Information Theory in Molecular Biology

Christoph Adami[1,2]

[1]*Keck Graduate Institute of Applied Life Sciences,*
535 Watson Drive, Claremont, CA 91711

[2]*Digital Life Laboratory 139-74,*
California Institute of Technology, Pasadena, CA 91125

Abstract

This article introduces the physics of information in the context of molecular biology and genomics. Entropy and information, the two central concepts of Shannon's theory of information and communication, are often confused with each other but play transparent roles when applied to statistical *ensembles* (i.e., identically prepared sets) of symbolic sequences. Such an approach can distinguish between entropy and information in genes, predict the secondary structure of ribozymes, and detect the covariation between residues in folded proteins. We also review applications to molecular sequence and structure analysis, and introduce new tools in the characterization of resistance mutations, and in drug design.

In a curious twist of history, the dawn of the age of genomics has both seen the rise of the science of bioinformatics as a tool to cope with the enormous amounts of data being generated daily, and the decline of the *theory* of information as applied to molecular biology. Hailed as a harbinger of a "new movement" (Quastler 1953) along with Cybernetics, the principles of information theory were thought to be applicable to the higher functions of living organisms, and able to analyze such functions as metabolism, growth, and differentiation (Quastler 1953). Today, the metaphors and the jargon of information theory are still widely used (Maynard Smith 1999a, 1999b), as opposed to the mathematical formalism, which is too often considered to be inapplicable to biological information.

Clearly, looking back it appears that too much hope was laid upon this theory's relevance for biology. However, there was well-founded optimism that information theory ought to be able to address the complex issues associated with the storage of information in the genetic code, only to be repeatedly questioned and rebuked (see, e.g., Vincent 1994, Sarkar 1996). In this article, I outline the concepts of entropy and information (as defined by Shannon) in the context of molecular biology. We shall see that not only are these terms well-defined and useful, they also coincide precisely with what we intuitively mean when we speak about information stored in genes, for example. I then present examples of applications of the theory to measuring the information content of biomolecules, the identification of polymorphisms, RNA and protein secondary structure prediction, the prediction and analysis of molecular interactions, and drug design.

1 Entropy and Information

Entropy and information are often used in conflicting manners in the literature. A precise understanding, both mathematical and intuitive, of the notion of information (and its relationship to entropy) is crucial for applications in molecular biology. Therefore, let us begin by outlining Shannon's original entropy concept (Shannon, 1948).

1.1 Shannon's Uncertainty Measure

Entropy in Shannon's theory (defined mathematically below) is a measure of uncertainty about the identity of objects in an ensemble. Thus, while "en-

tropy" and "uncertainty" can be used interchangeably, they can *never* mean information. There is a simple relationship between the entropy concept in information theory and the Boltzmann-Gibbs entropy concept in thermodynamics, briefly pointed out below.

Shannon entropy or uncertainty is usually defined with respect to a particular observer. More precisely, the entropy of a system represents the amount of uncertainty *one particular observer* has about the state of this system. The simplest example of a system is a *random variable*, a mathematical object that can be thought of as an N-sided die that is uneven, i.e., the probability of it landing in any of its N states is not equal for all N states. For our purposes, we can conveniently think of a polymer of fixed length (fixed number of monomers), which can take on any one of N possible states, where each possible sequence corresponds to one possible state. Thus, for a sequence made of L monomers taken from an alphabet of size D, we would have $N = D^L$. The uncertainty we calculate below then describes the observer's uncertainty about the true identity of the molecule (among a very large number of identically prepared molecules: an *ensemble*), given that he only has a certain amount of probabilistic knowledge, as explained below.

This hypothetical molecule plays the role of a random variable if we are given its *probability distribution*: the set of probabilities $p_1, ..., p_N$ to find it in its N possible states. Let us thus call our random variable (random molecule) "X", and give the names $x_1, ..., x_N$ to its N states. If X will be found in state x_i with probability p_i, then the entropy H of X is given by Shannon's formula

$$H(X) = -\sum_{i=1}^{N} p_i \log p_i .\qquad(1)$$

I have not here specified the basis of the log to be taken in the above formula. Specifying it assigns units to the uncertainty. It is sometimes convenient to use the number of possible states of X as the base of the logarithm (in which case the entropy is between zero and one), in other cases base 2 is convenient (leading to an entropy in units "bits"). For biomolecular sequences, a convenient unit obtains by taking logarithms to the basis of the alphabet size, leading to an entropy whose units we shall call "mers". Then, the maximal entropy equals the length of the sequence in mers.

Let us examine Eq. (1) more closely. If measured in bits, a standard interpretation of $H(X)$ as an uncertainty function connects it to the smallest number of "yes-no" questions necessary, on average, to identify the state of

2

random variable X. Because this series of yes/no questions can be thought of as a *description* of the random variable, the entropy $H(X)$ can also be viewed as the *length of the shortest description of X* (Cover and Thomas, 1991). In case nothing is known about X, this entropy is $H(X) = \log N$, the maximal value that $H(X)$ can take on. This occurs if all states are equally likely: $p_i = 1/N \, ; i = 1, ..., N$. If something (beyond the possible number of states N) is known about X, this reduces our necessary number of questions, or the length of tape necessary to describe X. If I know that state $X = x_7$, for example, is highly unlikely, then my uncertainty about X is going to be smaller.

How do we ever learn anything about a system? There are two choices. Either we obtain the probability distribution using *prior knowledge* (for example, by taking the system apart and predicting its states theoretically) or by making measurements on it, which for example might reveal that not all states, in fact, are taken on with the same probability. In both cases, the difference between the maximal entropy and the remaining entropy after we have either done our measurements or examined the system, is the amount of information we have about the system. Before I write this into a formula, let me remark that, by its very definition, information is a *relative* quantity. It measures the *difference of uncertainty*, in the previous case the entropy before and after the measurement, and thus can never be absolute, in the same sense as potential energy in physics is not absolute. In fact, it is not a bad analogy to refer to entropy as "potential information", because potentially all of a system's entropy can be transformed into information (for example by measurement).

1.2   Information

In the above case, information was the difference between the maximal and the actual entropy of a system. This is not the most general definition as I have alluded to. More generally, information measures the amount of *correlation* between two systems, and reduces to a difference in entropies in special cases. To define information properly, let me introduce another random variable or molecule (call it "Y"), which can be in states $y_1, ..., y_M$ with probabilities $p_1, ..., p_M$. We can then, along with the entropy $H(Y)$, introduce the joint entropy $H(XY)$, which measures my uncertainty about the joint system XY (which can be in $N \cdot M$ states). If X and Y are *independent* random variables (like, e.g., two dice that are thrown independently)

3

the joint entropy will be just the sum of the entropy of each of the random variables. Not so if X and Y are somehow connected. Imagine, for example, two coins that are glued together at one face. Then, heads for one of the coins will always imply tails for the other, and vice versa. By gluing them together, the two coins can only take on two states, not four, and the joint entropy is equal to the entropy of one of the coins.

