

CS840a
Fall 2006
Learning and Computer Vision
Prof. Olga Veksler

Lecture 3

SVM

Information Theory (a little BIT)
Some pictures from C. Burges

SVM

- Said to start in 1979 with Vladimir Vapnik's paper
- Major developments throughout 1990's
- Elegant theory
 - Has good generalization properties
- Have been applied to diverse problems very successfully in the last 10-15 years
- One of the most important developments in pattern recognition in the last 10 years



Today


- Support Vector Machines
- Mutual Information
- Preparation for the next time:
 - papers: "Object Recognition with Informative Features and Linear Classification" by M. Naquet and S. Ullman
 - Ignore section of tree-augmented network
 - "Face Recognition with Support Vector Machines: Global vs. Component-based Approach"

Linear Discriminant Functions

- A discriminant function is linear if it can be written as

$$g(x) = w^T x + w_0$$


$$\begin{aligned} g(x) > 0 &\Rightarrow x \in \text{class 1} \\ g(x) < 0 &\Rightarrow x \in \text{class 2} \end{aligned}$$



- which separating hyperplane should we choose?

Linear Discriminant Functions


- Training data is just a subset of all possible data
- Suppose hyperplane is close to sample x_i
- If we see new sample close to sample i , it is likely to be on the wrong side of the hyperplane



- Poor generalization (performance on unseen data)

SVM


- Idea: maximize distance to the closest example



- For the optimal hyperplane
 - distance to the closest negative example = distance to the closest positive example

Linear Discriminant Functions


- Hyperplane as far as possible from any sample



- New samples close to the old samples will be classified correctly
- Good generalization


SVM: Linearly Separable Case

- SVM: maximize the **margin**



- margin** is twice the absolute value of distance b of the closest example to the separating hyperplane
- Better generalization (performance on test data)
 - in practice
 - and in theory

SVM: Linearly Separable Case



- **Support vectors** are the samples closest to the separating hyperplane
 - they are the most difficult patterns to classify
 - Optimal hyperplane is completely defined by support vectors
 - of course, we do not know which samples are support vectors without finding the optimal hyperplane


SVM: Formula for the Margin

- For uniqueness, set $|w^T x_i + w_0| = 1$ for any example x_i closest to the boundary
- now distance from closest sample x_i to $g(x) = 0$ is

$$\frac{|w^T x_i + w_0|}{\|w\|} = \frac{1}{\|w\|}$$

- Thus the margin is


$$m = \frac{2}{\|w\|}$$



SVM: Formula for the Margin

- $g(x) = w^T x + w_0$
- absolute distance between x and the boundary $g(x) = 0$

$$\frac{|w^T x + w_0|}{\|w\|}$$



- distance is unchanged for hyperplane $g_1(x) = \alpha g(x)$

$$\frac{|\alpha w^T x + \alpha w_0|}{\|\alpha w\|} = \frac{|w^T x + w_0|}{\|w\|}$$

- Let x_i be an example closest to the boundary. Set $|w^T x_i + w_0| = 1$
- Now the largest margin hyperplane is unique

SVM: Optimal Hyperplane

- Maximize margin $m = \frac{2}{\|w\|}$
- subject to constraints

$$\begin{cases} w^T x_i + w_0 \geq 1 & \text{if } x_i \text{ is positive example} \\ w^T x_i + w_0 \leq -1 & \text{if } x_i \text{ is negative example} \end{cases}$$

- Let $\begin{cases} z_i = 1 & \text{if } x_i \text{ is positive example} \\ z_i = -1 & \text{if } x_i \text{ is negative example} \end{cases}$

- Can convert our problem to

$$\begin{aligned} &\text{minimize } J(w) = \frac{1}{2} \|w\|^2 \\ &\text{constrained to } z_i (w^T x_i + w_0) \geq 1 \quad \forall i \end{aligned}$$

- $J(w)$ is a quadratic function, thus there is a single global minimum

SVM: Optimal Hyperplane

- Use Kuhn-Tucker theorem to convert our problem to:

