How the immune system generates diversity: Pathogen space coverage with random and evolved antibody libraries

Mihaela Oprea
Dept. of Computer Science
University of New Mexico
Albuquerque, NM 87131-1386
mihaela@cs.unm.edu

Stephanie Forrest
Dept. of Computer Science
University of New Mexico
Albuquerque, NM 87131-1386
forrest@cs.unm.edu

Abstract

The immune system uses many strategies to generate its enormous repertoire of diverse antibodies, but their relative importance is not understood. Here we address the contribution of antibody gene libraries to the antibody repertoire. We introduce a general framework, in which we can study many antibody-pathogen matching rules, including the widely-used shape-space model (Perelson and Oster, 1979). We use the genetic algorithm as a model of evolution to investigate the type of antibody repertoires that might evolve in relation to a given pathogenic environment. For the antibody/pathogen matching rules that we studied, the scaling relation between fitness and the size of the evolved antibody library is only a shifted variant of the scaling relation that we obtain with random libraries of the same size. We discuss how our results compare to the antibodies that are expressed in newborns, and we discuss the implications of our results for recent experiments with phage antibody libraries.

1 INTRODUCTION

In order to respond effectively to a wide variety of pathogens, the immune system must generate a diverse set of immune receptors. This is accomplished by a number of diversity generating mechanisms which have been identified experimentally (Tonegawa et al., 1975; Gilfillan et al., 1993; Weigert et al., 1970). At the same time, the immune system only has finite resources, and we assume that there is some evolutionary pressure to use them efficiently. As the immune system cannot allocate one particular antibody for each possible pathogen that it might encounter, a natural hypothesis is that the antibody gene libraries reflect the evolutionary history of antigenic exposures of the species. However, immune responses have been induced in mice to artificially-produced molecules, suggesting that the immune system is able to recognize more than the antigens that the species encountered in its evolution. The following question then arises: What type of information do immune receptor libraries encode?

Using a model based on the shape-space concept introduced by Perelson and Oster (1979), we previously argued that the scaling of the survival probability of an organism with the size of its antibody repertoire makes it unlikely that germline diversity is the major contributor to immune system diversity (Oprea and Forrest, 1998). We argued that the germline repertoire induces a coarse-graining of the pathogen space, mapping the regions of this space that are essential for the survival of the organism. Here we extend our earlier analysis to antibody-pathogen matching rules that might be more realistic. Our approach is sufficiently general that it can be extended as more data on antibody-pathogen interactions becomes available. In the shape-space model, individual fitness is determined by finding the antibody with the smallest Hamming distance from the pathogen. This fitness landscape is highly structured. However, we are interested in what happens in the case where the fitness landscape has a completely different structure, one that is possibly more closely related to what is known about how biological molecules interact with one another. In this paper, we explore what happens when fitness is based on the idea of a random energy model, introduced by Derrida (1984), in the context of spin glasses. In this model, each bit string is assigned an energy drawn from a Gaussian distribution. We use the random-energy model to approximate the details of intermolecular interaction, as will be apparent in the following section, and we also extend our results to energy distributions other than Gaussian.

2 BASIC MODEL

If we view the the antigen-antibody interaction from a biochemical standpoint, the strength of the bond is given by the difference of the free energies of the complex on one hand, and of the two molecules in their unbound state, on the other hand. A realistic representation of the energy landscape as a function of the sequence of the molecules is beyond our current knowledge and computational power. Therefore, we use the following abstraction. We assume that each molecule has an "energy," which is a random deviate from a Gaussian distribution. The antigen-antibody complex also has an energy corresponding to it, which is likewise a random deviate from a Gaussian distribution. The difference between the energy of the complex and the energy of unbound molecules gives the strength of the bond between them. We use a genetic algorithm to evolve libraries of different sizes on a large pathogen set, and determine how the maximum fitness of an evolved individual scales with the size of its antibody library. One might argue that the landscape that we have constructed does not have any obvious structure for the genetic algorithm to work with, given that the energies assigned to closely related genotypes are random deviates from the Gaussian distribution. The landscape does, however, has some structure, as the antibodies with high energy have a better chance of lowering this energy by binding to pathogens. These are exactly the antibodies that the genetic algorithm discovers, as we will see.

