

2004s-29

Locally Weighted Full Covariance Gaussian Density Estimation

Yoshua Bengio, Pascal Vincent

Série Scientifique
Scientific Series

Montréal
Mai 2004

© 2004 Yoshua Bengio, Pascal Vincent. Tous droits réservés. *All rights reserved.* Reproduction partielle permise avec citation du document source, incluant la notice ©.

Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source.

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du ministère de la Recherche, de la Science et de la Technologie, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère de la Recherche, de la Science et de la Technologie, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

PARTENAIRE MAJEUR

- . Ministère du développement économique et régional et de la recherche [MDERR]

PARTENAIRES

- . Alcan inc.
- . Axa Canada
- . Banque du Canada
- . Banque Laurentienne du Canada
- . Banque Nationale du Canada
- . Banque Royale du Canada
- . Bell Canada
- . BMO Groupe Financier
- . Bombardier
- . Bourse de Montréal
- . Caisse de dépôt et placement du Québec
- . Développement des ressources humaines Canada [DRHC]
- . Fédération des caisses Desjardins du Québec
- . GazMétro
- . Hydro-Québec
- . Industrie Canada
- . Ministère des Finances du Québec
- . Pratt & Whitney Canada Inc.
- . Raymond Chabot Grant Thornton
- . Ville de Montréal

- . École Polytechnique de Montréal
- . HEC Montréal
- . Université Concordia
- . Université de Montréal
- . Université du Québec à Montréal
- . Université Laval
- . Université McGill
- . Université de Sherbrooke

ASSOCIE A :

- . Institut de Finance Mathématique de Montréal (IFM²)
- . Laboratoires universitaires Bell Canada
- . Réseau de calcul et de modélisation mathématique [RCM²]
- . Réseau de centres d'excellence MITACS (Les mathématiques des technologies de l'information et des systèmes complexes)

Les cahiers de la série scientifique (CS) visent à rendre accessibles des résultats de recherche effectuée au CIRANO afin de susciter échanges et commentaires. Ces cahiers sont écrits dans le style des publications scientifiques. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

Locally Weighted Full Covariance Gaussian Density Estimation

Yoshua Bengio^{*}, Pascal Vincent[†]

Résumé / Abstract

Nous décrivons une application du principe d'apprentissage local à l'estimation de densité. Le lissage pondéré localement d'une gaussienne utilisant une matrice de covariance pleine et régularisée conduit à un estimateur de densité ayant un comportement amélioré lorsque la masse de probabilité est concentrée le long d'une variété de basse dimension. Même si l'estimateur proposé n'est pas garanti d'intégrer à 1 sur un ensemble de données fini, nous prouvons la convergence asymptotique de la vraie densité. Les résultats expérimentaux illustrant les avantages de cet estimateur sur les estimateurs non paramétriques classiques sont présentés.

Mots clés : estimation de densité, modèles non paramétriques, modèles de variétés, convergence des estimateurs de densité.

We describe an interesting application of the principle of local learning to density estimation. Locally weighted fitting of a Gaussian with a regularized full covariance matrix yields a density estimator which displays improved behavior in the case where much of the probability mass is concentrated along a low dimensional manifold. While the proposed estimator is not guaranteed to integrate to 1 with a finite sample size, we prove asymptotic convergence to the true density. Experimental results illustrating the advantages of this estimator over classic non-parametric estimators are presented.

Keywords: *density estimation, non-parametric models, manifold models, convergence of density estimators.*

^{*} Yoshua Bengio, Université de Montréal et CIRANO, tél. : (514) 343-6804, Yoshua.Bengio@cirano.qc.ca

[†] Pascal Vincent, ApSTAT Technologies Inc., tél. : (514) 343-9119 #219, vincentp@apstat.com