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $\alpha = \{\alpha_1, \dots, \alpha_n\}$ are new variables, one for each sample
- Can rewrite $L_D(\alpha)$ using n by n matrix H :

$$L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}^t H \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}$$

- where the value in the i th row and j th column of H is $H_{ij} = z_i z_j x_i^t x_j$

SVM: Optimal Hyperplane

- After finding the optimal $\alpha = \{\alpha_1, \dots, \alpha_n\}$
 - For every sample i , one of the following must hold
 - $\alpha_i = 0$ (sample i is not a support vector)
 - $\alpha_i \neq 0$ and $z_i(w^t x_i + w_0 - 1) = 0$ (sample i is support vector)
 - can find w using $w = \sum_{i=1}^n \alpha_i z_i x_i$
 - can solve for w_0 using any $\alpha_i > 0$ and $\alpha_i [z_i (w^t x_i + w_0) - 1] = 0$

$$w_0 = \frac{1}{z_i} - w^t x_i$$

- Final discriminant function:

$$g(x) = \left(\sum_{i \in S} \alpha_i z_i x_i \right)^t x + w_0$$

- where S is the set of support vectors

$$S = \{x_i \mid \alpha_i \neq 0\}$$

SVM: Optimal Hyperplane

- Use Kuhn-Tucker theorem to convert our problem to:

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $\alpha = \{\alpha_1, \dots, \alpha_n\}$ are new variables, one for each sample
- $L_D(\alpha)$ can be optimized by quadratic programming
- $L_D(\alpha)$ formulated in terms of α
 - it depends on w and w_0 indirectly


SVM: Optimal Hyperplane

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $L_D(\alpha)$ depends on the number of samples, not on dimension of samples
- samples appear only through the dot products $x_i^t x_j$
- This will become important when looking for a **nonlinear** discriminant function, as we will see soon
- Code available on the web to optimize

SVM: Non Separable Case

- Data is most likely to be not linearly separable, but linear classifier may still be appropriate



- Can apply SVM in non linearly separable case
 - data should be "almost" linearly separable for good performance

SVM: Non Separable Case

- Would like to minimize

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n I(\xi_i > 0)$$

of samples not in ideal location

- where $I(\xi_i > 0) = \begin{cases} 1 & \text{if } \xi_i > 0 \\ 0 & \text{if } \xi_i \leq 0 \end{cases}$
- constrained to $z_i(w^T x_i + w_0) \geq 1 - \xi_i$ and $\xi_i \geq 0 \quad \forall i$
- β is a constant which measures relative weight of the first and second terms
 - if β is small, we allow a lot of samples not in ideal position
 - if β is large, we want to have very few samples not in ideal position


SVM: Non Separable Case

- Use non-negative slack variables ξ_1, \dots, ξ_n (one for each sample)

- Change constraints from $z_i(w^T x_i + w_0) \geq 1 \quad \forall i$ to $z_i(w^T x_i + w_0) \geq 1 - \xi_i \quad \forall i$

- ξ_i is a measure of deviation from the ideal for sample i


- $\xi_i > 1$ sample i is on the wrong side of the separating hyperplane
- $0 < \xi_i < 1$ sample i is on the right side of separating hyperplane but within the region of maximum margin



SVM: Non Separable Case

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n I(\xi_i > 0)$$

of examples not in ideal location



large β , few samples not in ideal position

small β , a lot of samples not in ideal position

SVM: Non Separable Case

- Unfortunately this minimization problem is NP-hard due to discontinuity of functions $I(\xi_i)$

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n I(\xi_i > 0)$$


of examples
not in ideal location

- where $I(\xi_i > 0) = \begin{cases} 1 & \text{if } \xi_i > 0 \\ 0 & \text{if } \xi_i \leq 0 \end{cases}$
- constrained to $z_i(w'x_i + w_0) \geq 1 - \xi_i$ and $\xi_i \geq 0 \quad \forall i$


Non Linear Mapping

- Cover's theorem:
 - "pattern-classification problem cast in a high dimensional space non-linearly is more likely to be linearly separable than in a low-dimensional space"

