

The Accuracy of DNA Sequences: Estimating Sequence Quality

BU-1104-M

36 pp.

February 1991

Gary A. Churchill
Biometrics Unit
337 Warren Hall
Cornell University
Ithaca, NY 14853

Michael S. Waterman
Departments of Mathematics and Molecular Biology
University of Southern California
Los Angeles, CA 90089-1113

This paper was typed at USC.

The Accuracy of DNA Sequences: Estimating Sequence Quality*

Gary A. Churchill
Biometrics Unit
337 Warren Hall
Cornell University
Ithaca, NY 14853

Michael S. Waterman
Departments of Mathematics and Molecular Biology
University of Southern California
Los Angeles, CA 90089-1113

*Supported by grants from the National Institutes of Health and the National Science Foundation.

Abstract

In this paper we describe a method for the statistical reconstruction of a large DNA sequence from a set of sequenced fragments. We assume that the fragments have been assembled and address the problem of determining the degree to which the reconstructed sequence is free from errors, i.e. its accuracy. A consensus distribution is derived from the assembled fragment configuration based upon the rates of sequencing errors in the individual fragments. The consensus distribution can be used to find a minimally redundant consensus sequence which meets a prespecified confidence level, either base by base or across any region of the sequence. A likelihood based procedure for the estimation of the sequencing error rates is described which utilizes an iterative EM algorithm. Prior knowledge of the error rates is easily incorporated into the estimation procedure. The methods are applied to a set of assembled sequence fragments from the human G6PD locus. We close the paper with a brief discussion of the considerations involved in maintaining a database of sequence accuracy information.

1 Introduction

Central DNA sequence databases came into existence about 1982, and have undergone very rapid growth. Today the three major databases, DDBJ, EMBL and GenBank, which contain about 50×10^6 nucleotides, are virtually identical for all practical purposes and are positioning themselves for the era of genomic sequencing. In Waterman (1990) a number of issues regarding genomic sequence databases are raised. The present paper is focused on one of these issues: DNA sequence accuracy. It is commonly understood that the existing data is not of uniform accuracy, but generally this is not apparent from the database entries. In fact, what is acceptable accuracy and how to estimate accuracy is the subject of considerable recent debate (Roberts, 1990). Sequences appear as AGCTG..., for example, with no comments about their accuracy. In this paper we propose a method to estimate DNA sequence accuracy from the assembled fragments of a shotgun sequencing project. The result of our analysis is an estimate of the likelihood of the base in each position or a simple estimate of an entire consensus sequence at a given level of confidence. While the reliability of the consensus or finished sequence usually exceeds that of an individual fragment, it is valuable to quantify this. Similar approaches could be taken to determine the accuracy of other types of physical map data. See Branscomb et al. (1990) for the first approach to physical mapping based

on likelihood of overlap.

In a shotgun sequencing project, a large DNA molecule is broken into a collection of fragments. The fragments are cloned into a suitable vector and sequenced individually. The fragment sequences are then assembled by determining their relative orientations and overlaps and aligned to form a column-by-column correspondence (Churchill et al. 1991, Kececioglu and Myers 1990, Staden 1980). The accuracy of a finished sequence produced by this method will vary across sites. The depth of coverage varies for statistical and biological reasons (Lander and Waterman 1988, Edwards *et al.* 1990). It is intuitively obvious, and we shall make this notion precise, that increasing the depth of coverage should improve the accuracy of the finished sequence. The fragment assembly will inevitably contain columns with discrepancies due to errors in the fragment sequences. Sequencing errors will show some clustering due to chance but errors can also be caused by properties of the sequence itself. Homopolymeric runs and repetitive sequences in general are prone to errors. Sequences with a tendency to form secondary structures are suspected of causing compression errors. Sometimes all the fragments covering a region will have the same orientation, making these strand specific systematic errors hard to detect.

There are a number of reasons why it is important to assess the accuracy of

DNA sequence data. The level of accuracy required may depend on the types of analyses for which the sequence will be used. Sequence comparison algorithms exist which are robust to small numbers of errors. However, some aspects of sequence analysis are likely to be very sensitive to errors. The problem of finding coding regions is an important example. Current methods rely in large part on the ability to detect open reading frames (Fields and Soderland 1990). The effect of a single insertion or deletion error is a frameshift which can mask an open reading frame.

The study of natural variation in DNA sequences is another area where accuracy must be considered. It will be important to distinguish genuine variation from sequencing errors. The frequency of polymorphic DNA sites, although it may vary between species and across genomic regions within species, may be on the order of 1 base in 250-2000 (Gusella 1986). The level of accuracy should be known and accounted for in studies of sequence variation.

Sequence accuracy should be considered in the planning of experiments which require the design of specific oligonucleotides. Oligonucleotides selected from the most reliable regions of a sequence will maximize the probability of success. This probability could be calculated using the summary statistics we describe below. These considerations increase in importance for experiments which require a large number of specific hybridizations.

Finally, in order to design an effective strategy for genomic sequencing we must consider the tradeoff between accuracy and the cost of redundant sequencing. A reasonable strategy for a large project might be to produce the bulk of the sequence data rapidly and at low cost. Careful re-sequencing could then be focused on selected regions to attain a desired level of accuracy (E. Chen, Personal Communication).

This is the first attempt to quantify sequence quality. We will consider DNA sequence accuracy as a function of the redundancy of coverage and the frequencies of random errors in the fragment sequences. We take the term accuracy to mean the probability that a given DNA base or sequence of bases in a finished sequence is identical to the corresponding base in the actual DNA molecule. We assume implicitly that there is a single true sequence and not a population of sequences being studied. The problems of errors which occur at the DNA preparation or cloning stages and of errors in data transcription are not addressed here . Systematic errors which occur in sequence determination are also beyond the current scope of these methods. Thus, the statistics described represent a lower bound on sequence accuracy. We have made a number of assumptions to facilitate the statistical analysis. These assumptions are described next before we give an overview of our methods for estimating accuracy.

