

DO DYNAMIC ALLOSTERIC EFFECTS OCCUR IN IGG4 ANTIBODIES?

by

Lonnie Baker

A thesis submitted to the faculty of
The University of North Carolina at Charlotte
in partial fulfillment of the requirements
for the degree of Master of Science in
Applied Physics

Charlotte

2020

Approved by:

Dr. Donald Jacobs

Dr. Jun-tao Guo

Dr. Yuri Nesmelov

ABSTRACT

LONNIE BAKER. Do Dynamic Allosteric Effects Occur in IgG4 Antibodies?.
(Under the direction of DR. DONALD JACOBS)

Antibodies are large, multi-chain proteins which play a critical roll in the functioning of vertebrate immune systems. Immune effector cells which are responsible for the destruction of foreign antigens do not randomly target molecules but rely on antibodies to recognize and bind antigens with very high specificity before initiating the process of antigen destruction. This process, known as antibody-dependent cellular cytotoxicity (ADCC), involves coordination of binding sites on the surfaces of antibodies that are distant from each other implying an allosteric mechanism. From molecular dynamics data, an effective elastic network model (ENM) of an IgG4 monoclonal antibody is constructed from the inverse of a noise-decorrelated covariance matrix. A perturbation-response method is then applied which demonstrates long range correlations exist between antibody residues. The correlations between these distant sites are then shown to have potential site specific cooperative binding free energies which strongly suggest the presence of dynamic allosteric effects. Of particular significance, our results confirm the putative allosteric pathways directly connecting the functionally important CH2 domains to the complementarity determining regions (CDRs). Our results also show that this functional allosteric pathway is highly sensitive to mutations. This finding launches a potential paradigm shift in how antibodies use molecular cooperativity as part of their function.

TABLE OF CONTENTS

| | |
|---|----|
| LIST OF TABLES | v |
| LIST OF FIGURES | vi |
| LIST OF ABBREVIATIONS | 1 |
| CHAPTER 1: INTRODUCTION | 1 |
| 1.1. Allostery in Antibodies | 2 |
| 1.2. Dynamic Allostery Model | 5 |
| 1.3. Molecular Dynamics Simulations of Antibodies | 6 |
| 1.4. Principal Component and Normal Mode Analysis | 6 |
| 1.5. Effective Elastic Network Model | 9 |
| 1.6. Noise Decorrelation | 11 |
| 1.7. Perturbation-Response Model | 13 |
| 1.8. Cooperative Binding and Allosteric effects | 16 |
| 1.9. Eigenvalue Perturbations | 17 |
| CHAPTER 2: RESULTS | 19 |
| 2.1. Analysis of dynamic couplings | 19 |
| 2.2. Shrinkage of Estimated Covariance Matrices. | 24 |
| 2.3. New Noise Decorrelation Method. | 25 |
| 2.4. Modeling cooperative binding effects | 32 |
| CHAPTER 3: CONCLUSIONS | 46 |
| REFERENCES | 49 |

LIST OF TABLES

| | |
|--|----|
| TABLE 2.1: Correlations between SS and DD profiles generated with different perturbation types at CDR-L1a. Numbers in the upper triangle represent correlations between SS profiles (blue), while the numbers in the lower triangle represent DD profiles (red). | 23 |
| TABLE 2.2: Correlations between SS and DD profiles generated with different perturbation types at CDR-H3a. Numbers in the upper triangle represent correlations between SS profiles (blue), while the numbers in the lower triangle represent DD profiles (red). | 24 |
| TABLE 2.3: Correlations between SS (blue) and DD (red) global response profiles generated using the old and new methods of noise decorrelation (ND). | 30 |
| TABLE 2.4: Correlations between SD (red) and DS (blue) response profiles generated using the old and new methods of noise decorrelation (ND). | 30 |
| TABLE 2.5: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-L1a and all other residues for ball, star and triad at k=1 native simulations | 38 |
| TABLE 2.6: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-L1a and all other residues for all k values: k = 0.1, 0.5, 1, 5, 10, 20. These values indicate that the relative regions of cooperativity and anti-cooperativity remain fairly constant as the perturbation strength is increased. They also indicate that $\Delta\Delta G$ does not respond in a perfectly linear way to increasing perturbation strength. | 40 |
| TABLE 2.7: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-H3a and all other residues for all k values: k = 0.1, 0.5, 1, 5, 10, 20 | 43 |
| TABLE 2.8: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-H3a and all other residues for 3 simulations which were started from different molecular configurations of the same protein. | 45 |

LIST OF FIGURES

| | |
|--|----|
| FIGURE 1.1: An IG antibody monomer is composed of two heavy chains (grey) and two light chains (orange). Each heavy chain is composed of three constant domains (CH1, CH2, CH3) and one variable domain (VH). Each light chain is composed of a single constant domain (CL) and a single variable domain (VL). The variable regions of each chain have three complementarity determining regions, or CDRs (purple). | 4 |
| FIGURE 1.2: Hessian matrices at varying levels of noise decorrelation. (A) Zero percent noise decorrelation. (B) 5% noise decorrelation. (C) 10% noise decorrelation. (D) 20% noise decorrelation. | 12 |
| FIGURE 1.3: Triad, star and ball perturbation types. Spring constants 'k' are added to the model between sets of atoms defined by the perturbation type. The perturbation radius 'R' defines a cutoff distance between the target residue (red) and nearby residues which are to be included in the perturbation (blue). | 14 |
| FIGURE 1.4: All residue perturbation-response matrix. This matrix was generated using a positive/stabilizing force constant. Every blue point on the matrix represents a residue that was stabilized as a result of of stabilizing perturbation. Every red point represents a residue that was destabilized. The location of the perturbation is given by the y axis. The location of the response is given by the x axis. | 16 |
| FIGURE 2.1: An example of global perturbation and response profiles organized on the same plot. The three colors indicate perturbation scan conducted on each of three ENMs generated from different native structure simulations. The points above the axis represent the global response profile while points below the axis represent the global perturbation profile as defined above. | 20 |
| FIGURE 2.2: Stabilizing perturbation-Stabilizing response (SS) perturbation profiles for all three perturbation types. These images show the magnitudes of <i>stabilizing</i> responses to <i>stabilizing</i> perturbations. Blue regions indicate that perturbations placed here cause a higher than average global <i>decrease</i> in MSF. | 20 |

- FIGURE 2.3: Stabilizing perturbation-Destabilizing response (SD) perturbation profiles for all three perturbation patterns. These images show the magnitudes of *destabilizing* responses to *stabilizing* perturbations. Red regions indicate that perturbations placed here cause a higher than average global *increase* in MSF. 21
- FIGURE 2.4: Stabilizing perturbation-Stabilizing response (SS) response profiles for all three perturbation patterns. These images show the magnitudes of *stabilizing* responses to *stabilizing* perturbations. Blue regions indicate that these residues experience a higher than average *decrease* in mobility when other residues across the entire protein are perturbed. 21
- FIGURE 2.5: Stabilizing perturbation-Destabilizing response (SD) response profiles for all three perturbation patterns. These images show the magnitudes of *destabilizing* responses to *stabilizing* perturbations. Red regions indicate that these residues experience a higher than average *increase* in mobility when other residues across the entire protein are perturbed.. 21
- FIGURE 2.6: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the L1a CDR loop (purple). Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. 23
- FIGURE 2.7: Stabilizing-Destabilizing (SD) response profiles for targeted perturbations of the H3a CDR loop (purple). Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. 24
- FIGURE 2.8: The overlap matrix of all eigenvectors from Hessian matrices derived using both ND methods. A diagonal overlap matrix implies a very high degree of similarity between the two sets of vectors. 27
- FIGURE 2.9: Descending eigenvalue scree plots for covariance matrices de-noised using the old (left) and new (right) noise decorrelation methods. 27
- FIGURE 2.10: Ascending Hessian matrix eigenvalue plots for the old (left) and new (right) noise decorrelation methods. These plots show that the old noise decorrelation method introduces a sharp jump in the low frequency modes which is not present in the new method. 28

- FIGURE 2.11: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the L1a CDR loop (purple) using the old and new noise decorrelation methods. Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. From this figure, the new noise decorrelation method does seem to have minor effects on the dynamic coupling between the CDR-L1a perturbation target and other residues. 31
- FIGURE 2.12: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the H3a CDR loop (purple) using the old and new noise decorrelation methods. Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. The new noise decorrelation method seems to have minor effects on the dynamic coupling between the H3a-L1a perturbation target and other residues. 32
- FIGURE 2.13: The locations of three perturbation targets are shown above in purple. These include CDR loops L1a, H3a, and L1b. Residue 470 is colored red and indicates the residue that was most commonly the location of highest cooperativity (most negative $\Delta\Delta G$) across all perturbations. 34
- FIGURE 2.14: (A) shows a plot of $\Delta\Delta G$ values which associate a measure of cooperativity between an initial binding site at CDR-L1a (purple) and all other residues. (B) Orange to red colored sites show anti-cooperativity with CDR-L1a, while green to blue sites show positive cooperativity (more negative $\Delta\Delta G$ values). 35
- FIGURE 2.15: Values of $\Delta\Delta G$ per mode between CDR-L1a and residues 467-472. Across all combinations of parameters, these residues showed the highest degree of binding cooperativity (lowest $\Delta\Delta G$) when the initial perturbation targets were either CDR-L1a or CDR-H3a. 35
- FIGURE 2.16: Binding cooperativity $\Delta\Delta G$ between CDR-L1a and all other residues using all three perturbation types (ball, star, and triad) at a spring constant of $k = 10$. 37
- FIGURE 2.17: Binding cooperativity $\Delta\Delta G$ between CDR-L1a and all other residues at various values of spring constant k and a ball type perturbation. 39

FIGURE 2.18: Binding cooperativity $\Delta\Delta G$ between CDR-H3a and all other residues at various values of spring constant k and a ball type perturbation. 42

FIGURE 2.19: Binding cooperativity $\Delta\Delta G$ between CDR-H3a and all other residues using a spring constant $k = 10$ compared using three different simulations of a native IgG4 structure. 44

CHAPTER 1: INTRODUCTION

A typical protein interacts with multiple chemicals in the course of its functions within a cell. These chemical interactions take place at binding sites on the surface of the protein which can often be distant from each other. When a binding event at one binding site effects the binding affinity of another site which is distant from the first site, the phenomenon is known as allostery.

The existence of multiple binding sites on a protein comes as no surprise given the large size of most proteins. As different functions of proteins are measured experimentally, such as the ability of an enzyme to produce a product or a transporter protein to bind its cargo, it can be easily theorized that each function is facilitated by a set of independently operating binding sites. But by the 1960s, evidence of certain protein functions had been collected that could not be adequately explained by a completely functionally independent set of binding sites.

One of the earliest clear examples of this is a study on a bacterial enzyme L-threonine deaminase[1]. In this study, the function of the enzyme L-threonine deaminase is inhibited by an end product which is produced several chemical steps later. A simple model of competitive binding at the enzymatic cleft between the enzyme target and the inhibitor was not supported by experimental evidence and the conclusion drawn was that the binding sites for the enzyme target and the inhibitor were in fact distant from each other. This implied the existence of a phenomenon in which binding events at one site were affecting the function of another distinct site which did not overlap the first site. This began a search to explain the mechanism behind the new phenomenon called allostery.

At the time, the only means of determining the structure of proteins was x-ray

crystallography. A consequence of this was that early studies of protein function were focused on highly ordered and relatively rigid proteins which could be more easily crystallized. This led to the earliest models of allostery which were based on structural changes such as the KNF (Koshland–Nemethy–Filmer)[2] and MWC (Monod–Wyman–Changeux)[3] models. These models were highly successful and were in good agreement with the evidence available at the time. However, as new experimental methodologies such as nuclear magnetic resonance were developed which could be applied to proteins in motion, it became clear that the allosteric mechanism can involve changes in protein dynamics within a native state basin rather than a large scale change in structure.

