

Support Vector Machines: Applications

64-360 Algorithmic Learning, part 3b

Norman Hendrich

University of Hamburg
MIN Faculty, Dept. of Informatics
Vogt-Kölln-Str. 30, D-22527 Hamburg
hendrich@informatik.uni-hamburg.de

22/06/2011



Outline

Overview (Review)

Example SVM applications

- SVM for clustering

- SVM for multi-class classification

- SVM for visual pattern recognition

- SVM for text classification

- SVM for function approximation

Summary

References



Support Vector Machines

- ▶ a.k.a. maximum margin classifiers
- ▶ a family of related
 - ▶ supervised
 - ▶ learning methods
 - ▶ for classification and regression
- ▶ try to minimize the classification error
- ▶ while maximizing the geometric margin



Support Vector Machine

- ▶ based on the linear classifier

Four new main concepts:

- ▶ maximum margin classification
- ▶ soft-margin classification for noisy data
- ▶ introduce non-linearity via feature maps
- ▶ Kernel trick: implicit calculation of feature maps

- ▶ use Quadratic Programming for training
- ▶ polynomial or gaussian kernels often work well




Overall concept and architecture

- ▶ select a feature space \mathcal{H} and a mapping function $\Phi : x \mapsto \Phi(x)$
- ▶ select a classification (output) function σ

$$y(x) = \sigma(\sum_i \vartheta_i \langle \Phi(x), \Phi(x_i) \rangle)$$

- ▶ during training, find the *support-vectors* $x_1 \dots x_n$
- ▶ and weights ϑ which minimize the classification error
- ▶ map test input x to $\Phi(x)$
- ▶ calculate dot-products $\langle \Phi(x) \Phi(x_i) \rangle$
- ▶ feed linear combination of the dot-products into σ
- ▶ get the classification result

Maximum margin and support vectors



- ▶ the (linear) classifier with the largest margin
- ▶ data points that limit the margin are called the *support vectors*

Soft-margin classification

- ▶ allow some patterns to violate the margin constraints
- ▶ find a compromise between large margins
- ▶ and the number of violations

Idea:

- ▶ introduce slack-variables $\xi = (\xi_1 \dots \xi_n)$, $\xi_i \geq 0$
- ▶ which measure the margin violation (or classification error) on pattern x_i : $y(x_i)(w \cdot \Phi(x_i) + b) \geq 1 - \xi_i$
- ▶ introduce one global parameter C which controls the compromise between large margins and the number of violations

Soft-margin classification

- ▶ introduce slack-variables ξ_i
- ▶ and global control parameter C

$$\max_{w,b,\xi} \mathcal{P}(w, b, \xi) = \frac{1}{2} w^2 + C \sum_{i=1}^n \xi_i$$

subject to:

$$\forall i : y(x_i)(w \cdot \Phi(x_i) + b) \geq 1 - \xi_i$$

$$\forall i : \xi_i \geq 0$$

- ▶ problem is now very similar to the hard-margin case
- ▶ again, the dual representation is often easier to solve



Nonlinearity through feature maps

General idea:

- ▶ introduce a function Φ which maps the input data into a higher dimensional *feature space*

$$\Phi : x \in X \mapsto \Phi(x) \in \mathcal{H}$$

- ▶ similar to hidden layers of multi-layer ANNs
- ▶ explicit mappings can be expensive in terms of CPU and/or memory (especially in high dimensions)
- ▶ “*Kernel functions*” achieve this mapping *implicitly*
- ▶ often, very good performance



Common SVM feature maps

kernels

- ▶ $z_k = (\text{polynomial terms of } x_k \text{ of degree 1 to } q)$
- ▶ $z_k = (\text{radial basis functions of } x_k)$
- ▶ $z_k = (\text{sigmoid functions of } x_k)$
- ▶ ...
- ▶ combinations of the above, e.g.
- ▶ $K(x, z) = K_1(x, z) + K_2(x, z);$
- ▶ $K(x, z) = K_1(x, z) \cdot K_2(x, z);$

Note:

- ▶ feature map Φ only used in inner products
- ▶ for training, information on pairwise inner products is sufficient

Quadratic polynomial map: scalar product

- ▶ Calculating $\langle \Phi(x), \Phi(y) \rangle$ is $O(m^2)$
- ▶ For comparison, calculate $(x \cdot y + 1)^2$:
- ▶ $(x \cdot y + 1)^2 = ((\sum_{i=1}^m x_i \cdot y_i) + 1)^2$


$$= (\sum_{i=1}^m x_i y_i)^2 + 2(\sum_{i=1}^m x_i y_i) + 1$$

$$= \sum_{i=1}^m \sum_{j=1}^m x_i y_i x_j y_j + 2 \sum_{i=1}^m x_i y_i + 1$$

$$= \sum_{i=1}^m (x_i y_i)^2 + 2 \sum_{i=1}^m \sum_{j=1}^m x_i y_i x_j y_j + 2 \sum_{i=1}^m x_i y_i + 1$$

$$= \Phi(x) \cdot \Phi(y)$$
- ▶ We can replace $\langle \Phi(x), \Phi(y) \rangle$ with $(x \cdot y + 1)^2$, which is $O(m)$