Our genetic algorithm resembles the one introduced by Hightower (1996) to study the shape-space model of antibody library evolution. We consider a population of M individuals, called hosts, which are evolving in an environment of hostile pathogens, each pathogen represented as a bit string. Each individual in the population consists of an antibody library, containing A antibodies, each antibody represented as a bit string of length L. For the experiments described below, we chose $L = 16$. Pathogens are also represented as bit strings of length L. We evolve the antibody libraries on a pathogen set $\mathcal{P}$, of size 2^9, setting the 7 high order bits to 0 in all pathogen strings. We chose these parameters to match the setting in our previous study (Oprea and Forrest, 1998). Our representation of antibody libraries is reminiscent of the so-called "Pitt" approach to classifier systems in that we concatenate A antibodies together to form a single chromosome. Under this analogy, each library (one individual's genome) is analogous to a classifier system if we consider each encoded antibody to take the role of a single classifier rule. It is interesting that this aspect of our representation corresponds quite directly to, for example, V-region genes in humans.

The essence of the complicated antibody-pathogen interaction in the real world, which we try to capture in our model, is that for each pathogen in the environment there is at least one antibody in the individual's library that can bind to it. Moreover, the antibody with the highest affinity for a given pathogen will be the one that dominates the response to that pathogen. This phenomenon is known as clonal selection (see for example Takahashi, 1998). We use this property as the basis for our fitness function. To each individual, consisting of a single library $\mathcal{A}$, we assign a score σ in matching a pathogen p, which we define as

$$\sigma(p) = \max_{a \in \mathcal{A}} \beta(a, p),$$

where $\beta(a, p)$ is the strength of the bond between the antibody a and the pathogen p. To calculate the bond strength, we first determine the "energy" of the antibody in its unbound state, the "energy" of the pathogen in its unbound state, and, finally, the "energy" of the antibody-pathogen complex. The difference between the sum of the first two quantities and the last one of them gives the bond strength. The energy of each pathogen (antigen) and antibody is drawn from a Gaussian distribution with mean 50, and variance 2.5, whereas the energy of the complex was chosen from a Gaussian distribution with mean 100 and variance 10. The exact choice of the mean and variance of the energy of an individual molecule is clearly somewhat arbitrary, a topic that we hope to address in future work.

To determine the energy of each "molecule," we seed the random number generator with the integer representation of the bit string representing that "molecule," and then calculate a pseudo-random Gaussian deviate according to the algorithm given in Numerical Recipes (Press et al., 1988). We assign such an energy to each antigen and each antibody. To obtain the antigen-antibody complex, we take the XOR between the bit strings representing the antigen and the antibody, and then use the integer representation of the XOR string to calculate its energy, as described above. The bond strength, given by the difference in energy between the (sum of) unbound molecules and the complex, will be distributed as a Gaussian with mean 0 and variance 15.

In Hightower (1996) the fitness f of an individual was identified with its average score $\langle\sigma\rangle$ over all pathogens that it encountered. We use the same definition of fitness here. This choice is justified because the survival probability of an individual depends on all pathogen challenges it encounters (Oprea and Forrest, 1998). Thus, the fitness f is defined as:

$$f = \frac{1}{P} \sum_{p \in \mathcal{P}} \sigma(p) \equiv \langle\sigma\rangle.$$

Let us briefly summarize the genetic algorithm we used to evolve the libraries. We construct the initial population of $M = 50$ random libraries, of identical size, A. Each individual, then, consists of a single library. In the framework of the random energy model, we may, in fact, view the antibody library as exactly the antibody repertoire. Adding more realism to the model by using multiple libraries for each individual would not affect the results. A population size of 50 is sufficiently large to allow convergence to relatively high fitness solutions, given the mutation rate of 0.002 per bit that we used in evolving the libraries. We use rank selection as follows: If r is the rank of the fitness of an individual in the population, the chance of that individual being selected as a parent is, on average, $w_r = \frac{2(M-r)}{M(M-1)}$. To create one library of the new generation, we select, with replacement, two libraries of the old population. We generate two new libraries by crossing over the two chosen libraries. The number of crossover points n is chosen from a binomial distribution with mean $0.01A$. This is because chromosomal crossover in real genetic systems is not a deterministic process. Assuming that there is a constant crossover rate per gene, the number of crossover points per individual will then obey a binomial distribution. The crossover points are chosen at the boundary between antibodies, so individual antibodies are not disrupted by crossover. We then choose one of the new crossover products, mutate it, and add it to the new population. 1000 generations of the genetic algorithm constitute a run. At the end of the run, we take the library with the highest fitness in the population and use it for subsequent analysis.