In 1984, Cooper produced a model of dynamic allostery which could explain the cooperativity of distant binding sites through alterations of normal modes[4]. Importantly, this model of dynamic allostery did not rely on any shift in the mean configuration of the protein. Since then, experimental evidence has been collected that shows allosteric effects can occur without conformational change[5, 6].

1.1 Allostery in Antibodies

The vertebrate immune system is responsible for the identification and elimination of foreign pathogens in the host body. From the earliest days of immunology, it was known that the immune system produces a class of substances which can identify and bind to invading pathogens and trigger their destruction. It is now known that large proteins called antibodies are responsible for the recognition and binding of these pathogens while many other parts of the immune system including B cells and phagocytes interact with the complexes formed between antibodies and pathogens in order to produce a range of immune responses.

Determining the system of interaction between pathogens, antibodies and immune cells was a long and arduous process involving many scientists and opposing theories which eventually led to the current model of the immune system. In this model, anti-

bodies possess a highly specific binding affinity for their target pathogen, also known as an antigen. Upon binding to this invasive antigen, the antibody can have one of several possible effects depending on the type of antibody ranging from directly disabling the pathogen to marking the pathogen for destruction by immune cells known as phagocytes which specialize in the consumption and destruction of pathogens. This later process is known as antibody-dependent cell-mediated cytotoxicity (ADCC).

The discovery that antibodies were necessary for the activation of other immune system cells in addition to their ability to bind pathogens immediately suggested the presence of multiple binding sites on the surface of an antibody which could each serve different roles. Experimental evidence eventually followed that confirmed these distant binding sites.




Figure 1.1: An IG antibody monomer is composed of two heavy chains (grey) and two light chains (orange). Each heavy chain is composed of three constant domains (CH1, CH2, CH3) and one variable domain (VH). Each light chain is composed of a single constant domain (CL) and a single variable domain (VL). The variable regions of each chain have three complementarity determining regions, or CDRs (purple).

An antibody monomer is a large Y-shaped protein which consists of four chains bound by inter-chain disulfide bonds. Two of these chains are known as heavy chains and are colored grey in the figure. Each heavy chain consists of a highly variable domain at one end (VH), and three consecutive constant domains (CH1, CH2, and CH3) which are the same across all members of the same antibody class (also known as an isotype) such as IgG4. The other two chains are known as light chains and are colored orange in the figure. Each light chain consists of a single variable domain (VL) and a single constant domain (VC). The high specificity of antigen-antibody

binding is achieved through a highly developed system of directed mutations which allow sections of the variable domains (VH and VL) to be altered in order to form binding sites with affinities for specific antigens. These antigen binding sites are also known as complementarity determining regions (CDRs). Binding sites which act as immune cell activators are located in the CH2 and CH3 domains. Different isotypes of antibody such as IgG1 and IgG4 have constant regions with different binding affinities which control the function of that isotype in the body.

While the possibility of allosteric communication between these sites was theorized early on, the general consensus given the evidence on hand was that non-allosteric mechanisms could explain the apparent cooperativity between these sites[7]. These non-allosteric mechanisms rely on the tendency of antibodies to cluster together as multiple antibodies bind to the same antigen. The non-allosteric model proposed that this increase in local antibody concentration can in turn be detected by immune cells.

As time progressed however, more evidence was produced that allostery might be at play in the function of antibodies[8]. One key experimental finding was that antibodies with identical variable Fab regions but different Fc constant regions could have different binding affinities.

1.2 Dynamic Allostery Model

The model we will use to quantify long range dynamic allosteric effects will involve several distinct stages. First, molecular dynamics simulations of the antibodies will be conducted in order to generate trajectories from which a large sampling of conformations can be drawn. Second, an effective elastic network model will be constructed which will allow the use of a perturbation-response model based on normal mode analysis. Third, the results from a full perturbation-response scan of each protein will be analyzed to determine if long range dynamic couplings can be detected which indicate possible dynamic allosteric effects. These stages will be elaborated below.

1.3 Molecular Dynamics Simulations of Antibodies

An IgG4 antibody is a very large molecule which consists of 1324 amino-acids with 10,268 atoms. To increase the speed of computation and allow for a greater sampled time, the Martini force field was chosen to simulate the antibodies[9]. The Martini force field features a course grained model which maps groups of 3-5 heavy atoms from the all-atom protein structure onto single interaction centers. Parameters for this force field have been determined by running extensive simulations of large bio-molecules in various solvents and adjusting the parameters in order to match experimentally determined thermodynamic values such as partitioning free energies[10]. All simulations were run with an explicit Martini water solvent. The conformational ensembles produced by these simulations are then used to construct covariance matrices of atomic positions which serve as components in the analysis of large scale collective motions which will take place in later steps.

1.4 Principal Component and Normal Mode Analysis

A typical protein contains thousands of atoms. Analyzing the individual movements of each atom therefore requires the use of a very high number of degrees of freedom. Principal component analysis (PCA) is a dimensionality reduction technique that has long been a tool used to study protein dynamics[11]. In PCA, the true dynamics of a protein are decomposed into a linear combination of orthogonal modes[12]. Normal mode analysis (NMA) is another method that can also be used to decompose molecular motions into orthogonal normal modes. Typically in NMA, a harmonic approximation is made regarding the potential energy functions underlying the molecular dynamics.

The first theoretical model of **dynamic** allostery was based on normal mode analysis of protein motions[4]. The relationship between PCA and NMA is critical for an understanding of our model. At issue is the ability to model molecular driving

forces with harmonic potentials and the method used to determine the values of the harmonic force constants.

The true energy landscape of a protein contains many local minima which mark semi-stable conformations between which proteins diffuse over time. The inter atomic forces which define this landscape are definitely not harmonic in nature. Inside an energy minimum however, it is often reasonably accurate to approximate any stable potential as harmonic. The reason for this can be seen from the Taylor expansion of any arbitrary potential energy function of a single variable $V(x)$:

$$V(x) = V(0) + V'(0)x + \frac{V''(0)}{2!}x^2 + \frac{V'''(0)}{3!}x^3 + \dots$$

[1]

Where x is the displacement from equilibrium. At the local energy minimum located at $x = 0$, the first term $V(0)$ is zero as we are free to define the lowest point of the potential as zero. The slope at a minimum is also zero therefore $V'(0)$ is also zero. Later terms involving x^n at higher powers will all approach zero as $x \rightarrow 0$, but terms with higher n will approach much faster than terms with lower n . For this reason, the arbitrary potential function $V(x)$ can be approximated as a harmonic potential, $\frac{V''(0)}{2!}x^2$, when x is sufficiently close to 0. This logic carries over to potential functions of more than one variable in regions sufficiently close to an energy minimum.

If a protein is in a conformation that is sufficiently close to a free energy minimum, then the potential energy functions which drive the protein's motion can be approximated as harmonic potentials. Our method takes into consideration a full range of harmonic and anharmonic potentials and differs from other methods enough to be dubbed a quasi-harmonic approximation[13].

Within this approximation, the motions of the protein can be decomposed into an orthogonal set of normal modes and their associated fractions of captured variance

given by the eigenvectors $|n\rangle$ and eigenvalues λ of a covariance matrix of atomic positions \mathbf{Q} :

$$\mathbf{Q} = \sum_{n=1}^N |n\rangle \lambda \langle n| \quad [2]$$

For each matrix \mathbf{Q} , a $2\mu s$ molecular dynamics (MD) simulation was conducted on a full IgG4 antibody using the Martini course grained force field. From each simulation, 10,000 configurations were sampled in order to create a matrix of atomic positions where $M = 10,000$ were the number of samples, the index k is the sample number, and the indexes i and j are degrees of freedom used to define the elements of \mathbf{Q} :

$$Q_{i,j} = cov(x^{(i)}, x^{(j)}) = \frac{1}{M-1} \sum_{k=1}^M (x_k^{(i)} - \langle x^{(i)} \rangle)(x_k^{(j)} - \langle x^{(j)} \rangle) \quad [3]$$

This work will use only the coordinates of the carbon alpha C_α atoms of the protein backbone. For a protein consisting of N amino acids (residues), the degrees of freedom are therefore $3N$. This level of course graining is thought to capture the configurational and dynamic effects which are important for dynamic allostery while leaving behind data regarding the movements of individual atoms. Small amplitude, high frequency motions of individual atoms have been found to have less long range correlations than lower frequency motions[11, 12].

Due to the incomplete sampling of atomic configurations, the covariance matrix \mathbf{Q} is only an approximation of the true covariance matrix. It has been shown that a process known as shrinkage removes most biases, approaching a true unbiased estimator, which improves the quality of estimated covariance matrices generated from a limited sample space[14]. This process has been implemented and its effect will be monitored and benchmarked because we obtain different results from the modified covariance

matrices. This will be done by comparing the new normal modes derived from these matrices to the old modes based on covariance matrices without shrinkage applied.

1.5 Effective Elastic Network Model

For a true ENM which assumes harmonic potentials as the only driving forces, the Hessian \mathbf{H} is proportional to the inverse of the covariance matrix \mathbf{Q}^{-1} [15].

$$\mathbf{H} = RT\mathbf{Q}^{-1} = RT \sum_{n=1}^N |n\rangle \frac{1}{\lambda} \langle n|$$

[4]

Our method uses this relationship in reverse by first determining the covariance matrix \mathbf{Q} from an MD simulation and then using the above equation to derive an approximation of \mathbf{H} . This approximation is dubbed a quasi-harmonic approximation because \mathbf{Q} was determined by an MD simulation which used anharmonic as well as harmonic potentials[13]. The force constants H_{ij} which are determined by inverting \mathbf{Q} in this way capture the anharmonic potentials of the MD simulation. The force constants of a true ENM approximate the forces between atoms as harmonic potentials without considering long range effects from anharmonic forces such as van der Waals forces. The force constants of our effective ENM approximate the potentials of mean force derived from an MD simulation which implicitly include these anharmonic forces[16].

The force constants which compose an effective ENM form a symmetric matrix which is mathematically equivalent to a Hessian matrix \mathbf{H} of second partial derivatives of potential energy V with each element H_{ij} of \mathbf{H} given by:

$$H_{ij} = \frac{\partial^2 V}{\partial x_i \partial x_j}$$

[5]

A true elastic network model (ENM) treats an entire protein as if it were a collection of masses connected to each other by quadratic Hookean (harmonic) potentials (springs). In most cases the masses consist of all C_α atoms. The force constants that make up an ENM are highly arbitrary but are usually parameterized to fit some experimental data. This simplified model allows for analytical calculation of the normal modes of vibration. Much research has been conducted on the accuracy of ENM predictions of protein motions[17, 15, 18]. A major advantage of ENMs is that the global effects of small perturbations to the network can be computed analytically through normal mode analysis.

The most glaring flaw in a true ENM is the fact that the forces driving atomic motions are not entirely harmonic. The effective ENM which our method uses differs from a true ENM in that anharmonic forces are taken into account when \mathbf{H} is being generated.

The elements of \mathbf{H} are analogous to force constants in a highly interconnected, multi-dimensional elastic network. For a mass on a spring moving in one dimension x , $\frac{\partial^2 V}{\partial x^2} = k$. Where k is the force or spring constant. If each mass in the network has 3 degrees of freedom and there are N masses, then H will have a dimension of $3N$. The eigenvectors $|n\rangle$ of \mathbf{H} are the normal modes of the models motion while the eigenvalues $\frac{1}{\lambda} = \omega$ are the frequencies of the associated normal modes squared.