1 Introduction

Most machine-learning problems, as they occur in nature, are posed in a very high dimensional space. However, overcoming the curse of dimensionality has been an open problem since it was first described in (Bellman, 1961). The popularity of the new generation of kernel methods, in particular the Support Vector Machines (Boser, Guyon and Vapnik, 1992; Vapnik, 1995), is due in part to relatively good performance on high dimensional problems, while the traditional kernel methods (s.a. Parzen windows (Parzen, 1962)) often perform more poorly. Another, recently revived, and very promising research trend in dealing with the curse, is that of manifold learning. It is based on the idea that the data lives on (or close to) a non-linear manifold of much lower dimensionality, embedded in the high dimensional space. This trend is exemplified by Locally Linear Embedding (Roweis and Saul, 2000) and Isomap (Tenenbaum, de Silva and Langford, 2000) but also underlies the idea of mixtures of factor analyzers and similar algorithms (Ghahramani and Hinton, 1996; Hinton, Revow and Dayan, 1995; Tipping and Bishop, 1999; Ghahramani and Beal, 2000).

Our line of research attempts to integrate the notions of manifold modeling with the traditional non-parametric kernel and distance based methods such as k-nearest-neighbors and Parzen windows. We have already proposed improved algorithms for classification (Vincent and Bengio, 2002) and density estimation (Vincent and Bengio, 2003) but the latter one suffers from serious practical difficulties¹. In this paper, we propose a different approach to density estimation which does not pose the same memory requirement problems as (Vincent and Bengio, 2003), and is based on a general principle taking the point of view of local learning.

Local learning (Bottou and Vapnik, 1992; Atkeson, Moore and Schaal, 1997; Ormoneit and Hastie, 2000). can be understood as a general principle that allows to extend learning techniques designed for simple models, to the case of complex data for which the model's assumptions would not necessarily hold globally, but can be thought as valid *locally*. A simple example is the assumption of linear separability, which in general is *not* satisfied globally in classification problems with rich data. Yet any classification algorithm able to find only a linear separation, can be used inside a local learning procedure, yielding an algorithm able to model complex non-linear class boundaries.

Similarly, for density estimation, while it is in general unreasonable to assume that the data follows a Gaussian distribution globally, the Gaussian approximation holds locally. Note that if the data lies close to a low dimensional manifold, then the shape of that local Gaussian will be a flattened pancake, and it's crucial to use a non-spherical Gaussian, to capture the local principal directions of the manifold.

Traditional parametric density estimation can be formulated as the question: "What is the likelihood of a test point x under a model fitted to the whole training data". We formulate the principle of locally weighted density estimation in a similar manner as "What is the

¹Its memory requirement scales in $\mathcal{O}(n.d^2)$ where d is the input dimensionality, and n is the number of training samples, making it impossible to use with large, high dimensional data sets.

likelihood of a test point x under a *simple model* fitted only to the *local* training data in the neighborhood of x .”

Notice that locally weighted density estimation yields an unnormalized density estimate: in general it won't integrate to 1, as is also the case of several classical non-parametric density estimators, similar in spirit, like the nearest neighbor density estimator (Fix and Hodges, 1951; Loftsgaarden and Quesenberry, 1965). (see (Inzenman, 1991) for a survey of non-parametric density estimation techniques).

Local learning typically comes in two flavors, depending on the notion of “neighborhood”. The neighborhood is always based on some a-priori measure of locality (such as the Euclidean distance in input space), but it can be either defined as a “hard” neighborhood (the set of k nearest neighbors of x for instance), or as a “soft”, weighted neighborhood (the set of all training points, but with an associated weight given by a prior, continuous weighting kernel \mathcal{K} , centered on x). The former can be seen as a special case of the latter, with a particular discontinuous weighting kernel giving only weights of 0 or 1). We would like to stress the importance of using a “soft” neighborhood to avoid discontinuities in the estimate, a problem that plagues the nearest neighbor density estimator. Indeed, while *some* statistics of the set of k neighbors (such as the distance to the k^{th} neighbor) vary smoothly with x , the set of k neighbors doesn't: a small variation in x may yield a totally different k^{th} neighbor, and thus lead to a discontinuous estimate, if a “hard” neighborhood is used. Consequently the local model fitting procedure should accommodate sample weights.