Gaussian kernels



- ▶ SVM with Gaussian kernel can classify every input function
- ▶ if the Gaussian kernels are “narrow” enough
- ▶ Gaussian-kernel SVM has infinite VC dimension

Applications of SVM

- ▶ data clustering
- ▶ multi-class classification
- ▶ visual pattern (object) recognition
- ▶ text classification: string kernels
- ▶ DNA sequence classification
- ▶ function approximation
- ▶ ...

- ▶ “streamlined” kernels for each domain
- ▶ let’s take a look at a few examples...

SVC - support vector clustering

Ben-Hur, Horn, Siegelmann, Vapnik (2001)

- ▶ map data points to high-dimensional feature space
- ▶ using the Gaussian kernel
- ▶ look for the smallest sphere that encloses the data

- ▶ map back to data space
- ▶ to get the set of contours which enclose the cluster(s)

- ▶ identifies valleys in the data probability distribution
- ▶ use soft-margin SVM to handle outliers

SVC: Example data set and results

Gaussian kernel $K(x, z) = e^{-q(x-z)^2}$, radius q




Figure 1: Clustering of a data set containing 183 points using SVC with $C = 1$. Support vectors are designated by small circles, and cluster assignments are represented by different grey scales of the data points. (a): $q = 1$ (b): $q = 20$ (c): $q = 24$ (d): $q = 48$.

SVC: Example data set, number of support vectors

Gaussian kernel $K(x, z) = e^{-q(x-z)^2}$, radius q




Figure 2: Number of SVs as a function of q for the data of Figure 1. Contour splitting points are denoted by vertical lines.

SVC: Noisy data

- ▶ Use soft-margin SVM learning algorithm
- ▶ with control parameter C
- ▶ and slack variables $\xi_j \geq 0$

- ▶ non-support vectors: inside the cluster
- ▶ support vectors: on the cluster boundary

- ▶ **bounded support vectors**: outside the boundary (violation)
- ▶ number of bounded support vectors is $n_{\text{bsv}} < 1/C$
- ▶ fraction of outliers: $p = 1/NC$

SVC: Noisy data and bounded support vectors

Soft-margin SVM learning with control parameter C




Figure 3: Clustering with and without BSVs. The inner cluster is composed of 50 points generated from a Gaussian distribution. The two concentric rings contain 150/300 points, generated from a uniform angular distribution and radial Gaussian distribution. (a) The rings cannot be distinguished when $C = 1$. Shown here is $q = 3.5$, the lowest q value that leads to separation of the inner cluster. (b) Outliers allow easy clustering. The parameters are $p = 0.3$ and $q = 1.0$.

SVC: Noisy data and bounded support vectors

Soft-margin SVM learning with control parameter C




Figure 4: Clusters with overlapping density functions require the introduction of BSVs.

- ▶ Remember: larger C implies larger margin
- ▶ at the cost of more *bounded support vectors*
- ▶ with $y_i \langle w_j, x_i \rangle \geq 1 - \xi_j$

SVC: Strongly overlapping clusters





Figure 5: In the case of significant overlap between clusters the algorithm identifies clusters according to dense cores, or maxima of the underlying probability distribution.

Multi-class classification

- ▶ many practical classification problems involve more than just two classes!
- ▶ for example, clustering (see above), object recognition, handwritten digit and character recognition, audio and natural speech recognition, etc.
- ▶ but standard SVM handles only exactly two classes
- ▶ “hard-coded” in the SVM algorithm
- ▶ What to do?

Example: multi-class classification



- ▶ exmulticlassall from SVM-KM toolbox (3 classes)
- ▶ demo

One versus the rest classification

To get an M -class classifier:

- ▶ construct a set of M binary classifiers f^1, \dots, f^M
- ▶ each trained to separate one class from the rest
- ▶ combine them according to the maximal individual output before applying the `sgn`-function

$$\arg \max_{j=1, \dots, M} g^j(x), \quad \text{where} \quad g^j(x) = \sum_{i=1}^m \left(y_i \alpha_i^j k(x, x_i) + b^j \right)$$

- ▶ the *winner-takes-all* approach

One versus the rest: winner takes all

- ▶ the above algorithm looks for $\arg \max g^j(x)$
- ▶ the M different classifiers have been trained
 - ▶ on the same training data
 - ▶ but with different binary classification problems
- ▶ unclear whether the $g^j(x)$ are on comparable scales
 - ▶ a problem, when several (or none) classifiers claim the pattern
 - ▶ try to balance/scale the $g^j(x)$
- ▶ all classifiers trained on very unsymmetrical problems
 - ▶ many more negative than positive patterns
 - ▶ (e.g. digit-7 vs. all handwritten characters and digits)

