HOCGTC

# NRC Publications Archive Archives des publications du CNRC 

Theoretical Analyses of Cross-Validation Error and Voting in InstanceBased Learning<br>Turney, Peter

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

## Publisher's version / Version de l'éditeur:

Journal of Experimental and Theoretical Artificial Intelligence (JETAI), 6, 1994

NRC Publications Record / Notice d'Archives des publications de CNRC:<br>https://hrc-publications.canada.ca/eng/view/object/?id=82cd15e5-9759-403c-b307-f0ea24165970<br>https://publications-cnrc.canada.ca/fra/voir/objet/?id=82cd15e5-9759-403c-b307-f0ea24165970

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at https://nrc-publications.canada.ca/eng/copyright
READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site
https://publications-cnrc.canada.ca/fra/droits
LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

# Theoretical Analyses of Cross-Validation Error and Voting in Instance-Based Learning 

Peter Turney<br>Knowledge Systems Laboratory Institute for Information Technology<br>National Research Council Canada<br>Ottawa, Ontario, Canada<br>K1A 0R6

613-993-8564
peter@ai.iit.nrc.ca

Running Head: Error and Voting

# Theoretical Analyses of Cross-Validation Error and Voting in Instance-Based Learning 


#### Abstract

This paper begins with a general theory of error in cross-validation testing of algorithms for supervised learning from examples. It is assumed that the examples are described by attribute-value pairs, where the values are symbolic. Cross-validation requires a set of training examples and a set of testing examples. The value of the attribute that is to be predicted is known to the learner in the training set, but unknown in the testing set. The theory demonstrates that cross-validation error has two components: error on the training set (inaccuracy) and sensitivity to noise (instability).

This general theory is then applied to voting in instance-based learning. Given an example in the testing set, a typical instance-based learning algorithm predicts the designated attribute by voting among the $k$ nearest neighbors (the $k$ most similar examples) to the testing example in the training set. Voting is intended to increase the stability (resistance to noise) of instance-based learning, but a theoretical analysis shows that there are circumstances in which voting can be destabilizing. The theory suggests ways to minimize cross-validation error, by insuring that voting is stable and does not adversely affect accuracy.


## 1 Introduction

This paper is concerned with cross-validation testing of algorithms for supervised learning from examples. It is assumed that the examples are described by attribute-value pairs, where the attributes can have a finite number of symbolic values. The learning task is to predict the value of one of the attributes, given the value of the remaining attributes. It is assumed that each example is described by the same set of attributes.

Without loss of generality, we may assume that the attributes are restricted to boolean values. Let us suppose that each example has $r+1$ attributes. We may think of an example as a boolean vector in the space $\{0,1\}^{r+1}$. The learning task is to construct a function from $\{0,1\}^{r}$ to $\{0,1\}$, where $\{0,1\}^{r}$ is the space of the predictor attributes and $\{0,1\}$ is the space of the attribute that is to be predicted.

Cross-validation testing requires a set of training examples and a set of testing
examples. The learning algorithm (the student) is scored by a teacher, according to the student's performance on the testing set. The teacher knows the values of all of the attributes of all of the examples in the training set and the testing set. The student knows all of the values in the training set, but the student does not know the values of the attribute that is to be predicted in the testing set. The student uses the training set to develop a model of the data. The student then uses the model to make predictions for the testing set. The teacher compares the student's predictions with the actual values in the testing set. A mismatch between the student's prediction and the actual value is counted as an error. The student's goal is to minimize the number of errors that it makes on the testing set.

In a recent paper (Turney, 1993), a general theory of error in cross-validation testing of algorithms for predicting real-valued attributes is presented. ${ }^{1}$ Section 2 of this paper extends the theory of Turney (1993) to algorithms for predicting boolean-valued attributes. This section shows that cross-validation error has two components: error on the training set and sensitivity to noise. Error on the training set is commonly used as a measure of accuracy (Fraser, 1976). Turney (1990) introduces a formal definition of stability as a measure of the sensitivity of algorithms to noise. Section 2 proves that crossvalidation error is bounded by the sum of the training set error and the instability. The optimal cross-validation error is established, and it is proven that the strategy of minimizing error on the training set (maximizing accuracy) is sub-optimal.

Section 3 examines the cross-validation error of a simple form of instance-based learning (Aha et al., 1991; Kibler et al., 1989). Instance-based learning is not a single learning algorithm; it is a paradigm for a class of learning algorithms. It is related to the nearest neighbor pattern classification paradigm (Dasarathy, 1991). In instance-based learning, the student's model of the data consists of simply storing the training set. Given an example from the testing set, the student makes a prediction by looking for similar examples in the training set. Section 3 proves that this simple form of instance-based learning produces sub-optimal cross-validation error.

Looking for the single most similar example is sub-optimal because it is overly sensitive to noise in the data. Section 4 deals with instance-based learning algorithms that look for the $k$ most similar examples, where $k \geq 1$. In general, the $k$ most similar examples will not all agree on the value of the attribute that is to be predicted. This section examines algorithms that resolve this conflict by voting (Dasarathy, 1991). For example, suppose $k=3$ and the value of the attribute that is to be predicted is 1 for two of the three examples and 0 for the remaining example. If the algorithm uses majority voting, then the prediction is that the value is 1 .

Section 4 presents a detailed theoretical analysis of voting in instance-based learning. One motivation for voting is the desire to make instance-based learning more stable (more resistant to noise in the data). This section examines the stability of voting in the best case, the worst case, and the average case. It is shown that, in the worst case, voting can be destabilizing. That is, in certain circumstances, voting can actually increase the sensitivity to noise.

Section 5 discusses the practical application of this theory. It is shown how it is possible to estimate the expected cross-validation error from the training set. This section presents estimators that can indicate whether voting will increase or decrease stability for a particular set of data. This gives a method for choosing the value of $k$ that will minimize the cross-validation error.

Section 6 compares the work presented here with related work. The most closely related work is Turney (1993), which first introduced many of the concepts used here. There are interesting differences, which arise because Turney (1993) involves real-valued attributes and classes, while this paper involves boolean-valued attributes and classes. This theory is closely related to Akaike Information Criterion statistics (Sakamoto et al., 1986), as is discussed elsewhere (Turney, 1993). There is also an interesting connection with some prior work in nearest neighbor pattern classification (Cover \& Hart, 1967).

Finally, Section 7 considers future work. One weakness of this general theory of crossvalidation error is that it does not model interpolation and extrapolation. Another area for future research is applying the theory to problems other than voting in instance-based learning.

## 2 Cross-Validation Error

This section presents a general theory of error in cross-validation testing of algorithms for predicting symbolic attributes. In order to make the exposition simpler, the discussion is restricted to boolean-valued attributes. It is not difficult to extend the results presented here to $n$-valued attributes, where $n$ is any integer larger than one.

### 2.1 Exclusive-Or

This paper makes extensive use of the boolean exclusive-or operator. Suppose $x$ and $y$ are boolean variables. That is, $x$ and $y$ range over the set $\{0,1\}$. We may write $x \oplus y$ for " $x$ exclusive-or $y$ ". The expression $x \oplus y$ has the value 1 if and only if exactly one of $x$ and $y$ has the value 1 . Here are some of the properties of exclusive-or:

$$
\begin{gather*}
x \oplus 0=x  \tag{1}\\
x \oplus 1=\neg x  \tag{2}\\
x \oplus y=y \oplus x  \tag{3}\\
x \oplus x=0  \tag{4}\\
x \oplus \neg x=1  \tag{5}\\
x \oplus(y \oplus z)=(x \oplus y) \oplus z \tag{6}
\end{gather*}
$$

These properties can easily be verified with a truth-table.
The following results use exclusive-or with boolean vectors. Suppose $\vec{x}$ and $\vec{y}$ are boolean vectors in the space $\{0,1\}^{n}$. The expression $\vec{x} \oplus \vec{y}$ represents the vector that results from applying exclusive-or to corresponding elements of $\vec{x}$ and $\vec{y}$ :

$$
\begin{align*}
& \vec{z}=\vec{x} \oplus \vec{y}  \tag{7}\\
& \vec{x}=\left[\begin{array}{llll}
x_{1} & \ldots & x_{n}
\end{array}\right]  \tag{8}\\
& \vec{y}=\left[\begin{array}{lll}
y_{1} & \ldots & y_{n}
\end{array}\right]  \tag{9}\\
& \stackrel{\rightharpoonup}{z}=\left[\begin{array}{llll}
x_{1} \oplus y_{1} & \ldots & x_{n} \oplus y_{n}
\end{array}\right] \tag{10}
\end{align*}
$$

### 2.2 Accuracy and Stability

Suppose we have a black box with $r$ inputs and one output, where the inputs and the output can be represented by boolean values. Let the boolean vector $\vec{v}$ represent the inputs to the black box:

$$
\vec{v}=\left[\begin{array}{lll}
x_{1} & \ldots & x_{r} \tag{11}
\end{array}\right]
$$

Let $y$ represent the output of the black box. Suppose that the black box has a deterministic component $f$ and a random component $z$. The deterministic component $f$ is a function that maps from $\{0,1\}^{r}$ to $\{0,1\}$. The random component $z$ is a random boolean variable. The probability that $z$ is 1 is $p$. The probability that $z$ is 0 is $1-p$. The black box is represented by the equation:

$$
\begin{equation*}
y=f(\vec{v}) \oplus z \tag{12}
\end{equation*}
$$

The variable $p$ is a probability in the range $[0,1]$.

If $p$ is close to 0 , then $z$ is usually 0 , so the deterministic component dominates and the output $y$ is usually $f(\vec{v}$ ) (see equation (1) above). If $p$ is close to 1 , then $z$ is usually 1 , so the output is usually $\neg f(\vec{v})$ (see (2)). We may think of $p$ as the probability that $f(\vec{v})$ will be randomly negated. If $p$ is 0.5 , then the random component $z$ completely hides the deterministic component $f$. When $p$ is between 0.5 and $1.0, f(\vec{v})$ will be negated more often than not. When we know that $p$ is greater than 0.5 , we can negate the output of the black box:

$$
\begin{equation*}
y=\neg(f(\vec{v}) \oplus z) \tag{13}
\end{equation*}
$$

This counteracts the expected negation of $f(\vec{v})$ by the random component $z$. In (12) the output $y$ is $f(\vec{v})$ with probability $1-p$. In (13) the output $y$ is $f(\vec{v})$ with probability $p$.

Suppose we perform a series of $n$ experiments with the black box. We may represent the inputs in the $n$ experiments with a matrix $X$ :

$$
X=\left[\begin{array}{ccc}
x_{1,1} & \ldots & x_{1, r}  \tag{14}\\
\ldots & x_{i, j} & \ldots \\
x_{n, 1} & \ldots & x_{n, r}
\end{array}\right]
$$

Let the $i$-th row of the matrix $X$ be represented by the vector $\vec{v}_{i}$, where:

$$
\begin{equation*}
\vec{v}_{i}=\left[x_{i, 1} \ldots x_{i, r}\right] \quad i=1, \ldots, n \tag{15}
\end{equation*}
$$

The vector $\vec{v}_{i}$ contains the values of the $r$ inputs for the $i$-th experiment. Let the $j$-th column of the matrix $X$ be represented by the vector $\vec{x}_{j}$, where:

$$
\vec{x}_{j}=\left[\begin{array}{c}
x_{1, j}  \tag{16}\\
\ldots \\
x_{n, j}
\end{array}\right] \quad j=1, \ldots, r
$$

The vector $\vec{x}_{j}$ contains the values of the $j$-th input for the $n$ experiments. Let the $n$ outputs be represented by the vector $\vec{y}$, where:

$$
\vec{y}=\left[\begin{array}{l}
y_{1}  \tag{17}\\
\ldots \\
y_{n}
\end{array}\right]
$$

The scalar $y_{i}$ is the output of the black box for the $i$-th experiment.