For previous work on local learning applied to non-parametric density estimation, see also (Hjort and Jones, 1996).

2 The Locally Weighted Density Estimator

Let $D = \{x_1, \dots, x_n\}$ a data set with x_i sampled i.i.d. from an unknown distribution with continuous density $f(x)$.

Let $\mathcal{K}(x_i; x)$ a weighting Kernel centered on x , used to give a weight to every x_i .

In addition, we suppose that it is easy to fit a simple parametric model \mathcal{M} to a data set endowed with sample weights (e.g. the maximum of the log-likelihood can be found analytically or cheaply).

The locally weighted density estimation principle computes an estimate $\hat{f}(x)$ of the density at a given test point x as follows:

- Associate a weight $\mathcal{K}(x_i; x)$ to every $x_i \in D$
- Fit model \mathcal{M} to the weighted data.
- The estimate is $\hat{f}(x) = \frac{1}{Z} \mathcal{M}(\xi)$ with Z a normalization factor to try to make $\hat{f}(x)$ integrate to 1 over the domain of x (at least asymptotically).

In our particular case we use $\mathcal{K}(x_i; x) = \mathcal{N}(x_i; x, \sigma_n^2(x)I)$, with $\sigma_n(x) = \alpha d(x, x_{v_k})$, where x_{v_k} denotes the k^{th} neighbor of x according to the Euclidean distance d . i.e. the *width* of our weighting kernel is proportional to the distance from its center to its k^{th} neighbor.

The model \mathcal{M} we fit to the weighted training samples is a Gaussian with a regularized full covariance matrix, i.e.: $\mathcal{M}(\xi) = \mathcal{N}(\xi; \mu_\backslash(\xi), \mathcal{C}_\backslash(\xi))$, where $\mu_n(x) = \frac{1}{\sum_i \mathcal{K}(x_i; x)} \sum_i \mathcal{K}(x_i; x)x_i$ is the weighted sample mean, and $C_n(x) = \frac{1}{\sum_i \mathcal{K}(x_i; x)} \sum_i \mathcal{K}(x_i; x)(x_i - \mu_n(x))(x_i - \mu_n(x))' + \gamma_n^2 I$ is the weighted sample covariance with an additional regularization parameter γ_n .

$\mathcal{N}(x; \mu, \Sigma)$ denotes the density of x under a multivariate Normal density with center μ and covariance matrix Σ :

$$\mathcal{N}(x; \mu, \Sigma) = \frac{e^{-\frac{1}{2}(x-\mu)'\Sigma^{-1}(x-\mu)}}{(2\pi)^{d/2}|\Sigma|^{1/2}} \quad (1)$$

We denote the resulting estimator f^{WN} ,

$$f^{WN}(x) = \frac{1}{Z} \mathcal{N}(x; \mu_n(x), C_n(x)), \quad (2)$$

and will prove its asymptotic convergence to the true density $f(x)$, when the normalization factor is estimated with $Z = \frac{nV_d(1)}{k_n(2\pi)^{d/2}\alpha^d}$, with $V_d(r)$ being the volume of a d -ball of radius r in \mathbf{R}^d , i.e. $V_d(r) = r^d V_d(1) = r^d \frac{\pi^{d/2}}{\Gamma(1+\frac{d}{2})}$.

A slight variant of the above f^{WN} will also be considered (in particular for the convergence proof) in which $\mu_n(x)$ is fixed on $\mu_n(x) = x$ rather than on the weighted sample average.

3 Other Classical Estimators Considered

In addition we'll consider the following classical non-parametric estimators, using the above definitions.

1. The fixed width Parzen density estimator (Parzen, 1962):

$$f^{FP}(x) = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(x; x_i, r_n^2 I). \quad (3)$$

2. The fixed width nearest neighbor density estimator (Parzen, 1962):

$$f^{FN}(x) = \frac{|N(x, r_n)|}{nV_d(r_n)} \quad (4)$$

where $N(x, r_n) = \{x_i \in D : \|x - x_i\| < r_n\}$ are the neighbors of x closer than r_n and $\frac{|N(x, r_n)|}{n}$ estimates unbiasedly the probability of falling in the d -ball of radius r_n around x .

