Rational design of vaccine targets and strategies for HIV: a crossroad of statistical physics, biology, and medicine

IOP Publishing Reports on Progress in Physics
Rep. Prog. Phys. **80** (2017) 032601 (19pp) doi:10.1088/1361-6633/aa574a

Key Issues Review

Rational design of vaccine targets and strategies for HIV: a crossroad of statistical physics, biology, and medicine

Arup K Chakraborty[1,2,3,4,5,6] **and John P Barton**[2,5,6]

[1] Departments of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America
[2] Departments of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America
[3] Departments of Chemistry, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America
[4] Departments of Biological Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America
[5] Institute for Medical Engineering & Science, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America
[6] Ragon Institute of MIT, MGH, & Harvard, Cambridge, MA 02139, United States of America

E-mail: arupc@mit.edu

Received 6 January 2015, revised 20 November 2016
Accepted for publication 6 January 2017
Published 13 February 2017

Corresponding Editor Professor Robert H Austin

Abstract

Vaccination has saved more lives than any other medical procedure. Pathogens have now evolved that have not succumbed to vaccination using the empirical paradigms pioneered by Pasteur and Jenner. Vaccine design strategies that are based on a mechanistic understanding of the pertinent immunology and virology are required to confront and eliminate these scourges. In this perspective, we describe just a few examples of work aimed to achieve this goal by bringing together approaches from statistical physics with biology and clinical research.

Keywords: immunology, statistical physics, virology

(Some figures may appear in colour only in the online journal)

Introduction

No medical procedure has saved more lives or money than vaccination. Successful vaccination programs have resulted in the eradication of smallpox, which had caused devastation since antiquity, and the near-eradication of polio. Vaccines for numerous childhood diseases are also a major contributor to the reduction of infant mortality. Vaccination manipulates our immune system in a way that enables it to act rapidly upon infection to eliminate a specific pathogen, thereby preventing disease.

The immune system of higher organisms [1] can be roughly partitioned in to two linked parts. The first is the innate immune system, which is made up of many types of cells and molecules that can bind to ('recognize') molecules which are often displayed on the surface of many pathogens, but not on the surface of host cells. Successful recognition can result in elimination of the pathogen. The innate immune system is very effective as evidenced by the fact that most of the time we are not sick despite constant exposure to infectious pathogens. However, innate immune responses are not pathogen specific. Many viruses and bacteria have evolved

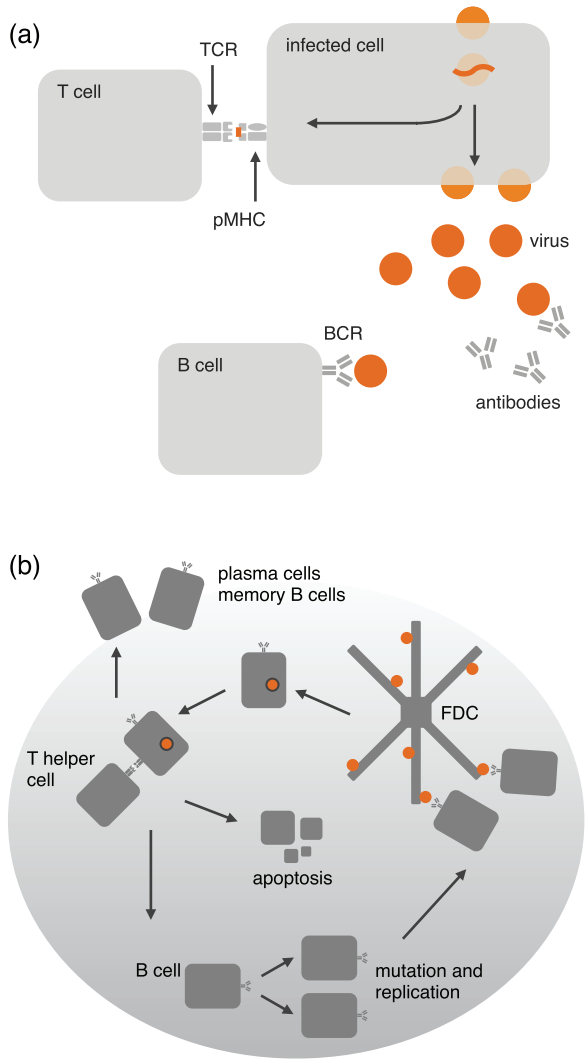

Figure 1. (a) Schematic depiction of the adaptive immune response to infection. Certain kinds of T cells can kill infected cells by recognizing virus-derived peptides in complex with MHC molecules on the surface of the infected cell. B cells generate antibodies, soluble forms of its B cell receptor, which can primarily neutralize viruses in extracellular spaces. (b) Schematic of the affinity maturation cycle. B cells proliferate and develop mutated versions of their receptors in the dark zone. In the light zone they encounter FDCs which present antigens. B cells that are successfully able to bind to and engulf antigens receive survival signals from T helper cells that allow them to continue to proliferate. A fraction of successful B cells exit the germinal center as plasma cells or memory B cells, while others are recycled for further rounds of mutation and selection. B cells that are not positively selected undergo apoptosis.

strategies to evade innate immunity, and when they do, an infection is established. The adaptive immune system gets involved when this happens, and it enables us to mount pathogen-specific responses against a diverse and evolving world of microbes. This is remarkable because it is not specificity to a pre-determined list of pathogens; in fact, pathogen-specific responses can be mounted against microbes which had not evolved when a person was born.

To roughly see how this works, consider infection with a virus (figure 1). A virus hijacks the transcriptional machinery of the host cell enabling the synthesis of viral proteins and the assembly of new virus particles that can infect other cells. B lymphocytes (B cells) are important players in adaptive immunity. They display a receptor on their surface called the B cell receptor (BCR). The gene that codes for the BCR is inherited as various flavors of different gene segments. During synthesis of a B cell, a flavor of each gene segment is chosen stochastically (subject to some constraints), and the segments are joined to form the gene encoding this B cell's receptor [2]. Thus, the BCR on a particular B cell is likely to be distinct from the receptor on most other B cells. If the BCR on a particular B cell can bind strongly to the surface proteins of a particular virus, then biochemical reactions ensue in the B cell that can result in its activation. Activated B cells then undergo a process called affinity maturation [3] (see below), which ultimately results in the secretion of a soluble form of a mutated BCR (antibody) that has an even higher affinity for the viral surface proteins. Antibodies can bind to the virus and dispose of it in numerous ways which require the action of the components of the innate immune system.

B cells primarily act on virus particles that are in blood or extracellular spaces. T lymphocytes (T cells), the other important arm of adaptive immunity, are largely responsible for the control of intracellular pathogens. Some of the viral proteins in infected cells are chopped up into short peptide fragments by molecular machinery in human cells. These peptides (p) are displayed on the surface of infected cells in complex with a human protein called major histocompatibility complex (MHC; in humans this protein is also referred to as human leukocyte antigen or HLA). We each have six to twelve types of MHC molecules, but there are thousands of variants in the human population [4]. Different types of MHCs can bind different peptides. For reasons similar to that described for B cells, most T cells express a distinct surface receptor, called the T cell receptor (TCR). If the TCR on a particular T cell can bind strongly to a particular virus-derived p-MHC molecule displayed on an infected cell, biochemical reactions in the T cell ensue which result in activation and proliferation [1, 5]. Activated T cells can then coordinate an immune response in a number of different ways. For example, certain kinds of T cells, called cytotoxic T lymphocytes (CTL), are such that, if an activated CTL sees another cell displaying the same viral peptide that originally activated it, products are secreted that kill the infected cell and the virus particles it harbors. Another kind of T cells, called T helper cells, play a key role in antibody production (see below) and the secretion of chemicals that are important for mounting immune responses.

As was mentioned above, upon activation due to interactions with virus particles (which takes place in lymph nodes), B cells undergo affinity maturation. Activated B cells can nucleate structures within lymph nodes called germinal centers (GCs) [6], where a Darwinian evolutionary process ensues in a short period of time. The activated B cells proliferate and mutations are introduced in to the binding site of the BCR at a high rate through a process known as somatic hypermutation. B cells with mutated receptors then migrate to another part of the GC, where cells of the innate immune system (follicular dendritic cells) display the virus. B cells compete to

bind to the spikes on the virus. If the binding strength exceeds a threshold, B cells can internalize the virus. B cells with a BCR that binds more strongly to the virus' spike have a higher probability of binding and internalizing more virus particles. The B cells that successfully internalize virus particles then display virus-derived p-MHC molecules on their surface. B cells compete with each other to bind to the TCR on T helper cells, which are present at limited numbers in GCs [7]. This interaction can result in a survival signal (intracellular biochemical reactions) for the B cell, allowing it to continue the affinity maturation process. B cells that internalize more virus particles have a higher probability of receiving this survival signal because they display more p-MHC molecules to which TCRs can bind. B cells that do not internalize virus particles or fail to receive a T helper cell-mediated survival signal die. Some of the positively selected B cells emigrate from GCs in to the blood and tissues as plasma cells that secrete the BCR in soluble form (i.e. antibodies) and as memory cells (see below). Most of the surviving B cells are recycled for further rounds of mutation and selection [7–9]. Thus, as time ensues, more potent antibodies with higher affinity for the infecting virus are produced.

After an infection is cleared, most of the T cells and B cells that proliferated in response to this particular virus die. But, a few remain as so called memory cells that can mount robust and rapid responses upon reinfection with the same virus. This pathogen-specific immunological memory is the basis for vaccination. A vaccine aims to induce memory T cells, B cells, and antibodies which are specific for the pathogen against which protection is desired.

Perhaps as early as the tenth century, the Chinese tried to protect people from smallpox by rubbing the fluids in scabs from diseased people onto the skin of healthy people, resulting in the protection of some and death of others [10]. Modern vaccination roughly follows the paradigm pioneered by Jenner and Pasteur in the nineteenth century. A dead or weakened form of the pathogen is injected into humans with the goal of inducing effective memory immune responses. It is interesting that this empirical protocol was developed before we knew much about the immune system, and Jenner did not know about microbial pathogens. In the twentieth century, additives (adjuvants) were added to vaccines to help stimulate innate immune responses that are critical for the development of potent adaptive immune responses. The design of adjuvants remains largely an art and many new formulations fail to work.

In recent years, some pathogens have evolved which have defied successful vaccination using the traditional empirical paradigm. Prominent examples are HIV, HCV, tuberculosis, and malaria, many of which are wreaking havoc around the world. We do not have a broadly effective vaccine against influenza either, and attempts to predict the right vaccine target for the ensuing year, and thus the vaccine, often fail. Many of these pathogens share two features: (1) They present themselves in different guises, thus making them hard to target with specific immune memory responses. (2) They often degrade, or hide from, the immune system.

HIV has characteristics which are extreme examples of both these features. It is a highly mutable virus with a very rapid

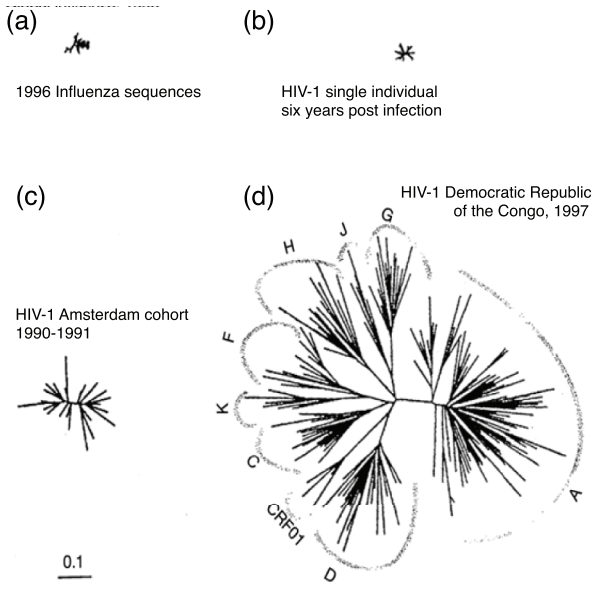

Figure 2. Comparison of phylogenetic trees generated from collections of HIV and influenza sequences. The global diversity of influenza sequences in a given year (a) is comparable to the diversity of HIV sequences within a single chronically infected individual (b), and completely dwarfed by the observed diversity of HIV sequences in a single region in Africa (d). Figure adapted with permission from [11]. Copyright 2001 Oxford University Press.

replication rate. Thus, it generates many mutant strains when it infects even a single person. For example, the diversity of HIV strains in a single infected person is not that different from the diversity of circulating influenza strains in the entire world in a particular year, and the latter is dwarfed by the diversity of circulating strains in a single region in Africa during the same time period [11] (figure 2). The high mutability allows HIV to evade natural or vaccine-induced immune responses [12]. Furthermore, HIV principally infects and eventually kills human T helper cells, thus degrading the adaptive immune system. This is the reason why acquired immunodeficiency syndrome (AIDS), the disease associated with HIV infections, results in a severe state of immunodeficiency allowing many normally easy to control infections to afflict patients. Other pathogens listed above also mutate (e.g. influenza and tuberculosis) and malaria uses a different strategy in that it expresses different interchangeable proteins on the surface. The bacterium that causes TB hides from the immune system in red blood cells. Vaccine design against HIV is particularly daunting because, unlike other pathogens for which vaccines exist, HIV infection is not known to have ever been successfully cleared by natural human immune responses.