The function $f$ can be extended to a vector function $\dot{f}(X)$, where:

$$
\dot{f}(X)=\left[\begin{array}{c}
f\left(\vec{v}_{1}\right)  \tag{18}\\
\ldots \\
f\left(\vec{v}_{n}\right)
\end{array}\right]
$$

Our model for the $n$ experiments is:

$$
\begin{equation*}
\vec{y}=\hat{f}(X) \oplus \vec{z} \tag{19}
\end{equation*}
$$

The vector $\vec{z}$ is a sequence of $n$ independent random boolean variables, each having the value 1 with probability $p$ :

$$
\overrightarrow{\mathbf{z}}=\left[\begin{array}{l}
z_{1}  \tag{20}\\
\ldots \\
z_{n}
\end{array}\right]
$$

Let us suppose that we are trying to develop a model of $f$. We may write $m(\vec{v} \mid X, \vec{y})$ to represent the model's prediction for $f(\vec{v})$, given that the model is based on the data $X$ and $\vec{y}$. We can extend $m(\vec{v} \mid X, \vec{y})$ to a vector function:

$$
\vec{m}(X \mid X, \vec{y})=\left[\begin{array}{c}
m\left(\vec{v}_{1} \mid X, \vec{y}\right)  \tag{21}\\
\ldots \\
m\left(\vec{v}_{n} \mid X, \vec{y}\right)
\end{array}\right]
$$

Thus $\vec{m}(X \mid X, \vec{y})$ represents the model's prediction for $f(X)$, given the data $X$ and $\vec{y}$.
Suppose we repeat the whole sequence of $n$ experiments, holding the inputs $X$ constant:

$$
\begin{align*}
& \vec{y}_{1}=\dot{f}(X) \oplus \vec{z}_{1}  \tag{22}\\
& \vec{y}_{2}=\dot{f}(X) \oplus \vec{z}_{2} \tag{23}
\end{align*}
$$

The outputs of the first set of $n$ experiments are represented by $\vec{y}_{1}$ and the outputs of the second set of $n$ experiments are represented by $\vec{y}_{2}$ :

$$
\vec{y}_{1}=\left[\begin{array}{c}
y_{1,1}  \tag{24}\\
\ldots \\
y_{1, n}
\end{array}\right] \quad \vec{y}_{2}=\left[\begin{array}{c}
y_{2,1} \\
\ldots \\
y_{2, n}
\end{array}\right]
$$

Although the inputs $X$ are the same, the outputs may have changed, due to the random component. That is, assuming $p$ is neither 0 nor 1 , it is possible that $\vec{z}_{1} \neq \vec{z}_{2}$ :

$$
\vec{z}_{1}=\left[\begin{array}{c}
z_{1,1}  \tag{25}\\
\ldots \\
z_{1, n}
\end{array}\right] \quad \vec{z}_{2}=\left[\begin{array}{c}
z_{2,1} \\
\ldots \\
z_{2, n}
\end{array}\right]
$$

Let the data $\left(X, \dot{y}_{1}\right)$ be the training set and let the data $\left(X, \vec{y}_{2}\right)$ be the testing set in crossvalidation testing of the model $m$. The cross-validation error vector $\vec{e}_{c}$ is:

$$
\begin{gather*}
\vec{e}_{c}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{2}  \tag{26}\\
\vec{e}_{c}=\left[\begin{array}{c}
e_{c, 1} \\
\ldots \\
e_{c, n}
\end{array}\right]=\left[\begin{array}{c}
m\left(\vec{v}_{1} \mid X, \vec{y}_{1}\right) \oplus y_{2,1} \\
\ldots \\
m\left(\vec{v}_{n} \mid X, \vec{y}_{1}\right) \oplus y_{2, n}
\end{array}\right] \tag{27}
\end{gather*}
$$

$e_{c, i}$ is 1 if and only if the prediction of the model $m\left(\vec{v}_{i} \mid X, \vec{y}_{1}\right)$ differs from the output $y_{2, i}$ in the testing data. We may write $\left\|\vec{e}_{c}\right\|$ for the length of the cross-validation error vector:

$$
\begin{equation*}
\left\|\stackrel{e}{c}_{c}\right\|=\sum_{i=1}^{n} e_{c, i} \tag{28}
\end{equation*}
$$

In other words, $\left\|\vec{e}_{c}\right\|$ is the number of errors that the model $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ makes on the testing set.

The model $m\left(\vec{v} \mid X, \vec{y}_{1}\right)$ has been tuned to perform well on the training set $\left(X, \vec{y}_{1}\right)$. Therefore the number of errors made on the training set may be deceptively low. To get a good indication of the quality of the model, we must test it on an independent set of data. Some authors call this error measure "cross-validation error", while other authors call it "train-and-test error" (Weiss and Kulikowski, 1991). In this paper, the former terminology is used. ${ }^{2}$

It is assumed that our goal is to minimize the expected number of errors $E\left(\left\|\vec{e}_{c}\right\|\right)$ that the model makes on the testing set. $E(\ldots)$ is the expectation operator from probability theory (Fraser, 1976). If $t(x)$ is a function of a random variable $x$, where $x \in S$ ( $S$ is the set of possible values of $x$ ) and the probability of observing a particular value of $x$ is $p(x)$, then $E(t(x))$ is:

$$
\begin{equation*}
E(t(x))=\sum_{x \in S} t(x) p(x) \tag{29}
\end{equation*}
$$

The expected cross-validation error $E\left(\left\|\vec{e}_{c}\right\|\right)$ depends on the random boolean vectors $\vec{z}_{1}$ and $\vec{z}_{2}$.

The next theorem gives some justification for the assumption that $X$ is the same in the training set $\left(X, \dot{y}_{1}\right)$ and the testing set $\left(X, \dot{y}_{2}\right)$.

Theorem 1: Let $D$ be an arbitrary probability distribution on $\{0,1\}^{r}$. Let the $n$ rows $\vec{v}_{1}, \ldots, \vec{v}_{n}$ of $X$ be independently randomly selected according to the distribution $D$. Let the vector $\vec{v}$ be randomly selected according to the distribution $D$. Let $p^{*}$ be the probability that $\vec{v}$ is not equal to any of $\vec{v}_{1}, \ldots, \vec{v}_{n}$. Then:

$$
\begin{equation*}
E\left(p^{*}\right) \leq\left(1-2^{-r}\right)^{n} \tag{30}
\end{equation*}
$$

Proof: Let $D(\vec{v})$ be the probability of randomly selecting $\vec{v}$ with the distribution $D$. The probability that $\vec{v}$ is not equal to any of $\vec{v}_{1}, \ldots, \vec{v}_{n}$ is thus:

$$
\begin{equation*}
p^{*}=(1-D(\vec{v}))^{n} \tag{31}
\end{equation*}
$$

The expected value of $p^{*}$ is:

$$
\begin{align*}
& E\left(p^{*}\right)=\sum_{\vec{v} \in\{0,1\}^{r}} p^{*} D(\vec{v})  \tag{32}\\
= & \sum_{\vec{v} \in\{0,1\}^{r}}(1-D(\vec{v}))^{n} D(\vec{v}) \tag{33}
\end{align*}
$$

The expected value is at its maximum when $D$ is the uniform distribution on $\{0,1\}^{r}$. For the uniform distribution, $D(\vec{v})=2^{-r}$. Therefore:

$$
\begin{align*}
E\left(p^{*}\right) & \leq \sum_{\vec{v} \in\{0,1\}^{r}}\left(1-2^{-r}\right)^{n} 2^{-r}  \tag{34}\\
& =2^{r}\left(1-2^{-r}\right)^{n} 2^{-r}  \tag{35}\\
& =\left(1-2^{-r}\right)^{n} \tag{36}
\end{align*}
$$

The implication of Theorem 1 is that, as $n$ increases, the probability that the inputs in the training set and the testing set are the same approaches 1 . Thus the assumption that $X$ is the same in the training set and the testing set is reasonable for large values of $n$.

Although it is assumed that the inputs are boolean, only Theorem 1 uses this assumption; none of the following results depend on the assumption that the inputs are boolean. The inputs could just as well be real numbers, so that $f$ maps from $\mathfrak{R}^{r}$ to $\{0,1\}$. However, it is necessary to assume that the output is boolean. If the output is real-valued, then we may turn to the analysis in Turney (1993).

Cover and Hart (1967) prove the following Lemma (their symbols have been changed to be consistent with the notation used here):

Lemma 1: Let $\vec{v}$ and $\vec{v}_{1}, \ldots, \vec{v}_{n}$ be independent identically distributed random variables taking values in a separable metric space $S$, with metric $d$. Let $\vec{v}_{i}$ be the nearest neighbor, according to the metric $d$, to $\vec{v}$ in the set $\vec{v}_{1}, \ldots, \vec{v}_{n}$. Then, as $n$ approaches infinity, $d\left(\vec{v}^{\prime}, \vec{v}_{i}\right)$ approaches 0 , with probability 1 .

Proof: See Cover and Hart (1967) $\square$.
Cover and Hart (1967) assume that $n$ is large and the probability distribution for $P(f(\vec{v}) \oplus z=1)$ is a continuous function of $\vec{v}$. It follows from Lemma 1 that, as $n$
approaches infinity, $P(f(\vec{v}) \oplus z=1)$ approaches $P\left(f\left(\vec{v}_{i}\right) \oplus z=1\right)$. This is similar to Theorem 1 here, except that $S$ is continuous. Note that, if the probability distribution for $P(f(\vec{v}) \oplus z=1)$ is a continuous function of $\vec{v}$, then $z$ must be a function of $\vec{v}, z(\vec{v})$, in order to compensate for the fact that $f(\vec{v})$ is discontinuous. Thus the assumption that the distribution is continuous is relatively strong.

For the rest of this paper, let us assume that $f$ maps from $\{0,1\}^{r}$ to $\{0,1\}$. We may assume that $X$ is the same in the training set and the testing set. This assumption may be justified with Theorem 1. Alternatively, we could assume that $f$ maps from $\mathfrak{R}^{r}$ to $\{0,1\}$ and use Lemma 1 to justify the assumption that $X$ is the same in the training set and the testing set. However, the analysis is simpler when $z$ is not a function of $\vec{v}$.

The error on the training set $\vec{e}_{t}$ is:

$$
\begin{gather*}
\vec{e}_{t}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{1}  \tag{37}\\
\vec{e}_{t}=\left[\begin{array}{c}
e_{t, 1} \\
\ldots \\
e_{t, n}
\end{array}\right]=\left[\begin{array}{c}
m\left(\vec{v}_{1} \mid X, \vec{y}_{1}\right) \oplus y_{1,1} \\
\ldots \\
m\left(\vec{v}_{n} \mid X, \dot{y}_{1}\right) \oplus y_{1, n}
\end{array}\right] \tag{38}
\end{gather*}
$$

$e_{t, i}$ is 1 if and only if the prediction of the model $m\left(\vec{v}_{i} \mid X, \vec{y}_{1}\right)$ differs from the output $y_{1, i}$ in the training data.

There is another form of error that is of interest. We may call this error instability:

$$
\begin{gather*}
\vec{e}_{s}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right)  \tag{39}\\
\vec{e}_{s}=\left[\begin{array}{c}
e_{s, 1} \\
\ldots \\
e_{s, n}
\end{array}\right]=\left[\begin{array}{c}
m\left(\vec{v}_{1} \mid X, \vec{y}_{1}\right) \oplus m\left(\vec{v}_{1} \mid X, \vec{y}_{2}\right) \\
\cdots \\
m\left(\vec{v}_{n} \mid X, \vec{y}_{1}\right) \oplus m\left(\vec{v}_{n} \mid X, \vec{y}_{2}\right)
\end{array}\right] \tag{40}
\end{gather*}
$$

Instability is a measure of the sensitivity of the model to noise in the data. If the model resists noise, then $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ will be relatively similar to $\vec{m}\left(X \mid X, \vec{y}_{2}\right)$, so $E\left(\left\|\vec{e}_{s}\right\|\right)$ will be relatively small. If the model is sensitive to noise, then $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ will be relatively dissimilar from $\vec{m}\left(X \mid X, \vec{y}_{2}\right)$, so $E\left(\left\|\vec{e}_{s}\right\|\right)$ will be relatively large.

The following theorem shows the relationship between cross-validation error, error on the training set, and instability.