- One dimensional space, not linearly separable



- Lift to two dimensional space with $\phi(x) = (x, x^2)$



SVM: Non Separable Case

- Instead we minimize

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n \xi_i$$

a measure of
of misclassified
examples

- constrained to $\begin{cases} z_i(w'x_i + w_0) \geq 1 - \xi_i & \forall i \\ \xi_i \geq 0 & \forall i \end{cases}$


- Can use Kuhn-Tucker theorem to converted to

$$\begin{aligned} &\text{maximize} && L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^T x_j \\ &\text{constrained to} && 0 \leq \alpha_i \leq \beta \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- find w using $w = \sum_{i=1}^n \alpha_i z_i x_i$
- solve for w_0 using any $0 < \alpha_i < \beta$ and $\alpha_i [z_i (w'x_i + w_0) - 1] = 0$

Non Linear Mapping

- To solve a non linear classification problem with a linear classifier
 - Project data x to high dimension using function $\phi(x)$
 - Find a linear discriminant function for transformed data $\phi(x)$
 - Final nonlinear discriminant function is $g(x) = w^T \phi(x) + w_0$




- In 2D, discriminant function is linear

$$g\left(\begin{bmatrix} x^{(1)} \\ x^{(2)} \end{bmatrix}\right) = [w_1 \quad w_2] \begin{bmatrix} x^{(1)} \\ x^{(2)} \end{bmatrix} + w_0$$

- In 1D, discriminant function is not linear $g(x) = w_1 x + w_2 x^2 + w_0$

Non Linear Mapping: Another Example



Non Linear SVM: Kernels

- Recall SVM optimization

$$\text{maximize } L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j$$
- Note this optimization depends on samples x_i only through the dot product $x_i^t x_j$
- If we lift x_i to high dimension using $\phi(x)$, need to compute high dimensional product $\phi(x_i)^t \phi(x_j)$

$$\text{maximize } L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j \phi(x_i)^t \phi(x_j)$$

$K(x_i, x_j)$
- Idea: find **kernel** function $K(x_i, x_j)$ s.t.

$$K(x_i, x_j) = \phi(x_i)^t \phi(x_j)$$

Non Linear SVM

- Can use any linear classifier after lifting data into a higher dimensional space. However we will have to deal with the “curse of dimensionality”
 - poor generalization to test data
 - computationally expensive
- SVM avoids the “curse of dimensionality” problems by
 - enforcing largest margin permits good generalization
 - It can be shown that generalization in SVM is a function of the margin, independent of the dimensionality
 - computation in the higher dimensional case is performed only implicitly through the use of **kernel** functions

Non Linear SVM: Kernels

- $$\text{maximize } L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j \phi(x_i)^t \phi(x_j)$$

$K(x_i, x_j)$
- Then we only need to compute $K(x_i, x_j)$ instead of $\phi(x_i)^t \phi(x_j)$
 - “kernel trick”: do not need to perform operations in high dimensional space explicitly

Non Linear SVM: Kernels


- Suppose we have 2 features and $K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^t \mathbf{y})^2$
- Which mapping $\phi(\mathbf{x})$ does it correspond to?

$$K(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^t \mathbf{y})^2 = \left(\begin{bmatrix} \mathbf{x}^{(1)} & \mathbf{x}^{(2)} \end{bmatrix} \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \end{bmatrix} \right)^2 = (\mathbf{x}^{(1)} \mathbf{y}^{(1)} + \mathbf{x}^{(2)} \mathbf{y}^{(2)})^2$$

$$= (\mathbf{x}^{(1)} \mathbf{y}^{(1)})^2 + 2(\mathbf{x}^{(1)} \mathbf{y}^{(1)})(\mathbf{x}^{(2)} \mathbf{y}^{(2)}) + (\mathbf{x}^{(2)} \mathbf{y}^{(2)})^2$$

$$= \begin{bmatrix} (\mathbf{x}^{(1)})^2 & \sqrt{2} \mathbf{x}^{(1)} \mathbf{x}^{(2)} & (\mathbf{x}^{(2)})^2 \end{bmatrix} \begin{bmatrix} (\mathbf{y}^{(1)})^2 & \sqrt{2} \mathbf{y}^{(1)} \mathbf{y}^{(2)} & (\mathbf{y}^{(2)})^2 \end{bmatrix}^t$$
- Thus

$$\phi(\mathbf{x}) = \begin{bmatrix} (\mathbf{x}^{(1)})^2 & \sqrt{2} \mathbf{x}^{(1)} \mathbf{x}^{(2)} & (\mathbf{x}^{(2)})^2 \end{bmatrix}$$



