

Fast low-rank metric learning

Dan-Theodor Oneață



Master of Science
Artificial Intelligence
School of Informatics
University of Edinburgh
2011

(Revised: October 2011)

Abstract

In this thesis we focus on Neighbourhood Component Analysis (NCA; Goldberger et al., 2004) — an elegant metric learning framework developed in the context of k nearest neighbours (k NN). NCA learns a (low-rank) metric by maximizing the expected number of correctly classified points. The objective function depends on the pairwise distances between the points. Hence, the learning procedure scales quadratically with the number of samples, making the algorithm impractical when applied on large data sets.

We propose two families of algorithms for reducing the computational cost of the NCA method. First, we consider ideas inspired by the sub-sampling methods. We investigate a mini-batch method that forms mini-batches by clustering. Our results indicate that this approach works well for medium-sized mini-batches. Next, we present a sub-set learning algorithm that is theoretically justified by stochastic optimization arguments. Our experiments demonstrate that this method offers similar results to classical NCA.

The second family of algorithms includes variants of approximate methods. We derive these methods by first interpreting NCA as a class-conditional kernel-density estimation (CC-KDE) problem. This formulation offers several advantages: (i) it allows us to adapt existing algorithms for fast kernel-density estimation into the context of NCA and (ii) it offers more flexibility; for example, we develop a compact support version NCA method that achieves considerable speed-ups when combined with the stochastic learning procedure.

Acknowledgements

I would like to thank my supervisor, Iain Murray, for accepting me as his student, for his continuous support and for his patience. I benefited greatly from his careful advice and feedback throughout the development of this project.

I am grateful to acknowledge that my studies were supported by the “Dinu Patriciu” Foundation — it was a great opportunity that they have offered me.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Dan-Theodor Oneață)

Table of Contents

1	Introduction	1
1.1	Neighbourhood component analysis	2
1.2	Reducing the computational cost	3
2	Background	5
2.1	Theoretical background	5
2.2	Related methods	7
3	Neighbourhood component analysis	10
3.1	General presentation	10
3.2	Class-conditional kernel density estimation interpretation	14
3.3	Practical notes	16
3.3.1	Optimization methods	16
3.3.2	Initialization	19
3.3.3	Numerical issues	22
3.3.4	Regularization	23
3.3.5	Doing classification	24
3.3.6	Dimensionality annealing	24
4	Reducing the computational cost	27
4.1	Sub-sampling	27
4.2	Mini-batches	28
4.3	Stochastic learning	31
4.4	Approximate computations	33
4.4.1	k -d trees	34
4.4.2	Approximate kernel density estimation	37
4.4.3	Approximate KDE for NCA	39
4.5	Exact computations	40

4.5.1	NCA with compact support kernels and background distribution	43
5	Evaluation	46
5.1	Evaluation setup	46
5.2	Baseline	48
5.3	Mini-batches methods	50
5.3.1	Sub-sampling	51
5.3.2	Mini-batches	52
5.3.3	Stochastic learning	52
5.3.4	Comparison	53
5.4	Approximate computations	54
5.4.1	NCA with k -d trees	54
5.4.2	NCA with compact support kernels	57
5.5	Recommendations	59
6	Conclusions	60
6.1	Further directions	62
A	MATLAB code	63
A.1	NCA cost function	63
B	Mathematical derivations	65
B.1	The gradient of $\mathcal{N}(\mathbf{Ax} \mathbf{A}\boldsymbol{\mu}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^T)$ with respect to the linear transformation \mathbf{A}	65
C	Further results	67
	Bibliography	76

List of Figures

2.1	Motivating the need for a distance metric — different tasks are solved by different metrics	7
2.2	Equivalence between the Mahalanobis metric and its associated linear projection	8
3.1	NCA as a class-conditional kernel density estimation model	14
3.2	Illustration of the “early-stopping” technique	19
3.3	Effects of initialization on <code>wine</code> data set	20
3.4	Dimensionality annealing example	26
4.1	Associated problems with sub-sampling method	28
4.2	Evolution of the stochastic assignments p_{ij} during training	33
4.3	k -d tree illustration	35
4.4	NCA with compact support kernels and normal background distribution	44
5.1	Test accuracy as function of reduced dimensionality d on the <code>landsat</code> data	51
5.2	Time vs. accuracy plots for four of the methods proposed: sub-sampling, mini-batches, stochastic learning and stochastic learning for compact support NCA	55
5.3	Two dimensional projections of <code>usps</code> data set using two variants of NCA learning: mini-batches and stochastic learning	56
5.4	Two dimensional projections of <code>magic</code> data set using two variants of NCA learning: mini-batches and stochastic learning	56
5.5	Fraction of the points visited during the training of the compact support NCA	58

5.6	Two dimensional projections of the landsat data set using the stochastic learning for the compact support NCA	58
C.1	Initialization experiments on iris data set.	67
C.2	Initialization experiments on balance data set.	68
C.3	Initialization experiments on ecoli data set.	69

List of Tables

5.1	List of the data sets used for evaluation	47
5.2	NCA accuracy on four small data sets	49
5.3	Accuracy of PCA, LDA and RCA on large data sets	49
5.4	Accuracy for the sub-sampling method on large data sets	51
5.5	Accuracy for the mini-batch method on large data sets	52
5.6	Accuracy for the stochastic learning method on small data sets . .	53
5.7	Accuracy for the stochastic learning method on large data sets . .	53
5.8	Accuracy for the approximated version of NCA trained using stochastic learning on the landsat data set	57
5.9	Accuracy for the approximated version of NCA trained using stochastic learning on large data sets	57
5.10	Accuracy scores for compact support NCA trained using stochastic learning on large data sets	59
C.1	Accuracy for NCA optimized with conjugate gradient method and gradient ascent with “bold driver” heuristic (part 1 of 2)	70
C.2	Accuracy for NCA optimized with conjugate gradient method and gradient ascent with “bold driver” heuristic (part 2 of 2)	71
C.3	Accuracy for NCA trained using stochastic learning on small and medium sized data sets (part 1 of 2)	72
C.4	Accuracy for NCA trained using stochastic learning on small and medium sized data sets (part 2 of 2)	73
C.5	Accuracy for NCA, the compact support NCA and the compact support NCA with background distribution (part 1 of 2)	74
C.6	Accuracy for NCA, the compact support NCA and the compact support NCA with background distribution (part 2 of 2)	75

List of Algorithms

3.1	Gradient ascent (batch version)	17
3.2	Dimensionality annealing	25
4.1	NCA learning algorithm using mini-batches (NCA MB)	29
4.2	Farthest point clustering (FPC; Gonzalez, 1985)	30
4.3	Recursive projection clustering (RPC; Chalupka, 2011)	30
4.4	Stochastic learning for NCA (NCA SL)	32
4.5	k -d tree building algorithm	37
4.6	Approximate NCA	40
4.7	Approximate class-conditional kernel density estimation	41

Chapter 1

Introduction

k nearest neighbours (k NN) is one of the oldest and simplest classification methods. It has its origins in an unpublished technical report by Fix and Hodges (1951) and since then it has become standard textbook material (Russell et al., 1996; Mitchell, 1997; Bishop, 2006). In the last 50 years, k NN was present in most of the machine learning related fields (pattern recognition, statistical classification, data mining, information retrieval, data compression) and it plays a role in many applications (e.g., face recognition, plagiarism detection, vector quantization).

The idea behind k NN is intuitive and straightforward: classify a given point according to a majority vote of its neighbours; the selected class is the one that is the most represented amongst the k nearest neighbours. This is easy to implement and it usually represents a good way to approach new problems or data sets. Despite its simplicity k NN is still a powerful tool, performing surprisingly well in practice (Holte, 1993).

There are also other characteristics that make k NN an interesting method. First, k NN makes no assumptions about the underlying structure of the data. No a priori knowledge is needed beforehand, but we let the data “speak for itself”. The accuracy increases with the number of points in the data set and, in fact, it approaches Bayes optimality as the cardinality of the training set goes to infinity and k is sufficiently large (Cover and Hart, 1967). Second, k NN is able to represent complex functions with non-linear decision boundaries by using only simple local approximations. Lastly, k NN operates in a “lazy” fashion. The training data set is just stored and its use is delayed until testing. The quasi-inexistent training allows the easy addition of new training examples.

k NN has some drawbacks that influence both the computational performance

and the quality of its predictions. Since k NN has to store all the exemplars, the memory requirements are directly proportional to the number of instances in the training set. The cardinality of the data also influences the method's speed. All computations are done at testing time, making k NN painfully slow when applied on large data sets.

The accuracy of k NN is closely related to how we define what “close” and “far” mean for our data set and task. Mathematically, the notion of dissimilarity is incorporated into k NN by using different distance metrics. Usually the standard Euclidean metric is not satisfactory and the aim is to find the particular metric that gives the best results on our data set for a given task (section 2.1). There is an entire literature that tries to come up with possible solutions (section 2.2).

1.1 Neighbourhood component analysis

This thesis focuses on neighbourhood component analysis (NCA; Goldberger et al., 2004). The NCA method learns a metric that maximizes the expected number of correctly classified points (section 3.1). Using the NCA metric with k NN usually improves performance over simple k NN, since we use the label information to construct a suitable metric that selects the relevant attributes.

In its initial formulation, NCA considers only a family of distance metrics known as Mahalanobis (or quadratic) metrics; these are characterized by having associated a linear transformation¹ (section 2.1). By restricting the transformation to be non-square we obtain the corresponding low-rank metric. Such a projection results in a low dimensional representation of the original data that reduces the storage needs and the computational expense at test time. Also it alleviates some of the concerns that have been raised regarding the usefulness of nearest neighbours methods for high dimensional data (Beyer et al., 1999; Hinneburg et al., 2000). The curse of dimensionality arises for k NN, because inter-point distances can become indiscernible in many dimensions. For a given distribution the maximum and the minimum distance between points become equal in the limit of many dimensions. NCA proves to be an elegant answer for the above issues and, consequently, it was successfully used in a variety of applications: face recognition (Butman and Goldberger, 2008), hyper-spectral image classification

¹In the original paper, Goldberger et al. (2004) present also a non-linear version of NCA. In this thesis, we will restrict ourselves to linear transformations.

(Weizman and Goldberger, 2007), acoustic modelling (Singh-Miller, 2010) and even reinforcement learning (Keller et al., 2006).

However, the method introduces additional training time. The objective function needs the pairwise distances of the points, so the function evaluation is quadratic in the number of data points. Also the optimization process is iterative. For large data sets, NCA training becomes prohibitively slow. For example, Weinberger et al. (2006) reported that the original implementation ran out of RAM and Weizman and Goldberger (2007) had to use only 10% of their data set in order to successfully train the model. There is only little previous work that uses NCA for large scaled applications. One example is by Singh-Miller (2010), who parallelized the computations across multiple computers and adopted various heuristics to prune terms of the objective function and the gradient.

1.2 Reducing the computational cost

The main aim of this project is to reduce NCA training time without significant losses in accuracy (chapter 4). We start our investigation with the most simple ideas: use only a subset of the data set (section 4.1) or train the metric on different mini-batches subsampled from the data set (section 4.2). This last idea can be further refined by using clustered mini-batches. Also we present an alternative mini-batch algorithm (section 4.3) that decreases the theoretical cost.

These methods can achieve further speedings if we use approximations of the objective function. Simple approximation ideas (such as ignoring the points which are farther away than a certain distance from the current point) were mentioned in the original paper (Goldberger et al., 2004) and they are re-iterated by Weinberger and Tesauro (2007) and Singh-Miller (2010). We present a more principled approach that borrows ideas from fast kernel density estimation problems (Deng and Moore, 1995; Gray and Moore, 2003; Alexander Gray, 2003). We first recast NCA into a class-conditional kernel density estimation framework (section 3.2) and next we present how fast density estimation is done using a space partitioning structure, such as k -d trees (section 4.4). Similar work of fast metric learning was done by (Weinberger and Saul, 2008) which uses ball trees (Omohundro, 1989) to accelerate a different distance metric technique, large margin nearest neighbour (LMNN; Weinberger et al., 2006). However, LMNN is different from NCA and the way in which ball trees are applied also differs from our approach.

An alternative for the approximation method is to change the model such that it allows exact and efficient computations. In the kernel density estimation model, we can use compact support kernels instead of the standard Gaussian functions. The expense is reduced because only the points that lie in the compact support are used for estimating the density (section 4.5).

We evaluate the proposed techniques in terms of accuracy and speed (chapter 5). Also we provide low dimensional representations of the projected data. The results are promising, showing that we can obtain considerable speed-ups for large data sets while retaining almost the same accuracy as in the classic case. For small and medium sized data sets the accuracy and the time spent are similar to the standard NCA.

Chapter 2

Background

This chapter sets the theoretical basis for the following discussion on low-rank metrics. We start with a brief mathematical introduction into distance metrics. Then we outline the characteristics of the Mahalanobis-like metric. At the end of the chapter we review some of the most prominent work in the area. The method we study, neighbourhood component analysis (NCA), is presented in chapter 3.

2.1 Theoretical background

A distance metric is a mapping from a pair of inputs to a scalar. Let $d(x, y)$ denote the distance function between the arguments x and y . Given a set \mathcal{S} , we say that d is a metric on \mathcal{S} if it satisfies the following properties for any $x, y, z \in \mathcal{S}$:

- Non-negativity: $d(x, y) \geq 0$.
- Distinguishability: $d(x, y) = 0 \Leftrightarrow x \equiv y$.
- Symmetry: $d(x, y) = d(y, x)$.
- Triangle Inequality: $d(x, y) \leq d(x, z) + d(z, y)$.

A metric can be defined on any set \mathcal{S} of objects. The distance metric is valid as long as the above properties hold for any elements in the set. However, for the scope of our applications we limit ourselves to points in D dimensional real space. The inputs will be represented by column vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^D$. In this case the distance d is a function from $\mathbb{R}^D \times \mathbb{R}^D$ to \mathbb{R} . A well known example of such a function is the Euclidean metric. This is defined by:

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^T(\mathbf{x} - \mathbf{y})}, \quad \text{with } \mathbf{x}, \mathbf{y} \in \mathbb{R}^D. \quad (2.1)$$

Albeit useful in many scenarios, the Euclidean metric has two properties that are problematic in machine learning applications (Dwinnell, 2006):

Sensitivity to variable scaling. In geometrical situations, the values across all D dimensions are measured in the same unit of length. In practical situations however, each dimension encodes a different feature of the data set. The Euclidean metric is sensitive to scaling and it returns different results if we change the measurement units for one of the D variables. We compensate for these differences by dividing the values on dimension i by the standard deviation s_i on that direction:

$$\begin{aligned} d(\mathbf{x}, \mathbf{y}) &= \sqrt{\left(\frac{x_1 - y_1}{s_1}\right)^2 + \dots + \left(\frac{x_D - y_D}{s_D}\right)^2} = \\ &= \sqrt{(\mathbf{x} - \mathbf{y})^T \mathbf{S}^{-1} (\mathbf{x} - \mathbf{y})}, \end{aligned} \quad (2.2)$$

where $\mathbf{S} = \text{diag}(s_1^2, \dots, s_D^2)$ and $s_i^2 = \text{var}[x_i], \forall i = 1, \dots, D$.

Invariance to correlated variables. Euclidean distance does not take into account any associations between variables. We can take into account the correlations between attributes if we replace \mathbf{S} in equation (2.2) with the covariance matrix of the data:

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^T \mathbf{S}^{-1} (\mathbf{x} - \mathbf{y})}, \quad (2.3)$$

where $\mathbf{S} = \text{cov}[\mathcal{D}]$, \mathcal{D} is the data set. The metric defined in equation (2.3) is known as the Mahalanobis distance.

The Mahalanobis metric extends the Euclidean distance by incorporating data specific information. But a good metric should be even more flexible than that. Apart from the data set, we are also given a task to perform. We want our metric to extract and discriminate only between information that is relevant to our given task. Take the example of images and face recognition. The pixels from the image background are highly correlated and they reflect the same information: the colour of the background. If the two pictures of the same subject have different backgrounds, we get a high distance between the images. Since we are not interested at all in the background colour, the metric should identify all the pixels in the background as non-relevant and down-weight their importance. In figure 2.1, we illustrate why we need our metric should depend on a given task and why for a given data set it is not optimal to use a fixed metric.

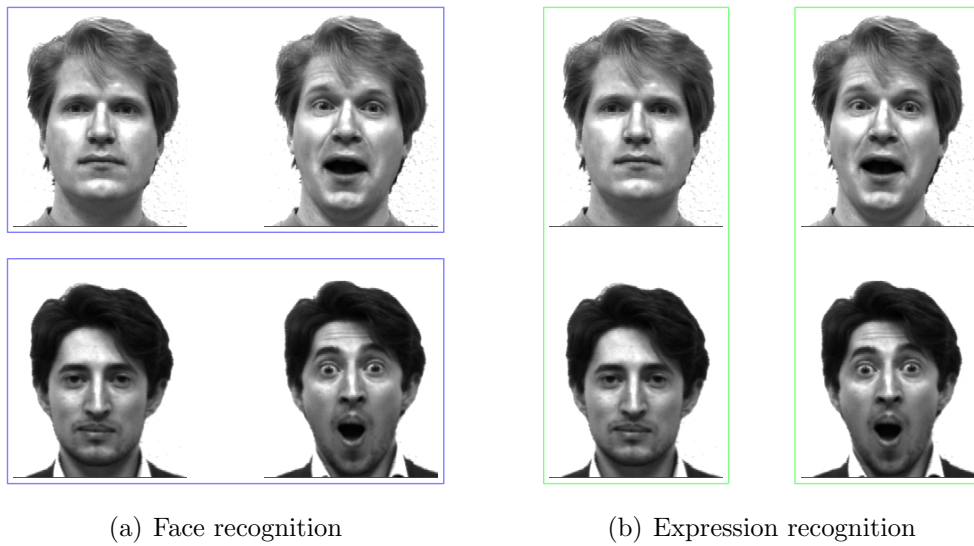


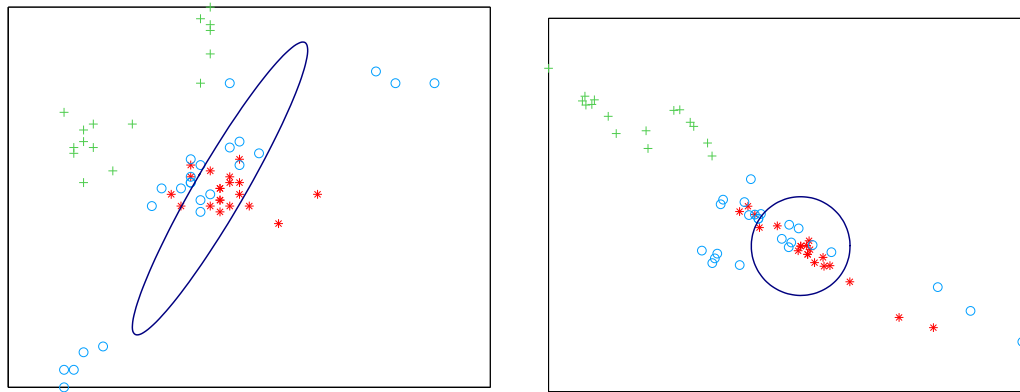
Figure 2.1: On a given data set we might need to perform multiple tasks; in this example: face recognition and expression recognition. The Euclidean metric is not optimal since we want different pairs of equivalence for the two tasks. A good distance metric should be task dependent. These images are an extract from Yale face database.

We can further generalize the result from equation (2.3). Instead of setting \mathbf{S} as the covariance of the data, we let \mathbf{S} be a different matrix. The problem of determining a suitable matrix \mathbf{S}^{-1} is called Mahalanobis (or quadratic) distance metric learning. We note that \mathbf{S}^{-1} ought to be positive semi-definite $\mathbf{x}^T \mathbf{S}^{-1} \mathbf{x} \geq 0, \forall \mathbf{x}$. This ensures that the norm is a real value: $d(\mathbf{x}, \mathbf{0}) \equiv \|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{S}^{-1} \mathbf{x}} \in \mathbb{R}$.

There is an interesting equivalence between a Mahalanobis-like metric and a linear transformation. Using the Cholesky decomposition or other matrix square root we can write $\mathbf{S}^{-1} = \mathbf{A}^T \mathbf{A}$, which gives $\|\mathbf{x}\| = \sqrt{(\mathbf{A}\mathbf{x})^T (\mathbf{A}\mathbf{x})}$, where \mathbf{A} represents a linear transformation of the data. Hence, the problem of finding a suitable distance metric \mathbf{S} is equivalent to the problem of finding a good linear transformation \mathbf{A} and then applying Euclidean distance in the projected space (figure 2.2). We will discuss these two variants interchangeably and we will often parametrize our models in terms of \mathbf{A} .

2.2 Related methods

The first attempts in metric learning were simple alterations of the Euclidean metric to improve k NN performance. For example, Mitchell (1997) suggests weighting



(a) Mahalanobis metric in the original space (b) Euclidean metric in the projected space

Figure 2.2: Illustration of the equivalence between a Mahalanobis metric in the original data space and an Euclidean metric in the transformed data space. The metrics are denoted by an ellipse and, respectively, a circle whose contours indicate equal distance from the centre point to the rest of the points.

each attribute of the data points such that the cross validation error is minimized. This technique is equivalent to stretching the axis of the relevant attributes and compressing the axis of the attributes that are less informative. The approach is a special case of learning a diagonal metric, equation (2.2). Another method (Moore and Lee, 1994) is to eliminate the least relevant attributes: they are given a weight of zero. Albeit useful, these techniques are restrictive as the feature scaling does not reflect how the data covary.