One versus the rest: reject decisions

- ▶ The values of $g^j(x)$ can be used for reject decisions in the classification of x
- ▶ consider the difference between the two largest $g^j(x)$ as a measure of confidence
- ▶ if the measure falls short of a threshold θ , the classifier rejects the pattern
- ▶ can often lower the error-rate on other patterns
- ▶ can forward un-classified patterns to human experts

Pairwise classification

- ▶ train a classifier for each possible pair of classes
- ▶ for M classes, requires $M(M - 1)/2$ binary classifiers
 - ▶ digit-0-vs-digit-1, digit-0-vs-digit-2, . . . , digit-8-vs-digit-9
- ▶ (many) more classifiers than one-vs-the-rest for $M > 3$
- ▶ and probably, longer training times
- ▶ but each individual pairwise classifier is (usually) much simpler than each one-vs-the-rest classifier



Pairwise classification: tradeoff

- ▶ requires $(M - 1)M/2$ classifiers vs M one-vs-the-rest
- ▶ each individual classifier much simpler
- ▶ smaller training sets (e.g. digit-7 vs. digit-8)
 - ▶ for super-linear learning complexity like $O(n^3)$, the shorter training times can outweigh the higher number of classifiers
- ▶ usually, fewer support vectors
 - ▶ training sets are smaller
 - ▶ classes have less overlap



Pairwise classification: tradeoff

- ▶ requires $(M - 1)M/2$ classifiers vs M one-vs-the-rest
- ▶ but fewer support vectors per classifier

- ▶ if M is large, will be slower than M one-vs-the-rest
- ▶ example: digit-recognition task and the following scenario:
 - ▶ after evaluating the first few classifiers,
 - ▶ digit 7 and digit 8 seem unlikely (“lost” in the first rounds)
 - ▶ rather pointless to run the digit-7-vs-digit-8 classifier

- ▶ embed the pairwise classifiers into a directed acyclic graph
 - ▶ each classification run corresponds to a graph traversal
 - ▶ much faster than running all pairwise classifiers

Error-correcting output coding

- ▶ train a number of L binary classifiers f^1, \dots, f^L
- ▶ on subproblems involving subsets of the M classes
 - ▶ e.g. separate digits 0..4 from 5..9, 0..2 from 3..9, etc.
- ▶ If the set of binary classifiers is chosen correctly,
- ▶ their responses $\{\pm 1\}^L$ determine the output class of a test pattern
- ▶ e.g., $\log_2(M)$ classifiers on binary-encoding ...
- ▶ use error-correcting codes to improve robustness against individual mis-classifications
 - ▶ note: newest schemes also use the margins of the individual classifiers for decoding

Error-correcting output coding

- ▶ example: Hamming (7,4) code
 - ▶ linear error-correcting block-code
 - ▶ 4 databits, 3 parity bits
 - ▶ detects and corrects 1-bit errors
 - ▶ generator matrix G and decoding matrix H

- ▶ more efficient codes
 - ▶ BCD (Bose, Chaudhuri, Hocquenghem)
e.g. BCH(15,7,5) corrects 2-bit errors
 - ▶ RS (Reed-Solomon)
 - ▶ ...
 - ▶ large block-sizes required for low overhead

$$G' := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 \end{pmatrix}$$

$$H' := \begin{pmatrix} 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 \end{pmatrix}$$

Multi-class objective functions

LWK p.213

- ▶ re-design the SVM algorithm to directly handle multi-classes
- ▶ $y_i \in \{1, \dots, M\}$ the multi-class label of pattern x_i
- ▶ $m \in \{1, \dots, M\}$

$$\text{minimize}_{w_r \in \mathcal{H}, \xi_i^r \in \mathbb{R}^m, b_r \in \mathbb{R}} \frac{1}{2} \sum_{r=1}^M \|w_r\|^2 + \frac{C}{m} \sum_{i=1}^m \sum_{r \neq y_i} \xi_i^r$$

subject to $\langle w_{y_i}, x_i \rangle + b_{y_i} \geq \langle w_r, x_i \rangle + b_r + 2 - \xi_i^r$, with $\xi_i^r \geq 0$.