Non Linear SVM

- search for separating hyperplane in high dimension

$$\mathbf{w} \phi(\mathbf{x}) + w_0 = 0$$
- Choose $\phi(\mathbf{x})$ so that the first ("0"th) dimension is the augmented dimension with feature value fixed to 1

$$\phi(\mathbf{x}) = \begin{bmatrix} 1 & \mathbf{x}^{(1)} & \mathbf{x}^{(2)} & \mathbf{x}^{(1)} \mathbf{x}^{(2)} \end{bmatrix}^t$$
- Threshold parameter w_0 gets folded into the weight vector \mathbf{w}

$$\begin{bmatrix} w_0 & \mathbf{w} \end{bmatrix} \begin{bmatrix} 1 \\ \phi(\mathbf{x}) \end{bmatrix} = 0$$

Non Linear SVM: Kernels

- How to choose kernel function $K(\mathbf{x}_i, \mathbf{x}_j)$?
 - $K(\mathbf{x}_i, \mathbf{x}_j)$ should correspond to product $\phi(\mathbf{x}_i)^t \phi(\mathbf{x}_j)$ in a higher dimensional space
 - Mercer's condition tells us which kernel function can be expressed as dot product of two vectors
 - Kernel's not satisfying Mercer's condition can be sometimes used, but no geometrical interpretation
- Some common choices (satisfying Mercer's condition):
 - Polynomial kernel $K(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i^t \mathbf{x}_j + 1)^p$
 - Gaussian radial Basis kernel (data is lifted in infinite dimension)

$$K(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{x}_j\|^2\right)$$

Non Linear SVM

- Will not use notation $\mathbf{a} = [w_0 \ \mathbf{w}]$, we'll use old notation \mathbf{w} and seek hyperplane through the origin

$$\mathbf{w} \phi(\mathbf{x}) = 0$$
- If the first component of $\phi(\mathbf{x})$ is not 1, the above is equivalent to saying that the hyperplane has to go through the origin in high dimension
 - removes only one degree of freedom
 - But we have introduced many new degrees when we lifted the data in high dimension

Non Linear SVM Receptie

- Start with data $\mathbf{x}_1, \dots, \mathbf{x}_n$ which lives in feature space of dimension d
- Choose kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ or function $\phi(\mathbf{x}_i)$ which takes sample \mathbf{x}_i to a higher dimensional space
- Find the largest margin linear discriminant function in the higher dimensional space by using quadratic programming package to solve:

$$\text{maximize } L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j K(\mathbf{x}_i, \mathbf{x}_j)$$

constrained to $0 \leq \alpha_i \leq \beta \quad \forall i$ and $\sum_{i=1}^n \alpha_i z_i = 0$

Non Linear SVM

- Nonlinear discriminant function

$$g(\mathbf{x}) = \sum_{\mathbf{x}_i \in S} \alpha_i z_i K(\mathbf{x}_i, \mathbf{x})$$

$$g(\mathbf{x}) = \sum_{\substack{\text{weight of support} \\ \text{vector } \mathbf{x}_i}} \alpha_i z_i \underbrace{K(\mathbf{x}_i, \mathbf{x})}_{\substack{\text{"inverse distance"} \\ \text{from } \mathbf{x} \text{ to} \\ \text{support vector } \mathbf{x}_i}}$$

most important training samples, i.e. support vectors

$$K(\mathbf{x}_i, \mathbf{x}) = \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{x}\|^2\right)$$