The same is true for two molecules that can bind to each other. First, remark that random molecules do not bind. Second, binding is effected by mutual specificity, which requires that part of the sequence of one of the molecules is interacting with the sequence of the other, so that the joint entropy of the pair is much less than the sum of entropies of each. Quite clearly, this binding introduces strong correlations between the states of X and Y: if I know the state of one, I can make strong predictions about the state of the other. The information that one molecule has *about* the other is given by

$$I(X:Y) = H(X) + H(Y) - H(XY) , \qquad (2)$$

i.e., it is the difference between the sum of the entropies of each, and the joint entropy. The colon between X and Y in the notation for the information is standard; it is supposed to remind the reader that information is a symmetric quantity: what X knows about Y, Y also knows about X. For later reference, let me introduce some more jargon. When more than one random variable is involved, we can define the concept of *conditional entropy*. This is straightforward. The entropy of X conditional on Y is the entropy of X *given Y*, that is, if I know which state Y is in. It is denoted by $H(X|Y)$ (read "H of X given Y") and is calculated as

$$H(X|Y) = H(XY) - H(Y) . \qquad (3)$$

This formula is self-explanatory: the uncertainty I have about X if Y is known is the uncertainty about the joint system minus the uncertainty about Y alone. The latter, namely the entropy of Y without regard to X (as opposed to "conditional on X") is sometimes called a *marginal* entropy. Using the concept of conditional entropy, we can rewrite Eq. (2) as

$$I(X:Y) = H(X) - H(X|Y) . \qquad (4)$$

We have seen earlier that for independent variables $H(XY) = H(X) + H(Y)$, so information measures the *deviation* from independence. In fact, it

4

measures exactly the amount by which the entropy of X or Y is reduced by knowing the other, Y or X. If I is non-zero, knowing one of the molecules allows you to make more accurate predictions about the other: quite clearly this is exactly what we mean by information in ordinary language. Note that this definition reduces to the example given earlier (information as difference between entropies), if the only possible correlations are *between* X and Y, while in the absence of the other each molecule is equiprobable (meaning that any sequence is equally likely). In that case, the marginal entropy $H(X)$ must be maximal ($H = \log N$) and the information is the difference between maximal and actual (i.e., conditional) entropy, as before.

1.3   Entropy in Thermodynamics

I will briefly comment about the relationship between Shannon's theory and thermodynamics (Adami and Cerf 1999). For the present purpose it should suffice to remark that Boltzmann-Gibbs thermodynamic entropy is just like Shannon entropy, only that the probability distribution p_i is given by the Boltzmann distribution of the relevant degrees of freedom (position and momentum):

$$\rho(p,q) = \frac{1}{Z} e^{-E(p,q)/kT} \ , \tag{5}$$

and the thermodynamic quantity is made dimensional by multiplying Shannon's dimensionless uncertainty by Boltzmann's constant. It should not worry us that the degrees of freedom in thermodynamics are continuous, because any particular measurement device that is used to measure these quantities will have a finite resolution, rendering these variables effectively discrete through coarse-graining. More importantly, equilibrium thermodynamics assumes that all entropies of isolated systems are at their maximum, so there are no correlations in equilibrium thermodynamic systems, and therefore there is *no information*. This is important for our purposes, because it implies, a fortiori, that the information stored in biological genomes guarantees that living systems are far away from thermodynamical equilibrium. Information theory can thus be viewed as a type of non-equilibrium thermodynamics.

Before exploring the uses of these concepts in molecular biology, let me reiterate the most important points which tend to be obscured when discussing information. Information is defined as the amount of correlation between two systems. It measures the amount of entropy *shared* between two systems, and

this shared entropy is the information that one system has *about the other*. Perhaps this is the key insight that I would like to convey: Information is always *about something*. If it cannot be specified what the information is about, then we are dealing with entropy, not information. Indeed, entropy is sometimes called, in what borders on an abuse of language, "useless information". The previous discussion also implies that information is only defined *relative* to the system it is information about, and is therefore *never* absolute. This will be particularly clear in the discussion of the information content of genomes, which we now enter.

2   Information in Genomes

There is a long history of applying information theory to symbolic sequences. Most of this work is concerned with the randomness, or, conversely, regularity, of the sequence. Ascertaining the probabilities with which symbols are found on a sequence or message will allow us to estimate the entropy of the *source of symbols*, but not what they stand for. In other words, information cannot be accessed in this manner. It should be noted, however, that studying *horizontal* correlations, i.e., correlations between symbols along a sequence rather than across sequences, can be useful for distinguishing coding from non-coding regions in DNA (Grosse et al., 2000), and can serve as a distance measure between DNA sequences that can be used to assemble fragments obtained from shotgun-sequencing (Otu and Sayood, 2003).

In terms of the jargon introduced above, measuring the probabilities with which symbols (or groups of symbols) appear *anywhere* in a sequence will reveal the *marginal* entropy of the sequence, i.e., the entropy *without* regard to the environment or context. The entropy *with* regard to the environment is the entropy *given* the environment, a conditional entropy, which we shall calculate below. This will involve obtaining the probability to find a symbol at a *specific* point in the sequence, as opposed to anywhere on it. We sometimes refer to this as obtaining the *vertical* correlations between symbols.

2.1   Meaning from Correlations

Obtaining the marginal entropy of a genetic sequence can be quite involved (in particular if multi-symbol probabilities are required), but a very good approximative answer can be given without any work at all: This entropy

6

(for DNA sequences) is about two bits per base. There are deviations of interest (for example in GC-rich genes, etc.) but overall this is what the (non-conditional) entropy of most of DNA is (see, e.g., Schmitt and Herzel 1997). The reason for this is immediately clear: DNA is a *code*, and codes do not reveal information from sequence alone. Optimal codes, e.g., are such that the encoded sequences cannot be compressed any further (Cover and Thomas, 1991). While DNA is not optimal (there are some correlations between symbols along the sequence), it is nearly so. The same seems to hold true for proteins: a random protein would have $\log_2(20) = 4.32$ bits of entropy per site (or 1 mer, the entropy of a random monomer introduced above), while the actual entropy is somewhat lower due to biases in the overall abundance (leucine is over three times as abundant as tyrosine, for example), and due to pair and triplet correlations. Depending on the data set used, the protein entropy per site is between 2.5 (Strait and Dewey, 1996) and 4.17 bits (Weiss et al., 2000), or between 0.6 and 0.97 mers. Indeed, it seems that protein sequences can only be compressed by about 1% (Weiss et al. 2000). This is a pretty good code! But this entropy per symbol only allows us to quantify our uncertainty about the sequence identity, but it will not reveal to us the *function* of the genes. If this is all that information theory could do, we would have to agree with the critics that information theory is nearly useless in molecular biology. Yet, I have promised that information theory *is* relevant, and I shall presently point out how. First of all, let us return to the concept of information. How should we decide whether or not *potential information* (a.k.a entropy) is in *actuality* information, i.e., whether it is shared with another variable?