Theorem 2: The expected cross-validation error is less than or equal to the sum of the expected error on the training set and the expected instability:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right) \leq E\left(\left\|\vec{e}_{t}\right\|\right)+E\left(\left\|\vec{e}_{s}\right\|\right) \tag{41}
\end{equation*}
$$

Proof: Let us introduce a new term $\vec{e}_{\omega}$ :

$$
\begin{equation*}
\vec{e}_{\omega}=\vec{m}\left(X \mid X, \vec{y}_{2}\right) \oplus \vec{y}_{1} \tag{42}
\end{equation*}
$$

Due to the symmetry of the training set (22) and the testing set (23), we have (from (26) and (42)):

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{\omega}\right\|\right) \tag{43}
\end{equation*}
$$

Recall the definition of $\vec{e}_{t}(37)$ :

$$
\begin{equation*}
\vec{e}_{t}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{1} \tag{44}
\end{equation*}
$$

It follows from this definition that $\vec{y}_{1}$ differs from $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ in $\left\|\vec{e}_{t}\right\|$ locations. Recall the definition of $\vec{e}_{s}(39)$ :

$$
\begin{equation*}
\vec{e}_{s}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right) \tag{45}
\end{equation*}
$$

It follows from this definition that $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ differs from $\vec{m}\left(X \mid X, \vec{y}_{2}\right)$ in $\left\|\vec{e}_{s}\right\|$ locations. Therefore $\vec{y}_{1}$ differs from $\vec{m}\left(X \mid X, \vec{y}_{2}\right)$ in at most $\left\|\vec{e}_{t}\right\|+\left\|\vec{e}_{s}\right\|$ locations. Thus:

$$
\begin{equation*}
\left\|\vec{e}_{\omega}\right\| \leq\left\|\vec{e}_{t}\right\|+\left\|\vec{e}_{s}\right\| \tag{46}
\end{equation*}
$$

Finally:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{\omega}\right\|\right) \leq E\left(\left\|\vec{e}_{t}\right\|+\left\|\vec{e}_{s}\right\|\right)=E\left(\left\|\vec{e}_{t}\right\|\right)+E\left(\left\|\vec{e}_{s}\right\|\right) \tag{47}
\end{equation*}
$$

The next theorem is a variation on Theorem 2.
Theorem 3: If $\vec{e}_{t}$ and $\vec{e}_{s}$ are statistically independent, then:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{t}\right\|\right)+E\left(\left\|\vec{e}_{s}\right\|\right)-\frac{2}{n} E\left(\left\|\vec{e}_{t}\right\|\right) E\left(\left\|\vec{e}_{s}\right\|\right) \tag{48}
\end{equation*}
$$

Proof: Assume that $\vec{e}_{t}$ and $\vec{e}_{s}$ are statistically independent. Let us introduce a new term
$\vec{e}_{\omega}:$

$$
\begin{equation*}
\vec{e}_{\omega}=\vec{m}\left(X \mid X, \vec{y}_{2}\right) \oplus \vec{y}_{1} \tag{49}
\end{equation*}
$$

We see that:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{\omega}\right\|\right) \tag{50}
\end{equation*}
$$

Recall the definition of $\vec{e}_{t}$ (37):

$$
\begin{equation*}
\vec{e}_{t}=\vec{m}\left(X \mid X, \stackrel{y}{y}_{1}\right) \oplus \vec{y}_{1} \tag{51}
\end{equation*}
$$

Recall the definition of $\vec{e}_{s}$ (39):

$$
\begin{equation*}
\vec{e}_{s}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right) \tag{52}
\end{equation*}
$$

We have:

$$
\begin{gather*}
\vec{e}_{t} \oplus \vec{e}_{s}=\left(\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{1}\right) \oplus\left(\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right)\right)  \tag{53}\\
=\left(\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{1}\right)\right) \oplus\left(\vec{y}_{1} \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right)\right)  \tag{54}\\
=\vec{y}_{1} \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right)  \tag{55}\\
=\vec{e}_{\omega} \tag{56}
\end{gather*}
$$

Let us introduce the following terms:

$$
\begin{align*}
p_{t} & =E\left(\left\|\vec{e}_{t}\right\|\right) / n  \tag{57}\\
p_{s} & =E\left(\left\|\vec{e}_{s}\right\|\right) / n  \tag{58}\\
p_{\omega} & =E\left(\left\|\vec{e}_{\omega}\right\|\right) / n \tag{59}
\end{align*}
$$

Since $\vec{e}_{t}$ and $\vec{e}_{s}$ are statistically independent and $\vec{e}_{\omega}=\vec{e}_{t} \oplus \vec{e}_{s}$, we have:

$$
\begin{equation*}
p_{\omega}=p_{t}\left(1-p_{s}\right)+p_{s}\left(1-p_{t}\right)=p_{t}+p_{s}-2 p_{t} p_{s} \tag{60}
\end{equation*}
$$

Therefore:

$$
\begin{gather*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{\omega}\right\|\right)  \tag{61}\\
=n p_{\omega}  \tag{62}\\
=n\left(p_{t}+p_{s}-2 p_{t} p_{s}\right)  \tag{63}\\
=E\left(\left\|\vec{e}_{t}\right\|\right)+E\left(\left\|\vec{e}_{s}\right\|\right)-\frac{2}{n} E\left(\left\|\vec{e}_{t}\right\|\right) E\left(\left\|\vec{e}_{s}\right\|\right) \tag{64}
\end{gather*}
$$

The assumptions of Theorem 3 are stronger than the assumptions of Theorem 2, but Theorem 3 gives a closer relationship between $E\left(\left\|\vec{e}_{c}\right\|\right), E\left(\left\|\vec{e}_{t}\right\|\right)$, and $E\left(\left\|\vec{e}_{s}\right\|\right)$.

The next theorem shows the optimal expected cross-validation error.
Theorem 4: Suppose that $f$ and $p$ are known to the modeler. To minimize the expected cross-validation error, the modeler should set the model as follows:

$$
\vec{m}\left(X \mid X, \stackrel{\rightharpoonup}{y}_{i}\right)=\left[\begin{array}{cc}
\stackrel{\rightharpoonup}{f}(X) & \text { if } p \leq 0.5  \tag{65}\\
\stackrel{\rightharpoonup}{f}^{f}(X) & \text { if } p>0.5
\end{array}\right]
$$

With this model, we have:

$$
\begin{gather*}
E\left(\left\|\vec{e}_{s}\right\|\right)=0  \tag{66}\\
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{t}\right\|\right)=\left[\begin{array}{cc}
n p & \text { if } p \leq 0.5 \\
n(1-p) & \text { if } p>0.5
\end{array}\right] \tag{67}
\end{gather*}
$$

Proof: Suppose that $p \leq 0.5$. By (4), (39), and (65):

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=E(\|\hat{f}(X) \oplus \mathfrak{f}(X)\|)=0 \tag{68}
\end{equation*}
$$

Thus the model is perfectly stable. This is natural, since the model is not based on the data; it is based on the a priori knowledge of $f$. Consider the cross-validation error:

$$
\begin{gather*}
\stackrel{\rightharpoonup}{e}_{c}=\dot{f}(X) \oplus \vec{y}_{2}  \tag{69}\\
=\stackrel{\grave{\prime}}{f}(X) \oplus\left(\dot{f}(X) \oplus \vec{z}_{2}\right)  \tag{70}\\
=(\dot{f} f(X) \oplus \dot{f}(X)) \oplus \vec{z}_{2}  \tag{71}\\
=\vec{z}_{2} \tag{72}
\end{gather*}
$$

Thus:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{z}_{2}\right\|\right)=E\left(\sum_{i=1}^{n} z_{2, i}\right)=\sum_{i=1}^{n} E\left(z_{2, i}\right)=\sum_{i=1}^{n} p=n p \tag{73}
\end{equation*}
$$

Consider the error on the training set:

$$
\begin{equation*}
\vec{e}_{t}=f(X) \oplus \vec{y}_{1}=\vec{z}_{1} \tag{74}
\end{equation*}
$$

Therefore:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{t}\right\|\right)=E\left(\left\|\vec{z}_{1}\right\|\right)=E\left(\left\|\vec{z}_{2}\right\|\right)=E\left(\left\|\vec{e}_{c}\right\|\right)=n p \tag{75}
\end{equation*}
$$

Now, suppose that $p>0.5$. For instability, we have:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=E\left(\left\|\neg \stackrel{\jmath}{f}(X) \oplus \neg^{\prime}(X)\right\|\right)=0 \tag{76}
\end{equation*}
$$

Consider the cross-validation error:

$$
\begin{gather*}
\vec{e}_{c}=\neg \dot{f}(X) \oplus \vec{y}_{2}  \tag{77}\\
=\neg \dot{f}(X) \oplus\left(\grave{f}(X) \oplus \vec{z}_{2}\right)  \tag{78}\\
=(\neg \dot{f}(X) \oplus \dot{f}(X)) \oplus \vec{z}_{2}  \tag{79}\\
=\neg \vec{z}_{2} \tag{80}
\end{gather*}
$$

Thus:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\neg \vec{z}_{2}\right\|\right)=E\left(\sum_{i=1}^{n} \neg z_{2, i}\right)=\sum_{i=1}^{n} E\left(\neg z_{2, i}\right)=n(1-p) \tag{81}
\end{equation*}
$$

Consider the error on the training set:

$$
\begin{equation*}
\vec{e}_{t}=\neg f(X) \oplus \vec{y}_{1}=\neg \vec{z}_{1} \tag{82}
\end{equation*}
$$

Therefore:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{t}\right\|\right)=E\left(\left\|\neg \vec{z}_{1}\right\|\right)=E\left(\left\|\neg \vec{z}_{2}\right\|\right)=E\left(\left\|\vec{e}_{c}\right\|\right)=n(1-p) \tag{83}
\end{equation*}
$$

It is clear that no model can have a lower $E\left(\left\|\vec{e}_{c}\right\|\right)$ than this model, since $\vec{z}_{2}$ cannot be predicted $\square$.

The next theorem considers models that minimize the error on the training set.
Theorem 5: Let our model be as follows:

$$
\begin{equation*}
\vec{m}\left(X \mid X, \vec{y}_{i}\right)=\vec{y}_{i} \tag{84}
\end{equation*}
$$

Then we have:

$$
\begin{gather*}
E\left(\left\|\vec{e}_{t}\right\|\right)=0  \tag{85}\\
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{s}\right\|\right)=2 n p-2 n p^{2} \tag{86}
\end{gather*}
$$

Proof: Consider the expected error on the training set:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{t}\right\|\right)=E\left(\left\|\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{1}\right\|\right)=E\left(\left\|\vec{y}_{1} \oplus \vec{y}_{1}\right\|\right)=0 \tag{87}
\end{equation*}
$$

Consider the cross-validation error:

$$
\begin{gather*}
\vec{e}_{c}=\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{y}_{2}  \tag{88}\\
=\vec{y}_{1} \oplus \vec{y}_{2}  \tag{89}\\
=\left(\dot{f}(X) \oplus \vec{z}_{1}\right) \oplus\left(\stackrel{\rightharpoonup}{f}(X) \oplus \vec{z}_{2}\right)  \tag{90}\\
=(\stackrel{\rightharpoonup}{f}(X) \oplus \stackrel{\rightharpoonup}{f}(X)) \oplus\left(\vec{z}_{1} \oplus \vec{z}_{2}\right)  \tag{91}\\
=\vec{z}_{1} \oplus \vec{z}_{2} \tag{92}
\end{gather*}
$$

Therefore:

$$
\begin{gather*}
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{z}_{1} \oplus \vec{z}_{2}\right\|\right)  \tag{93}\\
=E\left(\sum_{i=1}^{n} z_{1, i} \oplus z_{2, i}\right)  \tag{94}\\
=\sum_{i=1}^{n} E\left(z_{1, i} \oplus z_{2, i}\right)  \tag{95}\\
=2 n p-2 n p^{2} \tag{96}
\end{gather*}
$$

Consider the expected instability:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=E\left(\left\|\vec{m}\left(X \mid X, \vec{y}_{1}\right) \oplus \vec{m}\left(X \mid X, \vec{y}_{2}\right)\right\|\right)=E\left(\left\|\vec{y}_{1} \oplus \vec{y}_{2}\right\|\right)=E\left(\left\|\vec{e}_{c}\right\|\right) \tag{97}
\end{equation*}
$$

These theorems are variations on theorems that first appeared in Turney (1993). Theorems 2, 3, 4, and 5 in this paper correspond to Theorems 1, 12, 2, and 3 in Turney (1993). The difference is that the theorems in Turney (1993) cover real-valued variables, while the theorems here cover boolean-valued variables. A comparison will show that there are interesting contrasts between the continuous case and the discrete case.

