayesian approach does not require model selection!

- One formulates all plausible models under consideration and specifies a prior probability for those models

$$P(\mathcal{M}_i)$$

- The posterior prediction becomes

$$P(y|\mathbf{x}) = \sum_i P(\mathcal{M}_i|D) \int P(y|\mathbf{x}, \mathbf{w}, \mathcal{M}_i) P(\mathbf{w}|D, \mathcal{M}_i) d\mathbf{w}$$

Bayesian Model Selection

- The principled Bayesian approach is sometimes impractical and a model selection is performed

- A posteriori model probability

$$P(\mathcal{M}|D) \propto P(\mathcal{M})P(D|\mathcal{M})$$

- If one assumes that all models have the same prior probability, and the important term is the so-called marginal likelihood, or model evidence

$$P(D|\mathcal{M}) = \int P(D|\mathbf{w}, \mathcal{M})P(\mathbf{w}|\mathcal{M})d\mathbf{w}$$

Our Favorite Linear Model

- Fortunately we can sometimes calculate the evidence without solving complex integrals. From Bayes formula we get

$$P(\mathbf{w}|D, \mathcal{M}) = \frac{P(D|\mathbf{w}, \mathcal{M})P(\mathbf{w}|\mathcal{M})}{P(D|\mathcal{M})}$$

  and thus

$$P(D|\mathcal{M}) = \frac{P(D|\mathbf{w}, \mathcal{M})P(\mathbf{w}|\mathcal{M})}{P(\mathbf{w}|D, \mathcal{M})}$$

- This equation must be true for any $\mathbf{w}$. Let's substitute $\mathbf{w}_{MAP}$ and take the log

$$\log P(D|\mathcal{M}) =$$

$$\log P(D|\mathbf{w}_{MAP}, \mathcal{M}) + \log P(\mathbf{w}_{MAP}|\mathcal{M}) - \log P(\mathbf{w}_{MAP}|D, \mathcal{M})$$

- The first term is the log-likelihood and the second the prior. Both are readily available. So we only need to take care of the last term

Our Favorite Linear Model (cont'd)

- Recall from a previous lecture that for a Bayesian approach to linear regression, we get

$$P(\mathbf{w}|D, \mathcal{M}) = \mathcal{N}\left(\mathbf{w}; \mathbf{w}_{MAP}, cov(\mathbf{w}|D)\right)$$

Here,

$$\text{cov}(\mathbf{w}|D, \mathcal{M}) = \sigma^2 \left(\mathbf{X}^T\mathbf{X} + \frac{\sigma^2}{\alpha^2}I\right)^{-1}$$

Our Favorite Linear Model (cont'd)

- Thus at $\mathbf{w}_{MAP}$ the exponent is zero $(\exp(0) = 1)$ and we are left with the *normalization term*

$$\log P(\mathbf{w}_{MAP}|D, \mathcal{M}) = \log \frac{1}{\sqrt{(2\pi)^{M_p} \det \mathrm{cov}(\mathbf{w}|D, \mathcal{M})}}$$

$$= -\frac{M_p}{2} \log(2\pi) - \frac{1}{2} \log \det \mathrm{cov}(\mathbf{w}|D, \mathcal{M})$$

Our Favorite Linear Model (cont'd)

- Thus,

$$\log P(D|\mathcal{M}) = \log P(D|\mathbf{w}_{MAP}, \mathcal{M}) + \log P(\mathbf{w}_{MAP}, \mathcal{M})$$

$$+\frac{M_p}{2}\log(2\pi) + \frac{1}{2}\log \det \mathrm{cov}(\mathbf{w}|D, \mathcal{M})$$

- For large N, one can approximate

$$\log \det \mathrm{cov}(\mathbf{w}|D) \approx -M_p \log N + constants$$

(for large N, $\mathrm{cov}(\mathbf{w}|D)$ becomes diagonal and the diagonal entries become proportional to $1/N$; thus $\det(\mathrm{cov}(\mathbf{w}|D, \mathcal{M})) \propto (1/N)^{M_p} = N^{-M_p}$ is then the product over the diagonals)

Our Favorite Linear Model (cont'd)

