

# INFERRING PROPERTY SELECTION PRESSURE FROM POSITIONAL RESIDUE CONSERVATION

Rose Hoberman\*, Judith Klein-Seetharaman<sup>†</sup>\*, and Roni Rosenfeld\*

\*School of Computer Science, Carnegie Mellon University, Pittsburgh, PA; 412-268-7678

<sup>†</sup>University of Pittsburgh Medical School; Pittsburgh, PA; 412-383-7325

{roseh, roni, judithks}@cs.cmu.edu

**Keywords:** protein sequence analysis, variability, evolutionary conservation, physical-chemical properties, sequence alignment, selective pressure, entropy, Nef, Env, HIV-1

## Abstract

This work attempts to understand and explain positional selection pressure in terms of underlying physical and chemical properties. We propose a set of constraining assumptions about how these pressures behave, then describe a procedure for analyzing and explaining the distribution of residues at a particular position in a multiple sequence alignment. In contrast to previous approaches, our model takes into account both amino acid frequencies and a large number of physical-chemical properties. By analyzing each property separately, we are able to identify positions where an unusual conservation pattern is present. In addition, our model can easily incorporate sequence weights that adjust for bias in the sample sequences. Finally, we provide a measure of statistical significance for our conservation measure. We demonstrate the applicability of our method on two HIV-1 proteins: Nef and Env. Access to the data and results presented in this paper are available at <http://flan.blm.cs.cmu.edu>

## 1 Introduction

### 1.1 Motivation

Multiple sequence alignments (MSAs) contain a wealth of information about the structure, function, and evolution of proteins. In particular, residue conservation has long been associated with functional and structural significance. MSAs therefore provide important guidance in designing laboratory experiments such as studying the effects of site-directed mutants at the conserved sites. Due to the explosion in the amount of protein data that is available, quantitative measures of positional conservation have been devised. Measures such as entropy are used to systematically evaluate the distribution of amino acids in each column in a MSA, assigning a numerical value to each position to quantify its degree of conservation. These quantitative measures have been used to identify functionally or structurally important residues and to predict the structure and function of the entire protein [1, 2, 3]. Other uses for these measures include building motif-based models of protein families or domains [4]. Measures of conservation have also been used to detect functional surfaces, such as active sites, ligand binding sites or protein-protein interaction domains [5, 6, 7]. Finally, conservation measures can provide a way to evaluate and refine multiple sequence alignments [8].

While many measures have been proposed for quantifying positional conservation (see [9] for a review), identifying positional conservation is only a first step. A more difficult task is to understand the *specific* selective pressures that have influenced the distribution of residues at each position. Different positions have different functional and structural roles, and thus the amino acid in each position will be subject to different constraints. Biologists often manually analyze a MSA by using software packages that color positions in an alignment according to the types of amino acids present. These packages use simple heuristics to determine the color of each position, such as grouping amino acids by one or a few properties, and coloring the column according to the property that is most often represented. However, these methods cannot be used in any systematic way. They consider only a few possible conservation patterns and do not provide any statistical methodology for testing whether the apparent conservation of a particular property in a position is statistically significant.

The goal of this paper is to develop a methodology for identifying and quantifying the selection pressures acting on amino acid properties in each position in a MSA. We propose a model that analyzes the residue distribution at a

particular position in a MSA and explains sequence conservation in terms of selective pressures on a specific set of underlying physical and chemical properties. Our method requires only a MSA as input, and determines the most highly conserved properties in each position, as well as their statistical significance.

## 1.2 Related Work

The most common measure of positional conservation is entropy [10, 11]. However, like other methods based purely on symbol frequencies, entropy assumes that each amino acid is an independent symbol with no relation to any of the other symbols. This assumption is clearly invalid for amino acids. Furthermore, an explicit evaluation of an entropy score showed little correlation with structural conservation [12].

A second class of measures considers only the physical-chemical properties of each amino acid. These methods [13, 14] classify amino acids based on their position in a Venn diagram of overlapping sets representing different physical and chemical properties. For each aligned position, such a method finds the smallest set of properties that explains the observed amino acids at that site. Since the number of combinations of properties is large, only a limited number of combinations is considered. The ad hoc selection of a subset of combinations is necessarily subjective, and limiting. In addition, these methods only work with binary properties and, perhaps more importantly, fail to account for amino acid frequency. For example, they do not distinguish between a position with one leucine and 100 arginine from a position with 100 leucine and 100 arginine.

A third hybrid class of approaches tries to derive entropy scores which are sensitive to the physical and chemical relationships between amino acids [9]. These methods divide the 20 amino acids into a small number of physically or chemically related clusters (about 5-10) and calculate entropy on this reduced symbol set. Again, this clustering is ad hoc, and by choosing a handful of clusters, the number of properties modeled is necessarily limited. In addition, each amino acid in a cluster is considered to be uniformly distant from the amino acids in every other cluster. Consequently, even the properties that are used for clustering are modeled at only a very coarse level.

More sophisticated approaches use mutation-based substitution matrices to quantify the physical-chemical similarity between amino acids [9]. The values in these matrices are derived by averaging over many positions and types of proteins, which can at best capture the selection pressures corresponding to the most dominant properties. While this sort of averaging yields a reasonable measure of *overall* conservation, it will blur specific pressures, making it difficult to distinguish the effect of one property over another. So, although calculating a similarity measure for each pair of amino acids is more informative than a simple clustering approach, it will inevitably ignore similarities based on particular properties that are only infrequently selected for. Furthermore, amino acids may be similar with respect to one property, but not with respect to another. Since different positions select for different properties, certain amino acids are more likely to be aligned at one position than another. However, any approach based on a single similarity matrix cannot represent such position-dependent distinctions.

Position-specific patterns of amino acid conservation have been estimated for use in building better profile HMMs of protein families. One step in building a profile HMM is to learn the probability of observing each amino acid at a given site in the protein. When the number of training sequences is small, however, the observed frequency of each amino acid is a poor estimate of the true probability. To improve these estimates Sjolander et. al. [15] have developed a regularization method based on mixtures of Dirichlets. They estimate a small number of Dirichlet distributions and mixing weights from a large set of protein alignments. These distributions often appear to be correlated with physical and chemical properties such as size, hydrophobicity, and charge. Nevertheless, the mixture of Dirichlets is used merely as a prior to smooth the posterior probability estimates, not as a means of explaining the observed amino acid counts. Given a set of Dirichlets, one can calculate the posterior probability that each of them generated the position-specific conservation pattern from which the observed amino acids were sampled. Although the Dirichlet assigned the highest probability can be said to provide the best explanation of the data, no statistical test is provided to determine if that Dirichlet provides a significant fit to the data. In addition, the number of components must be pre-selected, and even once the number is fixed there is no guarantee that the components learned are optimal.