3. The k-nearest neighbor density estimator (Fix and Hodges, 1951; Loftsgaarden and Quesenberry, 1965):

$$f^{KN}(x) = \frac{k_n}{nV_d(\sigma_n(x))} \quad (5)$$

where $\frac{k_n}{n}$ estimates the probability of falling in the d -ball of radius $\sigma_n(x)$ around x , and we select $\alpha = 1$ in $\sigma_n(x)$.

4 Asymptotic Convergence Analysis

To prove convergence, a number of the quantities defined above will be subscripted with n . We have to impose conditions on the rate of modification of k_n which must increase with n , and the way it affects σ_n i.e. $\lim_{n \rightarrow \infty} k_n = \infty$ but slower than n , i.e. $\lim_{n \rightarrow \infty} \frac{k_n}{n} = 0$. We let k_n increase at a slow enough rate such that $n\sigma_n(x)^d \rightarrow \infty$. Besides, γ_n must go to zero faster than σ_n .

We already know from the classical literature (Parzen, 1962; Fix and Hodges, 1951; Loftsgaarden and Quesenberry, 1965) that f^{FP} , f^{FN} and f^{KN} are consistent i.e. that their limit as $n \rightarrow \infty$ is f for every x (with the hyperparameters converging at a proper rate, as defined above).

Lemma 1 . Let r_n be a probabilistic lower bound on σ_n , i.e. $\sigma_n(x) > r_n$ with high probability $1 - \delta$. Let $g(x, y)$ be a function that is $O(y^m)$ and is $O(\sigma_n(x)^{-p})$. The limit of a locally weighted average of $g(x, x_i)$ with $\sigma_n(x)$ converges to the same limit as $f(x) \int g(x, y)\mathcal{N}(y; x, \sigma_n^2(x)I)dy$:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_i g(x, x_i)\mathcal{N}(x_i; x, \sigma_n^2(x)I) = f(x) \lim_{n \rightarrow \infty} \int g(x, y)\mathcal{N}(y; x, \sigma_n^2(x)I)dy \quad (6)$$

To obtain convergence, then r_n must decrease at a slow enough rate such that $nr_n^{d+2(p-m)} \rightarrow \infty$.

Sketch of the Proof

Consider the left-hand-side of 6 with $\sigma_n(x)$ replaced by r_n (even in g). Its expected value is $\int f(y)g(x, y)\mathcal{N}(y; x, r_n(x)^2I)dy$. Let us show that the average converges to its expected value. For this we will show that the average error (i.e. the variance) goes to 0. Below we will show that for a data-independent spread r_n in the local weighting kernel \mathcal{N} , the variance of the average is inversely proportional to a power of r_n . Therefore with probability $1 - \delta$, the variance of the average using $\sigma_n(x)$ is less than the variance of the average using r_n . Since the latter will be shown to converge to zero we will obtain convergence to zero of the variance of the desired average. Consider the variance of the average with data-independent spread r_n :

$$\begin{aligned} \text{Var}\left[\frac{1}{n} \sum_i g(x, x_i)\mathcal{N}(x_i; x, r_n^2I)\right] &= \frac{\text{Var}[g(x, x_i)\mathcal{N}(x_i; x, r_n^2I)]}{n} \\ &< \frac{\int f(y)g(x, y)^2\mathcal{N}(y; x, r_n^2I)^2dy}{n} \end{aligned}$$

where the variance is over D , and we simply dropped $E[g(x, x_i)\mathcal{N}(x_i; x, r_n^2I)]^2$. Let us now make the change of variable $z = \sqrt{2}(y - x)/r_n$ (this is a vector in \mathbf{R}^d , i.e. $dy_1 \dots dy_d = 2^{-d/2}r_n^d dz_1 \dots dz_d$), yielding the bound on the variance