- Thus

$$\log P(D|\mathcal{M}) \approx \log P(D|\mathbf{w}_{MAP}, \mathcal{M}) + \log P(\mathbf{w}_{MAP}, \mathcal{M})$$

$$+ \frac{M_p}{2} \log(2\pi) - \frac{1}{2} M_p \log N + constants$$

- If we consider models with different number of parameters M_p, then $\log P(D|w_{MAP}) +$ $\log P(w_{MAP})$ might produce a larger value (better fit) for the model with the larger M_p. But for the larger model, we subtract a larger $M_p \log N$, so we obtain a compromise between both terms at the optimum

Laplace Approximation of the Marginal Likelihood

- By simplifying the previous equation (we only keep terms that depend on N) and using the ML (Maximum Likelihood) estimate instead of the MAP estimate one obtains

$$\log P(D|\mathcal{M}) \approx \log P(D|\hat{\mathbf{w}}_{ML}, \mathcal{M}) - \frac{M_p}{2} \log N$$

- The *Bayesian information criterion* (BIC) is -2 times this expression (definition in Wikipedia)

$$BIC = -2 \log P(D|\hat{\mathbf{w}}_{ML}, \mathcal{M}) + M_p \log N$$

(a better model has a smaller BIC)

- This approximation is generally applicable (not just for regression)
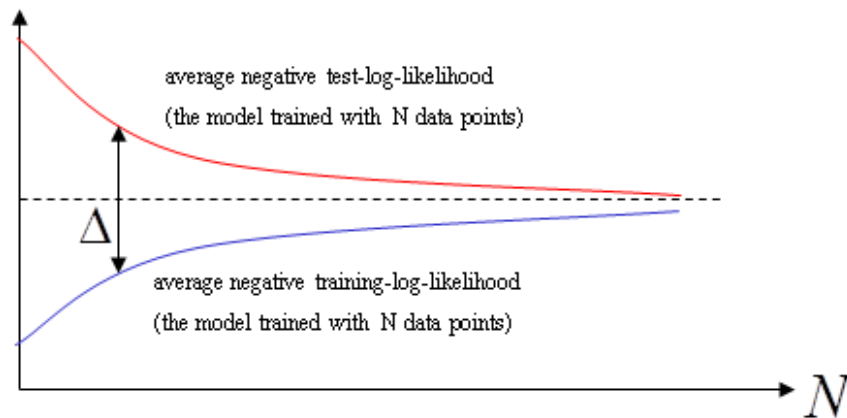
Bayesian Information Criterion (BIC)

- We get

$$\frac{BIC}{2N} = \text{cost}^l_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + \frac{1}{2}\frac{M_p}{N}\log N$$

Compare

$$\frac{AIC}{2N} = \text{cost}^l_{\text{train}}[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + \frac{M_p}{N}$$

- $\frac{1}{2}\frac{M_p}{N}\log N$ is an estimate of the difference between the average test likelihood and the average training log-likelihood

- BIC corection is by a factor $\frac{1}{2}\log N$ larger than the AIC correction and decreases more slowly $(\log N)/N$ with the number of training examples

Comparison: AIC and BIC

average negative test-log-likelihood

(the model trained with N data points)

Δ

average negative training-log-likelihood

(the model trained with N data points)

N

average negative test-log-likelihood

average negative training-log-likelihood

Δ

M

Schätzung von D:

AIC: $\quad \Delta = \dfrac{M}{N}$

BIC: $\quad \Delta = \dfrac{M}{N}\dfrac{1}{2}\log N$

With the number of data points N, Δ decreases, qith increasing complexity M,

Δ increases

C: Modern Frequentist Approaches

Minimum Description Length (MDL)

- Based on the concept of algorithmic complexity (Kolmogorov, Solomonoff, Chaitin)

- Based on these ideas: Rissanen (and Wallace, Boulton) introduced the principal of the minimum description length (MDL)

- Under simplifying assumptions the MDL criterion becomes the BIC criterion

Statistical Learning Theory

- The Statistical Learning Theory (SLT) is in the tradition of the Russian mathematicians Andrey Kolmogorov and Valery Ivanovich Glivenko and the Italian mathematician Francesco Paolo Cantelli

- SLT was founded by Vladimir Vapnik and Alexey Chervonenkis (VC-Theory)

- Part of *Computational Learning Theory* (COLT); similar to PAC (*probably approximately correct*) Learning (Leslie Valiant)

Function Classes

- As before: data is generated according to some distribution $P(\mathbf{x}, y)$. This distribution is fixed but otherwise arbitrary

