# Nonparametric Methods for Density Estimation Nearest Neighbour Classification 

Lecturer:
Ondřej Drbohlav

Authors:<br>Ondřej Drbohlav, Jiří Matas<br>Centre for Machine Perception<br>Czech Technical University, Prague http://cmp.felk.cvut.cz<br>Last update: Oct 2022

## Probability Density Estimation

## Parametric Methods for Density Estimation

- Have been dealt with in the previous lecture
- Advantage: Low number of parameters to estimate
- Disadvantage: The resulting estimated density can be arbitrarily wrong if the underlying distribution does not agree with the assumed parametric model.


## Non-Parametric Methods for Density Estimation

- Histogram
- Nearest Neighbor approach


## Histogram as piecewise constant density estimate: Task formulation

Consider the following distribution $q(x)$ on the interval $[0,1]$, and i.i.d. sampling from it. We will fit the distribution by a 'histogram' with $B$ bins. More precisely, we will estimate a piecewise-constant function on the interval $[0,1]$ with $B$ segments of the same width. For a given $B$, the parameters of this piecewise-constant function are the heights $d_{1}, d_{2}, \ldots, d_{B}$ of the individual bins. This function is denoted $p\left(x \mid\left\{d_{1}, d_{2}, \ldots, d_{B}\right\}\right)$.


$$
\begin{gathered}
p\left(x \mid\left\{d_{1}, d_{2}, . ., d_{B}\right\}\right) \text { to be } \\
\text { estimated }
\end{gathered}
$$



For the given number of bins $B, d_{1}, d_{2}, \ldots, d_{B}$ must conform to the constraint that the area under the function must sum up to one,

$$
1=\int_{-\infty}^{\infty} p\left(x \mid\left\{d_{1}, d_{2}, \ldots, d_{B}\right\}\right) \mathrm{d} x=\sum_{i=1}^{B} \int_{\frac{i-1}{B}}^{\frac{i}{B}} d_{i} \mathrm{~d} x=\sum_{i=1}^{B} \stackrel{\text { bin width }}{\downarrow} d_{i}{ }^{B}=\sum_{i=1}^{B} \frac{d_{i}}{B}
$$

## Histogram as piecewise constant density estimate:

## Finding $d_{i}$ 's using Maximum Likelihood

Let us estimate $\left\{d_{i}, i=1,2, \ldots, B\right\}$ by Maximum Likelihood (ML) approach. Let $N_{i}$ denote the number of samples which belong the $i$-th bin (thus clearly, $\sum_{i=1}^{B} N_{i}=N$ ). The likelihood $L(\mathcal{T})$ of observing the samples $\mathcal{T}=\left\{x_{1}, x_{2}, \ldots, x_{N}\right\}$ given the parameters $\boldsymbol{\theta}=\left\{d_{1}, d_{2}, \ldots, d_{B}\right\}$ is

$$
\begin{equation*}
L(\mathcal{T})=p(\mathcal{T} \mid \boldsymbol{\theta})=\prod_{i=1}^{N} p\left(x_{i} \mid \boldsymbol{\theta}\right)=\prod_{j=1}^{B} \overbrace{\left(\prod_{k=1}^{N_{j}} d_{j}\right)}=\prod_{j=1}^{B} d_{j}^{N_{j}} . \tag{2}
\end{equation*}
$$

The maximization task is then

$$
\begin{equation*}
\ell(\mathcal{T})=\sum_{j=1}^{B} N_{j} \log d_{j} \rightarrow \max , \quad \text { subject to } \frac{1}{B} \sum_{j=1}^{B} d_{j}=1 \tag{3}
\end{equation*}
$$

where maximization has been formulated using the $\log$-likelihood $\ell(\mathcal{T})$. The Lagrangian of the optimization task and the conditions of optimality (using the derivative $\partial / \partial d_{k}$ ) are then:

$$
\begin{align*}
& \text { Lagrangian: } \sum_{j=1}^{B} N_{j} \log d_{j}+\lambda\left(\frac{1}{B} \sum_{j=1}^{B} d_{j}-1\right)  \tag{4}\\
& \frac{N_{k}}{d_{k}}+\frac{\lambda}{B}=0 \Rightarrow \frac{d_{k}}{N_{k}}=\text { const. } \Rightarrow d_{k}=B \frac{N_{k}}{N} . \tag{5}
\end{align*}
$$