Successful vaccination against pathogens that have evolved sophisticated strategies to evade human immune responses will benefit from the development of firm scientific principles that can guide rational vaccine design [10]. At least two important issues must be studied in this regard: (1) What are the appropriate targets for vaccine-induced immune responses that will limit or eliminate the ability of these pathogens to evade such responses while simultaneously maintaining their viability/virulence? (2) How can such immune responses be induced by vaccination in humans with diverse genotypes?

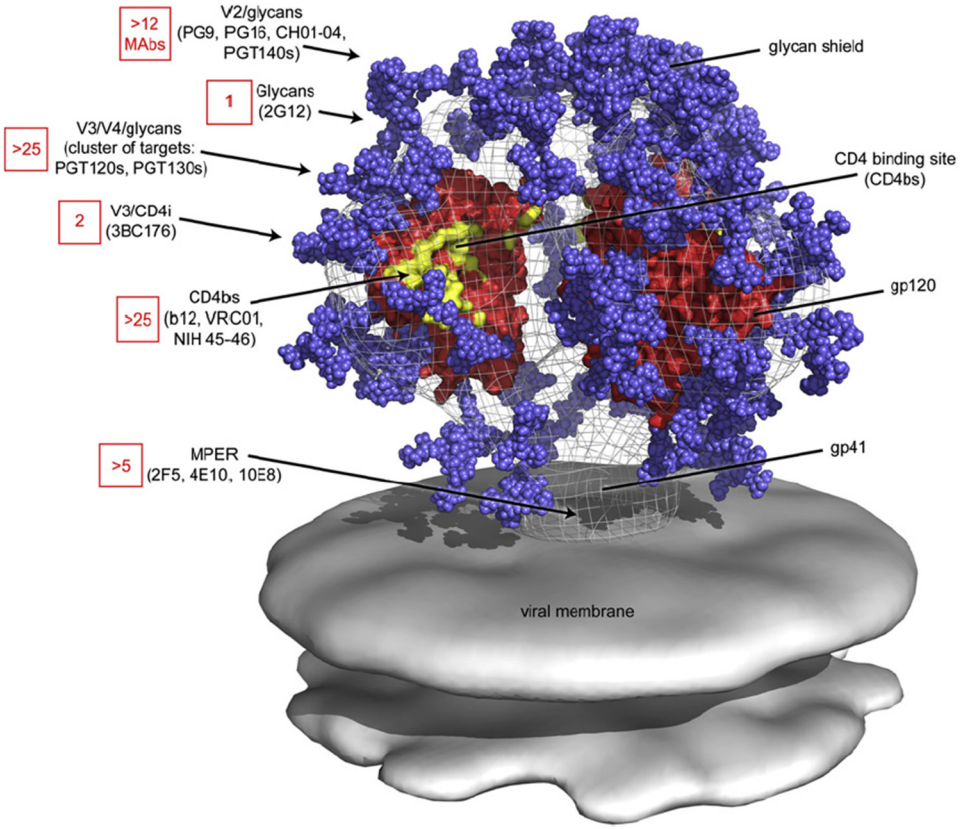

Figure 3. Representation of the HIV envelope protein in its native shape, as a trimer of gp120 and gp41 subunit proteins. Glycans surround the CD4 binding site region, hindering its access by antibodies. Binding sites of common broadly neutralizing antibodies are indicated by arrows. Figure adapted with permission from [111]. Copyright 2012 Elsevier.

A convergence of several factors is beginning to enable us to address these questions. Biologists and clinicians can collect enormous amounts of data on sequences of mutant strains of pathogens, and more recently, it is also becoming possible to interrogate the immune system on an unprecedented scale. Both the immune system and pathogens function by collective processes that involve myriad individual components, thus making mechanistic interpretation of this data complex. Physicists, especially statistical physicists, have begun to play a role in translating this type of data to mechanistic knowledge that addresses the questions noted above. The goal of developing mechanistic principles that can be harnessed for rational design of vaccines is bringing physicists, biologists, clinicians, and engineers to work together.

In this perspective piece, we will focus primarily on our efforts at the intersection of physics and biology that aim to address the first question noted above—i.e. what are optimal vaccine targets? But, we will also discuss a few issues pertinent to the second question. For each topic that is discussed in some detail, we will briefly review the current state of the science and note the many open questions. To make our discussions concrete, we will consider primarily one virus, HIV. However, to contrast the biology and challenges posed by different viruses, we will sometimes refer to the challenges posed by influenza, for which some significant progress has been made in recent years. We emphasize that this is a perspective piece, not a comprehensive review.

Brief description of the biology of HIV and definition of the key challenges

HIV is thought to have been transmitted to humans from monkeys, and is estimated to have been circulating in small populations of humans for nearly a century [13]. The first well-documented cases were reported in 1981. To date, HIV has infected over 78 million people and almost 40 million people have died from complications associated with AIDS. In developed nations, HIV infections can be controlled by daily expensive medication, but it cannot be cured. In other parts of the world, HIV continues to wreak havoc, with sub-Saharan Africa being the epicenter of the disease. For example, each day there are 1000 new HIV infections in South Africa alone. A vaccine or cure is urgently needed, but no successes have been reported over more than thirty years since it became known that the causative agent of AIDS was this virus.

HIV is a retrovirus, which carries its genome in the form of RNA. The virus has a lipid membrane through which proteins protrude. These Envelope proteins are gp120 and gp41, which form a non-covalently bonded trimer which constitutes the viral spike (figure 3). The outer membrane surrounds a capsid made up of structural proteins which encloses the RNA genome and other key proteins important for viral function. The trimeric spike binds to host cell surface proteins to initiate infection [14]. For example, the spike can bind to a receptor called CD4, which is primarily expressed on the surface of T helper cells [15, 16]. This binding event leads to

a conformational change in gp120, resulting in the dissociation of gp41 which then forms a six-helix bundle. This also allows gp120 to bind to a second receptor on the surface of the host cell, which can be either CCR5 or CXCR3. These binding events and the free energy gained from 'crystallization' of gp41 result in fusion of the virus' membrane with that of the host cell membrane, resulting in release of the capsid in to the cytoplasm. The viral capsid is then uncoated, releasing its contents. A viral protein, Reverse Transcriptase, then converts the RNA strands in to DNA. This viral DNA, also called a provirus, is transported in to the nucleus of the host cell along with a viral protein called integrase, which inserts the viral DNA in to the genome of the host cell, thus infecting this cell for life. The transcriptional machinery of the host cell can then help synthesize viral proteins.

HIV has nine genes, with four of them, Gag, Pol, Nef, and ENV being the most important. Gag codes for the proteins that form the structure of the virus (such as the ones that make up the capsid), Pol codes for Reverse Transcriptase, Integrase, and Protease, ENV codes for gp120 and gp41, and Nef plays a role in downregulation of HLA proteins and CD4 (downregulation of which helps the virus bud out of infected cells). Gag, Pol, and ENV genes code for polyproteins, and Protease cuts them in to the right individual proteins. In a series of steps, the virus' proteins are properly assembled at the membrane of the host cell [17]. A part of the host cell membrane becomes the membrane for a new virus particle as it buds out.

Reverse transcription is not a very high fidelity process, and for HIV, mutations are introduced at the rate of 3×10^{-5} per base pair per replication cycle [18]. HIV's genome is about 10^4 in length, and so this implies that during every replication cycle there is a 0.3 probability that a mutant will result. Moreover, Reverse Transcriptase can hop from one RNA molecule to the other, thus generating more options for creating diversity, and when two different RNA genomes are available, recombination occurs. Fitting parameters in ordinary differential equations to data from patients treated with drugs revealed that HIV replicates very rapidly [19], producing 10^{10} to 10^{11} virus particles per day in patients [20]. Of course, many of the mutant strains that are produced do not grow as they cannot form infective virus particles. But, taken together, the high mutation and replication rates are the main reasons underlying the extraordinary diversity of HIV strains circulating in the population and in individuals.

Upon successful infection, the virus replicates rapidly, and the viral load (that is, the number of viruses circulating in the host) increases. As immune responses develop, the viral load decreases and then stabilizes at a steady state where the immune system and the virus are in balance [21]. The steady state viral load varies widely between patients [22]. Because HIV kills the host cell when new virus particles bud out, the number of CD4 T cells in an infected patient declines at first and then increases again after the immune response develops. The period when the viral load and CD4 T cell counts are stable is called the asymptomatic phase as no disease manifests. HIV belongs to a class of viruses called lentiviruses (slow viruses) which cause disease slowly. Without treatment the immune system ultimately loses the battle, viral load goes up, and CD4 T cells decrease in numbers. At this point, the individual's immune system is severely compromised, and many opportunistic infections ensue leading ultimately to death.

The high mutability of HIV is a major reason why an effective vaccine does not exist [23]. An enormous diversity of possible mutant strains could infect a person, and the virus could mutate within a person after infection to evade specific vaccine-induced (or natural) immune responses. So, the challenge is to understand how to effectively attack diverse HIV strains through immune responses that can be induced in people with different genotypes. Although some immune responses are known to be more efficacious either in patients or *in vitro* [24–27], the HIV population has been met overwhelmingly with immune responses that are ineffective in controlling infection. In contrast, Influenza has been historically subjected to specific classes of effective natural and vaccine-induced immune responses. The challenge for vaccination against influenza is to determine which virulent strains that have not been previously subjected to effective immune responses (and for which no immune memory exists in the population) are likely to evolve and infect people in the ensuing year [28]. The implications of this difference in evolutionary history for vaccine development will be noted in the sections below.

In this perspective, we will focus on only two aspects of the challenge of creating an effective HIV vaccine. First, we will ask whether the data on thousands of sequences of HIV proteins derived from virus samples extracted from patients can be translated in to knowledge of the mutational vulnerabilities of the virus; i.e. which types of mutations is the virus unable to make to evade immune responses and still remain viable. This knowledge can be used to design the active component of a potentially effective T cell-based therapeutic vaccine. We will also briefly describe one aspect of how the temporal pattern in which active components of a vaccine are administered may be important for inducing the desired antibody responses for a prophylactic vaccine.

Determining the mutational vulnerabilities of HIV

At first glance, determining the mutational vulnerabilities of a virus like HIV appears to be a simple problem given that today we can sequence large numbers of viruses derived from patients. Just lining up all these sequences and looking for residues that appear relatively conserved (e.g. measured by low entropy) should reveal the answer. Focusing a vaccine-induced immune response to target such residues should be an effective strategy to control the virus. This is because to evade such an immune response, HIV would have to evolve a mutation at one of these residues, and that should hurt viral function as the residues are relatively conserved for a reason. This strategy is blunted because HIV can evolve other mutations, so-called compensatory mutations, which can partially restore the fitness cost incurred by making the primary immune-evading mutation [29, 30].

Many properties of an evolving virus population can be described by assuming that each residue in the virus' proteins evolves independently. But, if one wishes to determine

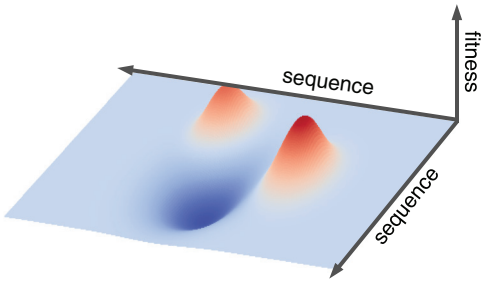

Figure 4. Schematic of a fitness landscape. The true sequence space is high-dimensional and discrete, but here for simplicity, we represent the landscape on a continuous 2D sequence space. Ideally, vaccine-induced immune responses could be generated that force escape mutations corresponding to a 'valley' on the fitness landscape, while blocking off pathways of low fitness cost between adjacent fitness 'peaks'.

the mutational vulnerabilities of a virus like HIV, one needs to define the collective mutational pathways that HIV uses to evade human immune responses in order to avoid targeting the involved residues with a vaccine-induced immune response. This is because, even though such compensatory interactions can be relatively rare, the high mutability and replication rate of HIV implies that they can be sampled, especially when immune responses are selecting for one (or more) of the involved mutations. One also needs to determine the combinations of mutations that the virus cannot make and remain viable, so as to target the involved residues with vaccine-induced immune responses and corner the virus between the immune responses and evolving unfit mutant strains. In short, one needs to determine the fitness landscape [31, 32] of the virus (figure 4). Knowledge of the fitness landscape can be extremely useful for vaccine design. With a vaccine-induced immune response, one wishes to target residues such that combinations of mutations therein correspond to a fitness valley. One also wishes to block the mountain passes corresponding to the collective compensatory mutational pathways that HIV uses to go from one fitness hill to another adjacent one when the first one is under immune attack (figure 4). This practical end motivates the goal of translating sequence data in to knowledge of the fitness landscape of HIV proteins. The methods that we describe below rely on sequence information alone; others have attempted to infer information about HIV fitness by combining sequence information with data from *in vitro* experiments [33, 34], an approach with its own set of potential complications [35].

