

Feature selection for genetic sequence classification

Nadia A. Chuzhanova, Antonia J. Jones and Steve Margetts

Published in: *Bioinformatics* **14**(2):139-143, 1998.

Abstract. *Motivation:* Most of existing methods for genetic sequence classification are based on a computer search for homologies in nucleotide or amino acid sequences. The standard sequence alignment programs scale very poorly as the number of the sequences increases or the degree of sequence identity is less than 30%. Some new computationally inexpensive methods based on nucleotide or amino acid compositional analysis have been proposed, but prediction results are still unsatisfactory and depend on the features chosen to represent the sequences.

Results: In this paper a feature selection method based on the Gamma (or near neighbour) test is proposed. If there is a continuous or smooth map from feature space to the classification target values the Gamma test gives an estimate for the mean-squared error of the classification, despite the fact that one has no *a priori* knowledge of the smooth mapping. We can search a large space of possible feature combinations for a combination which gives a smallest estimated mean-squared error using a genetic algorithm. The method was used for feature selection and classification of the large subunits of ribosomal RNA according to RDP (Ribosomal Database Project) phylogenetic classes. The sequences were represented by dinucleotide frequency distribution. The nearest-neighbour criterion has been used to estimate the predictive accuracy of the classification based on the selected features. For examples discussed we found that the classification according to the first nearest neighbour is correct for 80% of the test samples. If we consider the set of the 10 nearest-neighbours then 94% of the test samples are classified correctly.

Availability: The principal novel component of this method is the Gamma test and this can be downloaded compiled for Unix Sun 4, Windows 95, and MS-DOS from <http://www.cs.cf.ac.uk/ec/>

Contact: E-mail address: s.margetts@cs.cf.ac.uk

Keywords: feature selection, classification, nearest neighbour, large subunits of rRNA.

Communication regarding this paper should be addressed to:

Antonia J. Jones

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF WALES, CARDIFF
PO BOX 916
Cardiff CF24 3XF

Telephone: ..+44-292-087-5490

Telefax: ..+44-292-087-4598

Date/version: 18 January 2002

Accepted: *Bioinformatics*, 14: 1998

Copyright © 1997 Nadia A. Chuzhanova, Antonia J. Jones and Steve Margetts

Feature selection for genetic sequence classification

CONTENTS

INTRODUCTION	1
SYSTEM AND METHODS	2
DATA REPRESENTATION	2
FEATURE SELECTION VIA THE Γ -TEST	2
ACCURACY OF FEATURE SELECTION	4
IMPLEMENTATION	4
RIBOSOMAL LSU RNA CLASSIFICATION	4
CONCLUSIONS	5
REFERENCES	5

List of Tables

Table I LSU rRNA phylogenetic classes.	7
Table I Training, embedding, and prediction results.	8

Feature selection for genetic sequence classification

Nadia A. Chuzhanova¹

Institute of Mathematics, Siberian Branch of Russian Academy of Science, 630090, Novosibirsk, Russia.

Antonia J. Jones*

Department of Computer Science, University of Wales, Cardiff, PO Box 916, Cardiff CF2 3XF, UK.

Steve Margetts

Department of Computer Science, University of Wales, Cardiff, PO Box 916, Cardiff CF2 3XF, UK.

Abstract.

Motivation: Most of existing methods for genetic sequence classification are based on a computer search for homologies in nucleotide or amino acid sequences. The standard sequence alignment programs scale very poorly as the number of the sequences increases or the degree of sequence identity is less than 30%. Some new computationally inexpensive methods based on nucleotide or amino acid compositional analysis have been proposed, but prediction results are still unsatisfactory and depend on the features chosen to represent the sequences.

Results: In this paper a feature selection method based on the Gamma (or near neighbour) test is proposed. If there is a continuous or smooth map from feature space to the classification target values the Gamma test gives an estimate for the mean-squared error of the classification, despite the fact that one has no *a priori* knowledge of the smooth mapping. We can search a large space of possible feature combinations for a combination which gives a smallest estimated mean-squared error using a genetic algorithm. The method was used for feature selection and classification of the large subunits of ribosomal RNA according to RDP (Ribosomal Database Project) phylogenetic classes. The sequences were represented by dinucleotide frequency distribution. The nearest-neighbour criterion has been used to estimate the predictive accuracy of the classification based on the selected features. For examples discussed we found that the classification according to the first nearest neighbour is correct for 80% of the test samples. If we consider the set of the 10 nearest-neighbours then 94% of the test samples are classified correctly.