Histogram as piecewise constant density estimate: Example, different number of bins
$d_{k}=B \frac{N_{k}}{N}$
This result is in line with the common use of histograms for approximating pdf's. Results for different $B$ 's:










## Histogram as piecewise constant density estimate: <br> What number of bins produces closest pdf approximation?

Let us measure the differences between the (actual) source distribution $q(x)$ and the piecewise-constant density estimate $p(x)=p\left(x \mid\left\{d_{1}, d_{2}, \ldots, d_{B}\right\}\right)$ from the $N=1000$ samples, using $B$ bins.

Measures used:
Kullback-Leibler divergence $D_{\mathrm{KL}}$ :

$$
\begin{equation*}
D_{\mathrm{KL}}(p \| q)=\int_{-\infty}^{\infty} p(x) \log \frac{p(x)}{q(x)} \mathrm{d} x \tag{7}
\end{equation*}
$$

(Note that KL div. is not a metric.)
Sum of squared differences $D_{\mathrm{SSD}}$ :

$$
\begin{equation*}
D_{\mathrm{SSD}}(p, q)=\int_{-\infty}^{\infty}(p(x)-q(x))^{2} \mathrm{~d} x \tag{8}
\end{equation*}
$$




Histogram as piecewise constant density estimate: Choosing the number of bins $B$ by ML

How can we find the optimal number of bins $B$ ? Let us try to employ the ML approach again: find the $B$ which maximizes the likelihood. Recall that:
parameters $d_{j}: \quad d_{j}=B \frac{N_{j}}{N} \quad(\mathrm{ML}$ estimate $)$
likelihood $L(\mathcal{T}): \quad L(\mathcal{T})=p\left(\mathcal{T} \mid\left\{d_{1}, d_{2}, \ldots, d_{B}\right\}\right)=\prod_{j=1}^{B} d_{j}^{N_{j}}=\prod_{j=1}^{B}\left(\frac{B N_{j}}{N}\right)^{N_{j}}$
log-likelihood $\ell(\mathcal{T}): \quad \ell(\mathcal{T})=\sum_{j=1}^{B} N_{j} \log d_{j}=\sum_{j=1}^{B} N_{j} \log \frac{B N_{j}}{N}$


For $B=4000$, the log-likelihood $\ell$ is the highest.
But the pdf estimate with this $B$ is poor, and very different from the source distribution as measured by $D_{\mathrm{KL}}$ or $D_{\mathrm{SSD}}$.
For $B=10^{5}, \ell(\mathcal{T}) \sim 4600$.
What went wrong?

## Histogram, choosing the number of bins $B$ : ML overfits and produces $B=\infty$

When $B$ grows, eventually it will reach a number $\hat{B}$ such that there is either no or one point in every bin (assuming no two points in the data are identical), and this will stay true for any $B>\hat{B}$.

In such cases,

$$
d_{j}=\left\{\begin{array}{l}
\frac{B}{N} \text { if the bin is populated by a point }  \tag{12}\\
0 \text { if the bin is not populated. }
\end{array}\right.
$$

As the number of bins $B$ grows, the widths of occupied bins get narrower and the heights $d_{j}$ 's higher. If $B \rightarrow \infty$ then also $d_{j} \rightarrow \infty$ for the occupied bins, and therefore also $\ell(\mathcal{T}) \rightarrow \infty$. Thus, such an approach cannot produce a "reasonable" answer to choosing $B$, as the solution it provides is $B=\infty$.


The problem is that the log-likelihood $\ell$ is computed using the same data used for fitting the model (computing $d_{i}$ 's). This is a similar concept to training a classifier on certain data and testing on the same data, which is prone to over-fitting and poor generalization.