Simple calculations can reveal the importance of collective effects of multiple mutations

Given an alignment of sequences of any protein, one can compute the covariance matrix that describes correlations between mutations. For illustrative purposes, let us imagine that the amino acids at each residue come in two flavors, wild type or most frequent (denoted by 0) and mutant (denoted by 1). Such a representation is reasonable when the number of types of amino acids observed at each residue is small. In such an Ising representation, if z_i represents the amino acid at residue, i, the covariance matrix, C, is defined as

$$C_{ij} = \frac{\langle z_i z_j \rangle - \langle z_i \rangle \langle z_j \rangle}{\sqrt{V_i V_j}}, \tag{1}$$

where angular brackets represent an average over all sequences and V_i and V_j are the variances of the distribution of mutations at residues i and j, respectively. This covariance matrix can be diagonalized, and the eigenmodes are the simplest reflection of collective mutational pathways. We expect that such eigenmodes can reveal functionally coupled sites. As an example, compensatory mutations can arise in order to restore fitness lost through escape mutations, and so mutations at these combinations of residues occur together more frequently than would be expected by chance.

But, the problem is complicated by at least two effects. First, even though we have many samples of the sequences of HIV proteins, it is a finite sample. Even if two variables are completely uncorrelated, given a finite sample, the covariance matrix will exhibit spurious correlations. Second, some correlations simply reflect the fact that all sequences belong to the same family of HIV viruses, and mutational fidelity leads to correlations between sequences (a phylogenetic effect).

Random matrix theory (RMT) was developed in the context of nuclear physics to attempt to describe statistical properties of the energy eigenspectrum of complex nuclei, which could not be exactly determined. It has since been used profitably in diverse areas of physics, economics [36–38], and more recently, to study families of proteins across species [39], a single HIV polyprotein called Gag [40], and certain proteins of the Hepatitis C Virus [41]. Given a finite set of independent random variables, the covariance matrix has a distribution of eigenvalues that range from [42]

$$\lambda_- \leqslant \lambda \leqslant \lambda_+, \qquad \lambda_\pm = 1 + \frac{1}{Q} \pm 2\sqrt{\frac{1}{Q}}, \tag{2}$$

where λ_+ is the largest eigenvalue and λ_- is the smallest, and $Q = M/N$ is the ratio of the number of samples (number of sequences M) to the number of variables (sequence length N). Dahirel *et al* found that the distribution of eigenmodes of the matrix, C, obtained from aligning sequences of Gag is continuous up to a value slightly larger than λ_+ corresponding to the data (M and N), followed by a discrete spectrum of eigenmodes. One could conclude that only the discrete eigenmodes corresponding to eigenvalues greater than λ_+ contain information on real collective correlations between mutations. But, (2) applies only in the limit of M and N tending to infinity; neither condition is true for analyses of protein sequences. Finite size effects can be studied using a method suggested by Leibler, Ranganathan and co-workers [39]. Take each residue in the sequence alignment and randomly permute the amino acids at that residue across all sequences. This results in one realization of a random sequence alignment that preserves the average amino acid composition at each residue. Diagonalizing the covariance matrix, C, corresponding to this randomized sequence alignment, and repeating this operation many times produces a histogram representing the distribution

of eigenvalues for randomized sequences with the same values of M and N as the real sequence alignment. The result for Gag is that finite size effects result in a tail of the eigenvalue distribution that extends a bit beyond the λ_+ theoretical limit of 2.42 for the Gag sequence alignment; there were no eigenvalues beyond 3 for the distribution obtained from the randomized sequences. Therefore, one could conclude that the eigenmodes corresponding to eigenvalues larger than 3 reflect collective correlations that are not spurious.

For bacterial proteins [39], HIV Gag [40], a key protein in HCV [41], and analyses of financial markets [38], one finds that there is a very large eigenvalue which corresponds to a coherent eigenmode. For the biological examples, this mode has been interpreted to be one that largely reflects phylogeny. Analogously, for analyses of correlations between stock price fluctuations, this eigenmode has been interpreted to be one that reflects overall market forces that affect the price of all stocks.

The matrix, C, can be decomposed in terms of its eigenvalues as

$$C = \sum_k \lambda_k |k\rangle\langle k| \,, \tag{3}$$

where $|k\rangle$ represents the kth eigenvector. The eigenvalues below λ_- usually have a negligible contribution, although they have been recently interpreted to reflect direct contacts in protein structures [43]. The eigenvalues between λ_- and λ_+ reflect noise, and so the matrix, C, 'cleaned' of noise can be decomposed as

$$C_{\text{clean}} = \sum_{k|\lambda_k > \lambda_+} \lambda_k |k\rangle\langle k| \,. \tag{4}$$

In order to remove the phylogenetic contributions, one could simply exclude the top eigenvalue from the sum in (4). A more principled way has been suggested by Plerou *et al* [38]. For Gag, either approach leads to the same qualitative results [40].

With the relevant eigenmodes in hand, one can construct a description of the collective modes as determined by this simple analysis. The crudest way to do this is by graphing the values of the coefficients corresponding to each residue for pairs of eigenmodes, and then performing a visual clustering analysis. More sophisticated approaches are available [41]. Carrying out the simpler analysis, Dahirel *et al* identified five co-evolving groups of residues in Gag. Two of these groups were contained in the p24 protein. When these groups of residues were superimposed on the structure of p24 alone, no reason for co-evolution was apparent. However, six p24 proteins form hexamers through non-covalent interactions, and the hexamers form interfaces with other hexamers [44]. This honeycomb-like structure tiles the viral capsid. One group of co-evolving residues is largely involved in forming intra-hexamer and inter-hexamer interfaces between p24 proteins. Another is largely made up of residues that form the core of the hexamer. Given the structural basis of these groups of residues, one can ask whether multiple mutations at these co-evolving residues are likely to be deleterious or compensatory by examining the proportion of negative to positive correlations between pairs of residues within a group as well as the larger scale correlations (e.g. embodied as positive and negative lobes on certain eigenvector projections). Based on these considerations, Dahirel *et al* concluded that these two groups were more likely to be vulnerable to multiple simultaneous mutations than others. Structurally, this could be rationalized by surmising that multiple mutations at co-evolving residues that formed key protein-protein interfaces critical to capsid assembly may cause the capsid to become unstable. Thus, simultaneously targeting such residues with a T cell response might be effective as immune evasive mutations would be likely to be deleterious.

Clinical data seems to support this surmise. Some patients, called elite controllers, do not progress to AIDS because their immune systems can maintain low levels of HIV without any therapy [45]. The T cell responses in these patients seem to disproportionately target the two groups of residues deemed to be vulnerable to multiple simultaneous mutations by the analysis noted above. The analyses outlined above also predicted that, while both groups of co-evolving residues are equally resistant to single mutations, one group was more vulnerable to multiple simultaneous mutations than the other. Sequencing of virus samples extracted from elite controllers seem to support this prediction as the frequency of observed single mutations in residues in these groups are roughly the same, but multiple mutations are more restricted in one group of residues as predicted [40].

The type of analysis outlined above, which resembles principal component analysis with the eigenspectrum cleaned of noise, can be helpful in identifying some collective correlations. However, it does not provide a quantitative metric that differentiates relative fitness costs incurred by the virus by making one set of combinations of mutations versus another. The absence of a metric of relative fitness also does not allow for the calculation of evolutionary dynamics in response to different immune pressures to establish which types of immune responses will be able to keep the virus cornered for the longest times before escape can occur. These are the types of immune responses that one would wish to induce by vaccination. The inability to make predictions like the ones noted above also makes it difficult to test predictions against a wide range of *in vitro* and clinical data.

Another important point is that the virus samples used to carry out the analysis are derived from patients in whom a host-pathogen battle had ensued forcing certain mutations that evade the immune response. The analysis noted above does not allow for a principled way to analyze and deconvolve the effects of the immune response on the correlations between mutations.

Inference of prevalence and fitness landscapes

One can begin by seeking to construct a model for the prevalence landscape, or the probability $P(z)$ of observing a sequence, z, of an HIV protein in the circulating virus population [46]. The sequence data contains information on the probability of observing single mutations at every residue of a protein, double mutations at every pair of residues, triple mutations at every triplet of residues, etc. Any mathematical model for $P(z)$ that can recapitulate these mutational

correlation functions can be said to accurately describe the prevalence landscape. One way to approach this inference problem is to ask: what is the least biased model for $P(z)$ that recovers the one- and two-point mutational correlations observed in the available sequence data. Exploiting the connection between statistical mechanics and information theory, one may interpret 'least biased' to mean the probability distribution $P(z)$ which has the maximum entropy subject to these constraints on the correlations [47]. A similar approach has been used to infer contacts in protein structures [48, 49], correlations between the firing of neurons [50], etc. Related methods have also been employed to study structural properties of HIV protease [51] and inter-protein interactions [52].

We seek to maximize the entropy of $P(z)$, subject to the constraints that the probability distribution is normalized and that the one- and two-point mutational correlations predicted by the model are those that match the observed correlations. For most HIV proteins, we use a Potts model to represent the amino acids at each residue (see details later). For ease of illustration, we use an Ising representation to simplify the formulas noted below. In this Ising representation, the residue at site i in a given sequence corresponds to $z_i = 0$ if it is the wild type amino acid and $z_i = 1$ otherwise. With this simplified notation, the quantity we wish to maximize is

$$-\sum_z P(z)\log P(z) + \alpha\left(\sum_z P(z) - 1\right) + \sum_i h_i\left(p_i - \sum_z z_i P(z)\right) + \sum_{i<j} J_{ij}\left(p_{ij} - \sum_z z_i z_j P(z)\right), \tag{5}$$

where the constraints are enforced through the Lagrange multipliers α, h_i, and J_{ij}. Here p_i and p_{ij} represent the observed one- and two-point mutational correlations, respectively. Maximizing this functional yields

$$P(z) = \frac{\mathrm{e}^{-H(z)}}{Z}, \qquad H(z) = -\sum_i h_i z_i - \sum_{i<j} J_{ij} z_i z_j, \tag{6}$$

where the fields h_i and couplings J_{ij} are those that constrain the one- and two-point correlation functions to be the observed ones. The partition function Z ensures that the probability distribution is properly normalized. The couplings J_{ij} can have both positive and negative signs and so there are some analogies between the form of the Hamiltonian in (6) and the Hopfield model of neural networks [53] which can be profitably exploited [54].

Another way to see that the fields and couplings should be those that fit the observed one- and two-point correlation functions is to note that, once we choose the form of the Hamiltonian to be as in (6), the average log-likelihood of the observed sequence data is

$$\ell = \frac{1}{B}\log\left(\prod_{k=1}^{B} P(z^{(k)})\right) = -\log Z + \frac{1}{B}\sum_{k=1}^{B}\left(\sum_i h_i z_i^{(k)} + \sum_{i<j} J_{ij} z_i^{(k)} z_j^{(k)}\right), \tag{7}$$

where B is the number of sequences in the data, and $z^{(k)}$ represents the kth sequence. Maximizing this function with respect to the parameters (the fields and the couplings) yields

$$\frac{\partial \ell}{\partial h_i} = p_i - \sum_z z_i \frac{\mathrm{e}^{-H(z)}}{Z} = 0\,, \qquad \frac{\partial \ell}{\partial J_{ij}} = p_{ij} - \sum_z z_i z_j \frac{\mathrm{e}^{-H(z)}}{Z} = 0, \tag{8}$$

noting here that $p_i = \Sigma_k\, z_i^{(k)}/B$, $p_{ij} = \Sigma_k\, z_i^{(k)}\, z_j^{(k)}/B$. That is, the parameters that maximize the likelihood of the data are also ones that fit the observed one- and two-point mutational correlation functions.

Inferring the fields and couplings that maximize the likelihood is referred to as the inverse Ising (or Potts) problem. This problem is challenging to approach directly through traditional optimization methods because of the difficulty of evaluating the likelihood function. The likelihood depends on the partition function, which contains a number of terms that grow exponentially with the system size N. It cannot be computed directly except for very small systems. As a result, a number of inference schemes have been developed to provide approximate solutions to the inverse Ising problem.