Instead of deriving a similarity matrix indirectly from mutation data, attempts have been made to make direct use of a large set of continuous, experimentally-determined properties. These include methods that build a similarity matrix directly from physical-chemical properties and methods that use dimensionality reduction to reduce the large number of properties to a small set of prototypical, “aggregate” properties. Due to their aggregation of multiple properties, these direct approaches suffer from similar weaknesses as those based on substitution matrices. One specific model of this type is the exponential fitness model of Koshi *et al.* [16]. This method not only aggregates multiple columns, but

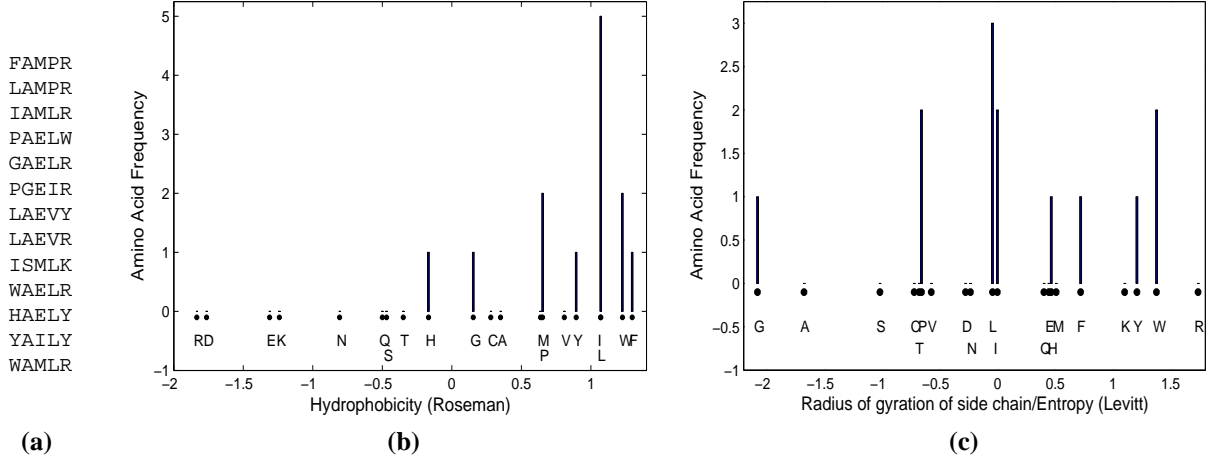


Figure 1: (a) A small sample alignment (b,c) Distribution of residues in the first position of the sample alignment, plotted against two different scales: hydrophobicity (b) and radius of gyration of side chain (c). Filled dots indicate the property values assigned to the 20 amino acids. Histogram bar heights represent frequency of the corresponding amino acid.

uses only two hybrid properties, corresponding roughly to hydrophobicity and size.

All of the previous approaches are based solely on the identities of the amino acids observed at each position in the alignment. However, the likelihood of seeing a particular amino acid at a particular position depends not only on the fitness of that amino acid, but on its mutational distance from the ancestral amino acid. As a result, any method based solely on observed amino acid frequencies will unintentionally confound mutation effects with fitness. Yampolsky and Stoltzfus [17] have attempted to decouple these two effects. A measure of amino acid exchangeability has been derived based solely on fitness measurements obtained by experimental replacement of amino acids. This work, however, is based on learning a global exchangeability matrix, and thus cannot make predictions in a position-specific manner.

## 2 Method

### 2.1 Introduction

In contrast to previous methods, our model takes into account both amino acid frequencies and a large number of experimentally-derived physical-chemical properties. Our approach is position- and property-specific: the degree of conservation of each property is tested independently at each position. We do not aggregate over many positions, and therefore our method can detect property constraints specific to single positions. We do not aggregate multiple properties into a similarity matrix or a small set of prototypical properties, and therefore our method can identify conservation of even those properties that only occasionally affect fitness.

To maximize the sensitivity of our conservation test, we use continuous instead of binary properties. Each property scale assigns a numerical value to each of the twenty amino acids. Frequencies of amino acids in a specific position in the alignment can be mapped directly into frequencies along the property scale. For example, the amino acid frequencies in the first position of the alignment shown in Figure 1(a) are shown plotted against two different properties in Figure 1. In Figure 1(b) these frequencies are mapped onto the (normalized) Roseman hydrophobicity scale. Figure 1(c) shows the same frequencies mapped onto the Levitt scale which measures radius of gyration of side chains. It is not difficult to see from these graphs that hydrophobicity is more conserved in this position than size.

### 2.2 Two Simple Measures of Property Conservation

A simple way to identify conserved properties is to look for properties with low variance in particular positions in a protein family. More specifically, let  $N$  be the number of sequences in the alignment, and  $v_{k,i}$  be the numerical value

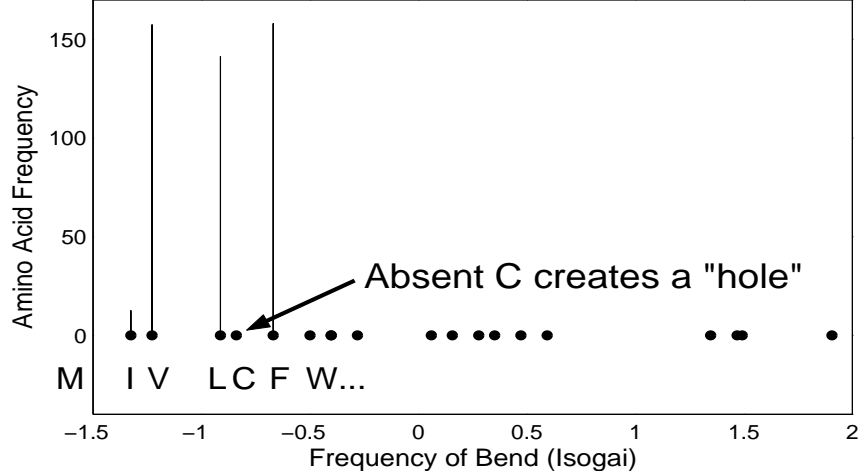


Figure 2: Distribution of residues in position 85 of HIV-1 Nef protein, plotted against a scale that quantifies the normalized frequency of a bend.

assigned to amino acid  $i$  by the  $k^{th}$  property scale. If  $n_{i,j}$  is the number of times amino acid  $i$  occurs in column  $j$  in the alignment, then the mean and variance of property  $k$  in the  $j^{th}$  position are:

$$\mu_{k,j} = \frac{1}{N} \sum_{i=1}^{20} n_{j,i} * v_{k,i}; \quad \sigma_{k,j}^2 = \frac{1}{N-1} \sum_{i=1}^{20} n_{j,i} * (v_{k,i} - \mu_{k,j})^2$$

Since property scales are measured in a variety of different units, the expected variance for a random sample of amino acids can differ substantially for different properties. An observed variance of 0.5 along a scale that ranges from 0 to 100 is not equivalent to the same variance measured on a scale that ranges from -1 to 1. To ensure that variances of different scales are comparable, each scale is normalized to have unit variance.