## Histogram, choosing the number of bins $B$ : Employing cross-validation

Let us compute the log likelihood using the following procedure: remove a given point from the dataset for computing $d_{i}$ 's and evaluate its contribution to the log-likelihood. Do this for all the points. This approach is related to cross-validation technique (leave-one-out) for choosing parameters of a classifier.

Let the point in question belong to the $j$-th bin. The ML estimate for $d_{j}$, after removing this point from the dataset, is

$$
\begin{equation*}
d_{j}=B \frac{N_{j}-1}{N-1}, \quad\left(N_{j} \geq 1\right) \tag{13}
\end{equation*}
$$

where the subtractions of 1 reflect the fact that the considered point is not used for estimating $d_{j}$. Computing the log likelihood $\ell$ this way produces the following result:


$$
\begin{aligned}
& \ell=\sum_{\substack{j=1 \\
N_{j} \geq 1}}^{B} N_{j} \log d_{j}, \\
& \text { with } d_{j}=B \frac{N_{j}-1}{N-1}
\end{aligned}
$$

The 'failure' for $B>7$ is caused by singly-occupied bins ( $N_{j}=1$ ) for which the modified ML estimate for $d_{j}$ becomes zero. This will be fixed by using different estimates for $d_{j}$ 's.

## Histogram, choosing the number of bins $B$ : More suitable estimates for $d_{j}$ 's

The problem of $d_{i}$ being estimated as 0 is similar to the one encountered previously: Recall the example of tossing a coin three times, always getting heads $(\mathcal{T}=\{H, H, H\})$. The ML estimate is a fully unfair coin (probability of getting heads is $1, \pi_{\text {head }}=1$ ), thus making the likelihood of any sequence containing tails zero. We have seen before that employing the prior for the parameters to be estimated can mitigate this problem.
A (conjugate) prior for the histogram bin counts is the Dirichlet Distribution, with the pdf $p\left(d_{1}, d_{2}, \ldots, d_{B} \mid \alpha_{1}, \alpha_{2}, \ldots, \alpha_{B}\right) \sim \prod d_{i}^{\alpha_{i}-1}$.

MAP Estimate:
Bayes Estimate:

$$
\begin{equation*}
d_{i}=B \frac{N_{i}+\alpha_{i}-1}{N+\sum_{i=1}^{B} \alpha_{i}-B} \tag{14}
\end{equation*}
$$

$$
\begin{equation*}
d_{i}=B \frac{N_{i}+\alpha_{i}}{N+\sum_{i=1}^{B} \alpha_{i}} \tag{15}
\end{equation*}
$$

Interpretation: The parameters $\alpha_{i}$ 's can be interpreted as 'virtual' observations, as if $\alpha_{k}$ points have already been assigned to the $k$-th bin.

Example: The Bayes estimate using $\alpha_{i}=1$ for all $i=1,2, \ldots, B$ is

$$
\begin{equation*}
d_{i}=B \frac{N_{i}+1}{N+B} \tag{16}
\end{equation*}
$$

Using this estimate will enable us to make reasonable computation of likelihood for all $B$ 's.

# Histogram, choosing the number of bins $B$ : ML to find $B$, cross-validation, Bayes esimate for $d_{j}$ 's 

Let us now return to the previous task. Compute the log likelihood using the following procedure: remove a given point from the dataset for computing $d_{i}$ 's and evaluate its contribution to the log-likelihood. Do this for all the points.

Use the Bayes estimate for $d_{j}$ from the previous example, $d_{j}=B \frac{N_{j}+1}{N+B}$. The modified estimation of $d_{j}$ (omitting the point in question) will become

$$
\begin{equation*}
d_{j}=B \frac{N_{j}}{N-1+B} . \tag{17}
\end{equation*}
$$

This leads to the following result:


$$
\begin{aligned}
& \ell=\sum_{j=1}^{B} N_{j} \log d_{j}, \\
& \text { with } d_{j}=B \frac{N_{j}}{N-1+B}
\end{aligned}
$$

This result is in agreement with distribution differences as measured by $D_{\mathrm{KL}}$ or $D_{\text {SSD }}$. In particular, $B=30$ is identified as the best-approximating number of bins.