Non Linear SVM Recipe

- Weight vector \mathbf{w} in the high dimensional space:

$$\mathbf{w} = \sum_{\mathbf{x}_i \in S} \alpha_i z_i \phi(\mathbf{x}_i)$$
 - where S is the set of support vectors $S = \{\mathbf{x}_i \mid \alpha_i \neq 0\}$
- Linear discriminant function of largest margin in the high dimensional space:

$$g(\phi(\mathbf{x})) = \mathbf{w}' \phi(\mathbf{x}) = \left(\sum_{\mathbf{x}_i \in S} \alpha_i z_i \phi(\mathbf{x}_i) \right)' \phi(\mathbf{x})$$
- Non linear discriminant function in the original space:

$$g(\mathbf{x}) = \left(\sum_{\mathbf{x}_i \in S} \alpha_i z_i \phi(\mathbf{x}_i) \right)' \phi(\mathbf{x}) = \sum_{\mathbf{x}_i \in S} \alpha_i z_i \phi'(\mathbf{x}_i) \phi(\mathbf{x}) = \sum_{\mathbf{x}_i \in S} \alpha_i z_i K(\mathbf{x}_i, \mathbf{x})$$
- decide class 1 if $g(\mathbf{x}) > 0$, otherwise decide class 2


SVM Example: XOR Problem

- Class 1: $\mathbf{x}_1 = [1, -1]$, $\mathbf{x}_2 = [-1, 1]$
- Class 2: $\mathbf{x}_3 = [1, 1]$, $\mathbf{x}_4 = [-1, -1]$
- Use polynomial kernel of degree 2:
 - $K(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i' \mathbf{x}_j + 1)^2$
 - This kernel corresponds to mapping

$$\phi(\mathbf{x}) = \left[1 \quad \sqrt{2}x^{(1)} \quad \sqrt{2}x^{(2)} \quad \sqrt{2}x^{(1)}x^{(2)} \quad (x^{(1)})^2 \quad (x^{(2)})^2 \right]$$
- Need to maximize

$$L_D(\alpha) = \sum_{i=1}^4 \alpha_i - \frac{1}{2} \sum_{i=1}^4 \sum_{j=1}^4 \alpha_i \alpha_j z_i z_j (\mathbf{x}_i' \mathbf{x}_j + 1)^2$$

constrained to $0 \leq \alpha_i \quad \forall i$ and $\alpha_1 + \alpha_2 - \alpha_3 - \alpha_4 = 0$




SVM Example: XOR Problem

- Can rewrite $L_D(\alpha) = \sum_{i=1}^4 \alpha_i - \frac{1}{2} \alpha' H \alpha$
 - where $\alpha = [\alpha_1 \ \alpha_2 \ \alpha_3 \ \alpha_4]'$ and $H = \begin{bmatrix} 9 & 1 & -1 & -1 \\ 1 & 9 & -1 & -1 \\ -1 & -1 & 9 & 1 \\ -1 & -1 & 1 & 9 \end{bmatrix}$
- Take derivative with respect to α and set it to 0

$$\frac{d}{d\alpha} L_D(\alpha) = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} 9 & 1 & -1 & -1 \\ 1 & 9 & -1 & -1 \\ -1 & -1 & 9 & 1 \\ -1 & -1 & 1 & 9 \end{bmatrix} \alpha = 0$$
- Solution to the above is $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0.25$
 - satisfies the constraints $\forall i, 0 \leq \alpha_i$ and $\alpha_1 + \alpha_2 - \alpha_3 - \alpha_4 = 0$
 - all samples are support vectors

SVM Example: XOR Problem




SVM Example: XOR Problem

- $\phi(x) = [1 \ \sqrt{2}x^{(1)} \ \sqrt{2}x^{(2)} \ \sqrt{2}x^{(1)}x^{(2)} \ (x^{(1)})^2 \ (x^{(2)})^2]'$
- Weight vector w is:

$$w = \sum_{i=1}^4 \alpha_i z_i \phi(x_i) = 0.25(\phi(x_1) + \phi(x_2) - \phi(x_3) - \phi(x_4)) = [0 \ 0 \ 0 \ -\sqrt{2} \ 0 \ 0]$$
 - Thus the nonlinear discriminant function is:

$$g(x) = w\phi(x) = \sum_{i=1}^6 w_i \phi_i(x) = -\sqrt{2}(\sqrt{2}x^{(1)}x^{(2)}) = -2x^{(1)}x^{(2)}$$