Hastie and Tibshirani (1996) proposed discriminant adaptive nearest neighbours (DANN), a method that constructs a local metric for each particular query. The metric is chosen such that it provides the best discrimination amongst classes in the neighbourhood of the query point. This increases the already costly testing time of k NN. However, DANN method provides a non-linear metric if we take into account all the possible resulting local linear distance metrics. Even though non-linear transformations are an interesting domain, we will concentrate our attention on linear metrics as they are less prone to over-fitting and easier to learn and represent.

Xing et al. (2003) proposed learning a Mahalanobis-like distance metric for clustering in a semi-supervised context. There are specified pairs of similar data points and the algorithm finds the linear projection that minimizes the distance

between these pairs, while imposing the total distance between points that are not similar to be larger than a constant. The solution is attained using an iterative procedure, such as Newton-Raphson. This method assumes that the underlying density of the data is unimodal which leverages the assumption-free k NN. Another drawback is the training time, because the iterative process is costly for large data sets.

Relevant component analysis (RCA; Bar-Hillel et al., 2003; Shental et al., 2002) is another method that makes use of the labels of the classes in a semi-supervised way to provide a distance metric. More precisely, RCA the so-called chunklet information. A chunklet contains points from the same class, but the class label is not known; data points that belong to the same chunklet have the same label, while data points from different chunklets do not necessarily belong to different classes. The main idea behind RCA is to find a linear transformation which “whitens” the data with respect to the averaged within-chunklet covariance matrix (Weinberger et al., 2006). Compared to the method of Xing et al. (2003), RCA has the advantage of presenting a closed form solution, but it has even stricter assumptions about the data: it assumes that the data points in each class are normally distributed so they can be described using only second order statistics (Goldberger et al., 2004).

Large margin nearest neighbour (LMNN; Weinberger et al., 2006) is a metric learning method that was designed especially for k NN classification. LMNN has some similar characteristics to the support vector machine (SVM) technique. As SVM, LMNN brings together points from the same class trying to separate the classes by a certain margin. Also LMNN inherits the convexity property and can be optimized as a semi-definite programming. Weinberger and Saul (2008) worked on fast metric learning by introducing ball trees to speed up the computations. In (Weinberger and Saul, 2008), they extended LMNN to learn local metrics further improving its performance.

Metric learning is an ongoing research area in machine learning with an abundance of different techniques. For a good report that covers many related methods, the reader is referred to (Yang and Jin, 2006).

Chapter 3

Neighbourhood component analysis

This chapter can be regarded as a self-contained tutorial on neighbourhood component analysis (NCA; Goldberger et al., 2004). Much of the material is especially useful for the reader interested in applying the algorithm in practice. We start with a review of the NCA method (section 3.1). Next we recast the NCA model into a class conditional kernel density estimation problem (section 3.2). This new formulation will prove useful later when we want to alter the model in a principled way (section 4.5.1). Lastly, in section 3.3 we present the lessons learnt from our experience with NCA. Both sections 3.2 and 3.3 offer new insights into the NCA method and they also bring together related ideas from previous work.

3.1 General presentation

Neighbourhood component analysis learns a Mahalanobis metric that improves the performance of k nearest neighbours (k NN). From a practical point of view, NCA is regarded as a desirable additional step before doing classification with k NN. It improves the classification accuracy and it can provide a low-dimensional representation of the data.

Because the goal is to enhance k NN's performance, the first idea Goldberger et al. (2004) had was to maximize the leave one out cross validation performance with respect to a linear projection \mathbf{A} . The procedure can be described as follows: apply the linear transformation \mathbf{A} to the whole data set, then take each point \mathbf{Ax}_i and classify it using k NN on the remainder of the transformed data set $\{\mathbf{Ax}_j\}_{j \neq i}$. The matrix \mathbf{A} that achieves the highest number of correct classifications will be used for testing. However, finding the optimal \mathbf{A} is not easy. Any

objective function based on the k nearest neighbours is piecewise constant and discontinuous and, hence, hard to optimize. The reason is that there does not exist a direct relation between \mathbf{A} and the neighbours of a given point: a small perturbation in \mathbf{A} might cause strong changes or, conversely, it might leave the neighbours unchanged.

The authors' solution lies in the concept of *stochastic* nearest neighbours. Remember that in the classical 1-NN scenario, a query point gets the label of the closest point. In the stochastic nearest neighbour case, the query point inherits the label of a neighbour with a probability that decreases with distance. The stochastic function is reminiscent of the generalized logistic function or of the softmax activation used for neural networks. Let p_{ij} denote the probability that the point j is selected as the nearest neighbour of the point i . The value of p_{ij} is given by:

$$p_{ij} = \frac{\exp(-d_{ij}^2)}{\sum_{k \neq i} \exp(-d_{ik}^2)}, \quad (3.1)$$

where d_{ij} represents the distance between point i and point j in the projected space, i.e. $d_{ij} = d(\mathbf{A}\mathbf{x}_i; \mathbf{A}\mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{x}_i - \mathbf{x}_j)$. Also we have to set $p_{ii} = 0$: point i cannot pick itself as the nearest neighbour since we assume its label is not known.

Now we can construct a continuous objective function using the stochastic assignments p_{ij} which are differentiable with respect to \mathbf{A} . A suitable quantity to optimize is the average probability of each point getting correctly classified. Point i is correctly classified when it is selected by a point j that has the same label as i :

$$p_i = \sum_{j \in c_i} p_{ij}. \quad (3.2)$$

The objective function considers each point in the data set and incorporates its probability of being labeled with the true class:

$$\begin{aligned} f(\mathbf{A}) &= \sum_{i=1}^N p_i \\ &= \sum_{i=1}^N \sum_{j \in c_i} \frac{\exp(-d_{ij}^2)}{\sum_{k \neq i} \exp(-d_{ik}^2)}. \end{aligned} \quad (3.3)$$

The score given by the objective function can be interpreted as the expected number of the correctly classified points.

To obtain the best linear projection \mathbf{A} , we maximize $f(\mathbf{A})$ using an iterative gradient based solver such as gradient ascent, conjugate gradients or delta-bar-delta (subsection 3.3.1). For these methods, we can use the analytical expression of the gradient. If we differentiate with respect to \mathbf{A} , we obtain:

$$\frac{\partial f}{\partial \mathbf{A}} = 2\mathbf{A} \sum_{i=1}^N \left(p_i \sum_{k=1}^N p_{ik} \mathbf{x}_{ik} \mathbf{x}_{ik}^T - \sum_{j \in c_i} p_{ij} \mathbf{x}_{ij} \mathbf{x}_{ij}^T \right), \quad (3.4)$$

where $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$.

At test time, we use the learnt matrix \mathbf{A} to transform the points before doing classification with k NN. Given a query point \mathbf{x}^* , we assign it the label c_j of the closest point j in the projected space, where $j = \operatorname{argmin}_i d(\mathbf{A}\mathbf{x}^*, \mathbf{A}\mathbf{x}_i)$.

Further remarks

There are some interesting observations that can be made about NCA. It is useful to note that equations (3.1) to (3.4) hold and NCA model can still be applied even if we set the matrix \mathbf{A} to be rectangular, i.e. $\mathbf{A} \in \mathbb{R}^{d \times D}$, with $d < D$. In this case the linear transformation \mathbf{A} is equivalent to a low dimensional projection. Restricting \mathbf{A} to be non-square gives rise to a couple of advantages. Firstly, the optimization process is less prone to overfitting since a rectangular matrix has fewer parameters than a square one. Secondly, the low dimensional representation of the data set reduces the memory requirements and the computational cost at test time. In this thesis we concentrate mostly on low-rank metrics partly motivated by the above reasons, but also because many techniques that are subsequently developed (chapter 4) work best in low dimensions.

We need to underline that NCA's objective function is not convex. The first consequence is that the parameter space is peppered with local optima and care must be taken to avoid poor solutions. Factors such as the choice of initialization or the optimization procedure affect the quality of the final result. Section 3.3 discusses these practical issues in detail. The other side of non-convexity is that it allows NCA to cope well with data sets that have multi-modal class densities. This is an important advantage over classical methods such as principal component analysis, linear discriminant analysis or relevant component analysis. When treated as probabilistic models, all of these methods assume that the data is normally distributed which harms the classification performance on complicated data sets.

Another important characteristic of NCA is its computational cost. For each point, we need to calculate the distances to all the other points. So the number of operations is quadratic in the number of data points. The evaluation of the objective function is done in two steps. First we find the point projections $\{\mathbf{A}\mathbf{x}_i\}_{i=1}^N$; this operation has $\mathcal{O}(dDN)$ cost. Next we compute all the distances d_{ij} in the low dimensional space; this is done in $\mathcal{O}(dN^2)$ flops. The total cost is then $\mathcal{O}(dDN + dN^2)$. Another way of evaluating the function f is to parametrize it in terms of the Mahalanobis metric $\mathbf{S} = \mathbf{A}^T \mathbf{A}$. Then we can compute the distances in the D -dimensional space in $\mathcal{O}(DN^2)$ flops. This option is slower than the previous case for $d < D$ and $D \ll N$.

Since the optimization is based on the gradient, we are also interested in its cost. The evaluation of gradient given in the original paper and in equation (3.4) scales in D^2N^2 . We improve this by “pushing” the matrix \mathbf{A} into the sum as in (Singh-Miller, 2010):

$$\frac{\partial f}{\partial \mathbf{A}} = 2 \sum_{i=1}^N \left(p_i \sum_{k=1}^N p_{ik} (\mathbf{A}\mathbf{x}_{ik}) \mathbf{x}_{ik}^T - \sum_{j \in c_i} p_{ij} (\mathbf{A}\mathbf{x}_{ij}) \mathbf{x}_{ij}^T \right), \quad (3.5)$$

If we consider that the projections $\mathbf{A}\mathbf{x}_i$ were already calculated for the function evaluation, the gradient has a cost of $\mathcal{O}(dDN^2)$. However, in theory, evaluating the gradient should have the same complexity as evaluating the objective function. Automatic differentiation (AD; Rall, 1981) is a technique that achieve this theoretical result. AD uses the chain rule to derivative the objective function and split the gradient into elementary functions that can be readily computed. So instead of evaluating the complicated analytical gradient, we work on a series of cheap operations. This method is very useful in practice and there exist libraries for this in most programming languages. In our implementation we used the analytical differentiation, because we were not aware of AD until recently.

The computational drawback represents an important setback in applying NCA in practice. For example, Weinberger and Tesauro (2007) reported that the method ran out of memory on large data sets and Weizman and Goldberger (2007) used only 10% of an hyper-spectral data set to train the NCA classifier. We present a number of solutions to this problem in chapter 4.




Figure 3.1: NCA as a class-conditional kernel density estimation model. The figure represents a schematic illustration of a mixture of Gaussians. The red circles denote isotropic Gaussians. These are centred onto the points that belong to a certain class. A point i is generated by either of the components of the class with a probability inverse proportional with the distance.

3.2 Class-conditional kernel density estimation interpretation

In this section we recast NCA into a class-conditional kernel density estimation framework. This interpretation will allow us to understand the assumptions behind the NCA method. After this we are in position to alter the model in a suitable way that is efficient for computations (sections 4.4 and 4.5). Similar ideas were previously presented by Weinberger and Tesauro (2007) and Singh-Miller (2010), but the following were derived independently and they offer different insights. Our following interpretation was inspired by the probabilistic k NN presented by Barber (2011).

We start with the basic assumption that each class can be modelled by a mixture of Gaussians. For each of the N_c data points in class c we consider a Gaussian “bump” centred around it (figure 3.1). From a generative perspective, we can imagine that each point j can generate a point i with a probability given by an isotropic normal distribution with variance σ^2 :

$$p(\mathbf{x}_i|\mathbf{x}_j) = \mathcal{N}(\mathbf{x}_i|\mathbf{x}_j, \sigma^2\mathbf{I}_D) \quad (3.6)$$

$$= \frac{1}{(2\pi)^{D/2}\sigma^D} \exp\left\{-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mathbf{x}_j)^\top(\mathbf{x}_i - \mathbf{x}_j)\right\}. \quad (3.7)$$

By changing the position of the points through a linear transformation \mathbf{A} , the probability changes as follows:

$$p(\mathbf{A}\mathbf{x}_i|\mathbf{A}\mathbf{x}_j) \propto \exp \left\{ -\frac{1}{2\sigma^2} (\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{x}_i - \mathbf{x}_j) \right\}. \quad (3.8)$$

We note that this is similar to the p_{ij} from NCA. Both $p(\mathbf{A}\mathbf{x}_i|\mathbf{A}\mathbf{x}_j)$ and p_{ij} are directly proportional to the same quantity.

Using the mixture of Gaussians assumption, we have that the probability of a point of being generated by class c is equal to the sum of all Gaussians in class c :

$$p(\mathbf{x}_i|c) = \frac{1}{N_c} \sum_{\mathbf{x}_j \in c} p(\mathbf{x}_i|\mathbf{x}_j) \quad (3.9)$$

$$= \frac{1}{N_c} \sum_{\mathbf{x}_j \in c} \mathcal{N}(\mathbf{x}_i|\mathbf{x}_j, \mathbf{I}_D). \quad (3.10)$$

However, we are interested in the inverse probability, given a point \mathbf{x}_i what is the probability of \mathbf{x}_i belonging to class c . We obtain the expression for $p(c|\mathbf{x}_i)$ using Bayes' theorem:

$$p(c|\mathbf{x}_i) = \frac{p(\mathbf{x}_i|c)p(c)}{p(c|\mathbf{x}_i)} = \frac{p(\mathbf{x}_i|c)p(c)}{\sum_c p(\mathbf{x}_i|c)p(c)}. \quad (3.11)$$

Now if we further consider the classes to be equally probable, we arrive at a result that resembles the expression of p_i (equation (3.2)):

$$p(c|\mathbf{A}\mathbf{x}_i) = \frac{\frac{1}{N_c} \sum_{\mathbf{x}_j \in c} \exp \left\{ -\frac{1}{2\sigma^2} (\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{x}_i - \mathbf{x}_j) \right\}}{\frac{1}{N_{c'}} \sum_{c'} \sum_{\mathbf{x}_k \in c'} \exp \left\{ -\frac{1}{2\sigma^2} (\mathbf{x}_i - \mathbf{x}_k)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{x}_i - \mathbf{x}_k) \right\}} \quad (3.12)$$

We are interested in predicting the correct class of the point i . So we try to find the linear transformation \mathbf{A} that maximizes the class conditional probability of this point $p(c_i|\mathbf{A}\mathbf{x}_i)$ to its true class c_i .

$$f(\mathbf{A}) = \sum_i p(c_i|\mathbf{A}\mathbf{x}_i). \quad (3.13)$$

As in section 3.1, we can optimize the function using its gradient information. The gradient is the following:

$$\frac{\partial f}{\partial \mathbf{A}} = \sum_i \left\{ \frac{\frac{\partial p(\mathbf{A}\mathbf{x}_i|c)}{\partial \mathbf{A}} p(c)}{\sum_c p(\mathbf{x}_i|c)p(c)} - \frac{p(\mathbf{A}\mathbf{x}_i|c)p(c)}{\underbrace{\sum_c p(\mathbf{A}\mathbf{x}_i|c)p(c)}_{p(c|\mathbf{A}\mathbf{x}_i)}} \frac{\sum_c \frac{\partial p(\mathbf{A}\mathbf{x}_i|c)}{\partial \mathbf{A}} p(c)}{\sum_c p(\mathbf{A}\mathbf{x}_i|c)p(c)} \right\}. \quad (3.14)$$

The main advantage of the above formulation is that we can construct different models simply by modifying the class-conditional probability (equation (3.9)). Once $p(\mathbf{x}_i|c)$ is altered in a suitable way, we use Bayes rule in conjunction with equations (3.13) and (3.14) to get the new function and its corresponding gradient. We optimize the parameter \mathbf{A} in exactly the same way as for classical NCA.

Examples of this technique are in sections 4.5 and 4.5.1. These alterations allow considerable speed-ups as we shall see in chapter 5.

Other benefits of the CC-KDE model are the possibility to incorporate prior information about the class distributions $p(c)$ and the possibility to classify a point \mathbf{x}^* using $p(c|\mathbf{A}\mathbf{x}^*)$ (subsection 3.3.5).

3.3 Practical notes

While NCA is not that hard to implement (appendix A.1), certain care is needed in order to achieve good solutions. In this section we try to answer some of the questions that can be raised when implementing NCA.

3.3.1 Optimization methods

The function $f(\mathbf{A})$, equation (3.3), can be maximized using any gradient based optimizer. We considered two popular approaches for our implementation: gradient ascent and conjugate gradients. These are only briefly presented here. The interested reader is pointed to (Bishop, 1995).

Gradient ascent

Gradient ascent is one of the simplest optimization methods. It is an iterative algorithm that aims to find the maximum of a function by following the gradient direction at each step. The entire procedure is summarized by algorithm 3.1.

Besides this version, an alternative is to compute the gradient on fewer points than the entire data set. This variant is known as stochastic gradient ascent and has the advantage of being much faster than the batch version on redundant data sets (redundancy is common in practice and especially in large data sets). Different versions of stochastic gradient ascent were tried in chapter 4 to reduce the computational cost.

For the algorithm 3.1, there are three aspects we need to consider:

Algorithm 3.1 Gradient ascent (batch version)

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$.

- 1: Get initial \mathbf{A}_0
 - 2: **repeat**
 - 3: Update parameter: $\mathbf{A}_{t+1} \leftarrow \mathbf{A}_t + \eta \frac{\partial f(\mathbf{A}_t, \mathcal{D})}{\partial \mathbf{A}}$
 - 4: Update learning rate η
 - 5: $t \leftarrow t + 1$
 - 6: **until** convergence
 - 7: **return** \mathbf{A}_t .
-

Initialization. The algorithm starts the search in the parameter space from an initial parameter \mathbf{A}_0 . Different values for \mathbf{A}_0 will give different final solutions. In the NCA context, initialization is related to finding a good initial linear transformation; this is discussed separately in subsection 3.3.2. For the moment, we can assume the values of \mathbf{A}_0 are randomly chosen.

Learning rate. The parameter η is called the *step size* or, in the neural networks literature, it also known as the *learning rate*. As its name suggests, η controls how much we move the parameters in the gradient direction.

The learning rate can be either fixed or adaptive. For the first case, choosing the correct value for η is critical. If η is set too large, the algorithm diverges. On the other hand, a small η results in slow convergence.

Usually, a variable learning rate η is preferred. “Bold driver” (Vogl et al., 1988) is a reasonable heuristic for automatically modifying η during training. The idea is to gradually increase the learning rate after each iteration until the objective function stops improving. Then we go back to the previous position and decrease the learning rate until a new increase in the function is found. The common way to increase the step size η is by multiplying it with a constant $\rho > 1$, usually $\rho = 1.1$. To decrease the step size, we can multiple it with a constant $\sigma < 1$, for example $\sigma = 0.5$.

Another way is to change the learning rate η using a decreasing rule proportional to the current iteration number: $\eta = \frac{\eta_0}{t+t_0}$. Now we have two parameters instead of one. Intuitively, η_0/t_0 can be regarded as the initial learning rate and t_0 as the number of iterations after which η_0 starts to drop off. Using a learning rate of such form is especially motivated in the

stochastic gradient case. If $\eta \propto 1/t$, the algorithm converges to the true minimum as t goes to infinity. However, in practice the convergence is very slow; so setting the constants η_0 and t_0 is important for reaching a good solution in a short time. Leon Bottou¹ recommends setting η_0 to be the regularization value and then select t_0 such that the resulting update will modify \mathbf{A} by a sensible amount. Since we did not use regularization, we followed some of the tricks presented in (LeCun et al., 1998). We fixed the value for η_0 and searched for t_0 on an exponential scale from 0.1 to 1000. We kept that value of t_0 that achieved the highest accuracy on a cross-validation set.

Convergence. The gradient ascent algorithm can be stopped either when we hit a maximum number of iterations or when the learning rate η falls below a certain threshold. Also we can look at the variation in the objective function. If the value of the objective function stops improving we can halt the algorithm. Another popular choice to check convergence is “early stopping”. We monitor the performance of the new \mathbf{A} on a cross validation set. If the objective function does not improve the best score for more than a pre-set number of iterations we stop and return to the previous best solution. This solution prevents overfitting since we stop when we achieve a minimum on a test error and not on the training error. We illustrate this in figure 3.2.

Gradient ascent has known convergence problems and determining good learning rates is often difficult. But it has the advantage of being quite flexible, especially in the stochastic optimization scenario, where we can select any noisy estimate of the gradient in order to replace the total gradient.