- ▶ optimization problem has to deal with all SVMs at once
- ▶ large number of support vectors
- ▶ results comparable with one-vs-the-rest approach

Summary: multi-class classification

LWK p.214

- ▶ Basic SVM algorithm only supports binary classifications
- ▶ Several options for M -class classification
 - 1 M one-versus-the-rest classifiers
 - 2 $M(M - 1)/2$ pairwise binary classifiers
 - 3 suitably chosen subset classifiers (at least, $\log_2 M$), plus error-correcting codes for robustness
 - 4 redesigned SVM with multi-class objective function
- ▶ no approach outperforms all others
- ▶ often, one-vs-the-rest produces acceptable results



SVM for visual pattern recognition

- ▶ one very popular application of SVMs
- ▶ can work on raw pixels
- ▶ or handcrafted feature maps


- ▶ MNIST handwritten digit recognition
- ▶ NORB object recognition
- ▶ histogram based classification

MNIST handwritten characters data set

- ▶ set of handwritten digits
- ▶ based on NIST database 1 and -3 (b&w)
- ▶ 28x28 pixel grayscale images (interpolated from 28x28 b&w)
- ▶ used as a benchmark for (multi-class) classifiers
- ▶ training set with 60.000 patterns
- ▶ test set with 10.000+ patterns

- ▶ <http://yann.lecun.com/exdb/mnist/>
- ▶ current best classifier achieves 0.38% error

MNIST data set: example '4'



http://www.cvl.isy.liu.se/ImageDB/images/external_images/MNIST_digits/mnist_train4.jpg

MNIST benchmark

- ▶ typical published results on raw MNIST
 - ▶ 0% percent errors on training set
 - ▶ about 3% errors on the test set

- ▶ apparently, training and test set don't match perfectly
- ▶ some test patterns quite different from training patterns
- ▶ difficult to achieve very good error rates

- ▶ several approaches based on extra training patterns

MNIST results overview

Table 14.1 Test error rates of various learning models on the MNIST dataset. Many results obtained with deslanted images or hand-designed feature extractors were left out. NN: fully connected neural network; ConvNN: convolutional neural network; k -NN: k -nearest neighbors.

Classifier	Error	Reference
Knowledge-free methods		
2-layer NN, 800 hidden units	1.60%	Simard et al. 2003
3-layer NN, 500+300 units	1.53%	Hinton et al. 2006
SVM, Gaussian kernel	1.40%	DeCoste and Schölkopf 2002
EVM, Gaussian kernel	1.03%	Haffner 2002
DBM + final backpropagation	0.95%	Hinton et al. 2006
Convolutional networks		
ConvNN LeNet-5	0.80%	Ranzato et al. 2007
ConvNN LeNet-6	0.70%	Ranzato et al. 2007
ConvNN LeNet-6 + unsupervised learning	0.60%	Ranzato et al. 2007
Training set augmented with <i>affine distortions</i>		
2-layer NN, 800 hidden units	1.10%	Simard et al. 2003
Virtual SVM, degree 9 polynomial kernel	0.80%	DeCoste and Schölkopf 2002
ConvNN,	0.60%	Simard et al. 2003
Training set augmented with <i>elastic distortions</i>		
2-layer NN, 800 hidden units	0.70%	Simard et al. 2003
SVM Gaussian Kernel + online training	0.67%	this book, chapter 13
Shape context features + elastic k -NN	0.63%	Belongie et al. 2002
ConvNN	0.40%	Simard et al. 2003
ConvNN LeNet-6	0.49%	Ranzato et al. 2007
ConvNN LeNet-6 + unsupervised learning	0.39%	Ranzato et al. 2007



MNIST experiments with SVMs

- ▶ G. Loosli, L. Bottou, S. Canu, *Training Invariant SVMs Using Selective Sampling*, in *Large-Scale Kernel Machines*, 2007
- ▶ an approach to improve the classification error rate
- ▶ by increasing the training set
- ▶ with automatically synthesized patterns
- ▶ derived from the original training patterns

- ▶ 100 random deformations of each original image
- ▶ 6 million training images. . .

Virtual training set

- ▶ synthesize new training patterns
- ▶ by applying deformations on each original pattern
 - ▶ affine transformations (sub-pixel accuracy)
translations, rotations, scaling
 - ▶ deformation-fields (elastic transformation)
 - ▶ thickening
 - ▶ ...

Goal:

- ▶ a *transformation invariant* classifier
- ▶ more robust to slight variations in the test patterns
- ▶ but handling transformations can also increase the test set error
- ▶ of course, much higher training effort and time

MNIST modified training patterns




Figure 13.4 This figure shows 16 variations of a digit with all the transformations cited here.

MNIST effect of transformations

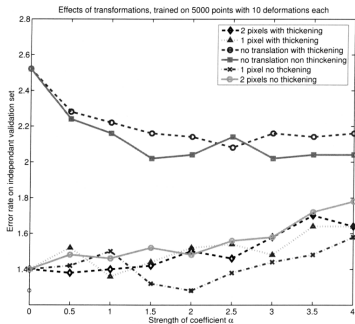



Figure 13.5 Effects of transformations on the performance. This graph is obtained on a validation set of 10,000 points, trained on 5000 points and 10 transformations for each (55,000 training points in total) with an RBF kernel with bandwidth $\gamma = 0.006$. The best configuration is elastic deformation without thickening, for $\eta = 2$ and one-pixel translations, which gives 1.28%. Note that $\eta = 0$ is equivalent to no elastic deformation. The baseline result for the validation set is thus 2.52%.