3 RESULTS

3.1 SCALING RELATION BETWEEN FITNESS AND LIBRARY SIZE.

Our previous study showed that, for the shape-space model, the scaling relation between fitness and library size for evolved libraries is only a shifted variant of the relation obtained for a random library of identical size. For both cases (evolved and random libraries), the scaling relation indicates a sub-logarithmic dependence of fitness on library size. We interpreted this result as showing that the germline-encoded antibodies are not a large contributor to the overall fitness of an individual and that other sources of diversity are likely more important. We hypothesized that the role of the germline-encoded repertoire is more likely to extract essential features of the pathogen space that the species has encountered in evolution. However, in order to draw such strong conclusions we need to show that our scaling result also holds for the more general case of the random energy model described above.

Let us first determine the fitness of a random library as a function of the library size. We write the derivation in terms of the density distribution of the bond strength, $g(x)$, and its corresponding cumulative density function, $G(x)$, and we will then apply it to the particular Gaussian distribution described above. For every pathogen, the fitness is given by the maximum of A random variables drawn from the distribution G, A being the size of the antibody library. The probability that the bond strength between a random pathogen and all of the antibodies in the library is less than or equal to a value, x, is $G(x)^A$, and the derivative of this gives the probability density of fitness x:

$$\begin{aligned} g_A(x) &= \frac{d}{dx}\left[(G(x))^A\right] \\ &= A \times g(x)(G(x))^{A-1}. \end{aligned} \tag{1}$$

Now, the fitness of a random library of A antibodies on the complete pathogen space, given the probability density function of the fitness, $g_A(x)$, is

$$\begin{aligned} f_g(A) &= \int_0^\infty x g_A(x) \\ &= \int_0^\infty x \frac{d}{dx}\left[G_A(x)\right]. \end{aligned} \tag{2}$$

Let $y = G_A(x)$, taking values between 0 and 1. Then $\frac{d}{dx}\left[G_A(x)\right] = dy$ and Eq. 2 can be rewritten in terms of y as

$$f_g(A) = \int_0^1 x(y)dy, \tag{3}$$

where $x(y)$ denotes the fact that x has to be expressed now as a function of y. But $y = G_A(x) = (G(x))^A$, thus $G(x) = y^{\frac{1}{A}}$, and $x = G^{-1}(y^{\frac{1}{A}})$, where G^{-1} denotes the inverse function of G. With this, Equation 3 becomes

$$f_g(A) = \int_0^1 G^{-1}(y^{\frac{1}{A}})dy. \tag{4}$$

In the case of the Gaussian distributed bond strengths, mentioned above, we cannot derive an analytical form for the fitness dependency on antibody library size, as we cannot analytically invert the error function, which is the integral of the normal distribution. We may, however, compute the values numerically, and this is how we generated the data for random antibody libraries shown in Fig. 1 (the dashed line). As mentioned above, for the case that we studied, the bond strengths are Gaussian distributed, with mean $\mu = 0$, and variance $\sigma^2 = 15$.

Fig. 1 shows how fitness scales with the library size A for the Gaussian distribution discussed above. As was the case for the shape-space model, the evolved libraries attain a fitness that has a similar functional dependency on the library size as the random libraries. The dependency is sublogarithmic, that is, the fitness increases more slowly than linear as a function of the logarithm of the library size. Thus, the shape-space model, with a binomial distribution of bond strengths, is well approximated by the Gaussian distributed bond strengths, as we expected.