The key to information lies in its use to make predictions *about* other systems. Only in *reference* to another ensemble can entropy become information, i.e., be promoted from useless to useful, from potential to actual. Information therefore is clearly not stored *within* a sequence, but rather in the *correlations* between the sequence and what it describes, or what it *corresponds to*. What do biomolecular sequences correspond to? What is the *meaning* of a genomic sequence, what information does it represent? This depends, quite naturally, on what environment the sequence is to be interpreted within. According to the arguments advanced here, no sequence has an intrinsic meaning, but only a relative (or conditional) one with respect to an environment. So, for example, the genome of *Mycoplasma pneumoniae* (a bacterium that causes pneumonia-like respiratory illnesses) has an entropy of almost a million base pairs, which is its genome length. Within the soft

7

tissues that it relies on for survival, most of these base pairs (about 89%) are information (Dandekar et al., 2000). Indeed, Mycoplasmas are obligate parasites in these soft tissues, having shed from 50% to three quarters of the genome of their bacterial ancestors (the *Bacillae*). Within these soft tissues that make many metabolites readily available, what was information for a Bacillus had become entropy for the Mycoplasma. With respect to *other* environments, the Mycoplasma information might mean very little, i.e., it might not *correspond* to anything there. Whether or not a sequence means something in its environment determines whether or not the organism hosting it lives or dies there. This will allow us to find a way to distinguish entropy from information in genomes.

2.2   Physical Complexity

In practice, how can we determine whether a particular base's entropy is shared, i.e., whether a nucleotide carries entropy or information? At first glance one might fear that we would have to know a gene's function (i.e., know what it corresponds to within its surrounding) before we can determine the information content; that, for example, we might need to know that a gene codes for an alcoholdehydrogenase before we can ascertain which base pairs code for it. Fortunately, this is not true. What is clear, however, is that we may never distinguish entropy from information if we are only given a *single* sequence to make this determination, because, in a single sequence, symbols that carry information are indistinguishable from those that do not. The trick lies in studying *functionally equivalent sets* of sequences, and the substitution patterns at each aligned position. In an equilibrated population, i.e, one where sufficient time has passed since the last evolutionary innovation or bottleneck, we expect a position that codes for information to be nearly uniform *across* the population (meaning that the same base pair will be found at that position in all sequences of that population), because a mutation at that position would detrimentally affect the fitness of the bearer, and, over time, be purged from the ensemble (this holds in its precise form only for asexual populations). Positions that do not code for information, on the other hand, are selectively neutral, and, with time, will take on all possible symbols at that position. Thus, we may think of each position on the genome as a four-sided die. A priori, the uncertainty (entropy) at each position is

two bits, the maximal entropy:

$$H = -\sum_{i=\text{G,C,A,T}} p(i) \log_2 p(i) = \log_2 4 = 2 \text{ bits} \tag{6}$$

because, a priori, $p(i) = 1/4$. For the *actual* entropy, we need the actual probabilities $p_j(i)$, for each position j on the sequence. In a pool of N sequences, $p_j(i)$ is estimated by counting the number $n_j(i)$ of occurrences of nucleotide i at position j, so that $p_j(i) = n_j(i)/N$. This should be done for all positions $j = 1, ..., L$ of the sequence, where L is the sequence length. Ignoring correlations *between* positions j on a sequence (so-called "epistatic" correlations, to which we shall return below), the information stored in the sequence is then (with logs to base 2)

$$I = H_{\max} - H = 2L - H \text{ bits} , \tag{7}$$

where

$$H = -\sum_{j=1}^{L} \sum_{i=\text{G,C,A,T}} p_j(i) \log_2 p_j(i) . \tag{8}$$

Note that this estimate, because it relies on the difference of maximal and actual entropy, does not require us to know which variables in the environment cause some nucleotides to be uniform, or "fixed". These probabilities are set by mutation-selection balance in the environment. I have argued earlier (Adami and Cerf 2000, Adami et al. 2000) that the information stored in a sequence is a good proxy for the sequences's complexity (called "physical complexity"), which itself might be a good predictor of functional complexity. And indeed, it seems to correspond to the quantity that increases during Darwinian evolution (Adami 2002a). We will encounter below an evolutionary experiment that seems to corroborate these notions.

In general (for sequences taken from any monomer alphabet of size D), the information stored in the sequence is

$$I = H_{\max} - H = L - \left(-\sum_{j=1}^{L} \sum_{i=1}^{D} p_j(i) \log_D p_j(i) \right) \tag{9}$$

$$= L - J \text{ mers} , \tag{10}$$

where J can be thought of as the number of *non-functional* (i.e., "junk") instructions, and I remind the reader that we defined the "mer" as the entropy of a random monomer, normalized to lie between zero and one.

2.3   Application to DNA and RNA

In the simplest case, the environment is essentially given by the intra-cellular binding proteins, and the measure (7) can be used to investigate the information content of DNA binding sites (this use of information theory was pioneered by Schneider et al., 1986). Here, the sample of sequences can be provided by a sample of equivalent binding sites within a single genome. For example, the latter authors aligned the sequences of 149 *E. coli* and coliphage ribosome binding sites in order to calculate the substitution probabilities at each position of a 44 base pair region (which encompasses the 34 positions that can be said to constitute the binding site). Fig. 1 shows the information
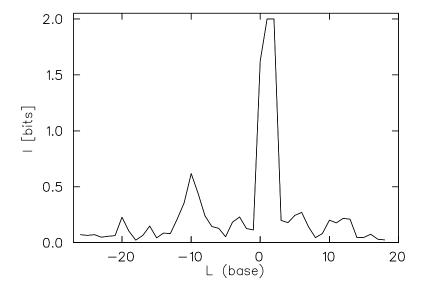

Figure 1: Information content (in bits) of an *E. coli* ribosome binding site, aligned at the fMet-tRNA$_f$ initiation site (L=0), from Schneider et al. (1986).

content as a function of position (Schneider et al. 1986), where position $L = 0$ is the first base of the initiation codon. The information content is highest near the initiation codon, and shows several distinct peaks. The peak at $L = -10$ corresponds to the Shine-Dalgarno sequence (Shine and Dalgarno, 1974).

When the information content of a base is zero we must assume that it has no function, i.e., it is neither expressed nor does anything bind to

10

it. Regions with positive information content[1] carry information about the binding protein, just as the binding protein carries information about the binding site.

It is important to emphasize that the reason that sites $L = 1$ and $L = 2$, for example, have maximal information content is a consequence of the fact that their *conditional* entropy Eq. (8) vanishes. The entropy is conditional because only *given* the environment of binding proteins in which it functions in *E. coli* or a coliphage, is the entropy zero. If there were, say, two different proteins which could initiate translation at the same site (two different environments), the conditional entropy of these sites could be higher. Intermediate information content (between zero and 2 bits) signals the presence of *polymorphisms* implying either non-specific binding to one protein or competition between more than one protein for that site.

A polymorphism is a deviation from the consensus sequence that is not, as a rule, detrimental to the organism carrying it. If it was, we would call it a "deleterious mutation" (or just "mutation"). The latter should be very infrequent as it implies disease or death for the carrier. On the contrary, polymorphisms can establish themselves in the population, leading either to no change in the phenotype whatsoever, in which case we may term them "strictly neutral", or they may be deleterious by themselves but neutral if associated with a commensurate (compensatory) mutation either on the same sequence or somewhere else.