We shall call the model of Theorem $4 m_{\alpha}$ :

$$
\vec{m}_{\alpha}\left(X \mid X, \vec{y}_{i}\right)=\left[\begin{array}{cc}
\stackrel{\rightharpoonup}{f}(X) & \text { if } p \leq 0.5  \tag{98}\\
\vec{\prime}(X) & \text { if } p>0.5
\end{array}\right]
$$

We shall call the model of Theorem $5 m_{\beta}$ :

$$
\begin{equation*}
\vec{m}_{\beta}\left(X \mid X, \vec{y}_{i}\right)=\vec{y}_{i} \tag{99}
\end{equation*}
$$

Let us compare $E\left(\left\|\vec{e}_{c}\right\|\right)$ for $m_{\alpha}$ and $m_{\beta}$. Let $p_{\alpha}$ be $E\left(\left\|\vec{e}_{c}\right\|\right) / n$ for $m_{\alpha}$ :

$$
p_{\alpha}=\left[\begin{array}{cl}
p & \text { if } p \leq 0.5  \tag{100}\\
(1-p) & \text { if } p>0.5
\end{array}\right]
$$

Let $p_{\beta}$ be $E\left(\left\|\vec{e}_{c}\right\|\right) / n$ for $m_{\beta}$ :

$$
\begin{equation*}
p_{\beta}=2 p-2 p^{2} \tag{101}
\end{equation*}
$$

Figure 1 is a plot of $E\left(\left\|\vec{e}_{c}\right\|\right) / n$ as a function of $p$, for $m_{\alpha}$ and $m_{\beta}$. $m_{\alpha}$ has lower expected cross-validation error than $m_{\beta}$, except at the points $0,0.5$, and 1 , where $m_{\alpha}$ has the same expected cross-validation error as $m_{\beta}$. When $p$ equals 0 or 1 , there is no noise in the black box, so $m_{\alpha}$ and $m_{\beta}$ have the same expected cross-validation error. When $p$ equals 0.5 , the noise completely hides the deterministic component $f$, so again $m_{\alpha}$ and $m_{\beta}$ have the same expected cross-validation error. In general, when $p$ is neither $0,0.5$, nor 1 , models that minimize error on the training set (such as $m_{\beta}$ ) will give sub-optimal cross-validation error.


Figure 1. Plot of $E\left(\left\|\vec{e}_{c}\right\|\right) / n$ as a function of $p$.

The next theorem shows the relation between $p_{\alpha}$ and $p_{\beta}$.

Theorem 6: For $0 \leq p \leq 1$ :

$$
\begin{equation*}
p_{\beta}=2 p_{\alpha}-2 p_{\alpha}^{2} \tag{102}
\end{equation*}
$$

Proof: When $0 \leq p \leq 0.5, p_{\alpha}=p$. Therefore we have:

$$
\begin{equation*}
p_{\beta}=2 p_{\alpha}-2 p_{\alpha}^{2} \tag{103}
\end{equation*}
$$

When $0.5<p \leq 1, p_{\alpha}=1-p$. Therefore we have:

$$
\begin{gather*}
2 p_{\alpha}-2 p_{\alpha}^{2}=2(1-p)-2(1-p)^{2}  \tag{104}\\
=2-2 p-2+4 p-2 p^{2}  \tag{105}\\
=p_{\beta} \tag{106}
\end{gather*}
$$

Theorem 6 shows that, for small values of $p, p_{\beta}$ is approximately twice the optimal $p_{\alpha}$.
At this point, it may be worthwhile to summarize the assumptions that are made in this theory:

1. It is assumed that the inputs (the features; the attributes; the matrix $X$ ) are the same in the training set $\left(X, \vec{y}_{1}\right)$ and the testing set $\left(X, \vec{y}_{2}\right)$. This assumption is the weakest element in this theory. It is discussed in detail earlier in this section and also in Section 7. It is also discussed in more depth in Turney (1993). Note that the outputs $\vec{y}_{1}$ and $\vec{y}_{2}$ are not the same in the training and testing sets.
2. It is assumed that the inputs are boolean-valued. This assumption is only used in Theorem 1. The purpose of Theorem 1 is to increase the credibility of the above assumption 1. Theorem 1 can easily be adapted to the case of multi-valued symbolic attributes. Lemma 1 covers the case of real-valued attributes. Thus there is nothing essential about the assumption that the inputs are boolean-valued.
3. It is assumed that the noise vector $\vec{z}$ is a sequence of $n$ independent random boolean variables, each having the value 1 with probability $p$. That is, $\vec{z}$ is a sequence of samples from a $\operatorname{Bernoulli}(p)$ distribution (Fraser, 1976). This is a very weak assumption, which is likely to be (approximately) satisfied in most real-world data (given assumption 6 below).
4. It is assumed that there is noise in the class attribute (assumption 3 above), but noise in the input attributes is not addressed (due to assumption 1 above).
5. It is assumed that the data consist of a deterministic component $f$ and a random component $z$. This assumption is expressed in the model $y=f(\vec{v}) \oplus z$. One implication of this assumption is that the target concept $f$ does not drift with time or shift with context.
6. It is assumed that the output $y$ (the class attribute) is boolean-valued. The theory can easily be extended to handle multi-valued symbolic class attributes. Turney (1993) discusses real-valued outputs (and real-valued noise vectors).

With the exception of the first assumption, these assumptions are relatively weak. ${ }^{3}$
This section has presented some general results that apply to any algorithm for predicting boolean-valued attributes. The rest of this paper focuses on a particular class of algorithms: instance-based learning algorithms. This section has presented a general theory and the remainder of the paper will demonstrate that the theory can be fruitfully applied to a real, concrete machine learning algorithm.

## 3 Single Nearest Neighbor

Instance-based learning may be used either for predicting boolean-valued attributes (Aha et al., 1991) or for predicting real-valued attributes (Kibler et al., 1989). Instance-based learning is a paradigm for a class of learning algorithms; it is not a single algorithm. It is related to the nearest neighbor pattern recognition paradigm (Dasarathy, 1991).

With instance-based learning, the model $m(\vec{v} \mid X, \vec{y})$ is constructed by simply storing the data $(X, \vec{y})$. These stored data are the instances. In order to make a prediction for the input $\vec{v}$, we examine the row vectors $\vec{v}_{1}, \ldots, \vec{v}_{n}$ of the matrix $X$. In the simplest version of instance-based learning, we look for the row vector $\vec{v}_{i}$ that is most similar to the input $\vec{v}$. The prediction for the output is $m(\vec{v} \mid X, \vec{y})=y_{i}$, the element of $\vec{y}$ that corresponds to the row vector $\vec{v}_{i}$.

There are many ways that one might choose to measure the similarity between two vectors. It is assumed only that we are using a reasonable measure of similarity. Let us say that a similarity measure $\operatorname{sim}\left(\vec{u}_{1}, \vec{u}_{2}\right)$ is reasonable if:

$$
\begin{equation*}
\vec{u}_{1} \neq \vec{u}_{2} \rightarrow \operatorname{sim}\left(\vec{u}_{1}, \vec{u}_{2}\right)<\operatorname{sim}\left(\vec{u}_{1}, \vec{u}_{1}\right) \tag{107}
\end{equation*}
$$

That is, the similarity between distinct vectors is always less than the similarity between identical vectors.

The $k$ row vectors in $X$ that are most similar to the input $\vec{v}$ are called the $k$ nearest neighbors of $\vec{v}$ (Dasarathy, 1991). We may use $m_{k}$ to represent instance-based learning with $k$ nearest neighbors. This section focuses on $k=1$.

The next theorem shows that instance-based learning $m_{1}$ using the single nearest neighbor gives sub-optimal cross-validation error.

Theorem 7: Let the model $m_{1}$ use instance-based learning with the single nearest neighbor. That is, if the row vector $\vec{v}_{i}$ is the most similar to the input $\vec{v}$ of all the row vectors $\vec{v}_{1}, \ldots, \vec{v}_{n}$ of the matrix $X$, then $m_{1}(\vec{v} \mid X, \vec{y})=y_{i}$. If $m_{1}$ uses a reasonable measure of similarity and no two rows in $X$ are identical, then $\vec{m}_{1}(X \mid X, \vec{y})=\vec{y}$.

Proof: Since $m_{1}$ is based on a reasonable measure of similarity and no two rows in $X$ are identical, no vector is more similar to $\vec{v}_{i}$ than $\vec{v}_{i}$ itself, so $m_{1}\left(\vec{v}_{i} \mid X, \vec{y}\right)=y_{i}$. Thus it follows from (21) that $\vec{m}(X \mid X, \vec{y})=\vec{y} \square$.

Theorem 7 shows that $m_{1}$ satisfies the assumptions of Theorem 5. Therefore $m_{1}$ has the following properties:

$$
\begin{gather*}
E\left(\left\|\vec{e}_{t}\right\|\right)=0  \tag{108}\\
E\left(\left\|\vec{e}_{c}\right\|\right)=E\left(\left\|\vec{e}_{s}\right\|\right)=2 n p-2 n p^{2} \tag{109}
\end{gather*}
$$

In other words, $m_{1}$ is equivalent to $m_{\beta}$. We know from Section 2 that $m_{1}$ gives suboptimal cross-validation error.

If there are two (or more) identical rows in $X$, then (109) still holds, but (108) is not necessarily true. $m_{1}$ will, in general, no longer be equivalent to $m_{\beta}$. The proof is a small variation on Theorem 5.

The problem with $m_{1}$ is that it is unstable. It is natural to consider increasing the stability of instance-based learning by considering $k>1$.

## 4 Voting

This section examines $k \geq 1$. Suppose the model $m_{k}$ is given the input vector $\vec{v}$. Let $\vec{v}_{1}, \ldots, \vec{v}_{k}$ be the $k$ nearest neighbors to $\vec{v}$. That is, let $\vec{v}_{1}, \ldots, \vec{v}_{k}$ be the $k$ row vectors in $X$ that are most similar to the input vector $\vec{v}$. Let us assume a reasonable measure of similarity. Let $y_{1}, \ldots, y_{k}$ be the outputs corresponding to the rows $\vec{v}_{1}, \ldots, \vec{v}_{k}$. The model $m_{k}$ predicts the output for the input $\vec{v}$ to be the value of the majority of $y_{1}, \ldots, y_{k}$ (Fix \& Hodges, 1951):

$$
m_{k}(\vec{v} \mid X, \vec{y})=\left[\begin{array}{cc}
0 & \text { if }  \tag{110}\\
\sum_{i=1}^{k} y_{i}<k / 2 \\
1 & \text { if } \\
\sum_{i=1}^{k} y_{i}>k / 2
\end{array}\right]
$$

Let us assume that $k$ is an odd number, $k=3,5,7, \ldots$, so there will never be a tie in majority voting with $y_{1}, \ldots, y_{k}$. Of course, we require that $k \leq n$. In general, $k<n$.

More generally, let us consider the following model (Tomek, 1976):

$$
m_{k}(\vec{v} \mid X, \vec{y})=\left[\begin{array}{cc}
0 & \text { if }  \tag{111}\\
\sum_{i=1}^{k} y_{i}<t k \\
1 & \text { if } \sum_{i=1}^{k} y_{i}>t k
\end{array}\right]
$$

In this model, $t$ is a threshold, such that $0<t<1$. With majority voting, $t=0.5$. When $t \neq 0.5$, it may be appropriate to consider values of $k$ other than the odd numbers.

The motivation for voting is the belief that it will increase stability; that it will make the model more resistant to noise in the data. The following sections give a formal analysis of the stability of instance-based learning with voting. They examine the best case, the worst case, and the average case.