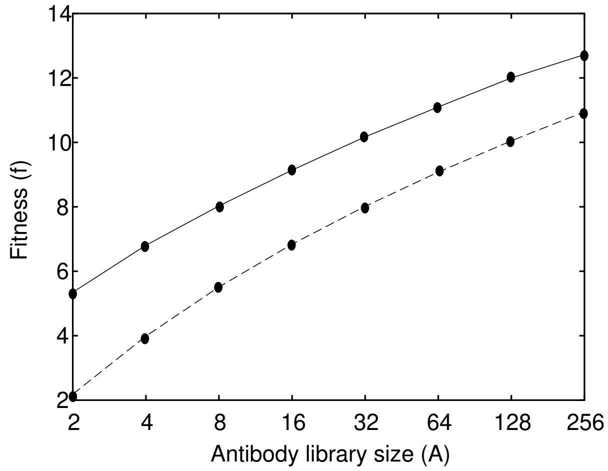

Figure 1: Scaling Of Fitness On A Random Pathogen Set With The Antibody Set Size A. The solid line shows the fitness of the best library evolved in 1000 steps of the genetic algorithm, and the dashed line the expected fitness of a random library. For the solid line, the points on the curve are averages over 100 (for library size $A = 2, 4, 8, 16, 32, 64$) or 10 (for library size $A = 128$ and 256) independent runs, in which we took the best fitness in the population at the end of the run. The line is obtained by interpolating between these points.

Let us analyze the structure of the evolved libraries. Given the fitness function, we would expect that antibodies that have a high free energy in the unbound state would have the highest chance of lowering their free energy through intermolecular binding. Recall that the energy of the free antibodies was a random deviate from a Gaussian distribution. It turns out that the evolved antibodies have higher than average energy. To assess the significance of this difference, we calculate the *z statistic* for the evolved antibodies, that is $z = \frac{x-\mu}{\sigma}$, where x is the energy of an evolved antibody, μ is the mean energy of the antibody molecules, and σ is the standard deviation of the mean. The evolved antibodies have a z-statistic centered around 2 standard deviations higher than the mean, clearly different from the mean. This result tells us that, as expected, the antibodies that were evolved are the equivalent of "sticky" antibodies, of high interconnectivity and multispecificity, such as those commonly seen in the immune systems of newborns (Kearney et al., 1992). These antibodies bind not only to pathogens, but to many other molecules normally present in the body, including DNA and molecules on the surface of lymphoid cells. Thus, the genetic algorithm was indeed able to evolve a property known to characterize the immune systems of newborns.

3.2 SCALING RELATION FOR OTHER DISTRIBUTIONS OF BOND STRENGTHS

Fig. 1 suggests that evolving the antibody libraries allows us to reach higher fitness values than we would have with random libraries, even though the functional form of the dependency between fitness and library size does not change. Let us then explore what this functional form might be for a random library, under assumptions about the fitness of individual antigen-antibody interactions that may have biological relevance.

Let us assume again the random energy model, with all antibody-antigen interactions being characterized by a bond strength distributed according to a density function, g. The cumulative distribution of a single bond strength will be again denoted by G. For example, assume that the bond strength of an antigen-antibody interaction is exponentially distributed, meaning that most interactions are of low energy, higher energy interactions being progressively more rare. Then $G(x) = 1 - e^{-\alpha x}$, with α constant. Correspondingly, $G^{-1}(x) = -\frac{1}{\alpha}\log(1-x)$. Let us denote $y^{\frac{1}{A}}$ by z. Then $y = z^A$, $\frac{dy}{dz} = Az^{A-1}$, and the average fitness over the complete pathogen space will be given by

$$\begin{aligned} f &= -\frac{1}{\alpha}\int_0^1 Az^{A-1}\log(1-z) \\ &= \frac{1}{\alpha}\left(\frac{d}{dz}\log(\Gamma(A+1)) + \gamma\right), \end{aligned}$$

which is approximated by

$$f \approx \frac{1}{\alpha}(\log(A) + \gamma),$$

with γ being Euler's constant, and Γ being the factorial function. Thus, in the case where antigen-antibody bond strengths are exponentially distributed, the fitness of a random antibody library scales logarithmically with the size of the library.