It can be seen that the eigenvectors $|n\rangle$ of \mathbf{H} are identical to the eigenvectors of the covariance matrix \mathbf{Q} . The approximation $\mathbf{H} \propto \mathbf{Q}^{-1}$ is true only as long as the approximation of harmonic potential functions holds. \mathbf{Q}^{-1} therefore approximates \mathbf{H} more accurately the closer the conformation of the protein is to a true free energy minimum. The eigenvectors of a Hessian matrix are always the normal modes. The eigenvectors of a covariance matrix are estimates of the normal modes which rely on the approximation of harmonic potential functions.

1.6 Noise Decorrelation

The eigenvalues ω of \mathbf{H} are the inverses of the eigenvalues λ of \mathbf{Q} , $\omega = \frac{1}{\lambda}$. A problem with the computation of \mathbf{H} naturally arises from extremely small values of λ . The computation of λ relies on measurements of atomic positions which have some degree of uncertainty. There is a value $\Delta\lambda$, below which an estimate of λ becomes inaccurate due to the presence of these uncertainties. As a result, extremely small values of λ will cause noise to dominate the construction of \mathbf{H} .

The noise present in \mathbf{Q} creates numerical instability (for small lambda), and is the main reason why inverse covariance matrices have not been previously used to construct effective ENMs. Physically, any eigenvalue λ of a covariance matrix \mathbf{Q} with a magnitude below a certain limit is not accurate due to the uncertainty of the positional measurements used to derive \mathbf{Q} . To remedy this issue, a noise decorrelation step has been implemented[19] which replaces all eigenvalues of \mathbf{Q} below a certain threshold value $\lambda_n < \lambda_c$ with a mean value according to:

$$\lambda_n \rightarrow \frac{1}{3N - c} \sum_{n=c}^{3N} \lambda_n$$

[6]

This transformation reduces the effect that extremely small and uncertain eigenvalues of \mathbf{Q} have on the construction of \mathbf{H} while at the same time maintaining the value of the trace of \mathbf{Q} . The sum of the eigenvalues that are below λ_c divided by the sum of all eigenvalues gives the fraction of total variance that de-noised:

$$\text{Percent noise decorrelation} = \frac{\sum_{n=c}^{3N} \lambda_n}{\sum_{n=1}^{3N} \lambda_n} * 100$$

[7]

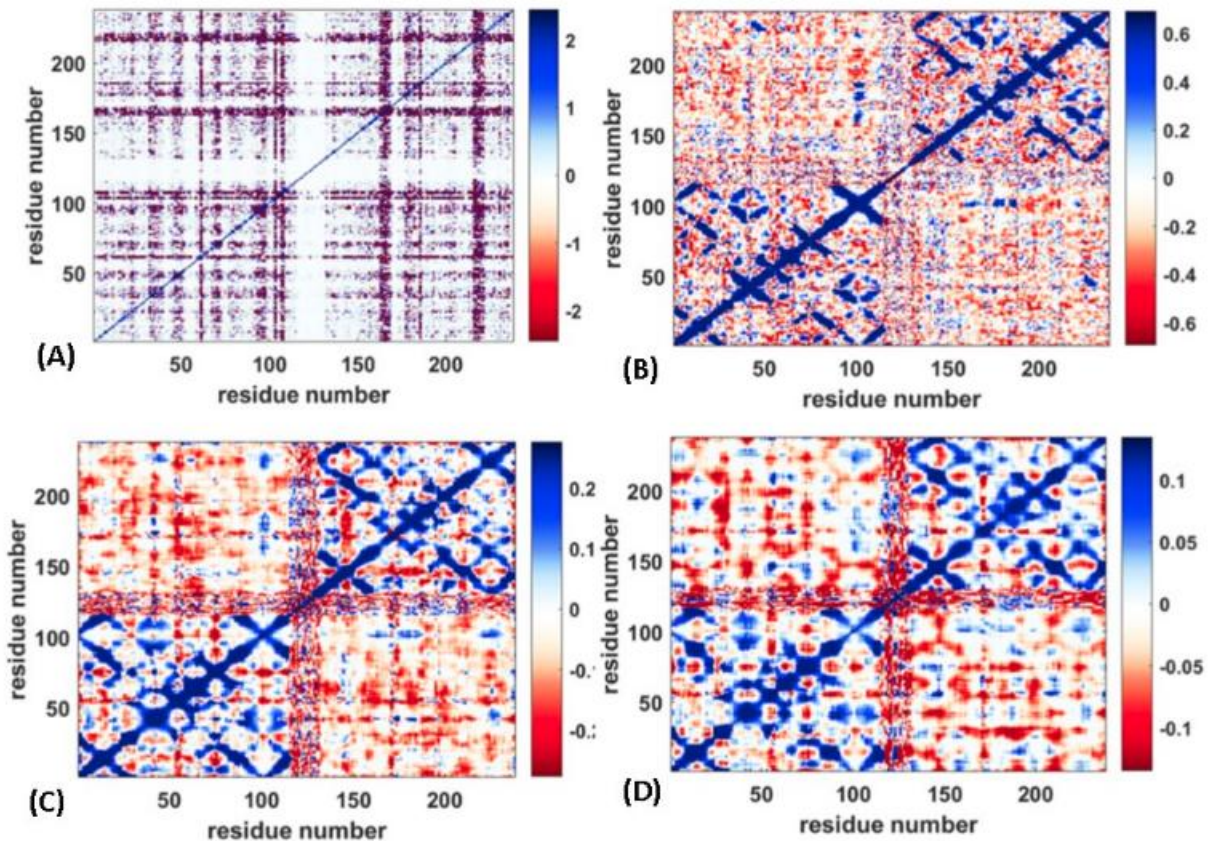


Figure 1.2: Hessian matrices at varying levels of noise decorrelation. (A) Zero percent noise decorrelation. (B) 5% noise decorrelation. (C) 10% noise decorrelation. (D) 20% noise decorrelation.

Figure 2 shows four examples of Hessian matrices which represent a 238 residue scFv antibody fragment[19]. Each Hessian was derived from the inversion of a covariance matrix of atomic positions. The matrix on the top left is a Hessian derived from a covariance matrix with zero noise decorrelation. The remaining three matrices are Hessians derived from the same covariance matrix with 5%, 10% and 20% noise decorrelation by replacement of the lowest eigenvalues with their mean value. 10% indicates that the sum of the replaced eigenvalues account for 10% of the total variance.

An alternative noise decorrelation method is introduced in this thesis which replaces all eigenvalues below λ_c with a fixed minimum value of λ_c which is based on

the estimated accuracy limit of λ . This new method has the advantage of a better physical basis than the previous method which simply set an arbitrary percentage of noise decorrelation as a user defined parameter. This alternative method will be implemented as a part of this thesis and the new results will be compared to results generated using the previously published method.

1.7 Perturbation-Response Model

A model has been developed which measures possible dynamic coupling between specific binding sites on the protein which are distant from each other[19]. The denoised Hessian matrix \mathbf{H} can be decomposed into a set of eigenvectors $|n\rangle$ which are the normal modes of vibration of the protein. Each normal mode indicates the amplitudes of motion of every mass along each degree of freedom. By perturbing \mathbf{H} in a way that models a binding event and then tracking alterations to $|n\rangle$, dynamic effects across the entire protein can be measured.

Binding between a protein and a ligand is almost always accompanied by a decrease in the proteins heat capacity. This observation suggests that internal vibrational degrees of freedom are being reduced as a result of the binding event[4]. By adding new force constants or "springs" to the elastic network model, localized reductions to vibrational freedom can be induced to model binding events at that location.

The addition of new springs to the Hessian \mathbf{H} perturbs its normal modes $|n\rangle$. These perturbed modes can be used to calculate the changes in mean squared fluctuation (MSF) for every mass in the ENM which in this case is every residue or C_α atom. Changes in the MSF of residues which are distant from the perturbation site indicate that a long range dynamic coupling has occurred which may imply a dynamic allosteric effect.

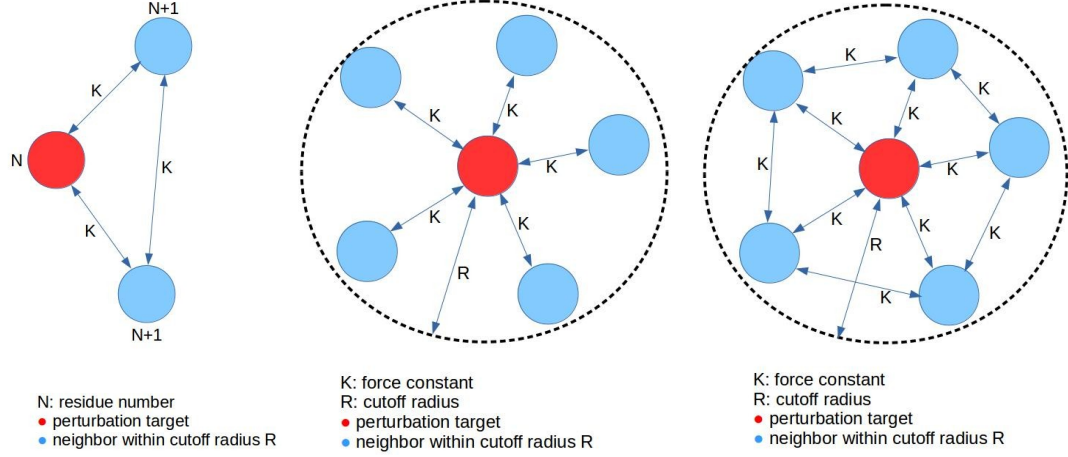


Figure 1.3: Triad, star and ball perturbation types. Spring constants 'k' are added to the model between sets of atoms defined by the perturbation type. The perturbation radius 'R' defines a cutoff distance between the target residue (red) and nearby residues which are to be included in the perturbation (blue).

Several types of perturbations have been tested which vary in the way they add springs to the model. Triad, star and ball perturbations are shown in figure 3. All results and analysis presented here use the ball perturbation type. This perturbation adds springs to the model such that all masses within a user defined radius R_p of the target residue have a spring placed between them with a uniform spring constant value of k . While previous work focused on the ball perturbation type, triad and star perturbations will be included as a part of this thesis. A triad perturbation at a target residue p adds springs between residue pairs $(p-1, p)$, $(p+1, p)$, $(p-1, p+1)$ while a star perturbation adds springs only between the target residue and each other residue within the cutoff radius R . Springs with either positive or negative constants can be added to explore the effects of binding events which either decrease or increase the vibrational freedom of the targeted residues.

The addition of new springs to the model does not alter the equilibrium position of the protein but instead alters the normal modes $|n\rangle$ by changing fluctuations around the equilibrium conformation. As the normal modes change, the MSFs of residues across the protein can potentially be altered. Each perturbation produces a new

set of normal modes which are determined by a new eigenvalue decomposition of the perturbed Hessian \mathbf{H}_p . If p is the index of the residue where a perturbation is centered then a new set of n normal modes and eigenvalues $\lambda_{p,n}$ will be associated with that perturbation due to the additional force constants in the elastic network. Each element of a normal mode gives the amplitude of motion in that degree of freedom. Each residue has three degrees of freedom: x, y and z . For an N residue protein, the normal modes will have $3N$ elements.

For a perturbation at target residue p , the calculation of MSF for residue r is as follows:

$$MSF(p, r) = \sum_n \lambda_{p,n} [V_{x,p,n}^2(r) + V_{y,p,n}^2(r) + V_{z,p,n}^2(r)]$$

[8]

Where $\lambda_{p,n}$ is the eigenvalue of the n th normal mode associated with the perturbed Hessian \mathbf{H}_p and $V_{x,p,n}(r)$ is the element of the n th eigenvector of \mathbf{H}_p that represents the x degree of freedom for residue r . In this way, the MSFs of residues which are distant from the perturbation target are measured and a possible dynamic coupling effect has been detected. Changes in MSF due to a perturbation are referred to as "responses".