Degree 3 Polynomial Kernel



- In linearly separable case (on the left), decision boundary is roughly linear, indicating that dimensionality is controlled
- Nonseparable case (on the right) is handled by a polynomial of degree 3

SVM Summary

- Advantages:
 - Based on nice theory
 - excellent generalization properties
 - objective function has no local minima
 - can be used to find non linear discriminant functions
 - Complexity of the classifier is characterized by the number of support vectors rather than the dimensionality of the transformed space
- Disadvantages:
 - tends to be slower than other methods
 - quadratic programming is computationally expensive
 - Not clear how to choose the Kernel

Information theory

- Suppose we toss a **fair** die with 8 sides
 - need 3 bits to transmit the results of each toss
 - 1000 throws will need 3000 bits to transmit
- Suppose the die is biased
 - side A occurs with probability 1/2, chances of throwing B are 1/4, C are 1/8, D are 1/16, E are 1/32, F 1/64, G and H are 1/128
 - Encode A= 0, B = 10, C = 110, D = 1110, ..., so on until G = 1111110, H = 1111111
 - We need, on average, $1/2+2/4+3/8+4/16+5/32+6/64+7/128+7/128 = 1.984$ bits to encode results of a toss
 - 1000 throws require 1984 bits to transmit
 - Less bits to send = less "information"
 - Biased die tosses contain less "information" than unbiased die tosses (know in advance biased sequence will have a lot of A's)
 - What's the number of bits in the best encoding?
- Extreme case: if a die always shows side A, a sequence of 1,000 tosses has no information, 0 bits to encode

Information theory

- Information Theory regards information as only those symbols that are uncertain to the receiver
 - only information essential to understand must be transmitted**
- Shannon made clear that uncertainty is the very commodity of communication
- The amount of information, or uncertainty, output by an information source is a measure of its entropy
- In turn, a source's entropy determines the amount of bits per symbol required to encode the source's information
- Messages are encoded with strings of 0 and 1 (bits)

Information theory

- if a die is fair (any side is equally likely, or uniform distribution), for any toss we need $\log(8) = 3$ bits
- Suppose any of n events is equally likely (uniform distribution)
 - $P(x) = 1/n$, therefore $-\log P = -\log(1/n) = \log n$
- In the "good" encoding strategy for our biased die example, every side x has $-\log p(x)$ bits in its code
- Expected number of bits is

$$-\sum_x p(x) \log p(x)$$

Shannon's Entropy

$$H[p(x)] = -\sum_x p(x) \log p(x) = \sum_x p(x) \log \frac{1}{p(x)}$$

- How much randomness (or uncertainty) is there in the value of signal x if it has distribution $p(x)$
 - For uniform distribution (every event is equally likely), $H[x]$ is maximum
 - If $p(x) = 1$ for some event x , then $H[x] = 0$
 - Systems with one very common event have less entropy than systems with many equally probable events
- Gives the expected length of optimal encoding (in binary bits) of a message following distribution $p(x)$
 - doesn't actually give this optimal encoding

Mutual Information of X and Y

$$I[x, y] = H(x) - H(x | y)$$

- Measures the average reduction in uncertainty about x after y is known
- or, equivalently, it **measures the amount of information that y conveys about x**
- Properties
 - $I(x, y) = I(y, x)$
 - $I(x, y) \geq 0$
 - If x and y are independent, then $I(x, y) = 0$
 - $I(x, x) = H(x)$

Conditional Entropy of X given Y

$$H[x | y] = \sum_{x,y} p(x, y) \log \frac{1}{p(x | y)} = -\sum_{x,y} p(x, y) \log p(x | y)$$

- Measures average uncertainty about x when y is known
- Property:
 - $H[x] \geq H[x|y]$, which means after seeing new data (y), the uncertainty about x is not increased, on average

MI for Feature Selection

$$I[x, c] = H(c) - H(c | x)$$

- Let x be a proposed feature and c be the class
- If $I[x, c]$ is high, we can expect feature x be good at predicting class c