Even after this normalization, however, low variance in itself does not provide sufficient evidence that a property exerts pressure on a particular position. If the entropy is low (i.e. only a few different amino acids are present in this position), then *many* properties will have low variance. In the extreme case of only a single amino acid present, *all* properties will have zero variance. In addition, a variance-based measure is inadequate because it fails to penalize for residues that have not been observed. Figure 2 illustrates this point for an example protein. Shown is the distribution of amino acids at position 85 of the HIV-1 Nef protein, plotted on a scale that quantifies the normalized frequency of a bend. Four amino acids occur in this position (I, V, L, and F). All four of these amino acids have a similar frequency of occurring in a bend region of a protein, with the result that this property has low variance in this position. Nonetheless, the striking absence of cysteine raises a red flag. This “hole” causes us to doubt whether this property is actually being selected for. Thus, we make an additional assumption: when selection is based on a single property, fitness is unimodal (i.e. holes constitute negative evidence).

This new assumption suggests another very simple strategy: look for properties where all the residues at a particular position are adjacent in property space. Adjacency requires that if two residue types both occur in a particular position of the MSA, then any residue which has a property value lying between their property values must also occur in that position. Thus, properties for which all the amino acids in a specific position are adjacent will be considered significantly conserved. Clearly for positions with only a few amino acids this method will have a high false discovery rate. A position with only a single residue will appear equally conserved with respect to almost all properties. Even two-residue positions will appear to be equally conserved with respect to many properties. The false discovery rate for a position with only  $i$  amino acids can be quantified precisely; it is just the probability that these  $i$  amino acids are adjacent according to a random property:  $(21 - i) / \binom{20}{i}$ . Thus the false discovery rate is 10% for a position with only two amino acids present, 1.6% for three amino acids, 0.35% for four, and 0.1% for five.

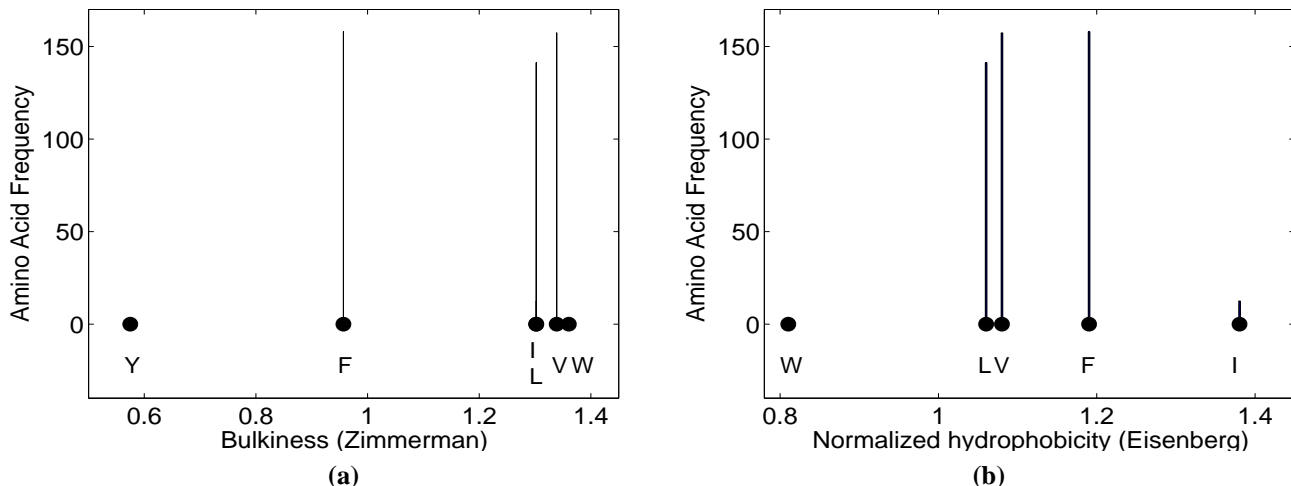


Figure 3: Distribution of residues in position 85 of HIV-1 Nef protein, plotted against two different scales: bulkiness (a) and hydrophobicity (b).

## 2.3 Gaussian Fitting

Unlike variance, the adjacency measure does not explicitly evaluate the degree of dispersion of the property values of residues at a specific position. Furthermore, it only considers whether an amino acid is present or absent. This leads to a strict definition of a hole and a subsequent failure to reject property distributions that are clearly not unimodal. For example in Figure 2, if cysteine was not completely absent in this position, but occurred in only one or two sequences, then the adjacency method would view this amino acid as “present” and fail to detect the very clear presence of a hole.

Adjacency measures also ignore the distances between the property values of different amino acids. For instance, consider the two property plots of the HIV-1 Nef protein at position 85 (Figure 3). The two graphs show the same data plotted against two different property scales, bulkiness and normalized hydrophobicity, shown in Figure 3(a) and 3(b), respectively. In both cases all the observed residues are adjacent (in Figure 3(a), I and L have *identical* property values). Nevertheless, in Figure 3(a), V is closer to W than it is to any of the other amino acids occurring at this position (F, I, and L). Given the relative distances between these amino acids, the total absence of W in this alignment position makes it doubtful that bulkiness is truly conserved, but does not effect our confidence in the conservation of hydrophobicity.

We would like our model to take these relative distances into account, and so we propose using a Gaussian distribution to model property conservation. A Gaussian is a good choice because it is unimodal, and will thus penalize heavily for missing amino acids with property values near the mean value. In addition, distributions such as that in Figure 3(a) will score badly because the Gaussian density decreases steeply only when the variance is small.

Given the amino acid frequencies for a position in the MSA, for each property we could select the parameters of the Gaussian that give maximum likelihood to the data. However, maximum likelihood estimation is not robust: a single amino acid with a property value far from the mean will result in fitting a Gaussian with a large variance. In addition, since the space of possible observations is not continuous, but discrete (with at most 20 values), the sample mean and covariance are not actually the maximum likelihood estimates. Indeed, when the space of possible observations is discrete and fixed, the maximum likelihood estimates for the parameters cannot be found analytically. Consequently, we smooth the likelihood function by interpolating it with a uniform background distribution over the 20 amino acids, then use a simple heuristic search procedure to efficiently select the parameters of the Gaussian to fit the observed data.

To measure property conservation we then apply a goodness-of-fit test to the learned model. We can calculate the number of times we expect to observe amino acid  $i$  in column  $j$ :

$$E_j[i] = (1 - \lambda) \frac{f_i}{\sum_{j=1}^{20} f_j} + \lambda \left( \frac{1}{20} \right) \quad \text{where} \quad f_i = \exp \left( \frac{(v_{k,i} - \mu)^2}{-2\sigma^2} \right)$$

where  $\mu$  and  $\sigma$  are the parameters of the Gaussian,  $\lambda$  is the interpolation weight of the background distribution<sup>1</sup>, and  $v_{k,i}$  is the numerical value assigned to amino acid  $i$  by the  $k^{th}$  property scale. These expected counts are compared to the observed counts with a  $\chi^2$  goodness-of-fit test. We don't expect an absolute fit of the data to a Gaussian, but we can use the relative fit to rank different hypotheses.