The starting point of our analysis is the assembled fragment configuration. A crucial assumption is that the fragment assembly is correct. We realize that this will never be absolutely true. Some ambiguities will exist in the alignments and the final product may depend on the assembly algorithm and/or the judgment of the investigator. Still, with low error rates in the fragment sequences and moderate to deep coverage, it is possible to obtain highly reliable assemblies. We assume that all fragment sequences have been carefully examined prior to assembly and that all traces of vector have been removed. The investigator may wish to re-examine certain gels that are questionable and correct the fragment sequences. This is expected and acceptable. However, it is not recommended that fragment sequences be altered to conform with an alignment after the assembly. Altering the fragment sequences in light of the alignment may distort the estimated error probabilities and may lead to unreliable estimates of accuracy.

We assume that sequencing errors are independent of their local context, i.e., that error probabilities depend only on the true base at a position and not on adjacent or more distant bases. This assumption can be relaxed somewhat at the cost of increasing the number of model parameters. Error rates that depend on the bases immediately to their 5' side in the fragment can be introduced into the model. However, the true situation is likely to be more

complex than a simple one step dependence. For example, compression sites are a common source of errors which may be related to the formation of local secondary structures extending over several bases.

We assume that the composition of the sequence is independent, both of adjacent bases and over large regions. Markovian dependence between bases is well established as are local variations in composition (Churchill 1989). However, the effect of these assumptions on the inferred final sequence is likely to be minor. Additional modeling of the DNA sequence composition could be incorporated, again at the cost of additional parameters.

We assume that the sequencing error rates are constant across the entire sequence. Constancy could be checked by a sliding window plot of the expected error frequencies. Some clustering of errors can be expected due to chance. The assembly itself may be incorrect in regions with a high rate of errors and should be re-examined. Otherwise, high error rates may be due to some unusual aspect of the local sequence.

We assume that all fragment sequences are equally reliable. Any number of factors may influence the reliability of individual fragments. If these factors are known to the investigator, a subjective weighting scheme could be imposed on the fragment set. It may also be possible to detect “bad” fragments after the assembly and consensus computations. A quality-of-fragment statistic

could be computed such as a chi-squared statistic based on the observed and expected numbers of errors in a fragment. However, a fragment sequence should not be removed from the assembly unless it is clearly anomalous, i.e., aligned incorrectly. A multiple testing problem is involved here and we refer the interested reader to Arratia and Gordon (1989) for a discussion of extreme values and the binomial distribution.

We assume that all regions within a fragment are equally reliable. However, the resolution of small fragments on a gel is generally better than the resolution of large fragments. Hence, one end of a fragment sequence will be more reliable than the other. It may be possible to explicitly model the decay of accuracy which occurs as the gel is read out to greater lengths. This would allow one to include data which are still informative although less accurate without compromising the quality of the final sequence. A simple solution to this problem would be to use two sets of error rate parameters, one for the first, more reliable part of the sequence, and a second set of rates for the more error prone portions.

In the sections below, we begin with a description of the assembled fragments. Then we examine the problem of determining a consensus when the sequencing error rates are known. A consensus distribution is defined which allows one to compute the most probable base at each position (Procedure A,

Section 3) or a redundant set of bases which exceeds a specified level of probability (Procedure B, Section 3). The latter procedure is extended in Procedure C, Section 3, to produce a collection of sequences which constitutes a global confidence interval for the true DNA sequence. Next, we turn to the problem of estimating the sequencing error rates using the method of maximum likelihood. The result is a simple iterative solution (Procedure D, Section 4), which is a special case of the EM algorithm (Dempster et al. 1977). Incorporation of prior knowledge about error rates into the estimation procedure is described. The methods are illustrated with an example from a large sequencing project (Chen et al. 1990).

2 The Assembled Fragments

A set of fragment sequences (fragments are indexed by $j = 1, \dots, m$) is aligned by some procedure. The result of the alignment procedure is a matrix with m rows. Each row contains the ordered sequence of bases in a particular fragment written in either direct or reverse complemented orientation. Gaps may be inserted internally and each fragment is offset to produce a column-by-column correspondence among the entire set of fragments. The column index $i = 1, \dots, n$ runs from the leftmost base in the assembly to the rightmost.

Let s_i denote the true state of the DNA sequence corresponding to column

i in the fragment assembly. The true state may take any value in the set $\mathcal{A} = \{A, C, G, T, \Delta\}$. The symbol Δ is included to allow for extra columns in the fragment assembly which do not correspond to any base in the true DNA sequence. The alphabet of redundant bases is defined to be the set of all non-empty subsets of \mathcal{A} . A partial notation for this set is given by the standard IUPAC DNA alphabet. When a redundant set includes one or more bases and Δ , we will append a * to the IUPAC symbol, i.e. $R^* = \{A, G, \Delta\}$. Redundant bases will be used below in the definition of consensus sequences.

The elements of the fragment assembly matrix, denoted by x_{ij} , may take values in the set $\mathcal{B} = \{A, C, G, T, \Delta, X, \phi\}$, where Δ denotes an internal gap and X denotes any ambiguous determination of a base. The set of ambiguous bases can be expanded, but the methods described below are essentially unchanged. The null symbol ϕ is used as a place holder for non-aligned positions beyond the ends of a fragment.

All sequenced bases x_{ij} are recorded in a standard orientation relative to the assembly. However, many (typically half) of the bases are sequenced on fragments which are reverse complemented relative to the standard orientation. For each fragment we will define the reverse complement indicator to be

$$r_j = \begin{cases} 0 & \text{Fragment } j \text{ is direct} \\ 1 & \text{Fragment } j \text{ is reverse complemented.} \end{cases} \quad (1)$$

We will use the notation a^c to denote the base complementary to a . Note that a^c is well defined even for redundant bases: $A^c = T$, $\Delta^c = \Delta$, $\{A, G\}^c = \{C, T\}$.

The depth of coverage at position i is defined to be the number of fragments contributing sequence information (including internal gaps),

$$d_i = \sum_{j=1}^m \mathbf{1}(x_{ij} \neq \phi). \quad (2)$$

The notation $\mathbf{1}(E)$ is used to denote the indicator function for the event E . This function takes the value one when E is true and otherwise zero.