### 4.1 Best Case

The next theorem concerns the best case, when stability is maximal.
Theorem 8: Let $N_{k}\left(\vec{v}_{i}\right)$ be the set of the $k$ nearest neighbors $\vec{v}_{1}, \ldots, \vec{v}_{k}$ to $\vec{v}_{i}$. Let $\rho_{i}$ be
defined as follows:

$$
\begin{equation*}
\rho_{i}=\frac{1}{k}\left(\sum_{\vec{v}_{j} \in N_{k}\left(\vec{v}_{i}\right)} f\left(\vec{v}_{j}\right)\right) \quad i=1, \ldots, n \tag{112}
\end{equation*}
$$

That is, $\rho_{i}$ is the frequency with which $f$ takes the value 1 , in the neighborhood $N_{k}\left(\vec{v}_{i}\right)$ of $\vec{v}_{i}$. Assume that we are using majority voting, $t=0.5$, and that $k=3,5,7, \ldots$ Suppose that:

$$
\begin{equation*}
(\forall i)\left(\left(\rho_{i}=0\right) \vee\left(\rho_{i}=1\right)\right) \tag{113}
\end{equation*}
$$

This is the most stable situation; other values of $\rho_{i}$ are less stable. $E\left(\left\|\vec{e}_{s}\right\|\right)$ for $m_{k}$ is determined by the following formula:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=2 n P_{k}-2 n P_{k}^{2} \tag{114}
\end{equation*}
$$

In this formula, $P_{k}$ depends on $m_{k} . P_{k}$ has the following values for $m_{1}, m_{3}, m_{5}$, and $m_{\infty}$ :

$$
\begin{gather*}
P_{1}=p  \tag{115}\\
P_{3}=3 p^{2}-2 p^{3}  \tag{116}\\
P_{5}=6 p^{5}-15 p^{4}+10 p^{3}  \tag{117}\\
P_{\infty}=\left[\begin{array}{cc}
0 & \text { if } p<0.5 \\
0.5 & \text { if } p=0.5 \\
1 & \text { if } p>0.5
\end{array}\right] \tag{118}
\end{gather*}
$$

Proof: Let $\vec{v}_{1}, \ldots, \vec{v}_{k}$ be the $k$ nearest neighbors to $\vec{v}$. Let $y_{1,1}, \ldots, y_{1, k}$ be the corresponding elements in $\vec{y}_{1}$. Let $y_{2,1}, \ldots, y_{2, k}$ be the corresponding elements in $\vec{y}_{2}$. The stability of $m_{k}$ is determined by the probability that:

$$
\begin{equation*}
\sum_{i=1}^{k} y_{1, i}<k / 2 \quad \text { and } \quad \sum_{i=1}^{k} y_{2, i}>k / 2 \tag{119}
\end{equation*}
$$

If we can find this probability for each row vector in $X$, then we can find the stability of $m_{k}$. Suppose that $p_{i}$ is the probability of (119) when the input to $m_{k}$ is $v_{i}$. Then $2 p_{i}$ is the probability that:

$$
\begin{equation*}
m_{k}\left(\vec{v}_{i} \mid X, \vec{y}_{1}\right) \oplus m_{k}\left(\vec{v}_{i} \mid X, \vec{y}_{2}\right)=1 \tag{120}
\end{equation*}
$$

Therefore:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=\sum_{i=1}^{n} E\left(m_{k}\left(\vec{v}_{i} \mid X, \vec{y}_{1}\right) \oplus m_{k}\left(\vec{v}_{i} \mid X, \vec{y}_{2}\right)\right)=\sum_{i=1}^{n} 2 p_{i} \tag{121}
\end{equation*}
$$

The best case - the case with minimal expected instability $E\left(\left\|\vec{e}_{s}\right\|\right)$ - arises when $\rho_{i}=0$ or $\rho_{i}=1$. Without loss of generality, we may assume that $\rho_{i}=0$. In this case, we have:

$$
\begin{equation*}
y_{i, j}=f\left(\vec{v}_{j}\right) \oplus z_{i, j}=0 \oplus z_{i, j}=z_{i, j} \quad i=1,2 \quad j=1, \ldots, k \tag{122}
\end{equation*}
$$

Therefore (119) becomes:

$$
\begin{equation*}
\sum_{i=1}^{k} z_{1, i}<k / 2 \quad \text { and } \quad \sum_{i=1}^{k} z_{2, i}>k / 2 \tag{123}
\end{equation*}
$$

Let us introduce the term $P_{k}$ :

$$
\begin{equation*}
P_{k}=P\left(\sum_{i=1}^{k} z_{1, i}>k / 2\right) \tag{124}
\end{equation*}
$$

That is, $P_{k}$ is the probability that the majority of $y_{1,1}, \ldots, y_{1, k}$ are 1 , given that $f\left(\vec{v}_{1}\right), \ldots, f\left(\vec{v}_{k}\right)$ are all 0 . The probability of (119) is:

$$
\begin{equation*}
P_{k}\left(1-P_{k}\right) \tag{125}
\end{equation*}
$$

If the best case holds for every row vector in $X$ - that is, if (113) is true - then we have:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=\sum_{i=1}^{n} 2 P_{k}\left(1-P_{k}\right)=2 n P_{k}\left(1-P_{k}\right) \tag{126}
\end{equation*}
$$

Let us consider the case $k=3$. We have:

$$
\begin{equation*}
P_{3}=p^{3}+3 p^{2}(1-p)=p^{3}+3 p^{2}-3 p^{3}=3 p^{2}-2 p^{3} \tag{127}
\end{equation*}
$$

Recall that for $k=1$ we have (109):

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=2 n p-2 n p^{2} \tag{128}
\end{equation*}
$$

Thus $P_{1}=p$. For $k=5$, we have:

$$
\begin{equation*}
P_{5}=6 p^{5}-15 p^{4}+10 p^{3} \tag{129}
\end{equation*}
$$

To find the behavior of the model in the limit, $m_{\infty}$, we can use the Central Limit Theorem (Fraser, 1976). Consider the following sum:

$$
\begin{equation*}
\sum_{i=1}^{k} z_{1, i} \tag{130}
\end{equation*}
$$

The mean and variance of this sum are:

$$
\begin{gather*}
E\left(\sum_{i=1}^{k} z_{1, i}\right)=k p  \tag{131}\\
\operatorname{var}\left(\sum_{i=1}^{k} z_{1, i}\right)=k p(1-p) \tag{132}
\end{gather*}
$$

Consider the following expression:

$$
\begin{equation*}
\frac{\sum_{i=1}^{k} z_{1, i}-k p}{\sqrt{k p(1-p)}} \tag{133}
\end{equation*}
$$

By the Central Limit Theorem, the distribution of (133) approaches a standard normal distribution as $k$ approaches infinity. Recall the definition of $P_{k}$ :

$$
\begin{equation*}
P_{k}=P\left(\sum_{i=1}^{k} z_{1, i}>k / 2\right) \tag{134}
\end{equation*}
$$

Let $z_{\alpha}$ be a random variable with a standard normal distribution. As $k$ approaches infinity, $P_{k}$ approaches:

$$
\begin{equation*}
P\left(z_{\alpha}>\frac{(k / 2)-k p}{\sqrt{k p(1-p)}}\right) \tag{135}
\end{equation*}
$$

We see that:

$$
\lim _{k \rightarrow \infty}\left(\frac{(k / 2)-k p}{\sqrt{k p(1-p)}}\right)=\left[\begin{array}{r}
\infty  \tag{136}\\
\text { if } p<0.5 \\
0 \\
\text { if } p=0.5 \\
-\infty
\end{array} \text { if } p>0.5\right] ~
$$

Therefore:

$$
\left.\begin{array}{c}
P_{\infty}=\left[\begin{array}{cc}
0 & \text { if } p<0.5 \\
0.5 \text { if } p=0.5 \\
1 & \text { if } p>0.5
\end{array}\right] \\
\frac{E\left(\left\|\vec{e}_{s}\right\|\right)}{n}=2 P_{\infty}\left(1-P_{\infty}\right)=\left[\begin{array}{c}
0 \\
\text { if } p \neq 0.5 \\
0.5
\end{array} \text { if } p=0.5\right. \tag{138}
\end{array}\right] .
$$

Theorem 8 implies that, in the best case, $m_{k}$ grows increasingly stable as $k$ increases, unless $p=0.5$. Figure 2 is a plot of $E\left(\left\|\vec{e}_{S}\right\|\right) / n$ as a function of $p$, for the models $m_{1}, m_{3}$, $m_{5}$, and $m_{\infty}$.


Figure 2. A plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$.

When comparing Figure 1 with Figure 2, note that Figure 1 is a plot of $E\left(\left\|\vec{e}_{c}\right\|\right) / n$, while Figure 2 is a plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$. In general, we cannot plot $E\left(\left\|\vec{e}_{c}\right\|\right) / n$ for $m_{k}$, because $E\left(\left\|\vec{e}_{t}\right\|\right)$ will depend on the data. Theorems 2 and 3 show that $E\left(\left\|\vec{e}_{c}\right\|\right)$ depends on both $E\left(\left\|\vec{e}_{s}\right\|\right)$ and $E\left(\left\|\vec{e}_{t}\right\|\right)$. The exception is $m_{1}$, where we know that $E\left(\left\|\vec{e}_{t}\right\|\right)=0$.
(Recall that $m_{1}=m_{\beta}$. )
Theorem 8 does not necessarily mean that we should make $k$ as large as possible, even if we assume that the best case (113) holds. It is assumed that our goal is to minimize the expected cross-validation error. Theorems 2 and 3 show that the expected cross-validation error has two components, the error on the training set and the instability. In the best case, increasing $k$ will increase stability, but increasing $k$ is also likely to increase the error on the training set. We must find the value of $k$ that best balances the conflicting demands of accuracy (low error on the training set) and stability (resistance to noise). This value will, of course, depend on $f$ and $p$.

### 4.2 Worst Case

The next theorem concerns the worst case, when stability is minimal.
Theorem 9: Assume that we are using majority voting, $t=0.5$, and that $k=3,5,7, \ldots$. Suppose that:

$$
\begin{equation*}
(\forall i)\left[\left[\rho_{i}=\frac{1}{2}+\frac{1}{2 k}\right] \vee\left[\rho_{i}=\frac{1}{2}-\frac{1}{2 k}\right]\right] \tag{139}
\end{equation*}
$$

This is the least stable situation; other values of $\rho_{i}$ are more stable. $E\left(\left\|\vec{e}_{s}\right\|\right)$ for $m_{k}$ is determined by the following formula:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=2 n P_{k}-2 n P_{k}^{2} \tag{140}
\end{equation*}
$$

In this formula, $P_{k}$ depends on $m_{k} . P_{k}$ has the following values for $m_{1}, m_{3}, m_{5}$, and $m_{\infty}$ :

$$
\begin{gather*}
P_{1}=p  \tag{141}\\
P_{3}=2 p-3 p^{2}+2 p^{3}  \tag{142}\\
P_{5}=3 p-9 p^{2}+16 p^{3}-15 p^{4}+6 p^{5}  \tag{143}\\
P_{\infty}=\left[\begin{array}{rcc}
0 & \text { if } & p=0 \\
0.5 & \text { if } & 0<p<1 \\
1 & \text { if } & p=1
\end{array}\right] \tag{144}
\end{gather*}
$$

Proof: Let $\vec{v}_{1}, \ldots, \vec{v}_{k}$ be the $k$ nearest neighbors to $\vec{v}$. Let $y_{1,1}, \ldots, y_{1, k}$ be the corresponding elements in $\vec{y}_{1}$. Let $y_{2,1}, \ldots, y_{2, k}$ be the corresponding elements in $\vec{y}_{2}$. The stability of
$m_{k}$ is determined by the probability that:

$$
\begin{equation*}
\sum_{i=1}^{k} y_{1, i}<k / 2 \quad \text { and } \quad \sum_{i=1}^{k} y_{2, i}>k / 2 \tag{145}
\end{equation*}
$$