## 2.4 Significance Testing

Clearly, it is not difficult to fit a Gaussian to the amino acid frequencies if only a few of them are non-zero. Consider, for example, a position with almost perfect conservation. The best Gaussian for each property will have a very small variance and a mean centered at the dominant amino acid. As a result, almost every property will fit the data equally well. In these cases, we will have a high false discovery rate. We therefore designed a Monte-Carlo type significance test that determines which goodness-of-fit values are statistically significant. The procedure to determine the significance of a property in a particular position is as follows:

A large set of "pseudo properties" is generated by randomly permuting the numerical vectors corresponding to the physical-chemical properties. A Gaussian is then fit to each of the randomized properties and its  $\chi^2$  goodness-of-fit statistic determined. Cumulatively, these  $\chi^2$  values thus provide an estimate of the empirical distribution of  $\chi^2$  values for each position under the null (=randomness) assumption. This distribution is then used to convert the  $\chi^2$  statistic of the original physical-chemical property into a significance score, which plays the role of an empirically estimated p-value. This score approximates the probability of obtaining a  $\chi^2$  value smaller than the observed value by chance—small scores represent a highly significant property. Unlike the original  $\chi^2$  statistics, these significance scores are comparable across positions with significantly different entropies.

Since we are conducting a large number of tests at each position in the MSA, any p-value threshold should be corrected to account for the number of comparisons being performed. For instance, if we want the experiment-wide chance of a type I error to be no more than 5%, then a Bonferroni correction for multiple testing (of 240 properties) yields a significance threshold of  $\alpha = 0.05/240 = 0.0002$ . The Bonferroni correction is an overly conservative adjustment, especially since there are significant correlations between many of the properties we are testing.

Indeed, it is often the case that for a single site, the p-values of multiple property scales are significant. In this case, we report all of them as likely alternative explanations of the observed positional conservation. In most cases, these scales will be highly correlated, as they are often just different experimental measurements of the same underlying property. When this is not the case and the scales truly measure different properties, our results can give guidance in designing mutation experiments to fully determine the physical or chemical constraints imposed on the amino acid at this position in the protein.

## 2.5 Adjusting for Sequence Bias

Sequences in a multiple sequence alignment are not independent. In most cases they have diverged from a common ancestor. In other cases the sampling of sequences is biased toward certain sub-families. As a result, a typical alignment contains clusters of similar sequences. If these related clusters are large, they can significantly bias our estimates of amino acid frequency, effectively obscuring the variability exhibited by the less common sub-families.

Many different methods have been proposed for reducing this bias, the majority of which assign a numerical weight to each sequence with the goal of down-weighting closely related sequences. Although there are many algorithms for determining sequence weights (see [18] for a review), most measures of positional conservation provide no logical way to incorporate these weights. Since our methods are based solely on the frequency of each amino acid in each position, it is quite simple to incorporate sequence weights. We merely weight each position by its corresponding weight before counting frequencies. If  $W_i(a_j)$  is the cumulative weight of all sequences that have amino acid  $a_j$  in position  $i$ , the weighted frequency of the  $j^{th}$  amino acid is:

$$f_i(a_j) = \frac{W_i(a_j)}{\sum_{j'=1}^{20} W_i(a_{j'})}$$

The tree-weighting algorithm of Gerstein/Sonnhammer/Chothia (as implemented in [19]) was used in this analysis. This algorithm first constructs a guide tree using pairwise distances between sequences, and then assigns weights

---

<sup>1</sup>The weight of the uniform background component was set to 0.05 for these experiments.

according to the branch lengths and local density of the tree. The method essentially works by downweighting those sequences that appear to have diverged from each other recently.

Sequence weighting methods are generally simple, fast to compute, and effectively adjust for sample bias. They are not, however, able to fully account for the influence of phylogenetic relationships among the sequences [20]. Fully correcting for phylogenetic correlations requires a more sophisticated approach, and is discussed further in Section 5 on future work.

## 3 Experiments

### 3.1 Data Sources

A set of 240 property scales was obtained from the online databases PDbase [21] and ProtScale [22]. This set includes both experimentally-derived scales, such as average accessible surface area, and scales derived computationally from data, such as relative frequency in an alpha-helix. Each scale assigns a numerical value to each of the twenty amino acids.

The dataset we used to develop and test our method is comprised of proteins from HIV, for the following reasons. HIV, like other RNA viruses, has a high rate of mutation. This leads to a large and variable collection of protein sequences which provides a wealth of information about the selective constraints acting at particular positions in the HIV proteins. Multiple sequence alignments for two HIV proteins were downloaded from the HIV Molecular Immunology Sequence Database [23]. These alignments include many isolates of each HIV-1 subtype and inter-subtype recombinants, but include only one sequence from any one individual. The alignments were built by scientists at the Los Alamos National Laboratory. They used HMM models and hand-editing to construct very accurate alignments. Furthermore, since the functions of the proteins in the alignment are identical, and the sequences are derived from the same organism, the amino acid substitutions in the MSAs are not affected by possible contributions from organism-specific amino acid usage preferences, nor substitutions that slightly modulate function (such as would be the case in protein families where individual members bind different ligands, for example). The envelope protein (Env) alignment was constructed from 388 protein sequences. The alignment for the negative factor protein (Nef) is slightly larger, encompassing 484 sequences. The Nef sequences are about 200 residues long and the Env proteins are over 800 residues long.

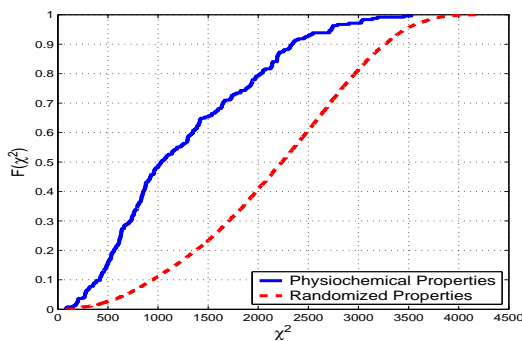
A large number of gaps often indicates that a position is not important to the structure or function of the protein. In addition, there is no way to assign property values to gaps. Hence, we analyze only those positions where a majority of sequences do not contain gaps. For efficiency reasons we only test positions with entropy  $> 0.5$ , since our method will not be able to identify the cause of conservation in such highly conserved positions. After filtering, there were 145 columns in the Nef alignment and 625 in the Env alignment.

In the following sections, unless otherwise stated, all position numbers refer to the HIV reference proteins with Swiss-prot IDs NEF\_HV112 and ENV\_HV1H2.