The simplest approach to the inverse Ising problem is based on iterative Monte Carlo simulations [55]. Given a starting value for the model parameters, one can estimate the gradient of the likelihood (8) by computing the correlations from Monte Carlo simulation and comparing them to the ones from the data. Iteratively updating the model parameters based on the difference between the true and simulated correlations yields models that do progressively better at reproducing the data. The drawback of this approach is that it tends to be slow, as Monte Carlo simulation can be particularly costly for large systems.

Alternative methods replace the intractable likelihood (7) with one that is easier to evaluate. One such example is the Gaussian approximation: note that, if the spin variables z_i were real-valued rather than binary, the coupling values J_{ij} would be given simply by the entries of the inverse of the sample covariance matrix. Similar expressions can also be obtained from a systematic expansion assuming weak couplings [56, 57], and this approach has been fruitfully applied to study the structural properties of protein families [58]. Pseudolikelihood methods replace the full expression in (7) with an array of likelihood expressions, one for each variable, which are maximized in parallel. Here other spin variables are treated as quenched random variables, simplifying the inference problem [59, 60]. While these sorts of approaches are much faster than iterative Monte Carlo learning, they do not always yield a model that accurately reproduces the data [61, 62].

In recent work, the adaptive cluster expansion (ACE) method for inverse Ising inference [62–64] has been applied to infer the HIV prevalence landscape. In this approach, one maximizes the likelihood exactly, but only on small subsets (or clusters) of spins at a time. An approximation for the set of couplings and fields for the entire system can then be built up from the parameter values inferred on these smaller clusters.

The success of this approach relies on the ability to accurately estimate the model parameters using only a subset of the correlation data (i.e. that the inverse covariance matrix $(\mathrm{d}\mathbf{p}/\mathrm{d}\mathbf{J})^{-1}$ is sparse). ACE is typically slower than approximate methods such as the Gaussian and pseudolikelihood approximations, but it consistently yields models that accurately reproduce the original data [62]. Efforts to improve the speed and accuracy of this and other methods of inverse Ising inference remain an active topic of current research.

Connection between prevalence and fitness of HIV strains

The model that we infer by following the procedure above describes the prevalence of circulating strains of HIV. One can argue that the more prevalent strains are also the ones that are intrinsically more fit, but this can be proven to be true only in some circumstances [65–67]. In the case of a real human pathogen, such as HIV, the sequences of strains that are used to infer the model are derived from patients whose immune systems have battled the virus. Thus, the virus sequences in individual patients are likely to bear the imprint of the host-pathogen riposte; mutations that make the virus intrinsically less fit may allow the virus to evade the immune response, thus making a strain bearing these mutations effectively more fit and prevalent in a particular host. How does the host-pathogen riposte affect the relationship between the prevalence and intrinsic viral fitness? To explore this question, a number of efforts have been undertaken recently for both HIV and influenza. We will focus here on HIV, but will note the contrast with influenza.

Shekhar *et al* carried out computer simulations with a coarse-grained model that aimed to mimic the way in which the virus samples were collected [68]. An in silico person is infected with N copies of a virus, and a host-pathogen battle ensues (see below for details). At a randomly chosen time, one of the viruses from this in silico person is transmitted to another host, and the infection begins anew with N copies of the transmitted virus strain. HIV proteins are subjected to T cell responses which target peptides bound to MHC molecules. Because the newly infected person is likely to have a different MHC genotype and antibody responses also vary across hosts, in our calculations the virus evolves in response to a different host immune response. At a randomly chosen time, a virus sample from this person infects another new host, and this process is continued. At separate randomly chosen times, a virus sample from the current population within each host is recorded, and these samples can be thought to represent a sequence database.

Within each person, the fitness of the virus is modulated by the immune response. In our representation, we assume that there is a Hamiltonian that reflects the intrinsic fitness of the virus, and that the immune response, which acts on particular residues of viral proteins, can be modeled as external fields that act at a few points to promote mutations. Again, for simplicity, in an Ising representation, the effective Hamiltonian (or fitness) can be written as

$$H_{\mathrm{eff}}(z) = -\sum_{i=1}^{N} h_i z_i - \sum_{i=1}^{N-1}\sum_{j=i+1}^{N} J_{ij} z_i z_j + \sum_{i=1}^{k} b_{s_i} z_{s_i}\,. \qquad (9)$$

The fields, b_i, act on residues in a peptide that is being presented by an individual with particular MHC molecules. So, there are short-range correlations between the locations at which these fields act. For simplicity, Shekhar *et al* ignored these correlations. Here we use s_i to denote the (set of) targeted sites and to emphasize that they may not be contiguous. The number and location of the residues at which immune pressure is applied in each individual comes from a statistical distribution, whose choice was guided by clinical data. Human T cell responses are extraordinarily diverse [69] because of the enormous diversity of MHC genes in the population [4]. Thus the same epitopes are not consistently targeted among different hosts. For example, of the 363 residues in the immunogenic structural proteins, p17 and p24, only 46 are targeted by more than 10% of humans, none by more than 23%, and 146 residues are not targeted at all [40]. Thus, Shekhar *et al* considered the number of targeted sites, k, within each host to be chosen randomly between 0 and $n_{\max}$ (=6). Their locations were selected from a uniform distribution across the entire protein, mimicking the high diversity of targeted epitopes. At each targeted site the value of the field b_i was chosen from a Gaussian distribution, whose mean and variance are the same as for the inferred h_i. For all other sites b_i is zero.

Viral dynamics within each host can be simulated using a model similar in spirit to Wright–Fisher models in evolutionary biology. An attempt is made to mutate each residue in proteins in each viral strain in an individual with a certain probability per site (e.g. that of HIV). The new viral strains thus produced survive with a probability equal to

$$P_{\mathrm{surv}}(z) = \frac{1}{1+\exp(-H_{\mathrm{eff}}(z))}\,, \qquad (10)$$

where the consensus sequence is assumed to have the best fitness. The number of viruses is then scaled back to a total of N strains, and the calculation is repeated. The assumption of constant population size may be appropriate as it is expected that most samples of viruses were drawn from patients who were in the chronic stage of infection when viral load does not fluctuate much. Moreover, several studies have shown that, for large enough population sizes, Wright–Fisher like dynamics asymptotically approach results of calculations that do not impose constant population size. Shekhar *et al* studied population sizes that ranged from 2×10^3–5×10^5.

Shekhar *et al* assumed that the Hamiltonian representing the intrinsic fitness of the virus (represented by a single HIV protein, p17) is that inferred from the prevalence of circulating strains for a particular HIV protein, and obtained a set of sequences of this protein derived from the in silico patients after the host-pathogen riposte. They calculated the mutational correlation functions from this set of sequences and compared them to those corresponding to the sequences derived from the real patients (those used to infer the prevalence landscape). They reasoned that if these mutational correlations were the same, then the prevalence landscape is the same as the fitness landscape. This is because the assumed intrinsic fitness landscape is the

prevalence landscape which, by construction, fits the mutational correlation functions observed in the actual sequence database. So, if the sequences obtained from the in silico patients (post host-pathogen riposte) exhibited the same mutational correlations, the prevalence landscape inferred from these sequences would be statistically the same as the assumed intrinsic fitness landscape. The two sets of mutational correlation functions are not the same, but are statistically monotonically correlated [68]. If the Hamiltonian in (9) was of the ferromagnetic form, Griffiths theorem [70] would lead to the conclusion that the fields and coupling constants are also monotonically correlated. Thus, the fitness and prevalence landscapes would be related by a simple shift. However, the Hamiltonian is not of this form, and for a spin glass or Hopfield-like Hamiltonian, an analog of Griffiths theorem does not exist.

Eigen's deterministic equation for viral evolution [71] in the limit of a large number of viral strains is reproduced below

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = r_i x_i - \sum_{k\neq i} W_{ki} x_i + \sum_{k\neq i} W_{ik} x_k - x_i \sum_k r_k x_k \,. \qquad (11)$$

Here, x_i is the frequency of strain i, W_{ki} is the rate of mutation from strain i to k, and r_i is the replication rate of strain i. Leuthäusser showed that this equation describing a non-equilibrium process is isomorphic with the equilibrium statistical mechanics of a 2D Ising model [72]. Each row z^α of this Ising model corresponds to a particular strain, and the next row $z^{\alpha+1}$ is its progeny. Each configuration of this Ising model is a particular evolutionary trajectory, analogous to a path integral formulation. Following developments by Leuthäusser adapted to our problem, the weight of each evolutionary trajectory Σ is given by the Hamiltonian

$$H(\Sigma) = -J_\mu \sum_{\alpha=1}^{n-1}(1-2z^\alpha)(1-2z^{\alpha+1}) + \sum_{\alpha=1}^{n} H_{\mathrm{eff}}(z^\alpha). \qquad (12)$$

The in-row couplings are the intrinsic fitness of a viral strain modulated by the immune response, and the nearest-neighbor ferromagnetic coupling J_μ across rows reflects mutational fidelity originating from the fact that the mutation rate is less than half. A Hamiltonian formulation enables exploitation of the machinery of statistical physics. For example, Shekhar *et al* asked: if one inferred the optimal quadratic Hamiltonian (such as the one that is inferred for the prevalence), how would it correlate with the intrinsic fitness landscape? The quadratic Hamiltonian is of the form

$$H_{\mathrm{T}}(\Sigma) = \sum_{\alpha=1}^{n}\left(-\sum_i a_i z_i^\alpha - \sum_{i<j} K_{ij} z_i^\alpha z_j^\alpha\right), \qquad (13)$$

where K_{ij} and a_i are the couplings and fields, respectively. So, the question above is tantamount to answering what is the fitness order predicted by using (13) with the optimal values of K_{ij} and a_i, and that resulting from J_{ij} and h_i?

To answer this question, Shekhar *et al* used Feynman's variational theory to estimate the optimal values of K_{ij} and a_i for the p17 protein of HIV, such that (13) is as close as possible to (12). They found that, with high statistical accuracy, the fitness and prevalence of circulating strains appear to be monotonically correlated.

The above analysis suggests that while human immunity is an important driver of HIV evolution, its overall effect on prevalence is perturbative and therefore does not significantly compromise the connection between prevalence and fitness. The principal biological reasons that underlie this result are as follows. HIV is a chronic infection and is subject to T cell responses. As we have argued above, human T cell responses are extraordinarily diverse [69] because of the enormous diversity of HLA genes in the population, and so most regions of the viral proteome are targeted by a small fraction of people. Thus the same epitopes are not consistently targeted among different hosts. Furthermore, deleterious escape mutations can revert when the virus is transmitted to a new host [73]. Importantly, although a few HLA-epitope combinations have been associated with better outcome in infected persons [24], the circulating HIV population has not been persistently subjected to classes of effective natural or vaccine-induced memory immune responses. In contrast, the evolution of families of influenza has been strongly directed by persistent and effective natural and vaccine-induced immune responses. Thus, unlike influenza [11, 28], at the population level, HIV evolution is not narrowly directed and strongly biased over time due to persistent effective human immune responses. For these reasons, the prevalence of strains in the HIV population is related simply to intrinsic fitness. Using very different methods than our own, recent work has also shown that patterns of HIV diversity over long times are mirrored across different infected individuals, supporting the claim that universal information about HIV fitness can be derived from prevalence data [74].

Of course, in individual hosts the virus evolves to evade host immunity, forcing HIV to adapt and explore the sequence space. If a mutation that evades host immunity comes at a substantial fitness cost to the virus, compensatory mutations often arise to restore lost fitness, and so mutations at these combinations of residues are observed more frequently than by chance in the circulating virus population. Similarly, some combinations of mutations that are especially deleterious for the virus are observed less frequently than by chance. These correlations, which reflect intrinsic viral fitness effects observed because of the host-pathogen riposte, are the key inputs to our inference procedure. Thus, our landscape describes the collective mutational pathways that HIV uses to evade host immunity and those that it does not. Because of the great diversity of HLA genes, that more than one HLA can target the same HIV epitope, more than one type of HLA molecule is present in each human, and recombination, specific sets of collective mutational correlations observed at the population level, which inform our inference procedure, cannot be uniquely assigned to individual HLA molecules alone [54].

Recent work has correlated clinical markers such as viral load and rate of $CD4^+$ T cell decline with a measure of HIV adaptation to different HLAs [75]. The degree of adaptation as defined therein depends upon the relative frequency of mutations in HIV sequences isolated from individuals with different HLA genes, quantified in terms of a Potts-like model (but one without coupling parameters). Sequences that carry more and stronger HLA-associated mutations [76] are considered

to be more adapted to that HLA. These findings are consistent with our expectations that mutations in more conserved regions, where the difference between the frequency of mutations with and without particular HLAs can be maximized, tend to be deleterious for the virus unless compensated for by other mutations. This also shows that immune responses mediated by a particular HLA are likely to be less effective at controlling infection if pre-adapted escape and compensatory mutations are already present in conserved, immunodominantly targeted HIV epitopes.