Availability: The principal novel component of this method is the Gamma test and this can be downloaded compiled for Unix Sun 4, Windows 95, and MS-DOS from <http://www.cs.cf.ac.uk/ec/>

Contact: E-mail address: s.margetts@@cs.cf.ac.uk

INTRODUCTION

The exponential growth of molecular sequencing data requires developing advanced computational methods for rapid comparison of new sequences with known genetic material in order to make a decision about their taxonomic (evolutionary) and/or functional relatedness.

Most of existing methods for genetic sequence classification are based on a computer search for homologies in nucleotide or amino acid sequences. The standard sequence alignment programs have been designed to provide a compromise between the speed (it is known that the problem of multiple alignment is NP-hard) and accuracy of the search. As a result, they work well only when there

¹ Present address: Department of Computer Science, University of Wales, Cardiff, PO Box 916, Cardiff CF2 3XF, UK.

* To whom correspondence should be addressed.

is a reasonably high degree of sequence identity, usually of the order 30% or more.

Some new computationally inexpensive methods based on nucleotide or amino acid compositional analysis with or without biochemical parameters, e.g. molecular weight (if known) and isoelectric point have been proposed. Different metrics and distances have been used to express the similarity between sequences, but prediction results are still unsatisfactory and depend on the features chosen to represent the sequences.

The frequency of the oligonucleotides or amino acids are commonly used as a classificational features. Nussinov has found that there are stable and statistically significant asymmetries in some dinucleotide frequency distribution for different taxonomic groups [Nussinov 1984]. According to the "genome hypothesis" formulated by Grantham *et al* [Grantham 1980], the choice between synonymous codons varies from one gene to another and depends on the type of genome the gene occurs in. So codon usage may be regarded, to some extent, as a measure of "genome" similarity and gene expressivity [Cowe 1991]. Even the amino acid composition is not random and, as has been shown in [Hobohm 1994], can be used for protein identification.

On the other hand, as has been shown in [Hobohm 1994], some classification procedures may perform better, if the number of features is reduced according to some biochemical rationale and only a part of the frequency distribution is used. This corresponds also to current trends in neural networks, the principle of minimal architecture.

It is clear that the ideal feature selection for sequence classification should satisfy the following conditions:

- similar sequences should be represented by similar feature vectors;
- similar feature vectors should give similar or equal classification value.

Methods for feature selection proposed in [Petrokovski 1990], [Wu 1996] satisfied the first of these conditions. They do not incorporate classification values at the stage of feature selection.

In this paper the formal feature selection method based on the Gamma (or near neighbour) test [Stefánsson 1997], [Končar 1997] is proposed. This procedure appears to give accurate (probabilistic) estimates for the mean-squared error of the classification variable, for a wide class of feature vectors (not especially for oligonucleotide or amino acid frequency distribution), independently of any detailed knowledge of the function from input feature space to output classification target except that it should be smooth (bounded first and second order partial derivatives). The process of finding the best subset of the features is speeded up by using genetic algorithms [Holland 1975] and a *kd*-tree technique for the construction of the nearest-neighbour lists [Friedman 1977].

The method was used for feature selection and classification of the large subunit of ribosomal RNA (LSU rRNA) according to RDP (Ribosomal Database Project) phylogenetic classes [Maidak 1994]. The sequences were represented by dinucleotide frequency distribution. The nearest-neighbour criterion has been used to estimate the predictive accuracy of the classification based on a selected features.

SYSTEM AND METHODS

DATA REPRESENTATION

Denote by Σ a finite alphabet of cardinality $|\Sigma|$. Let T be a string $a_1 a_2 \dots a_N$ of length N over the alphabet Σ . A substring of length l such as $a_i a_{i+1} \dots a_{i+l-1}$ is called an *oligonucleotide* or *l-gram*. There are $|\Sigma|^l$ possible *l-grams* over the alphabet Σ and $N-l+1$ *l-grams* in the text of length N .