If the fragment sequences could be determined without error, our knowledge of the true sequence would be exact and, except for the problem of correct assembly, there would be no need for inference. However, sequencing errors do occur and we propose to describe them by the following probability model. An error occurs when the true base $s_i = a$ is misread in fragment j to yield $x_{ij} = b$, $b \neq a$. Recall the assumptions that the error probabilities are constant across all positions, that sequencing errors occur independently within and across columns of the fragment assembly, and that the error rate depends only on the values of a and b . The sequencing error probabilities will be denoted by

$$p(b|a) = \Pr(x_{ij} = b \mid s_i = a); a \in \mathcal{A}, b \in \mathcal{B}. \quad (3)$$

It will also be necessary to define a probability distribution for the com-

position of the DNA sequence. We have assumed bases occur independently with identical distribution across the assembly. The composition probabilities will be denoted by

$$p(a) = \Pr(s_i = a); a \in \mathcal{A}. \quad (4)$$

Note that our definition of sequence composition includes the gap frequencies. More complex models for the DNA base composition and the sequencing error probabilities could be incorporated at this stage.

3 Consensus with Confidence

Our goal is to make the best possible determination of s_1, \dots, s_n given the data in the fragment assembly. We will initially assume that the sequencing error probabilities and DNA base composition are known exactly. In section 4, we will address the practical problem of estimating these quantities while simultaneously estimating the consensus sequence.

The *consensus distribution* is a probability distribution over the set \mathcal{A} defined by

$$\pi_i(a) = \Pr(s_i = a \mid x_{ij}; j = 1, \dots, m). \quad (5)$$

It is the probability that the true base is a given the fragment assembly data.

These probabilities can be computed using Bayes' rule,

$$\pi_i(a) = \frac{p(a) \prod_{j=1}^m [(1 - r_j)p(x_{ij}|a) + r_j p(x_{ij}^c|a^c)]}{\sum_{b \in \mathcal{A}} p(b) \prod_{j=1}^m [(1 - r_j)p(x_{ij}|b) + r_j p(x_{ij}^c|b^c)]} \quad (6)$$

where we set $p(\phi|a) = 1$ for all $a \in \mathcal{A}$.

We will present three different approaches to the definition of a consensus sequence. First, we will compute the most likely values of s_i at each position. Second, we will allow redundancy in the consensus and compute a subset of \mathcal{A} such that the probability that s_i is a member of this set exceeds $1 - \alpha$ at each position. Finally, we will describe the construction of a consensus sequence with redundant bases such that the probability that the entire sequence s_1, \dots, s_n or region of interest s_a, \dots, s_b is contained within the redundant consensus sequence with probability at least $1 - \alpha$.

Procedure A: The most likely value of s_i is the base a that maximizes $\pi_i(a)$ over a . Denote this value by c_i , then

$$\max_{a \in \mathcal{A}} \pi_i(a) = \pi_i(c_i). \quad (7)$$

Procedure B: Now we wish to allow redundant bases into the consensus. Our goal is to find the character with minimum redundancy which has probability in excess of a specified level $1 - \alpha$. The probability that a redundant base includes the true base is the sum over all bases represented by the sym-

bol c_i . The choice of consensus sequence is made so that at each position the probability $\pi_i(c_i)$ exceeds a specified level,

$$\sum_{a \in c_i} \pi_i(a) \geq 1 - \alpha. \quad (8)$$

To find the best redundant set, we begin by choosing the most likely base at position i , call it a_i . If $\pi_i(a_i) > 1 - \alpha$ we stop. Otherwise, we add to the redundant set, the next most likely base, call it b_i . If $\pi_i(a_i) + \pi_i(b_i) > 1 - \alpha$ we stop. The process continues until the level exceeds $1 - \alpha$. It is guaranteed to stop because $\pi_i(A) + \pi_i(C) + \pi_i(G) + \pi_i(T) + \pi_i(\Delta) = 1$.

Procedure C: A redundant sequence can be thought of as a collection of sequences with a different member for each unique expansion of its redundant characters. Formally this is a *cylinder set* in sequence space (Feller p. 130). We wish to construct such a collection of sequences which will contain the true sequence with probability exceeding $1 - \alpha$. The probability that the true sequence s_1, \dots, s_n lies within a redundant sequence c_1, \dots, c_n is given by

$$\prod_{i=1}^n \Pr(s_i \subset c_i \mid x_{ij}, j = 1, \dots, m), \quad (9)$$

where c_i is taken from the redundant sequence alphabet and the event $s_i \subset c_i$ indicates that s_i is a member of the set represented by c_i .

A consensus sequence with global coverage probability $1 - \alpha$ can be constructed by the following procedure. Set a threshold level T and find the

marginal consensus for this level according to Procedure B. Compute the global coverage probability for this consensus (equation 9). If this probability is less than $1 - \alpha$, increase T . Otherwise T may be decreased until the probability of containment just exceeds $1 - \alpha$. Because the containment probability is a monotone function of T , we can solve the problem quickly using a bisection algorithm.

One practical problem with this procedure is that many positions may have equal distribution π_i and hence will all be set to the same level of redundancy. This problem could be avoided by randomization. We prefer to choose the smallest value of T such that the coverage exceeds $1 - \alpha$ and report the actual coverage achieved.

4 Estimation of the Sequencing Error Rates

If the true DNA sequence were known, it would be trivial to estimate the composition and sequencing error rates. One would simply count the number of times that a base a occurs in the sequence

$$n_a = \sum_{i=1}^n \mathbf{1}(s_i = a) \quad (10)$$

and the number of times the base a was recorded as the base b in a fragment

$$n_{ab} = \sum_{i=1}^n \sum_{j=1}^m [(1 - r_j) \mathbf{1}(x_{ij} = b) \mathbf{1}(s_i = a) + r_j \mathbf{1}(x_{ij}^c = b) \mathbf{1}(s_i^c = a)] \quad (11)$$

for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$. Maximum likelihood estimates of the base composition and error rate parameters are given by

$$\hat{p}(a) = n_a/n \quad (12)$$

and

$$\hat{p}(b|a) = n_{ab}/n_a. \quad (13)$$

This situation suggests the following algorithm for the simultaneous estimation of the error rates and the consensus distribution.

Procedure D:

1. Initialize the consensus distribution:

Set $\pi_i(x) = 1.0$ where x is the most frequently occurring letter at column i .