Conjugate gradients

The conjugate gradient (CG; Hestenes and Stiefel, 1952) method solves some of the convergence issues. Instead of selecting the search directions to be the gradient directions, the CG method selects the next direction depending on both the previous directions and on the curvature of the surface. CG reaches the maximum of a quadratic energy function in a number of iterations equal to the number of dimensions the function has support on. Also CG eliminates the

¹Oral communication: http://videlectures.net/mmdss07_bottou_lume/

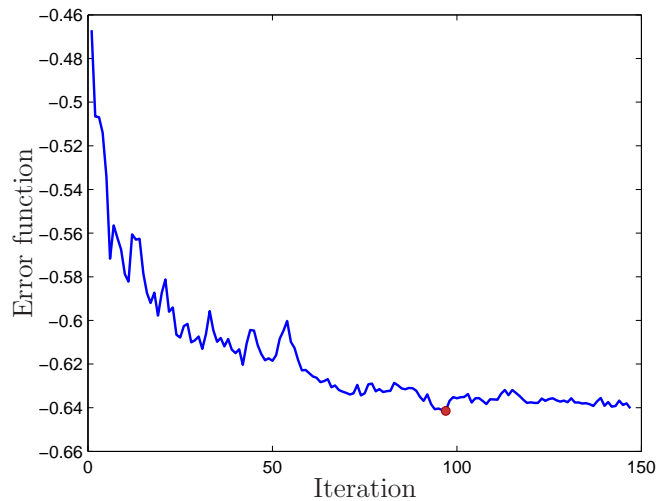


Figure 3.2: This figure illustrates the early stopping procedure. During training we keep a small part of the data set apart, for cross-validation. On this cross-validation data set we compute the error function at each iteration. When the error function stops decreasing we stop the learning procedure and return to the parameter that achieved the best value. This figure resulted while applying NCA on the `mnist` data set. The best value was obtained at iteration 98, as indicated by the red dot.

manual set-up for the learning rates. As Bishop (1995) notes, CG can be regarded as a more complex version of gradient ascent that automatically chooses the parameters. The details of the conjugate gradient method are out of the scope of this thesis; a gentle introduction into the subject is given by Shewchuk (1994).

In practice it is considerably easier to use an existing implementation of CG (for example, `minimize.m` by Rasmussen, 2010) than trying to find the optimal parameters for the learning rate for gradient ascent. For small and medium-sized data sets, the “bold driver” heuristic gives comparable scores to the more sophisticated CG method (tables C.1 and C.2).

3.3.2 Initialization

Initialization is important because the function $f(\mathbf{A})$ is not convex. The quality of the final solution relies on the starting point of the optimization algorithm. A general rule of thumb is to try multiple initial seeds and select the final \mathbf{A} that gives the highest score.

We have already mentioned random initialization in subsection 3.3.1. If the projection is full-rank, $\mathbf{A} \in \mathbb{R}^{D \times D}$, another obvious initialization is the identity

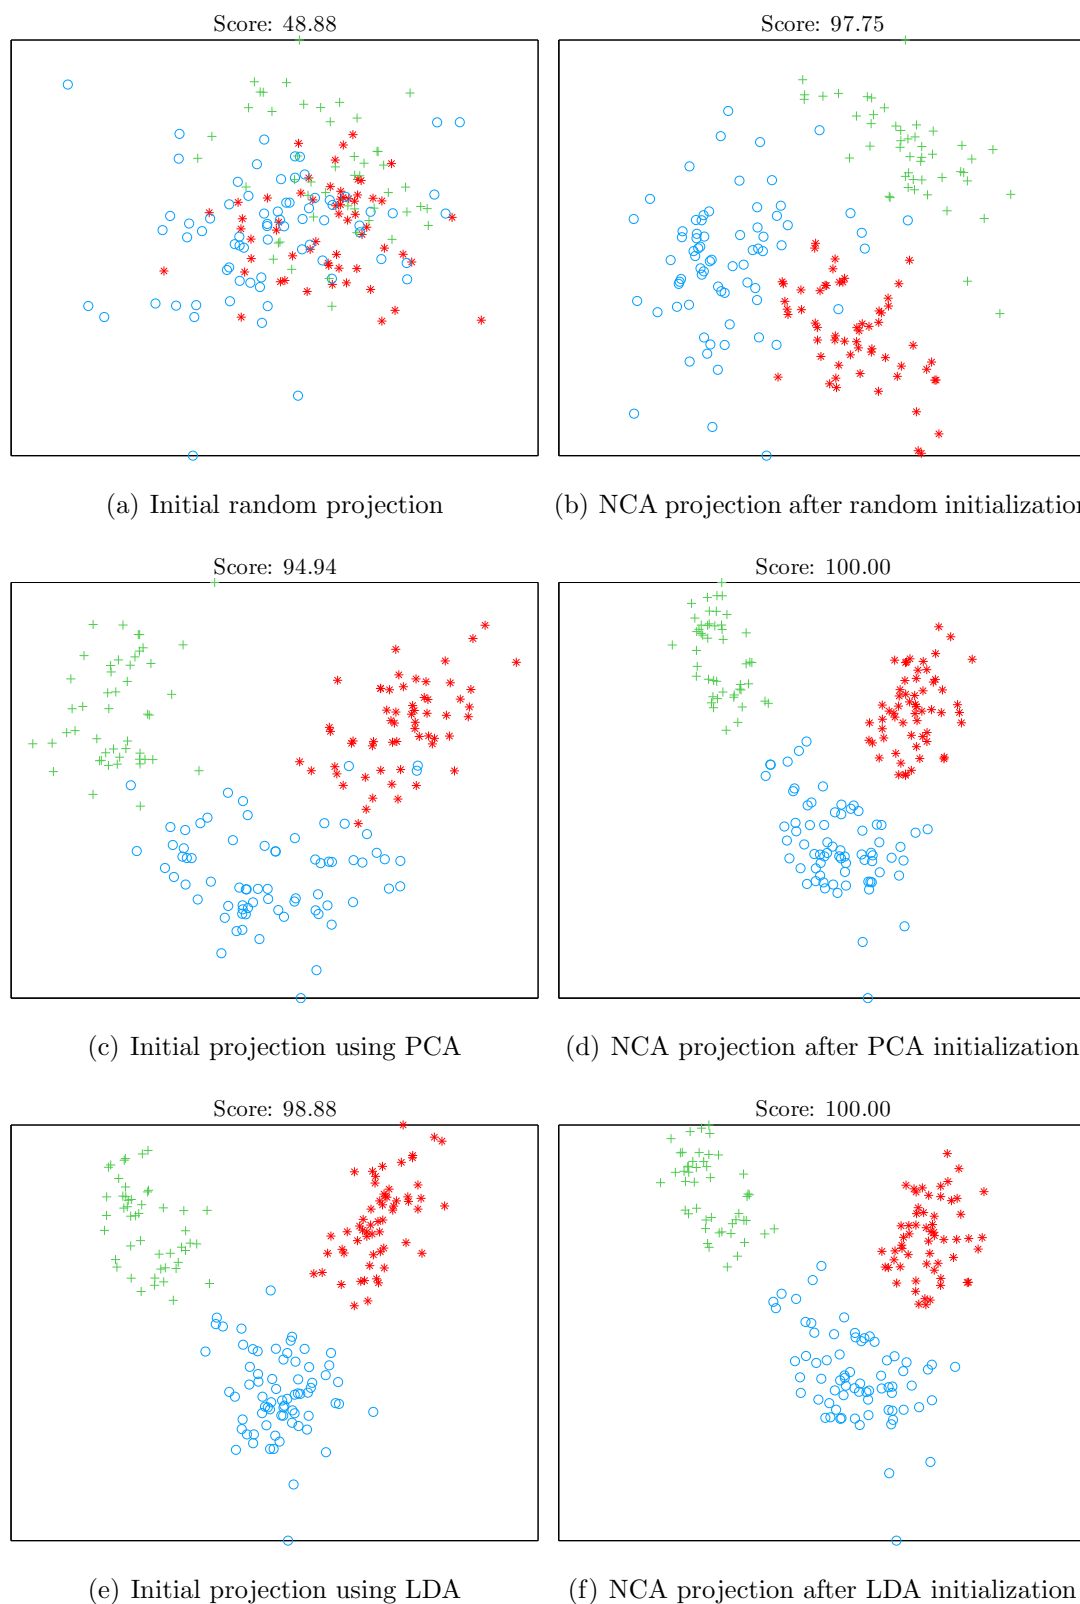


Figure 3.3: Results for different initializations (random, PCA and LDA) on wine data set. The images on the left side represent the projections of the data set using the initial \mathbf{A} . The images on the right side are data set projection with the final \mathbf{A} . Above each figure there is presented the LOOCV score which is normalized to 100.

matrix $\mathbf{A} = \mathbf{I}_D$. Beside these, we can try linear transformations that are cheap to compute. Such examples include principal component analysis (PCA; Pearson, 1901), linear discriminant analysis (LDA; Fisher, 1936) and relevant component analysis (RCA; Bar-Hillel et al., 2003). For completeness, we give here the equations and also include further notes:

- PCA finds an orthogonal linear transformation of the data. This is obtained by computing the eigendecomposition of the outer covariance matrix:

$$\mathbf{S} = \frac{1}{N} \sum_{i=1}^N (\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T. \quad (3.15)$$

The projection \mathbf{A} is then given by $\mathbf{S}^{-1/2}$. If we want to learn a low-rank projection, $\mathbf{A} \in \mathbb{R}^{d \times D}$, $d < D$, we construct \mathbf{A} using only the top d most discriminative eigenvectors, i.e., those eigenvectors that have the highest eigenvalues associated.

- LDA finds a linear transformation \mathbf{A} by maximizing the variance between classes \mathbf{S}_B relative to the amount of within-class variance \mathbf{S}_W :

$$\mathbf{S}_B = \frac{1}{C} \sum_{c=1}^C \boldsymbol{\mu}_c \boldsymbol{\mu}_c^T \quad (3.16)$$

$$\mathbf{S}_W = \frac{1}{N} \sum_{c=1}^C \sum_{i \in c} (\mathbf{x}_i - \boldsymbol{\mu}_c)(\mathbf{x}_i - \boldsymbol{\mu}_c)^T. \quad (3.17)$$

The projection matrix \mathbf{A} that achieves this maximization consists of the eigenvectors of $\mathbf{S}_W^{-1} \mathbf{S}_B$.

Unlike PCA, LDA makes use of the class labels and this usually guarantees a better initial projection.

- RCA finds a linear transformation \mathbf{A} that decorrelates the points within a chunklet, i.e. it makes the within-chunklet covariance equal to the identity matrix. Because for NCA we restrict ourselves to fully labeled data, the within-chunklet covariance is the within-class covariance \mathbf{S}_W , equation (3.17). The “whitening” transformation \mathbf{A} is then $\mathbf{A} = \mathbf{S}_W^{-1/2}$.

From our experiments, we conclude that a good initialization results in a good solution and a faster convergence; these benefits are more evident on large data sets. As advertised by Butman and Goldberger (2008), we found RCA to work

the best. Occasionally random initialization gives better final scores than the other techniques (figure C.2). Figure 3.3 depicts initialization effects on a small data set. More illustrations are shown in appendix C.

3.3.3 Numerical issues

Numerical problems can easily appear when computing the soft assignments p_i . If a point \mathbf{x}_i is far away from the rest of the points, the stochastic probabilities $p_{ij}, \forall j$, are all 0 in numerical precision. Consequently, the result p_i is undetermined: $\frac{0}{0}$. To give an idea of how far \mathbf{x}_i has to be for this to happen, let us consider an example in MATLAB. The answer to $\exp(-d^2)$ is 0 whenever d exceeds 30 units. This is problematic in practice, since distances larger than 30 often appear. Some common cases include data sets that contain outliers or data sets that present a large scale.

The large scale effect can be partially alleviated if we initialize \mathbf{A} with small values. This idea was used by van der Maaten (2010) in his NCA implementation: $\mathbf{A} = 0.01 * \text{randn}(d, D)$. However, this fixed coefficient of 0.01 does not reduce the scale variation of any data set enough to guarantee no numerical issues.

A more robust solution is to normalize the data, i.e., centre and make it have unit marginal variances:

$$x_{ij} \leftarrow \frac{x_{ij} - \mu_j}{\sigma_j}, \quad i = \{1, \dots, N\}, j = \{1, \dots, D\}. \quad (3.18)$$

In this case, we have to store the initial mean and variance of the data and, at test time, transform the test data accordingly: subtract the mean and scale it using the variance coefficients. The variance scaling can be regarded as a linear transformation. We can combine this transformation with the learnt transformation \mathbf{A} and get:

$$\mathbf{A}_{\text{total}} \leftarrow \mathbf{A} \cdot \begin{pmatrix} \sigma_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_D \end{pmatrix}. \quad (3.19)$$

In general, data normalization avoids numerical issues for the first iterations. But during training, the scale of \mathbf{A} increases and data points can be “thrown” away. We adopt a rather simplistic approach to avoid any numerical problems: replace p_i with a very small value whenever we are in the $\frac{0}{0}$ case, as in the

implementation of van der Maaten (2010). In MATLAB, this is done using the following command `max(p_i, realmin)`.

A more rigorous way of dealing with this by multiplying both the numerator and denominator of p_i with a certain quantity $\exp(L)$:

$$p_i = \frac{\sum_{j \in c_i} \exp(L - d_{ij}^2)}{\sum_{k \neq i} \exp(L - d_{ik}^2)}, \quad (3.20)$$

where $L = \min_{k \neq i} d_{ik}^2$. This value of L ensures that at least one term in the denominator does not underflow.

3.3.4 Regularization

Although the original paper (Goldberger et al., 2004) claimed that there were no problems with overfitting, we observed that the NCA objective function has the tendency to increase the scale of the linear projection \mathbf{A} . Butman and Goldberger (2008) pointed out that this is especially problematic for data sets whose size N is significantly smaller than the dimensionality D . In such a case, the linear transformation \mathbf{A} can be chosen to project each point \mathbf{x}_i to a pre-specified a location \mathbf{y}_i . The values of \mathbf{A} are obtained by solving the linear system:

$$\mathbf{A}\mathbf{x}_i = \mathbf{y}_i, \quad i = \{1, \dots, N\}. \quad (3.21)$$

Because the system has more equations dN than unknowns dD , it has exact solutions. If we set the same positions \mathbf{y}_i for points in the same class, we can virtually increase the scale of the projection to infinity and get an error-free classification. This degeneracy can be corrected with the help of regularization. We alter the objective function by subtracting a regularization term proportional to the magnitude of \mathbf{A} . This gives:

$$g(\mathbf{A}) = f(\mathbf{A}) - \lambda \sum_{i=1}^d \sum_{j=1}^D A_{ij}^2. \quad (3.22)$$

$$\frac{\partial g}{\partial \mathbf{A}} = \frac{\partial f}{\partial \mathbf{A}} - 2\lambda \mathbf{A}, \quad (3.23)$$

where λ is a positive constant that is tuned via cross-validation.

Another effect of a large scale linear projection \mathbf{A} is the fact that only the nearest neighbour is considered for each point. This is not usually the optimal thing to do, so regularization is also used to prevent this (Singh-Miller, 2010).

In our implementation, we use an “early stopping” technique which has a similar effect to regularization.

3.3.5 Doing classification

For testing, Goldberger et al. (2004) considered only the k NN decision rule. Since we are optimizing a particular function $f(\mathbf{A})$, another sensible tool for classification is an NCA based function. If we are using the probabilistic interpretation (section 3.2) a query point \mathbf{x}_i is labeled with the most probable class c :

$$c = \operatorname{argmax}_c p(c|\mathbf{x}_i). \quad (3.24)$$

This can be re-written in NCA specific notation as:

$$c = \operatorname{argmax}_c \frac{\sum_{j \in c} p_{ij}}{\sum_k p_{ik}}. \quad (3.25)$$

In our experiments, 1-NN and the NCA classification rule described by equation (3.25) give very similar results if we train \mathbf{A} on the un-regularized function $f(\mathbf{A})$. When we used early stopping, the NCA classification rule yielded better performances (chapter 5). In this case, k NN with $k > 1$ should also perform better than simple 1-NN.

3.3.6 Dimensionality annealing

Dimensionality annealing is an approach to learning a low-rank metric in an way that avoids local optima (Hinton via Murray, oral communication). We start with a full rank projection \mathbf{A} and gradually reduce the rank by introducing regularization terms on the rows of \mathbf{A} . Compared with the classical regularization procedure, subsection 3.3.4, here we use a regularization term on each dimension d . The new objective function and its gradient are given by the following equations:

$$g(\mathbf{A}) = f(\mathbf{A}) - \sum_{i=1}^d \lambda_i \sum_{j=1}^D A_{ij}^2 \quad (3.26)$$

$$\frac{\partial g}{\partial \mathbf{A}} = \frac{\partial f}{\partial \mathbf{A}} - 2 \begin{pmatrix} \lambda_1 A_{11} & \cdots & \lambda_1 A_{1D} \\ \vdots & \ddots & \vdots \\ \lambda_d A_{d1} & \cdots & \lambda_d A_{dD} \end{pmatrix}. \quad (3.27)$$

A large value of λ_d will impose a small magnitude of \mathbf{A} on dimension d . We increase λ_d slowly; this permits the rest of the values of \mathbf{A} to readjust. We clarify these steps in algorithm 3.2.

Algorithm 3.2 Dimensionality annealing

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, initial linear transformation $\mathbf{A} \in \mathbb{R}^{D \times D}$ and low dimension d .

- 1: **for** $i = 1, \dots, D - d$ **do**
 - 2: Select dimension d to anneal
 - 3: **for** $j = 1, \dots, P$ **do**
 - 4: Increase regularization coefficient $\lambda_d \leftarrow \lambda_d + \Delta\lambda$
 - 5: Optimize function g : $\mathbf{A} = \mathbf{A} + \eta \frac{\partial g(\mathbf{A}, \mathcal{D}, \boldsymbol{\lambda})}{\partial \mathbf{A}}$
 - 6: **end for**
 - 7: Remove row d from \mathbf{A}
 - 8: **end for**
-

There are several notes to be made. There are alternatives for selecting the dimension d we want to anneal, step 2 of algorithm 3.2. We can choose d to be the direction that has:

- minimum variance in our data set $\{\mathbf{x}_i\}_{i=1}^N$
- minimum variance in the projected data set $\{\mathbf{A}\mathbf{x}_i\}_{i=1}^N$
- smallest magnitude in \mathbf{A} .

The function optimization step can be done either by gradient ascent or conjugate gradients. It is possible to do conjugate gradients for a small number of iterations, typically 2 or 3. After a dimension is annealed, we reach the end of the inner for loop, then we can remove that particular dimension; for example set the elements on the d -th row equal to 0. A further idea would be to run conjugate gradients until convergence, initialized with low dimensional \mathbf{A} returned by algorithm 3.2.

Figure 3.4 shows a successful example of dimensionality annealing. The two dimensional projection given by the dimensionality annealing technique, figure 3.4(b), is better than the one provided by simple NCA transformation, figure 3.4(a).

Summary

In this chapter we have concentrated on the classic NCA method presenting its theoretical and practical aspects. In section 3.2 we have interpreted NCA

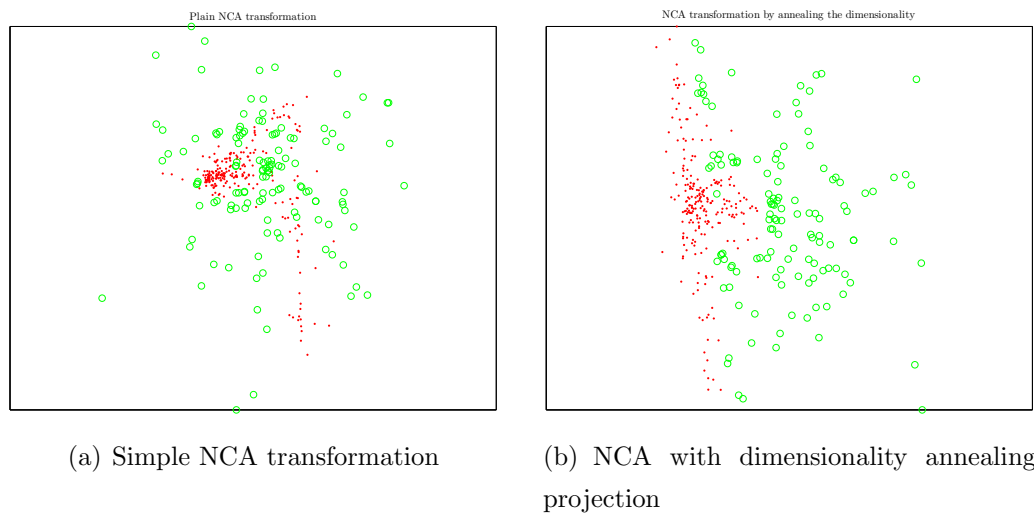


Figure 3.4: Example of applying the dimensionality annealing procedure to ionosphere data set. The method using dimensionality annealing is more visually appealing than the classical approach.

in a more flexible framework. This allows us to easily incorporate prior class probabilities $p(c)$. Usually we can estimate $p(c) = \frac{N_c}{N}$ and this might prove useful for unbalanced class. Furthermore, by modeling the posterior $p(c|\mathbf{x})$ and then using decision theory, we can obtain optimal solutions. Regarding the practical side of NCA, there does not seem to exist a technique that clearly dominates the others, but a combination of RCA for initialization and CG for optimization can be applied with little effort on any small and medium-sized data set. For high dimensional data sets, we recommend regularization or early stopping and then using an NCA-based classification function (equation 3.25 or 3.24). In the next chapter, we will see how we can apply NCA on large data sets.