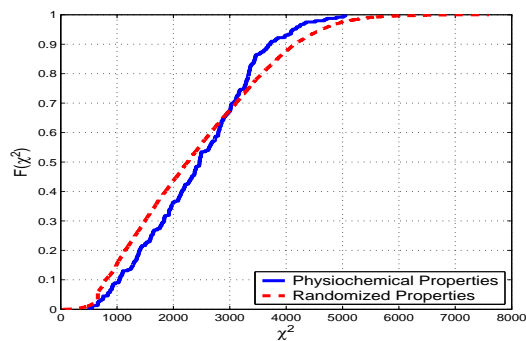
### 3.2 Results

Applying the Gaussian-fitting method to the two HIV protein alignments described in the previous section, we observe that for certain positions the method is clearly identifying properties that fit the data much better than would be expected by chance. The use of randomized properties for significance testing of property conservation is illustrated by Figure 4, which shows the empirical cumulative distribution of  $\chi^2$  goodness-of-fit statistics for three different positions in the HIV-1 Nef Protein.

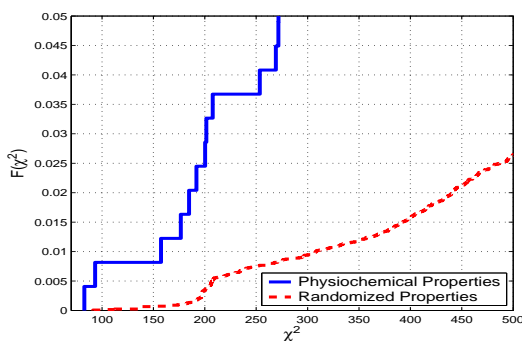
Figure 4(a) shows a position for which a small number of properties fit the amino acid distribution much more closely than would be expected by chance. The detail of this graph in Figure 4(c) illustrates how the distribution of random  $\chi^2$  values can be used to establish a significance cutoff. If we want the experiment-wide chance of a type I error to be no more than 5%, then a Bonferroni correction yields a significance threshold of  $\alpha = 0.05/240 = 0.0002$ , which for this cumulative distribution function yields a critical value of  $\chi^2 < 102$ . Any property with a goodness-of-fit that is better than this cutoff is highly unlikely to have fit the data well purely by chance. In this case there are two (almost identical) property scales that pass this threshold: 'Percentage of buried residues/Janin' and 'Free energy of transfer from inside to outside of a globular protein/Janin'.



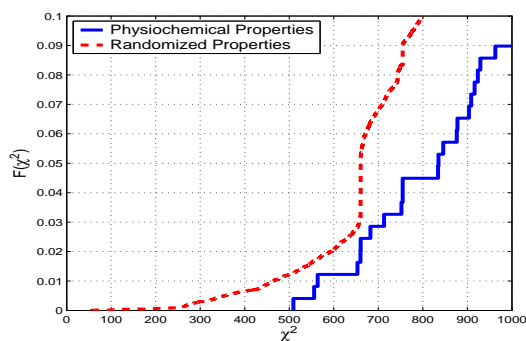
(a) Nef position 85



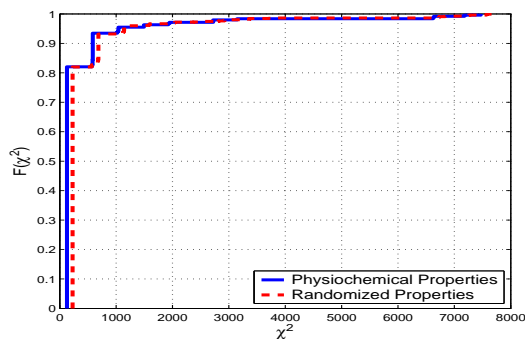
(b) Nef position 54



(c) Nef position 85 (Detail)



(d) Nef position 54 (Detail)



(e) Nef position 52

Figure 4: Empirical cumulative distribution function  $F$  of the  $\chi^2$  goodness-of-fit statistic for physical-chemical properties compared to randomized properties. The x-axis shows the value of the  $\chi^2$  statistic, where a value of 0 represents a perfect fit to a Gaussian. The y-axis shows the percentage of properties whose goodness-of-fit statistic is less than or equal to that value. The solid line shows the distribution of  $\chi^2$  values for the set of 240 physical-chemical properties. The dashed line shows the distribution of  $\chi^2$  values for a set of 12,250 randomly shuffled properties. The subfigures illustrate the significant differences between the distribution of  $\chi^2$  values at different positions in the Nef alignment. Subfigure (a) and its detail (c) illustrate a more variable site, with a small number of significantly conserved properties. Subfigure (b) and its detail (d) also illustrate a variable site, but in this position there are no physiochemical properties with a  $\chi^2$  value significantly smaller than would be expected by chance. Subfigure (e) shows a highly conserved site, so the average  $\chi^2$  statistic is close to zero and there are no properties with a  $\chi^2$  value significantly smaller than expected under the null model.



Protein	Alignment Position	Reference Position	Conserved property (p-value < 0.0002)
Nef	17	16	Percentage of exposed residues/Janin
Nef	22	21	Hydropathy index/Kyte-Doolittle
Nef	62	45	Relative mutability/Jones
Nef	63	46	Retention coefficient in HPLC, pH7.4/Meek
Nef	106	85	Percentage of buried residues/Janin
Nef	110	89	Average flexibility indices/Bhaskaran-Ponnuswamy
Nef	119	98	Net charge/Klein
Nef	162	135	Hydrophobicity from HPLC peptide retention times/Wilson
Nef	177	149	Hydrophobicity at pH 3.4 determined by HPLC/Cowan
Nef	214	184	Partition energy/Guy
Nef	223	192	Optimized transfer energy parameter/Oobatake
Env	48	32	Normalized flexibility/B-values
Env	81	63	Flexibility parameter for no rigid neighbors/Karplus-Schulz
Env	203	165	Hydrophobicity/Prabhakaran
Env	212	174	Absolute entropy/Hutchens
Env	356	291	Side chain torsion angle phi/Levitt
Env	427	352	Recognition factors
Env	442	363	Long range non-bonded energy per atom/Oobatake-Ooi
Env	595	494	Polarity/Grantham
Env	739	620	Hydrophobic parameter pi/Fauchere-Pliska

Table 1: Properties found to be most significantly conserved at specific positions in the Nef and Env proteins, where the significance level is  $\alpha = 0.05$ . The threshold after a Bonferroni correction for multiple testing thus requires p-value < 0.0002.

Figure 4(b) and its detail (Figure 4(d)), on the other hand, show a position where none of the 240 properties achieve a significant fit. There is little difference between the distribution of the physiochemical properties and that of random properties. Either this position has very little conservation, or the selective pressures are acting on more than one property, with the result that our method fails to detect a single property which is conserved.

In Figure 4(e) an almost perfectly conserved position is shown. Since almost all properties (including random properties) can be fit well by a Gaussian, none of the 240 properties achieve a fit that is significantly better than random. Thus our method is not applicable for highly conserved positions where sequence conservation is a better descriptor for evolutionary pressure than property conservation.

Table 1 lists some of the most highly conserved properties at specific positions in the Nef and Env proteins. For each of the reported positions, either the property listed in the table is the only property that passes the given significance threshold, or it is highly correlated with all other scales found to be significantly conserved at that position.