Any text T may be represented by its *l-gram* frequency distribution or, in other words, by the *l-gram* spectrum [Chuzhanova 1989] for fixed l which consists of pairs $\{l\text{-gram}, \text{the number of occurrences of this } l\text{-gram in text } T\}$. When the order of *l-grams* is fixed, the spectrum contains only the frequencies of corresponding *l-grams*.

FEATURE SELECTION VIA THE Γ -TEST

Let a data sample be represented by $((x_1, \dots, x_N), y) = (\mathbf{x}, y)$ in which we think of the vector $\mathbf{x} = (x_1, \dots, x_N)$ as the *input*, confined to a closed bounded set $C \subseteq \mathbb{R}^m$ and the scalar y as the *output*. In the present case, the *input* is the *l-gram* spectrum of a data sample and the *output* is the class number to which the data sample belongs.

We assume that training and testing data are different sample sets in which:

- the training set inputs are non-sparse in input-space;
- each output is determined from the inputs by a deterministic process which is the same for both training and test sets;
- each output is subjected to statistical noise with finite variance whose distribution may be different for different outputs but which is the same in both training and test sets for corresponding outputs.

We focus on the case where samples are generated by a suitably continuous but unknown function $f: \mathbb{C} \subseteq \mathbb{R}^m \rightarrow \mathbb{R}$ and $y = f(x_1, \dots, x_N) + r$ where r represents an indeterminable part, which may be due to real noise or might be due to lack of functional determination in the posited input/output relationship, i.e. an element of 'one \rightarrow many-ness' present in the data. The Gamma test is designed to give a data-derived estimate for the variance $\text{Var}(r)$.

Suppose (\mathbf{x}, y) is a data sample. Let (\mathbf{x}', y') be a data sample such that Euclidean distance $|\mathbf{x}' - \mathbf{x}| > 0$ is minimal and the minimum is taken over the set of all sample points different from \mathbf{x} . Thus \mathbf{x}' is the nearest neighbour to \mathbf{x} (in any ambiguous case we just pick one of the several equidistant points arbitrarily). The Gamma test (or near neighbour technique) is based on the statistic

$$\gamma = \frac{1}{2M} \sum_{i=1}^M (y'(i) - y(i))^2$$

where M is the number of input/output training pairs.

It can be shown that $\gamma \rightarrow \text{Var}(r)$ in probability as the nearest neighbour distances approach zero. In a finite data set we cannot have nearest neighbour distances arbitrarily small so the Gamma test is designed to estimate this limit by means of a linear correlation.

Given data samples $(\mathbf{x}(i), y(i))$, where $\mathbf{x}(i) = (x_1(i), \dots, x_m(i))$, $1 \leq i \leq M$, let $\mathbf{x}(N(i, p))$ be the p -th nearest neighbour to $\mathbf{x}(i)$. Nearest neighbour lists for $p \leq 20$ (say) nearest neighbours can be found in $O(M \log M)$ time using a kd -tree technique developed by Bentley [Friedman 1977].

We write

$$\Delta(p) = \frac{1}{p} \sum_{h=1}^p \frac{1}{M} \sum_{i=1}^M \|\mathbf{x}(N(i, h)) - \mathbf{x}(i)\|^2$$

and

$$\Gamma(p) = \frac{1}{p} \sum_{h=1}^p \frac{1}{2M} \sum_{i=1}^M (y(N(i, h)) - y(i))^2$$

then $\Delta(p)$ is the mean square distance of the $h \leq p$ nearest neighbours and $\Gamma(p)$ is an estimate for the statistic γ based on the $h \leq p$ nearest neighbours. Both $\Delta(p)$ and $\Gamma(p)$ are easily computed from the data set.

As p increases we might expect to find that $\Gamma(p)$ grows approximately linearly with $\Delta(p)$ for small p . We certainly found this correlation to hold when function f was known to be smooth. The linear correlation of $\Gamma(p)$ with $\Delta(p)$ suggests a possible method for dealing with the fact that in a finite data set we cannot have $\max |\mathbf{x}' - \mathbf{x}|$ arbitrary small. We can estimate $\Gamma(p)$ and $\Delta(p)$ for the first several values of p and extrapolate the regression line to $\Delta = 0$. The intercept $\Gamma = \lim \Gamma$, see Figure 1, will usually give an improved estimate for $\text{Var}(r)$.