2. Estimate $p(a)$ and $p(b|a)$ for all a, b :

Set the counts n_a and n_{ab} equal to their conditional expected values

$$\hat{n}_a = \sum_{i=1}^m \pi_i(a), \quad (14)$$

$$\hat{n}_{ab} = \sum_{i=1}^n \sum_{j=1}^m [(1 - r_j) \mathbf{1}(x_{ij} = b) \pi_i(a) + r_j \mathbf{1}(x_{ij}^c = b) \pi_i(a^c)] \quad (15)$$

and estimate $p(a)$ and $p(b|a)$ as before (equations 12 and 13).

3. Recompute $\pi_i(a)$ for all i and a :

according to equation (6), with $p(a)$ and $p(b|a)$ replaced by their current estimates.

4. Continue.

If the change in $\hat{p}(a|b)$ and $\hat{p}(a)$ are less than ϵ , for all a and b , stop. Otherwise go to Step 2.

This estimation procedure is a special case of the EM algorithm for a mixture of multinomial distributions (Dempster et al. 1977).

5 Using Prior Information

When a sequencing project has been ongoing, a great deal of information will accumulate regarding the types and frequencies of sequencing errors (see Krametz (1989)). It will be to our advantage to incorporate this information into estimates of the error rates. When error rates are small any one assembly will contain only a few errors and probabilities will not be estimated accurately from the data. In general, our revised estimate of $p(b|a)$ will be a mixture of

the prior expectation $\mu(b|a)$ and the current maximum likelihood estimate $\hat{p}(b|a)$,

$$\tilde{p}(b|a) = \alpha\mu(b|a) + (1 - \alpha)\hat{p}(b|a), \quad (16)$$

for some mixing proportion $0 \leq \alpha \leq 1$. In the extreme cases, for $\alpha = 1$ we use only prior information and for $\alpha = 0$ we use only the current data.

An analytically convenient representation of prior information is the Dirichlet distribution (Lindley 1972, pp. 59). For each $a \in \mathcal{A}$ we have a prior distribution over $p(\cdot|a)$ which is proportional to $\prod_{b \in \mathcal{B}} p(b|a)^{\beta_{ab}}$. The parameters β_{ab} determine the prior means and the mixing proportions. The *weight* of the prior is defined as

$$k_a = \sum_{b \in \mathcal{B}} \beta_{ab} \quad (17)$$

and the prior means are

$$\mu(b|a) = \beta_{ab}/k_a. \quad (18)$$

The error rate estimates corresponding to equation (16) will be

$$\tilde{p}(b|a) = \left(\frac{k_a}{n_a + k_a} \right) \mu(b|a) + \left(\frac{n_a}{n_a + k_a} \right) \hat{p}(b|a). \quad (19)$$

Thus, the prior information effectively adds $k = \sum_{a \in \mathcal{A}} k_a$ observations to the current data.

To implement these estimates using the EM algorithm (Procedure D),

augment equations 14 and 15 as follows:

$$\tilde{n}_a = \hat{n}_a + k_a \quad (20)$$

$$\tilde{n}_{ab} = \hat{n}_{ab} + \beta_{ab}. \quad (21)$$

These quantities are used to estimate $p(b|a)$ (equation 13) but the estimates of $p(a)$ (equation 12) are unchanged. A convenient choice for the prior distribution parameters is to take β_{ab} equal or proportional to \tilde{n}_{ab} from previous sequence assemblies.

6 Coverage and Expected Accuracy

Before a sequencing project is undertaken, the investigators should make some decision about the degree of accuracy which is desired. The size of the region to be sequenced and the purposes for which the sequence will be used should be considered. The need to quantify the relationship between cost of sequencing and accuracy was discussed by Waterston (Roberts, 1990). For example, if the target region is 10^6 base pairs and the maximum acceptable number of errors is 100, the probability of a sequencing error should be less than 10^{-4} . The desired degree of accuracy can be attained by adjusting the depth of coverage.

An approximate relationship between accuracy and the depth of coverage will be derived. An exact result depends on knowledge of the detailed structure

of the error rates. We avoid this complication by the following assumption.

Let

$$\sum_{\substack{b \in \mathcal{B} \\ b \neq a}} p(b|a) = \bar{p} \quad \text{for all } a \in \mathcal{A}. \quad (22)$$

We think of \bar{p} as the average error rate in sequenced fragments.

Clearly, increasing depth increases the accuracy of columns with no discrepancies, eg. a column with all “A”s. The probability that such a column is generated in error is approximately

$$\Pr(\text{error}) = 1 - \exp \left\{ (\bar{p}/1 - \bar{p})^d \right\} \quad (23)$$

which goes to zero very rapidly as d gets larger. However, the probability that a column of depth d contains no discrepancies is $(1 - \bar{p})^d$ which decreases as d gets larger. Thus, the number of columns without discrepancies will decrease as depth of coverage increases. To show that increasing depth is desirable, we will compute the expected accuracy, first for fixed depth d and then as an average over the distribution of d .

For a position which is covered to a depth d , we will compute the probability that at most $1/2$ of the observed bases are correct. This procedure for determining consensus is more conservative than procedure A and has the advantage of being much easier to compute. This is a binomial probability

given by

$$P(\text{error}|d) = \sum_{k=0}^{\lfloor d/2 \rfloor} \binom{d}{k} (\bar{p})^{d-k} (1-\bar{p})^k. \quad (24)$$

When sequence fragments are randomly located, there will be some variation in depth of coverage. In fact the distribution of d in a shotgun assembly is approximately Poisson with some mean λ . We ignore positions not covered by any fragments. Thus the mean coverage is $c = \lambda/(1 - e^{-\lambda})$. We can compute the average error probability conditional on the depth being at least 1 as:

$$P(\text{error}) = \frac{1}{1 - e^{-\lambda}} \sum_{k=1}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} P(\text{error}|k). \quad (25)$$

This quantity is summarized in Figure 1. Variations in the coverage will decrease the accuracy relative to a uniformly covered sequence with the same mean coverage c . In either case, increasing depth of coverage increases the expected accuracy of the finished sequence.

Now we are able to provide an answer to the question, what is the depth of coverage required to attain a given level of accuracy in the finished sequence. If our goal is 10^{-4} errors per base, then for $\bar{p} = 0.01$ we will require a mean coverage of $c = 6.8$. If the error rate in fragment sequences can be kept down to $\bar{p} = 0.001$ we will need $c = 3.8$. To achieve 10^{-6} errors per base at $\bar{p} = 0.01$ we will need $c = 12.5$ and at $\bar{p} = 0.001$ we will need $c = 4.2$.