## 4 Discussion

To obtain a qualitative evaluation and interpretation of the results obtained with our method, we analyzed in detail the multiple sequence alignments of the HIV-1 Nef and Env proteins in relation to the existing knowledge about their structure and function. HIV-1 is a Lentivirus, which as retroviruses encode the prototypic Gag, Pol and Env genes, plus a number of accessory genes including Nef.

### 4.1 Property Conservation in Nef Proteins

Nef is a 27-kDa myristoylated protein, observed in the cytoplasm of cells infected by lentiviruses [25]. Nef is involved in viral signaling processes and has been shown to be associated with cellular membranes and the cytoskeleton, depending on myristoylation of the N-terminal glycine [26, 27, 28]. Cleavage between W57 and L58 separates the protein into two domains [29]: the C-terminal well folded core domain (amino acids 58 to 206) with structural similar-

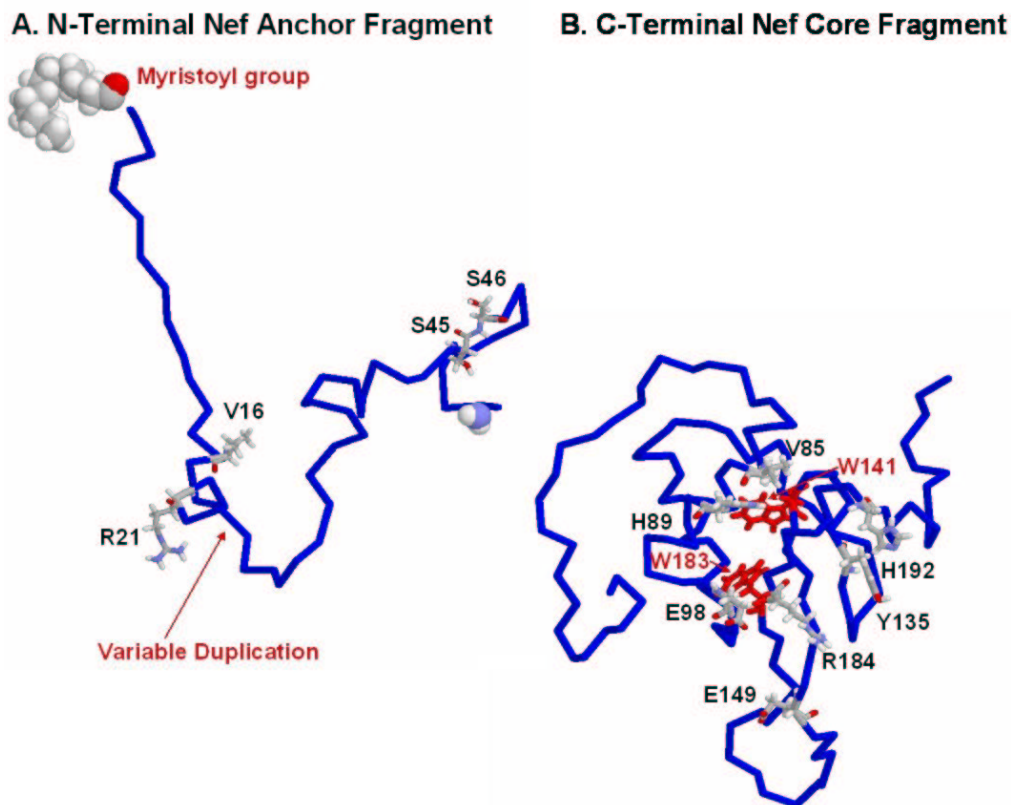


Figure 5: Structures of two fragments of the HIV-1 Nef protein obtained from NMR spectroscopy. Residues with conserved properties (Table 1) are highlighted in stick representation using CPK color coding. The amino acid identities and positions are shown in black. Two other amino acids are highlighted in red, W141 and W183. The locations of the myristoyl group and the variable duplication region are shown. A. N-terminal fragment myristate-G2 to W57, PDB code 1QA5 [24]. B. C-terminal fragment 40-206 with a deletion from 159-173 and C206A point mutation, PDB code 2NES [30]. The amino acids from 40-56 are unstructured and are not shown. The structure was drawn using the Non-covalent bond finder (MDL).




Figure 6: Crystal structure model of gp120 (blue) in complex with CD4 and neutralizing antibody (green and yellow). Three amino acids with conserved properties are modeled in this structure, shown as stick representations. The structures were drawn using the Non-covalent bond finder (MDL).

ities to DNA binding proteins containing a helix-turn-helix motif [30, 31] and the flexible N-terminal sequence which is proposed to function as a membrane anchor [29, 32]. The N-terminal fragment is largely unstructured in solution but in the presence of myristoylation, there are two relatively well defined helices separated by a turn [24]. The structures of fragments corresponding to the two domains obtained by NMR spectroscopy are shown in Figure 5A and B, for the N-terminal and C-terminal fragments, respectively. The structures are shown as predicted to be the correct relative orientation between the two domains in the full-length protein [33].

Table 1 lists properties found to be most significantly conserved at specific positions in the Nef proteins, where the significance level is  $\alpha = 0.05$ , and thus the corrected p-value threshold is 0.0002. These positions are highlighted in the structures shown in Figure 5, labeled in black with the amino acid identity and position of the HIV-1 Nef reference protein. Previous structure-function studies have correlated surface accessibility with secondary structure elements and sequence conservation [34]. It was found that conserved motifs involved in Nef-mediated CD4 and MHC I down-regulation are located in flexible regions of Nef, suggesting that the formation of the transient trafficking complexes depends on the recognition of primary sequences. In contrast, the interaction sites for signaling molecules that contain SH3 domains or the p21-activated kinases are associated with the well folded core domain, suggesting the recognition of highly structured protein surfaces [34]. Critical for the folding of the C-terminal domain are residues W141 and W183 which form the scaffold of the protein and assemble the two core helices to the  $\beta$ -pleated sheet (Figure 5, [33]). As one can see, several of the amino acids with conserved properties are located in close proximity to these two amino acids, highlighted in red in Figure 5B.

An interesting observation that came out of our analysis is that a single property, retention coefficient in  $\text{NaH}_2\text{PO}_4$ , appears to be conserved in seven amino acids<sup>2</sup> that are all located in a not well conserved, variable duplication region in the loop following the first helix of the N-terminal domain, shown in red in Figure 5A. A likely reason for conservation of retention coefficient in  $\text{NaH}_2\text{PO}_4$ , a measure of hydrophobicity, at these sites is suggested by the main function of the N-terminal Nef domain in membrane anchoring. The insertion of amino acids with the same property therefore would be predicted to enhance the interactions of Nef with the membrane. This is a hypothesis that can be tested directly through wet-lab experiments. Thus, this example highlights the type of hypothesis that can be generated using property conservation revealing evolutionary pressure that is not captured by sequence conservation.