The relationship between intrinsic fitness and prevalence can be complicated if one compares fitness differences between strains that differ by many mutations. This is because even if two such strains are reasonably close in fitness, it would be more difficult to access the strain that is more mutations away from the closest common ancestor. Thus, it would be less prevalent, and our inference method would incorrectly predict it to be less fit. Indeed, the simple variational theory reported in Shekhar *et al* related h_i and a_i as follows:

$$a_i^\alpha = h_i - b_i^\alpha + 2J_\mu(1 - 2\langle z_i^\alpha \rangle). \tag{14}$$

The superscripts refer to rows in the 2D Ising mapping of Eigen's equation, as described above. Here the immune pressure b_i makes a particular residue appear more mutable than intrinsic fitness would indicate. The third term reflects mutational fidelity and its size grows with the number of mutations separating two strains. For the influence of this term to be small, one needs to compare strains that are separated by only a few mutations for the simple monotonic relationship between prevalence and fitness to hold. However, it is worth noting that the properties of the HIV reverse transcriptase enzyme enable substantial amounts of recombination [74, 77, 78] and a high mutation rate [18], promoting wide exploration of sequence space and attenuating the effects of phylogeny.

The arguments made above imply that, within some mutational distance, HIV strains have reached a quasi-steady state and that the effects of immune responses are perturbative. Therefore, the relationship between prevalence and fitness differences between strains that differ by a modest number of mutations is simple. Notice that this is not true for influenza, which is driven very far away from equilibrium by persistent effective memory immune responses [28]. In the future, it may be useful to make these arguments more precise by examining Fokker–Planck equations corresponding to viruses mimicking the biological and evolutionary history differences between HIV and influenza. Such analyses may also reveal the range of conditions over which the fitness and prevalence landscapes of a virus are expected to be simply related ('HIV-like'), and how as parameters describing the virus and the immune system change one interpolates between flu-like and HIV-like viruses.

Tests against *in vitro* experiments and clinical data

The extent to which the approximate calculations and biological arguments noted above are valid can only be determined by comparing predictions with *in vitro* experiments and clinical data. Consider *in vitro* experiments first. For example, one could make predictions for the 'energy' corresponding to particular mutant sequences relative to a reference sequence. The model would predict that the replicative fitness of the mutant strain relative to that of the reference sequence should scale negatively with the energy difference between the mutant strain and the reference sequence (equation (6)). The mutant sequences for which predictions are made can be generated through site-directed mutagenesis, and then their relative fitness can be measured by assaying their growth rates when placed in culture with human cells that HIV can infect. Notice that such experiments are carried out in the absence of immune pressure, and so reflect the intrinsic fitness of mutant strains. Figure 5 shows such a comparison between experiments and model predictions for 38 strains of HIV with mutations in the Gag polyprotein [46, 79]. As is evident, the comparison is reasonably good. In published and unpublished work, we also see robust correlations with similar data for HIV strains bearing mutations in the Nef, Pol, and ENV proteins.

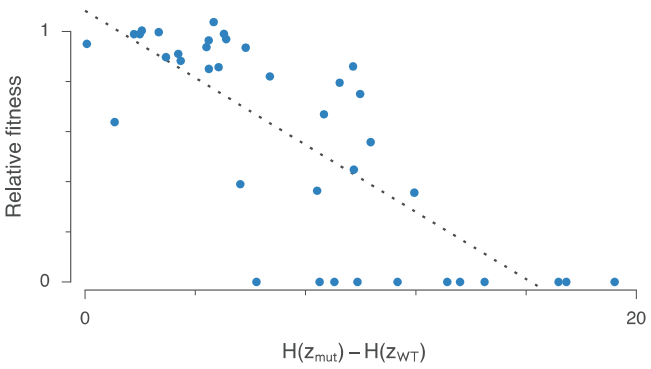

Figure 5. Comparison between experimental measurements of viral growth, shown relative to the growth rate of the wild-type sequence (relative fitness), versus the corresponding difference in energy between the mutant and wild-type sequences. We observe a strong negative correlation between energy and replicative fitness *in vitro* (Pearson's $r = -0.78$, $p = 8.2 \times 10^{-9}$, $n = 36$). Analysis of experimental data from [79].

Another study highlights the ability of the fitness model to capture the effects of interactions between mutations [80]. HIV protease plays an important role in viral replication because this molecule cuts up the polyproteins for which the HIV genes code into the individual proteins required for proper replication and the virus' function. Thus, it has been the target of antiretroviral drug therapy through a class of drugs known as protease inhibitors. The virus is able to make mutations that increase its resistance to protease inhibitors, but these typically come at substantial fitness costs to the virus. The drug resistance mutations that are most likely to be relevant, then, should be the ones whose fitness costs can be compensated by other mutations. One may then ask whether our fitness model can identify potential drug resistance mutations based on the values of the coupling constants in the Hamiltonian. Specifically, when the positive coupling constants exceed a certain value (here the sign of the interaction is important because it should have a compensatory effect), the strain bearing both the mutation that confers drug resistance and an additional mutation becomes sufficiently fit. This

condition $J_{ij} = -h_i - h_j$, can be thought of as a level-crossing phenomenon in statistical mechanics.

It could be argued, however, that predictions emerging from the model could be good simply because the model reflects the presence of drug-resistance mutations in the population of circulating viruses today. Therefore, we inferred a model for protease only using sequences that were obtained prior to 1996 (the year that antiretroviral therapy was introduced). Model predictions compare increasingly well with observed drug-resistance mutations as the cut-off value for the coupling constants increases [80]. These results again suggest that the model can capture effects of coupling between mutations that reflect the virus' intrinsic fitness.

In addition to the tests of predictions with clinical data on elite controllers noted earlier, more recently, predictions from the fitness model have been confronted with data on virus evolution in individual patients. In these patients, viral sequences have been obtained as a function of time, and the first detectable T cell responses and their targets are also known [81]. In these patients the virus makes mutations to evade these T cell responses. Is it possible to predict the residues at which these escape mutations emerged, and the relative times required for this to happen in these patients? Toward this end, Barton *et al* combined Wright–Fisher like evolutionary dynamics with the model for fitness landscapes to try to predict the evolution of a real pathogen under immune pressure in humans [82]. As in the simulations of Shekhar *et al* described above, these simulations involved virus populations of constant population size. Here, however, the logistic map from energy to survival probability,

$$P_{\text{surv}}(z) = \frac{1}{1 + \exp(\beta(H(z) - \overline{H}))}, \qquad \overline{H} = \sum_z H(z), \quad (15)$$

was further refined by including an 'inverse temperature' parameter $\beta \approx 0.07$, computed from comparisons between energy values and *in vitro* replicative capacity measurements [79]. Note that the sum in (15) runs only over the set of sequences in the current viral population. Because the true T cell responses were measured experimentally, it was also possible to include the effects of each individual's immune system in the model. To mimic the killing of infected cells by epitope-specific T cells, all viruses that contained a targeted epitope had their energy increased (i.e. fitness decreased) by a fixed amount, chosen large enough so that escape was favored for all epitopes observed in the clinical data. Barton *et al* then conservatively assumed that any nonsynonymous mutation within a targeted epitope would be sufficient to allow the virus to avoid detection by epitope-specific T cells. Though this assumption is not always correct, there is experimental evidence that most mutations within an epitope tend to substantially impair T cell recognition [83]. Overall, the effect of these assumptions appears to be mild as model predictions for the most likely and second most likely locations for escape mutations matched the clinical data in roughly 85% of the cases.

Interestingly, some of the patients studied provided illustrations of the importance of the effects of coupling between

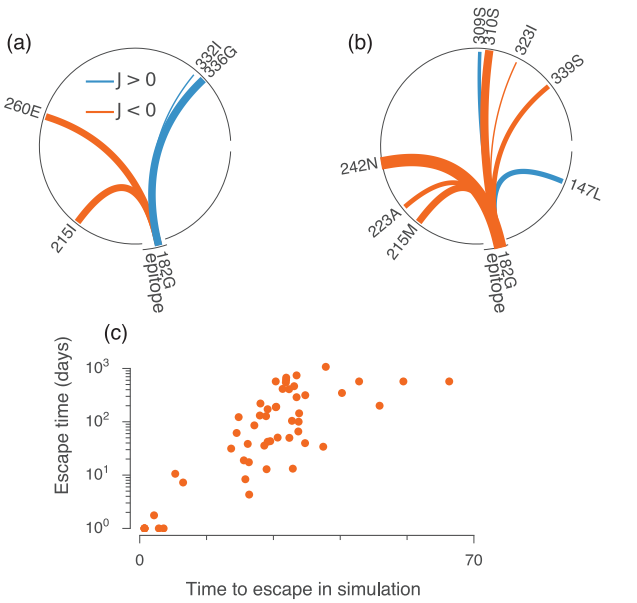

Figure 6. Comparison of identical epitopes targeted by two different individuals reveals effects of the viral sequence background on escape [82]. (a) In one individual, escape occurred through the mutation 182G fairly rapidly (120 d). (b) In another, strong negative interactions between the 182G mutation and residues of the viral sequence background in this particular individual suppress escape, which is not observed to occur even after years of monitoring (>1103 d). (c) Overall, we find good correlation between times to escape in simulation and those found in clinical data (Pearson's $r = 0.81$, $p = 3 \times 10^{-13}$, $n = 53$). Analysis of clinical data from [81].

mutations. This is most vivid when examining pairs of patients who target the same epitope (peptide). Figure 6 shows an example of two such patients. In one individual, the virus escaped T cell immunity fairly quickly, while in the other escape mutations were never observed at high frequency, even after years of infection. The circles in this figure show the rest of the sequence of the protein containing the targeted peptide. Marked residues represent ones that had a mutant amino acid (compared to the reference sequence for the fitness model) that strongly interacts with the ultimate escape mutation. Blue curves represent that the mutation was predicted to be synergistically ($J > 0$) coupled to the escape mutation, and orange curves imply a predicted antagonistic ($J < 0$) interaction. The thickness of the curves corresponds to the strength of these couplings. The patient in whom escape is suppressed was infected with a virus that contained many more mutations at residues that were strongly negatively coupled with the escape mutation compared to the patient in whom the virus escaped quickly. Cases like this suggest that couplings between mutations can play a key role in determining the evolution of HIV in individual patients.

In the evolutionary dynamics that we carry out, the steps of replication, mutation, and selection occur in a single step. However, biologically these steps are separate and involve different time scales. Mutation occurs rapidly during the reverse transcription of the viral RNA into DNA, while selection effectively operates at the level of infected cells (which may or may not successfully produce new viruses, and which may

be killed during the replication process by cytotoxic T cells); a typical lifetime of infected cells is around 2 d when productively infected [20]. Also, our simulations operate in the regime of very strong selection for escape, so that escape is favored even at epitopes where the fitness cost of mutation is high. For these reasons, it is difficult to precisely connect generations of evolution in the simulations of virus evolution to real time. Thus, the predicted generations of evolution required for escape mutations to emerge reflect relative rates. Figure 6(c) shows that the clinically measured escape times compare reasonably well with the predicted evolutionary generations. However, there is considerable scatter, especially for cases where escape is neither very fast or slow; but, adding error bars to the calculated results shows that the rough order is correctly predicted by the simulations. An important question for future exploration is to determine the extent to which the discrepancies between clinical data and predictions can be ascribed to different approximations made in the simulations of evolutionary dynamics or the errors in statistical inference of the fitness landscape. For example, is the comparison with data better if we do not make the approximation of constant population size, or do not encapsulate mutation/replication and selection in one effective time step (by treating infection of new cells as a separate step)?

Based on the progress made so far and on positive correlations between predictions and data, we have designed an immunogen that can be tested in monkeys in the context of a therapeutic vaccine. A major engineering challenge here is devising carriers and adjuvants that can efficiently deliver long peptide immunogens, an issue that also confronts the development of cancer vaccines.

Evolution driven by potentially conflicting selection forces: guiding the choice of strategies for inducing broadly neutralizing antibodies

As noted earlier, an effective prophylactic vaccine against highly mutable pathogens, such as influenza and HIV, would be one that stimulates the immune system to produce antibodies that can neutralize diverse strains (broadly neutralizing antibodies or bnAbs). Recently, such bnAbs have been isolated from some patients [84–87]. In the case of HIV, bnAbs typically only emerge several years after infection, if they are observed at all. Nonetheless, their existence shows that the human immune system can evolve bnAbs. This finding thus raises the tantalizing possibility that, by properly tuning the vaccination strategy, bnAbs could be induced rapidly and efficiently in many people.