## K-Nearest Neighbor Approach to Density Estimation

Find $K$ neighbors, the density estimate is then $p \sim 1 / V$ where $V$ is the volume of a minimum cell containing $K$ NNs. Example ( $p \sim$ inverse distance to $K$-th NN, same 1000 samples as before):


## K-Nearest Neighbor Approach to Classification

Outline:

- Definition
- Properties
- Asymptotic error of NN classifier
- Error reduction by edit operation on the training class
- Fast NN search


## K-NN Classification Definition

## Assumption:

- Training set $\mathcal{T}=\left\{\left(x_{1}, k_{1}\right),\left(x_{2}, k_{2}\right), \ldots,\left(x_{N}, k_{N}\right)\right\}$. There are $R$ classes (letter $K$ is reserved for $K-\mathrm{NN}$ in this lecture)
- A distance function $d: X \times X \mapsto \mathbb{R}_{0}^{+}$


## Algorithm:

1. Given $x$, find $K$ points $S=\left\{\left(x_{1}^{\prime}, k_{1}^{\prime}\right),\left(x_{2}^{\prime}, k_{2}^{\prime}\right), \ldots,\left(x_{K}^{\prime}, k_{K}^{\prime}\right)\right\}$ from the training set $\mathcal{T}$ which are closest to $x$ in the metric $d$ :

$$
\begin{align*}
S= & \left\{\left(x_{1}^{\prime}, k_{1}^{\prime}\right),\left(x_{2}^{\prime}, k_{2}^{\prime}\right), \ldots,\left(x_{K}^{\prime}, k_{K}^{\prime}\right)\right\} \equiv\left\{\left(x_{r_{1}}, k_{r_{1}}\right),\left(x_{r_{2}}, k_{r_{2}}\right), \ldots,\left(x_{r_{K}}, k_{r_{K}}\right)\right\}  \tag{18}\\
& r_{i}: \text { the rank of }\left(x_{i}, k_{i}\right) \in \mathcal{T} \text { as given by the ordering } d\left(x, x_{i}\right) \tag{19}
\end{align*}
$$

2. Classify $x$ to the class $k$ which has majority in $S$ :

$$
\begin{equation*}
k=\underset{l \in R}{\operatorname{argmax}} \sum_{i=1}^{K} \llbracket k_{i}^{\prime}=l \rrbracket \quad\left(x_{i}^{\prime}, k_{i}^{\prime}\right) \in S \tag{20}
\end{equation*}
$$

## K-NN Example (1)



Consider the two distributions shown. The priors are assumed to be the same,

$$
p(1)=p(2)=0.5
$$

Bayesian optimal decision boundary is shown by the black circle.

Bayesian error is $\epsilon_{B}=0.026$.

## K-NN Example (2)


$K=7$, error $\epsilon=0.030$


$$
K=9, \text { error } \epsilon=0.031
$$



$$
K=5, \text { error } \epsilon=0.032
$$



$$
K=11, \text { error } \epsilon=0.032
$$


$N=100$ samples for each class. Bayes error $\epsilon_{B}=0.026$.

K-NN Example (3)


The results depend on the training set (result of a random process.)
Each of the training sets $\mathcal{T}_{1}, \mathcal{T}_{2}, \mathcal{T}_{3}$ contain 100 points for each class.

## K-NN Example (4)

$K$-NN error for different $K$ and different sizes of the training set ( $N$ samples per class). 10 training sets have been generated randomly for each setting of $K$ and $N$. Average error and its std is shown. Minimum average error is highligted for each $N$. Bayes err. $\epsilon_{B}=2.58 \%$.