## 4.2 Property Conservation in Env Proteins

The HIV-1 envelope glycoprotein (Env) mediates the fusion of viral and cellular membranes during HIV-1 infection of cells. Env is synthesized as the precursor protein gp160, which is proteolytically cleaved into two subunits, the surface subunit gp120 and the transmembrane subunit gp41. Both fragments remain non-covalently associated with each other after cleavage in a trimeric structure. HIV-1 infection is initiated by binding of gp120 primarily to CD4 with co-receptors of the chemokine receptor family on the host surface. Subsequent conformational changes in Env result in exposure of a hydrophobic N-terminal fusion peptide within the gp41 domain, thus initiating membrane fusion.

Only one of the amino acid with highly conserved properties is located in gp41, in the disulfide-bonded loop region. This region is thought to act as a hinge in the formation of the trimer-of-hairpins structure of gp41 that brings the cellular and viral membranes in proximity and to play a key role in the interaction between gp120 and gp41 [35]. Thus, the conservation of hydrophobicity at amino acid at position 739 in the alignment may play an important role in this interaction. All other amino acids are located in the gp120 fragment. The structure of the gp120 fragment has been solved by X-ray crystallography in complex with CD4 and a neutralizing antibody [36]. Only three of the amino acids are located in the fragment that has been crystallized, alignment positions 356, 427 and 442 corresponding to S291, Q363 and Q352 in the HIV-1 reference protein. These three amino acids are highlighted in Figure 6. All are located on different parts of the surface of the molecule. Different properties are conserved at all three sites, however, all of them are conceivably connected to possible recognition of protein partners. All other amino acids with conserved properties are located in regions that are flexible and thus not amenable to X-ray crystallographic analysis. This includes the N-terminus (positions 48 and 81), for which both properties are related to its flexibility supporting the hypothesis that these amino acids contribute to the observed flexibility. The other two amino acids, at positions 203 and 212, are located in the V1/V2 loop, a sequence variable region that was deleted in the gp120 crystal structure to allow crystallization. This loop has been shown to rescue changes in the important V3 loop on the opposite side of the molecule, and are a potential interaction site between individual molecules of the three units of

<sup>2</sup>These seven amino acids are at positions 29, 31, 32, 33, 34, 37, and 38 in the alignment. The first amino acid corresponds to residue 28 in the reference protein.

the trimer. The lack of sequence conservation in this region, but high conservation of absolute entropy (position 212) and hydrophobicity (position 203) suggest that these residues may be involved in the inter-monomer interactions, an experimentally testable hypothesis.

## 5 Conclusions and Future Work

We have described a framework for systematically identifying the property-based selection pressures affecting each position in a MSA. We first calculate a distance from the observed distribution to an idealized distribution (currently a Gaussian) assuming a particular property is conserved at this position. We then use a Monte-Carlo procedure to estimate the probability of achieving such a good fit by chance. Our method requires only a MSA as input, and determines the most highly conserved properties in each position, as well as their statistical significance. In our biological evaluation we discuss the significance of our results with respect to current understanding of the structure and function of two proteins from HIV-1, Nef and Env. We make predictions regarding previously unknown sites of property conservation, providing guidance for future site-directed mutagenesis studies by experimental biologists.

Our current test evaluates whether a particular property scale explains the observed data better than we would expect by chance. However, it is often the case that more than one property scale is determined to provide a significant fit to the data. This usually occurs because there are many highly correlated scales which are just different experimental measurements of the same underlying property (such as hydrophobicity or size). In this case, we do not expect to be able to discriminate between these highly correlated scales. In the less common case where two properties which are not highly correlated are both scored as significant, we would like to test whether the different properties explain the data equally well or whether one property is significantly better than the other.

Currently, we try to limit confounding effects due to the lack of independence between sequences by using sequence weights. However, it has been shown by Bruno [20] that since sequence weights assume that the effects of phylogenetic correlation are identical at every position in the alignment, they cannot fully correct for phylogenetic dependencies. A more thorough, but also more computationally intensive solution to this problem would either use Bruno's Rind program to estimate site-specific residue frequencies or explicitly incorporate the phylogenetic tree into a generative model of protein evolution that includes selection of physical and chemical properties.

We chose to use a Gaussian to model property conservation effects because it is unimodal and has other desirable properties (discussed in Section 2.3). However, we do not necessarily believe that selective pressures are truly Gaussian in nature. Another option would be to not assume any particular unimodal distribution, but just test the hypothesis that the distribution of property values in a position is unimodal. A dip test is one way to determine whether a distribution is unimodal [37].

The current method is only designed to find evidence for positional conservation of a *single* property, but we know that in many cases selection exerts pressure on multiple properties simultaneously. In particular, there are many cases where we find a property that is more conserved than would be expected by chance, but this property still has unavoidable "holes" which could be explained by the intersection of two property pressures. We would like to extend our approach to use multi-dimensional fitting of properties in these cases where a single property is insufficient to explain the observed data. Although testing for combinations of properties could be carried out by a simple multivariate extension of the Gaussian-fitting test, the larger search space makes multiple testing a more problematic issue. In order to avoid false positives we will have to reduce the number of tests conducted, possibly by considering only combinations of properties which were conserved (to some extent) individually.

## Acknowledgements

This research was supported by National Science Foundation Information Technology Research grant NSF0225656 and by the American Society for Engineering Education (ASEE) through a National Defense Science & Engineering Graduate (NDSEG) fellowship. We are grateful to Bill Bruno for his detailed and helpful comments and to Larry Wasserman for numerous helpful discussions.

## References

- [1] L. Oliveira, P.B. Paiva, A.C. Paiva, and G. Vriend, "Sequence analysis reveals how G protein-coupled receptors transduce the signal to the G protein," *Proteins*, vol. 52, no. 4, pp. 553–60, Sep 2003.