The Gamma test can be used to estimate the best embedding dimension: the minimal Γ corresponding to the more important and informative subset of features. On data sets which are not excessively large and the number of features m is less than 20 the test is sufficiently fast to be run on a complete examination of all possible $2^m - 1$ subsets of features. For a larger number of features we use a Genetic Algorithm (in the sense of Holland [Holland 1975]). Every feature selection S_j , $1 \leq j \leq 2^m - 1$ is represented by a binary string of length m , where '0' in the i -th position means that the i -th feature is excluded from current selection and '1' indicates its presence otherwise. The fitness of each member (individual) S_j of the population $\{S_1, \dots, S_K\}$ ($K \leq 2^m - 1$) is a function of the Γ -value: $\text{Fitness}(S_j) = 1/(1 + e^{\Gamma(j)/\tau})$, where $\tau > 0$ is a constant which controls the shape of the fitness curve, and $\Gamma(j)$ is the Γ -value found by running the Gamma test with features selected according to S_j .

An initial population of unique random bitstrings is generated and the fitness of each is found as above. In each generation, a set number of breeding events take place. Two parents P_1 and P_2 are randomly selected from the current population: P_1 with a probability proportional to its fitness, P_2 uniformly from all individuals. These two parents are combined using a one-point crossover operation, by picking a random cut-point c , where $1 \leq c \leq m$. A 'child' is formed from the first c bits of P_1 and the last $m - c$ bits of P_2 . This child is then subjected to a mutation operator which flips each bit in the child according to a small mutation probability (the experiments performed used a probability of 0.01). If the resulting child is unique and not the all-zero string, its fitness is calculated as above and it is inserted into the population by overwriting the least fit individual. If not, another cut-point is chosen and a new child generated. If no unique child can be formed after 10 iterations, a new pair of parents is chosen. Testing the uniqueness of the children in this way takes very little efforts and ensures that the population remains diverse.

ACCURACY OF FEATURE SELECTION

To estimate the accuracy of feature selection the proposed method was tested on a set of sequences with known feature preferences.

Let us consider Markov chains of order 0 where the symbols from the alphabet $\{A, C, G, U\}$ occur with probabilities p_A, p_C, p_G and p_U . It is clear, that the probabilities of corresponding bigrams will be $p_{AA} = p_A \cdot p_A, p_{AC} = p_A \cdot p_C$ and so on. Suppose that the training set includes two classes of the sequences over the alphabet $\{A, C, G, U\}$ generated with the probabilities $p_A = 0.25, p_C = 0.25, p_G = 0.25, p_U = 0.25$ and $p_A = 0.2, p_C = 0.2, p_G = 0.3, p_U = 0.3$. In other words, the sequences from the second class have more letters G or U and correspondingly, more bigrams GU or UG ($p_{GU} = p_{UG} = 0.09$) than the sequences from the first class ($p_{GU} = p_{UG} = 0.0625$). Thus the frequency of bigrams GU or UG might be a good classification feature for these particular classes.

We were therefore encouraged to find that the feature selection method proposed herein when applied to this problem does indeed select the frequency of bigram UG as the most informative feature.

The analysis of the Markov chains of order 0 with other symbol distributions has shown that if there are asymmetries in bigram frequencies then the feature selection obtained by the method discussed here is the same as expected according to the probabilities of bigrams.

IMPLEMENTATION

RIBOSOMAL LSU RNA CLASSIFICATION




Figure 1 Gamma test regression line on the LSU rRNA dataset based on 2-gram's, where selected features are included, see [Končar 1997].

To illustrate its possibilities, the method was used for feature selection and LSU rRNA classification according to RDP (Ribosomal Database Project) phylogenetic classes [Maidak 1994]. The RDP database was chosen for the following reasons. First, the RDP database is one of the few databases that organize their entries according to phylogenetic relationship with a high level of accuracy. Second, its size is relatively small. Third, classification and annotation of unknown rRNA sequences is now very important as a method of biosphere monitoring.