Chapter 4

Reducing the computational cost

As emphasized in section 3.1, neighbourhood component analysis (NCA) is a computationally expensive method. The evaluation of its objective function is quadratic in the size of the data set. Given that the optimization is done iteratively, NCA becomes prohibitively slow when applied on large data sets. There is only little previous work that uses NCA for large scaled applications. One example is Singh-Miller (2010), who parallelizes the computations across multiple computers and adopts various heuristics to prune terms of the objective function and the gradient.


This chapter aims to continue the existing work and to present a series of methods that can improve NCA's speed. Most of these ideas are new in the context of NCA. We start with simple approaches like sub-sampling (section 4.1) and mini-batch methods (sections 4.2 and 4.3). For further speed-ups, we proceed with more sophisticated ideas (starting from section 4.4).

4.1 Sub-sampling

Sub-sampling is the simplest idea that can help speeding up the computations. For the training procedure we use a randomly selected sub-set \mathcal{D}_n of the original data set \mathcal{D} :

$$\mathcal{D}_n = \{\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_n}\} \subseteq \mathcal{D}.$$

If n is the size of the sub-set then the cost of the gradient, equation (3.5), is reduced to $\mathcal{O}(dDn^2)$. After the projection matrix \mathbf{A} is learnt, we apply it to the whole data set $\{\mathbf{x}_i\}_{i=1}^N$. Then all the new data points $\{\mathbf{A}\mathbf{x}_i\}_{i=1}^N$ are used for



(a) Learnt projection \mathbf{A} on the sub-sampled data set \mathcal{D}_n . (b) The projection \mathbf{A} applied to the whole data set \mathcal{D} .

Figure 4.1: Result of sub-sampling method on `wine`. One third of the original data set were used for training, i.e., $n = N/3$. We note that the points that belong to the sub-set \mathcal{D}_n are perfectly separated. But after applying the metric to the whole data, misclassification errors appear. The effects are more acute if we use smaller sub-sets.

classification. The cost of the classification is $\mathcal{O}(dN)$ which is linear in the total number of points N .

While easy to implement, this method discards a lot of the information available. Also it is affected by the fact the sub-sampled data has a thinner density than the real data. The distances between the randomly selected points are larger than they are in the full data set. This causes the scale of the projection matrix \mathbf{A} not to be large enough. In figure 4.1 we illustrate that a seemingly perfect linear transformation of the sub-sampled data set does not correspond to a good transformation of the entire data set.

4.2 Mini-batches

The next obvious idea is to use sub-sets of data in an iterative manner, similar to the stochastic gradient descent method: split the data into mini-batches and train on them successively. Again the cost for one evaluation of the gradient will be $\mathcal{O}(dDn^2)$ if the mini-batch consists of n points.

A possible way of using mini-batches for learning an NCA projection is summarized by algorithm 4.1. This algorithm should be used in conjunction with the advice on convergence, the learning rate or initialization that was discussed in

sections 3.3.1 and 3.3.2.

Algorithm 4.1 NCA learning algorithm using mini-batches (NCA MB)

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and initial linear transformation \mathbf{A} .

- 1: **repeat**
 - 2: Project each data point using \mathbf{A} : $\mathcal{D}_{\mathbf{A}} = \{\mathbf{A}\mathbf{x}_1, \dots, \mathbf{A}\mathbf{x}_N\}$.
 - 3: Use either algorithm 4.2 or 4.3 on $\mathcal{D}_{\mathbf{A}}$ to split \mathcal{D} into K mini-batches $\mathcal{M}_1, \dots, \mathcal{M}_K$.
 - 4: **for all** \mathcal{M}_i **do**
 - 5: Update parameter: $\mathbf{A} \leftarrow \mathbf{A} + \eta \frac{\partial f(\mathbf{A}, \mathcal{M}_i)}{\partial \mathbf{A}}$.
 - 6: Update learning rate η .
 - 7: **end for**
 - 8: **until** convergence.
-

We considered two different choices for splitting the data-set:

Random selection. In this case the points are assigned randomly to each mini-batch. After one pass through the whole data set, we repeat the random allocation procedure. As in section 4.1, this suffers from the thin distribution problem. In order to alleviate this problem and achieve convergence, we should use large-sized mini-batches (as in the NCA implementation of van der Maaten, 2010). The algorithm is similar to algorithm 4.1, but lines 2 and 3 will be replaced with a simple random selection.

Clustering. Constructing mini-batches by clustering ensures that the point density in each mini-batch is conserved. In order to maintain a low computational cost, we consider cheap clustering methods, e.g., farthest point clustering (FPC; Gonzalez, 1985) and recursive projection clustering (RPC; Chalupka, 2011).

FPC gradually selects cluster centres until it reaches the desired number of clusters K . The point which is the farthest away from all the current centres is selected as the new centre. The precise procedure is given by algorithm 4.2.

The computational cost of this method is $\mathcal{O}(NK)$. However, we do not have any control over the number of points in each cluster, so we might end up with very unbalanced clusters. A very uneven split has a couple of

Algorithm 4.2 Farthest point clustering (FPC; Gonzalez, 1985)

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and K number of clusters.

- 1: Randomly pick a point that will be the first centre \mathbf{c}_1 .
 - 2: Allocate all the points in the first cluster $\mathcal{M}_1 \leftarrow \mathcal{D}$.
 - 3: **for** $i = 1$ to K **do**
 - 4: Select the i -th cluster centre \mathbf{c}_i as the point that is farthest away from any cluster centre $\mathbf{c}_1, \dots, \mathbf{c}_{i-1}$.
 - 5: Move to the cluster \mathcal{M}_i those points that are closer to its centre than to any other cluster centre: $\mathcal{M}_i = \{\mathbf{x} \in \mathcal{D} \mid d(\mathbf{x}; \mathbf{c}_i) < d(\mathbf{x}; \mathbf{c}_j), \forall j \neq i\}$
 - 6: **end for**
-

obvious drawbacks: too large mini-batches will maintain high cost, while on too small clusters there is not too much to learn.

An alternative is RPC which was especially designed to mitigate this problem. It constructs the clusters similarly to how the k -d trees are built, (subsection 4.4.1). But instead of splitting the data set across axis aligned directions it chooses the splitting directions randomly (algorithm 4.3). Because RPC uses the median value it will result in similar sized clusters and we can easily control the dimension of each cluster.

Algorithm 4.3 Recursive projection clustering (RPC; Chalupka, 2011)

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and n size of clusters.

- 1: **if** $N < n$ **then**
 - 2: New cluster: $i \leftarrow i + 1$.
 - 3: **return** current points as a cluster: $\mathcal{M}_i \leftarrow \mathcal{D}$.
 - 4: **else**
 - 5: Randomly select two points \mathbf{x}_j and \mathbf{x}_k from \mathcal{D} .
 - 6: Project all data points onto the line defined by \mathbf{x}_j and \mathbf{x}_k .
 - 7: Select the median value $\tilde{\mathbf{x}}$ from the projected points.
 - 8: Recurse on the data points above and below $\tilde{\mathbf{x}}$: $\text{RPC}(\mathcal{D}_{>\tilde{\mathbf{x}}})$ and $\text{RPC}(\mathcal{D}_{\leq\tilde{\mathbf{x}}})$.
 - 9: **end if**
-

In algorithm 4.1, we are re-clustering in the transformed space after one sweep through the whole data set. There are other alternatives. For example, we could cluster in the original space, either once or periodically. The proposed variant

works well in the case of a low-rank projection matrix \mathbf{A} : we obtain better clusters using RPC on low dimensional data sets than on high dimensional data sets.

4.3 Stochastic learning

The following technique is theoretically justified by stochastic approximation arguments. The main idea is to get an unbiased estimator of the gradient by looking only at a few points and how they relate to the *entire* data set. More precisely, at each iteration we randomly select n data points and find their stochastic nearest neighbours assignments using all the data. We maximize the probability of these n points belonging to the true class:

$$f_{\text{NCA-SL}}(\mathbf{A}) = \sum_{l=1}^n p_{i_l}, \quad (4.1)$$

where i_1, \dots, i_n denote the indices of the randomly selected points and p_i is the average probability of the point i of getting correctly classified:

$$p_i = \sum_{\substack{j=1 \\ j \in c_i}}^N p_{ij}. \quad (4.2)$$

The gradient of the new objective function is given by:

$$\frac{\hat{\partial} f}{\partial \mathbf{A}} = \frac{\partial f_{\text{NCA-SL}}}{\partial \mathbf{A}} = \sum_{l=1}^n \frac{\partial p_{i_l}}{\partial \mathbf{A}} \quad (4.3)$$

$$= 2 \sum_{l=1}^n \left(p_{i_l} \sum_{k=1}^N p_{i_l k} (\mathbf{A} \mathbf{x}_{i_l k}) \mathbf{x}_{i_l k}^T - \sum_{j \in c_{i_l}} p_{i_l j} (\mathbf{A} \mathbf{x}_{i_l j}) \mathbf{x}_{i_l j}^T \right). \quad (4.4)$$

The evaluation of the objective function and its gradient scales with nN . Algorithm 4.4 presents a stochastic learning procedure for NCA. As before, we refer the reader to the previous subsections 3.3.1 and 3.3.2 for advice regarding the convergence and update of the learning rate.

It might be useful to contrast the NCA SL algorithm with the simple NCA. In the classical NCA learning setting, we update our parameter \mathbf{A} after we have considered each point i in the data set and how it relates to the rest of the training set. In the stochastic learning procedure, we update \mathbf{A} more frequently by considering only n randomly selected points and how they relate to the whole training set. The use of the stochastic gradient is motivated by the fact that at

Algorithm 4.4 Stochastic learning for NCA (NCA SL)

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, n number of points to consider for the gradient estimation, \mathbf{A} initial linear transformation.

- 1: **repeat**
 - 2: Split data \mathcal{D} into groups \mathcal{M}_i of size n .
 - 3: **for all** \mathcal{M}_i **do**
 - 4: Update parameter using gradient given by equation 4.4:

$$\mathbf{A} \leftarrow \mathbf{A} + \eta \frac{\partial f_{\text{NCA-SL}}(\mathbf{A}, \mathcal{M}_i)}{\partial \mathbf{A}}.$$
 - 5: Update learning rate η .
 - 6: **end for**
 - 7: **until** convergence.
-

the initial stages we are far from the optimum solution. If we follow a direction that makes an angle smaller than $\pi/2$ with the true gradient, we get closer to a maximum of the function. Estimated gradients often represent a cheap and practical way of optimizing the parameters.

NCA SL also differs from the mini-batch method because for the mini-batch method the contributions p_i are calculated only between the n points that belong to the mini-batch.

This stochastic learning method comes with an additional facility. It can be used for on-line learning. Given a new point \mathbf{x}_{N+1} we update \mathbf{A} using the derivative $\frac{\partial p_{N+1}}{\partial \mathbf{A}}$.

Interlude

The algorithms we presented are characterized by data sub-sampling. These are easy to apply once we decide on the convergence conditions and a learning rate update rule.

In the following sections we describe variants of algorithms that use approximations. When computing p_i we can ignore some of the individual point contributions p_{ij} . While useful on their own, this family of algorithms can achieve better accelerations when combined with the previous methods.

In section 4.4 we use k -d trees to select the significant stochastic assignments p_{ij} that are significant. To clearly present this method, we need to introduce the k -d tree data structure (subsection 4.4.1) and how it is used for fast kernel density

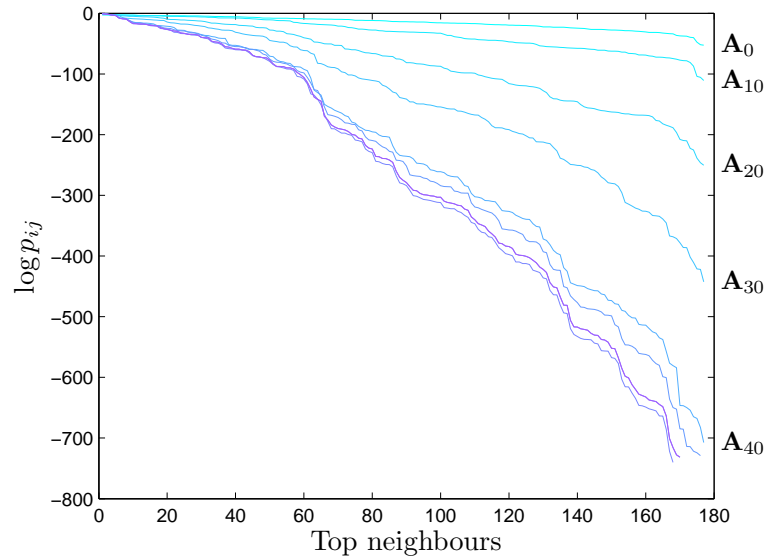


Figure 4.2: Evolution of the stochastic assignments p_{ij} during training for a given point i . The y -axis specifies the contributions p_{ij} on a logarithmic scale. On x -axis we sorted the neighbours in descending order of their contributions. Each curve is associated to a certain linear projection \mathbf{A}_t , where t denotes the iteration number from the optimization algorithm.

estimation (subsection 4.4.2). Subsection 4.4.3 offers details of adapting existing algorithms specifically for NCA. In section 4.5, we change the NCA model into a compact support version of NCA, such that only the points within a certain radius from the query point are considered, while the rest are given 0 weight. Lastly, we present a more robust version of the compact support NCA that avoids possible numerical problems (subsection 4.5.1).

4.4 Approximate computations

A straightforward way of speeding up the computations was previously mentioned in the original paper (Goldberger et al., 2004) and in some NCA related work (Weinberger and Tesauro, 2007; Singh-Miller, 2010). The observations involve pruning small terms in the original objective function. We then use an approximated objective function and its corresponding gradient for the optimization process.

The motivation lies in the fact that the contributions p_{ij} decay very quickly

with distance:

$$p_{ij} \propto \exp\{-d(\mathbf{A}\mathbf{x}_i; \mathbf{A}\mathbf{x}_j)^2\}.$$

The evolution of the contributions during the training period is depicted in figure 4.2. We notice that maybe in the first 5-10 iterations the contributions of more than 50 neighbours are significant; after we get out of this regime we can discard a large part of the neighbours and preserve the accuracy of our estimations.

Weinberger and Tesauro (2007) choose to use only the top $m = 1000$ neighbours for each point \mathbf{x}_i . Also they disregard those points that are farther away than $d_{\max} = 34$ units from the query point: $p_{ij} = 0, \forall \mathbf{x}_j$ such that $d(\mathbf{A}\mathbf{x}_i; \mathbf{A}\mathbf{x}_j) > d_{\max}$. While useful in practical situations, these suggestions lack a principled description: how can we optimally choose m and d_{\max} in a general setting? We would also like to be able to estimate the error introduced by the approximations.

We correct those drawbacks by making use of the KDE formulation of NCA (see section 3.2) and adapting existing ideas for fast KDE (Deng and Moore, 1995; Gray and Moore, 2003) to our particular application. We will use a class of accelerated methods that are based on data partitioning structures (e.g., k -d trees, ball trees). As we shall see shortly, these provide us with means to quickly find only the neighbours \mathbf{x}_j that give significant values p_{ij} for any query point \mathbf{x}_i .

4.4.1 k -d trees

The k dimensional tree structure (k -d tree; Bentley, 1975) organises the data in a binary tree using axis-aligned splitting planes. The k -d tree places points that live nearby in the original geometrical space close in the tree. This makes such structures efficient mechanisms for nearest neighbour searches (Friedman et al., 1977) or range searches (Moore, 1991).

There are different flavours of k -d trees. We choose for our application a variant of k -d tree that uses bounding boxes to describe the position of the points. Intuitively, we can imagine each node of the tree as a bounding hyper-rectangle in the D dimensional space of our data. The root node will represent the whole data set and it can be viewed as a hyper-rectangle that contains all the data points, see figure 4.3(a). In the two-dimensional example presented, the points are enclosed by rectangles. Figures 4.3(a) to 4.3(d) show the existing bounding boxes at different levels of the binary tree. To understand how these are obtained,

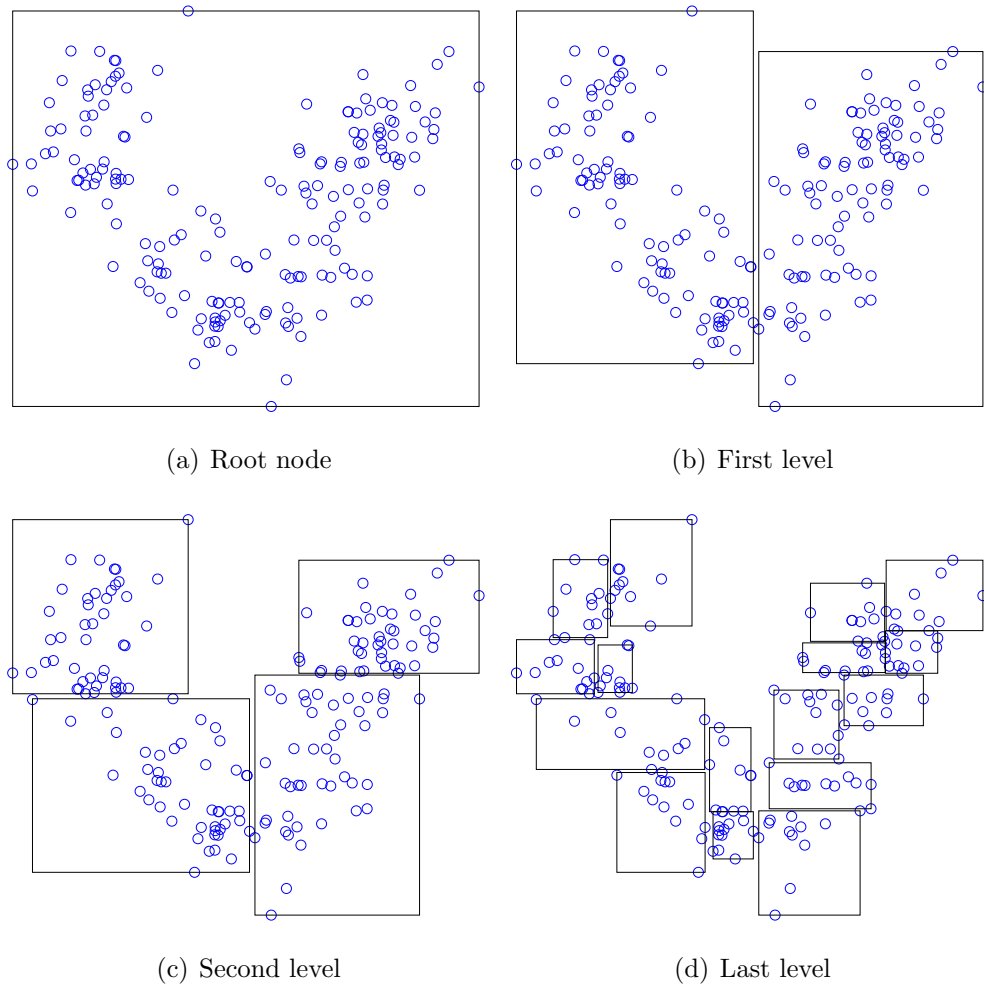


Figure 4.3: Illustration of the k -d tree with bounding boxes at different levels of depths. This figure also outlines the building phases of the tree.

we discuss the k -d tree construction.

We start building the tree from the root node and then proceed in a recursive manner. At each node we select which of the points from the current node will be allocated to each of the two children. Because these are also described by hyper-rectangles, we just have to select a splitting plane. Then the two successors will consist of the points from the two sides of the hyper-plane.

A splitting hyper-plane can be fully defined by two parameters: a direction \vec{d} on which the plane is perpendicular and a point P that is in the plane. Given that the splits are axis aligned, there are D possible directions \vec{d} . We can either choose this randomly or we can use each of the directions from 1 to D in a successive manner. A more common approach is to choose \vec{d} to be the dimension

that presents the largest range:

$$\vec{d} \leftarrow \operatorname{argmax}_d (\max_i x_{id} - \min_i x_{id}). \quad (4.5)$$

Splitting perpendicular on the direction with the greatest range results in a better clustering of the points and the shape of the bounding boxes will be closer to the shape of a square. Otherwise it might happen that points situated in the same node can still be further away.

Regarding the splitting value P , a usual choice is the median value \tilde{x}_d on the previously selected direction \vec{d} . This choice of P guarantees a balanced tree which offers several advantages. We can allocate static memory for the entire data structure. This is faster to access than dynamical allocation. Also a balanced tree has a better worst case complexity than an unbalanced one. Other useful implementation tricks that can be applied to balanced k -d trees are suggested by Lang (2009).