- As before: SLT considers functions out of a function class $f_{\mathbf{w}}(\mathbf{x}) \in \mathcal{M}$. Example: $\mathcal{M}$ is the class of all linear classifiers with $M_p = M + 1$ parameters (for simplicity, we consider that an element of the function class can be described by a parameter vector $\mathbf{w}$)

- SLT does not assume that the best possible (true or target) function $f_{true}(\mathbf{x})$ is contained in $\mathcal{M}$

- We now consider binary classification without noise (i.e., classes are in principal separable)

SLT Bounds

- As before: train is a random sample of $P(\mathbf{x}, y)$

- SLT considers the difference between $\text{cost}^m_{P(\mathbf{x},y)}[\mathbf{w}, \mathcal{M}]$ and $\text{cost}^m_{\text{train}}[\mathbf{w}, \mathcal{M}]$ (recall that cost^m denotes the misclassification cost); we are interested in the difference between generalization cost and average training cost for any $\mathbf{w}$ (not just for $\widehat{\mathbf{w}}(\text{train})$)

- As before: probabilities are calculated w.r.t. all training sets of size N

SLT Bounds (cont'd)

- SLT shows that with (large) probability $1 - \eta$ that for any $\mathbf{w}$,

$$\text{cost}_{P(\mathbf{x},y)}^m[\mathbf{w}, \mathcal{M}] \leq \text{cost}_{\text{train}}^m[\mathbf{w}, \mathcal{M}] + \epsilon$$

- Model selection is performed on

$$\text{cost}_{\text{train}}^m[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + \epsilon$$

(note; for model selection, we consider the particular parameter choice $\widehat{\mathbf{w}}(\text{train})$)

VC-Dimension

- The statement is trivially true for a large enough ϵ; the science (and the art) is now to find the smallest possible ϵ

- Of great importance is the so-called VC-dimension dim_{VC} of the model class

- Wikipedia: In Vapnik-Chervonenkis theory, the VC dimension (for Vapnik-Chervonenkis dimension) is a measure of the capacity (complexity, expressive power, richness, or flexibility) of a space of functions that can be learned by a statistical classification algorithm. It was originally defined by Vladimir Vapnik and Alexey Chervonenkis

- The VC-dimension can be finite or infinite. For linear classifiers $dim_{VC} = M_p = M + 1$, which means that the VC-dimension is simply the number of parameters

- For systems with a finite VC dimension, the bound decreases with N when $N > dim_{VC}$.

- In practice, $\text{cost}_{\text{train}}^m[\widehat{\mathbf{w}}(\text{train}), \mathcal{M}] + \epsilon$ is much larger than the average test set error, which limits the application of the theory in practice

Typical Bound

- A typical estimate is

$$\epsilon = \sqrt{\frac{1}{N}\left[dim_{VC}\left(\log\left(\frac{2N}{dim_{VC}}\right)+1\right)-\log\left(\frac{\eta}{4}\right)\right]}$$

Comparing SLT and a Frequentist Approach

- SLT does not make any assumption about the true function; in a frequentist view, this can mean that the bias can be large

- SLT makes a statement about the worst case (supremum); often the supremum corresponds to the weight vector that minimizes the training cost, i.e. $\widehat{\mathbf{w}}(\text{train})$

- Note: training data is are not worse case (they are not selected by a demon to fool you): they are assumed to be generated i.i.d. from a fixed $P(y|x)P(x)$

Informal Link to C_p

- In the C_p analysis of systems with fixed basis functions, we use

$$\text{cost}^q_{\text{train}}[\widehat{\mathbf{w}}(\text{train})] + 2\widehat{Var}$$

for model selection, with $\widehat{Var} = \frac{M_p}{N}\widehat{\sigma}^2$

- To get closer to the spirit of SLT, one could use for model selection

$$\text{cost}^q_{\text{train}}[\widehat{\mathbf{w}}(\text{train})] + 2\widehat{Var} + 2\sqrt{\widehat{Var(\widehat{Var})}}$$

So we also consider the uncertainty of the variance estimate

Conclusion

- Machine learning focusses on generalization costs and traditionally not as much on parameter estimation, although explainable AI is gaining interest

- Empirical model selection is most often used. If possible, cross validation results should be reported

- Frequentist approaches typically estimate $\mathbb{E}_{\text{train}}\text{cost}_{P(\mathbf{x},y)}[\hat{\mathbf{w}}(\text{train}), \mathcal{M}]$. We have studied C_P and the AIC