## $\boldsymbol{K}$-NN Properties

- Trivial implementation ( $\rightarrow$ good baseline method)
- 1-NN: Bayes error $\epsilon_{B}$ is the lower bound on error of classification $\epsilon_{N N}$ (in the asymptotic case $N \rightarrow \infty$.) Upper bounds can also be constructed, e.g. $\epsilon_{N N} \leq 2 \epsilon_{B}$
- Slow when implemented naively, but can be sped up (Voronoi, k-D trees)
- High computer memory requirements (but training set can be edited and its cardinality decreased)
- How to construct the metric $d$ ? (problem of scales in different axes)


## K-NN : Speeding Up the Classification

- Sophisticated algorithms for NN search:
- Classical problem in Comp. Geometry
- k-D trees
- Removing the samples from the training class $\mathcal{T}$ which do not change the result of classification
- Exactly: using Voronoi diagram
- Approximately: E.g. use Gabriel graph instead of Voronoi
- Condensation algorithm: iterative, also approximate.


## K-d Tree

k-d tree decomposition for the point set $(2,3),(5,4),(9,6),(4,7),(8,1),(7,2)$



## Condensation Algorithm

Input: The training set $\mathcal{T}$.

## Algorithm

1. Create two lists, $A$ and $B$. Insert a randomly selected sample from $\mathcal{T}$ to $A$. Insert the rest of the training samples to $B$.
2. Classify samples from $B$ using 1 NN with training set $A$. If an $x \in B$ is mis-classified, move it from $B$ to $A$.
3. If a move has been triggered in Step 2., goto Step 2.

Output: $A$ (the condensed training set for 1NN classification)

## Condensation Algorithm, Example



The training dataset


The dataset after the condensation. Shown with the new decision boundary.

## 1-NN Classification Error

Recall that a classification error $\bar{\epsilon}$ for strategy $q: X \rightarrow R$ is computed as

$$
\begin{equation*}
\bar{\epsilon}=\int \sum_{k: q(x) \neq k} p(x, k) \mathrm{d} x=\int \underbrace{\sum_{k: q(x) \neq k} p(k \mid x)}_{\epsilon(x)} p(x) \mathrm{d} x=\int \epsilon(x) p(x) \mathrm{d} x . \tag{21}
\end{equation*}
$$

We know that the Bayesian strategy $q_{B}$ decides for the highest posterior probability $q(x)=\operatorname{argmax}_{k} p(k \mid x)$, thus the partial error $\epsilon_{B}(x)$ for a given $x$ is

$$
\begin{equation*}
\epsilon_{B}(x)=1-\max _{k} p(k \mid x) . \tag{22}
\end{equation*}
$$

Assume the asymptotic case. We will show that the following bounds hold for the partial error $\epsilon_{N N}(x)$ and classification error $\bar{\epsilon}_{N N}$ in the 1-NN classification,

$$
\begin{array}{r}
\epsilon_{B}(x) \leq \epsilon_{N N}(x) \leq 2 \epsilon_{B}(x)-\frac{R}{R-1} \epsilon_{B}^{2}(x), \\
\bar{\epsilon}_{B} \leq \bar{\epsilon}_{N N} \leq 2 \bar{\epsilon}_{B}-\frac{R}{R-1} \bar{\epsilon}_{B}^{2}, \tag{24}
\end{array}
$$

where $\bar{\epsilon}_{B}$ is the Bayes classification error and $R$ is the number of classes.

## 1-NN Classification Error, Example (1)



Consider two distributions as shown, a small interval $\delta$ on an $x$-axis, and a point $s \in \delta$. Let the class priors be $p(1)=p(2)=0.5$. Assume $\delta \rightarrow 0$ and number of samples $N \rightarrow \infty$.

Observe the following:

$$
\begin{align*}
& p(1 \mid s)=0.8, \quad p(2 \mid s)=0.2  \tag{25}\\
& p(N N=1 \mid s)=p(1 \mid s)=0.8, \quad p(N N=2 \mid s)=p(2 \mid s)=0.2, \tag{26}
\end{align*}
$$

where $p(N N=k \mid s)$ is the probability that the 1 -NN of $s$ is from class $k(k=1,2)$ and thus $s$ is classified as $k$.