The worst case - the case with maximal expected instability $E\left(\left\|\vec{e}_{s}\right\|\right)$ - arises when:

$$
\begin{equation*}
\sum_{i=1}^{k} f\left(\stackrel{\rightharpoonup}{v}_{i}\right)=\frac{k+1}{2} \quad \text { or } \quad \sum_{i=1}^{k} f\left(\vec{v}_{i}\right)=\frac{k-1}{2} \tag{146}
\end{equation*}
$$

Without loss of generality, we may assume that:

$$
\begin{gather*}
s=\frac{k+1}{2}  \tag{147}\\
f\left(\vec{v}_{1}\right)=\ldots=f\left(\vec{v}_{s}\right)=1  \tag{148}\\
f\left(\vec{v}_{s+1}\right)=\ldots=f\left(\vec{v}_{k}\right)=0 \tag{149}
\end{gather*}
$$

We see that:

$$
\begin{gather*}
y_{i, j}=\neg z_{i, j}  \tag{150}\\
y_{i, j}=z_{i, j} \tag{151}
\end{gather*} \quad i=1,2 \quad j=1, \ldots, s, ~ j=s+1, \ldots, k
$$

Therefore:

$$
\begin{equation*}
P\left(\sum_{i=1}^{k} y_{1, i}>k / 2\right)=P\left(\sum_{i=1}^{s} z_{1, i} \leq \sum_{i=s+1}^{k} z_{1, i}\right) \tag{152}
\end{equation*}
$$

We may use the term $P_{k}$ to indicate the probability of (152). That is, $P_{k}$ is the probability that the majority of $y_{1,1}, \ldots, y_{1, k}$ are 1 , given (148) and (149). If the worst case holds for every row vector in $X$ - that is, (139) is true - then we have:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=\sum_{i=1}^{n} 2 P_{k}\left(1-P_{k}\right)=2 n P_{k}\left(1-P_{k}\right) \tag{153}
\end{equation*}
$$

Let us consider the case $k=3$. We have:

$$
\begin{equation*}
P_{3}=2 p-3 p^{2}+2 p^{3} \tag{154}
\end{equation*}
$$

For $k=5$, we have:

$$
\begin{equation*}
P_{5}=3 p-9 p^{2}+16 p^{3}-15 p^{4}+6 p^{5} \tag{155}
\end{equation*}
$$

Again, to find the behavior $m_{\infty}$ of the model in the limit, we can use the Central Limit Theorem. Let $z_{\alpha}$ and $z_{\beta}$ be independent random variables with a standard normal distribution. As $k$ approaches infinity, $P_{k}$ approaches:

$$
\begin{equation*}
P\left(z_{\alpha} \leq z_{\beta}\right) \tag{156}
\end{equation*}
$$

Therefore:

$$
\begin{gather*}
P_{\infty}=\left[\begin{array}{ccc}
0 & \text { if } & p=0 \\
0.5 & \text { if } & 0<p<1 \\
1 & \text { if } & p=1
\end{array}\right]  \tag{157}\\
\frac{E\left(\left\|\vec{e}_{s}\right\|\right)}{n}=2 P_{\infty}\left(1-P_{\infty}\right)=\left[\begin{array}{cl}
0 & \text { if }(p=0) \text { or }(p=1) \\
0.5 & \text { if } \\
0<p<1
\end{array}\right] \tag{158}
\end{gather*}
$$

Theorem 9 implies that, in the worst case, $m_{k}$ grows increasingly unstable as $k$ increases, unless $p=0$ or $p=1$. Figure 3 is a plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$, for the models $m_{1}, m_{3}, m_{5}$, and $m_{\infty}$.


Figure 3. A plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$.

### 4.3 Average Case

Sections 4.1 and 4.2 show that voting can be either very beneficial or very detrimental, depending on the relationship between the function $f$ and the voting threshold $t$. The function $f$ determines $\rho_{i}$, the frequency with which $f$ takes the value 1 , in the neighborhood $N_{k}\left(\vec{v}_{i}\right)$ of $\vec{v}_{i}$. In the previous sections, we assumed majority voting, $t=0.5$. When $\rho_{i}$ is near 0.5 , majority voting is detrimental. When $\rho_{i}$ is far from 0.5 , majority voting is beneficial. These results generalize to voting with other values of $t$. When $\rho_{i}$ is near $t$, voting is detrimental. When $\rho_{i}$ is far from $t$, voting is beneficial.

We would like to know how majority voting performs with an average $f$. For the average function $f$, is $\rho_{i}$ dangerously close to 0.5 , or is it usually safely far away? This is a difficult question: What exactly is an average $f$ ?

The next theorem uses a simple model of an average $f$.
Theorem 10: Assume that we are using majority voting, $t=0.5$. Let us suppose that we can treat $f\left(\vec{v}_{i}\right)$ as a random boolean variable, with probability $p^{\prime}$ that $f\left(\vec{v}_{i}\right)$ is 1 , and probability $1-p^{\prime}$ that $f\left(\vec{v}_{i}\right)$ is 0 . In the limit, as $k$ approaches infinity, $m_{k}$ tends to be stable, unless either $p=0.5$ or $p^{\prime}=0.5$.

Proof: We see that:

$$
\begin{equation*}
E\left(\sum_{i=1}^{k} f\left(\vec{v}_{i}\right)\right)=k p^{\prime} \tag{159}
\end{equation*}
$$

Recall that:

$$
\begin{equation*}
y_{i, j}=f\left(\stackrel{\rightharpoonup}{v}_{j}\right) \oplus z_{i, j} \quad i=1,2 \quad j=1, \ldots, k \tag{160}
\end{equation*}
$$

Therefore:

$$
\begin{align*}
& P\left(y_{i, j}=0\right)=\left(1-p^{\prime}\right)(1-p)+p^{\prime} p  \tag{161}\\
& P\left(y_{i, j}=1\right)=p\left(1-p^{\prime}\right)+p^{\prime}(1-p) \tag{162}
\end{align*}
$$

Thus:

$$
\begin{equation*}
E\left(\sum_{i=1}^{k} y_{1, i}\right)=k\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]=\mu \tag{163}
\end{equation*}
$$

$$
\begin{equation*}
\operatorname{var}\left(\sum_{i=1}^{k} y_{1, i}\right)=k\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]\left[\left(1-p^{\prime}\right)(1-p)+p^{\prime} p\right]=\sigma^{2} \tag{164}
\end{equation*}
$$

## Consider:

$$
\begin{equation*}
\frac{\sum_{i=1}^{k} y_{1, i}-\mu}{\sigma} \tag{165}
\end{equation*}
$$

By the Central Limit Theorem, as $k$ approaches infinity, (165) approaches a standard normal distribution. As before, we define $P_{k}$ as the probability:

$$
\begin{equation*}
P\left(\sum_{i=1}^{k} y_{1, i}>k / 2\right) \tag{166}
\end{equation*}
$$

Let $z_{\alpha}$ be a random variable with a standard normal distribution. As $k$ approaches infinity, $P_{k}$ approaches:

$$
\begin{equation*}
P\left(z_{\alpha}>((k / 2)-\mu) / \sigma\right) \tag{167}
\end{equation*}
$$

We see that:

$$
\lim _{k \rightarrow \infty}(((k / 2)-\mu) / \sigma)=\left[\begin{array}{ccc}
\infty & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]<0.5  \tag{168}\\
0 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]=0.5 \\
-\infty & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]>0.5
\end{array}\right]
$$

Therefore:

$$
\begin{gather*}
P_{\infty}=\left[\begin{array}{cc}
0 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]<0.5 \\
0.5 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]=0.5 \\
1 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]>0.5
\end{array}\right]  \tag{169}\\
\frac{E\left(\left\|\vec{e}_{s}\right\|\right)}{n}=2 P_{\infty}\left(1-P_{\infty}\right)=\left[\begin{array}{cc}
0 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right] \neq 0.5 \\
0.5 & \text { if }\left[p\left(1-p^{\prime}\right)+p^{\prime}(1-p)\right]=0.5
\end{array}\right] \tag{170}
\end{gather*}
$$

Thus $m_{k}$ tends to be stable, unless:

$$
\begin{equation*}
p\left(1-p^{\prime}\right)+p^{\prime}(1-p)=0.5 \tag{171}
\end{equation*}
$$

Some algebraic manipulation converts (171) to:

$$
\begin{equation*}
(2 p-1)\left(1-2 p^{\prime}\right)=0 \tag{172}
\end{equation*}
$$

Therefore $m_{k}$ tends to be stable, unless either $p=0.5$ or $p^{\prime}=0.5 \square$.
Let us now consider the general case (111), when the threshold for voting $t$ is not necessarily equal to 0.5 .

Theorem 11: Let $t$ be any value, such that $0<t<1$. Let us suppose that we can treat $f\left(\vec{v}_{i}\right)$ as a random boolean variable, with probability $p^{\prime}$ that $f\left(\vec{v}_{i}\right)$ is 1 , and probability $1-p^{\prime}$ that $f\left(\vec{v}_{i}\right)$ is 0 . Let us introduce the term $\tau$ :

$$
\begin{equation*}
\tau=p\left(1-p^{\prime}\right)+p^{\prime}(1-p)=p+p^{\prime}-2 p p^{\prime} \tag{173}
\end{equation*}
$$

In the limit, as $k$ approaches infinity, $m_{k}$ tends to be stable, unless $t=\tau$.
Proof: This is a straightforward extension of the reasoning in Theorem $10 \square$.
Presumably $p$ (the probability that $z$ is 1 ) and $p^{\prime}$ (the probability that $f$ is 1 ) are not under the control of the modeler. However, $t$ can be adjusted by the modeler. If the modeler can estimate $p$ and $p^{\prime}$, then $\tau$ can be estimated. Theorem 11 implies that the modeler should make sure that $t$ is relatively far from $\tau$.

The assumption that we can treat $f\left(\vec{v}_{i}\right)$ as a random boolean variable is somewhat unrealistic. Let us consider a more sophisticated analysis.

Theorem 12: Let us define $\tau_{i}$ as follows:

$$
\begin{equation*}
\tau_{i}=p+\rho_{i}-2 p \rho_{i} \quad i=1, \ldots, n \tag{174}
\end{equation*}
$$

In the limit, as $k$ approaches infinity, $m_{k}$ tends to be stable for $\vec{v}_{i}$, unless $t$ is close to $\tau_{i}$.
Proof: This is a straightforward extension of the reasoning in Theorem $10 \square$.
Let us define $p^{\prime}$ as the average of the $\rho_{i}$ :

$$
\begin{equation*}
p^{\prime}=\frac{1}{n} \sum_{i=1}^{n} \rho_{i} \tag{175}
\end{equation*}
$$

Note that:

$$
\begin{equation*}
p^{\prime} \neq \frac{1}{n} \sum_{i=1}^{n} f\left(\vec{v}_{i}\right) \tag{176}
\end{equation*}
$$

We can use (173) and (175) to calculate $\tau$. Theorem 12 implies that it is possible for $m_{k}$ to be stable even when $t=\tau$, if $t$ is far from each $\tau_{i}$.

Theorem 13: We can estimate $\tau_{i}$ from the data $\left(X, \vec{y}_{1}\right)$ as follows:

$$
\begin{equation*}
t_{i}=\frac{1}{k}\left(\sum_{\stackrel{\rightharpoonup}{v}_{j} \in N_{k}\left(\vec{v}_{i}\right)} y_{1, j}\right) \quad i=1, \ldots, n \tag{177}
\end{equation*}
$$

$t_{i}$ is an unbiased estimator for $\tau_{i}$.
Proof: We have:

$$
\begin{gather*}
E\left(t_{i}\right)=\frac{1}{k}\left(\sum_{\vec{v}_{j} \in N_{k}\left(\vec{v}_{i}\right)} E\left(y_{1, j}\right)\right)  \tag{178}\\
=\frac{1}{k}\left(\sum_{\vec{v}_{j} \in N_{k}\left(\vec{v}_{i}\right)} E\left(f\left(\vec{v}_{j}\right) \oplus z_{1, j}\right)\right)  \tag{179}\\
=p+\rho_{i}-2 p \rho_{i}  \tag{180}\\
=\tau_{i} \tag{181}
\end{gather*}
$$

Using (177), we can estimate $\tau_{i}$ for any $k$ and any $i . t_{i}$ is the frequency of the output 1 , in the neighborhood $N_{k}\left(\vec{v}_{i}\right)$ of $\vec{v}_{i}$.