As we described, upon natural infection or vaccination, bnAbs are produced by the Darwinian evolutionary process of affinity maturation (AM). It seems evident that one would have to vaccinate with multiple variants of the molecules comprising the viral spike of a pathogen in order to induce bnAbs. Otherwise, the Darwinian evolutionary process of AM should produce strain-specific antibodies.

If multiple variants are used in the vaccine, several new questions need to be addressed. First, what should be the variant pathogens? The part of the viral spike to which bnAbs bind contain residues that are conserved from one strain to another. One example is the CD4 binding site on the spikes of HIV, which is highly conserved because this region must bind to CD4 on human cells in order to enter and infect them. As shown in figure 3, the CD4 binding site is surrounded by highly variable residues and the spike is also decorated with sugars (glycans) that act like a protective shield that hinders antibody access to the conserved protein residues [88, 89]. The variant antigens that are used to induce bnAbs should share a set of conserved residues, and have different variable regions around these residues. The antigens must also mimic the steric constraints present on the actual viral spike. Much work is being done to design such antigens for both influenza and HIV. Much work is also underway to devise immunogens that can activate germline B cells that target regions containing conserved residues and have a chance of maturing into bnAbs [90–92]. We will not be concerned with these issues here.

Our focus here will be on the effects of the temporal order and concentrations of the variant antigens on the probability of inducing bnAbs, provided that the right germline B cells can be activated [92–94]. For example, should the variant antigens be administered as a cocktail, sequentially, or first one variant and then a cocktail? How does the concentration of antigen that is delivered influence AM when there are multiple variant antigens present? How many variant antigens should be used, and what should be the mutational distances that separate their variable regions? The answers to these questions are drawn from a large space of possibilities, and the availability of mechanistic principles that describe how AM occurs could guide promising choices to be explored for induction of bnAbs.

AM has been studied extensively ever since Eisen's seminal experiments describing the phenomenon. It is also the most studied problem in immunology using computation, following the seminal studies by Perelson and co-workers [9, 95–102]. But, most studies focused on AM induced by single model antigens. Thus, little is known about a basic problem in immunology—viz., how do antibodies evolve by affinity maturation when there are multiple variant antigens that share a common set of residues? Shedding light on this fundamental issue will also help address an important challenge to human health.

There have been a few recent studies of this problem using computation [103–106]; very little has been done using theory and simple analytical models with the goal of obtaining deep mechanistic insights.

Recently, Wang *et al* developed a coarse-grained computational model for AM in the presence of multiple variant antigens [103]. This agent-based stochastic simulation method executes a set of rules derived from experimental studies with a single model antigen (see earlier description of AM). However, calculations were carried out using different variant antigens. B cells multiply and mutations are introduced in to a coarse grained model of the receptor with a probability determined from experiments. Half the mutated B cells die because experiments suggest that this

is the probability that a mutation is lethal, another third of mutations are silent (do not affect affinity for the antigen), and the rest are affinity-affecting mutations. Based on experimental results, affinity-affecting mutations are more likely to be deleterious rather than beneficial. A coarse-grained Ising-like model is used to mimic the variant antigens and the BCR in order to compute affinities. Mutated B cells then undergo selection in two steps. Each B cell is successful in internalizing the antigen it encounters with a probability that grows with its BCR's affinity for this antigen and then saturates (Langmuir form). B cells that succeed in internalizing antigen then compete with each other for limiting T cell help, and receive a survival signal with a probability that is determined by the relative value of its affinity for the encountered antigen compared to the average value of the affinity for all B cells that internalized antigen. Most B cells that are positively selected are recycled for further rounds of mutation and selection [7, 9], and a few emerge from the GC as antibody producing cells and memory cells. The simulations of the GC processes end when all B cells in the GC die, when a threshold number of B cells is reached, or when a maximum time has elapsed. The second condition mimics the internalization of all the antigens by B cells, and the last may reflect antigen decay. When multiple variant antigens are simultaneously present, the simulations studied two scenarios: (a) Each B cell encounters only one type of variant antigen at a time on FDCs; (b) Each B cell interacts with all types of variant antigens simultaneously.

The simulations studied the situation where the variant antigens shared a set of conserved residues, but were distinguished from each other by a relatively large mutational distance in the variable residues around it (11 non-overlapping mutations). Known bnAbs against HIV focus their binding on the conserved residues [87]. The coarse-grained Ising-like model for the binding free energies allows three types of contacts of the BCR with the antigen—with conserved residues, with the variable residues, and the glycans. Another feature included in the free energy function is that if a mutation results in a weaker interaction with either a glycan or a mutated variable residue, the interaction with a randomly chosen conserved residue is increased. The opposite is also true. This feature aims to mimic the fact that mutations in the variable residues can insert loops that hinder access to the conserved residues, much like the glycans do.

Stochastic trajectories of the agent-based model for GC reactions are simulated and the affinity of each antibody present in the end for a large set of variant antigens is determined. The 'breadth' of coverage is defined as the fraction of these antigens to which this antibody binds with an affinity exceeding a threshold. When experimentalists measure breadth, they determine the ability of an antibody to prevent infection in a culture containing viruses and cells that can be infected. The relationship between this ability to 'neutralize' a virus and the affinity of the antibody to bind to it is complex. So, in the study by Wang *et al*, they were really studying how cross-reactivity depends upon the immunization scheme. In this study, trajectories representing GC dynamics are simulated many times for exactly the same conditions, and histograms representing the probability with which a typical vaccinated person will produce antibodies with a certain breadth are reported.

Wang *et al* predict that, if the variant antigens are administered as a cocktail, the probability of obtaining bnAbs is small and very few antibodies are produced. During the early stages of AM, the affinity for the antigen is not high, and so strong interactions with the conserved residues have not evolved. If a B cell is positively selected by a specific variant antigen in one round of AM, then mutates its BCR, and goes back for selection and encounters a different variant antigen, the likelihood of it binding sufficiently strongly is small. Therefore, apoptosis is the most likely outcome (figure 7(a)). Thus, the studies revealed that the variant antigens represent potentially conflicting selection forces during the evolutionary process, thus frustrating AM. If one immunizes first with one variant antigen followed by a cocktail of the others, bnAbs do emerge with low probability. This is because AM induced by the first antigen can lead to strong contacts with the conserved residues for some B cells, and this alleviates some of the frustration described above. Wang *et al* reported that sequential administration of the variant antigens results in the highest probability of evolving bnAbs as the conflicting selection forces represented by the three variants are temporally separated. The basic reason for this result is shown in figure 7(b), where the evolution of the footprint of the antibodies on the antigen is depicted. The first round of AM results in a broad footprint on both the conserved and variable residues of the administered variant antigen. When the second variant is introduced, the simulations suggested that roughly two types of lineages of antibodies tend to evolve: (1) Those that strengthen contacts with the conserved residues and not with the newly mutated variable sites or the other ones. (2) Those that evolved stronger contacts with the variable residues that were not mutated in the second antigenic variant. The latter types of lineages have a low probability of survival when the third variant with non-overlapping mutations (compared to the first two variants) is administered. The former types of lineages can continue to strengthen interactions with the conserved residues and thus acquire large breadth. Model experiments with mice show that, consistent with these predictions, cross-reactive antibodies are more likely to emerge upon sequential immunization with variant antigens rather than with a cocktail of the same. Wang *et al* also show that the degree of frustration could be tuned by manipulating the concentration of the antigens.

Note that Wang *et al* assumed that each B cell has only one shot at being selected. In reality, each B cell has a refractory time between attempts to be selected, but in the time a B cell spends in the light zone (where selection occurs), a B cell can make a few attempts to selection. So, in principle, the B cell could see the same variant antigen again with a higher probability than that considered by Wang *et al*. However, this alleviation of frustration would increase the likelihood of producing strain-specific antibodies, not bnAbs. Similarly, if each B cell encountered all types of variant antigens every time it encountered a FDC displaying antigens, it would be able to be selected in successive rounds by the same antigen, but this would not increase the chance of producing bnAbs. It is of course possible that since the B cells survive with greater

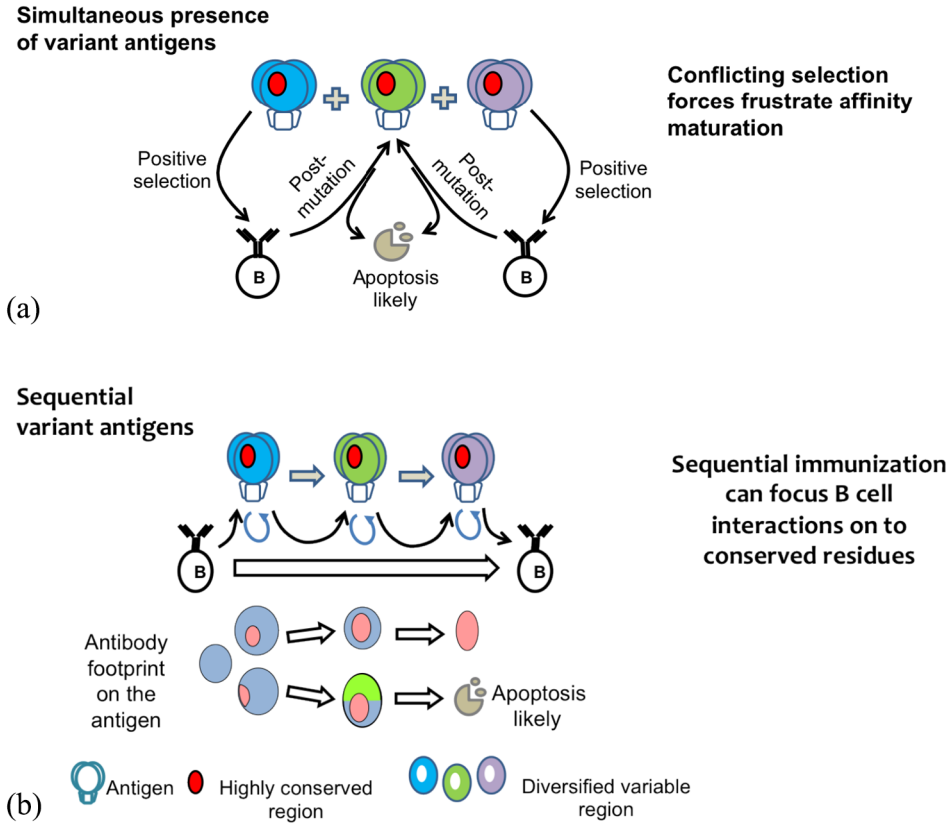

Figure 7. (a) Schematic depiction of frustrated affinity maturation upon immunization with a cocktail of variant antigens. The three variant antigens share conserved residues (shown in red), but have distinct mutations in the variable residues (shown in blue, green, and purple). Because at early stages of affinity maturation, strong interactions with conserved residues are unlikely to evolve, if a B cell is positively selected in one round by a particular variant antigen, then mutates, and encounters a different variant antigen in the next round of selection, it is unlikely to bind sufficiently strongly to it, resulting in apoptosis. (b) Alleviation of frustration by temporal separation of conflicting selection forces represented by the variant antigens. The footprints of the types of lineages of antibodies that tend to evolve in the simulations are shown. Figure adapted with permission from [103]. Copyright 2015 Elsevier.

probability that AM would continue for a while, thus increasing the odds of producing a B cell that might mature in to a bnAb, and this is a possibility that needs further study. Lowering the concentration can increase frustration and manipulation of many variables might lead to bnAbs in these cases.

Three other recent papers on this topic are noteworthy [104–106]. In a very interesting paper, Chaudhury *et al* [106] studied the generation of protective and cross-reactive antibodies against diverse strains of malaria. The malaria apical membrane antigen has a cluster of conserved epitopes and a cluster of highly variable ones. Experiments showed that including a few strains of malaria antigens seemed to give a cross-reactive response. But, the enhancement in cross-reactivity could not be explained simply by the increase in responses to the conserved residues. Chaudhury *et al* [106] described AM by considering the various events that occur in the GC as chemical reactions, and then employed the Gillespie algorithm to numerically solve the underlying Master equations. They employed a shape space representation [107] to model Ab–Ag interactions. Each Ag was comprised of a single conserved epitope (shared by all Ags) and one variable epitope. The relative immunogenicity of epitopes (i.e. the relative number of naïve B cells specific for the epitopes) and the antigenic distance between the variable epitopes were specified. Many other parameters specific to the GC reactions were estimated from the literature. The interesting finding in this paper is that, when the epitopes on an antigen are either conserved or variable, immunization with multiple antigens enhance the fraction of B cells that target the conserved epitope (compared to the use of one antigen), and importantly, also enhance the development of Abs that are more cross-reactive to the variable epitopes. Thus, the polyclonal response from all Abs together results in a protective response to multiple strains. Notice that in the model used by Chaudhury *et al* [106], which may be realistic for malaria, epitopes are comprised only of purely conserved residues or variable residues.