LASVM training algorithm

- ▶ due to the problem size, training is done using an iterative algorithm, one pattern a time.
- ▶ several choices to select the next training pattern:
 - ▶ random selection: picks a random unseen training example
 - ▶ gradient selection: pick the most poorly classified example (smallest value of $y_k f(x_k)$ among 50 randomly selected unseen training examples)
 - ▶ active selection: pick the training example that is closest to the decision boundary (smallest value of $|f(x_k)|$) among 50 randomly selected unseen training examples)
 - ▶ autoactive selection: randomly sample at most 100 unseen training examples, but stop as soon as 5 fall inside the margins (will become support vectors). Pick the one closest to the decision boundary.

MNIST benchmark results



MNIST benchmark results explanation

Figure 13.6 This figure compares the error rates (left), the numbers of support vectors (center), and the training times (right) of different LASVM runs. The first bar of each graph corresponds to training a regular SVM on the original 60,000 MNIST examples (NT: no transformation). The other three bars were obtained using 100 random deformations of each MNIST example, that is, 6 millions points. The second columns reports results for random selection (RS), the third for active selection (AS), and the last for auto-active selection (AAS). The deformation settings are set according to previous results (figure 13.5). The autoactive run gives the best compromise.

MNIST evolution of training and test errors




Figure 13.8 This figure shows the evolution of the training error on the 60,000 original points and the test error during training. The results are shown for class 2 against the others. We observe that training on deformed examples does not affect the performance on original examples but increases the performance on unseen data.

MNIST SVM experiment results

G. Loosli, L. Bottou, S. Canu, Training Invariant SVMs Using Selective Sampling

Table 13.2 Summary of our final experiment.

Number of binary classifiers	10
Number of examples for each binary classifier	8,100,000
Thickening transformation	no
Additional translations	1 pixel
RBF kernel bandwidth (γ)	0.006
Example selection criterion	auto-active
Finishing step	every 600,000 examples
Full training time	8 days
Test set error	0.67%

NORB

- ▶ Natural images of 3D-objects, 96x96 pixels
- ▶ 50 different toys in 5 categories
- ▶ 25 objects for training, 25 for testing
- ▶ each object captured in stereo from 162 viewpoints (9 elevations, 18 azimuths)

- ▶ objects in front of uniform background
- ▶ or in front of cluttered background (or missing object)

- ▶ <http://www.cs.nyu.edu/~ylclab/data/norb-v1.0/>

NORB normalized-uniform training set




Figure 14.6 The 25 testing objects in the *normalized-uniform* NORB set. The testing objects are unseen by the trained system.

NORB jittered-cluttered training set



Figure 14.7 Some of the 291,600 examples from the *jittered-cluttered* training set (left camera images). Each column shows images from one category. A 6-th background category is added

NORB data set

- ▶ recognition basically can only rely on the shape of the object
- ▶ all other typical clues eliminated or unusable

- ▶ different orientations (viewing angles)
- ▶ different lighting conditions
- ▶ no color information (grayscale only)
- ▶ no object texture
- ▶ different backgrounds (cluttered set)

- ▶ no hidden regularities

Error rates: normalized uniform set

Table 14.2 Testing error rates and training/testing timings on the *normalized-uniform* dataset of different methods. The timing is normalized to hypothetical 1GHz single CPU. The convolutional nets have multiple results with different training passes due to their iterative training procedure.

	SVM	Convolutional network		Hybrid	
Test error	11.6%	10.4%	6.0%	6.2%	5.9%
Train time (min×GHz)	480	64	448	3200	50+
Test time per sample (min×GHz)	0.95	0.03		0.04+	
Fraction of SV	28%				28%
Parameters	$\sigma=2000$	Step size		dim=80	
	$C=40$	$2 \times 10^{-5} - 2 \times 10^{-7}$		$\sigma=5$	$C=40$

- ▶ five binary SVMs, one for each class
- ▶ trained on the raw pixel images ($d = 96 \cdot 96 \cdot 2 = 18432$)
- ▶ convolutional network uses handcrafted feature map
- ▶ hybrid system trains SVMs on those features

Feature functions for the NORB set




Figure 14.5 The architecture of the convolutional net used for the NORB experiments. The input is an image pair, the system extracts 8 feature maps of size 92×92 , 8 maps of 23×23 , 24 maps of 18×18 , 24 maps of 6×6 , and a 100-dimensional feature vector. The feature vector is then transformed into a 5-dimensional vector in the last layer to compute the distance with target vectors.

Some feature functions for the NORB set




Figure 14.8 The learned convolution kernels of the C3 layer. The columns correspond to the 24 feature maps output by C3, and the rows correspond to the 8 feature maps output by the S2 layer. Each feature map draws from 2 monocular maps and 2 binocular maps of S2. Ninety-six convolution kernels are used in total.