A full perturbation-response matrix (PRM) tabulates the changes in MSF of every residue as a result of perturbations to all other residues. A full PRM therefore consists of an $N \times N$ matrix where N is the number of residues in the protein. Figure 4 is an example of a full PRM where the vertical axis indicates the perturbation target residue p and the horizontal axis indicates the response residue r .

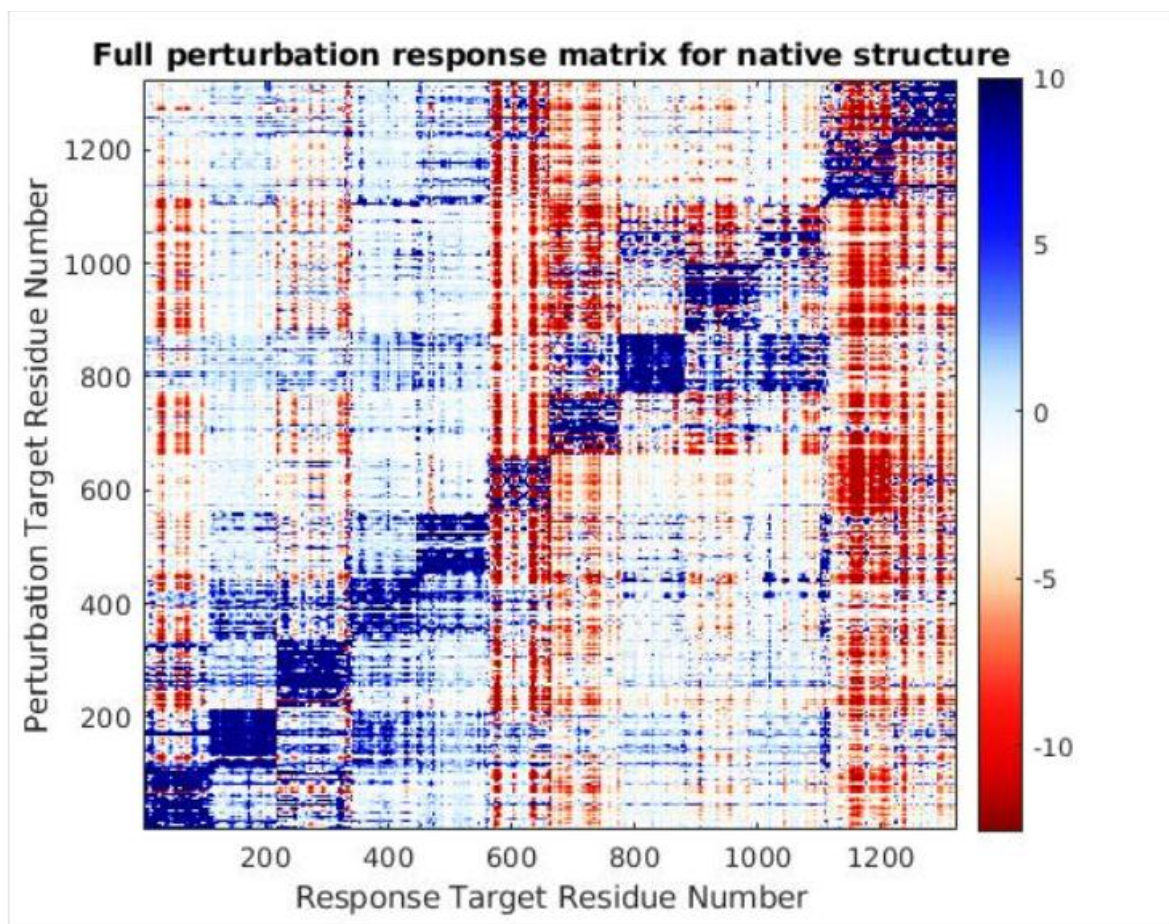


Figure 1.4: All residue perturbation-response matrix. This matrix was generated using a positive/stabilizing force constant. Every blue point on the matrix represents a residue that was stabilized as a result of stabilizing perturbation. Every red point represents a residue that was destabilized. The location of the perturbation is given by the y axis. The location of the response is given by the x axis.

This perturbation-response model will locate and measure long range dynamic couplings between residues. In addition to full PRMs which perturb and measure all residues, targeted PRMs can be made which select only specific residues at which to apply perturbations or measure responses.

1.8 Cooperative Binding and Allosteric effects

In order for dynamic allosteric effects to occur, the long range dynamic couplings described above must result in some sort of cooperativity in binding between distant

sites. The model of dynamic allostery proposed by Cooper and Dryden attempted to calculate the change in the free energy of binding ($\Delta\Delta G_{vib}$) which would occur if the binding of one ligand had an effect on the binding affinity of a second ligand[4]. This model was based on the observation that binding between a protein and its ligand usually resulted in a stiffening or loss of flexibility in the protein which would cause an increase in its normal mode frequencies ν . Their proposed equation took the form:

$$\Delta\Delta G_{vib} \approx -KT \cdot \ln\left(\frac{\nu_1^2}{\nu_0\nu_2}\right) \quad [9]$$

Here ν_0 is the frequency of the given normal mode when no ligand is bound, ν_1 is the frequency of this mode with the first ligand bound, and ν_2 is the frequency with both ligands bound. The condition for $\Delta\Delta G_{vib}$ to be negative is clearly $\nu_1^2 > \nu_0\nu_2$. If ligand binding is assumed to cause an increase in ν then a negative $\Delta\Delta G_{vib}$ implies a free energy driven cooperative effect between the first and second ligands.

1.9 Eigenvalue Perturbations

Our method models binding events by adding springs to the effective elastic network, altering the rigidity of the network at that location. This perturbs the Hessian matrix $\mathbf{H}_0 \rightarrow \mathbf{H}_p$. The perturbed hessian \mathbf{H}_p has a new set of normal mode frequencies given by its eigenvalues. In order for $\Delta\Delta G_{vib}$ to be computed using the above method, it is necessary to compare the frequencies of specific modes as they transition from $\nu \rightarrow \nu_1 \rightarrow \nu_2$. Separate eigenvalue decompositions of \mathbf{H}_0 and \mathbf{H}_p do not guarantee the proper alignment of eigenvalues for this method to be valid.

To solve this issue we used first order eigenvalue perturbation theory to track the changes in ν for each mode $\langle n|$ separately. For an unperturbed Hessian \mathbf{H}_0 the eigenvalues are the normal mode frequencies squared:

$$\mathbf{H}_0 |n\rangle = \omega_0 |n\rangle$$

Where $\omega_0 = \nu_0^2$. For a small initial perturbation V_1 followed by a small second perturbation V_2 , the original unperturbed frequencies are altered according to the equations:

$$\omega_1 = \omega_0 + \langle n | V_1 | n \rangle$$

[10]

$$\omega_2 = \omega_0 + \langle n | V_1 + V_2 | n \rangle$$

[11]

In this case, each perturbation V takes the form of a sparse matrix of force constants which are added to \mathbf{H} . Using this method, the frequencies ν_0 , ν_1 , and ν_2 are kept aligned for each mode. This enables us to compute $\Delta\Delta G_{vib}$ for each mode.

Antibodies are a natural choice on which to test these methods as there is significant evidence that distant binding sites on an antibody communicate allosterically[20, 8]. Despite this, no proposed mechanism has been put forth which conclusively explains allostery in such a large protein. If our methods predict binding cooperativity between CDR loops and residues in the CH2 regions when applied to antibodies, then the results will be in significant agreement with experimental evidence.

CHAPTER 2: RESULTS

2.1 Analysis of dynamic couplings

Our method has several parameters associated with each perturbation. These include the strength of the force constant k , the cutoff radius R_p , and the geometry which can be either ball, star, or triad (Figure 3). Many combinations of these parameters were sampled in order to explore potential differences in long range couplings. Global perturbation and response profiles are a good way to initially quantify the differences between the effects of perturbations.

A perturbation which adds positive spring constants to the model is referred to as a *stabilizing* perturbation due to its effect of reducing the motions of the perturbed residues. The addition of negative spring constants has the opposite effect and is therefore a *destabilizing* perturbation. The results below all involve stabilizing perturbations due to the commonly accepted belief that a binding event will decrease the mobility of atoms which are involved in the binding. The local reduction of mobility caused by a stabilizing perturbation can have either stabilizing or destabilizing effects on residues distant from the perturbation site.

A global perturbation scan begins by applying a perturbation to each residue in turn. As each residue is perturbed, every other residue in the protein will potentially experience an alteration in the magnitude of its MSF. The average of all these changes in MSF is recorded for each perturbation target. Each single residue in the protein is therefore associated with an average global change in MSF when it is perturbed. A *perturbation profile* is created in this way.

Alternately, as each residue is perturbed during the scan, the MSF of a single fixed residue can be measured. The average change in MSF at a single residue as a result

of all other residues being perturbed in turn creates a *response profile*. Below is a figure which has examples of both global perturbation and response profiles.

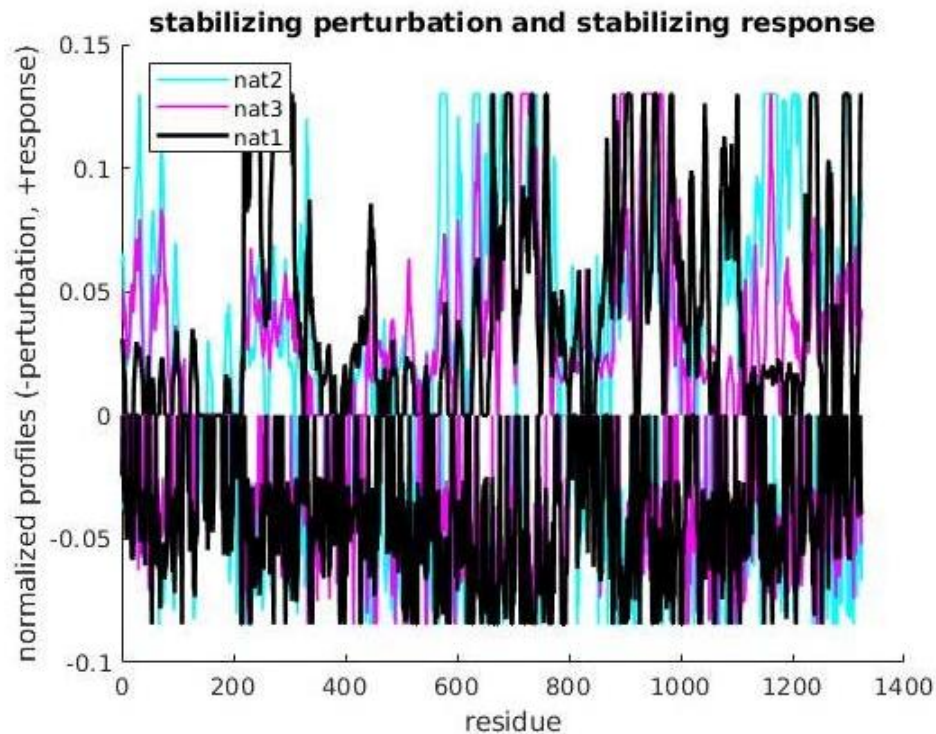


Figure 2.1: An example of global perturbation and response profiles organized on the same plot. The three colors indicate perturbation scan conducted on each of three ENMs generated from different native structure simulations. The points above the axis represent the global response profile while points below the axis represent the global perturbation profile as defined above.