Theorem 14: If we know the value of $p$, then we can also estimate $\rho_{i}$ from the data $\left(X, \vec{y}_{1}\right)$ :

$$
\begin{equation*}
r_{i}=\frac{t_{i}-p}{1-2 p} \tag{182}
\end{equation*}
$$

$r_{i}$ is an unbiased estimator for $\rho_{i}$.
Proof: We have:

$$
\begin{gather*}
E\left(r_{i}\right)=\frac{E\left(t_{i}\right)-p}{1-2 p}  \tag{183}\\
=\left(\tau_{i}-p\right) /(1-2 p)  \tag{184}\\
=\left(p+\rho_{i}-2 p \rho_{i}-p\right) /(1-2 p)  \tag{185}\\
=\rho_{i} \tag{186}
\end{gather*}
$$

Using (177) and (182), we can estimate $\rho_{i}$ for any $k$ and any $i$. Once we have an estimate of $\rho_{i}$, we can establish whether we are closer to the best case (Section 4.1) or the worst case (Section 4.2).

We now have some understanding of the average behavior of $m_{k}$ in the limit. Let us consider the average behavior of $m_{k}$ for $k=3$ and $t=0.5$.

Theorem 15: Suppose that $n$ is even. Let $k=3$ and $t=0.5$. Suppose that:

$$
\begin{array}{cccc}
\rho_{i}=1 / 3 & \text { or } & \rho_{i}=2 / 3 & i=1, \ldots, n / 2 \\
\rho_{i}=0 / 3 & \text { or } & \rho_{i}=3 / 3 & i=n / 2+1, \ldots, n \tag{188}
\end{array}
$$

Then:

$$
\begin{equation*}
E\left(\left\|\vec{e}_{s}\right\|\right)=n\left(2 p-4 p^{2}+12 p^{3}-26 p^{4}+24 p^{5}-8 p^{6}\right) \tag{189}
\end{equation*}
$$

Proof: We know from Section 4.2 that $m_{3}$ is unstable when:

$$
\begin{equation*}
\rho_{i}=1 / 3 \quad \text { or } \quad \rho_{i}=2 / 3 \quad i=1, \ldots, n \tag{190}
\end{equation*}
$$

We know from Section 4.1 that $m_{3}$ is stable when:

$$
\begin{equation*}
\rho_{i}=0 / 3 \quad \text { or } \quad \rho_{i}=3 / 3 \quad i=1, \ldots, n \tag{191}
\end{equation*}
$$

For the unstable case, $i=1, \ldots, n / 2$, we have (154):

$$
\begin{equation*}
P_{u, 3}=2 p-3 p^{2}+2 p^{3} \tag{192}
\end{equation*}
$$

For the stable case, $i=n / 2+1, \ldots, n$, we have (127):

$$
\begin{equation*}
P_{s, 3}=3 p^{2}-2 p^{3} \tag{193}
\end{equation*}
$$

Therefore (126):

$$
\begin{align*}
E\left(\left\|\vec{e}_{s}\right\|\right)= & \sum_{i=1}^{n / 2} 2 P_{u, 3}\left(1-P_{u, 3}\right)+\sum_{i=n / 2+1}^{n} 2 P_{s, 3}\left(1-P_{s, 3}\right)  \tag{194}\\
= & \frac{n}{2} 2 P_{u, 3}\left(1-P_{u, 3}\right)+\frac{n}{2} 2 P_{s, 3}\left(1-P_{s, 3}\right)  \tag{195}\\
& =n\left[P_{u, 3}-P_{u, 3}^{2}+P_{s, 3}-P_{s, 3}^{2}\right]  \tag{196}\\
= & n\left(2 p-4 p^{2}+12 p^{3}-26 p^{4}+24 p^{5}-8 p^{6}\right) \tag{197}
\end{align*}
$$

Figure 4 is a plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ for $m_{3}$, using the average case analysis of Theorem 15. For comparison, $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ is also plotted for $m_{1} . m_{3}$ is slightly more stable than $m_{1}$, except at $0,0.5$, and 1 , where they are equal.


Figure 4. A plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$.

## 5 Application

Let us assume that $n, p, X$, and $\vec{y}$ are fixed, while $k$ and $t$ are under our control. Our goal is to minimize the expected cross-validation error. The data $X$ and $\vec{y}$ are the training set, and the testing set is not yet available to us. How do we determine the best values for $k$ and $t$ ?

Using Theorems 2 or 3 , we can estimate $E\left(\left\|\vec{e}_{c}\right\|\right)$ if we can estimate $E\left(\left\|\vec{e}_{t}\right\|\right)$ and $E\left(\left\|\vec{e}_{s}\right\|\right)$.
Obtaining an estimate of $E\left(\left\|\vec{e}_{t}\right\|\right)$ is relatively simple. We can use the actual error on the training set $\left\|\vec{e}_{t}\right\|$ as an estimator for the expected error on the training set $E\left(\left\|\vec{e}_{t}\right\|\right)$. By the Law of Large Numbers (Fraser, 1976), as $n$ increases, $\left\|\vec{e}_{t}\right\| / n$ becomes increasingly close to $E\left(\left\|\vec{e}_{t}\right\|\right) / n$.

Obtaining an estimate of $E\left(\left\|\vec{e}_{s}\right\|\right)$ is more difficult. Suppose $k=3, t=0.5$, and the value of $p$ is known by the modeler. We can use Theorem 14 to estimate $\rho_{i}$ for $i=1, \ldots, n$. We can then use the method of Theorem 15 to obtain an estimate of $E\left(\left\|\vec{e}_{s}\right\|\right)$.

To estimate $\rho_{i}$, we need to know $p$, the level of noise in the black box. We can estimate $p$ empirically from the data (by assuming that $f$ is usually constant within a small neighborhood, for example), or we can theoretically analyze the internal mechanism of the black box. A mixed approach, combining theoretical and empirical analysis, may be fruitful.

This approach to estimating $E\left(\left\|\vec{e}_{s}\right\|\right)$ is feasible for $k=3$ or $k=5$, but the algebra becomes exponentially messy for larger values of $k$. For large values of $k$, we can use Theorem 13 to estimate $\tau_{i}$. Theorem 12 suggests that $E\left(\left\|\vec{e}_{s}\right\|\right)$ will tend toward 0 , as long as $\tau_{i}$ is far enough from $t$, for all $i$.

Once we have an estimate of $E\left(\left\|\vec{e}_{t}\right\|\right)$ and $E\left(\left\|\vec{e}_{s}\right\|\right)$, we can use Theorem 2 to set an upper bound on $E\left(\left\|\vec{e}_{c}\right\|\right)$. Alternatively, we can use Theorem 3 to estimate $E\left(\left\|\vec{e}_{c}\right\|\right)$. (It is conjectured that $m_{k}$ satisfies the assumption of Theorem 3, that $\vec{e}_{t}$ and $\vec{e}_{s}$ are statistically independent, but there is not yet a proof for this.) We can estimate $E\left(\left\|\vec{e}_{c}\right\|\right)$ for varying values of $k$ and $t$. We should then choose the values of $k$ and $t$ that minimize our estimated expected cross-validation error.

The average-case analysis of the previous section suggests that increasing $k$ will tend to increase stability. Increasing $k$ will also tend to increase training set error. If we plot cross-validation error as a function of $k$, then we expect to see a curve with a single minimum at the optimal value of $k$, somewhere between the extremes $k=1$ and $k=n$. Empirical tests of nearest neighbor algorithms report exactly this result (Dasarathy, 1991). I have not included any empirical tests with real-world data in this paper, because it would
merely duplicate work that has already been done. That is, the average-case analysis presented here is consistent with typical empirical experience.

The worst-case analysis, on the other hand, yields a surprising result, that increasing $k$ can actually decrease stability. In the worst case, if we plot cross-validation error as a function of $k$, then we may see a curve with a minimum at $k=1$, that steadily rises as $k$ increases. It is sometimes reported that cross-validation error is minimal at $k=1$, but it is usually assumed that this is due to rapid increase in training set error with increasing $k$, rather than decrease in stability with increasing $k$. It seems likely that the worst case, where stability decreases as $k$ increases, is very rare in real-world data. Thus it would be difficult to find real-world data that illustrates the worst-case analysis. Simulated data could be used, but there would be no point in such an exercise, since the results would be determined by the assumptions embedded in the simulation. Therefore this paper does not include any empirical results.

## 6 Related Work

This work is most closely related to Turney (1993), which presents a general theory of cross-validation error for algorithms that predict real-valued attributes. The theory is then applied to linear regression and instance-based learning. The theory of cross-validation error for algorithms that predict boolean-valued attributes, presented in Section 2, is similar to the theory for real-valued attributes (Turney, 1993), but there are some differences. Real-valued noise is quite different from boolean-valued noise. Boolean noise $z$ is a random variable in $\{0,1\}$, such that the probability that $z$ is 1 is $p$. Real noise $z$ is a random variable in $\Re$, such that $z$ has mean 0 and variance $\sigma^{2}$. Equivalently, the noise is $\sigma z$, where $z$ has mean 0 and variance 1 . The role of $p$ in boolean noise is similar to the role of $\sigma$ in real noise, except that $p$ must be in the range [0,1], but the range of $\sigma$ is $\mathfrak{R}^{+}$. This makes a substantial difference in the details of the two cases, real-valued and booleanvalued.

Turney (1993) also examined the cross-validation error of instance-based learning with $k$ nearest neighbors. When instance-based learning is used to predict real-values, the $k$ nearest neighbors are averaged together (Kibler et al., 1989). Averaging the $k$ nearest neighbors for real-valued attributes is analogous to voting for boolean-valued attributes. There is a surprising difference, however. Turney (1993) proves that averaging, in the best case, improves stability and, in the worst case, does not affect stability (Theorem 10 in (Turney, 1993)). Voting, in the worst case, can be destabilizing (Section 4.2). It is surpris-
ing to discover this break in the analogy between the real-valued case and the booleanvalued case.

Turney (1993) mentions that the theory of cross-validation error is similar to the work in Akaike Information Criterion (AIC) statistics (Sakamoto et al., 1986). The similarity stems from the idea of finding the best model $\vec{m}\left(X \mid X, \vec{y}_{1}\right)$ for one set of outputs $\vec{y}_{1}$, then evaluating the model's performance with a second set of outputs $\vec{y}_{2}$. This is the definition of the expected cross-validation error. Akaike uses the same approach to define mean expected log likelihood (Sakamoto et al., 1986). For a more thorough comparison of the approach described here with AIC, see Turney (1993).

There is an interesting connection between the work here and previous work in nearest neighbor pattern classification. Cover and Hart (1967) assume that $n$ is large and the probability distribution for $P(f(\vec{v}) \oplus z=1)$ is a continuous function of $\vec{v}$. Using Lemma 1, they prove the following Theorem:

$$
\begin{equation*}
R^{*} \leq R \leq R^{*}\left(2-M R^{*} /(M-1)\right) \tag{198}
\end{equation*}
$$

Here, $R$ is the probability of error of $m_{1}, R^{*}$ is the Bayes probability of error, and $M$ is the number of categories. The Bayes probability of error is the minimum probability of error over all decision rules taking underlying probability structure into account (Cover \& Hart, 1967). This paper assumes boolean-valued attributes, so $M=2$ :

$$
\begin{equation*}
R^{*} \leq R \leq 2 R-2\left(R^{*}\right)^{2} \tag{199}
\end{equation*}
$$

Compare this with Theorem 6:

$$
\begin{equation*}
p_{\beta}=2 p_{\alpha}-2 p_{\alpha}^{2} \tag{200}
\end{equation*}
$$

There are clearly some similarities between (199) and (200). However, there are also some differences. (200) is an exact equality, while (199) is an inequality. The reason for this difference is that Cover and Hart's (1967) noise, $z(\vec{v})$, is a function of $\vec{v}$, while the noise $z$ discussed here is not a function of $\vec{v}$. To remove the dependence on $\vec{v}$, Cover and Hart (1967) define $R$ and $R^{*}$ as the expected probabilities, where the expectation is taken over $\vec{v}$. This makes their proof more complex than mine.