Error rates: jittered cluttered set

Table 14.3 Testing error rates and training/testing timings on the *jittered-cluttered* dataset of different methods. The timing is normalized to hypothetical 1GHz single CPU. The convolutional nets have multiple results with different training passes due to their iterative training procedure.

	SVM	Convolutional network			Hybrid
Test error	43.3%	16.38%	7.5%	7.2%	5.9%
Train time (min×GHz)	10,944	420	2100	5880	330+
Test time per sample (min×GHz)	2.2	0.04			0.06+
Fraction of SV	5%				2%
Parameters	$\sigma = 10^4$ C=40	Step size $2 \times 10^{-5} - 1 \times 10^{-6}$		dim=100 $\sigma=5$ C=1	

- ▶ again, SVM trained on raw pixels
- ▶ convolutional network uses handcrafted feature maps
- ▶ hybrid system trains SVM on those feature maps



Error rates: SVM training and setup

- ▶ basic SVM fails with 43% error-rate. . .
- ▶ six one-vs-the-rest binary SVMs (one per category)
- ▶ training samples are raw 108x108 pixel images
 - ▶ again, use a virtual training set
 - ▶ ± 3 pixel translations
 - ▶ scaling from 80% to 110%
 - ▶ rotations $\pm 5^\circ$
 - ▶ changed brightness ± 20 and contrast
 - ▶ a total of 291.600 images
- ▶ overall, a 23 328-dimensional input vector
- ▶ only the Gaussian width σ as a free parameter



Poor performance of SVM on raw pixels?

- ▶ Gaussian kernel basically computes matching score (based on Euclidean distance) between training templates and the test pattern
- ▶ very sensitive to variations in registration, pose, illumination
- ▶ most of the pixels in NORB are background clutter
- ▶ hence, template matching dominated by background irregularities

- ▶ a general weakness of standard kernel methods: their inability to select relevant input features
- ▶ feature maps must be hand-crafted by experts



Summary: visual pattern recognition

- ▶ SVM can be trained on and applied to raw pixel data
- ▶ use *virtual training set* for better generalization
- ▶ but no performance guarantees

- ▶ good results on MNIST
- ▶ but total breakdown on NORB

- ▶ must use appropriate feature maps
- ▶ or hybrid architectures

Text classification

LSKM p0.41

- ▶ another high dimensional problem. . .
- ▶ e.g. Reuters RCV1 text corpus
 - ▶ 810.000 news stories from 1996/1997
 - ▶ partitioned and indexed in 135 categories
 - ▶ <http://trec.nist.gov/data/reuters/reuters.html>
- ▶ represent word frequencies, e.g. *bag of words*
- ▶ or represent substring correlations
- ▶ train a SVM on the corpus
- ▶ classify the given input texts

Bag of words representation

LWK 13.2

Sparse vector kernel

- ▶ map the given text into a sparse vector
- ▶ each component corresponds to a word
- ▶ a component is set to one when the word occurs
- ▶ dot products between such vectors are fast

- ▶ but ignores the ordering of the words
- ▶ no vicinity information (e.g. words in one sentence)
- ▶ only detects exact matches (e.g. mismatch on mathces)

String kernel example

efficient kernel that computes the dot product

- ▶ in the feature space spanned by *all* substrings of documents
- ▶ computational complexity is linear in the document length and the length of the substring
- ▶ allows to classify texts on similarity of substrings
- ▶ the more substrings match, the higher the output value
- ▶ use for text classification, DNA sequence analysis, ...

String kernel: definitions

- ▶ a finite alphabet Σ , so Σ^n the set of all strings of length n
- ▶ $\Sigma^* = \bigcup_{n=0}^{\infty} \Sigma^n$ the set of all finite strings
- ▶ length of a string $s \in \Sigma^*$ is $|s|$
- ▶ string elements are $s(1) \dots s(|s|)$
- ▶ st is the string concatenation of s and t

subsequences u of strings:

- ▶ index sequence $i := (i_1, \dots, i_{|u|})$ with $1 \leq i_1 < \dots < i_{|u|} \leq |s|$
- ▶ define $u := s(i) := s(i_1) \dots s(i_{|u|})$
- ▶ $l(i) := i_{|u|} - i_1 + 1$ the length of the subsequence in s

String kernel: feature space

- ▶ feature space $\mathcal{H} := \mathbb{R}^{(\Sigma^n)}$ built from strings of length n
- ▶ one dimension (coordinate) for each element of Σ^n

- ▶ feature map

$$[\Phi_n(s)]_u := \sum_{i:s(i)=u} \lambda^{l(i)}$$

- ▶ with *decay parameter* λ , $0 < \lambda < 1$
- ▶ the larger the length of the subsequence in s , the smaller its contribution to $[\Phi_n(s)]_u$.
- ▶ sum over all subsequences of s which equal u