Luo and Perelson [104] used a similar model for the epitopes—i.e. either the epitope is completely conserved or completely variable in the context of HIV. B cells target either the conserved epitope or the variable epitopes. They studied immunization with cocktails of such antigens that were separated by variable degrees of dissimilarity between the variable epitopes. They used a string like model for Ab–Ag interactions, rather than a shape-space model, but otherwise the situation studied by Luo and Perelson [104] is very similar to that studied by Chaudhury *et al* [106]. Luo and Perelson considered broadly cross-reactive antibodies to be those that targeted the conserved epitope. They reported that, upon immunization

with a cocktail of variant antigens, the probability of obtaining cross-reactive antibodies increases if the mutational distances separating the variable epitopes and the number of variable epitopes are larger. This result is most likely due to the fact that Luo and Perelson [104] assume that there is one perfectly conserved epitope shared by all antigens and different antigens contain different epitopes made up of variable residues only and B cells target either the conserved epitope or a variable epitope. The structure of the HIV virus spike [108, 109] indicates that the conserved residues, such as those that comprise the CD4 binding site, are surrounded by highly variable residues and protective glycans, which can be part of an epitope that also contains the conserved residues (as is the case in the study by Wang *et al* [103]).

Childs *et al* [105] have also studied a problem pertinent to the development of cross-reactive antibodies. They used the NK-model, which was first used by Deem and co-workers [101] to study AM. The main result from Childs *et al* [105] is that in the presence of multiple antigenic sites, the breadth of the Abs that evolve is reduced.

In addition, Nourmohammad *et al* [110] recently considered a model of virus-antibody coevolution inspired by HIV. Like Luo and Perelson, they modeled Ab–Ag interactions through a string model with variable and conserved regions for the antigen (viral spike). Using methods from population genetics they modeled the evolution of the mean antibody binding affinity over time, under the assumption that the strength of selection for a B cell is directly related to its affinity. Coevolutionary dynamics reveal a general relationship in the susceptibility of viruses to neutralization by antibodies over time: viruses at any given time are less susceptible to antibodies from previous times, which they have been selected to escape from, while being vulnerable to the antibodies that will arise in the future to neutralize them. Though they do not focus specifically on the development of bnAbs, Nourmohammad *et al* also identify diversity in the viral population as an important condition to promote the emergence of antibodies with a broad neutralization spectrum.

While recent computational work has revealed some interesting concepts and suggested that sequential immunization with variant antigens separated from each other by relatively large mutational distances may be more successful at inducing bnAbs against HIV than a cocktail of the same antigens, many key questions were not addressed. These questions include: (1) What should be the variant antigens that are used as immunogens in a real vaccine against HIV or influenza? (2) What is an optimal vaccination protocol? Of the many permutations, only three scenarios were studied. (3) One could imagine that the optimal vaccination protocol might change if the mutational distance between the variant antigens and their concentrations were manipulated. How does the chance of getting bnAbs depend upon the number of variant antigens used and the mutational distances between them? Could a cocktail with many similar variants followed by a sequence of a few not so similar variants, for example, be helpful? (4) Even if we get the right germline B cells that target an epitope containing conserved residues when activated by a simpler antigen, when the trimeric antigen is introduced many other epitopes that do not have any chance of inducing bnAbs are also available, which can activate new naïve B cells. How does the competition between these B cells and those with a potential for evolving in to bnAbs work? How many distracting epitopes can you have and still evolve bnAbs? Addressing these questions may require a study that is a hybrid of those conducted by Chaudhury *et al* and Wang *et al*. (5) What are the principles that describe how the degree of frustration is influenced by the complex interplay between concentration, mutational distance between antigens, and the number of antigens, and how does this influence the probability of evolving bnAbs? How should we manipulate conditions so that dynamical evolutionary trajectories that result in bnAbs become highly probable?

Our understanding of how we might manipulate conditions to get bnAbs is at an early stage, and much remains to be done. Experimental, clinical, and theoretical studies should ultimately culminate in a general theory for AM in the presence of different spatio-temporal patterns of variant antigen presentation. Such a theory will provide a deep understanding how evolution occurs when there are different spatio-temporal arrangements of the potentially conflicting selection forces, which can then be harnessed toward practical ends.

Acknowledgments

AKC would like to thank Bruce Walker, Mehran Kardar, and numerous current and former lab members (especially, Karthik Shekhar, Vincent Dahirel, Andrew Ferguson, and Shenshen Wang) for fruitful collaborations that resulted in the work summarized in this perspective. We also acknowledge Joy Louveau for help in preparing figure 2. Financial support provided by the Ragon Institute of MGH, MIT, & Harvard (AKC, JPB) and the NIH through Center for HIV/AIDS Vaccine Immunology and Immunogen Discovery grant UM1AI100663 (AKC).

References

[1] Janeway C A, Travers P, Walport M and Capra J D 2005 *Immunobiology: The Immune System in Health and Disease* (New York: Garland Science)

[2] Hozumi N and Tonegawa S 1976 Evidence for somatic rearrangement of immunoglobulin genes coding for variable and constant regions *Proc. Natl Acad. Sci. USA* **73** 3628–32

[3] Eisen H N and Siskind G W 1964 Variations in affinities of antibodies during the immune response *Biochemistry* **3** 996–1008

[4] Robinson J, Halliwell J A, Hayhurst J D, Flicek P, Parham P and Marsh S G 2014 The IPD and IMGT/HLA database: allele variant databases *Nucl. Acids Res.* **43** D423–31

[5] Chakraborty A K and Weiss A 2014 Insights into the initiation of TCR signaling *Nat. Immunol.* **15** 798–807

[6] Victora G D and Nussenzweig M C 2012 Germinal centers *Annu. Rev. Immunol.* **30** 429–57

[7] Victora G D, Schwickert T A, Fooksman D R, Kamphorst A O, Meyer-Hermann M, Dustin M L and Nussenzweig M C 2010 Germinal center dynamics revealed by multiphoton microscopy with a photoactivatable fluorescent reporter *Cell* **143** 592–605

[8] Kepler T B and Perelson A S 1993 Cyclic re-entry of germinal center B cells and the efficiency of affinity maturation *Immunol. Today* **14** 412–5

[9] Oprea M and Perelson A S 1997 Somatic mutation leads to efficient affinity maturation when centrocytes recycle back to centroblasts *J. Immunol.* **158** 5155–62

[10] De Gregorio E and Rappuoli R 2014 From empiricism to rational design: a personal perspective of the evolution of vaccine development *Nat. Rev. Immunol.* **14** 505–14

[11] Korber B, Gaschen B, Yusim K, Thakallapally R, Keşmir C and Detours V 2001 Evolutionary and immunological implications of contemporary HIV-1 variation *Br. Med. Bull.* **58** 19–42

[12] Phillips R E *et al* 1991 Human immunodeficiency virus genetic variation that can escape cytotoxic T cell recognition *Nature* **354** 453–9

[13] Faria N R *et al* 2014 The early spread and epidemic ignition of HIV-1 in human populations *Science* **346** 56–61

[14] Allan J, Lee T H, McLane M F, Sodroski J, Haseltine W and Essex M 1983 Identification of the major envelope glycoprotein product of HTLV-III *Science* **228** 1091–4

[15] Dalgleish A G, Beverley P C, Clapham P R, Crawford D H, Greaves M F and Weiss R A 1983 The CD4 (T4) antigen is an essential component of the receptor for the AIDS retrovirus *Nature* **312** 763–7

[16] Klatzmann D, Champagne E, Chamaret S, Gruest J, Guetard D, Hercend T, Gluckman J-C and Montagnier L 1983 T-lymphocyte T4 molecule behaves as the receptor for human retrovirus LAV *Nature* **312** 767–8

[17] Freed E O 2015 HIV-1 assembly, release and maturation *Nat. Rev. Microbiol.* **13** 484–96

[18] Sanjuán R, Nebot M R, Chirico N, Mansky L M and Belshaw R 2010 Viral mutation rates *J. Virol.* **84** 9733–48

[19] Wei X *et al* 1995 Viral dynamics in human immunodeficiency virus type 1 infection *Nature* **373** 117–22

[20] Perelson A S, Neumann A U, Markowitz M, Leonard J M and Ho D D 1996 HIV-1 dynamics *in vivo*: virion clearance rate, infected cell life-span, and viral generation time *Science* **271** 1582–6

[21] McMichael A J, Borrow P, Tomaras G D, Goonetilleke N and Haynes B F 2009 The immune response during acute HIV-1 infection: clues for vaccine development *Nat. Rev. Immunol.* **10** 11–23

[22] Fraser C, Lythgoe K, Leventhal G E, Shirreff G, Hollingsworth T D, Alizon S and Bonhoeffer S 2014 Virulence and pathogenesis of HIV-1 Infection: an evolutionary perspective *Science* **343** 1243727

[23] Barouch D H 2008 Challenges in the development of an HIV-1 vaccine *Nature* **455** 613–9

[24] Goulder P J R and Walker B D 2012 HIV and HLA class I: an evolving relationship *Immunity* **37** 426–40

[25] Fellay J *et al* 2007 A whole-genome association study of major determinants for host control of HIV-1 *Science* **317** 944–7

[26] Pereyra F *et al* 2008 Genetic and immunologic heterogeneity among persons WHO control HIV infection in the absence of therapy *J. Infect. Dis.* **197** 563–71

[27] The International HIV Controllers Study 2010 The major genetic determinants of HIV-1 control affect HLA class I peptide presentation *Science* **330** 1551–7

[28] Łuksza M and Lässig M 2014 A predictive fitness model for influenza *Nature* **507** 57–61

[29] Martinez-Picado J *et al* 2006 Fitness cost of escape mutations in p24 Gag in association with control of human immunodeficiency virus type 1 *J. Virol.* **80** 3617–23

[30] Brockman M A *et al* 2007 Escape and compensation from Early HLA-B57-mediated cytotoxic T-lymphocyte pressure on human immunodeficiency virus type 1 Gag alter capsid interactions with cyclophilin A *J. Virol.* **81** 12608–18

[31] Wright S 1932 The roles of mutation, inbreeding, crossbreeding, and selection in evolution *Proc. 6th Int. Congress of Genetics* vol 1, ed D F Jones (na) pp 356–66

[32] de Visser J A G M and Krug J 2014 Empirical fitness landscapes and the predictability of evolution *Nat. Rev. Genet.* **15** 480–90

[33] Hinkley T, Martins J, Chappey C, Haddad M, Stawiski E, Whitcomb J M, Petropoulos C J and Bonhoeffer S 2011 A systems analysis of mutational effects in HIV-1 protease and reverse transcriptase *Nat. Genet.* **43** 487–9

[34] Kouyos R D and the Swiss HIV Cohort Study *et al* 2011 Assessing predicted HIV-1 replicative capacity in a clinical setting *PLoS Pathog.* **7** e1002321

[35] Otwinowski J and Plotkin J B 2014 Inferring fitness landscapes by regression produces biased estimates of epistasis *Proc. Natl Acad. Sci. USA* **111** E2301–9

[36] Plerou V, Gopikrishnan P, Rosenow B, Amaral L A N and Stanley H E 1999 Universal and nonuniversal properties of cross correlations in financial time series *Phys. Rev. Lett.* **83** 1471

[37] Laloux L, Cizeau P, Bouchaud J-P and Potters M 1999 Noise dressing of financial correlation matrices *Phys. Rev. Lett.* **83** 1467–70

[38] Plerou V, Gopikrishnan P, Rosenow B, Amaral L, Guhr T and Stanley H 2002 Random matrix approach to cross correlations in financial data *Phys. Rev.* E **65** 066126

[39] Halabi N, Rivoire O, Leibler S and Ranganathan R 2009 Protein sectors: evolutionary units of three-dimensional structure *Cell* **138** 774–86

[40] Dahirel V *et al* 2011 Coordinate linkage of HIV evolution reveals regions of immunological vulnerability *Proc. Natl Acad. Sci. USA* **108** 11530–5

[41] Quadeer A A, Louie R H Y, Shekhar K, Chakraborty A K, Hsing I M and McKay M R 2014 Statistical linkage analysis of substitutions in patient-derived sequences of genotype 1a hepatitis C virus nonstructural protein 3 exposes targets for immunogen design *J. Virol.* **88** 7628–44

[42] Sengupta A M and Mitra P P 1999 Distributions of singular values for some random matrices *Phys. Rev.* E **60** 3389–92