Cover and Hart's (1967) Lemma is similar to Theorem 1 here. Their Theorem is similar to Theorem 6 here. They do not have any results comparable to Theorems 2 and 3. That is, they do not separate cross-validation error into two components, accuracy and
stability. It is this separation which enables me to achieve the results of Section 4. Section 4 does not correspond to any previous work with nearest neighbor pattern classification. Section 4 is evidence in favor of the two-component view of cross-validation error.

We know from Section 3 that $m_{1}$ has maximal accuracy, since $E\left(\left\|\vec{e}_{t}\right\|\right)=0$. In general, $m_{k}$ will be less accurate than $m_{1}$, but $m_{k}$ will also be more stable than $m_{1}$. As $k$ increases, we expect that accuracy decreases (error on the training set increases) and stability increases. At some point, as $k$ increases, further increase in stability will not sufficiently compensate for the decrease in accuracy. Section 5 shows how we can find the value of $k$ that gives the best balance between accuracy and stability. It is certainly possible, for a certain set of data, that the best balance is $k=1$. This is not a new insight. Cover and Hart (1967) prove that there are situations in which $m_{1}$ has a lower probability of error than $m_{k}$, where $k>1$. What is new is the observation that there are situations in which, as $k$ increases, stability may decrease (see Section 4.2). Thus both accuracy and stability decrease. Section 5 shows how we may detect such situations, using the estimators $r_{i}$ for $\rho_{i}$ and $t_{i}$ for $\tau_{i}$.

Aha et al. (1991) also examine $k$ nearest neighbor algorithms, but their analysis is a worst-case analysis, assuming certain probability distributions. They also assume numeric-valued attributes.

Langley (1993) also presents an average-case analysis of a nearest neighbor algorithm. Langley's (1993) analysis assumes a conjunctive target concept, noise-free boolean attributes, and a uniform distribution over the instance space. These assumptions do not overlap with the assumptions made here, which makes the two analyses highly complementary. A worthwhile project for the future would be to integrate the two analyses in a single framework.

## 7 Future Work

The weak point of this theoretical model of cross-validation error is the assumption that the inputs $X$ are the same in the training set $\left(X, \vec{y}_{1}\right)$ and the testing set $\left(X, \vec{y}_{2}\right)$. The reason for making this assumption is that it simplifies the mathematics. If the inputs in the training set were different from the inputs in the testing set, then we would need to make some assumptions about $f$ before we could prove any interesting results. When the inputs are the same in the training and testing sets, we can prove some interesting results about cross-validation error, without making any assumptions about the deterministic
component $f$ of the black box.
The main implication of the assumption, that $X$ is the same in the training set and the testing set, is that we may be underestimating the cross-validation error. We can expect a model to have more difficulty with prediction when the inputs in the testing set are different from the inputs in the training set. Theorem 1 shows that the assumption is reasonable for large $n$, but we would like to know what happens when $n$ is small. When a learning algorithm makes predictions for inputs in the testing set that do not appear in the training set, the algorithm is performing interpolation or extrapolation. Future work should extend the theory of cross-validation error to cover interpolation and extrapolation.

Another area for future work is proving that instance-based learning satisfies the assumptions of Theorem 3. It is conjectured that Theorem 3 is true for instance-based learning, or that it is true under certain weak conditions. However, there is not yet a proof for this.

Finally, it would be interesting to apply the general theory of cross-validation error to some other learning algorithms, such as ID3. Even when analytical results are not feasible, it is possible to use Monte Carlo techniques to evaluate the stability of learning algorithms. A Monte Carlo approach may also be required to evaluate the stability of some of the more elaborate instance-based learning or memory-based reasoning algorithms. The complexity of some of these algorithms makes a formal analysis difficult.

## 8 Conclusion

This paper introduces a general theory of cross-validation error for algorithms that predict boolean-valued attributes. It shows that cross-validation error has two components, accuracy and stability. The optimal cross-validation error is derived and it is shown that algorithms that maximize accuracy will give sub-optimal cross-validation error.

The theory was then applied to an analysis of voting in instance-based learning. It was proven that single nearest neighbor algorithms are unstable, because they maximize accuracy. The motivation for voting is to increase stability. In the best case, voting does indeed increase stability. However, in the worst case, voting can actually decrease stability. Techniques were provided for estimating the expected cross-validation error. These techniques can be used to find the best values for $k$ (the size of the neighborhood) and $t$ (the threshold for voting).

## Acknowledgments

I would like to thank Dr. David Aha for extensive, helpful comments on this paper. I would also like to thank Professor Malcolm Forster for pointing out the relation between my work and Akaike Information Criterion statistics.

## Notes

1. The work described here does not rely on any of the results in Turney (1993). However, to avoid duplication, some relevant issues that are discussed in Turney (1993) are not discussed here (and vice versa). Therefore, although the work here stands on its own, the interested reader may wish to also read Turney (1993).
2. Weiss and Kulikowski (1991) reserve the term "cross-validation" for $N$-fold cross-validation. In $N$-fold cross-validation, the data set is split into $N$ subsets of roughly equal size. The classification algorithm is then tested $N$ times, each time training with $N-1$ of the subsets and testing with the remaining subset. Therefore N -fold cross-validation is essentially $N$ applications of what is called "cross-validation" here. The essential concepts remain the same whether the process is done once or $N$ times.
3. It is a common practice in statistics to prove theorems that hold asymptotically, as the size of the data set increases to infinity. For example, Cover and Hart (1967) prove asymptotic theorems. Informally, we may say that asymptotic theorems assume an infinite-sized data set. The first assumption here, that the inputs are the same in the testing and training sets, is closely related to the assumption that the data sets are infinitely large. This is the point of Theorem 1 and Lemma 1 . However, assumption 1 may be satisfied by small, finite sets of data, when the data are generated by a planned, controlled series of experiments. Thus assumption 1 is weaker than the assumption of infinite data sets.

## References

Aha, D.W., Kibler, D., \& Albert, M.K. (1991) Instance-based learning algorithms, Machine Learning, 6:37-66.

Cover, T.M., \& Hart, P.E. (1967) Nearest neighbor pattern classification, IEEE Transactions on Information Theory, IT-13:21-27. Also in (Dasarathy, 1991).

Dasarathy, B.V. (1991) Nearest Neighbor Pattern Classification Techniques, Edited collection (California: IEEE Press).

Fix, E., \& Hodges, J.L. (1951) Discriminatory analysis: nonparametric discrimination: consistency properties, Project 21-49-004, Report Number 4, USAF School of Aviation Medicine, Randolph Field, Texas, 261-279. Also in (Dasarathy, 1991).

Fraser, D.A.S. (1976) Probability and Statistics: Theory and Applications (Massachusetts: Duxbury Press).

Kibler, D., Aha, D.W., \& Albert, M.K. (1989) Instance-based prediction of real-valued attributes, Computational Intelligence, 5:51-57.

Langley, P. (1993) Average-case analysis of a nearest neighbor algorithm, Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, Chambéry, France, in press.

Sakamoto, Y., Ishiguro, M., \& Kitagawa, G. (1986) Akaike Information Criterion Statistics (Dordrecht, Holland: Kluwer Academic Publishers).

Tomek, I. (1976) A generalization of the $k$-NN rule, IEEE Transactions on Systems, Man, and Cybernetics, SMC-6:121-126. Also in (Dasarathy, 1991).

Turney, P.D. (1990) The curve fitting problem: a solution, British Journal for the Philosophy of Science, 41:509-530.

Turney, P.D. (1993) A theory of cross-validation error. Submitted to the Journal of Experimental and Theoretical Artificial Intelligence.

Weiss, S.M., \& Kulikowski, C.A. (1991) Computer Systems that Learn: Classification and Prediction Methods from Statistics, Neural Nets, Machine Learning, and Expert Systems (California: Morgan Kaufmann).

## Appendix: List of Symbols

Note: Symbols that only occur in one theorem are omitted from this list.

| $\ldots \oplus \ldots$ | the exclusive-or operator |
| :---: | :---: |
| $\\|\ldots\\|$ | the length of a vector |
| $\mu$ | the mean of a probability distribution |
| $\rho_{i}$ | the frequency with which $f$ takes the value 1 in $N_{k}\left(\vec{v}_{i}\right)$ |
| $\sigma$ | the standard deviation of a probability distribution |
| $\sigma^{2}$ | the variance of a probability distribution |
| $\tau$ | the value of the threshold $t$ at which voting becomes unstable |
| $\tau_{i}$ | $\tau$ for a particular neighborhood $N_{k}\left(\vec{v}_{i}\right)$ |
| $\vec{e}_{c}$ | the cross-validation error vector |
| $\vec{e}_{s}$ | the instability vector |
| $\vec{e}_{t}$ | the training set error vector |
| $E(\ldots)$ | the expectation operator |
| $f(\vec{v})$ | a scalar function of the input vector $\vec{v}$; the target concept |
| $\mathfrak{f}(X)$ | a vector function of the input matrix $X$ |
| $k$ | the size of the neighborhood |
| $m(\vec{v} \mid X, \vec{y})$ | the model's prediction for $f(\vec{v})$, given data $X$ and $\vec{y}$; a scalar |
| $\vec{m}(X \mid X, \vec{y})$ | the model's prediction for $\hat{f}(X)$, given the data $X$ and $\vec{y}$; a vector |
| $m_{\alpha}$ | the optimal model, assuming that $f$ and $p$ are known to the modeler |
| $m_{\beta}$ | the model $\vec{m}\left(X \mid X, \vec{y}_{i}\right)=\vec{y}_{i}$ |
| $n$ | the number of samples; the number of observations |


| $N_{k}\left(\vec{v}_{i}\right)$ | the set of the $k$ nearest neighbors $\vec{v}_{1}, \ldots, \vec{v}_{k}$ to $\vec{v}_{i}$ |
| :---: | :---: |
| $p$ | the probability that a random boolean variable $z$ has the value 1 |
| $p^{\prime}$ | the probability that $f(\vec{v})$ is 1 , when $f(\vec{v})$ is treated as a random variable |
| $p_{\alpha}$ | the probability of error $E\left(\left\\|\vec{e}_{c}\right\\|\right) / n$ for $m_{\alpha}$ |
| $p_{\beta}$ | the probability of error $E\left(\left\\|\stackrel{e}{e}_{c}\right\\|\right) / n$ for $m_{\beta}$ |
| $P_{k}$ | the probability that the majority of $y_{1,1}, \ldots, y_{1, k}$ are 1 |
| $r$ | the number of inputs; the number of features |
| $r_{i}$ | an unbiased estimator for $\rho_{i}$ |
| $R$ | $p_{\beta}$ in the terminology of Cover and Hart (1967) |
| $R^{*}$ | $p_{\alpha}$ in the terminology of Cover and Hart (1967) |
| $\operatorname{sim}(\ldots, \ldots)$ | the similarity measure; a scalar measure of similarity between vectors |
| $t$ | a threshold for voting, where $0<t<1$; with majority voting, $t=0.5$ |
| $t_{i}$ | an unbiased estimator for $\tau_{i}$ |
| $\vec{v}=\left[\begin{array}{lll} x_{1} & \ldots & x_{r} \end{array}\right]$ | a vector of inputs; the attributes of an observation; features |
| $\vec{v}_{i}$ | the $i$-th row in $X$; the $i$-th input vector; the $i$-th observation |
| X | a matrix of inputs; rows are observations and columns are features |
| $y$ | the output; the class variable |
| $z$ | a random boolean variable; noise in the class variable $y$ |
| $\stackrel{\rightharpoonup}{z}$ | a vector of random boolean variables |

## Figures



Figure 1. Plot of $E\left(\left\|\stackrel{\rightharpoonup}{e}_{c}\right\|\right) / n$ as a function of $p$.


Figure 2. A plot of $E\left(\left\|\stackrel{\rightharpoonup}{e}_{s}\right\|\right) / n$ as a function of $p$.


Figure 3. A plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$.


Figure 4. A plot of $E\left(\left\|\vec{e}_{s}\right\|\right) / n$ as a function of $p$.