Polymorphisms are easily detected if we plot the per-site entropies of a sequence vs. residue or nucleotide number in an *entropy map* of a gene. Polymorphisms carry per-site entropies intermediate between zero (perfectly conserved locus) and unity (strictly neutral locus). Mutations, on the other hand, (because they are deleterious) are associated with very low entropy (Rogan and Schneider 1995), so polymorphisms stand out among conserved regions and even mutations. In principle, mutations can occur on sites which are themselves polymorphic; those can only be detected by a more in-depth analysis of substitution patterns such as suggested in Schneider (1997). Because polymorphic sites in proteins are a clue to which sites can easily be mutated, per-site entropies have also been calculated for the directed evolution of proteins and enzymes (Saven and Wolynes 1997, Voigt et al. 2001).

---

[1]Finite sampling of the substitution probabilities introduces a systematic error in the information content, which can be corrected (Miller 1954, Basharin 1959, Schneider et al. 1986). In the present case, the correction ensures that the information content is
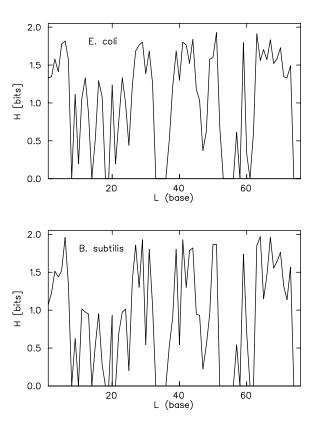
Figure 2: Entropy (in bits) of *E. coli* tRNA (upper panel) from 5' (L=0) to 3' (L=76), from 33 structurally similar sequences obtained from Sprinzl et al. (1996), where we arbitrarily set the entropy of the anti-codon to zero. Lower panel: Same for 32 sequences of *B. subtilis* tRNA.

As mentioned earlier, the actual function of a sequence is irrelevant for determining its information content. In the previous example, the region investigated was a binding site. However, any gene's information content can be measured in such a manner. In Adami and Cerf (2000), the information content of the 76 base pair nucleic acid sequence that codes for bacterial tRNA was investigated. In this case the analysis is complicated by the fact that the tRNA sequence displays secondary and tertiary structure, so that the entropy of those sites that bind in Watson-Crick pairs, for example, are shared, reducing the information content estimate based on Eq. (2) significantly. In Fig. 2, I show the entropy (in bits) derived from 33 structurally

---

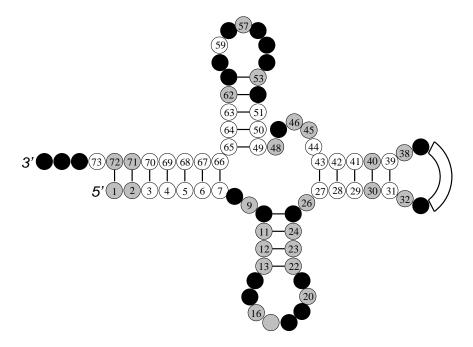approximately zero at the left and right edge of the binding site.

Figure 3: Secondary structure of tRNA molecule, with bases colored black for low entropy ($0 \leq H \leq 0.3$ mers), grey for intermediate ($0.3 < H \leq 0.7$ mers), and white for maximal entropy ($0.7 < H \leq 1.0$ mers), numbered 1-76 (entropies from *E. coli* sequences).

similar sequences of *E. coli* tRNA (upper panel) and 32 sequences of *B. subtilis* tRNA, respectively, obtained from the EMBL nucleotide sequence library (Sprinzl et al. 1996). Note how similar these entropy maps are across species (even though they last shared an ancestor over 1.6 billion years ago), indicating that the profiles are characteristic of the *function* of the molecule, and thus statistically stable.

Because of base-pairing, we should not expect to be able to simply sum up the per-site entropies of the sequence to obtain the (conditional) sequence entropy. The pairing in the stacks (the ladder-like arrangement of bases that bind in pairs) of the secondary structure (see Fig. 3) reduces the actual entropy, because two nucleotides that are bound together *share* their entropy. This is an example where *epistatic correlations* are important. Two sites (loci) are called epistatic if their contributions to the sequence's fitness are not independent, in other words, if the probability to find a particular base at one position depends on the identity of a base at another position. Watson-

Crick-binding in stacks is the simplest such example; it is also a typical example of the maintenance of polymorphisms in a population because of functional association. Indeed, the fact that polymorphisms are correlated in stacks makes it possible to deduce the secondary structure of an RNA molecule from sequence information alone. Take, for example, nucleotide
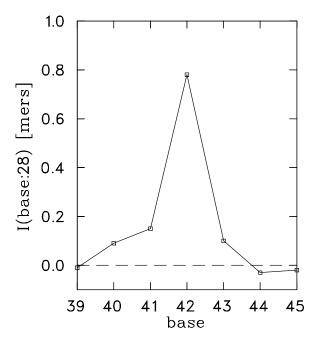

Figure 4: Mutual entropy (information) between base 28 and bases 39 to 45 (information is normalized to $I_{max} = 1$ by taking logarithms to base 4). Because finite sample size corrections of higher order have been neglected, the information estimate can appear to be negative by an amount of the order of this error.

$L = 28$ (in the anti-codon stack) which is bound to nucleotide $L = 42$, and let us measure entropies in mers (by taking logarithms to the base 4). The mutual entropy between $L = 28$ and $L = 42$ (in *E. coli*) can be calculated using Eq. (4):

$$I(28 : 42) = H(28) + H(42) - H(28, 42) = 0.78 . \qquad (11)$$

Thus indeed, these two bases share almost all of their entropy. We can see furthermore that they share very little entropy with any other base. Note that, in order to estimate the entropies in Eq. (11), we applied a first-order

correction that takes into account a bias due to the finite size of the sample, as described in Miller (1954). This correction amounts to $\Delta H_1 = 3/(132 \ln 2)$ for single nucleotide entropies, and $\Delta H_2 = 15/(132 \ln 2)$ for the joint entropy. In Fig. 4, I plot the mutual entropy of base 28 with bases 39 to 45 respectively, showing that base 42 is picked out unambiguously. Such an analysis can be carried out for all pairs of nucleotides, so that the secondary structure of the molecule is revealed unambiguously (see, e.g., Durbin et al. 1998). In Fig. 5, I show the entropy (in bits) for all pairs of bases of the set of *E. coli* sequences used to produce the entropy map in Fig. 2, which demonstrates how the paired bases in the four stems stand out.
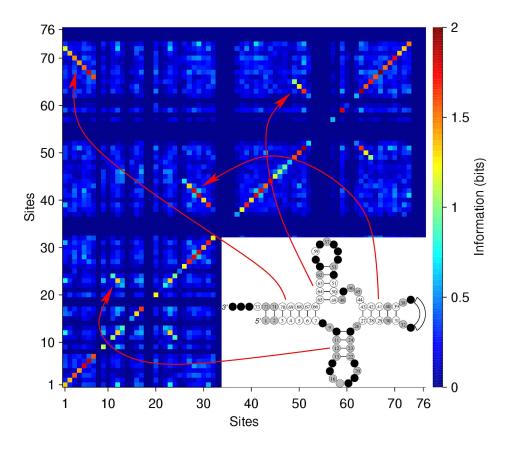

Figure 5: Mutual entropy (information) between all bases (in bits), colored according to the color bar on the right, from 33 sequences of *E. coli* tRNA. The four stems are readily identified by their correlations as indicated.