We may also consider a long-tailed distribution, such as a power law $G(x) = 1 - x^{-\alpha}$, with α constant. The inverse of this function is $G^{-1}(x) = (1-x)^{\frac{-1}{\alpha}}$. With the same notation, $z = y^{\frac{1}{A}}$, the average fitness over the complete pathogen space is given by

$$f = \int_0^1 Az^{A-1}(1-z)^{\frac{-1}{\alpha}} = \frac{\Gamma(A+1)\Gamma(1-\frac{1}{\alpha})}{\Gamma(A+1-\frac{1}{\alpha})}.$$

Expanding $\frac{\Gamma(A+1)}{\Gamma(A+1-\frac{1}{\alpha})}$, we obtain for the average fitness

$$f = A^{\frac{1}{\alpha}}\left(1 - \frac{1}{\alpha}\left(1-\frac{1}{\alpha}\right)\frac{1}{2A} + O\left(\frac{1}{A^2}\right)\right).$$

Summarizing, when the bond strengths are exponentially distributed, fitness grows logarithmically with the antibody library size; when the distribution is Gaussian, with faster than exponential tail, the fitness grows more slowly than logarithmically; and for a power law, the fitness is also a power law of the library size. The average fitness, then, as a function of the library size, has a functional form that is the inverse of the density function for the bond strength between an antibody and an antigen. We can use this framework to treat any distribution of antibody-pathogen bond strengths, as more data on this type of molecular interactions becomes available. This is an important feature, as the shape-space based models (and the results that depend on them) have often been criticized for being too restricted, and possibly unrealistic for analyzing biological data.

4 DISCUSSION

It is not yet understood what role the diversity of immune receptor libraries plays in the immune response. Based on the results that we presented here, together with our previous study (Oprea and Forrest, 1998), we argue that adding more and more antibodies to the genome-encoded repertoire improves the survival probability of the individual by smaller and smaller amounts. This may be an explanation for why the V-region libraries in various species do not seem to number more than approximately one hundred genes. But if the selection pressure for increasing library size is small, what would keep evolution from producing even smaller libraries than the ones that we observe? One possible explanation is that there is a hard threshold in antibody/pathogen binding, below which recognition will not occur at all. In this case, some minimal number of antibodies would be required to ensure that at least one has minimal affinity for any given pathogen. Alternatively, one can imagine that the pathogen set is structured as a distribution of clusters, such that different antibodies in the library would reflect different clusters of pathogens. We hypothesize that the antibody genes encode antibodies that are "strategically" placed in the space of possible receptors. The data on what antibody genes are involved in immune responses to virulent pathogens is sparse. In the response to Hemophilus influenzae in humans (Insel et al., 1992), and to Streptococcus pneumoniae in mice (Lee et al., 1974), preferential involvement of a small number of V region genes (and light-heavy chain combinations) has been reported, adding credence to our hypothesis.

Recently, Davis et al. (1998) proposed that the diversity of the repertoire for T cell, as well as for B cell receptors, resides in the third complementarity determining region (known as CDR3) of the immune receptor. In contrast with other complementary determining regions (CDR1 and CDR2), which are exclusively encoded by the V-region gene, CDR3 receives contributions from one or two more gene fragments. These additional gene fragments associate randomly with the V-region gene fragment to form the gene for the antigen-binding part of a functional immune receptor. The authors of the study proposed that CDR3 is sufficient for an initial binding of the immune receptor to the antigen, and that somatic mutation of CDR1 and CDR2 further improves the affinity/specificity of the interaction. In contrast, our hypothesis emphasizes that antibody gene libraries (which code for CDR1 and CDR2) might be the basis for evolutionary learning about the pathogenic environment of the species.