- Bayesian approaches do model averaging instead of model selection. The BIC criterion is useful, if model selection needs to be performed

- An advantage of the SLT is that the true function does not need to be included in the function class; the derived bounds are typically often rather conservative

- SLT has been developed in the Machine Learning community, whereas frequentist and Bayesian approaches originated in Statistics

Appendices

P-norm Cost Functions

- A p-norm is

$$\|x\|_p := \left(\sum_{i=1}^{N} |x_i|^p \right)^{1/p}$$

- Let $\mathbf{y}$ be targets on the training data and $\widehat{\mathbf{y}}$ the vector of predictions of the classifier

- Then the quadratic cost can be written as $\|\mathbf{y} - \widehat{\mathbf{y}}\|_2^2$, the absolute value cost as $\|\mathbf{y} - \widehat{\mathbf{y}}\|_1$

- $\|\mathbf{y} - \widehat{\mathbf{y}}\|_\infty$ is the infinity norm which is equal to the maximum absolute value of its components

- ℓ^p-norms generalize the p-norm to sequence spaces

- L^p-norms generalize the p-norm to function spaces

Evaluation

- The cost functions defined before can also be used to evaluate the performance on the test data; but measures like the cross-entropy on test data are difficult to interpret

- Note that the evaluation of the test set performance reduces to the evaluation of the scatter plot of y and $\widehat{y} = 0.5(1 + \text{sign}(f(\mathbf{x}, \mathbf{w}, \mathcal{M}))$

- For regression: mean-squared test set cost, explained variance, and Pearson correlation are popular choices

- For classification all is contained in the 2×2 table of $y \in \{0, 1\}$ and $\widehat{y} \in \{0, 1\}$ on the test data

- For balanced classes one typically reports mutual information and accuracy

- For unbalanced classes one focusses on measures derived from the co-occurrence distribution

Measures for Performance Evaluation*

- True Positive (a.k.a. hit)($y \in \{0, 1\}$):

$$\mathsf{TP}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = y\widehat{y}$$
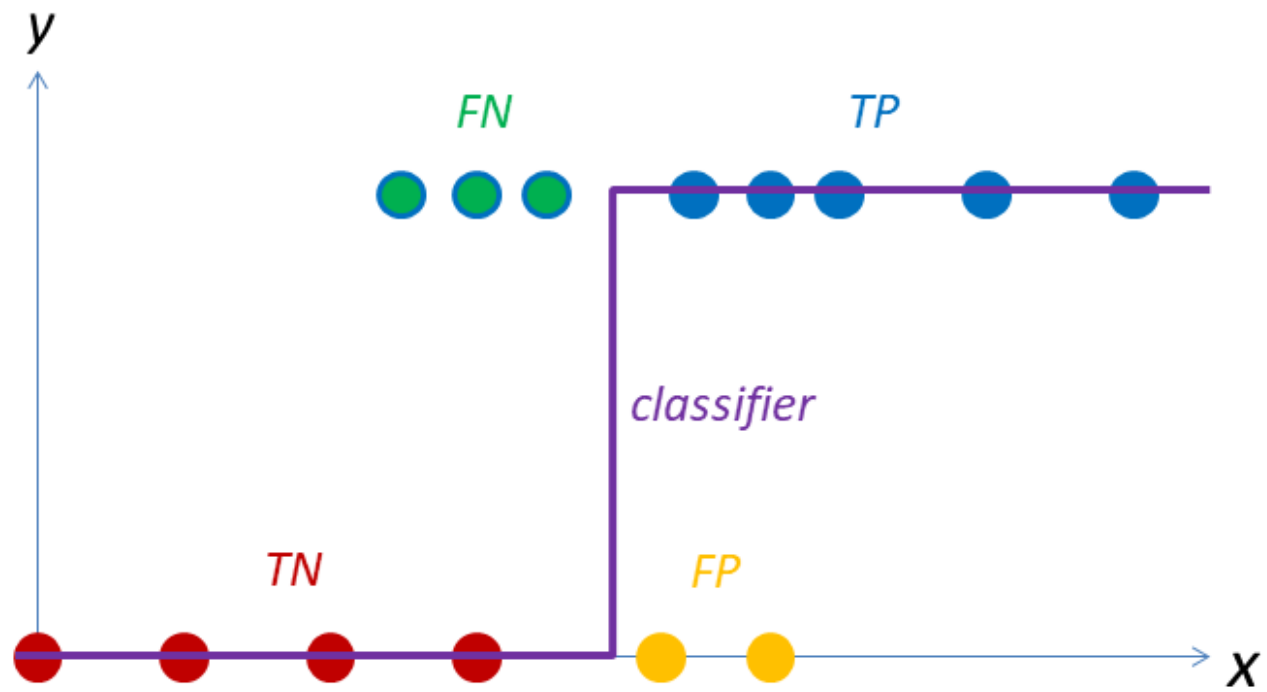
- True Negative ($y \in \{0, 1\}$):