7 The Human G6PD Locus

Recently Chen et al. (1990) presented the sequence of 20,114 nucleotides of human DNA which includes the human glucose-6-phosphate dehydrogenase gene (G6PD). Defective G6PD genes can cause hemolytic anemia but they also offer partial protection against malaria. Consequently many variants exist and over 300 have been described. Chen et al. (1990) sequenced 15,860 nucleotides of the transcribed region from the mRNA start site to the polyadenylation site. There are 13 exons, and the 20,114 nucleotides are about 25% Alu sequence.

The total sequence was determined from 3 EcoRI fragments. Chen et al. (1990) kindly provided us with their assembly of the largest fragment, 11,791 nucleotides in length. The fragment was isolated in lambda clones, subcloned in pUC18, and sequenced in fragments randomly subcloned in M13. The assembly has 165 fragments with mean length 283. The average depth of coverage is therefore about 4. The fragments are very accurate with small error probabilities. We ran Procedure D on the fragment set and estimated π_i and $p(a|b)$. (See Figure 2) To summarize the analysis we present two graphs. In Figure 3a is a graph of depth of coverage vs. nucleotide position. In addition to depth, we graph entropy e_i in Figure 3b to show the variability of π_i . The

entropy of the distribution π_i is

$$e_i = - \sum_{a \in A} \pi_i(a) \log_2 (\pi_i(a)). \quad (26)$$

When $\pi(a) = 1$ for some a , $e = 0$, the minimum value e can attain. The function e is maximized when $\pi = \left(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}\right)$: then $e = \log_2(5)$. Therefore uncertainty in the consensus distribution is indicated by large values of e_i .

As an illustration of procedure C, we have constructed a series of cylinder sets around the 605bp region containing the Alu cluster 8^R98^L (see Chen *et al.* 1991). This region was selected because of its intrinsic interest and because the depth of coverage is 1 at 165 positions and ranges up to 5 with a mean coverage $c = 2.94$. There are 11 columns with discrepancies involving insertion/deletion of bases and only 2 mismatched bases. A portion of this region is shown in Fig. 4. A cylinder set with containment probability $\geq 95\%$ comes at the high price of 255 ambiguous bases, leaving 350 bases well defined. This result should not be surprising, as the precise determination of a large number of independent quantities is a very strong inference. By reducing the containment probability, we can increase the number of well defined bases and/or reduce the ambiguity of the remaining bases. Notice, for example, the progression for * to X to R to A in the depth 1 portions of the sequence as the containment probability is decreased. The tightest cylinder set computed contains a single ambiguous base corresponding to a C,T mismatch in the alignment. This cylinder contains the

true sequence with 12% probability.

8 Discussion

Although these results emphasize the importance of redundancy in achieving accurate sequence data, they rely upon certain assumptions which will not always be met in practice. In particular the problem of handling systematic errors must be addressed more fully. The increased reliability of sequence data which results from using multiple independent clones and/or different sequencing technologies to cover the same region are not taken into account. We consider this study to be a starting point for future work on DNA sequence accuracy. Much work remains to be done both in the analytic aspects of modeling and in the empirical realm of cataloging sequence errors (see Krawetz, 1989). The procedures described here can be implemented as a first step.

The task of maintaining and updating a sequence quality database is straightforward. A quality index can be associated with each base. The quantities which must be stored in the database are $\pi_i(\cdot)$ and d_i , five floating point numbers and an integer. As new sequence information is added to the database, the accuracy information can be updated as follows:

$$\pi_i^*(a) = \frac{d_i^{(1)}\pi_i^{(1)}(a) + d_i^{(2)}\pi_i^{(2)}(a)}{d_i^{(1)} + d_i^{(2)}}, \quad (27)$$

where superscripts refer to the existing sequence information and the newly added information. The indices i should be adjusted to give the correct overlap of the data sets. It may be necessary to add columns internally to allow for insertions in one sequence relative to the other. The cost of tracking accuracy is a modest increase in storage space requirements which, in light of current technologies, is not a serious concern. If accuracy information were to be stored in a central database, simple sequence data could be distributed to the majority of users and the more detailed accuracy information could be made available upon request.

As a further refinement to sequence accuracy statistics, we would recommend that automated sequencing devices output a probability distribution reflecting the accuracy of individual base determinations. The methods described here can be readily modified to handle such data. The base values in fragment sequences are not simply single letters $x_{ij} \in \mathcal{B}$ but instead are probability distributions over the set \mathcal{A} , denoted $p_{ij}(x)$ ($i = 1, \dots, m$; $j = 1, \dots, n$; $x \in \mathcal{A}$). The error probabilities $p(b|a)$ are still well defined. The new version of the consensus distribution is

$$\pi_i(a) \propto p(a) \prod_{j=1}^m \left(\sum_{b \in \mathcal{A}} p(b|a) p_{ij}(b) \right), \quad (28)$$

where the denominator and the reversed-complement bookkeeping have been suppressed to simplify the notation.

There may be additional advantages to this approach. For example, fragment assembly algorithms could be modified to accept such input and would be less likely to get stuck trying to fit “bad” data. It may also be possible to read gels beyond the current limits of accurate resolution. One could extract information which is currently discarded without compromising the accuracy of the final sequence.

Although the procedures we have described are based on a number of assumptions which might not always be satisfied, we feel that many of the problems can be dealt with easily or safely ignored. A final and satisfactory answer to the problem of determining the accuracy of DNA sequences will eventually be achieved through experience with DNA sequencing and communication between statisticians and biologists. Practical experience with a simple system is the first step towards understanding its shortcomings. Eventually we will need to know much more of the details of sequence errors in order to deal with them effectively.

References

- Arratia, R.A. and Gordon, L. (1990) Tutorial on large deviations
for the binomial distribution. *Bulletin of Mathematical Biology*
51:125–131.