- [2] L. Oliveira, P.B. Paiva, A.C. Paiva, and G. Vriend, "Identification of functionally conserved residues with the use of entropy-variability plots," *Proteins*, vol. 52, no. 4, pp. 544–52, Sep 2003.
- [3] I.V. Grigoriev and S.H. Kim, "Detection of protein fold similarity based on correlation of amino acid properties," *Proc Natl Acad Sci*, vol. 96, no. 25, pp. 14318–23, Dec 1999.
- [4] V.S. Mathura, C.H. Schein, and W. Braun, "Identifying property based sequence motifs in protein families and superfamilies: application to dnase-1 related endonucleases," *Bioinformatics*, vol. 19, no. 11, pp. 1381–90, Jul 2003.
- [5] C. Ouzounis, C. Perez-Irratxeta, C. Sander, and A. Valencia, "Are binding residues conserved?," in *Pac Symp Biocomput*, 1998, pp. 401–12.
- [6] H.O. Villar and L.M. Kauvar, "Amino acid preferences at protein binding sites," *FEBS Lett*, vol. 349, no. 1, pp. 125–030, Jul 1994.
- [7] O. Lichtarge, H.R. Bourne, and F.E. Cohen, "An evolutionary trace method defines binding surfaces common to protein families," *J Mol Biol*, vol. 257, no. 2, pp. 342–58, Mar 1996.
- [8] J. Pei and N.V. Grishin, "AL2CO: calculation of positional conservation in a protein sequence alignment," *Bioinformatics*, vol. 17, no. 8, pp. 700–12, Aug 2001.
- [9] W.S. Valdar, "Scoring residue conservation," *Proteins*, vol. 48, no. 2, pp. 227–41, Aug 2002.
- [10] P.S. Shenkin, B. Erman, and L.D. Mastrandrea, "Information-theoretical entropy as a measure of sequence variability," *Proteins*, vol. 11, no. 4, pp. 297–313, 1991.
- [11] C. Sander and R. Schneider, "Database of homology-derived protein structures and the structural meaning of sequence alignment," *Proteins*, vol. 9, no. 1, pp. 56–58, 1991.
- [12] M. Gerstein and R.B. Altman, "Average core structures and variability measures for protein families: application to the immunoglobulins," *J Mol Biol*, vol. 251, no. 1, pp. 161–75, Aug 1995.
- [13] W.R. Taylor, "The classification of amino acid conservation," *J Theor Biol*, vol. 119, no. 2, pp. 205–18, Mar 1986.
- [14] M.J. Zvelebil, G.J. Barton, W.R. Taylor, and M.J. Sternberg, "Prediction of protein secondary structure and active sites using the alignment of homologous sequences," *J Mol Biol*, vol. 195, no. 4, pp. 957–61, Jun 1987.
- [15] K. Sjolander, K. Karplus, M. Brown, R. Hughey, A. Krogh, I.S. Mian, and D. Haussler, "Dirichlet mixtures: A method for improving detection of weak but significant protein sequence homology," *CABIOS*, vol. 12, no. 4, pp. 327–45, Aug 1996.
- [16] J.M. Koshi, D.P. Mindell, and R.A. Goldstein, "Beyond mutation matrices: Physical-chemistry based evolutionary models," *Genome Inform Ser Workshop Genome Inform*, vol. 7, no. 12, pp. 80–89, 1997.
- [17] L. Yampolsky and A. Stoltzfus, "Amino acid exchangeability from experimental data," Unpublished manuscript.
- [18] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis*, Cambridge University Press, Cambridge, 1998.
- [19] J.D. Thompson, D.G. Higgins, and T.J. Gibson, "CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice," *Nucleic Acids Res*, vol. 22, pp. 4673–4680, 1994.
- [20] W. Bruno, "Modeling residue usage in aligned protein sequences via maximum likelihood," *Mol Biol Evol*, vol. 13, no. 10, pp. 1368–74, 1996.
- [21] PDBase, ," [http://www.scsb.utmb.edu/comp\\_biol.html/venkat/prop.html](http://www.scsb.utmb.edu/comp_biol.html/venkat/prop.html).
- [22] Expasy's ProtScale, ," <http://us.expasy.org/cgi-bin/protscale.pl>.
- [23] HIV Molecular Immunology Database, ," <http://www.hiv.lanl.gov/content/immunology/>.
- [24] M. Geyer, C. E. Munte, J. Schorr, R. Kellner, and H. R. Kalbitzer, "Structure of the anchor-domain of myristoylated and non-myristoylated hiv-1 nef protein," *J Mol Biol*, vol. 289, no. 1, pp. 123–38, 1999.
- [25] M. Harris, "From negative factor to a critical role in virus pathogenesis: the changing fortunes of nef," *J Gen Virol*, vol. 77 ( Pt 10), no. 77 ( Pt 10), pp. 2379–92, 1996.
- [26] O. T. Fackler, N. Kienzle, E. Kremmer, A. Boese, B. Schramm, T. Klimkait, C. Kucherer, and N. Mueller-Lantzsch, "Association of human immunodeficiency virus nef protein with actin is myristoylation dependent and influences its subcellular localization," *Eur J Biochem*, vol. 247, no. 3, pp. 843–51, 1997.
- [27] G. Franchini, M. Robert-Guroff, J. Ghayeb, N. T. Chang, and F. Wong-Staal, "Cytoplasmic localization of the htlv-iii 3'; orf protein in cultured t cells," *Virology*, vol. 155, no. 2, pp. 593–9, 1986.
- [28] T. M. Niederman, W. R. Hastings, and L. Ratner, "Myristoylation-enhanced binding of the hiv-1 nef protein to t cell skeletal matrix," *Virology*, vol. 197, no. 1, pp. 420–5, 1993.

- [29] J. Freund, R. Kellner, T. Houthaeve, and H. R. Kalbitzer, "Stability and proteolytic domains of nef protein from human immunodeficiency virus (hiv) type 1," *Eur J Biochem*, vol. 221, no. 2, pp. 811–9, 1994.
- [30] S. Grzesiek, A. Bax, G. M. Clore, A. M. Gronenborn, J. S. Hu, J. Kaufman, I. Palmer, S. J. Stahl, and P. T. Wingfield, "The solution structure of hiv-1 nef reveals an unexpected fold and permits delineation of the binding surface for the sh3 domain of hck tyrosine protein kinase," *Nat Struct Biol*, vol. 3, no. 4, pp. 340–5, 1996.
- [31] C. H. Lee, K. Saksela, U. A. Mirza, B. T. Chait, and J. Kuriyan, "Crystal structure of the conserved core of hiv-1 nef complexed with a src family sh3 domain," *Cell*, vol. 85, no. 6, pp. 931–42, 1996.
- [32] J. Freund, R. Kellner, J. Konvalinka, V. Wolber, H. G. Krausslich, and H. R. Kalbitzer, "A possible regulation of negative factor (nef) activity of human immunodeficiency virus type 1 by the viral protease," *Eur J Biochem*, vol. 223, no. 2, pp. 589–93, 1994.
- [33] M. Geyer and B. M. Peterlin, "Domain assembly, surface accessibility and sequence conservation in full length hiv-1 nef," *FEBS Lett*, vol. 496, no. 2-3, pp. 91–5, 2001.
- [34] M. Geyer, O. T. Fackler, and B. M. Peterlin, "Structure–function relationships in hiv-1 nef," *EMBO Rep*, vol. 2, no. 7, pp. 580–5, 2001.
- [35] S. Wang, J. York, W. Shu, M. O. Stoller, J. H. Nunberg, and M. Lu, "Interhelical interactions in the gp41 core: implications for activation of hiv-1 membrane fusion," *Biochemistry*, vol. 41, no. 23, pp. 7283–92, 2002.
- [36] P. D. Kwong, R. Wyatt, J. Robinson, R. W. Sweet, J. Sodroski, and W. A. Hendrickson, "Structure of an hiv gp120 envelope glycoprotein in complex with the cd4 receptor and a neutralizing human antibody," *Nature*, vol. 393, no. 6686, pp. 648–59, 1998.
- [37] J. A. Hartigan and P. M. Hartigan, "The dip test of unimodality," *Annals Of Statistics*, vol. 13, pp. 70–84, 1985.