$$\frac{1}{(2\pi)^{d/2}n2^{d/2}r_n^d} \int f\left(\frac{r_n z}{\sqrt{2}} + x\right)g\left(x, \frac{r_n z}{\sqrt{2}} + x\right)^2\mathcal{N}(z; 0, 1)dz$$

using $\mathcal{N}(y; x, r_n^2I)^2 = \frac{e^{-\|y-x\|^2/r_n^2}}{(2\pi)^d r_n^{2d}} = \frac{1}{(2\pi)^{d/2}r_n^{2d}} \frac{e^{-\|z\|^2}}{(2\pi)^{d/2}} = \frac{1}{(2\pi)^{d/2}r_n^{2d}}\mathcal{N}(z; 0, 1)$. If g

varies with n in $\frac{1}{r_n^p}$ and in the power m of its second argument, then the integrand varies in $\frac{1}{r_n^{2(p-m)}}$: the condition on r_n is indeed $nr_n^{d+2(p-m)} \rightarrow \infty$.
Q.E.D.

Note that as a special case we obtain the convergence of f^{FP} , with $g(x, x_i) = 1$.

Lemma 2 *If $n\sigma_n^{d+1}(x) \rightarrow \infty$ and $\sigma_n(x) \rightarrow 0$ the locally weighted version of $\mu_n(x)$ converges to x and $\frac{\mu_n(x) - x}{\sigma_n(x)}$ converges to 0, in probability:*

$$\lim_{n \rightarrow \infty} \frac{\sum_i x_i \mathcal{N}(x_i; x, \sigma_n^2(x)I)}{\sum_i \mathcal{N}(x_i; x, \sigma_n^2(x)I)} = x$$

and

$$\lim_{n \rightarrow \infty} \frac{\mu_n(x) - x}{\sigma_n(x)} = 0,$$

in probability.

Proof

The denominator of $\mu_n(x)$ times $\frac{1}{n}$ is f^{FP} , which converges to $f(x)$. For the numerator, we apply Lemma 1 with $g(x, x_i) = \frac{x_i - x}{\sigma_n(x)}$ (i.e. $p = m = 1$) and obtain $f(x) \lim_{n \rightarrow \infty} \int \frac{y-x}{\sigma_n(x)} \mathcal{N}(y; x, \sigma_n(x)^2 I) dy$.

We apply the change of variable $z = \frac{y-x}{\sigma_n(x)}$ and obtain $f(x) \lim_{n \rightarrow \infty} \int \sigma_n(x)^d z \mathcal{N}(z; 0, I) dz = 0$.

Q.E.D.

Lemma 3 *The locally weighted covariance $C_n(x)$ has the same limit as $\sigma_n^2(x)I$.*

Proof

We rewrite the numerator and denominator in the first term of C_n as averages. Using Lemma 1 and Lemma 2 the numerator has the same limit as

$$f(x) \int (x - y)(x - y)' \mathcal{N}(y; x, \sigma_n^2(x)I) dy.$$

Apply the change of variable $z = y - x$, yielding

$$f(x) \int zz' \mathcal{N}(z; 0, \sigma_n(x)^2 I) dz \rightarrow f(x) \sigma_n(x)^2 I.$$

As in previous proofs, the denominator converges to $f(x)$.

$$C_n(x) \rightarrow (\sigma_n^2(x) f(x) + \gamma_n^2) I$$

Since we have assumed $\frac{\gamma_n}{\sigma_n} \rightarrow 0$, the second term can be ignored.

Q.E.D.