- Branscomb, E., Slezak, T., Pae, R., Galas, D., Carrano, A.V. and Waterman, M.S. (1990). Optimizing restriction fragment fingerprinting methods for ordering large genomic libraries. *Genomics*, **8**:351–366.
- Chen, E. *et al.* (1990). Sequence of human glucose-6-phosphate dehydrogenase cloned in plasmids and a Yeast Artificial Chromosome (YAC). Submitted to *Nucleic Acids Research*.
- Churchill, G.A. (1989) A stochastic model for heterogeneous DNA sequences. *Bull. Math. Biol.* **51**:79–94.
- Churchill, G.A., Burks, C., Eggert, M., Waterman. M.S. (1990) Fragment Assembly Methods for DNA Sequencing. Manuscript.
- Dempster, A.P., Laird, N.M., Rubin, D.B. (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. B* **39**:1–38.
- Edwards, A. *et al.* (1990). Automated DNA Sequencing of the Human HPRT Locus. *Genomics* **6**:593–608.
- Feller, W. (1968) *An Introduction to Probability Theory and Its Applications*. Vol. I, Third Edition, John Wiley & Sons.
- Fields and Soderland (1990). gm: a practical tool for automating

- DNA sequence analysis. *Cabois* **6**:263–270.
- Gusella, J.F. (1986) DNA Polymorphism and human disease. *Ann. Rev. Biochem.*, **55**:831–854.
- Kececioglu, J. and Myers, E. (1990). A robust automatic fragment assembly system. Preprint.
- Krawetz, S.A. (1989). Sequence errors described in GenBank: A means to determine the accuracy of DNA sequence interpretation. *Nucleic Acids Res.* **17**:3951–3957.
- Lander, E.S. and Waterman, M.S. (1988). Genomic mapping by fingerprinting random clones: A mathematical analysis. *Genomics* **2**:231–239.
- Lindley, D.V. (1972). *Bayesian Statistics, A Review*. SIAM Philadelphia, P.A.
- Roberts, L. (1990). Large-scale sequencing trials begin. *Science*, **250**: 1336–1338.
- Staden, R. (1980). A new computer method for the storage and manipulation of DNA gel reading data. *Nucleic Acids Res.* **8**:3673–2694.
- Waterman, M.S. (1990). Genomic Sequence Databases. *Genomics*

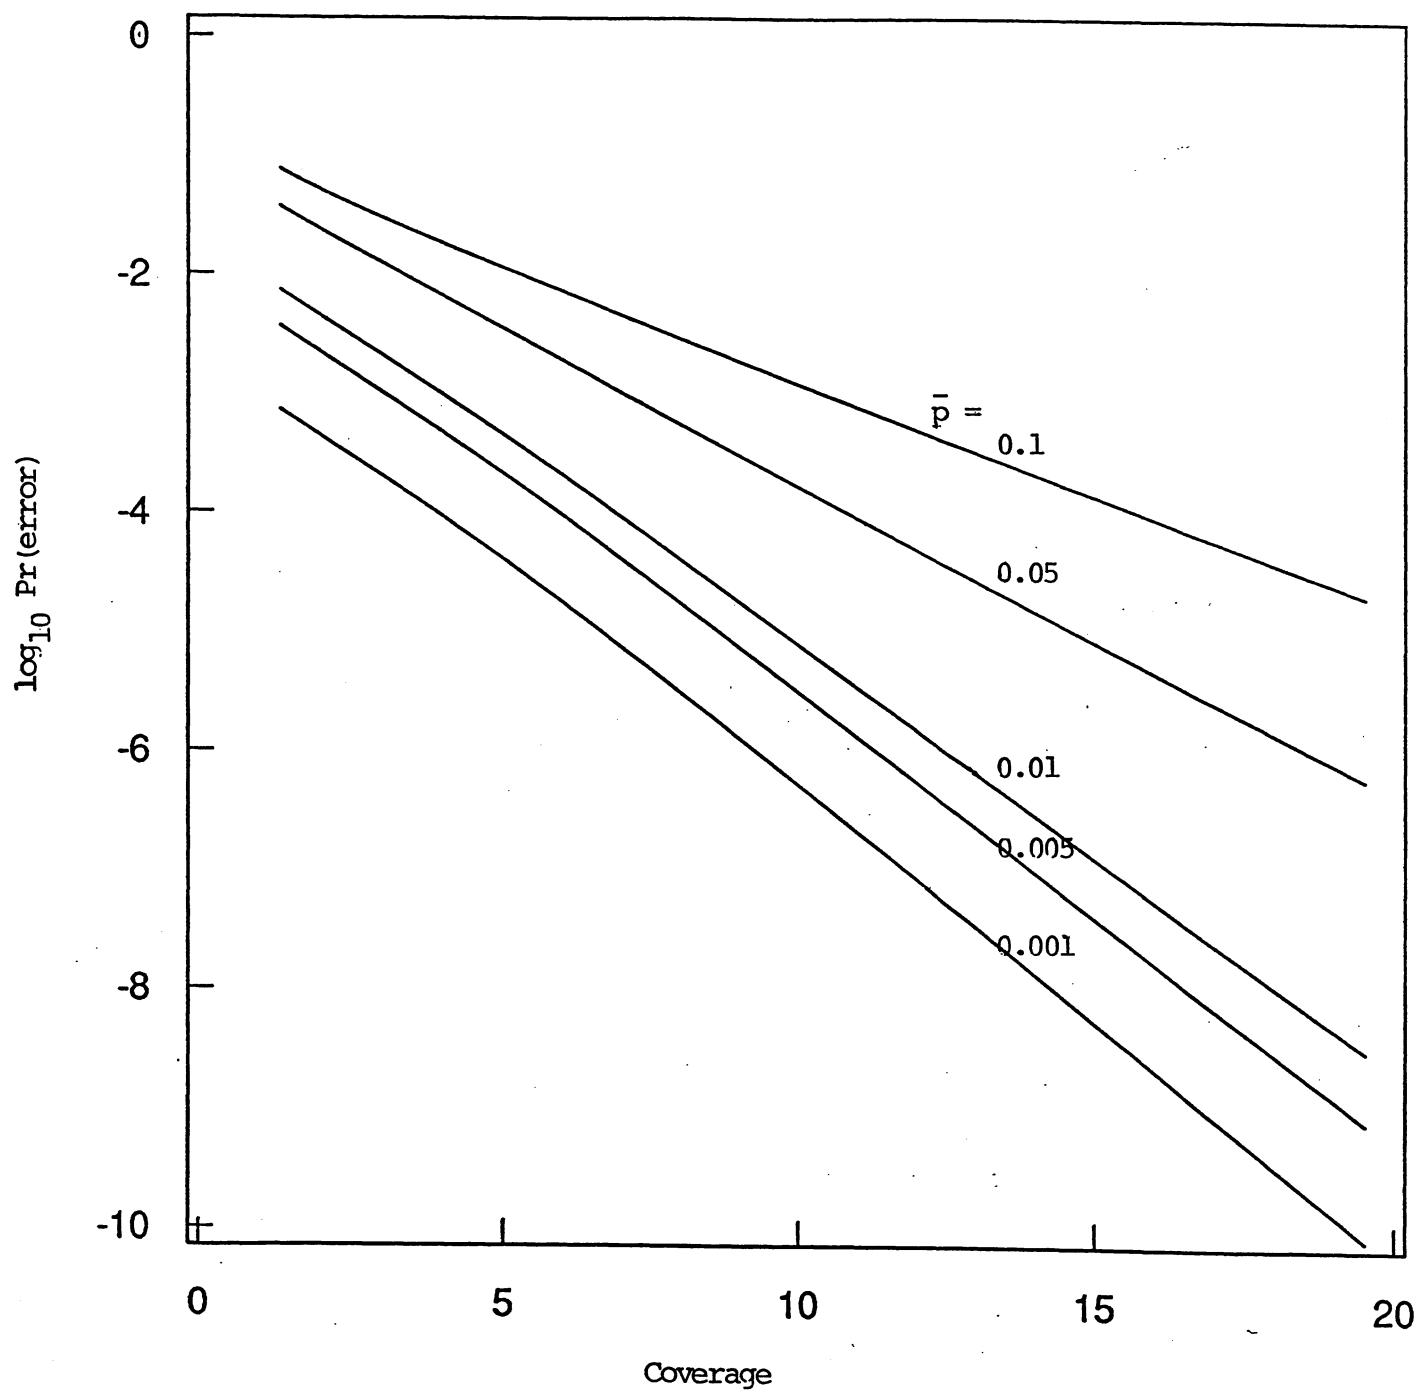
6:700–701.