## 1-NN Classification Error, Example (2)



The error $\epsilon_{N N}(s)$ at $s$ is

$$
\begin{align*}
\epsilon_{N N}(s) & =p(1 \mid s) p(N N=2 \mid s)+p(2 \mid s) p(N N=1 \mid s)  \tag{27}\\
& =1-p(1 \mid s) p(N N=1 \mid s)-p(2 \mid s) p(N N=2 \mid s)  \tag{28}\\
& =1-p^{2}(1 \mid s)-p^{2}(2 \mid s) . \tag{29}
\end{align*}
$$

Generally, for $R$ classes, the error will be

$$
\begin{equation*}
\epsilon_{N N}(s)=1-\sum_{k \in R} p^{2}(k \mid s) . \tag{30}
\end{equation*}
$$

## 1-NN Classification Error, Example (3)

The two distributions and the partial errors (the Bayesian error $\epsilon_{B}(x)$ and the 1-NN error $\epsilon_{N N}(x)$ )


## 1-NN Classification Error Bounds (1)

Let us now return to the inequalities and prove them:

$$
\begin{equation*}
\epsilon_{B}(x) \leq \epsilon_{N N}(x) \leq 2 \epsilon_{B}(x)-\frac{R}{R-1} \epsilon_{B}^{2}(x) \tag{31}
\end{equation*}
$$

The first inequality follows from the fact that Bayes strategies are optimal.
To prove the second inequality, let $P(x)$ denote the maximum posterior for $x$ :

$$
\begin{align*}
& P(x)=\max _{k} p(k \mid x)  \tag{32}\\
\Rightarrow \quad & \epsilon_{B}(x)=1-P(x) . \tag{33}
\end{align*}
$$

Let us rewrite the partial error $\epsilon_{N N}(x)$ using the Bayesian entities $P(x)$ and $q(x)$ :

$$
\begin{equation*}
\epsilon_{N N}(x)=1-\sum_{k \in R} p^{2}(k \mid x)=1-P^{2}(x)-\sum_{k \neq q(x)} p^{2}(k \mid x) . \tag{34}
\end{equation*}
$$

We know that $p(q(x) \mid x)=P(x)$, but the remaining posteriors can be arbitrary. Let us consider the worst case. i.e. set $p(k \mid x)$ for $k \neq q(x)$ such that Eq. (34) is maximized. This will provide the upper bound.

## 1-NN Classification Error Bounds (2)

There are the following constraints on $p(k \mid x)(k \neq q(x))$ :

$$
\begin{align*}
& \sum_{k \neq q(x)} p(k \mid x)+P(x)=1 \quad \text { (posteriors sum to } 1 \text { ) }  \tag{35}\\
& \sum_{k \neq q(x)} p^{2}(k \mid x) \rightarrow \min \tag{36}
\end{align*}
$$

It is easy to show that this optimization problem is solved by setting all the posteriors to the same number. Thus,

$$
\begin{equation*}
p(k \mid x)=\frac{1-P(x)}{R-1}=\frac{\epsilon_{B}(x)}{R-1} \quad(k \neq q(x)) \tag{37}
\end{equation*}
$$

The upper bound can then be rewritten in terms of the Bayes partial error $\epsilon_{B}(x)=1-P(x)$ :

$$
\begin{equation*}
\epsilon_{N N}(x) \leq 1-P^{2}(x)-\sum_{k \neq q(x)} p^{2}(k \mid x)=1-\left(1-\epsilon_{B}(x)\right)^{2}-(R-1) \frac{\epsilon_{B}^{2}(x)}{(R-1)^{2}} \tag{38}
\end{equation*}
$$

## 1-NN Classification Error Bounds (3)

$$
\begin{equation*}
\epsilon_{N N}(x) \leq 1-P^{2}(x)-\sum_{k \neq q(x)} p^{2}(k \mid x)=1-\left(1-\epsilon_{B}(x)\right)^{2}-\frac{\epsilon_{B}^{2}(x)}{R-1} \tag{39}
\end{equation*}
$$