$$\mathsf{TN}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = (1 - y)(1 - \widehat{y})$$

- False Positive (a.k.a. false alarm, type I error) ($y \in \{0, 1\}$):

$$\mathsf{FP}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = (1 - y)\widehat{y}$$

- False Negative (a.k.a. miss, type II error) ($y \in \{0, 1\}$):

$$\mathsf{FN}_{\mathbf{x},y}[\mathbf{w}, \mathcal{M}] = y(1 - \widehat{y})$$

Precision, Recall, and Specificity

Based on these, one defines

- Precision (positive predicted value):

$$\frac{\mathsf{TP}_{\text{test}}}{\mathsf{TP}_{\text{test}} + \mathsf{FP}_{\text{test}}} \approx P(Y = 1 | \hat{Y}(X) = 1)$$

- Recall (sensitivity, true positive rate, hit rate, or detection rate):

$$\frac{\mathsf{TP}_{\text{test}}}{\mathsf{TP}_{\text{test}} + \mathsf{FN}_{\text{test}}} \approx P(\hat{Y}(X) = 1 | Y = 1)$$

- Specificity (True Negative Rate):

$$\frac{\mathsf{TN}_{\text{test}}}{\mathsf{FP}_{\text{test}} + \mathsf{TN}_{\text{test}}} \approx P(\hat{Y}(X) = 0 | Y = 0)$$

F1 and Accuracy

- F1-score

$$\mathsf{F1} = 2\frac{\mathrm{Precision} \times \mathrm{Recall}}{\mathrm{Precision} + \mathrm{Recall}}$$

- Accuracy

$$\frac{\mathsf{TP}_{\mathrm{test}} + \mathsf{TN}_{\mathrm{test}}}{\mathsf{TP}_{\mathrm{test}} + \mathsf{TN}_{\mathrm{test}} + \mathsf{FP}_{\mathrm{test}} + \mathsf{FN}_{\mathrm{test}}}$$

Interpretation: Inverse Model

- Consider a fire detector where X is the degree of smoke associated with a fire $Y = 1$; the fire is the cause, the smoke is the effect, $Y \to X$

- The fire detector $\hat{Y}(X)$ tries to model the inverse: predict the probability of a fire Y given the degree of smoke X

- Recall (Sensitivity) is the probability that there is an alarm when there is a fire: $P(\hat{Y}(X) = 1 | Y = 1)$

- Specificity is the probability that there is no alarm when there is no fire: $P(\hat{Y}(X) = 0 | Y = 0)$

- **Recall and Specificity** condition on Y and are thus independent on the prevalence of fire in a neighborhood. They measure the **performance of a fire detector** and are the basis for the ROC curve

Interpretation: Inverse Model (cont'd)

- Precision is the probability that there is a fire when there is an alarm:
  $P(Y = 1 | \hat{Y}(X) = 1)$

- Usefulness in an application: **Recall and Precision** is what you typically care about in an application: you don't want to have false alarms (you want to have a high precision) and you don't want to miss a fire (you want to have a high recall)

- But Precision is not independent of the prevalence of fire in an environment!

Interpretation: Forward Model

- Y is the interest of a user (the effect), X is the web page (the cause) and $\hat{Y}(X)$ is the estimated interest, $X \rightarrow Y$

- The model tries to model the forward model: predict the probability of a user interest given X

- Recall (Sensitivity) is the probability that we label each relevant web page as being interesting (each relevant page should be shown to you):
$P(\hat{Y} = 1(X)|Y = 1)$

- Precision is the probability that a web page is of interest when we label it to be interesting (if a page is shown to you, it should be relevant):
$P(Y = 1|\hat{Y}(X) = 1)$

- Again, in the application one is interested in a high Recall and a high Precision