15

Since we found that most bases in stacks share almost all of their entropy with their binding partner, it is easy to correct formula (10) to account for the epistatic effects of stack-binding: We only need to subtract from the total length of the molecule (in mers) the number of bases involved in stack binding. In a tRNA molecule (with a secondary structure as in Fig. 3) there are 21 such bases, so the sum in Eq. (8) should only go over the 52 "reference positions"[2]. For *E. coli*, the entropy summed over the reference positions gives $H \approx 24$ mers, while the *B. subtilis* set gives $H \approx 21$ mers. We thus conclude that bacterial tRNA stores between 52 and 55 mers of information about its environment (104-110 bits).

This type of sequence analysis combining structural and complexity information has recently been used to quantify the information gain during in-vitro evolution of catalytic RNA molecules (ribozyme ligases) (Carothers et al. 2004). The authors evolved RNA aptamers that bind GTP (guanine triphosphate) with different catalytic effectiveness (different functional capacity) from a mutagenized sequence library. They found 11 different classes of ribozymes, whose structure they determined using the correlation analysis outlined above. They subsequently measured the amount of information in each structure [using Eq. (7) and correcting for stack binding as described above] and showed that ligases with higher affinity for the substrate had more complex secondary structure *and* stored more information. Furthermore, they found that the information estimate based on Eq. (7) was consistent with an interpretation in terms of the amount of information necessary to specify the particular structure in the given environment. Thus, at least in this restricted biochemical example, structural, functional, and informational complexity seem to go hand in hand.

2.4   Application to Proteins

If the secondary structure of RNA and DNA enzymes can be predicted based on correlations alone, what about protein secondary structure? Because proteins fold and function via the interactions among the amino acids they are made of, these interactions should, in evolutionary time, lead to correlations between residues so that the fitness effect of an amino acid substitution at

---

[2]We exclude the three anticodon-specifying bases from the entropy calculation because they have zero conditional entropy by *definition* (they cannot vary among a tRNA-type because it would change the type). However, the substitution probabilities are obtained from mixtures of *different* tRNA-types, and therefore appear to deviate from zero or one.

one position will depend on the residue at another position. (Care must be taken to avoid contamination from correlations that are due entirely to a common evolutionary path, see Wollenberg and Atchley 2000; Govindarajan et al. 2003.) Such an analysis has been carried out on a number of different molecule families, such as the V3 loop region of HIV-1 (Korber et al. 1993), which shows high variability (high entropy) and strong correlations between residues (leading to shared entropy) that are due to functional constraints. These correlations have also been modelled (Giraud et al. 1998).

A similar analysis for the homeodomain sequence family was performed by Clarke (1995), who was able to detect 16 strongly co-varying pairs in this 60 amino acid binding motif. However, determining secondary structure based on these correlations alone is much more difficult, because proteins do not fold neatly into stacks and loops as does RNA. Also, residue covariation does not necessarily indicate physical proximity (Clarke 1995), even though the strongest correlations are often due to salt-bridges. But the correlations can at least help in eliminating some models of protein structure (Clarke 1995).

Atchley et al. (2000) carried out a detailed analysis of correlated mutations in the bHLH (basic helix-loop-helix) protein domain of a family of transcription factors. Their set covered 242 proteins across a large number of vertebrates that could be aligned to detect covariation. They found that amino acid sites known to pack against each other showed low entropy, whereas exposed non-contact sites exhibited significantly larger entropy. Furthermore, they determined that a significant amount of the observed correlations between sites was due to functional or structural constraints that could help in elucidating the structural, functional, and evolutionary dynamics of these proteins (Atchley et al. 2000).

Some attempts have been made to study the *thermodynamics* of protein structures and relate it to the sequence entropy (Dewey 1997), by studying the mutual entropy between protein sequence and *structure*. This line of thought is inspired by our concept of the genotype-phenotype map, which implies that sequence should predict structure. If we hypothesize a structural entropy of proteins $H(\mathrm{str})$, obtained for example as the logarithm of the possible stable protein structures for a given chain length (and a given environment), then we can write down the mutual entropy between structure and sequence simply as

$$I(\mathrm{seq} : \mathrm{str}) = H(\mathrm{seq}) - H(\mathrm{seq}|\mathrm{str}) \; , \tag{12}$$

where $H(\text{seq})$ is the entropy of sequences of length L, given by L, and $H(\text{seq}|\text{str})$ is the entropy of sequences *given* the structure. If we assume that the environment perfectly dictates structure (i.e., if we assume that only one particular structure will perform any given function) then

$$H(\text{seq}|\text{str}) \approx H(\text{seq}|\text{env}) \tag{13}$$

and $I(\text{str} : \text{seq})$ is then roughly equal to the physical complexity defined earlier. Because $H(\text{str}|\text{seq}) = 0$ (per the above assumption that any given sequence gives rise to exactly one structure), we can rewrite (12) as

$$I(\text{seq} : \text{env}) \approx I(\text{seq} : \text{str}) = H(\text{str}) - \underbrace{H(\text{str}|\text{seq})}_{=0}, \tag{14}$$

i.e., the mutual entropy between sequence and structure only tells us that the thermodynamical entropy of possible protein structures is limited by the amount of information about the environment coded for by the sequence. This is interesting because it implies that sequences that encode more information about the environment are also potentially more complex, a relationship we discussed earlier in connection with ribozymes (Carothers et al. 2004). Note, however, that the assumption that only one particular structure will perform any given function need not hold. Szostak (2003), for example, advocates a definition of *functional information* that allows for different structures carrying out an equivalent biochemical function.

3 Molecular Interactions and Resistance

One of the more pressing concerns in bioinformatics is the identification of DNA protein-binding regions, such as promoters, regulatory regions, and splice junctions. The common method to find such regions is through *sequence identity*, i.e., known promoter or binding sites are compared to the region being scanned (e.g., via freely available bioinformatics software such as BLAST), and a "hit" results if the scanned region is sufficiently identical according to a user-specified threshold. Such a method cannot, of course, find *unknown* binding sites, nor can it detect interactions between proteins, which is another one of bioinformatics' holy grails (see, e.g., Tucker et al. 2001). Information theory can in principle detect interactions between different molecules (such as DNA-protein or protein-protein interactions) from *sequence heterogeneity*, because interacting pairs share *correlated mutations*, that arise as follows.