Figure 2.2: Stabilizing perturbation-Stabilizing response (SS) perturbation profiles for all three perturbation types. These images show the magnitudes of *stabilizing* responses to *stabilizing* perturbations. Blue regions indicate that perturbations placed here cause a higher than average global *decrease* in MSF.




Figure 2.3: Stabilizing perturbation-Destabilizing response (SD) perturbation profiles for all three perturbation patterns. These images show the magnitudes of *destabilizing* responses to *stabilizing* perturbations. Red regions indicate that perturbations placed here cause a higher than average global *increase* in MSF.




Figure 2.4: Stabilizing perturbation-Stabilizing response (SS) response profiles for all three perturbation patterns. These images show the magnitudes of *stabilizing* responses to *stabilizing* perturbations. Blue regions indicate that these residues experience a higher than average *decrease* in mobility when other residues across the entire protein are perturbed.




Figure 2.5: Stabilizing perturbation-Destabilizing response (SD) response profiles for all three perturbation patterns. These images show the magnitudes of *destabilizing* responses to *stabilizing* perturbations. Red regions indicate that these residues experience a higher than average *increase* in mobility when other residues across the entire protein are perturbed..

The data shown above involves a global perturbation scan which sequentially perturbs all residues in the protein. The results of this global scan can give information about the effects binding events have on the protein as a whole. To determine interactions between specific residues a targeted perturbation can be used. By targeting only a select few residues for perturbation, a response profile can be generated which shows the response each residue in the protein experiences as a result of the specific targeted residue being perturbed. This is important for determining the presence of dynamic allostery between known or suspected binding sites.

If allostery is present in antibodies, the most likely binding sites on the antibodies surface which will be involved will be the complementarity determining regions (CDRs) and the residues of the CH2 regions. The commonly accepted model of antibody function involves highly specific antigen binding at the CDRs. After binding to its antigens, the antibody gains the ability to activate immune effector cells by binding to receptors on the surface of these cells. The sites on the antibody which bind to immune cells surface receptors are located in the CH2 regions of the antibodies. The cooperativity between CDRs and CH2 region sites which is observed in the functioning of the immune system is the key evidence that allostery may be present in antibodies.

For this reason, our choice of perturbation targets in the following test runs are CDRs. Perturbations in these regions simulate antigen binding. These targeted perturbations will lead to response profiles which will indicate exactly which residues in the antibody respond dynamically to binding events at the perturbation targets.

For a targeted perturbation of CDR-L1a involving all three perturbation types, correlations between *response* profiles as well as visual maps of regions which responded to the perturbations are given below. The numbers in the upper triangle represent correlations between SS profiles, while the numbers in the lower triangle represent DD profiles:

Table 2.1: Correlations between SS and DD profiles generated with different perturbation types at CDR-L1a. Numbers in the upper triangle represent correlations between SS profiles (blue), while the numbers in the lower triangle represent DD profiles (red).

| | Ball | Star | Triad |
|-------|--------|--------|--------|
| Ball | x | 0.9854 | 0.9490 |
| Star | 0.9354 | x | 0.9681 |
| Triad | 0.9227 | 0.8903 | x |




Figure 2.6: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the L1a CDR loop (purple). Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target.

For a targeted perturbation of CDR-H3a involving all three perturbation types, correlations between *response* profiles as well as visual maps of regions which responded to the perturbations are given below. The numbers in the upper triangle represent correlations between SS profiles, while the numbers in the lower triangle represent DD profiles:

Table 2.2: Correlations between SS and DD profiles generated with different perturbation types at CDR-H3a. Numbers in the upper triangle represent correlations between SS profiles (blue), while the numbers in the lower triangle represent DD profiles (red).

| | Ball | Star | Triad |
|-------|--------|--------|--------|
| Ball | x | 0.9530 | 0.8597 |
| Star | 0.9305 | x | 0.8865 |
| Triad | 0.7971 | 0.7918 | x |




Figure 2.7: Stabilizing-Destabilizing (SD) response profiles for targeted perturbations of the H3a CDR loop (purple). Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target.

From the above figures, it can be seen that the residues most likely to be dynamically destabilized as a result of perturbations to either CDR-L1a or CDR-H3a are located in the CDRs of the opposite chain as well as the CH2 regions of both chains. The fact that small perturbations to a single CDR loop have significant destabilizing effects on these very distant regions is evidence that dynamic allosteric effects are possible.

2.2 Shrinkage of Estimated Covariance Matrices.

One of the improvements I implemented was the application of shrinkage to the MD positional covariance matrices. The usefulness of shrinkage arises from the problem of imperfect sampling of positions during the MD simulation. When these imperfect

measurements are used to construct an estimate of the true covariance matrix \mathbf{Q} , they introduce systematic errors. These errors can be minimized through a process known as shrinkage. When conducted at the optimal level, the process of shrinkage is thought to almost universally improve the quality of covariance matrices when the number of samples is not sufficiently large relative to the number of dimensions[14].

A quantitative measure of shrinkage required to optimize a covariance matrix is shrinkage factor δ . A large value of δ implies that the estimated covariance matrix in question is very likely to have systematic errors, while a small value of δ implies that the estimated covariance matrix most likely has enough samples to minimize these errors. When applied to our covariance matrices, we determined the optimal shrinkage factors were always around $\delta = 2.56 \times 10^{-4}$. This very low shrinkage factor has virtually zero impact on the estimated covariance matrices.

This is a very good result for the current data. In the future, when our methods are used on other protein, the optimal shrinkage will be implemented automatically. If the initial quality of estimated covariance matrices are lower in the new proteins, the level of shrinkage implemented will increase to a level which is guaranteed to minimize certain systematic errors.

2.3 New Noise Decorrelation Method.

The second improvement to our method was a new noise decorrelation method. Noise decorrelation is a necessary step when inverting a covariance matrix because very low eigenvalues can become overly large in the inverted matrix. These very low covariance eigenvalues represent very small amplitude motions which are often dominated by noise. The previous method was designed to remove these extremely low eigenvalues by replacing a set of the lowest eigenvalues with the average value of the set. The set of eigenvalues to be replaced was determined by choosing a threshold percentage. This percentage would be the fraction of total variance represented by

the replaced eigenvalues.

The new method is based on an estimate of the accuracy of positional data taken from an MD simulation. Eigenvalues which are below the square of this estimate are simply replaced with the threshold value. A threshold value of 0.01\AA^2 indicates that all eigenvalues below 0.01\AA^2 are replaced with 0.01\AA^2 . This implies an expected accuracy limit of 0.1\AA in the MD positional data because the eigenvalues of the covariance matrix are in units of Angstroms squared.

The degree to which both methods differ in their effects can be measured in several ways. Both ND methods produce de-noised Hessian matrices with complete sets of orthogonal eigenvectors (normal modes) and eigenvalues (mode frequencies). An overlap matrix can be computed from any two Hessian matrices by projecting every eigenvector from one Hessian onto all eigenvectors from the other Hessian. An overlap matrix between two identical sets of vectors will be perfectly diagonal, while sets of vectors which differ from each other will have proportionally greater off-diagonal elements. A single covariance matrix was de-noised using both methods and each resulting de-noised matrix was inverted to give a Hessian matrix. Below is an overlap matrix between the normal modes of these two Hessians.




Figure 2.8: The overlap matrix of all eigenvectors from Hessian matrices derived using both ND methods. A diagonal overlap matrix implies a very high degree of similarity between the two sets of vectors.

The virtual lack of off-diagonal elements in the overlap matrix above means that the normal modes derived using both ND methods are very similar to each other. Noise decorrelation directly changes the eigenvalues of the covariance matrices. The figure below compares scree plots (descending eigenvalues) from covariance matrices which have had both methods of noise decorrelation applied to them.




Figure 2.9: Descending eigenvalue scree plots for covariance matrices de-noised using the old (left) and new (right) noise decorrelation methods.

The scree plots above are nearly identical upon first examination. This is due to the fact that the top covariance eigenvalues which represent high amplitude, low frequency motions are not affected by the noise decorrelation. An more obvious effect of the new noise decorrelation method can be seen after inverting both of the de-noised covariance matrices from above and plotting the Hessian matrix eigenvalues in ascending order.




Figure 2.10: Ascending Hessian matrix eigenvalue plots for the old (left) and new (right) noise decorrelation methods. These plots show that the old noise decorrelation method introduces a sharp jump in the low frequency modes which is not present in the new method.

Using the old noise decorrelation method, there is a significant jump around the 35th eigenvalue of \mathbf{H} . This was caused by the fact that over 90% of the total variance of the original covariance matrix \mathbf{Q} came from the first 35 or so modes. With the old noise decorrelation threshold set to 10%, all modes of the covariance matrix \mathbf{Q} after the first 35 or so were replaced with a fixed number. In the old method, this number was the average of the replaced eigenvalues. In this case the value turned out to be approximately 0.0005.

Using the old ND method with a threshold of 10%, replacing all of the eigenvalues of \mathbf{Q} past the 90% variance threshold does not have a noticeable effect on a scree plot of descending eigenvalues. When \mathbf{Q} is inverted to form a Hessian matrix \mathbf{H} , the largest eigenvalues of \mathbf{Q} become the smallest eigenvalues of \mathbf{H} . Plotting the eigenvalues of

\mathbf{H} from smallest to largest, a significant step can be seen at the point where the noise decorrelation began replacing eigenvalues. Only the first 35 modes of \mathbf{H} were not affected because they correspond to the largest eigenvalues of \mathbf{Q} and were not replaced. All modes of \mathbf{H} after this threshold were increased significantly.

Using the new noise decorrelation method and with a threshold of 0.01, there is again no significant effect on the scree plot for \mathbf{Q} . The ascending plot of \mathbf{H} eigenvalues, however, no longer shows a sudden jump anywhere in the first 500 eigenvalues. This is the first evidence that that new method is significantly less invasive to the data than the old method. The eigenvalues of \mathbf{Q} do not drop below the new method’s threshold until after the first several thousand eigenvalues. From the plot of ascending \mathbf{H} eigenvalues above, we can see that the new method is a much more subtle correction than the old method. With the old method, only the first 35 out of 3972 eigenvalues were left uncorrected. This produced a sharp step in the lowest eigenvalues of \mathbf{H} . The new method only replaced the last few hundred eigenvalues of \mathbf{Q} with a fixed value, leaving far more of the data intact.

The old ND method replaced over 3900 of the lowest eigenvalues with the number 0.0005, beginning with eigenvalues as high as 0.0195. The new method only replaced a few hundred of the lowest eigenvalues with the number 0.01, and no eigenvalues over 0.01 were replaced at all. This establishes the new method as significantly less invasive to the original data when dealing with our system and using the old ND threshold of 10%. This could be expected from any system where the vast majority or the variance is captured in the first few modes of \mathbf{Q} . At much lower thresholds, the old method could be less invasive to the data for our system. The new method avoids the problem of determining the correct de-noising threshold for each system.

The Hessian matrices \mathbf{H} define the parameters of protein elastic network models (ENMs). These Hessian matrices are constructed by inverting covariance matrices \mathbf{Q} derived from MD simulations. The noise decorrelation step serves to remove un-

wanted effects from the ENMs which enter into the model at the matrix inversion step. Given two ENMs, one constructed with the old noise decorrelation method and one constructed with the new method, we can compare the global perturbation and response profiles quantitatively by computing the correlations between profiles.