Finally, large phage antibody are now used as a vehicle for rapidly producing high affinity antibodies to protein antigens. Their tentative use ranges from cancer therapy to studying the function of gene products identified by genome projects (Griffiths et al., 1994; Hoogenboom, 1997). Our results are relevant to this work, because they suggest what library sizes we can expect to construct before reaching a certain affinity range for a random antigen. In particular, if we know the distribution of affinities of the antibodies in the library to a random antigen, we can predict what library sizes we need to reach in order for the best antibody in the library to be within a certain affinity range. We suggest that preliminary affinity measurements on a subset of such a library, in conjunction with our analysis, could be a useful test in evaluating methods for generating the phage libraries. Conversely, we could use the data obtained from these large antibody libraries to gain insight into the energy landscape of antigen-antibody interactions.

In conclusion, it is only recently that biological data have become available which allow us to scrutinize the various theories that have been proposed for different mechanisms in the immune system. The role of germline diversity, discussed in this paper, is an example of such a theory. Understanding its role in the overall immune response has consequences, both for theoretical immunology and biotechnology and medicine. However, the availability of detailed biological data means that we need to refine many of the details of our models. Seemingly small details, such as how we model the interactions between antigens and antibodies, can have large impact on the validity of our results. Just as the advent of modern genetic technologies, such as knock-out techniques, stimulated a review of classical embryology, so must we continually revisit the theoretical underpinnings of our models as new biological data become available.

Acknowledgments

The authors gratefully acknowledge support from the National Science Foundation (grant IRI-9711199), the Office of Naval Research (grants N00014-95-1-0364 and N00014-99-1-0417), and the Santa Fe Institute. We also thank Alan Perelson, Tom Kepler, and Erik van Nimwegen for many stimulating discussions leading up to this study.

References

Davis, M., J. Boniface, Z. Reich, D. Lyons, J. Hampl, B. Arden, and Y. Chien, 1998. Ligand recognition by $\alpha\beta$ T cell receptors. *Annu. Rev. Immunol.* **16**: 523–544.

Derrida, B., 1984. Random-energy model: An exactly solvable model of disordered systems. *Phys. Rev. B* **24**: 2613–2626.

Gilfillan, S., A. Dierich, M. Lemeur, C. Benoist, and D. Mathis, 1993. Mice lacking TdT: mature animals with an immature lymphocyte repertoire. *Science* **261**: 1175–1178.

Griffiths, A., S. Williams, O. Hartley, I. Tomlinson, P. Waterhouse, W. Crosby, R. Kontermann, P. Jones, N. Low, and T. Allison, 1994. Isolation of high affinity human antibodies directly from large synthetic repertoires. *EMBO J.* **13**: 3245–3260.

Hightower, R., 1996. *Computational aspect of antibody gene families*. Ph.D. thesis, University of New Mexico.

Hoogenboom, H., 1997. Designing and optimizing library selection strategies for generating high−affinity antibodies. *Trends Biotechnol.* **15**: 62–70.

Insel, R., E. Adderson, and W. Carroll, 1992. The repertoire of human antibody to the *Haemophilus influenzae* type b capsular polysaccharide. *Int. Rev. Immunol.* **9**: 25–42.

Kearney, J., J. Bartels, A. Hamilton, A. Lehuen, N. Solvason, and M. Vakil, 1992. Development and function of the early B cell repertoire. *Intern. Rev. Immunol.* **8**: 247–257.

Lee, W., H. Cosenza, and H. Köhler, 1974. Clonal restriction of the immune response to phosphorylcholine. *Nature* **247**: 55–61.

Oprea, M. and S. Forrest, 1998. *Simulated evolution of antibody libraries under pathogen selection*. Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics, San Diego, CA.

Perelson, A. and G. Oster, 1979. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J. Theor. Biol.* **81**: 645–670.

Press, W., B. Flannery, S. Teukolsky, and W. Vetterling, 1988. *Numerical recipes in C*. Cambridge University Press, Cambridge.

Tonegawa, S., N. Hozumi, G. Matthyssens, and R. Schuller, 1975. Somatic changes in the content and the context of immunoglobulin genes. *Cold Spring Harbor Symposium on Quantitative Biology* **41**: 877–888.

Weigert, M., I. Cesari, S. Yonkovitch, and M. Cohn, 1970. Variability in the light chain sequences of mouse antibody. *Nature* **228**: 1045–1047.