String kernel: the actual kernel

- ▶ consider the dimension of \mathcal{H} for the string asd
- ▶ $[\Phi_n(\text{Nasdaq})]_{\text{asd}} = \lambda^3$ (one exact match of length 3)
- ▶ $[\Phi_n(\text{lass das})]_{\text{asd}} = 2\lambda^5$ (two matches of length 5: $\text{las}\square\square\square\text{d}\square\square$ and $\square\text{a}\square\text{s}\square\text{d}\square\square$)
- ▶ kernel corresponding to the map $\Phi(n)$ is:

$$k_n(s, t) = \sum_{u \in \Sigma^n} [\Phi_n(s)]_u [\Phi_n(t)]_u = \sum_{u \in \Sigma^n} \sum_{(i,j): s(i)=t(j)=u} \lambda^{l(i)} \lambda^{l(j)}$$

- ▶ normalize: use $k(s, t) / \sqrt{k(s, s)k(t, t)}$

DNA sequence classification

LWK table 13.2 p.417

- ▶ DNA sequence contains *coding sequences* which encode proteins, but also untranslated sequences
- ▶ find the start points of proteins (TIS: translation initiation sites)
- ▶ out of $\{A, C, G, T\}$, typically an ATG triplet
- ▶ certain local correlations are typical
- ▶ $\text{match}_{p+j}(x, x')$ is 1 for matching nucleotides at position $p + j$, 0 otherwise
- ▶ construct kernel that rewards nearby matches

$$\text{win}_p(x, x') = \left(\sum_{j=-l}^{+l} \nu_j \text{match}_{p+j}(x, x') \right)^{d_1}$$

$$k(x, x') = \left(\sum_{p=1}^l \text{win}_p(x, x') \right)^{d_2}$$

WD kernel with shifts



Figure 4.2 Given two sequences x_1 and x_2 of equal length, the WD kernel with shifts consists of a weighted sum to which each match in the sequences makes a contribution $\gamma_{k,p}$ depending on its length k and relative position p , where long matches at the same position contribute most significantly. The γ 's can be computed from the β 's and δ 's in (4.4). The spectrum kernel is based on a similar idea, but it only considers substrings of a fixed length and the contributions are independent of the relative positions of the matches to each other.

SVM-based regression

- ▶ use SVMs for function approximation?
- ▶ especially, for high-dimensional functions?

basic idea is very similar to classification:

- ▶ estimate linear functions $f(x) = \langle w, x \rangle + b$
- ▶ based on $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{H} \times \mathbb{R}$
- ▶ use a $\|w\|^2$ regularizer (“maximum margin”)
- ▶ use optimization algorithm similar to SVM training
- ▶ use feature-maps to generalize to the non-linear case

ϵ -insensitive loss function

Vapnik 1995


- ▶ need a suitable cost-function
- ▶ define the following ϵ -insensitive loss function:

$$|y - f(x)|_\epsilon = \max\{0, |y - f(x)| - \epsilon\}$$

- ▶ threshold $\epsilon \geq 0$ is chosen a-priori
- ▶ small ϵ implies high approximation accuracy
- ▶ no penalty, when error below some threshold
- ▶ similar to classification loss-function: no penalty for correctly-classified training patterns

The ϵ -tube

around the ϵ -insensitive loss function



- ▶ geometrical interpretation: allow a tube
- ▶ of width ϵ around the given function values

Goal of the SVR learning

Given:

- ▶ a dot product space \mathcal{H}
- ▶ (mapped) input patterns $(x_1, y_1), \dots, (x_m, y_m) \in \mathcal{H} \times \mathbb{R}$

Goal:

- ▶ find a function f with a small *risk* (or test error),

$$R[f] = \int c(f, x, y) dP(x, y)$$
- ▶ where P is the probability measure for the observations
- ▶ and c is a loss function, e.g. $c(f, x, y) = (f(x) - y)^2$
- ▶ loss function can be chosen depending on the application



Regularized risk functional

- ▶ find a function f with a small *risk* (or test error),
 $R[f] = \int c(f, x, y) dP(x, y)$
- ▶ where P is the probability measure for the observations
- ▶ and c is a loss function, e.g. $c(f, x, y) = (f(x) - y)^2$
- ▶ cannot minimize c directly, because P is not known
- ▶ Instead, minimize the *regularized risk functional*

$$\frac{1}{2} \|w\|^2 + C \cdot R_{\text{emp}}^\epsilon, \quad \text{where} \quad R_{\text{emp}}^\epsilon := \frac{1}{m} \sum_{i=1}^m |y_i - f(x_i)|_\epsilon$$

- ▶ R_{emp}^ϵ measures the ϵ -insensitive training error
- ▶ C controls the trade-off between margin and training error