After the splitting plane is chosen, the left child will contain the points that are on the left of the hyper-plane:

$$\mathcal{D}_{\leq \tilde{x}_d} = \{\mathbf{x} \in \mathcal{D}_{\mathbf{x}_i} | x_d \leq \tilde{x}_d\}, \quad (4.6)$$

where $\mathcal{D}_{\mathbf{x}_i}$ denotes the data points bounded by the current node \mathbf{x}_i . Similarly, the right child will contain the points that are placed on the right of the hyper-plane:

$$\mathcal{D}_{> \tilde{x}_d} = \{\mathbf{x} \in \mathcal{D}_{\mathbf{x}_i} | x_d > \tilde{x}_d\}. \quad (4.7)$$

This process is repeated until the number of points bounded by the current node goes below a threshold m . These nodes are the leaves of the tree and they store the data points. The other non-leaf nodes store information regarding the bounding box and the splitting plane. A hyper-rectangle is completely defined by only two D -dimensional points, one for the “top-right” corner and the other for the “bottom-left” corner.

The most common operation on a k -d tree is the nearest neighbour (NN) search. While we will not apply pure NN for the next method, we will use similar concepts. However, we can do NN retrieval with k -d trees after we applied NCA, as suggested by Goldberger et al. (2004). The search in the k -d tree is done in a depth-first search manner: start from the root and traverse the whole tree by selecting the closest node to the query point. In the leaf, we find the nearest neighbour from the m points and store it and the corresponding distance d_{\min} .

Algorithm 4.5 *k*-d tree building algorithm

Require: Data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, i position in tree and m number of points in leaves.

- 1: **if** $N < m$ **then**
 - 2: Mark node i as leaf: `splitting_direction(i) = -1`.
 - 3: Add points to leaf: `points(i) ← \mathcal{D}` .
 - 4: **return**
 - 5: **end if**
 - 6: Choose direction \vec{d} using equation (4.5): `splitting_direction(i) = d` .
 - 7: Find the median value \tilde{x}_d on the direction \vec{d} : `splitting_value(i) = \tilde{x}_d` .
 - 8: Determine the subsets of points that are separated by the resulting splitting plane: $\mathcal{D}_{\leq \tilde{x}}$ and $\mathcal{D}_{> \tilde{x}}$, see equations (4.6) and (4.7).
 - 9: Build left child `build_kdtree($\mathcal{D}_{\leq \tilde{x}}$, $2*i$)`.
 - 10: Build right child `build_kdtree($\mathcal{D}_{> \tilde{x}}$, $2*i+1$)`.
-

Then we recurse up the tree and look at the farther node. If this is situated at a minimum distance that is smaller than d_{\min} , we also have to investigate that node. Otherwise, we can ignore the node and all the points it contains. Usually, a large fraction of the points can be omitted, especially when the data is clustered. It is important to stress that the performance of *k*-d trees quickly degrades with the dimensionality of the data.

4.4.2 Approximate kernel density estimation

The next ideas are mostly inspired by previous work on fast kernel density estimators with *k*-d trees (Gray and Moore, 2001, 2003; Alexander Gray, 2003) and follows closely work on fast Gaussian Processes regression with *k*-d trees (Shen et al., 2006).

In kernel density estimation (KDE), we are given a data set \mathcal{D} and the goal is to compute the probability density at a given point using a sum of the contributions from all the points:

$$p(\mathbf{x}_i) = \frac{1}{N} \sum_{j=1}^N k(\mathbf{x}_i | \mathbf{x}_j). \quad (4.8)$$

A common class of KDE problems are the *N*-body problems (Gray and Moore, 2001) where we have to estimate $p(\mathbf{x}_i)$ for each data point $\{\mathbf{x}_i\}_{i=1}^N$ in the data set

\mathcal{D} . This operation is quadratic in the number of points, as it is the case of NCA.

If the number of samples N in \mathcal{D} is sufficiently large, we expect to accurately approximate $p(\mathbf{x}_i)$ by using only nearby neighbours of \mathbf{x}_i .

To illustrate how this works, let us assume we are given a query point \mathbf{x}_i and a group of points G . We try to reduce computations by replacing each individual contribution $k(\mathbf{x}_i|\mathbf{x}_j)$, $\mathbf{x}_j \in G$, with a fixed quantity $k(\mathbf{x}_i|\mathbf{x}_g)$. The value of $k(\mathbf{x}_i|\mathbf{x}_g)$ is group specific and since it is used for all points in G , it is chosen such that it does not introduce a large error. A reasonable value for $k(\mathbf{x}_i|\mathbf{x}_g)$ is obtained by approximating each point \mathbf{x}_i with a fixed \mathbf{x}_g , for example, the mean of the points in G . Then we compute the kernel value $k(\mathbf{x}_i|\mathbf{x}_g)$ using the estimated \mathbf{x}_g . A second possibility is to directly approximate $k(\mathbf{x}_i|\mathbf{x}_g)$. For example:

$$k(\mathbf{x}_i|\mathbf{x}_g) = \frac{1}{2}(k_{\min} + k_{\max}), \quad (4.9)$$

where k_{\min} and k_{\max} represent the minimum and maximum contributions that can arise for the given query point \mathbf{x}_i and the bounding box that encloses the group G . This last option does not introduce any further computational expense. Both the minimum and the maximum contributions are previously calculated to decide whether to prune or not. Also it does not need to store any additional statistics, such as the mean.

The error introduced by each approximation is bounded by the following quantity:

$$\epsilon_{\max} = \frac{1}{2}(k_{\max} - k_{\min}). \quad (4.10)$$

This can be controlled to be small if we approximate only for those groups that are far away from the query point or when the variation of the kernel value is small within the group. It is better still to consider the error relative to the total quantity $p(\mathbf{x}_i)$. Of course we do not know the total sum we want to estimate in advance, but we can use a lower bound: $p_{\text{SoFar}}(\mathbf{x}_i) + N_G k_{\min}$, where p_{SoFar} denotes the current estimate for $p(\mathbf{x}_i)$ and N_G is the number of points in the group G . Hence, a possible cut-off rule is:

$$\epsilon_{\max} N_G \leq \tau(p_{\text{SoFar}} + N_G k_{\min}), \quad (4.11)$$

where τ is a constant that controls the error introduced by the approximations: a small τ means the computations will more accurate, conversely, a large τ allows

more approximations. We use a pruning technique whenever the cut-off rule is true. So in this case it is important the order in which we accumulate. A large $p_{\text{SoFar}}(\mathbf{x}_i)$ in the early stages will allow more computational savings.

We use k -d trees to form groups of points G that will be described as hyper-rectangles. To compute the probability density function $p(\mathbf{x}_i)$, we start at the root, the largest group, and traverse the tree going through the nearest node each time until we reach the leaf. In this manner, we are able to add large contributions at the beginning. Then we recurse up the tree and visit other nodes only if necessary, when the cut-off condition is not satisfied. We give the exact algorithm for this procedure in the next subsection.

4.4.3 Approximate KDE for NCA

We recall that NCA was formulated as a class-conditional kernel density estimation problem (CC-KDE; section 3.2). By combining ideas from the previous two subsections, we can develop an NCA specific approximation algorithm.

There are some differences from the classical KDE approximation. In this case, we deal with class-conditional probabilities $p(\mathbf{Ax}_i|c), \forall c$. So each class c needs to be treated separately: we build a k -d tree with the projected data points $\{\mathbf{Ax}_j\}_{j \in c}$ and calculate the estimated probability $\hat{p}(\mathbf{Ax}_i|c)$ for each class. Another distinction is that for NCA our final interests are the objective function and its gradient. Algorithm 4.6 presents at a high-level how the NCA objective function and its gradient are computed in the CC-KDE scenario.

We demonstrate how we can obtain an approximated version of the objective function. We begin by replacing $p(\mathbf{Ax}_i|c)$ with the approximated $\hat{p}(\mathbf{Ax}_i|c)$ in equation (3.13):

$$\hat{f}(\mathbf{A}) = \sum_{i=1}^N \hat{p}(c_i|\mathbf{Ax}_i). \quad (4.12)$$

To obtain the gradient of this new objective function we use equation (3.14). Recalling that each individual contribution is given by a normal distribution:

$$k(\mathbf{Ax}|\mathbf{Ax}_i) \propto \exp \left\{ -(\mathbf{Ax} - \mathbf{Ax}_i)^T (\mathbf{Ax} - \mathbf{Ax}_i) \right\}, \quad (4.13)$$

The gradient of each individual contribution with respect to the linear transformation \mathbf{A} will be:

$$\frac{\partial}{\partial \mathbf{A}} k(\mathbf{Ax}|\mathbf{Ax}_i) \propto -2\mathbf{A} \cdot k(\mathbf{Ax}|\mathbf{Ax}_i)(\mathbf{x} - \mathbf{x}_i)(\mathbf{x} - \mathbf{x}_i)^T. \quad (4.14)$$

Algorithm 4.6 Approximate NCA

Require: Projection matrix \mathbf{A} , data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and error τ .

- 1: **for all** classes c **do**
 - 2: Build k -d tree for the points in class c .
 - 3: **end for**
 - 4: **for all** data points \mathbf{x}_i **do**
 - 5: **for all** classes c **do**
 - 6: Compute approximate class-conditional probability $\hat{p}(\mathbf{A}\mathbf{x}_i|c)$ and its corresponding gradient $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}_i|c)$ using algorithm 4.7
 - 7: **end for**
 - 8: Compute soft probability $\hat{p}_i \equiv \hat{p}(c|\mathbf{A}\mathbf{x}_i) = \frac{\hat{p}(\mathbf{A}\mathbf{x}_i|c_i)}{\sum_c \hat{p}(\mathbf{A}\mathbf{x}_i|c)}$.
 - 9: Update function value (equation 3.13) and gradient value (equation 3.14).
 - 10: **end for**
-

Now let us consider the derivative $\frac{\partial}{\partial \mathbf{A}} p(\mathbf{A}\mathbf{x}|G)$ for a group G and then approximate it using the criterion given in equation (4.9). We obtain:

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{A}} p(\mathbf{A}\mathbf{x}|G) &= \frac{\partial}{\partial \mathbf{A}} \sum_{j \in G} k(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{x}_j) \\
&\approx \frac{\partial}{\partial \mathbf{A}} \frac{1}{2} (k_{\min} + k_{\max}) \\
&= \frac{1}{2} \left\{ \frac{\partial}{\partial \mathbf{A}} k(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{x}_c) + \frac{\partial}{\partial \mathbf{A}} k(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{x}_f) \right\} \\
&= -\mathbf{A} \left\{ k_{\max}(\mathbf{x} - \mathbf{x}_c)(\mathbf{x} - \mathbf{x}_c)^{\mathbf{T}} + k_{\min}(\mathbf{x} - \mathbf{x}_f)(\mathbf{x} - \mathbf{x}_f)^{\mathbf{T}} \right\}, \quad (4.15)
\end{aligned}$$

where $\mathbf{A}\mathbf{x}_c$ and $\mathbf{A}\mathbf{x}_f$ denote the points associated with k_{\max} and k_{\min} . We observe that for the gradient computation we need to find these points and then to recover their positions in the original space. The class-conditional approximated gradient is obtained by adding the contributions for all the groups (algorithm 4.7).

4.5 Exact computations

Exact methods are the counterpart of approximate methods. We can have both efficient and exact computations just by modifying the NCA model. Again, the idea is motivated by the rapid decay of the exponential function. Instead of operating on very small values, we will make them exactly zero. This is achieved by replacing the squared exponential kernel with a compact support function. So,

Algorithm 4.7 Approximate class-conditional kernel density estimation

Require: Query point \mathbf{x} , projection matrix \mathbf{A} , data set $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and error τ .

- 1: Find d_{\max} and d_{\min} from \mathbf{x} to the current bounding box and then compute k_{\min} and k_{\max} .
 - 2: **if** $N_G(k_{\max} - k_{\min}) \leq 2\tau(p_{\text{SoFar}} + N_G k_{\min})$ **then**
 - 3: Compute group contribution: $\hat{p}(\mathbf{A}\mathbf{x}|G) = \frac{N_G}{2} (k_{\min} + k_{\max})$
 - 4: Compute derivative $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G)$ using equation (4.15)
 - 5: Update accumulated contribution: $p_{\text{SoFar}} \leftarrow p_{\text{SoFar}} + \hat{p}(\mathbf{A}\mathbf{x}|G)$
 - 6: **return** $\hat{p}(\mathbf{A}\mathbf{x}|G)$ and $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G)$.
 - 7: **end if**
 - 8: **if** current node is a leaf **then**
 - 9: Compute probability: $p(\mathbf{A}\mathbf{x}|G_{\text{leaf}}) = \sum_{j \in \text{leaf}} k(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{x}_j)$
 - 10: Compute derivative: $\frac{\partial}{\partial \mathbf{A}} p(\mathbf{A}\mathbf{x}|G_{\text{leaf}}) = \frac{\partial}{\partial \mathbf{A}} \sum_{j \in G_{\text{leaf}}} k(\mathbf{A}\mathbf{x}|\mathbf{A}\mathbf{x}_j)$
 - 11: Update p_{SoFar}
 - 12: **return** $p(\mathbf{A}\mathbf{x}|G_{\text{leaf}})$ and $\frac{\partial}{\partial \mathbf{A}} p(\mathbf{A}\mathbf{x}|G_{\text{leaf}})$.
 - 13: **end if**
 - 14: Recurse on nearest child to \mathbf{x} and compute the probability and the gradient contributions: $\hat{p}(\mathbf{A}\mathbf{x}|G_{\text{near}})$ and $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G_{\text{near}})$
 - 15: Recurse on farthest child to \mathbf{x} and compute the probability and the gradient contributions: $\hat{p}(\mathbf{A}\mathbf{x}|G_{\text{far}})$ and $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G_{\text{far}})$
 - 16: **return** $\hat{p}(\mathbf{A}\mathbf{x}|G_{\text{near}}) + \hat{p}(\mathbf{A}\mathbf{x}|G_{\text{far}})$ and $\frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G_{\text{near}}) + \frac{\partial}{\partial \mathbf{A}} \hat{p}(\mathbf{A}\mathbf{x}|G_{\text{far}})$.
-

the points that lie outside the support of the kernel are ignored and just a fraction of the total number of points is used for computing the contributions p_{ij} . The most significant cost savings come from the fact that we do not have to compute the expensive terms for the points that have $p_{ij} = 0$. Further gains in speed are obtained if the search for those points is done with k -d trees (the range search algorithm is suitable for this task; Moore, 1991).

The choice of the compact support kernel is restricted by a single requirement: differentiability. We will use the simplest polynomial function that has this property. This is given by the following expression:

$$k_{\text{CS}}(u) = \begin{cases} h(a^2 - u^2)^2 & \text{if } u \in [-a; +a] \\ 0 & \text{otherwise,} \end{cases} \quad (4.16)$$

where h is a constant that controls the height of the kernel and a is a constant that controls the width of the kernel. In the given context, the kernel will be a function of the distance between two points: $k_{\text{CS}}(u) = k_{\text{CS}}(d_{ij})$, where $d_{ij} = d(\mathbf{A}\mathbf{x}_i; \mathbf{A}\mathbf{x}_j)$.

We introduced a general expression for the compact support kernel because later we will need to carefully set the parameters a and h (subsection 4.5.1). For now, we can consider $a = 1$ and $h = 1$. We obtain the following simplified version of the kernel:

$$k_{\text{CS}}(d_{ij}) = (1 - d_{ij}^2)^2 \text{I}(|d_{ij}| \leq 1), \quad (4.17)$$

where $\text{I}(\cdot)$ denotes the indicator function: $\text{I}(\cdot)$ returns 1 when its argument is true and 0 when its argument is false.

Now we reiterate the steps of the NCA algorithm (presented in section 3.1), and replace each $\exp(-d_{ij}^2)$ with $k_{\text{CS}}(d_{ij})$. We obtain the following new stochastic neighbour assignments:

$$q_{ij} = \frac{k_{\text{CS}}(d_{ij})}{\sum_{k \neq i} k_{\text{CS}}(d_{ik})}. \quad (4.18)$$

These can be compared to the classical soft assignments given by equation (3.1). Next we do not need to change the general form of the objective function:

$$f_{\text{CS}}(\mathbf{A}) = \sum_i \sum_{j \in c_i} q_{ij}. \quad (4.19)$$

In order to derive the gradient of the function f_{CS} , we start by computing the

gradient of the kernel:

$$\begin{aligned}\frac{\partial}{\partial \mathbf{A}} k_{\text{CS}}(d_{ij}) &= \frac{\partial}{\partial \mathbf{A}} [(1 - d_{ij}^2)^2 \cdot \mathbf{I}(|d_{ij}| \leq 1)] \\ &= -4\mathbf{A}(1 - d_{ij}^2)\mathbf{x}_{ij}\mathbf{x}_{ij}^{\text{T}} \cdot \mathbf{I}(|d_{ij}| \leq 1),\end{aligned}\quad (4.20)$$

where $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$.

The gradient of the new objective function is:

$$\frac{\partial f_{\text{CS}}}{\partial \mathbf{A}} = 4\mathbf{A} \sum_{i=1}^N \left(q_i \sum_{k=1}^N \frac{q_{ik}}{1 - d_{ik}^2} \mathbf{x}_{ik}\mathbf{x}_{ik}^{\text{T}} - \sum_{j \in c_i} \frac{q_{ij}}{1 - d_{ij}^2} \mathbf{x}_{ij}\mathbf{x}_{ij}^{\text{T}} \right). \quad (4.21)$$

This method can be applied in same way as classic NCA: learn a metric \mathbf{A} that maximizes the objective function $f_{\text{CS}}(\mathbf{A})$. Since the function is differentiable, any gradient based method is suitable for optimization and can be used on equation (4.21).

There is one concern with the compact support version of NCA. There are situations when a point \mathbf{x}_i is placed outside the support of any other point in the data set. Intuitively, this means that the point \mathbf{x}_i is not selected by any point, hence it is not assigned any class label. Also this causes mathematical problems: as in subsection 3.3.3, the contributions p_{ij} will have an indeterminate value $\frac{0}{0}$. The advice from subsection 3.3.3 can be applied here as well. A more robust way of dealing with this is discussed in the next subsection.

4.5.1 NCA with compact support kernels and background distribution

We extend the previous model to handle cases where points fall outside the support of any other neighbours. The idea is to use for each class a background distribution that explains the unallocated points. The background distribution should have an infinite support and an obvious example is the normal distribution.

To introduce a background distribution in a principled manner, we return to the class conditional kernel density estimation (CC-KDE) formulation of NCA (section 3.2). First, we recast the compact support NCA in the probabilistic framework and consider each class as mixture of compact support distributions:

$$p(\mathbf{x}_i|c) = \frac{1}{N} \sum_{j \in c} k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j), \quad (4.22)$$

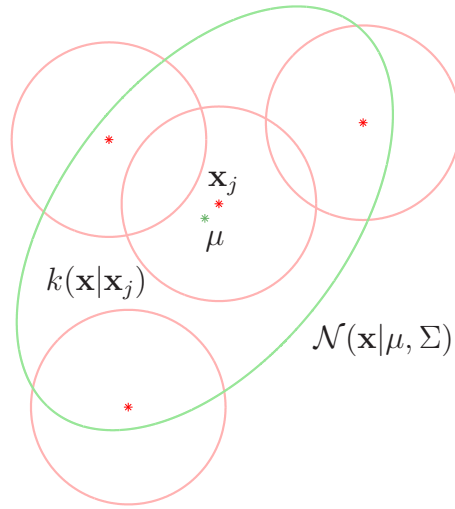


Figure 4.4: Neighbourhood component analysis with compact support kernels and background distribution. The main assumption is that each class is a mixture of compact support distributions $k(\mathbf{x}|\mathbf{x}_j)$ plus a normal background distribution $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

where $k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j) = k_{\text{CS}}(d_{ij})$ and is defined by equation (4.16). Because $k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j)$ denotes a distribution, it ought to integrate to 1. For $h = \frac{15}{16}$ and $a = 1$ the requirement is satisfied.

We further change the model and incorporate an additional distribution in the class-conditional probability $p(\mathbf{x}_i|c)$. From a generative perspective this can be interpreted as follows: a point \mathbf{x}_i is generated by either the compact support distribution from each point $k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j)$ or by a class-specific normal distribution $\mathcal{N}(\mathbf{x}_i|\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$. So, the distribution $p(\mathbf{x}_i|c)$ can be written as the sum of these components:

$$p(\mathbf{x}_i|c) = \beta \mathcal{N}(\mathbf{x}_i|\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c) + (1 - \beta) \frac{1}{N_c} \sum_{j \in c} k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j), \quad (4.23)$$

where $\beta \in [0, 1]$ is the mixing coefficient between the background distribution and the compact support model, $\boldsymbol{\mu}_c$ is the sample mean of the class c and $\boldsymbol{\Sigma}_c$ is the sample covariance of the class c . The constant β can be set to $\frac{1}{N_c+1}$. This will give equal weights to the background distribution and to each compact support distribution. It might be better to treat β as a parameter and fit it during training. We expect β to adapt to the data set: for example, β should increase for data sets with convex classes.