Correlations between global response profiles can give an initial indication of how much the new noise decorrelation methods affect our results. Correlations between SS (blue) and DD (red) global response profiles generated using the old and new methods of noise decorrelation (ND) are shown below:

Table 2.3: Correlations between SS (blue) and DD (red) global response profiles generated using the old and new methods of noise decorrelation (ND).

| ND method | New | Old |
|-----------|--------|--------|
| New | x | 0.9979 |
| Old | 0.9382 | x |

Correlations between SD (red) and DS (blue) response profiles generated using the old and new methods of noise decorrelation (ND) are shown below:

Table 2.4: Correlations between SD (red) and DS (blue) response profiles generated using the old and new methods of noise decorrelation (ND).

| ND method | New | Old |
|-----------|--------|--------|
| New | x | 0.8690 |
| Old | 0.9823 | x |

From these correlations between response profiles, we can see that the new noise decorrelation method has some minor but measurable effects on global responses to perturbation when compared to the old method. The most significant effect was seen in the SD response profiles with only 0.869 correlation.

The effects of different ND methods on the response profiles of targeted perturbations are also of interest. Two perturbation targets of interest, CDR-L1a and

CDR-h3a, were perturbed using the old and new noise decorrelation methods. These sites are very likely to be involved in allosteric mechanisms in antibodies if allostery is present at all. The results can be seen graphically in the figures below:




Figure 2.11: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the L1a CDR loop (purple) using the old and new noise decorrelation methods. Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. From this figure, the new noise decorrelation method does seem to have minor effects on the dynamic coupling between the CDR-L1a perturbation target and other residues.




Figure 2.12: Stabilizing perturbation-Destabilizing response (SD) response profiles for targeted perturbations of the H3a CDR loop (purple) using the old and new noise decorrelation methods. Regions colored yellow or red were destabilized to a significantly higher degree than others when a perturbation was applied to the target. The new noise decorrelation method seems to have minor effects on the dynamic coupling between the H3a-L1a perturbation target and other residues.

The original purpose of noise decorrelation was to remove the oversized effects of small, noise dominated eigenvalues from covariance matrices. While both methods succeeded in this task, the new method does so with much less adjustment to the original data. Additionally, the lowest eigenvalues of the Hessian matrices \mathbf{H} are required when computing binding cooperativity $\Delta\Delta G$, as will be seen in the next step. The sharp change in low \mathbf{H} eigenvalues which is seen when using the old method is expected to have significant effects on this calculation. From these results it would seem that the new method is preferable to the old method. If the old method is still used then a much lower threshold than 10% would be recommended to avoid overly modifying the original data beyond what is necessary.

2.4 Modeling cooperative binding effects

To calculate the binding cooperativity $\Delta\Delta G$ between two binding sites, our method follows the arguments outlined by Cooper and Dryden[4]. Note that the exact $\Delta\Delta G$

between two sites is not actually being computed here, but rather the $\Delta\Delta G_{vib}$ which is the contribution to $\Delta\Delta G$ from changes in vibrational modes. We refer to $\Delta\Delta G_{vib}$ as $\Delta\Delta G$ throughout this paper. As described above, the value of $\Delta\Delta G$ for a single normal mode is given by:

$$\Delta\Delta G_{vib} \approx -KT \cdot \ln\left(\frac{\nu_1^2}{\nu_0\nu_2}\right)$$

[9]

Our method modifies this approximation so that $\Delta\Delta G$ can be computed between any two sites by computing the changes in normal modes ν as perturbations are applied to each of the sites in turn. Sites can be either single residues or groups of residues. Our method is currently implemented such that an initial perturbation site is specified by the user which represents the binding location of the first ligand. At this point the normal modes of the unperturbed protein (ν_0 modes) and the normal modes of the single ligand-bound protein (ν_1 modes) have been computed. A scan is then performed which places a new perturbation at every residue in the protein in turn and computes the two ligand-bound normal mode frequencies (ν_2 modes) associated with each residue.

With each perturbation comes possible changes to all of the 3972 Hessian eigenvalues (3 degrees of freedom for each of 1324 residues). These eigenvalues give the frequencies of the normal modes. Each mode is associated with a value of $\Delta\Delta G$ as the frequency of the mode changes from ν_0 to ν_1 to ν_2 as a result of perturbations. The sum of $\Delta\Delta G$ for all modes with respect to a given pair of perturbations gives the total $\Delta\Delta G$ between the two perturbation sites. Changes to all eigenvalues can be computed and stored as requested by the user. In practice only the first few frequencies representing the lowest frequency modes are required to compute the value

The following results are based on series of cooperativity scans which began with

one of two initial perturbation sites: CDR-L1a and CDR-H3a. CDR loops are expected to play a major roll in the ability of an antibody to recognize and bind antigens and then subsequently bind to and activate immune cell receptors. The two chosen to be shown here are partly arbitrary although CDR-H3a is suspected of playing a more significant roll than other loops. The figure below shows the locations of these CDR loops.




Figure 2.13: The locations of three perturbation targets are show above in purple. These include CDR loops L1a, H3a, and L1b. Residue 470 is colored red and indicates the residue that was most commonly the location of highest cooperativity (most negative $\Delta\Delta G$) across all perturbations.

After initially perturbing CDR-L1a to simulate a binding event there, a secondary perturbation scan is conducted across all other residues to simulate secondary binding events at these other locations. The figures below show the results of this perturbation scan.

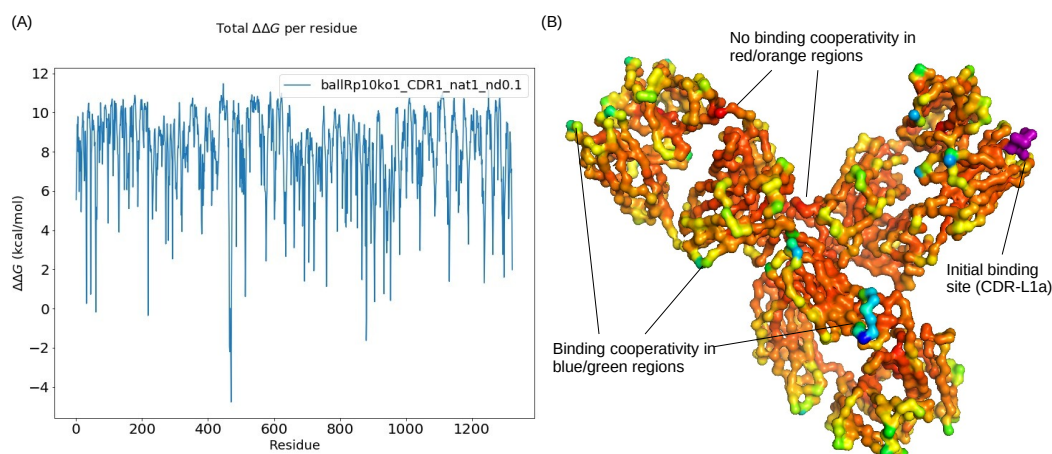


Figure 2.14: (A) shows a plot of $\Delta\Delta G$ values which associate a measure of cooperativity between an initial binding site at CDR-L1a (purple) and all other residues. (B) Orange to red colored sites show anti-cooperativity with CDR-L1a, while green to blue sites show positive cooperativity (more negative $\Delta\Delta G$ values).

Also of interest are the values of $\Delta\Delta G$ which belong to individual modes given a pair of perturbation sites. Below is a plot which shows $\Delta\Delta G$ values between an initial CDR-L1a perturbation site and residues 467-472.

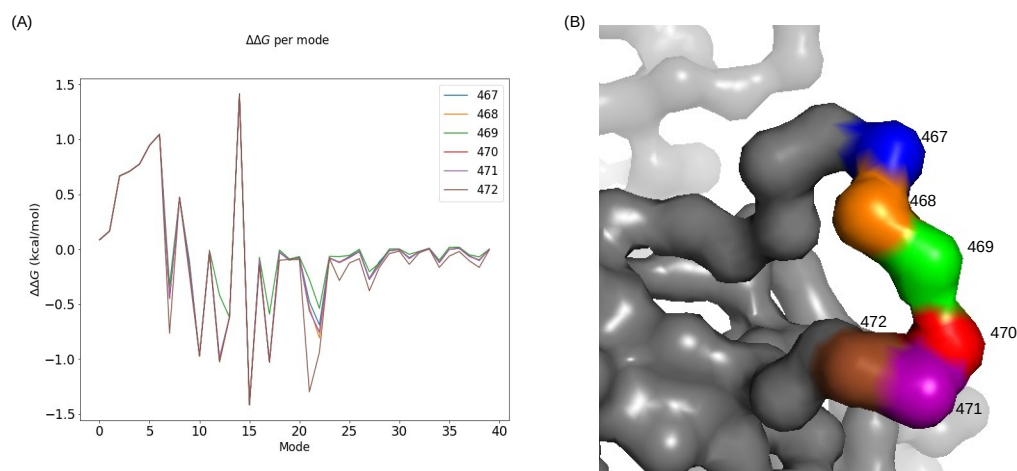


Figure 2.15: Values of $\Delta\Delta G$ per mode between CDR-L1a and residues 467-472. Across all combinations of parameters, these residues showed the highest degree of binding cooperativity (lowest $\Delta\Delta G$) when the initial perturbation targets were either CDR-L1a or CDR-H3a.

The first characteristic that stands out from these plots is the widely varying level of $\Delta\Delta G$ between different modes. Many modes $\Delta\Delta G$ values are above zero, indicating a lack of cooperativity. Some modes $\Delta\Delta G$ values are below zero and in these cases a cooperative effect is mediated by the mode. The locations of the peaks are of interest. If the maximum or minimum values of $\Delta\Delta G$ are always carried by the same modes, it would indicate that these modes are possibly responsible for some kind of important function. Since the most negative values of $\Delta\Delta G$ represent the greatest degree of cooperativity, the modes which have the lowest values of $\Delta\Delta G$ are of interest. Below is chart of the five modes which carry the lowest values of $\Delta\Delta G$ across five different arbitrary residues when CDR-L1a is perturbed with a spring constant of $k = 1$:

```
Residue 100, highest cooperativity modes:  [12 20 19 15 28]
Residue 200, highest cooperativity modes:  [19 15 39 33 28]
Residue 700, highest cooperativity modes:  [11  9  7 19 32]
Residue 1000, highest cooperativity modes: [10 12  7 19 11]
Residue 1200, highest cooperativity modes: [ 7 11 28 39 33]
```

At the moment, there is no clear pattern to this data. This is not surprising however, as a single mode which is always negative regardless of the residues in question would indicate highly non-specific binding cooperativity between a large number of residues. This kind of non-specific cooperativity is not known or suspected in antibodies.

The next most notable aspect of these results is that the levels of $\Delta\Delta G$ drop to zero around the 40th mode. This drop is found in every plot we have seen so far. The implications of this are that binding cooperativity between distant residues is mediated almost entirely by the first few modes as originally suggested by Cooper and Dryden. These modes represent the lowest frequency motions of the protein. These results are in agreement with the predictions of Cooper and Dryden who argued that

long range dynamic allosteric effects would be mediated by low frequency modes, while higher frequency modes are highly localized and tend to only involve small clusters of residues.

Our method computes $\Delta\Delta G$ values for all modes as default. This allows us to sum whatever number of modes we wish when computing the net binding $\Delta\Delta G$. While including all or even most modes does not appear to be necessary to compute $\Delta\Delta G$ in this case, the option exists.

The next results of interest are plots of which sum the $\Delta\Delta G$ values of many modes to give a total value of $\Delta\Delta G$ between the initial target and other residues. Below is a plot which shows the total $\Delta\Delta G$ per residue when the initial CDR-L1a perturbation site is perturbed using the three different perturbation patterns:




Figure 2.16: Binding cooperativity $\Delta\Delta G$ between CDR-L1a and all other residues using all three perturbation types (ball, star, and triad) at a spring constant of $k = 10$.