The data used for Gamma test contained 236 LSU rRNA from 20 phylogenetic classes (Table 1) derived from RDP database (Release 5.0, December 13, 1996). Five classes were rejected because each had only one sequence present in the database. The alphabet of RNA includes four symbols {A, C, G, U}. The approximate length of the sequences is 3200-5000 nucleotides. The sequences were represented by their bigram spectrum or, in other words, by 16 features. To make the sequence representation length-invariant, all values were scaled by N_i^{-l+1} where N_i is the length of the i -th sequence. Note that the choice of the length l is very important. Values $l = 2, 3$ are preferable both from the biological point of view (see introduction) and for the robustness of prediction: with $l > 3$ the training set is separated better but the robustness of prediction falls down [Gusev 1990].

To estimate the predictive accuracy of the proposed method, a set of sequences was randomly divided into three groups of approximately equal size. Two of them were used for feature selection (or training) and one for prediction. All three combinations have been employed. In the fourth case all 236 LSU rRNA were used for feature selection and 34 other sequences from 10 phylogenetic classes for predicting (i.e. testing).

The full embedding has been done for each case and subsets of features were selected using the Gamma test. During the prediction phase for each sequence from the test set n nearest neighbours in this selected feature space are computed. The output values of these n nearest neighbour sequences are then used as outputs for the new sequences. For "First fit" $n=1$ and the class number is predicted as simply the class of the first nearest neighbour sequence. For "Five fits" $n=5$ and correct classification lies within the set of the first five nearest neighbours, etc.

The information about training and testing sets, the number of selected features and prediction results according to one (first fit), five (5-fits) and ten nearest neighbours (10-fits) are shown in Table II.

As can be seen from Table II, the predictive accuracy according to ten nearest neighbours is about 94%. Of course, once the search has been narrowed with high probability to ten possibilities then a more detailed examination of these ten sequences can feasibly complete the classification. For the sequences which are not recognized correctly according to n nearest neighbours, where $1 \leq n \leq 10$, the predicted n classes although incorrect are nevertheless closely located on the phylogenetic tree. It would also be possible to train a neural network using these features and we shall report neural network classification results in a later paper. Here we are more interested in the feature selection procedure.

CONCLUSIONS

The method of feature selection and classification described here does not require searches for homologies, common subsequences or specific patterns all of which are very time consuming. It gives the opportunity to classify the new sequences into predefined classes on the phylogenetic tree without sequence alignment. The method is robust in the sense that when errors occur the incorrect classification is phylogenetically close to the correct classification.

Experiments with neural networks on the 72 LSU rRNA from 15 phylogenetic classes derived from the same database have been reported [Wu 1996]. The sequences were represented by their octagram spectrum ($l = 8$). The singular value decomposition method was used to reduce the number of octagrams to forty. In comparison the predictive accuracies are higher, from 92-100%. Nevertheless, the results of [Gusev 1990] suggest that methods based on the frequency of longer l -grams (in fact $l > 3$) are much less likely to be robust, in the sense that given a particular example of a class one can readily find an extended l -gram whose frequency characterizes this example uniquely. However, the same l -gram is relatively unlikely to occur without any mutations in another example of the same class. Thus frequencies of a diverse range of shorter l -grams taken together are likely to provide a much more robust classification.

This is a preliminary account of a new technique and is principally designed to show that the method has promise. One can envisage many improvements, for example the Euclidean metric may not be the most appropriate given the context and the weighting of

features could be continuous rather than discrete, but these are questions which can reasonably be addressed in future studies. Of course, it is also possible to improve the prediction accuracy by enlarging the training set and by increasing the quality of the samples.