To finalize this method, we just need to plug equation (4.23) into the set of equations (3.11), (3.13) and (3.14). The only difficulty is the gradient computation. We now give the derivatives for each individual component:

- The gradient of the compact support distribution $k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j)$ with respect to \mathbf{A} is very similar to what is given in equation (4.20). The only difference is that in this case we have everything multiplied by the constant $h = \frac{15}{16}$.
- For the gradient of the background distribution it is useful to note that projecting the points $\{\mathbf{x}_i\}_{i=1}^N$ into a new space $\{\mathbf{A}\mathbf{x}_i\}_{i=1}^N$ will change the sample mean $\boldsymbol{\mu}_c$ to $\mathbf{A}\boldsymbol{\mu}_c$ and the sample covariance $\boldsymbol{\Sigma}_c$ to $\mathbf{A}\boldsymbol{\Sigma}_c\mathbf{A}^\text{T}$. Hence, we have:

$$\begin{aligned} \frac{\partial}{\partial \mathbf{A}} \mathcal{N}(\mathbf{A}\mathbf{x}_i|\mathbf{A}\boldsymbol{\mu}_c, \mathbf{A}\boldsymbol{\Sigma}_c\mathbf{A}^\text{T}) &= \mathcal{N}(\mathbf{A}\mathbf{x}_i|\mathbf{A}\boldsymbol{\mu}_c, \mathbf{A}\boldsymbol{\Sigma}_c\mathbf{A}^\text{T}) \\ &\times \{-(\mathbf{A}\boldsymbol{\Sigma}_c\mathbf{A}^\text{T})^{-1}\mathbf{A}\boldsymbol{\Sigma}_c + \mathbf{v}\mathbf{v}^\text{T}\mathbf{A}\boldsymbol{\Sigma}_c - \mathbf{v}(\mathbf{x}_i - \boldsymbol{\mu}_c)^\text{T}\}, \end{aligned} \quad (4.24)$$

where $\mathbf{v} = (\mathbf{A}\boldsymbol{\Sigma}_c\mathbf{A}^\text{T})^{-1}\mathbf{A}(\mathbf{x}_i - \boldsymbol{\mu}_c)$. The interested reader can find the derivation in appendix B.

- If we also consider β a parameter, we also need the derivative of the objective function with respect to β . This can be easily obtained, if we use the derivative of the class conditional distribution with respect to β :

$$\frac{\partial}{\partial \beta} p(\mathbf{x}_i|c) = \mathcal{N}(\mathbf{x}_i|\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c) - \frac{1}{N_c} \sum_{j \in c} k_{\text{CS}}(\mathbf{x}_i|\mathbf{x}_j). \quad (4.25)$$

We obtain a suitable metric \mathbf{A} in the standard way, by using a gradient-based optimizer.

Summary

In this chapter we have shown two main families of algorithms that can help reducing the computational cost of NCA. First, we have considered mini-batch algorithms. These achieve speed-ups by using sub-sets of the data instead of the whole training set. Second, we have presented algorithms that are based on approximations. In section 4.4 we saw how the NCA function can be directly approximated using k -d trees, while in section 4.5 we propose a different model to get the same effect even if we did not introduce approximation errors. Finally, in subsection 4.5.1 we make the compact support classifier well defined across the whole space by using a Gaussian distribution to represent each class. The next chapter will present the results obtained by evaluating these methods and a comparison between different algorithms.

Chapter 5

Evaluation

This chapter contains the evaluation of the methods proposed in the previous part, chapter 4. In section 5.1, we present the data sets and the methodology used for testing. We give details about the parameter choices and set baseline scores obtained by either classical NCA or simple linear projections, such as PCA, LDA or RCA. Results for each individual method are presented in sections 5.3 and 5.4. A comparison of the methods is shown in subsection 5.3.4 using accuracy versus time plots.

We should mention that we did not include all the results in this chapter to prevent cluttering. Further experimentations can be found in appendix C.

5.1 Evaluation setup

The evaluation was done in terms of two metrics: accuracy and speed. An additional and more subjective criterion is to judge a method by visualizing low dimensional representations of various data sets. We provided 2D plots of the projected data where suitable.

The data sets selected for testing are listed in table 5.1. We note that the used data vary both in number of samples N and in dimensionality D . Even if we concentrate on large amounts of data, we need small data sets to assess the performance of the new models. The methods' speeds were tested on the large data sets (`usps`, `magic` and `mnist`). However, the diversity in size and complexity made it difficult to find a single optimal selection of parameters. We are aware that there is “no free lunch” in accurately solving widely different problems with a fixed model. When concentrating on a single task it is often

Data set name	Abbreviation	N	D	C
Balance scale	balance	625	4	3
Ecoli	ecoli	336	7	8
Glass identification	glass	214	9	6
Ionosphere	ionosphere	351	33	2
Iris	iris	150	4	3
Landsat satellite	landsat	6435	36	6
MAGIC Gamma telescope	magic	19020	10	2
MNIST digits	mnist	70000	784	10
Pima Indians diabetes	pima	768	8	2
Image segmentation	segment	2310	18	7
Blood transfusion	transfusion	748	4	2
USPS digits	usps	11000	256	10
Wine	wine	178	13	3
Yeast	yeast	1484	8	10

Table 5.1: This table presents the characteristics of the data sets used: number of samples N , dimensionality of the data D and number of classes C . The two digits data sets `mnist` and `usps` were downloaded from the following URL <http://cs.nyu.edu/~roweis/data.html>. All the others data sets are available in the UCI repository <http://archive.ics.uci.edu/ml/datasets.html>.

advised to include prior knowledge into the model.

For evaluation we used 70% of each data set for training and the remaining 30% was kept for testing. We report standard errors for the mean accuracy averaged over different splits. We made exceptions for two data sets: `landsat` and `mnist`. These are already split into standard training and testing sets.

The methods we test are: sub-sampling (SS; section 4.1), mini-batches (MB; section 4.2), stochastic learning (SL; section 4.3). For stochastic learning we included two approximated method: the fast kernel density estimation idea (SL-KDE; section 4.4) and compact support version of NCA (SL-CS; section 4.5). Each method has particular parameters that we discuss in its corresponding subsection.

There are some common choices for all the methods, related to the NCA implementation. For convenience, we review these choices here. We experimented

with three optimization methods: gradient ascent with the “bold driver” heuristic, conjugate gradients and variants of stochastic gradient ascent with “early stopping”. For initialization we used the techniques described in subsection 3.3.2: random initialization, principal component analysis (PCA; Pearson, 1901), linear discriminant analysis (LDA; Fisher, 1936) or relevant component analysis (RCA; Bar-Hillel et al. (2003)). At test time, we did classification using 1-NN or using an NCA based function as described in subsection 3.3.5. However, we usually present scores using both classification rules.

The experiments were carried out in MATLAB and most of the implementations are the authors’ own work. There are some exceptions however. For conjugate gradients (CG) we used the function `minimize.m`, written by Rasmussen (2010).¹ The RCA implementation was provided by Noam Shental on his webpage.² Also we used functions from Iain Murray’s MATLAB toolbox.³ Finally, we inspected previous implementations of NCA,⁴ even if we did not explicitly make use of them.

5.2 Baseline

We started by implementing the standard NCA algorithm (appendix A.1). This is the main baseline against which we compare new models. For our first series of experiments, we tried to replicate the work in the original article (Goldberger et al., 2004). We encountered some difficulties since no information about their implementation was provided in the paper. Our results are presented in table 5.2 (scores are averaged over 40 runs). We randomly initialized the matrix \mathbf{A} and optimized it using the conjugate gradients (CG) method. This was the easiest thing to do since no parameter tuning is need for CG. We note that the results are similar to those of Goldberger et al. (2004); for the ionosphere data set we obtained slightly worse results.

However, in order to achieve a robust implementation of NCA we had to carry

¹This is available for download at <http://www.gaussianprocess.org/gpml/code/matlab/util/minimize.m>.

²The code was downloaded from <http://www.openu.ac.il/home/shental/>.

³Iain Murray’s toolbox is available at <http://homepages.inf.ed.ac.uk/imurray2/code/imurray-matlab/>.

⁴Implementations of NCA are provided by Laurens van der Maaten in his MATLAB Toolbox for Dimensionality Reduction http://homepage.tudelft.nl/19j49/Matlab_Toolbox_for_Dimensionality_Reduction.html and by Charless C. Fowlkes on his website <http://www.ics.uci.edu/~fowlkes/software/nca/>.

Data set	d	Train score	Test scores		Baseline
		$f(\mathbf{A})$	1-NN	NCA	Eucl.
balance	2	92.86 ± 0.47	90.78 ± 0.53	90.61 ± 0.55	
	$D = 4$	95.36 ± 0.38	93.40 ± 0.47	93.04 ± 0.51	76.18
ionosphere	2	98.31 ± 0.14	79.86 ± 0.75	79.74 ± 0.78	
	$D = 33$	72.07 ± 0.71	86.22 ± 0.64	72.87 ± 0.71	85.38
iris	2	99.38 ± 0.11	94.94 ± 0.39	94.72 ± 0.39	
	$D = 4$	99.48 ± 0.10	95.10 ± 0.44	95.15 ± 0.44	95.53
wine	2	99.15 ± 0.14	92.4 ± 1.0	92.4 ± 1.0	
	$D = 13$	98.95 ± 0.15	95.36 ± 0.51	95.36 ± 0.51	74.53

Table 5.2: Accuracy of standard NCA on four small data sets. Scores are averaged over 40 runs. The second column presents the dimensionality d the data set is reduced to. The last column shows the leave one out cross validation performance on the data set using Euclidean metric.

Data set	PCA	LDA	RCA
usps	73.47 ± 0.13	87.44 ± 0.12	87.42 ± 0.13
magic	77.097 ± 0.080	76.17 ± 0.29	77.574 ± 0.078
mnist	70.13	9.96	79.11

Table 5.3: Accuracy of three linear transformation techniques applied on the large data sets. We used 1-NN for classification. Scores are averaged over 20 runs, except for `mnist` data set. We reduced each data set dimensionality to $d = 5$.

out additional experiments. The lessons learnt were summarized in section 3.3. We will further see the influence of the implementation tricks in the next part (section 5.3). Results were shown in section 3.3 and are also attached at the end of the thesis, appendix C:

- Tables C.1 and C.2 compare two of the optimization methods: conjugate gradients and gradient ascent with the “bold driver” heuristic. We observe that the two methods give close scores on most of data sets. However, gradient ascent takes longer until it reaches convergence.
- Figures C.1, C.2 and C.3 illustrate the initialization effect on three data sets: `iris`, `balance` and `ecoli`. RCA seems to be the best option for initialization and we used it in most of our comparisons. However, random projection can also be sometimes surprisingly good as we see in figure C.1(a).

We could not apply NCA on large data sets (`usps`, `magic`, and `mnist`) since it would have taken far too long. We decided to use as a baseline the linear transformations that we also use for initialization (PCA, LDA and RCA). The results are averaged over different splits of training and testing set and they are listed in table 5.3. For the `mnist` data set we have scores of classic NCA (Singh-Miller, 2010). For $d = 5$ the accuracy obtained was 91.1% using the k NN classifier and 90.9% using the NCA classifier. For the classification procedure Singh-Miller learnt the optimal value for k using cross-validation.

5.3 Mini-batches methods

The mini-batch methods were compared on the larger data sets: `usps`, `magic`, and `mnist`. We chose to reduce the dimensionality to $d = 5$. This decision was partly motivated by the fact that NCA is very effective for reducing the data dimensionality to small values, somewhere around 5 to 15. A more principled approach would have been to develop a model choosing algorithm: start with a projection to $d = 2$ dimensions and then increase the number of dimensions until the score stops improving. This idea is similar to the “early stopping” procedure, and is illustrated for the `landsat` data set in figure 5.1.




Figure 5.1: Evolution of test accuracy as dimensionality increases on `landsat` data set. We see that NCA operates almost as well in low dimensions ($d = 6, \dots, 10$) as in high dimensions ($d > 25$). This approach can be used for selecting a suitable dimension to which we project the data.

Data set	Train	1-NN	NCA
<code>usps</code>	99.112 ± 0.099	89.30 ± 0.18	89.41 ± 0.16
<code>magic</code>	82.71 ± 0.41	79.25 ± 0.45	80.22 ± 0.72
<code>mnist</code>	96.867	82.130	82.470

Table 5.4: Accuracy scores for the sub-sampling method on the larger data sets. We used RCA for initialization and CGs for optimization. We used a subset of $n = 3000$ data points for training and the whole data set for testing.

5.3.1 Sub-sampling

For sub-sampling we trained NCA on a subset of $n = 3000$ samples of the original data. We used conjugate gradients for optimization and the RCA linear transformation for initialization. As previously mentioned in section 4.1, the sub-sampled data has a thinner distribution than the original data which helps the method to obtain good scores at training. But the test performance is hindered because we do not use the true distribution. This is especially evident for the digits data `usps` and `mnist` (table 5.4).

Data set	Train	1-NN	NCA
usps	92.00 ± 0.43	91.26 ± 0.17	92.37 ± 0.17
magic	80.04 ± 0.48	79.14 ± 0.52	79.8 ± 1.1
mnist	87.47 ± 0.49	87.52 ± 0.36	89.88 ± 0.25

Table 5.5: Accuracy scores for mini-batch method on the larger data sets. We used RCA for initialization and the mini-batches were clustered in the low-dimensional space using RPC. The size of a mini-batch was of maximum $n = 2000$ data points.

5.3.2 Mini-batches

We trained NCA using the gradient ascent variant with clustered mini-batches (section 4.2). For the learning rate, we used an update rule of the form $\frac{\eta}{t+t_0}$. We fixed $\eta = 1$ and tuned the other free parameter t_0 using cross validation across an exponential scale from 0.1 to 1000. After that, a finer tuning was done on a linear scale around the best value of t_0 . We used 5% of the training set for cross validation to monitor the accuracy score at each iteration. If the performance on the cross validation set does not increase for 25 iterations, we stopped the learning process and returned to the previously best parameter.

To get significant gradients we used large mini-batches $n = 2000$. The points in a batch are selected via the recursive projection clustering (RPC; Chalupka, 2011) algorithm. The clustering was done in the low dimensional space, after projecting the points with the current linear transformation matrix \mathbf{A} . The results obtained by the mini-batch method can be found in table 5.5. We note similar scores on `magic` and better results on the other two data sets.

5.3.3 Stochastic learning

This method was trained using the variant of stochastic gradient ascent presented in section 4.3. We used the same parameters as in the previous section. We considered $n = 50$ neighbours to look at for each iteration and computed their contributions with respect to the whole data set.

Besides the results on the large data sets (table 5.7), we also present the performance of this method when used for small data sets (table 5.6). We note a considerable improvement on the `magic` data set compared to the previous two methods. For small data sets, we observe similar results as the baseline NCA. The

Data set	d	Train score		Test scores	
		$f(A)$	1-NN	NCA	
balance	2	88.35 ± 0.83	87.37 ± 0.49	90.45 ± 0.38	
	$D = 4$	94.70 ± 0.87	95.32 ± 0.34	96.14 ± 0.29	
ionosphere	2	89.0 ± 1.7	85.71 ± 0.94	87.08 ± 0.95	
	$D = 33$	92.6 ± 1.5	84.72 ± 0.57	84.34 ± 0.59	
iris	2	96.41 ± 0.94	96.33 ± 0.57	97.00 ± 0.46	
	$D = 4$	97.5 ± 1.3	95.67 ± 0.71	96.11 ± 0.66	
wine	2	98.80 ± 0.70	97.22 ± 0.65	97.50 ± 0.49	
	$D = 13$	99.25 ± 0.62	96.85 ± 0.41	96.85 ± 0.41	

Table 5.6: Accuracy scores for stochastic learning method on the small data sets. We used RCA for initialization. The scores are averaged after 20 iterations.

Data set	Train	1-NN	NCA
usps	90.23 ± 0.50	90.68 ± 0.22	92.64 ± 0.17
magic	78.39 ± 0.25	79.76 ± 0.13	84.49 ± 0.12
mnist	85.97 ± 0.37	86.07 ± 0.43	89.35 ± 0.39

Table 5.7: Accuracy scores for stochastic learning method on the larger data sets. We used RCA for initialization. At each iteration we considered $n = 50$ data points.

most important remark is that the classification done using the NCA objective function is better than 1-NN classification. This observation also applies for the large data sets results. More results for this method are in the appendix, table C.3 and C.4.

5.3.4 Comparison

To easily compare the three methods presented, we provide time-accuracy plots (figure 5.2). We did not record the time taken for smaller data sets, since classical NCA was usually fast enough. We also included in the plots the method based on the compact support kernels (we discuss its individual scores in subsection 5.4.2).

In terms of accuracy, all of the NCA variants improve the accuracy over PCA,

LDA or RCA (table 5.3). Sub-sampling gives comparable results to the other two methods only on the `magic` data set. This data set has only two classes and a sub-sampled data set does not remove too much information about the decision boundaries. Mini-batches and stochastic learning are similar in accuracy. The proposed methods obtain a slightly worse score than classic NCA on `mnist`. The performance reported by Singh-Miller (2010) was around 91%, while the mean accuracy for these two methods is about 90%.

The differences in time varied strongly. We used a logarithmic scale on the time to better discriminate the plotted scores. The time spent varies even for the same method because the number of iterations until convergence depends on random parameters. We notice a certain trade-off between time and accuracy, especially for the digits data sets.

We also show low dimensional representations of the data sets for $d = 2$ for `usps` (figure 5.3) and `magic` (figure 5.4).


5.4 Approximate computations

The following methods can be applied on classic NCA, but we tested them in conjunction with the stochastic learning procedure to further boost the speed.


For approximate computations we also tried a more simplistic approach, similar to that of Weinberger and Tesauro (2007): we selected only the first 100 neighbours for each point and discarded those points whose contribution p_{ij} is less than $\exp(-30)$. This simple approach gave surprisingly good scores, although we do not present them here.

5.4.1 NCA with k -d trees


For the class-conditional kernel density estimation idea we used the algorithm described in section 4.4. Our k -d tree implementation was written in `MATLAB` and, for this reason, it was pretty slow. The code was slower than the simple SL version, because we could not vectorize the recursive computation of the objective function and its gradient. We experimented with kernel density code written in `C` and `mex-ed` in `MATLAB` and we think that such an approach can improve the speed. Doing simple kernel density estimation with a `C` written code proved to be even faster than doing it in `MATLAB`. The short time did not permit us to



(a) usps



(b) magic



(c) mnist

Figure 5.2: Time vs. accuracy plots on larger data sets for four of the proposed methods: sub-sampling (SS), mini-batches (MB), stochastic learning (SL) and stochastic learning with compact support NCA (SL-CS). For SS we plotted the 1-NN score, while for the other three the points indicate the NCA score.

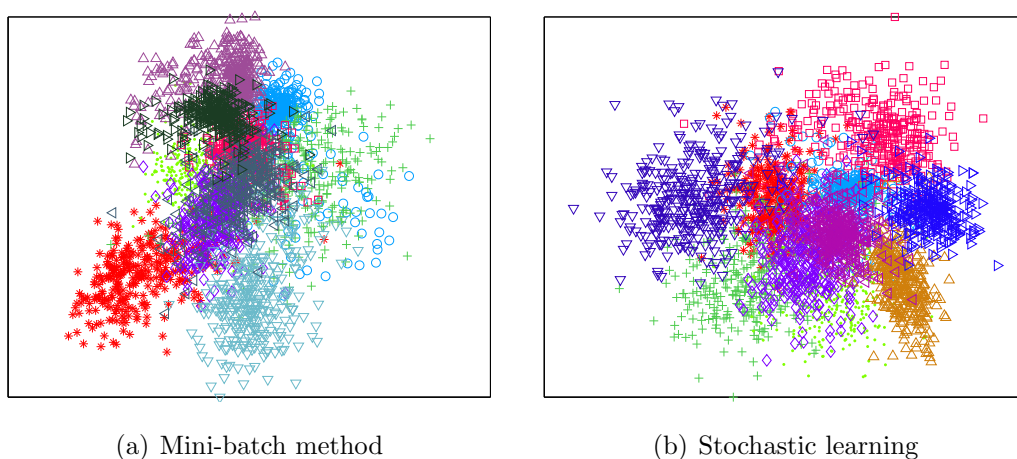


Figure 5.3: Two dimensional projections of `usps` data set using two variants of NCA learning. The linear transformation was learnt on a training set, and here is plotted the projection of a testing set.

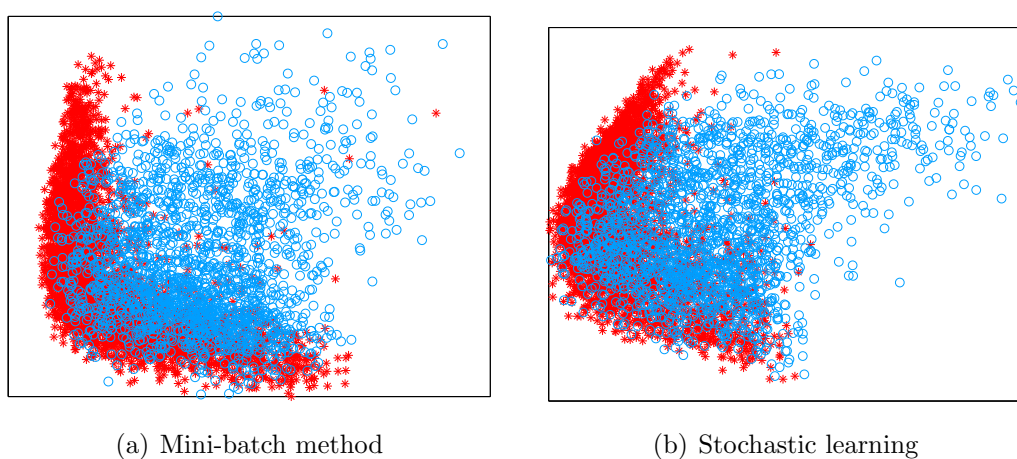


Figure 5.4: Two dimensional projections of `magic` data set using two variants of NCA learning. The linear transformation was learnt on a training set, and here is plotted the projection of a testing set.