The first thing to note in the above plots are the significantly different magnitudes of $\Delta\Delta G$ between different perturbation types. The next thing to note is the very

obvious location of certain peaks. For a perturbation applied to CDR-L1a with a spring constant $k = 1$ and either ball, star or triad perturbation patterns, the locations of the five most negative $\Delta\Delta G$ peaks from each plot are given below:

Ball, highest cooperativity residues: [471 467 470 469 466]

Star, highest cooperativity residues: [471 469 470 467 468]

Triad, highest cooperativity residues: [470 469 1 468 1324]

From this data, we can see that certain highly localized sites are the locations of maximum cooperativity (lowest $\Delta\Delta G$). Residue 470 and other very nearby residues show a very clear tendency to be the sites of maximum cooperative binding when CDR-L1a is perturbed. This is a very interesting result because these residues are located inside the CH2 region. This region is the location of binding sites which interact with immune cells.

Despite the highly varying average magnitudes of $\Delta\Delta G$ across the three perturbation types, there is a strong correlation in the plots shape between ball and star perturbation types:

Table 2.5: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-L1a and all other residues for ball, star and triad at $k=1$ native simulations

| Perturbation Type | Ball | Star | Triad |
|-------------------|------------|------------|------------|
| Ball | 1. | 0.91277062 | 0.22216516 |
| Star | 0.91277062 | 1. | 0.26581246 |
| Triad | 0.22216516 | 0.26581246 | 1. |

The effects of different spring constant values k were explored next. The plots below show results for a ball type perturbation applied to CDR-L1a with $k = 0.1, 0.5, 1, 5, 10, 20$:

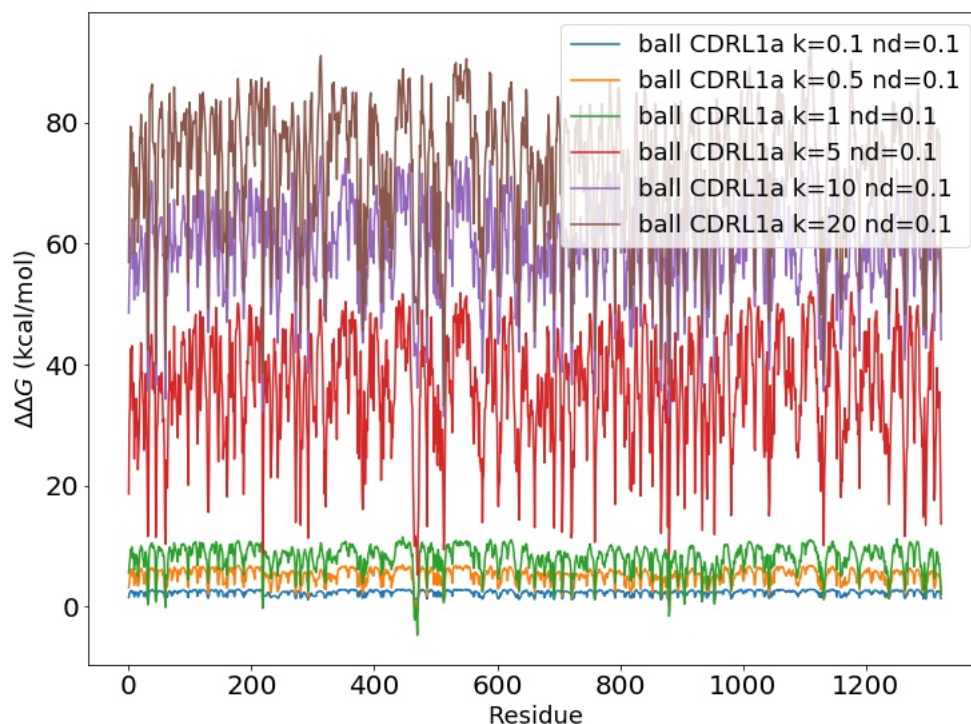


Figure 2.17: Binding cooperativity $\Delta\Delta G$ between CDR-L1a and all other residues at various values of spring constant k and a ball type perturbation.

There is an obvious trend in this data which shows that increasing the value of k increases the values of $\Delta\Delta G$, thereby decreasing or eliminating any cooperativity. Lower values of k are therefore recommended for future cooperativity scans. The most accurate or realistic value of k to use when attempting to emulate a binding event is currently speculative. These results make a case for lower values of k .

The locations of minimum $\Delta\Delta G$ peaks (highest cooperativity) in the above plots are given below for a ball type initial perturbation at CDR-L1a and varying values of k :

$k = 0.1$, highest cooperativity residues: [219 881 471 955 61]

$k = 0.5$, highest cooperativity residues: [471 881 469 470 467]

$k = 1$, highest cooperativity residues: [471 467 470 469 466]

$k = 5$, highest cooperativity residues: [471 881 219 514 469]

$k = 10$, highest cooperativity residues: [881 514 219 61 760]

$k = 20$, highest cooperativity residues: [514 515 881 1133 219]

Residues 470 and 880, along with nearby residues make a very strong showing in the above data. It should be noted that residue 880 is located in inside the CH2 region of the chain opposite to CDR-L1a (antibodies are composed of four chains). This is further confirmation of a strong cooperative binding phenomenon between CDR-L1a and residues in the CH2 region.

Another interesting result is the highly correlated shapes of the $\Delta\Delta G$ plots. Despite their vastly different magnitudes, the plots show a very high correlation:

Table 2.6: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-L1a and all other residues for all k values: $k = 0.1, 0.5, 1, 5, 10, 20$. These values indicate that the relative regions of cooperativity and anti-cooperativity remain fairly constant as the perturbation strength is increased. They also indicate that $\Delta\Delta G$ does not respond in a perfectly linear way to increasing perturbation strength.

| k | 0.1 | 0.5 | 1 | 5 | 10 | 20 |
|-----|------------|------------|------------|------------|------------|------------|
| 0.1 | 1. | 0.93645819 | 0.88480329 | 0.91617338 | 0.89082965 | 0.87302275 |
| 0.5 | 0.93645819 | 1. | 0.95266221 | 0.92870632 | 0.90928372 | 0.90654884 |
| 1 | 0.88480329 | 0.95266221 | 1. | 0.90503942 | 0.8873517 | 0.90383461 |
| 5 | 0.91617338 | 0.92870632 | 0.90503942 | 1. | 0.96634485 | 0.92892235 |
| 10 | 0.89082965 | 0.90928372 | 0.8873517 | 0.96634485 | 1. | 0.94053016 |
| 20 | 0.87302275 | 0.90654884 | 0.90383461 | 0.92892235 | 0.94053016 | 1. |

These correlations show that perturbations of very different strengths have similar $\Delta\Delta G$ profiles across the protein. This data also shows that $\Delta\Delta G$ is not linearly proportional to perturbation strength. A perfectly linear response would give a correlation of 1 as the perturbation strength was increased. This non-linear response can be reasoned mathematically from the way in which $\Delta\Delta G$ is calculated. First, we can see that individual normal mode frequencies ν respond quadratically to the strength

of the perturbation:

$$\nu_0^2 = \omega_0$$

$$\nu_1^2 = \omega_1 = \omega_0 + \langle n | V_1 | n \rangle$$

$$\nu_2^2 = \omega_2 = \omega_0 + \langle n | V_1 + V_2 | n \rangle$$

If the strength of a perturbation is proportional to k by some constant m we can write:

$$\nu_1^2 = b + m_1 k \rightarrow \nu_1^2 = \nu_0^2 + m_1 k$$

$$\nu_2^2 = b + m_2 k \rightarrow \nu_2^2 = \nu_0^2 + m_2 k$$

For a perturbation of strength 0, $\nu_1^2 = \nu_2^2 = \nu_0^2$. We can then rewrite $\Delta\Delta G$ as:

$$\Delta\Delta G \approx -KT \cdot \ln\left(\frac{\nu_0^2 + m_1 k}{\nu_0 \sqrt{\nu_0^2 + m_2 k}}\right) = -KT \cdot \ln\left(\frac{\nu_0^2(1 + \frac{m_1 k}{\nu_0^2})}{\nu_0^2 \sqrt{1 + \frac{m_2 k}{\nu_0^2}}}\right)$$

$$\Delta\Delta G \approx -KT \left[\ln\left(1 + \frac{m_1 k}{\nu_0^2}\right) - \frac{1}{2} \ln\left(1 + \frac{m_2 k}{\nu_0^2}\right) \right]$$

[12]

From the above equation we can see that $\Delta\Delta G$ of each mode will respond linearly only for values of k such that $\frac{m_1 k}{\nu_0^2} \ll 1$ and $\frac{m_2 k}{\nu_0^2} \ll 1$.

A second perturbation target, CDR-H3a was chosen due to its suggested role in mediating antibody-antigen binding specificity. The results of this cooperativity scan are shown below:




Figure 2.18: Binding cooperativity $\Delta\Delta G$ between CDR-H3a and all other residues at various values of spring constant k and a ball type perturbation.

Again we see the values of $\Delta\Delta G$ are highly dependent on the values of k . Below are the locations of highest cooperativity as a result of ball type perturbations to CDR-H3a at different values of k :

$k = 0.1$, highest cooperativity residues: [219 881 32 220 955]
 $k = 0.5$, highest cooperativity residues: [471 469 470 467 881]
 $k = 1$, highest cooperativity residues: [471 881 467 469 470]
 $k = 5$, highest cooperativity residues: [471 466 467 469 470]
 $k = 10$, highest cooperativity residues: [471 514 1133 515 881]
 $k = 20$, highest cooperativity residues: [881 1133 61 471 219]

Again the regions around residues 470 and 880 show the highest cooperativity, the same as for perturbations to CDR-L1a. This is significant because CDR-H3a is located on a different chain than CDR-L1a, although they are relatively close to

each other structurally. This suggests a conserved dynamic cooperative pathway.

Correlations between cooperativity plots are again very similar:

Table 2.7: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-H3a and all other residues for all k values: k = 0.1, 0.5, 1, 5, 10, 20

| k | 0.1 | 0.5 | 1 | 5 | 10 | 20 |
|-----|------------|------------|------------|------------|------------|------------|
| 0.1 | 1. | 0.88397599 | 0.8260738 | 0.8305561 | 0.8392471 | 0.8247699 |
| 0.5 | 0.88397599 | 1. | 0.94095141 | 0.90476379 | 0.88522425 | 0.89195969 |
| 1 | 0.8260738 | 0.94095141 | 1. | 0.88655648 | 0.88182968 | 0.89292393 |
| 5 | 0.8305561 | 0.90476379 | 0.88655648 | 1. | 0.88376531 | 0.8994009 |
| 10 | 0.8392471 | 0.88522425 | 0.88182968 | 0.88376531 | 1. | 0.94053016 |
| 20 | 0.8247699 | 0.89195969 | 0.89292393 | 0.89940098 | 0.94053016 | 1. |

Some arguments have been made that molecular dynamics simulations do not sample a complete configurational space and that covariance matrices constructed from these simulations may be inaccurate. To partly address this issue, three independent $2\mu s$ MD simulations were run for an IgG4 antibody. Each of these simulations began with a different starting configuration referred to as Native 1, Native 2 and Native 3. Each simulation was used to generate an ENM using the procedures above. These ENMs were then perturbed in the same ways so that the results could be compared. Cooperativity plots for these three structures are shown below with CDR-L1a as a perturbation target:

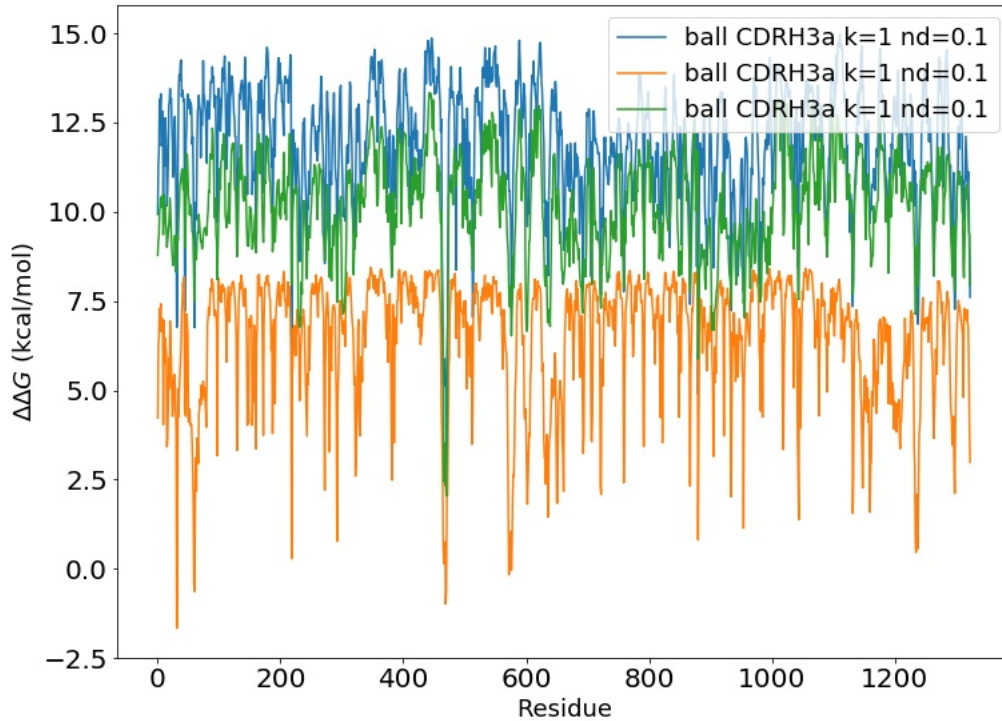


Figure 2.19: Binding cooperativity $\Delta\Delta G$ between CDR-H3a and all other residues using a spring constant $k = 10$ compared using three different simulations of a native IgG4 structure.

Some significant differences in the shapes and magnitudes of these plots is apparent. This is not entirely surprising as it has been suggested that allosteric pathways may only be present in certain configurations. The difference between active and inactive pathways could be mediated by a shift in the distribution of the configurational ensemble[21]. The differences between these plots can be quantitatively seen in their correlations:

Table 2.8: Correlation matrix of cooperativity ($\Delta\Delta G$) values between CDR-H3a and all other residues for 3 simulations which were started from different molecular configurations of the same protein.

| Starting Structure | Native 1 | Native 2 | Native 3 |
|--------------------|------------|------------|------------|
| Native 1 | 1. | 0.62773771 | 0.76928955 |
| Native 2 | 0.62773771 | 1. | 0.62222122 |
| Native 3 | 0.76928955 | 0.62222122 | 1. |

What is interesting however, is that the residues of peak cooperativity are conserved between the simulations which were started with different molecular configurations. Comparing the locations of minimum $\Delta\Delta G$ between the simulations for ball perturbation type with spring constant $k = 1$ at CDR-H3a:

Native 1, Minimum ddG residues: [471 881 467 469 470]

Native 2, Minimum ddG residues: [32 470 61 471 60]

Native 3, Minimum ddG residues: [471 470 469 467 472]

Despite the differences in shape and magnitude, the pathway linking the CDRs to residues in the CH2 regions is conserved. This could be an implication that the motions which mediate the cooperativity between these particular sites has been adequately sampled by all three simulations, and that these pathways are not artifacts of a particular simulation.

CHAPTER 3: CONCLUSIONS

The methods developed here have shown a plausible mechanism by which dynamic allostery could be mediated between distant binding sites on large IgG4 antibody proteins. Our methods begin with data from molecular dynamics simulations of an antibody. The conformations sampled during these simulations enable us to estimate positional covariance matrices which can be used to determine the essential dynamics of the antibody using PCA. Under certain circumstances, these essential dynamics are equivalent to the normal modes of vibration predicted by an elastic network model. By analyzing changes to these normal modes as the results of binding events, we have shown that dynamic coupling exists between distant residues which can account for long range dynamic allosteric effects.

The estimated positional covariance matrix encodes all of the atomic motions experienced by the protein during the MD simulation. Inversion of this positional covariance matrix is known to give a Hessian matrix which defines the force constants of an effective elastic network model (ENM). This process can be used to develop ENMs which have superior estimates of both short and long range forces when compared to simpler ENMs. The inversion step brings with it a potential problem, however, when extremely small eigenvalues of the covariance matrix create too much noise in the Hessian. To address this issue, we implemented two noise decorrelation methods which would reduce this noise while retaining data which was relevant to the computation of dynamic allosteric effects.

While both noise decorrelation methods successfully removed very low eigenvalues from the computation of the Hessian matrices, the newer of the two methods proved to be far more conservative when deciding how much of the original data should

be left intact. The old method could have compensated for this with a much lower percent decorrelation threshold which is useful to know if the old method is used in the future.

Using one of the original formulations of dynamic allostery[4], we showed that a model based on global normal mode perturbations could account for allosteric phenomena in an IgG4 antibody. The simulations we ran in this study focused on possible allosteric pathways which exist between two specific CDR loops and the rest of the antibody. Our results showed a level of agreement between our model and experimentally determined properties of antibodies. Notably, binding events at CDR loops L1a and H3a reduced the binding free energy at highly localized residues in the CH2 region.

Changes in locations and magnitudes of cooperative binding effects were modeled over a range of values for the perturbation strength (spring constant k) and perturbation type (ball, star or triad) parameters. A high degree of conservation was found in the shapes of the plots but the overall magnitudes of $\Delta\Delta G$ consistently rose or fell with different parameters. This seems to imply that the shape of the cooperativity landscape is strongly embedded in the protein while the absolute values of the cooperativity rise and fall roughly together.

Dependence on initial structure has been a known issue for MD simulations for some time. We accounted for this by running three separate simulations for our antibody with different starting structures. While the shapes of the $\Delta\Delta G$ plots differed to some extent, they all agreed upon the location of binding sites which showed maximum levels of cooperativity when initial binding sites were the CDR loops.

More experimentation on antibodies is planned in the future. The pipeline developed in this model can easily be applied to other proteins as well. Molecular dynamics data of any protein is all that is required for our method to be used. Our method does not inherently rely on previously determined knowledge of allosteric sites. This

gives our method the potential to identify novel sites which may be relevant to the design of drugs which bind to target proteins at sites distant from an active site. This work only sampled a fraction of possible allosteric sites in a single protein. Even so, the plausibility of dynamic allosteric effects in IgG4 antibodies was consistently demonstrated.

REFERENCES

- [1] J.-P. Changeux, “The feedback control mechanism of biosynthetic l-threonine deaminase by l-isoleucine,” in *Cold Spring Harbor symposia on quantitative biology*, vol. 26, pp. 313–318, Cold Spring Harbor Laboratory Press, 1961.
- [2] D. Koshland Jr, G. Nemethy, and D. Filmer, “Comparison of experimental binding data and theoretical models in proteins containing subunits,” *Biochemistry*, vol. 5, no. 1, pp. 365–385, 1966.
- [3] J. Monod, J. Wyman, and J.-P. Changeux, “On the nature of allosteric transitions: a plausible model,” *J Mol Biol*, vol. 12, no. 1, pp. 88–118, 1965.
- [4] A. Cooper and D. Dryden, “Allostery without conformational change,” *European Biophysics Journal*, vol. 11, no. 2, pp. 103–109, 1984.
- [5] C. M. Petit, J. Zhang, P. J. Sapienza, E. J. Fuentes, and A. L. Lee, “Hidden dynamic allostery in a pdz domain,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 43, pp. 18249–18254, 2009.
- [6] N. Popovych, S. Sun, R. H. Ebright, and C. G. Kalodimos, “Dynamically driven protein allostery,” *Nature structural & molecular biology*, vol. 13, no. 9, p. 831, 2006.
- [7] H. Metzger, “The effect of antigen on antibodies: recent studies,” in *Contemporary topics in molecular immunology*, pp. 119–152, Springer, 1978.
- [8] D. Yang, R. Kroe-Barrett, S. Singh, C. J. Roberts, and T. M. Laue, “IgG cooperativity—is there allostery? implications for antibody functions and therapeutic antibody development,” in *MAbs*, vol. 9, pp. 1231–1252, Taylor & Francis, 2017.
- [9] S. J. Marrink, H. J. Risselada, S. Yefimov, D. P. Tieleman, and A. H. De Vries, “The martini force field: coarse grained model for biomolecular simulations,” *The journal of physical chemistry B*, vol. 111, no. 27, pp. 7812–7824, 2007.
- [10] D. H. de Jong, G. Singh, W. D. Bennett, C. Arnarez, T. A. Wassenaar, L. V. SchaĀlfer, X. Periole, D. P. Tieleman, and S. J. Marrink, “Improved parameters for the martini coarse-grained protein force field,” *Journal of Chemical Theory and Computation*, vol. 9, no. 1, pp. 687–697, 2012.
- [11] B. Brooks and M. Karplus, “Harmonic dynamics of proteins: normal modes and fluctuations in bovine pancreatic trypsin inhibitor,” *Proceedings of the National Academy of Sciences*, vol. 80, no. 21, pp. 6571–6575, 1983.
- [12] A. Amadei, M. A. Ceruso, and A. Di Nola, “On the convergence of the conformational coordinates basis set obtained by the essential dynamics analysis of proteins’ molecular dynamics simulations,” *Proteins: Structure, Function, and Bioinformatics*, vol. 36, no. 4, pp. 419–424, 1999.

- [13] R. Levy, A. Srinivasan, W. Olson, and J. McCammon, “Quasi-harmonic method for studying very low frequency modes in proteins,” *Biopolymers: Original Research on Biomolecules*, vol. 23, no. 6, pp. 1099–1112, 1984.
- [14] J. Schäfer and K. Strimmer, “A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics,” *Statistical applications in genetics and molecular biology*, vol. 4, no. 1, 2005.
- [15] I. Bahar, A. R. Atilgan, and B. Erman, “Direct evaluation of thermal fluctuations in proteins using a single-parameter harmonic potential,” *Folding and Design*, vol. 2, no. 3, pp. 173–181, 1997.
- [16] J. A. McCammon, P. G. Wolynes, and M. Karplus, “Picosecond dynamics of tyrosine side chains in proteins,” *Biochemistry*, vol. 18, no. 6, pp. 927–942, 1979.
- [17] I. Bahar, T. R. Lezon, L.-W. Yang, and E. Eyal, “Global dynamics of proteins: bridging between structure and function,” *Annual review of biophysics*, vol. 39, pp. 23–42, 2010.
- [18] M. M. Tirion, “Large amplitude elastic motions in proteins from a single-parameter, atomic analysis,” *Physical review letters*, vol. 77, no. 9, p. 1905, 1996.
- [19] A. S. Ettayapuram Ramaprasad, S. Uddin, J. Casas-Finet, and D. J. Jacobs, “Decomposing dynamical couplings in mutated scfv antibody fragments into stabilizing and destabilizing effects,” *Journal of the American Chemical Society*, vol. 139, no. 48, pp. 17508–17517, 2017.
- [20] J. Zhao, R. Nussinov, and B. Ma, “Antigen binding allosterically promotes fc receptor recognition,” in *Mabs*, vol. 11, pp. 58–74, Taylor & Francis, 2019.
- [21] H. N. Motlagh, J. O. Wrabl, J. Li, and V. J. Hilser, “The ensemble nature of allostery,” *Nature*, vol. 508, no. 7496, pp. 331–339, 2014.