18

3.1 Detecting Protein-Protein and DNA-Protein Interactions

Imagine two proteins bound to each other, while each protein has some entropy in its binding motif (substitutions that do not affect structure). If a mutation in one of the proteins leads to severely reduced interaction specificity, the substitution is strongly selected against. It is possible, however, that a *compensatory* mutation in the binding partner restores specificity, such that the *pair* of mutations together is neutral (and will persist in the population), while each mutation by itself is deleterious. Over evolutionary time, such pairs of correlated mutations will establish themselves in populations and in homologous genes across species, and could be used to identify interacting pairs. This effect has been seen previously in the Cytochrome c/Cytochrome oxidase (CYC/COX) heterodimer (Rawson and Burton 2002) of the marine copepod *Tigriopus californicus*. In Rawson and Burton (2002), the authors performed crosses between the San Diego (SD) and Santa Cruz (SC) variants from two natural allopatric populations that have long, independent evolutionary histories. Inter-population crosses produced strongly reduced activity of the cytochrome complex, while intra-population crosses were vigorous. Indeed, the SD and SC variants of COX differ by at least 30 amino acid substitutions, while the smaller CYC has up to 5 substitutions. But can these correlated mutations be found from sequence data alone? This turns out to be a difficult computational problem unless it is known precisely which member of a set of N sequences of one binding partner binds to which member of a set of N of the other. Unless we are in possession of this N to N assignment, we cannot calculate the joint probabilities p_{ij} that go into the calculation of the mutual entropies such as Eq. (11) that reveal correlated mutations.

Of course, if we have one pair of sequences from N species of organisms with the same homologous gene, the assignment is automatically implied. In the absence of such an assignment, it may be possible to recover the correct matches from two sets of N sequences by searching for the assignment with the highest mutual entropy, because we can safely assume that the correct assignment maximizes the correlations (Adami and Thomsen 2004). However, this is a difficult search problem because the number of possible assignments scales like $N!$. Still, because correlated mutations due to coevolution seem to be relatively common (Bonneton et al. 2003), this would be a useful tool for revealing those residues involved in binding, or even in protein-protein

19

interaction prediction.

In principle, the information-theoretical method described above can potentially identify *unknown* binding sites by identifying complementary patterns (between binding sites and protein coding regions), if the binding regions are not well-conserved, i.e., when the binding site and the corresponding transcription factor carry a reasonable amount of polymorphisms, and if enough annotation exists to identify the genomic partners that correspond to each other in a set. If sufficient pairs of transcription-factor/binding-domain pairs can be sequenced, an information-theoretic analysis could conceivably reveal genomic regulatory regions that standard sequence analysis methods miss. For example, it was suggested recently (Brown and Callan, 2004) that the cAMP response protein (CRP, a transcription factor that regulates many *E. coli* genes) binds to a number of entropic sites in *E.coli*, i.e., sites that are not strictly conserved, but that still retain functionality (see also Berg and von Hippel, 1987).
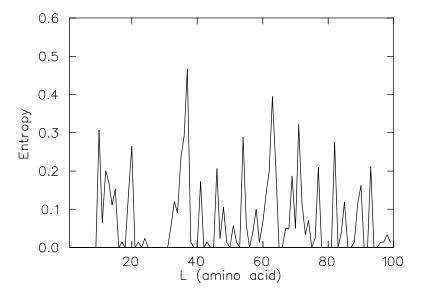

Figure 6: Normalized ($0 \leq H \leq 1$) entropy of HIV-1 protease in mers, as a function of residue number, using 146 sequences from patients exposed to a protease inhibitor drug (entropy is normalized to $H_{\max} = 1$ per amino acid by taking logarithms to base 20).

3.2   Tracking Drug Resistance

An interesting pattern of mutations can be observed in the protease of HIV-1, a protein that binds to particular motifs on a virus polyprotein, and then cuts it into functional pieces. Resistance to protease inhibitors (small molecules designed to bind to the "business end" of the protease, thereby preventing its function) occurs via mutations in the protease that do not change the protease's cutting function (proteolysis), while preventing the inhibitor to bind to it. Information theory can be used to study whether mutations are involved in drug resistance or whether they are purely neutral, and to discover correlated resistance mutations.

The emergence of resistance mutations in the protease after exposure to antiviral drugs has been well studied (Molla et al. 1996, Schinazi, Larder, and Mellors 1999). The entropy map of HIV protease in Fig. 6[3] (on the level of amino acids) reveals a distinctive pattern of polymorphisms and only two strictly conserved regions. HIV protease *not* exposed to inhibitory drugs, on the other hand, shows three such conserved regions (Loeb et al. 1989). It is believed that the polymorphisms contribute to resistance mutations involved in HAART (Highly Active Antiretroviral Therapy) failure patients (Servais et al. 2001b). But, as a matter of fact, many of the observed polymorphisms can be observed in treatment-naive patients (Kozal et al. 1996, and Lech et al. 1996) so it is not immediately clear which of the polymorphic sites are involved in drug resistance.

In principle, exposure of a population to a new environment can lead to fast adaptation if the mutation rate is high enough. This is certainly the case with HIV. The adaptive changes generally fall into two classes: mutations in regions that were previously conserved (true resistance mutations), and changes in the substitution pattern on sites that were previously polymorphic. In the case of HIV-1 protease, both patterns seem to contribute. In Fig. 7, I show the *changes* in the entropic profile of HIV-1 protease obtained from a group of patients before and six months after treatment with high doses of saquinavir (a protease inhibitor). Most spikes are positive, in particular the changes around residues 46-56, a region that is well-conserved in treatment-naive proteases, and that is associated with a *flap* in the molecule that must be flexible and that extends over the substrate binding cleft (Shao et al. 1997). Mutations in that region indeed appeared on sites that were

---

[3]The map was created using 146 sequences obtained from a cohort in Luxembourg, and deposited in GenBank (Servais et al. 1999 and 2001a).

previously uniform, while some changes occurred on polymorphic sites (negative spikes). For those, exposure to the new environment usually reduced the entropy at that site.

Some of the resistance mutations actually appear in pairs, indicating that they may be compensatory in nature (Leigh Brown et al. 1999, Hoffman et al. 2003, Wu et al. 2003). The strongest association occurs between residues 54 and 82, the former associated with the flap, and the latter right within the binding cleft. This association does not occur in treatment-naive patients, but stands out strongly after therapy (such correlations are easily detected by creating mutual entropy graphs such as Fig. 5, data not shown). The common explanation for this covariation is again compensation: while a mutation in the flap or in the cleft leads to reduced functionality of the protease, both together restore function while evading the inhibitor.
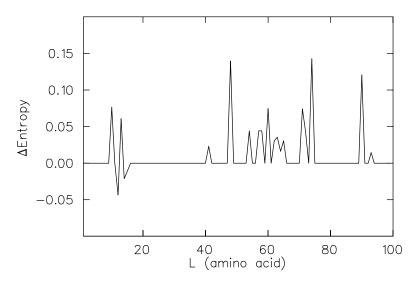

Figure 7: Change in per-site entropy of HIV-1 protease after six months of exposure to saquinavir, ΔEntropy$=H_{26} - H_0$, where H_{26} is the entropy after 26 weeks of exposure. The entropies were obtained from 34 sequences before and after exposure, available through GenBank (Schapiro et al. 1997). The three highest (positive) spikes are associated to the well-known resistance mutations G48V, T74(A,S), and L90M, respectively.