ϵ_{\max}	Train	1-NN	NCA	Visited points
0	87.05 ± 0.66	86.30 ± 0.32	87.87 ± 0.24	80.0 ± 5.9
10^{-50}	87.50 ± 0.31	86.26 ± 0.24	87.88 ± 0.20	55.1 ± 5.2
10^{-20}	85.49 ± 0.89	86.32 ± 0.28	87.87 ± 0.26	45.7 ± 4.1
10^{-5}	87.29 ± 0.49	85.95 ± 0.19	87.63 ± 0.29	22.8 ± 2.7
0.01	87.34 ± 0.76	86.14 ± 0.29	87.90 ± 0.24	22.1 ± 4.2
0.1	86.70 ± 0.39	86.12 ± 0.29	88.09 ± 0.19	20.0 ± 2.0

Table 5.8: Stochastic learning for the approximated version of NCA that uses k -d trees. The data set used is `landsat`. ϵ_{\max} denotes the maximum error that we accept while approximating the density for a point given a class $p(\mathbf{x}|c)$. ‘Visited points’ indicates the fraction of points that are used for computing the function and the gradient.

Data set	Train	1-NN	NCA
<code>usps</code>	89.68 ± 0.51	90.57 ± 0.20	92.67 ± 0.15
<code>magic</code>	78.09 ± 0.37	79.68 ± 0.30	84.50 ± 0.21
<code>mnist</code>	84.4 ± 1.4	85.5 ± 1.0	88.98 ± 0.75

Table 5.9: Stochastic learning for the approximated version of NCA that uses k -d trees. For these experiments we set $\epsilon_{\max} = 0.1$.

finalize this approach. We demonstrate that SL with k -d trees achieves good results even if we discard a large number of the points. Table 5.9 shows results for this method on the large data sets. We also present how the accuracy and the average number of prunings depend on the maximum error imposed (table 5.8).

5.4.2 NCA with compact support kernels

The compact support version of NCA is easy to implement and it is very fast as we already saw in section 5.3.4. Usually only a small fraction of the points is inspected. In figure 5.5, we see how this fraction evolves with the learning process. In the first iterations there might be a larger number of points, but soon this drops off and stabilizes. It might be tempting to start with all the points very close together to ensure that no point is outside the compact support of any


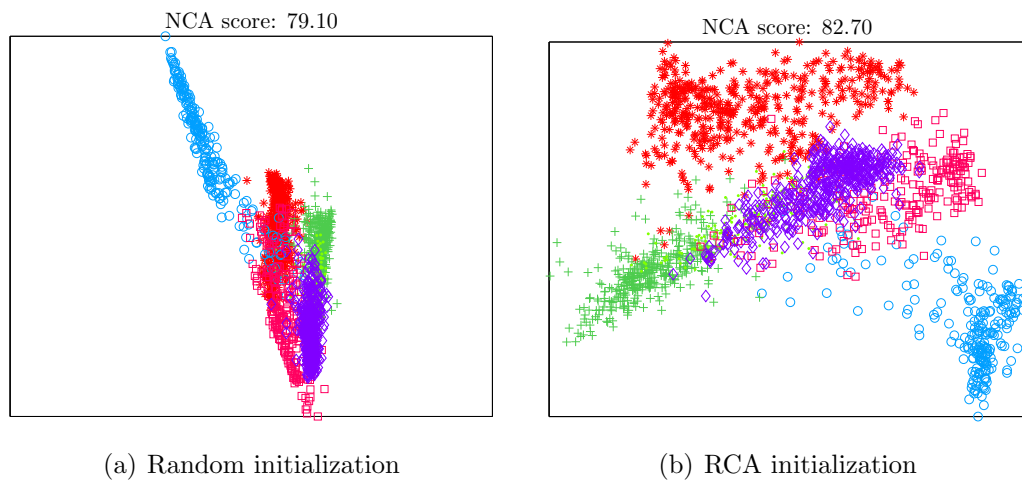


Figure 5.5: This figure illustrates how the fraction of the points inspected varies during the learning procedure. When we use random initialization, there are inspected only 0.1% of the points. If RCA is used to initialize the learning algorithm a fraction of about 10% is used. The data set used in this experiment was `landsat`.



(a) Random initialization

(b) RCA initialization

Figure 5.6: Illustration of final two-dimensional projections for the compact support NCA using stochastic learning. There were used two different initialization: random selection of parameters and RCA initialization. The data set used is `landsat`

Data set	Train	1-NN	NCA
usps	85.16 ± 0.33	87.02 ± 0.44	89.85 ± 0.30
magic	76.92 ± 0.56	79.09 ± 0.42	82.92 ± 0.36
mnist	80.55 ± 0.28	81.96 ± 0.26	86.36 ± 0.18

Table 5.10: Accuracy scores for the compact support NCA method trained using stochastic learning on large data sets.

other point. This approach has however two drawbacks. Since all the points are in the support of all the other points it means the first iteration will not present any gain in speed. In fact, the first iteration can easily be as expensive as the whole learning process. A second issue is that the gradients will be very large and might “throw away” points out of the existing support. An initialization with RCA provided a more stable evolution. We also see that the final projection looks better in the second case (figure 5.6).

Results that demonstrate performance in terms of accuracy on small data sets are attached in appendix, tables C.5 and C.6. The comparison has as baselines the classical NCA and the extended version of NCA with compact support kernels and background distribution (NCA CS BACK). The results on the large data set are available in table 5.10.

5.5 Recommendations

When dealing with a large data set, we suggest first trying the stochastic learning for the compact support NCA (NCA CS SL). This method has the advantage of being easy to implement and very fast even for large data sets. However, the speed does come at a cost. We note a slight decrease in accuracy for most experimentations. Nonetheless, applying the NCA CS SL first gives us an idea of how well NCA can perform on a given data set. If one is pleased with the score, one can also try the NCA optimized using stochastic learning (NCA SL) or the more sophisticated approximate version that uses k -d trees. Usually, these last two methods get an improvement of a couple of percentages in accuracy over NCA CS SL. However all the methods offer better scores than the simple projections such as LDA, PCA and RCA.

Chapter 6

Conclusions

This thesis presented our approach to fast low-rank metric learning. The need for a low rank metric was motivated in the context of k NN (chapter 1 and section 2.1), but we argue that learning a metric is useful whenever our algorithm relies on dissimilarities. Our efforts were directed towards the already established method, neighbourhood component analysis (NCA; section 3.1). We introduced a family of algorithms to mitigate NCA’s main drawback — the computational cost. In our attempt of speeding up NCA, we encountered other interesting theoretical and practical challenges. The answers to these issues represent an important part of the thesis (sections 3.2 and 3.3). We summarize here our main contributions and present the conclusions.

Section 3.2 The class-conditional kernel density estimation (CC-KDE) interpretation offers flexibility to NCA and opens a new range of methods. We can easily change the model in a suitable way by redefining the conditional probability $p(\mathbf{x}|c)$. Other advantages are (i) the possibility of including prior class distribution $p(c)$ into the model and (ii) we can use the posterior probability $p(c|\mathbf{x})$ for classification or in other scenarios that arise in decision theory.

Section 3.3 The practical advice in this section is helpful for those who are interested in applying the classic version of NCA. The ideas are particularly useful because we are optimizing a non-convex function. We believe that using relevant component analysis (RCA; Bar-Hillel et al., 2003) for initialization and conjugate gradients for optimization works best on small and medium-sized data sets. This combination does not require tuning any

parameters so it can be readily applied on any data set. Among others, we highlight the dimensionality annealing technique; this look promising: it obtained good projections that seem to avoid local optima.

Section 4.2 We gave several ideas of adapting mini-batch (MB) techniques in a suitable way for NCA method. We tested a version based on recursive projection clustering (RPC, Chalupka, 2011). The size of each mini-batch still has to be quite large to get significant gradients (we used $n = 2000$). For large data set this already represents a significant improvement in terms of cost, but we question whether there is a more principled method of selecting n , the size of a cluster.

Section 4.3 We presented a stochastic learning (SL) algorithm that is theoretically motivated by stochastic optimization arguments. Empirical investigation demonstrated that this method achieves very close scores to the classic NCA. The SL method scales with the number of the points N , but it can be further accelerated using technique presented in sections 4.4 and 4.5. A further advantage of SL is that can be used for online learning.

Section 4.4 Using the CC-KDE framework for NCA and inspired from previous work on fast KDE (Deng and Moore, 1995; Gray and Moore, 2003), we proposed an algorithm that efficiently evaluates an approximation of the NCA objective function and its corresponding gradient. This method is fastest when combined with the SL method. Our experiments suggested that even if we accept a large maximum error $\epsilon_{\max} = 0.1$, we do not seem to lose in classification accuracy. We noticed a similar behaviour when we did brute pruning.

Section 4.5 Further alterations of NCA method are the compact support version (NCA CS) and the compact support and background distribution version (NCA CS BACK). NCA CS achieves considerable speed-ups because we do not have to compute the gradient terms for all the points. Also we observe that it convergences faster than the other methods. However its speed does come at a cost: the accuracy performance is slightly worse than for the rest of the methods. These results might suggest to try a different approach: use a heavy tailed distribution

6.1 Further directions

This project involved many ideas. Unfortunately, the time was limited and we could only look at a part of the whole picture and we left some paths open. We aim to continue some of the work in the near future.

We plan to further investigate the automatic differentiation procedure and see whether this is a better idea than symbolic differentiation. Automatic differentiation will help reducing the number of operations needed for the gradient computation, but it might imply a huge memory cost. This trade-off can make automatic differentiation impractical.

Implementing efficient k -d tree code for approximate CC-KDE algorithm should prove beneficial. It will allow our methods to run on even larger data sets making possible more comparisons between the methods.

We think that dimensionality annealing approach can be fruitful. However, this idea is still in its infant stages and more experimentation is needed to see whether it can be applied successfully on a variety of data sets.

Further refinements of our work are still possible. It would be interesting to try a mini-batch and stochastic gradient combined algorithm: cluster points and compute mini-batch mini-batch contributions using a dual tree recursion (Gray and Moore, 2003).

We believe that other extensions of NCA are possible. For example, a low-dimensional feature extractor could be built using a two step NCA algorithm: first use a diagonal metric \mathbf{S} to select the most relevant attributes, followed by a low-rank metric.

Appendix A

MATLAB code

A.1 NCA cost function

```
function [f, df] = nca_obj(A, X, c)
%NCA_OBJ Neighbourhood Component Analysis cost function.
%Returns function value and the first order derivatives.
%This is a function that should be minimized; it is
%minus the objective function in paper.
%
% [f, df] = nca_obj(A, X, c)
%
% Inputs:
%   A dxD - projection matrix (d <= D).
%   X DxN - data.
%   c 1xN - class labels.
%
% Outputs:
%   f 1x1 - function value.
%   df 1xD*d - derivative values.
%
% Dan Oneata, June 2011

[D N] = size(X);
A = reshape(A, [], D);

df = zeros(size(A,1), D);

AX = A*X;
dist_all = square_dist(AX, AX);
kern_all = exp(-dist_all);
kern_all(1:N+1:end) = 0;

% Compute distance to the neighbours:
row_dict = logical(eye(max(c)));
```

```

neigh_mask = row_dict(c,c);
kern_neigh = sum(kern_all.*neigh_mask, 1);

kern_sum = sum(kern_all, 1);
p = kern_neigh ./ kern_sum;
p = max(p, eps);

% Compute function value:
f = - sum(p);

if nargin > 1,
    K = bsxfun(@rdivide, kern_all, kern_sum);
    K = max(K, eps);
    K = K';
    for i=1:N,
        x_ik = bsxfun(@minus, X(:,i), X);
        Ax_ik = bsxfun(@minus, AX(:,i), AX);
        x_ij = x_ik(:,c==c(i));
        Ax_ij = Ax_ik(:,c==c(i));

        % Update gradient value:
        df = df + p(i) * bsxfun(@times, K(i,:), Ax_ik) * x_ik' ...
            - bsxfun(@times, K(i,c==c(i)), Ax_ij) * x_ij';
    end
    df = -2*df;
    df = df(:);
end

end

```

Notes:

1. `square_dist` computes the pairwise distances between two D -dimensional sets of points. The function was written by Iain Murray and it can be found in his MATLAB Toolbox, available at the following URL: http://homepages.inf.ed.ac.uk/imurray2/code/imurray-matlab/square_dist.m.
2. This implementation is suitable for rather small data sets whose pairwise distance matrix can fit in the RAM. If you are dealing with large data sets, it is better to iterate through the data set and compute the distances successively: from a point to the rest of the data set.

Appendix B

Mathematical derivations

B.1 The gradient of $\mathcal{N}(\mathbf{Ax}|\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)$ with respect to the linear transformation \mathbf{A}

This sections shows the main steps of deriving the result presented in equation (4.24).

$$\begin{aligned} \frac{\partial}{\partial \mathbf{A}} \mathcal{N}(\mathbf{Ax}|\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top) &= \frac{\partial}{\partial \mathbf{A}} \frac{1}{(2\pi)^{D/2} \det(\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{1/2}} \\ &\times \exp \left\{ -\frac{1}{2} (\mathbf{Ax} - \mathbf{A}\boldsymbol{\mu})^\top (\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{-1} (\mathbf{Ax} - \mathbf{A}\boldsymbol{\mu}) \right\} \end{aligned} \quad (\text{B.1})$$

Using the product rule, one can easily see that the computation is reduced to the following two derivatives:

$$\begin{aligned} \mathbf{P} &= \frac{\partial}{\partial \mathbf{A}} \frac{1}{(2\pi)^{D/2} \det(\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{1/2}} \quad \text{and} \\ \mathbf{Q} &= \frac{\partial}{\partial \mathbf{A}} \left\{ -\frac{1}{2} (\mathbf{Ax} - \mathbf{A}\boldsymbol{\mu})^\top (\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{-1} (\mathbf{Ax} - \mathbf{A}\boldsymbol{\mu}) \right\}. \end{aligned}$$

Starting with \mathbf{P} , we have:

$$\mathbf{P} = -\frac{1}{2} (2\pi)^{-D/2} \det(\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{-3/2} \times 2 \det(\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top) (\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top)^{-1} \mathbf{A}\boldsymbol{\Sigma}, \quad (\text{B.2})$$

where we have used equation (47) from (Petersen and Pedersen, 2008).

For the second derivative, let $\mathbf{Y} = \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top$ and $\mathbf{b} = \mathbf{x} - \boldsymbol{\mu}$. From equation (53) from (Petersen and Pedersen, 2008), we have that:

$$\frac{\partial \mathbf{Y}^{-1}}{\partial a_{ij}} = -\mathbf{Y}^{-1} \frac{\partial \mathbf{Y}}{\partial a_{ij}} \mathbf{Y}^{-1}, \quad (\text{B.3})$$

where a_{ij} is the element on the i -th row and j -th column of the matrix \mathbf{A} .

Besides that, we also know the following (see Petersen and Pedersen, 2008, equation (72)):

$$\frac{\partial(\mathbf{A}\Sigma\mathbf{A}^T)}{\partial a_{ij}} = \mathbf{A}\Sigma\mathbf{J}^{ij} + \mathbf{J}^{ij}\Sigma\mathbf{A}^T, \quad (\text{B.4})$$

with \mathbf{J}^{ij} denoting the single-entry matrix — the matrix that has the value 1 on the position (i, j) and 0 elsewhere.

By combining equations (B.3) and (B.4), we obtain an entry of the matrix \mathbf{Q} :

$$\begin{aligned} q_{ij} &= \frac{\partial}{\partial a_{ij}} \left\{ -\frac{1}{2} \mathbf{b}^T \mathbf{A}^T (\mathbf{A}\Sigma\mathbf{A}^T)^{-1} \mathbf{A} \mathbf{b} \right\} \\ &= \frac{1}{2} \mathbf{b}^T \mathbf{A}^T \mathbf{Y}^{-1} (\mathbf{A}\Sigma\mathbf{J}^{ij} + \mathbf{J}^{ij}\Sigma\mathbf{A}^T) \mathbf{Y}^{-1} \mathbf{A} \mathbf{b} \\ &= \mathbf{v}^T \mathbf{A}\Sigma\mathbf{J}^{ij} \mathbf{v}, \end{aligned} \quad (\text{B.5})$$

where $\mathbf{v} = \mathbf{Y}^{-1} \mathbf{A} \mathbf{b}$ and we have used the fact that Σ is a covariance matrix and hence symmetric.

By using equation (431) from (Petersen and Pedersen, 2008), we obtain that:

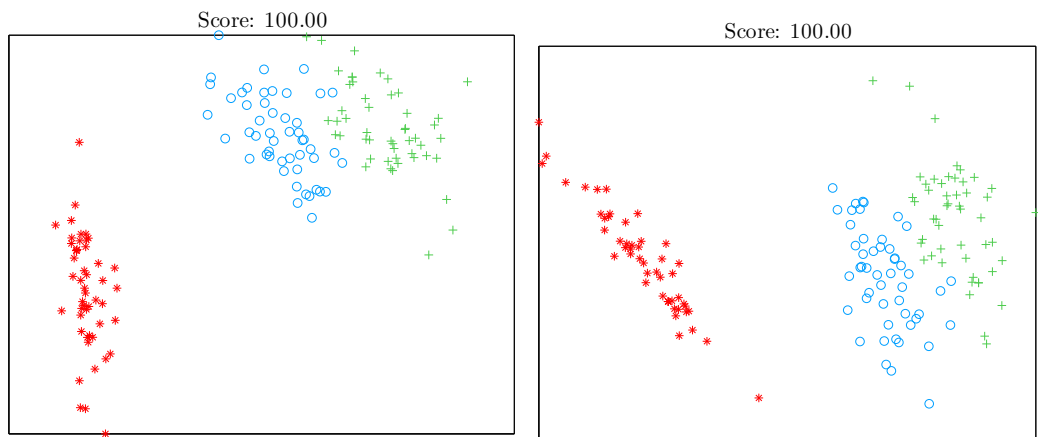
$$\mathbf{Q} = \Sigma \mathbf{A}^T \mathbf{v} \mathbf{v}^T. \quad (\text{B.6})$$

Finally, by bringing together the previous results (equations (B.2) and (B.6)) and replacing \mathbf{Y} with $\mathbf{A}\Sigma\mathbf{A}^T$, we arrive at the following identity:


$$\begin{aligned} \frac{\partial}{\partial \mathbf{A}} \mathcal{N}(\mathbf{A} \mathbf{x} | \mathbf{A} \boldsymbol{\mu}, \mathbf{A}\Sigma\mathbf{A}^T) &= \mathcal{N}(\mathbf{A} \mathbf{x} | \mathbf{A} \boldsymbol{\mu}, \mathbf{A}\Sigma\mathbf{A}^T) \\ &\times \{ -(\mathbf{A}\Sigma\mathbf{A}^T)^{-1} \mathbf{A}\Sigma + \mathbf{v} \mathbf{v}^T \mathbf{A}\Sigma - \mathbf{v}(\mathbf{x} - \boldsymbol{\mu})^T \}. \end{aligned} \quad (\text{B.7})$$

Appendix C

Further results

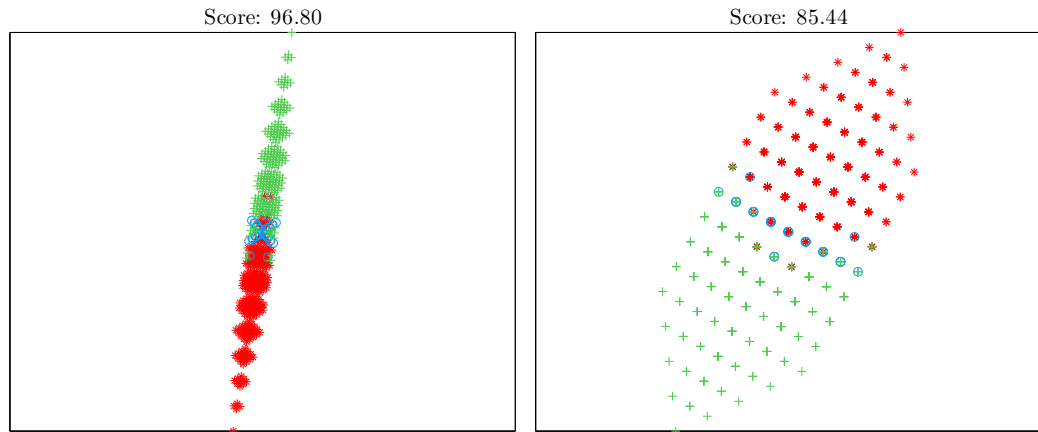


(a) NCA projection after random initialization (b) NCA projection after PCA initialization




(c) NCA projection after LDA initialization (d) NCA projection after RCA initialization

Figure C.1: Initialization experiments on `iris` data set.