After expanding this, we get

$$
\begin{align*}
\epsilon_{N N}(x) & \leq 1-\left(1-\epsilon_{B}(x)\right)^{2}-\frac{\epsilon_{B}^{2}(x)}{(R-1)}  \tag{40}\\
& =1-1+2 \epsilon_{B}(x)-\epsilon_{B}^{2}(x)-\epsilon_{B}^{2}(x) \frac{R}{R-1}  \tag{41}\\
& =2 \epsilon_{B}(x)-\epsilon_{B}^{2}(x) \frac{R}{R-1} \tag{42}
\end{align*}
$$

Note that for $R=2$, the bound is tight because using $\epsilon_{B}(x)=1-P(x)$ in Eq. (39) gives

$$
\begin{equation*}
\epsilon_{N N}(x) \leq 1-P^{2}(x)-\frac{(1-P(x))^{2}}{1}=\epsilon_{N N}(x) \tag{43}
\end{equation*}
$$

## 1-NN Classification Error Bounds (4)

The inequality for the local errors has been proven:

$$
\begin{equation*}
\epsilon_{N N}(x) \leq 2 \epsilon_{B}(x)-\epsilon_{B}^{2}(x) \frac{R}{R-1} \tag{44}
\end{equation*}
$$

Is there a similar upper bound for the classification error $\bar{\epsilon}_{N N}=\int \epsilon_{N N}(x) p(x) \mathrm{d} x$, based on the Bayes error $\bar{\epsilon}_{B}=\int \epsilon_{B}(x) p(x) \mathrm{d} x$ ?

Multiplying Eq. (45) by $p(x)$, and integrating, gives

$$
\begin{equation*}
\bar{\epsilon}_{N N} \leq 2 \bar{\epsilon}_{B}(x)-\frac{R}{R-1} \int \epsilon_{B}^{2}(x) p(x) \mathrm{d} x \tag{45}
\end{equation*}
$$

Let us use the known identity and inequality (where $E(\cdot)$ is the expectation operator)

$$
\begin{equation*}
\operatorname{var}(x)=E\left(x^{2}\right)-E^{2}(x), \operatorname{var}(x) \geq 0 \quad \Rightarrow \quad E\left(x^{2}\right) \geq E^{2}(x) \tag{46}
\end{equation*}
$$

Thus, $\int \epsilon_{B}^{2}(x) p(x) \mathrm{d} x \geq\left(\int \epsilon_{B}(x) p(x) \mathrm{d} x\right)^{2}$, and

$$
\begin{equation*}
\bar{\epsilon}_{N N} \leq 2 \bar{\epsilon}_{B}(x)-\frac{R}{R-1} \int \epsilon_{B}^{2}(x) p(x) \mathrm{d} x \leq 2 \bar{\epsilon}_{B}(x)-\frac{R}{R-1} \bar{\epsilon}_{B}^{2} \tag{47}
\end{equation*}
$$

## K-NN Classification Error Bound

It can be shown that for $K-\mathrm{NN}$, the following inequality holds:

$$
\begin{equation*}
\bar{\epsilon}_{K N N} \leq \bar{\epsilon}_{B}+\bar{\epsilon}_{1 N N} / \sqrt{K \mathrm{const}} \tag{48}
\end{equation*}
$$

## Edit algorithm

The primary goal of this method is to reduce the classification error (not the speed-up of classification.)

Input: The training set $\mathcal{T}$.

## Algorithm

1. Partition $\mathcal{T}$ to two sets, $A$ and $B(\mathcal{T}=A \cup B, A \cap B=\emptyset$.)
2. Classify samples in $B$ using $\mathbf{K}-\mathrm{NN}$ with training set $A$. Remove all samples from $B$ which have been mis-classified.

Output: $B$ the training set for 1 -NN classification.
Asymptotic property:

$$
\begin{equation*}
\bar{\epsilon}_{e d i t}=\bar{\epsilon}_{B} \frac{1-\bar{\epsilon}_{B}}{1-\bar{\epsilon}_{K N N}} \tag{49}
\end{equation*}
$$

If $\bar{\epsilon}_{K N N}$ is small (e.g. 0.05) then the edited 1 NN is quasi-Bayes (almost the same performance as Bayesian Classification.)