Theorem 1 *The locally weighted full covariance matrix estimator f^{WN} is consistent (converges to f) for both versions of $\mu_n(x)$.*

Proof

Consider the numerator of $\mathcal{N}(x; \mu_n(x), C_n(x))$ in eq. 1. When $\mu_n(x) = x$ it is simply equal to $e^0 = 1$. For the other versions, Lemma 2 shows that it converges to 1. Using

Lemma 3, and $|aB| = a^d|B|$ for a $d \times d$ matrix B and a scalar a , the denominator $(2\pi)^{d/2} \sqrt{|C_n(x)|}$ has the same limit as

$$(2\pi)^{d/2} \sqrt{|\sigma_n^2(x)I|} = (2\pi)^{d/2} \sigma_n^d(x).$$

Finally, we use the main result in (Loftsgaarden and Quesenberry, 1965), i.e.

$$\lim_{n \rightarrow \infty} \frac{k_n}{nV_d(1)\sigma_n^d(x)} = f(x).$$

when $\alpha = 1$, i.e. equal to $f(x)/\alpha^d$ when using a different value of α . Putting these together with the formula for f^{WN} , we obtain

$$\begin{aligned} \lim_{n \rightarrow \infty} f^{WN}(x) &= \lim_{n \rightarrow \infty} \frac{k_n \alpha^d (2\pi)^{d/2}}{nV_d(1)} \mathcal{N}(x; \mu_n(x), C_n(x)) \\ &= \lim_{n \rightarrow \infty} \frac{k_n \alpha^d (2\pi)^{d/2}}{nV_d(1) (2\pi)^{d/2} \sigma_n^d(x)} \\ &= f(x). \end{aligned}$$

Q.E.D.

5 Experiments

To assess the performance of the proposed algorithm, we performed the following experiment on a 2D spiral problem:

A training set of 300 points, a validation set of 300 points (reserved for tuning hyper-parameters), and a test set of 10000 points were generated from the following distribution of two dimensional (x, y) points:

$$x = 0.04 t \sin(t) + \epsilon_x, \quad y = 0.04 t \cos(t) + \epsilon_y$$

where $t \sim U(3, 15)$, $\epsilon_x \sim \mathcal{N}(0, 0.01)$, $\epsilon_y \sim \mathcal{N}(0, 0.01)$, $U(a, b)$ is uniform in the interval (a, b) and $\mathcal{N}(\mu, \sigma)$ is a normal density.

As the density estimators under consideration are not guaranteed to integrate to 1 (except for f^FP), we compute an approximation of their integral by sampling on a 300×300 regularly spaced grid (covering the x and y range of the training set plus 10%), and divide the raw estimators by the obtained integral approximation, yielding normalized estimators for our comparison study. Evaluation of the estimators over that same grid is also used to produce the graphs of Figure 1. The performance of each normalized estimator \hat{f} is then evaluated using the average log likelihood (ALL) over the test set T (of size $m = 10000$):

$$ALL(\hat{f}, T) = \frac{1}{m} \sum_{x \in T} \log \hat{f}(x)$$

For each estimator type, we tried several values of the hyper-parameters, keeping the choice that yielded the largest possible ALL over the validation set.

Table 1: Performance of various estimators on the spiral data, measured as average log-likelihood over the test set (standard errors are in parenthesis).

Estimator	Hyper-parameters used	ALL on test set
f^{FP}	$r = 0.0151$	1.292 (0.012)
f^{KN}	$k = 2$	1.058 (0.011)
f^{FN}	$r = 0.065$	0.739 (0.018)
f^{WN} fixed σ	$\sigma = 0.04, \gamma = 6e - 5$	1.461 (0.006)
f^{WN}	$\alpha = 0.23, k = 25, \gamma^2 = 3e - 5$	1.561 (0.008)
f^{WN} fixed $\sigma, \mu = x$	$\sigma = .045, \gamma^2 = 1e - 9$	0.759 (0.003)
$f^{WN}, \mu = x$	$\alpha = 1, k = 2, \gamma^2 = 1e - 5$	1.201 (0.010)

Results are reported in Table 1. f^{WN} (with local $\sigma(x)$ and using the local average for $\mu(x)$) appears to perform significantly better than the other classical estimators, indicating that it was able to more accurately capture and model the underlying true distribution. The difference between f^{FP} and f^{WN} can also be better appreciated qualitatively in Figure 1. For f^{WN} , the experiments suggest it is better to use $\alpha < 1$, which yields a less variable $\sigma(x)$, and it is better to use an adaptive $\sigma(x)$ than a fixed σ . The validation hyper-parameter selection also chooses a non-zero γ , which suggests that it is also useful.