FIGURE CAPTIONS

Figure 1: The approximate probability of an error in one base of finished sequence (equation 25) is shown as function of the mean depth of coverage. The scale for error probabilities is logarithm base 10.

Figure 3: The 11791 nucleotides of an EcoRI fragment from the G6PD region of human were sequenced by Chen et al. (1990). The depth of coverage is plotted against nucleotide position (3a). The consensus distribution (equation 6) for this assembly was computed using maximum likelihood parameter estimates given in figure 2. The uncertainty in the consensus distribution is summarized by plotting entropy (equation 26) against nucleotide position (3b).

Figure 4: Two segments of the fragment assembly from the region of Alu cluster 8R98L are shown (top). Symbols in the left hand column indicate the orientation of the sequenced fragments; R = reversed complement, D = direct. The depth of coverage is indicated. A series of cylinder sets is shown (bottom). Containment probabilities for the entire 605bp Alu cluster are given at left. The numbers of ambiguous bases are shown in brackets. Positions denoted by * are ambiguous including both a gap (-) and a base.



(2i)

	c	t	a	g	Δ	x
c	11918	6	4	2	65	3
t	9	10401	3	2	44	8
a	6	4	11035	3	30	5
g	5	1	11	12010	58	12
Δ	81	87	61	90	377	3