[43] Cocco S, Monasson R and Weigt M 2013 From principal component to direct coupling analysis of coevolution in proteins: low-eigenvalue modes are needed for structure prediction *PLoS Comput. Biol.* **9** e1003176

[44] Ganser-Pornillos B K, Yeager M and Sundquist W I 2008 The structural biology of HIV assembly *Curr. Opin. Struct. Biol.* **18** 203–17

[45] Deeks S G and Walker B D 2007 Human immunodeficiency virus controllers: mechanisms of durable virus control in the absence of antiretroviral therapy *Immunity* **27** 406–16

[46] Ferguson A L, Mann J K, Omarjee S, Ndung'u T, Walker B D and Chakraborty A K 2013 Translating HIV sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design *Immunity* **38** 606–17

[47] Jaynes E T 1982 On the rationale of maximum-entropy methods *Proc. IEEE* **70** 939–52

[48] Lapedes A, Giraud B and Jarzynski C 2012 Using sequence alignments to predict protein structure and stability with high accuracy (arXiv:1207.2484)

[49] Weigt M, White R A, Szurmant H, Hoch J A and Hwa T 2009 Identification of direct residue contacts in protein–protein interaction by message passing *Proc. Natl Acad. Sci. USA* **106** 67–72

[50] Schneidman E, Berry M J II, Segev R and Bialek W 2006 Weak pairwise correlations imply strongly correlated network states in a neural population *Nature* **440** 1007–12

[51] Haq O, Andrec M, Morozov A V and Levy R M 2012 Correlated electrostatic mutations provide a reservoir of stability in HIV protease *PLoS Comput. Biol.* **8** e1002675

[52] Flynn W F, Chang M W, Tan Z, Oliveira G, Yuan J, Okulicz J F, Torbett B E and Levy R M 2015 Deep sequencing of protease inhibitor resistant HIV patient isolates reveals patterns of correlated mutations in Gag and protease *PLoS Comput. Biol.* **11** e1004249

[53] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8

[54] Barton J P, Kardar M and Chakraborty A K 2015 Scaling laws describe memories of host-pathogen riposte in the HIV population *Proc. Natl Acad. Sci. USA* **112** 1965–70

[55] Ackley D H, Hinton G E and Sejnowski T J 1985 A learning algorithm for Boltzmann machines *Cogn. Sci.* **9** 147–69

[56] Plefka T 1982 Convergence condition of the TAP equation for the infinite-ranged Ising spin glass model *J. Phys. A: Math. Gen.* **15** 1971

[57] Georges A and Yedidia J S 1991 How to expand around mean-field theory using high-temperature expansions *J. Phys. A: Math. Gen.* **24** 2173

[58] Morcos F, Pagnani A, Lunt B, Bertolino A, Marks D S, Sander C, Zecchina R, Onuchic J N, Hwa T and Weigt M 2011 Direct-coupling analysis of residue coevolution captures native contacts across many protein families *Proc. Natl Acad. Sci. USA* **108** E1293–301

[59] Ravikumar P, Wainwright M J and Lafferty J D 2010 High-dimensional Ising model selection using ℓ1-regularized logistic regression *Ann. Stat.* **38** 1287–319

[60] Aurell E and Ekeberg M 2012 Inverse Ising inference using all the data *Phys. Rev. Lett.* **108** 090201

[61] Barton J P, Cocco S, De Leonardis E and Monasson R 2014 Large pseudocounts and L2-norm penalties are necessary for the mean-field inference of Ising and Potts models *Phys. Rev.* E **90** 012132

[62] Barton J P, De Leonardis E, Coucke A and Cocco S 2016 ACE: adaptive cluster expansion for maximum entropy graphical model inference *Bioinformatics* **32** 3089–97

[63] Cocco S and Monasson R 2011 Adaptive cluster expansion for inferring Boltzmann machines with noisy data *Phys. Rev. Lett.* **106** 090601

[64] Cocco S and Monasson R 2012 Adaptive cluster expansion for the inverse ising problem: convergence, algorithm and tests *J. Stat. Phys.* **147** 252–314

[65] Berg J, Willmann S and Lässig M 2004 Adaptive evolution of transcription factor binding sites *BMC Evol. Biol.* **4** 42

[66] Sella G and Hirsh A E 2005 The application of statistical physics to evolutionary biology *Proc. Natl Acad. Sci. USA* **102** 9541–6

[67] Brotto T, Bunin G and Kurchan J 2016 Population aging through survival of the fit and stable *J. Stat. Mech.* **2016** 033302

[68] Shekhar K, Ruberman C F, Ferguson A L, Barton J P, Kardar M and Chakraborty A K 2013 Spin models inferred from patient-derived viral sequence data faithfully describe HIV fitness landscapes *Phys. Rev.* E **88** 062705

[69] Goldrath A W and Bevan M J 1999 Selecting and maintaining a diverse T-cell repertoire *Nature* **402** 255–62

[70] Griffiths R B 1967 Correlations in Ising ferromagnets. II. External magnetic fields *J. Math. Phys.* **8** 484–9

[71] Eigen M 1971 Selforganization of matter and the evolution of biological macromolecules *Naturwissenschaften* **58** 465–523

[72] Leuthäusser I 1986 An exact correspondence between Eigen's evolution model and a two-dimensional Ising system *J. Chem. Phys.* **84** 1884

[73] Friedrich T C *et al* 2004 Reversion of CTL escape-variant immunodeficiency viruses *in vivo Nat. Med.* **10** 275–81

[74] Zanini F, Brodin J, Thebo L, Lanz C, Bratt G, Albert J and Neher R A 2015 Population genomics of intrapatient HIV-1 evolution *eLife* **4** 13239

[75] Carlson J M *et al* 2016 Impact of pre-adapted HIV transmission *Nat. Med.* **22** 606–13

[76] Brumme Z L *et al* 2009 HLA-associated immune escape pathways in HIV-1 subtype B Gag, Pol and Nef proteins *PLoS One* **4** e6687

[77] Neher R A and Leitner T 2010 Recombination rate and selection strength in HIV intra-patient evolution *PLoS Comput. Biol.* **6** e1000660

[78] Batorsky R, Kearney M F, Palmer S E, Maldarelli F, Rouzine I M and Coffin J M 2011 Estimate of effective recombination rate and average selection coefficient for HIV in chronic infection *Proc. Natl Acad. Sci. USA* **108** 5661–6

[79] Mann J K, Barton J P, Ferguson A L, Omarjee S, Walker B D, Chakraborty A K and Ndung'u T 2014 The fitness landscape of HIV-1 Gag: advanced modeling approaches and validation of model predictions by *in vitro* testing *PLoS Comput. Biol.* **10** e1003776

[80] Butler T C, Barton J P, Kardar M and Chakraborty A K 2016 Identification of drug resistance mutations in HIV from constraints on natural evolution *Phys. Rev.* E **93** 022412EP

[81] Liu M K P *et al* 2013 Vertical T cell immunodominance and epitope entropy determine HIV-1 escape *J. Clin. Invest.* **123** 380–93

[82] Barton J P, Goonetilleke N, Butler T C, Walker B D, McMichael A J and Chakraborty A K 2016 Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable *Nat. Commun.* **7** 11660

[83] Lee J K *et al* 2004 T cell cross-reactivity and conformational changes during TCR engagement *J. Exp. Med.* **200** 1455–66

[84] Binley J M *et al* 2004 Comprehensive cross-clade neutralization analysis of a panel of anti-human immunodeficiency virus type 1 monoclonal antibodies *J. Virol.* **78** 13232–52

[85] Li Y *et al* 2007 Broad HIV-1 neutralization mediated by CD4-binding site antibodies *Nat. Med.* **13** 1032–4

[86] Scheid J F *et al* 2009 Broad diversity of neutralizing antibodies isolated from memory B cells in HIV-infected individuals *Nature* **458** 636–40

[87] Kwong P D, Mascola J R and Nabel G J 2013 Broadly neutralizing antibodies and the search for an HIV-1 vaccine: the end of the beginning *Nat. Rev. Immunol.* **13** 693–701

[88] Wyatt R, Kwong P D, Desjardins E, Sweet R W, Robinson J, Hendrickson W A and Sodroski J G 1998 The antigenic structure of the HIV gp120 envelope glycoprotein *Nature* **393** 705–11

[89] Kwong P D *et al* 2002 HIV-1 evades antibody-mediated neutralization through conformational masking of receptor-binding sites *Nature* **420** 678–82

[90] Sanders R W *et al* 2013 A next-generation cleaved, soluble HIV-1 Env trimer, BG505 SOSIP.664 gp140, expresses multiple epitopes for broadly neutralizing but not non-neutralizing antibodies *PLoS Pathog.* **9** e1003618

[91] West A P Jr, Scharf L, Scheid J F, Klein F, Bjorkman P J and Nussenzweig M C 2014 Structural insights on the role of antibodies in HIV-1 vaccine and therapy *Cell* **156** 633–48

[92] Jardine J G *et al* 2016 HIV-1 broadly neutralizing antibody precursor B cells revealed by germline-targeting immunogen *Science* **351** 1458–63

[93] Kepler T B *et al* 2014 Reconstructing a B-Cell clonal lineage. II. Mutation, selection, and affinity maturation *Front. Immunol.* **5** 117

[94] Jardine J *et al* 2013 Rational HIV immunogen design to target specific germline B cell receptors *Science* **340** 711–6

[95] Kepler T B and Perelson A S 1993 Somatic hypermutation in B Cells: an Optimal control treatment *J. Theor. Biol.* **164** 37–64

[96] Zhang J and Shakhnovich E I 2010 Optimality of mutation and selection in germinal centers *PLoS Comput. Biol.* **6** e1000800
[97] Keşmir C and De Boer R J 2003 A spatial model of germinal center reactions: cellular adhesion based sorting of B cells results in efficient affinity maturation *J. Theor. Biol.* **222** 9–22
[98] Swerdlin N, Cohen I R and Harel D 2008 The lymph node B cell immune response: dynamic analysis in-silico *Proc. IEEE* **96** 1421–43
[99] Meyer-Hermann M E, Maini P K and Iber D 2006 An analysis of B cell selection mechanisms in germinal centers *Math. Med. Biol.* **23** 255–77
[100] Meyer-Hermann M, Mohr E, Pelletier N, Zhang Y, Victora G D and Toellner K-M 2012 A Theory of germinal center B cell selection, division, and exit *Cell Rep.* **2** 162–74
[101] Deem M and Lee H 2003 Sequence space localization in the immune system response to vaccination and disease *Phys. Rev. Lett.* **91** 068101
[102] Deem M W and Hejazi P 2010 Theoretical aspects of immunity *Annu. Rev. Chem. Biomol. Eng.* **1** 247–76
[103] Wang S, Mata-Fink J, Kriegsman B, Hanson M, Irvine D J, Eisen H N, Burton D R, Wittrup K D, Kardar M and Chakraborty A K 2015 Manipulating the selection forces during affinity maturation to generate cross-reactive HIV antibodies *Cell* **160** 785–97
[104] Luo S and Perelson A S 2015 Competitive exclusion by autologous antibodies can prevent broad HIV-1 antibodies from arising *Proc. Natl Acad. Sci. USA* **112** 11654–9
[105] Childs L M, Baskerville E B and Cobey S 2015 Trade-offs in antibody repertoires to complex antigens *Philos. Trans. R Soc.* B **370** 20140245
[106] Chaudhury S, Reifman J and Wallqvist A 2014 Simulation of B cell affinity maturation explains enhanced antibody cross-reactivity induced by the polyvalent malaria vaccine AMA1 *J. Immunol.* **193** 2073–86
[107] Perelson A S and Oster G F 1979 Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination *J. Theor. Biol.* **81** 645–70
[108] Julien J-P *et al* 2013 Crystal structure of a soluble cleaved HIV-1 envelope trimer *Science* **342** 1477–83
[109] Lyumkis D *et al* 2013 Cryo-EM structure of a fully glycosylated soluble cleaved HIV-1 envelope trimer *Science* **342** 1484–90
[110] Nourmohammad A, Otwinowski J and Plotkin J B 2016 Host-pathogen coevolution and the emergence of broadly neutralizing antibodies in chronic infections *PLoS Genet.* **12** e1006171
[111] Burton D R *et al* 2012 A blueprint for HIV vaccine discovery *Cell Host Microbe* **12** 396–407

Arup K Chakraborty is the Robert T Haslam Professor of Chemical Engineering, Physics, Chemistry, and Biological Engineering, and the founding Director of the Institute for Medical Engineering and Science at MIT.

John P Barton received his PhD in physics from Rutgers University in 2012, where he studied statistical physics under Prof Joel L Lebowitz. He is now a postdoctoral researcher at the Massachusetts Institute of Technology. Currently his research interests include statistical inference methods and computational immunology.