6 Conclusions

To summarize, we have introduced and analyzed a family of non-parametric locally weighted density estimators that are appropriate to model the local manifold structure of data, and experiments suggest that it performs well against classical estimators, in addition to being much more memory-efficient than the related estimator proposed in (Vincent and Bengio, 2003).

References

- Atkeson, C. G., Moore, A. W., and Schaal, S. (1997). Locally weighted learning for control. *Artificial Intelligence Review*, 11:75–113. 1
- Bellman, R. (1961). *Adaptive Control Processes: A Guided Tour*. Princeton University Press, New Jersey. 1
- Boser, B., Guyon, I., and Vapnik, V. (1992). A training algorithm for optimal margin classifiers. In *Fifth Annual Workshop on Computational Learning Theory*, pages 144–152, Pittsburgh. 1
- Bottou, L. and Vapnik, V. (1992). Local learning algorithms. *Neural Computation*, 4(6):888–900. 1
- Fix, E. and Hodges, J. (1951). Discriminatory analysis, non-parametric discrimination, consistency properties. Technical Report Report 21-49-004, USAF School of Aviation Medicine, Randolph Field, Texas. 2, 3, 4
- Ghahramani, Z. and Beal, M. J. (2000). Variational inference for Bayesian mixtures of factor analysers. In *Advances in Neural Information Processing Systems 12*, Cambridge, MA. MIT Press. 1

- Ghahramani, Z. and Hinton, G. (1996). The EM algorithm for mixtures of factor analyzers. Technical Report CRG-TR-96-1, Dpt. of Comp. Sci., Univ. of Toronto. [1](#)
- Hinton, G., Revow, M., and Dayan, P. (1995). Recognizing handwritten digits using mixtures of linear models. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems 7*, pages 1015–1022. MIT Press, Cambridge, MA. [1](#)
- Hjort, N. L. and Jones, M. C. (1996). Locally parametric nonparametric density estimation. *Annals of Statistics*, 24(4):1619–1647. [2](#)
- Inzenman, A. (1991). Recent developments in nonparametric density estimation. *Journal of the American Statistical Association*, 86(413):205–224. [2](#)
- Loftsgaarden, D. and Quesenberry, C. (1965). A nonparametric estimate of a multivariate density function. *Annals of Mathematical Statistics*, 36:1049–1051. [2](#), [3](#), [4](#), [6](#)
- Ormonet, D. and Hastie, T. (2000). Optimal kernel shapes for local linear regression. In Solla, S., Leen, T., and Müller, K.-R., editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press. [1](#)
- Parzen, E. (1962). On the estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1064–1076. [1](#), [3](#), [4](#)
- Roweis, S. and Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326. [1](#)
- Tenenbaum, J., de Silva, V., and Langford, J. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323. [1](#)
- Tipping, M. and Bishop, C. (1999). Mixtures of probabilistic principal component analysers. *Neural Computation*, 11(2):443–482. [1](#)
- Vapnik, V. (1995). *The Nature of Statistical Learning Theory*. Springer, New York. [1](#)
- Vincent, P. and Bengio, Y. (2002). K-local hyperplane and convex distance nearest neighbor algorithms. In Dietterich, T., Becker, S., and Ghahramani, Z., editors, *Advances in Neural Information Processing Systems*, volume 14, Cambridge, MA. MIT Press. [1](#)
- Vincent, P. and Bengio, Y. (2003). Manifold parzen windows. In Becker, S., Thrun, S., and Obermayer, K., editors, *Advances in Neural Information Processing Systems 15*, Cambridge, MA. MIT Press. [1](#), [7](#)

Fixed Parzen estimator f^{FP}

Locally weighted Gaussian f^{WN}