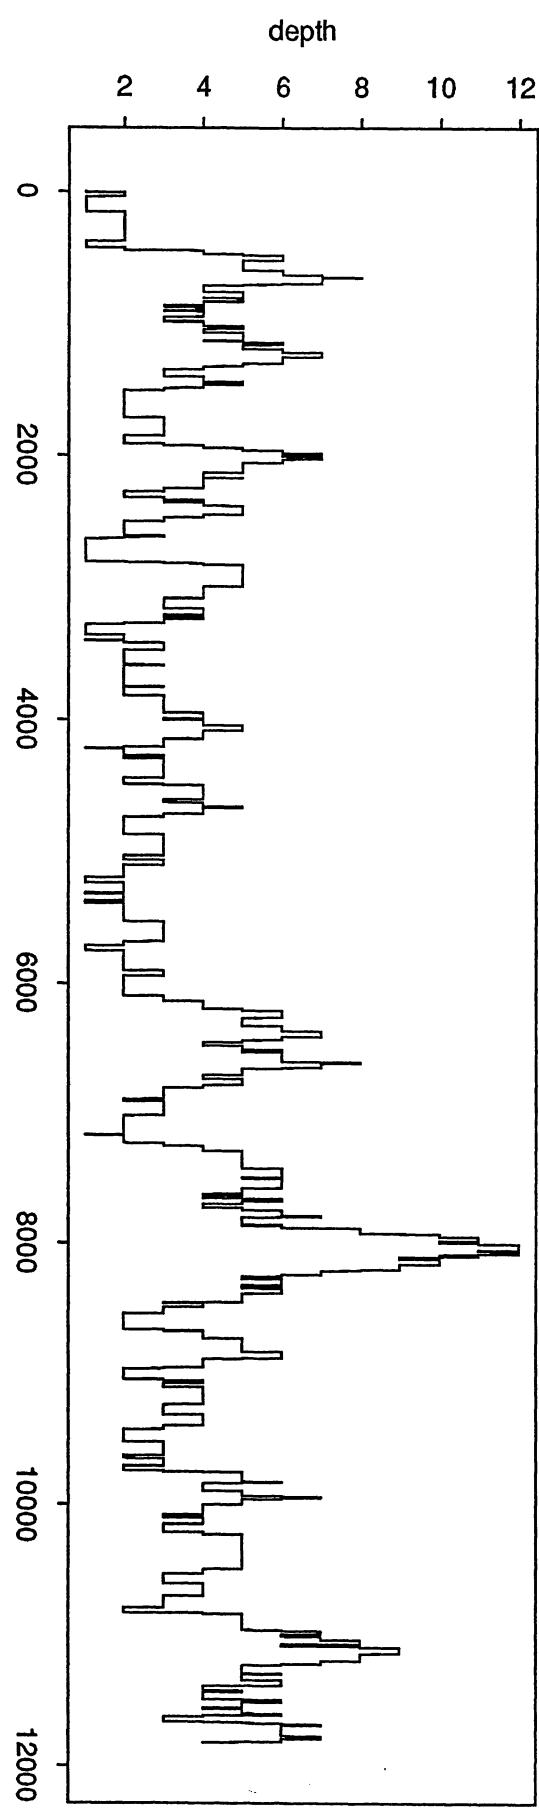
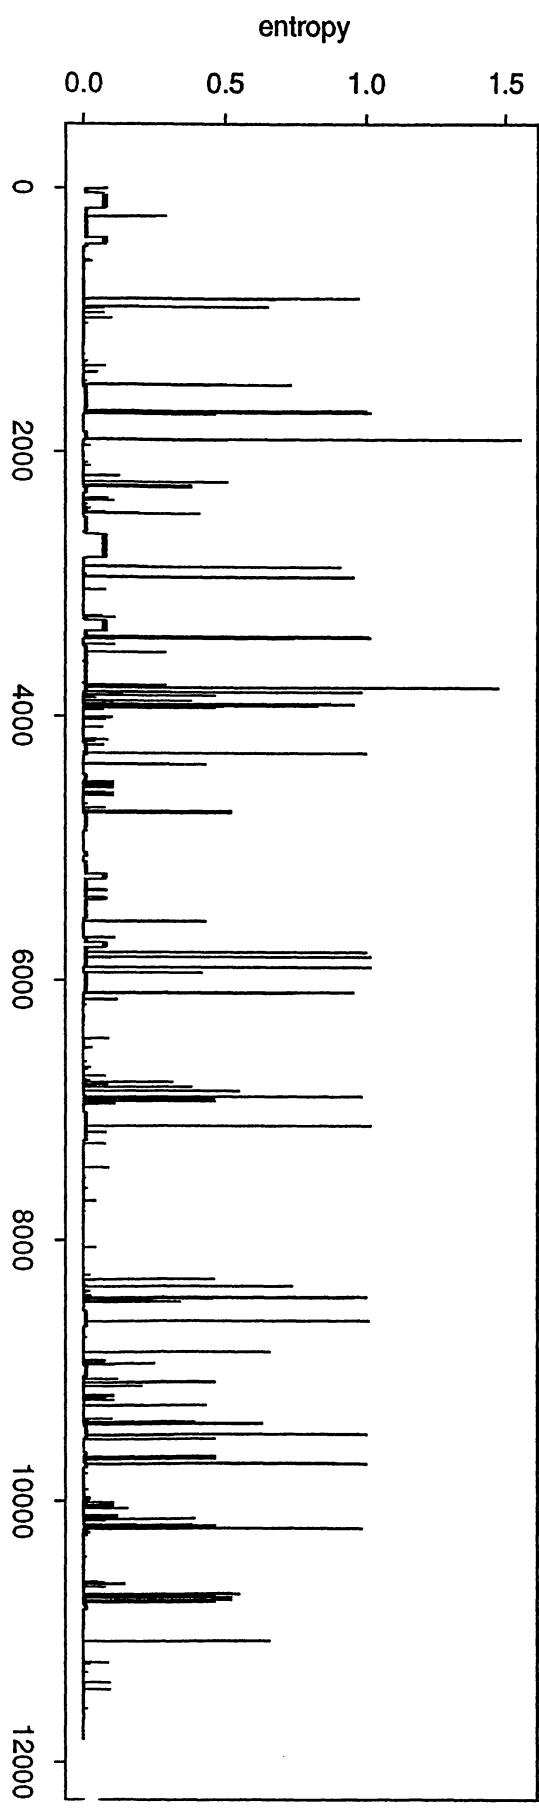
(2ii)

	c	t	a	g	Δ	x
c	0.993304	0.000504	0.000356	0.000178	0.005405	0.000252
t	0.000843	0.993742	0.000266	0.000188	0.004239	0.000721
a	0.000496	0.000402	0.995662	0.000266	0.002735	0.000438
g	0.000443	0.000086	0.000931	0.992712	0.004832	0.000996
Δ	0.116051	0.123898	0.087734	0.128809	0.539908	0.003599

(2iii)

c	t	a	g	Δ
0.258896	0.225850	0.239146	0.261041	0.015067

FIGURE 2 EM estimates from G6PD data (see section 7) with $N = 11865$ aligned positions. Figure (2i) is estimated error counts $= \hat{n}_{ab}$, Figure (2ii) is estimated error rates $\hat{p}(b|a)$, and Figure (3iii) is estimated sequence composition $\hat{p}(a)$.



R gagcaagactccatctaaaaaaaaaaaaaaaacaaaaattagctggatgtggcaggcacctggaatcccagctactcaggaggctaaggcagg
R aaattagctggatgtggatgtgg

tggaatcccagctactcaggaggctaaggcagg
aatcccagctactcaggaggctaaggcagg
ctactcaggaggctaaggcagg
taaggcagg

0.9562 (255) *****a*****a*****a*****aaa**a*c***a*****a*****a*****aatcccagctactcaggaggctaaggcagg
0.9153 (210) *****a*****a*****aaa**agc*gga*g*gg*gg*****gaaatcccagctactcaggaggctaaggcagg
0.8771 (170) *****a*****aaattagctggatgtggatgg*****tgaatcccagctactcaggaggctaaggcagg
0.6727 (170) *x***xx*x****x****xxxxxxxxx*xxxxxx*xxaaattagctggatgtggatgg*x***x**tgaatcccagctactcaggaggctaaggcagg
0.6524 (170) *r**rr*r****r****rrrrrrrr*rrrrrrr*rraaattagctggatgtggatgg*r***r**tgaatcccagctactcaggaggctaaggcagg
0.4372 (122) *a*xaa*ax*xxa*x*xaaaaaaaaaaaaaxaaaaattagctggatgtggatgg*a**xaxxtgaatcccagctactcaggaggctaaggcagg
0.3105 (73) gagyaagay*yya*y*yaaaaaaaaaaaaattagctggatgtggatgg*aggyayytgaatcccagctactcaggaggctaaggcagg
0.1215 (1) gagcaagactccatctaaaaaaaaaaaaacaaaaattagctggatgtggatggxaggcacctgaatcccagctactcaggaggctaaggcagg