3.3 Information-theoretic Drug Design

Because many of the protease polymorphisms are prevalent in treatment-naive patients, we must assume that they are either neutral, or that the steric changes they entail do not impede the protease's proteolytic activity while failing to bind the protease inhibitor. Thus, a typical protease population is a mixture of polymorphic molecules (polymorphic both in genotype and in structure, see Maggio et al. 2002) that can outsmart a drug designed for a single protease type relatively easily. An interesting alternative in drug design would therefore use an entropic mixture of polymorphisms, or "quasispecies" (Eigen 1971) as the drug target. Such a drug would *itself* form a quasispecies rather than a pure drug. Indeed, an analysis of the information content of realistic ensembles shows that consensus sequences are exceedingly rare in real populations (Schneider 1997), and certainly absent in highly variable ones such as HIV proteases. The absence of a consensus sequence is also predicted for molecules evolving at the *error threshold* (Eigen 1971), which is very likely in these viruses.

The ideal *superdrug* should represent a mixture of inhibitors that is perfectly tuned to the mixture of proteases. What this mixture is can be determined with information theory, by ensuring that the ensemble of inhibitors *co-varies* with the protease, such as to produce tight binding even in the presence of mutations (or more precisely *because* of the presence of mutations). The substitution probabilities of the inhibitor ensemble would be obtained by maximizing the mutual entropy (information) between the protease and an inhibitor library obtained by combinatorial methods, either on a nucleotide or on the amino acid level (Adami 2002b). If such a procedure could create a drug that successfully inhibits resistance mutations, we could no longer doubt the utility of information theory for molecular biology.

4 Conclusions

Information theory is not widely used in bioinformatics today even though, as the name suggests, it should be *the* relevant theory for investigating the information content of sequences. The reason for the neglect appears to be a misunderstanding of the concepts of entropy versus information throughout most of the literature, which has led to the widespread perception of its incompetence. Instead, I point out that Shannon's theory precisely defines

both entropy and information, and that our intuitive concept of information coincides with the mathematical notion. Using these concepts, it is possible in principle to distinguish information-coding regions from random ones in ensembles of genomes, and thus quantify the information content. A thorough application of this program should resolve the C-paradox, that is, the absence of a correlation between the size of the genome and the apparent complexity of an organism (Cavalier-Smith 1985), by distinguishing information that contributes to complexity from non-functional stretches that do not. However, this is a challenge for the future because of the dearth of multiply sequenced genomes.

Another possible application of information theory in molecular biology is the association of regulatory molecules with their binding sites or even protein-protein interactions, in the case where transcription factors and their corresponding binding site show a good amount of polymorphism (methods based on correlated heterogeneity), and the binding association between pairs can be established. This approach is complementary to sequence comparison of conserved regions (methods based on sequence identity), in which information theory methods cannot be used because zero (conditional) entropy regions cannot share entropy. Conversely, sequence comparison methods must fail if polymorphisms are too pronounced. Finally, the recognition of the polymorphic (or quasispecies) nature of many viral proteins suggests an information theory based approach to drug design in which the quasispecies of proteins—rather than the consensus sequence—is the drug target, by maximizing the information shared between the target and drug ensembles.

Acknowledgements

References

[] Adami, C. 2002a. What is complexity? BioEssays 24:1085-1094

[] Adami, C. 2002b. Combinatorial drug design augmented by information theory. NASA Tech Briefs 26:52.

Adami, C. and S.W. Thomsen 2004. Predicting protein-protein interactions from sequence data. In: Hicks MG, editor. *The Chemical Theatre of Biological Systems*. Proceedings of the Beilstein-Institute Workshop.

Adami, C. and N.J. Cerf 1999. Prolegomena to a non-equilibrium quantum statistical mechanics. Chaos, Solitons, and Fractals 10:1637-1650.

Adami, C., C. Ofria and T.C. Collier 2000. Evolution of biological complexity. Proc. Natl. Acad. Sci. USA 97:4463-4468.

Adami, C. and N.J. Cerf. 2000. Physical complexity of symbolic sequences. Physica D 137:62-69.

Atchley W.R., K.R. Wollenberg, W.M. Fitch, W. Terhalle, and A.D. Dress 2000. Correlations among amino acid sites in bHLH protein domains: An information-theoretic analysis. Mol. Biol. Evol. 17:164-178.

Basharin, G.P. 1959. On a statistical estimate for the entropy of a sequence of random variables. Theory Probability Appl. 4:333.

Berg, O.G. and P.H. von Hippel 1987. Selection of DNA binding sites by regulatory proteins. II. The binding specificity of cyclic AMP receptor protein to recognition sites. J. Mol. Biol. 200:709-723.

Bonneton, F., D. Zelus, T. Iwema, M. Robinson-Rechavi, and V. Laudet 2003. Rapid divergence of the Ecdysone receptor in Diptera and Lepidoptera suggests coevolution between ECR and USP-RXR. Mol. Biol. Evol. 20:541-553.

Brown, C.T. and C.G. Callan, Jr. 2004. Evolutionary comparisons suggest many novel CRP binding sites in *E. coli*. Proc. Natl. Acad. Sci. USA 101:2404-2409.

Carothers, J.M., S.C. Oestreich, J.H. Davis and J.W. Szostak 2004. Informational complexity and functional activity. J. Am. Chem. Society 126 (in press).

Cavalier-Smith, T. 1985. *The Evolution of Genome Size*, ed. Cavalier-Smith, T. (Wiley, N.Y.)

Clarke, N.D. 1995. Covariation of residues in the homeodomain sequence family. Prot. Sci. 4:2269-2278.

Cover, T.M. and J.A. Thomas 1991. *Elements of Information Theory.* (John Wiley, New York).

Dandekar, T. et al. 2000. Re-annotating the *Mycoplasma pneumoniae* genome sequence: Adding value, function, and reading frames. Nucl. Acids Res. 28:3278-3288.

Dewey, T.G. 1997. Algorithmic complexity and thermodynamics of sequence-structure relationships in proteins. Phys. Rev. E 56:4545-4552.

Durbin, R., S. Eddy, A. Krogh and G. Mitchison 1998. *Biological Sequence Analysis.* (Cambridge University Press, Cambridge MA).

Eigen, M. 1971. Self-organization of matter and the evolution of macromolecules. Naturwissenschaften 58:465.

Giraud B.G., A. Lapedes and L.C. Liu 1998. Analysis of correlations between sites of model protein sequences. Phys. Rev. E 58:6312-6322.

Govindarajan S., J.E. Ness, S. Kim, E.C. Mundorff, J. Minshull, and C. Gustafsson 2003. Systematic variation of amino acid substitutions for stringent assignment of pairwise covariation. J. Mol. Biol. 328:1061-2069.

Grosse, I., H. Herzel, S.V. Buldyrev, and H.E. Stanley 2000. Species independence of mutual information in coding and non-coding DNA. Phys. Rev. E 61:5624-5629.