(a) NCA projection after random initialization (b) NCA projection after PCA initialization




(c) NCA projection after LDA initialization (d) NCA projection after RCA initialization

Figure C.2: Initialization experiments on balance data set.



(a) NCA projection after random initialization (b) NCA projection after PCA initialization



(c) NCA projection after LDA initialization (d) NCA projection after RCA initialization

Figure C.3: Initialization experiments on `ecoli` data set.

Data set	d	Conjugate gradients		Bold driver	
		1-NN	NCA	1-NN	NCA
balance	2	92.74 ± 0.57	92.29 ± 0.56	92.05 ± 0.93	92.45 ± 0.81
	3	94.81 ± 0.52	94.87 ± 0.53	92.79 ± 0.95	92.87 ± 0.97
	4	95.16 ± 0.39	95.32 ± 0.33	94.18 ± 0.76	94.36 ± 0.69
glass	2	53.6 ± 1.4	54.2 ± 1.5	58.7 ± 1.5	62.5 ± 1.4
	3	61.8 ± 1.1	60.9 ± 1.2	66.3 ± 1.1	67.1 ± 1.1
	4	62.7 ± 1.4	62.2 ± 1.5	65.8 ± 1.0	66.6 ± 1.3
	5	64.5 ± 1.2	63.46 ± 0.97	66.69 ± 0.93	67.00 ± 0.87
	9	66.5 ± 1.8	65.5 ± 1.8	66.0 ± 1.7	65.6 ± 1.7
ionosphere	2	86.23 ± 0.80	86.18 ± 0.75	82.78 ± 0.98	82.78 ± 0.91
	3	87.69 ± 0.62	87.64 ± 0.62	86.04 ± 0.83	86.08 ± 0.85
	4	87.64 ± 0.55	87.64 ± 0.57	86.51 ± 0.73	86.51 ± 0.73
	5	89.39 ± 0.70	89.39 ± 0.70	87.88 ± 0.64	87.88 ± 0.64
	33	87.97 ± 0.68	87.97 ± 0.68	88.44 ± 0.59	88.44 ± 0.59

Table C.1: Comparison in terms of accuracy between two possible optimization methods for NCA: the conjugate gradient method and the gradient ascent with the “bold driver” heuristic. (Part 2 of 2)

Data set	d	Conjugate gradient		Bold-driver	
		1-NN	NCA	1-NN	NCA
iris	2	94.44 ± 0.51	94.11 ± 0.48	93.67 ± 0.82	93.78 ± 0.84
	3	95.44 ± 0.73	95.00 ± 0.75	95.22 ± 0.72	95.33 ± 0.74
	4	96.56 ± 0.46	96.44 ± 0.46	95.44 ± 0.76	95.56 ± 0.77
wine	2	96.67 ± 0.38	96.57 ± 0.38	96.30 ± 0.54	96.30 ± 0.54
	3	96.39 ± 0.48	96.39 ± 0.48	96.57 ± 0.59	96.57 ± 0.59
	4	96.20 ± 0.40	96.48 ± 0.41	97.13 ± 0.52	97.13 ± 0.52
	5	96.57 ± 0.60	96.57 ± 0.60	96.76 ± 0.55	96.76 ± 0.55
	13	96.20 ± 0.48	96.20 ± 0.48	97.22 ± 0.48	97.22 ± 0.48
yeast	2	42.37 ± 0.81	43.20 ± 0.84	44.65 ± 0.56	46.68 ± 0.57
	3	48.06 ± 0.67	48.30 ± 0.65	47.06 ± 0.53	47.95 ± 0.54
	4	49.06 ± 0.47	49.14 ± 0.47	49.48 ± 0.42	49.83 ± 0.38
	5	50.09 ± 0.39	50.09 ± 0.38	50.16 ± 0.61	50.41 ± 0.62
	8	50.66 ± 0.52	50.70 ± 0.52	50.19 ± 0.58	50.28 ± 0.60

Table C.2: Comparison in terms of accuracy between two possible optimization methods for NCA: conjugate gradients and gradient ascent with the “bold driver” heuristic. (Part 2 of 2)

Data set	d	Train	1-NN	NCA
balance	2	88.35 ± 0.83	87.37 ± 0.49	90.45 ± 0.38
	3	94.0 ± 1.0	94.44 ± 0.80	95.11 ± 0.56
	$D = 4$	94.70 ± 0.87	95.32 ± 0.34	96.14 ± 0.29
glass	2	55.0 ± 2.5	54.9 ± 1.7	59.0 ± 1.7
	3	65.4 ± 3.3	59.0 ± 1.2	62.5 ± 1.2
	4	69.2 ± 3.0	60.5 ± 1.2	63.77 ± 0.90
	5	68.2 ± 3.0	64.1 ± 1.3	65.7 ± 1.1
	$D = 9$	76.0 ± 2.3	68.0 ± 1.6	69.7 ± 1.6
ionosphere	2	89.0 ± 1.7	85.71 ± 0.94	87.08 ± 0.95
	3	91.0 ± 1.8	86.6 ± 1.1	86.8 ± 1.1
	4	89.6 ± 1.4	87.7 ± 1.1	87.74 ± 0.97
	5	91.4 ± 1.7	88.07 ± 0.73	88.35 ± 0.78
	$D = 33$	92.6 ± 1.5	84.72 ± 0.57	84.34 ± 0.59
iris	2	96.41 ± 0.94	96.33 ± 0.57	97.00 ± 0.46
	3	97.30 ± 0.82	95.67 ± 0.57	96.00 ± 0.62
	$D = 4$	97.5 ± 1.3	95.67 ± 0.71	96.11 ± 0.66
wine	2	98.80 ± 0.70	97.22 ± 0.65	97.50 ± 0.49
	3	98.17 ± 0.87	97.32 ± 0.58	97.69 ± 0.52
	4	98.51 ± 0.86	97.04 ± 0.53	97.32 ± 0.46
	5	99.22 ± 0.40	96.95 ± 0.45	97.13 ± 0.47
	$D = 13$	99.25 ± 0.62	96.85 ± 0.41	96.85 ± 0.41

Table C.3: Accuracy scores of NCA trained using stochastic learning on small and medium sized data sets. The scores are averaged over 20 repeats. RCA was used for initialization. (Part 1 of 2)

Data set	d	Train	1-NN	NCA
yeast	2	42.83 ± 0.93	45.85 ± 0.41	53.90 ± 0.47
	3	47.30 ± 0.85	46.49 ± 0.48	54.37 ± 0.40
	4	50.1 ± 1.2	48.69 ± 0.68	54.50 ± 0.63
	5	50.41 ± 0.82	50.19 ± 0.45	55.35 ± 0.49
	$D = 8$	52.09 ± 0.84	51.34 ± 0.47	55.44 ± 0.56
ecoli	2	74.1 ± 2.2	75.8 ± 1.0	81.98 ± 0.96
	3	81.2 ± 1.8	78.72 ± 0.77	81.49 ± 0.75
	4	79.9 ± 2.1	80.20 ± 0.73	83.12 ± 0.86
	5	80.9 ± 2.4	79.16 ± 0.77	83.62 ± 0.64
	$D = 7$	82.0 ± 2.1	80.89 ± 0.52	84.85 ± 0.48
pima	2	70.6 ± 1.1	69.09 ± 0.64	76.65 ± 0.44
	3	70.0 ± 1.3	68.57 ± 0.62	75.56 ± 0.82
	4	71.3 ± 1.2	69.41 ± 0.60	74.37 ± 0.57
	5	73.0 ± 1.2	69.03 ± 0.71	72.47 ± 0.81
	$D = 8$	76.2 ± 1.5	69.29 ± 0.88	70.48 ± 0.81
segment	2	84.62 ± 0.64	88.53 ± 0.39	88.80 ± 0.33
	3	92.27 ± 0.57	95.45 ± 0.31	94.46 ± 0.33
	4	94.24 ± 0.46	97.11 ± 0.18	96.09 ± 0.20
	5	94.98 ± 0.35	97.05 ± 0.16	96.64 ± 0.13
	$D = 18$	95.73 ± 0.57	97.01 ± 0.28	96.68 ± 0.30
transfusion	2	70.48 ± 0.99	73.60 ± 0.47	78.78 ± 0.42
	3	70.08 ± 0.92	73.49 ± 0.54	77.95 ± 0.39
	$D = 4$	66.0 ± 3.1	71.44 ± 0.56	75.4 ± 2.9

Table C.4: Accuracy scores of NCA trained using stochastic learning on small and medium sized data sets. The scores are averaged over 20 repeats. RCA was used for initialization. (Part 2 of 2)

Data set	d	NCA			NCA CS			NCA CS BACK		
		1-NN	NCA	1-NN	1-NN	NCA	1-NN	1-NN	NCA	1-NN
balance	2	92.74 ± 0.57	92.29 ± 0.56	87.90 ± 0.68	89.04 ± 0.52	89.04 ± 0.52	89.9 ± 1.3	89.9 ± 1.3	91.04 ± 0.94	
	3	94.81 ± 0.52	94.87 ± 0.53	89.60 ± 0.76	90.08 ± 0.60	90.08 ± 0.60	65.2 ± 3.8	65.2 ± 3.8	87.85 ± 0.53	
	4	95.16 ± 0.39	95.32 ± 0.33	90.82 ± 0.62	92.07 ± 0.54	92.07 ± 0.54	67.2 ± 3.4	67.2 ± 3.4	86.91 ± 0.30	
glass	2	53.6 ± 1.4	54.2 ± 1.5	55.8 ± 1.1	57.2 ± 1.7	57.2 ± 1.7	58.5 ± 1.2	58.5 ± 1.2	59.4 ± 1.5	
	3	61.8 ± 1.1	60.9 ± 1.2	63.2 ± 1.4	56.3 ± 1.8	56.3 ± 1.8	61.5 ± 1.2	61.5 ± 1.2	63.5 ± 1.1	
	4	62.7 ± 1.4	62.2 ± 1.5	62.0 ± 1.4	54.5 ± 1.6	54.5 ± 1.6	65.2 ± 1.2	65.2 ± 1.2	65.2 ± 1.0	
	5	64.5 ± 1.2	63.46 ± 0.97	64.9 ± 1.1	57.1 ± 1.3	57.1 ± 1.3	67.5 ± 1.2	67.5 ± 1.2	67.2 ± 1.3	
	9	66.5 ± 1.8	65.5 ± 1.8	68.9 ± 1.5	57.8 ± 1.5	57.8 ± 1.5	68.2 ± 1.6	68.2 ± 1.6	64.1 ± 1.3	
ionosphere	2	86.23 ± 0.80	86.18 ± 0.75	82.83 ± 0.74	85.33 ± 0.93	85.33 ± 0.93	85.90 ± 0.74	85.90 ± 0.74	84.20 ± 0.74	
	3	87.69 ± 0.62	87.64 ± 0.62	85.42 ± 0.62	83.44 ± 0.69	83.44 ± 0.69	88.11 ± 0.74	88.11 ± 0.74	84.91 ± 0.52	
	4	87.64 ± 0.55	87.64 ± 0.57	86.08 ± 0.90	83.6 ± 1.1	83.6 ± 1.1	88.73 ± 0.92	88.73 ± 0.92	83.7 ± 1.1	
	5	89.39 ± 0.70	89.39 ± 0.70	88.87 ± 0.61	81.46 ± 0.76	81.46 ± 0.76	86.65 ± 0.86	86.65 ± 0.86	82.08 ± 0.98	
	33	87.97 ± 0.68	87.97 ± 0.68	86.56 ± 0.92	66.09 ± 0.95	66.09 ± 0.95	87.26 ± 0.45	87.26 ± 0.45	84.76 ± 0.83	

Table C.5: Accuracy scores for NCA, the compact support NCA and the compact support NCA with background distribution (part 1 of 2).

Data set	d	NCA			NCA CS			NCA CS BACK		
		1-NN	NCA	1-NN	1-NN	NCA	1-NN	1-NN	NCA	1-NN
iris	2	94.44 ± 0.51	94.11 ± 0.48	95.22 ± 0.59	94.78 ± 0.72	95.22 ± 0.67	95.11 ± 0.60			
	3	95.44 ± 0.73	95.00 ± 0.75	95.33 ± 0.93	95.0 ± 1.0	94.89 ± 0.80	94.2 ± 1.1			
	4	96.56 ± 0.46	96.44 ± 0.46	95.56 ± 0.54	95.67 ± 0.53	95.2 ± 1.0	93.11 ± 0.95			
wine	2	96.67 ± 0.38	96.57 ± 0.38	97.04 ± 0.51	96.94 ± 0.44	97.22 ± 0.48	97.59 ± 0.46			
	3	96.39 ± 0.48	96.39 ± 0.48	96.48 ± 0.63	95.93 ± 0.65	97.22 ± 0.53	97.50 ± 0.40			
	4	96.20 ± 0.40	96.48 ± 0.41	98.06 ± 0.36	98.06 ± 0.28	98.24 ± 0.33	98.89 ± 0.30			
	5	96.57 ± 0.60	96.57 ± 0.60	96.20 ± 0.71	95.00 ± 0.81	97.78 ± 0.43	97.87 ± 0.40			
	13	96.20 ± 0.48	96.20 ± 0.48	95.09 ± 0.70	92.59 ± 0.92	94.63 ± 0.81	98.43 ± 0.30			
yeast	2	42.37 ± 0.81	43.20 ± 0.84	42.57 ± 0.63	48.31 ± 0.90	44.38 ± 0.66	48.24 ± 0.79			
	3	48.06 ± 0.67	48.30 ± 0.65	46.63 ± 0.72	51.6 ± 1.0	47.33 ± 0.61	50.67 ± 0.55			
	4	49.06 ± 0.47	49.14 ± 0.47	49.26 ± 0.51	53.06 ± 0.94	49.93 ± 0.36	53.08 ± 0.30			
	5	50.09 ± 0.39	50.09 ± 0.38	50.16 ± 0.59	52.96 ± 0.93	49.79 ± 0.62	54.54 ± 0.56			
	8	50.66 ± 0.52	50.70 ± 0.52	50.93 ± 0.44	54.89 ± 0.43	48.73 ± 0.65	56.93 ± 0.62			

Table C.6: Accuracy scores for NCA, the compact support NCA and the compact support NCA with background distribution (part 2 of 2).

Bibliography

- Alexander Gray, A. M. (2003). Rapid evaluation of multiple density models. In *Artificial Intelligence and Statistics*.
- Bar-Hillel, A., Hertz, T., Shental, N., and Weinshall, D. (2003). Learning distance functions using equivalence relations. In *Proceedings of the Twentieth International Conference on Machine Learning*, volume 20, page 11.
- Barber, D. (2011). *Bayesian Reasoning and Machine Learning*. Cambridge University Press. In press.
- Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18:509–517.
- Beyer, K., Goldstein, J., Ramakrishnan, R., and Shaft, U. (1999). When is “nearest neighbor” meaningful? *Database Theory ICDT 99*, pages 217–235.
- Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford University Press.
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- Butman, M. and Goldberger, J. (2008). Face recognition using classification-based linear projections. *EURASIP Journal on Advances in Signal Processing*, 2008:1–7.
- Chalupka, K. (2011). Empirical evaluation of Gaussian Process approximation algorithms. Master’s thesis, University of Edinburgh.
- Cover, T. and Hart, P. (1967). Nearest neighbor pattern classification. *Information Theory, IEEE Transactions on*, 13:21–27.
- Deng, K. and Moore, A. (1995). Multiresolution instance-based learning. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 1233–1239, San Francisco. Morgan Kaufmann.
- Dwinnell, W. (2006). Mahalanobis distance. <http://matlabdatamining.blogspot.com/2006/11/mahalanobis-distance.html>. (Date accessed: August 2011).

- Fisher, R. (1936). The use of multiple measurements in taxonomic problems. In *Annals of Eugenics*, volume 7, pages 179–188.
- Fix, E. and Hodges, J. (1951). Discriminatory analysis, nonparametric discrimination: Consistency properties. Technical Report 4, USAF School of Aviation Medicine, Randolph Field, Texas.
- Friedman, J. H., Bentley, J. L., and Finkel, R. A. (1977). An algorithm for finding best matches in logarithmic expected time. *ACM Trans. Math. Softw.*, 3:209–226.
- Goldberger, J., Roweis, S., Hinton, G., and Salakhutdinov, R. (2004). Neighbourhood components analysis. In *Advances in Neural Information Processing Systems*. MIT Press.
- Gonzalez, T. (1985). Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306.
- Gray, A. and Moore, A. (2001). ‘N-Body’ problems in statistical learning. *Advances in neural information processing systems*, pages 521–527.
- Gray, A. and Moore, A. (2003). Nonparametric density estimation: Toward computational tractability. In *SIAM International Conference on Data Mining*, volume 2003.
- Hastie, T. and Tibshirani, R. (1996). Discriminant adaptive nearest neighbor classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18:607–616.
- Hestenes, M. and Stiefel, E. (1952). Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49(6).
- Hinneburg, E., Aggarwal, C., Keim, D., and Hinneburg, A. (2000). What is the nearest neighbor in high dimensional spaces? In *Proceedings of the 26th International Conference on Very Large Data Bases*. Morgan Kaufmann Publishers Inc.
- Holte, R. (1993). Very simple classification rules perform well on most commonly used datasets. *Machine learning*, 11(1):63–90.
- Keller, P., Mannor, S., and Precup, D. (2006). Automatic basis function construction for approximate dynamic programming and reinforcement learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 449–456. ACM.
- Lang, D. (2009). *Astrometry.net: Automatic recognition and calibration of astronomical images*. PhD thesis, University of Toronto.
- LeCun, Y., Bottou, L., Orr, G., and Muller, K. (1998). Efficient backprop. In Orr, G. and K., M., editors, *Neural Networks: Tricks of the trade*. Springer.

- Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill, New York, NY, USA.
- Moore, A. (1991). A tutorial on kd-trees. Extract from PhD Thesis. Available from <http://www.cs.cmu.edu/~awm/papers.html>.
- Moore, A. and Lee, M. (1994). Efficient algorithms for minimizing cross validation error. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 190–198.
- Omohundro, S. (1989). Five balltree construction algorithms. Technical Report TR-89-063, International Computer Science Institute.
- Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine Series 6*, 2(11):559–572.
- Petersen, K. B. and Pedersen, M. S. (2008). The matrix cookbook. Available from <http://www2.imm.dtu.dk/pubdb/p.php?3274>. Version 2008/11/10.
- Rall, L. (1981). *Automatic differentiation: Techniques and applications*, volume 120. Springer-Verlag, Berlin.
- Rasmussen, C. E. (2010). GPML Matlab code. <http://www.gaussianprocess.org/gpml/code/>. (Date accessed: August 2011).
- Russell, S. J., Norvig, P., Candy, J. F., Malik, J. M., and Edwards, D. D. (1996). *Artificial intelligence: a modern approach*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Shen, Y., Ng, A., and Seeger, M. (2006). Fast Gaussian process regression using kd-trees. *Advances in neural information processing systems*, 18:1225.
- Shental, N., Hertz, T., Weinshall, D., and Pavel, M. (2002). Adjustment learning and relevant component analysis. In *Proceedings of the 7th European Conference on Computer Vision-Part IV*, pages 776–792. Springer.
- Shewchuk, J. (1994). An introduction to the conjugate gradient method without the agonizing pain. Technical Report CMU-CS-TR-94-125, Carnegie Mellon University.
- Singh-Miller, N. (2010). *Neighborhood Analysis Methods in Acoustic Modeling for Automatic Speech Recognition*. PhD thesis, Massachusetts Institute of Technology.
- van der Maaten, L. (2010). Matlab toolbox for dimensionality reduction. http://homepage.tudelft.nl/19j49/Matlab_Toolbox_for_Dimensionality_Reduction.html. (Date accessed: August 2011).
- Vogl, T., Mangis, J., Rigler, A., Zink, W., and Alkon, D. (1988). Accelerating the convergence of the back-propagation method. *Biological Cybernetics*, 59(4):257–263.

- Weinberger, K. and Tesauro, G. (2007). Metric learning for kernel regression. In *Eleventh international conference on artificial intelligence and statistics*, pages 608–615.
- Weinberger, K. Q., Blitzer, J., and Saul, L. K. (2006). Distance metric learning for large margin nearest neighbor classification. In *NIPS*. MIT Press.
- Weinberger, K. Q. and Saul, L. K. (2008). Fast solvers and efficient implementations for distance metric learning. In *Proceedings of the 25th international conference on Machine learning, ICML '08*, pages 1160–1167, New York, NY, USA. ACM.
- Weizman, L. and Goldberger, J. (2007). A classification-based linear projection of labeled hyperspectral data. In *Geoscience and Remote Sensing Symposium, 2007. IGARSS 2007. IEEE International*, pages 3202–3205. IEEE.
- Xing, E., Ng, A., Jordan, M., and Russell, S. (2003). Distance metric learning with application to clustering with side-information. In *Advances in Neural Information Processing Systems*, pages 521–528.
- Yang, L. and Jin, R. (2006). Distance metric learning: A comprehensive survey. *Michigan State University*, pages 1–51.