REFERENCES

- [Chuzhanova 1989] Chuzhanova,N. (1989) Inductive method of program synthesis for symbolic sequence processing. In: *Lecture Notes in Artificial Intelligence*, Springer-Verlag, **397**, 317--327.
- [Cowe 1991] Cowe,E. and Sharp,P.M. (1991) Molecular evolution of bacteriophages: discrete patterns of codon usage in T4 genes are related to the time of gene expression. *J. Mol. Evol.*, **33**, 13--22.
- [Friedman 1977] Friedman,J.H.,Bentley,J.L., and Finkel,R.A. (1977) An algorithm for finding best matches in logarithmic expected time. *ACM Trans. on Mathematical Software*, **3**,200--226.
- [Grantham 1980] Grantham,R., Gautier,C., Gouy,M., Mercier,R., and Pave,A. (1980) Codon catalog usage and genome hypothesis. *Nucleic Acids Res.*, **8**, 49--62.
- [Gusev 1990] Gusev,V.D. and Chuzhanova,N.A. (1990) The algorithms of recognition of the functional sites in genetic texts. In: *Proc. of the First Intern. Workshop on Algorithmic Learning Theory*, Tokyo, 109--119.
- [Hobohm 1994] Hobohm,U.,Houthaeve,T., and Sander,C. (1994) Amino acid analysis and protein database compositional search as a rapid and inexpensive method to identify proteins. *Analytical Biochem.*, **222**, 202--209.
- [Holland 1975] Holland,J. (1975) Adaptation in natural and artificial systems. MIT Press, Cambridge, MA.
- [Končar 1997] Končar, N. (1997) Optimisation methodologies for direct inverse neurocontrol. Ph.D. Thesis, Department of Computing, Imperial College, London.
- [Maidak 1994] Maidak,B. *et al.* (1994) The ribosomal database project. *Nucleic Acids Res.*, **22**, 3485--3487.
- [Nussinov 1984] Nussinov,R. (1984) Strong duplet preferences in nucleotide sequences and DNA geometry. *J. Mol. Evol.* **20**, 11--119.
- [Petrokovski 1990] Petrokovski, S., Hirshon, J. and Trifonov, E.N,. (1990) Linguistic measure of taxonomic and functional relatedness of nucleotide sequences. *J. Biomolec. Struct. Dyn.*, **7**, 1251--1268.
- [Stefánsson 1997] Aðalbjörn Stefánsson, Končar,N. and Jones, Antonia J. (1997) A note on the Gamma test. *Neural Computing and Applications*, **5**(3), 131--133.
- [Wu 1996] Wu,C.H. (1996) Gene classification artificial neural system. *Methods in Enzymology*, **266**, 71--88.

Table I LSU rRNA phylogenetic classes.

Class Number	Number of sequences	Class name
1	3	Archaea:Crenarchaeota
2	11	Archaea:Euryarchaeotata:Archaeoglobales
3	2	Bacteria:Flavobacteria and relatives
4	10	Bacteria:Gram Positives and relatives, High G+C
5	23	Bacteria:Gram Positives and relatives, Low G+C
6	5	Bacteria:Proteobacteria Alpha
7	8	Bacteria:Proteobacteria Beta
8	2	Bacteria:Proteobacteria Epsilon
9	7	Bacteria:Proteobacteria Gamma
10	2	Bacteria:Spirochetes
11	9	Eukarya:Animalia:Arthropoda:Uniramia:Insecta
12	5	Eukarya:Fungi:Eumycota:Ascomycotina Hemiascomycetes
13	6	Eukarya:Plantae:Magnoliophyta::Magnoliopsida
14	15	Eukarya:Protoctista:Zoomastigina::Kinetoplastida
15	69	Mitochondria:Animalia:Arthropoda:Uniramia:Insecta
16	4	Mitochondria:Fungi:Eumycota: Ascomycotina:Plectomycetes
17	3	Mitochondria:Plantae:Bryophyta::Marchantiopsida
18	21	Mitochondria:Protoctista:Rhizopoda::Lobosea
19	8	Plastids:Plantae:Magnoliophyta::Magnoliopsida
20	23	Plastids:Protoctista:Chlorophyta::Chlorophyceae

Table I Training, embedding, and prediction results.

Group ID	Number of training sequences	Number of selected features	F-value ($p \leq 20$)	Number of testing sequences	Prediction Accuracy (%)		
					First fit	5-fits	10-fits
1	165	8	1.92×10^{-7}	70	70	84.3	93
2	160	10	5.2×10^{-7}	75	74.7	92	94.6
3	161	8	2.8×10^{-6}	74	64	75	88
4	235	13	5.1×10^{-9}	34	80	91	94