Hoffman, N.G., C.A. Schiffer, and R. Swanstorm 2003. Covariation of amino acid positions in HIV-1 protease. Virology 314:536-548.

Korber B.T.M., R.M. Farber, D.H. Wolpert, and A.S. Lapedes 1993. Covariation of mutations in the V3 loop of human immunodeficiency virus type 1 envelope protein: An information-theoretic analysis. Proc. Natl. Acad. Sci. USA 90:7176-7180.

Kozal, M.J., et al. 1996. Extensive polymorphisms observed in HIV-1 clade B protease gene using high density oligonucleotide arrays. Nature Medicine 2:753-759.

Lech, W.J., et al. 1996. In vivo sequence diversity of the protease of human immunodeficiency virus type 1: Presence of protease inhibitor resistant variants in untreated subjects. J. Virol. 70:2038-2043.

[] Leigh Brown, A.J., B.T. Koerber, and J.H. Condra 1999. Associations between amino acids in the evolution of HIV type 1 protease sequences under indinavir therapy. AIDS Res. Hum. Retrovir. 15:247-253.

[] Loeb D.D., R. Swanstrom, L. Everitt, M. Manchester, S.E. Stamper, and C.A. Hutchison III 1989. Complete mutagenesis of the HIV-1 protease. Nature 340:397-400.

[] Maggio, E.T., M. Shenderovich, R. Kagan, D. Goddette, and K. Ramnarayan 2002. Structural pharmacogenetics, drug resistance, and the design of anti-infective superdrugs. Drug Discovery Today 7:1214-1220.

[] Maynard Smith, J. 1999a. The idea of information in biology. Quart. Rev. Biol. 74:395-400.

[] Maynard Smith, J. 1999b. The concept of information in biology. Philo. Sci. 67:177-194.

[] Miller, G.A. 1954. Note on the bias of information estimates. In H. Quastler, ed., *Information Theory and Psychology*, pp. 95-100 (The Free Press, Illinois).

[] Molla, A. et al. 1996. Ordered accumulation of mutations in HIV protease confers resistance to ritonavir. Nature Medicine 2:760-766.

[] Otu, H. H. and K. Sayood 2003. A divide-and-conquer approach to fragment assembly. Bioinformatics 19:22-29.

[] Quastler, H. (Ed.) 1953. *Information Theory in Biology* (University of Illinois Press, Urbana).

[] Rawson, P.D. and R.S. Burton 2002. Functional coadaptation between cytochrome c and cytochrome oxidase within allopatric populations of a marine copepod. Proc. Natl. Acad. Sci. USA 99:12955-12958.

[] Rogan, P.K. and T.D. Schneider 1995. Using information content and base frequencies to distinguish mutations from genetic polymorphisms. Hum. Mut. 6:74-76.

[] Sarkar, S. 1996. Decoding "coding"—information and DNA. BioScience 46:857-864.

[] Saven, J.G., and P.G. Wolynes 1997. Statistical mechanics of the combinatorial synthesis and analysis of folding macromolecules. J. Phys. Chem. B 101:8375-8389.

[] Schapiro, J.M., M.A. Winters, F. Stewart, B. Efron, J. Norris, M.J. Kozal, and T.C. Mergan 1997. The effect of high-dose saquinavir on viral load and CD4+ T-cell counts in HIV-infected patients (unpublished).

[] Schinazi, R.F., B. A. Larder, and J.W. Mellors 1999. Mutation in retroviral genes associated with drug resistance: 1999-2000 update. Intern. Antiviral News 7.4:46-68.

[] Schneider T.D., G.D. Stormo, L. Gold, and A. Ehrenfeucht 1986. Information content of binding sites on nucleotide sequences. J. Mol. Biol. 188:415-431.

[] Schmitt A.O. and H. Herzel 1997. Estimating the entropy of DNA sequences. J. theor. Biol. 188:369-377.

[] Schneider, T.D. 1997. Information content of individual genetic sequences. J. theor. Biol. 189:427-441.

[] Servais, J. et al. 1999. The natural polymorphism of the HIV-1 protease gene in treatment naive patients and response to combination therapy including a protease inhibitor. 4th European Conference on Experimental AIDS Research, June 18-21, Tampere, Finland.

[] Servais, J., et al. 2001a. Comparison of DNA sequencing and a line probe assay for detection of Human Immunodeficiency Virus type 1 drug resistance mutations in patients failing highly active antiretroviral therapy. J. Clin. Microbiol. 39:454-459.

[] Servais, J., et al. 2001b. Variant HIV-1 proteases and response to combination therapy including a protease inhibitor. Antimicrob. Agents Chemother. 45:893-900.

[] Shannon, C.E. 1948. A mathematical theory of communication. Bell Systems Technical Journal 27:379-423; *ibid* 27:623-656; reprinted in *C.E. Shannon: Collected Papers*, N.J.A. Sloane and A.D. Wyner, eds., IEEE Press (1993).

[] Shao, W., L. Everitt, M. Manchester, D.L. Loeb, C.A. Hutchison III, and R. Swanstrom 1997. Sequence requirements of the HIV-1 protease flap region determined by saturation mutagenesis and kinetic analysis of flap mutants. Proc. Natl. Acad. Sci. USA 94:2243-2248.

[] Shine, J., and L. Dalgarno 1974. The 3'-terminal sequence of *E. coli* 16s ribosomal RNA: Complementarity to nonsense triplets and ribosome binding sites. Proc. Natl. Acad. Sci. USA 71:1342-1346.

[] Sprinzl, M., C. Steegborn, F. Hübel, and S. Steinberg 1996. Compilation of tRNA sequences and sequences of tRNA genes. Nucl. Acids. Res. 23:68-72.

[] Strait, B.J. and T.G. Dewey 1996. The Shannon information entropy of protein sequences. Biophysical Journal 71:148-155.

[] Szostak, J.W. 2003. Molecular messages: Functional information. Nature 423:689.

[] Tucker, C.L., J.F. Gera, and P. Uetz 2001. Towards an understanding of complex protein interaction networks. Trends Cell Biol. 11:102-106.

[] Vincent, L.-M. 1994. Réflexions sur l'usage, en biologie, de la théorie de l'information. Acta Biotheoretica 42:167-179.

[] Voigt, C.A., S.L. Mayo, F.H. Arnold, and Z.-G. Wang. 2001. Computational method to reduce the search space for directed protein evolution. Proc. Natl. Acad. Sci. USA 98:3778-3783.

[] Weiss O., M.A. Jimenez-Montaño, and H. Herzel 2000. Information content of protein sequences. J. theor. Biol. 206:379-386.

[] Wollenberg, K.R. and W.R. Atchley 2000. Separation of phylogenetic and functional associations in biological sequences by using the parametric bootstrap. Proc. Natl. Acad. Sci. USA 97:3288-3291.

[] Wu T.D. et al. 2003. Mutation patterns and structural correlates in human immunodeficieny virus type 1 protease following different protease inhibitor treatments. J. Virol. 77:4836-4847.