Main idea

- ▶ minimize the *regularized risk functional*

$$\frac{1}{2} \|w\|^2 + C \cdot R_{\text{emp}}^\epsilon$$

- ▶ where $R_{\text{emp}}^\epsilon := \frac{1}{m} \sum_{i=1}^m |y_i - f(x_i)|_\epsilon$ measures the training error
- ▶ constant C determines the trade-off

To obtain a small risk

- ▶ control both the training error (R_{emp}^ϵ)
- ▶ and the model complexity ($\|w\|^2$)
- ▶ in short, “explain the data with a simple model”

ϵ -SVR objective function

- ▶ again, rewrite as (soft-margin) optimization problem:

$$\text{minimize}_{w_r \in \mathcal{H}, \xi^{(*)} \in \mathbb{R}^m, b \in \mathbb{R}} \frac{1}{2} \|w\|^2 + \frac{C}{m} \sum_{i=1}^m (\xi_i + \xi_i^*)$$

subject to

- ▶ $(\langle x, x_i \rangle + b) - y_i \leq \epsilon + \xi_i$
- ▶ $y_i - (\langle w, x_i \rangle + b) \leq \epsilon + \xi_i^*$
- ▶ $\xi_i^{(*)} \geq 0$
- ▶ where $(*)$ means both the variables with and without asterisks.



ϵ -SVR dual problem

- ▶ introduce two sets of Lagrange multipliers $\alpha_i^{(*)}$ and $\eta_i^{(*)}$
- ▶ and minimize

$$\begin{aligned}
 L := & \frac{1}{2} \|w\|^2 + \frac{C}{m} \sum_{i=1}^m (\xi_i + \xi_i^*) - \sum_{i=1}^m (\eta_i \xi_i + \eta_i^* \xi_i^*) \\
 & - \sum_{i=1}^m \alpha_i (\epsilon + \xi_i + y_i - \langle w, x_i \rangle - b) \\
 & - \sum_{i=1}^m \alpha_i^* (\epsilon + \xi_i^* + y_i - \langle w, x_i \rangle - b)
 \end{aligned}$$

- ▶ subject to $\alpha_i^{(*)}, \eta_i^{(*)} \geq 0$.



SV expansion

For the whole details, see LWK 9.2 (p. 254ff)


- ▶ solution of the optimization problem results in the SV expansion

$$f(x) = \sum_{i=1}^m (\alpha_i^* - \alpha_i) \langle x_i, x \rangle + b$$

- ▶ w can be written as a linear combination
- ▶ of a subset of the training patterns x_i
- ▶ algorithm described in terms of dot products between the data
- ▶ when evaluating $f(x)$, we need not to compute w explicitly




Example: function approximation




- ▶ `xreg1dls` from SVM-KM toolbox (gaussian basis function)
- ▶ only a few support-vectors

Example: function approximation



- ▶ `xreg1dls` from SVM-KM toolbox (4th-order polynomial)
- ▶ fails to approximate the target function $\sin(\exp(x))$

Example: function approximation



- ▶ `xreg1dls` from SVM-KM toolbox (ht radial basis function)
- ▶ very many support-vectors, but good approximation



Summary

- ▶ Support Vector Machine
 - ▶ maximum-margin linear classifier
 - ▶ concept of support vectors
 - ▶ soft-margin classifier
 - ▶ feature-maps and kernels to handle non-linearity
 - ▶ training via Quadratic Programming algorithms

- ▶ (multi-class) classification and clustering
- ▶ pattern and object recognition
- ▶ regression and function approximation
- ▶ algorithms and complexity estimates

- ▶ still an active research topic



References: web resources

- ▶ see full references in part one (AL 3a)
- ▶ L. Bottou, O. Chapelle, D. DeCoste, J. Weste (Eds), *Large-Scale Kernel Machines*, MIT Press, 2007
- ▶ C. J. C. Burges, *A Tutorial on Support Vector Machines for Pattern Recognition*, *Data Mining and Knowledge Discovery* 2, 121–167 (1998)
- ▶ A. Ben-Hur, D. Horn, H. T. Siegelmann, V. Vapnik, *Support Vector Clustering*, *Journal of Machine Learning Research* 2, 125–137 (2001)



Software: libsvm

- ▶ C.-C. Chang & C.-J. Lin, *libsvm*
<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>
Bindings to C/C++, Java, ...
- ▶ Alain Rakotomamonjy, Stephane Canu, *SVM and Kernel Methods Matlab Toolbox*, <http://asi.insa-rouen.fr/enseignants/~arakotom/toolbox/index.html>
- ▶ W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipes – The Art of Scientific Computing*, Cambridge University Press, 2007 (all algorithms on CD-ROM)
- ▶ several other software packages (Matlab, C/C++, ...)

Example datasets

- ▶ the libsvm page links to several training datasets
<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>
- ▶ MNIST handwritten digits
<http://yann.lecun.com/exdb/mnist/>
- ▶ NORB object recognition datasets
<http://www.cs.nyu.edu/~ylclab/data/norb-v1.0/>