

Quantitative Immunology for Physicists

Grégoire Altan-Bonnet,^{1,*} Thierry Mora,^{2,*} and Aleksandra M. Walczak^{2,*}

¹*Immunodynamics Section, Cancer & Inflammation Program,
National Cancer Institute, Bethesda MD 20892, USA*

²*Laboratoire de physique de l'École normale supérieure (PSL University),
CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France*

Abstract. The adaptive immune system is a dynamical, self-organized multiscale system that protects vertebrates from both pathogens and internal irregularities, such as tumours. For these reason it fascinates physicists, yet the multitude of different cells, molecules and sub-systems is often also petrifying. Despite this complexity, as experiments on different scales of the adaptive immune system become more quantitative, many physicists have made both theoretical and experimental contributions that help predict the behaviour of ensembles of cells and molecules that participate in an immune response. Here we review some recent contributions with an emphasis on quantitative questions and methodologies. We also provide a more general methods section that presents some of the wide array of theoretical tools used in the field.

CONTENTS

I. Introduction	2	A. Cytokine signaling and the JAK-STAT pathway	14
II. Physical chemistry of ligand-receptor interaction: specificity, sensitivity, kinetics.	4	1. Cytokine binding and signaling at equilibrium	15
A. Diffusion-limited reaction rates	4	2. Tunability of cytokine responses.	15
B. Extrapolating collision rates in solution to association rates in the physiological context	5	3. Regulation by cytokine consumption	17
C. Rates and numbers in the physiological context	5	4. Other regulations	17
1. Association rates	5	B. Communication across space and time, and cytokine niches	18
2. Dissociation rates	5	1. Cytokine dynamics as diffusion-degradation	18
3. Numbers of receptors per cell	5	2. Screening by cytokine-consuming cells	18
4. Typical binding rates and some biology	6	3. Probability of autocrine capture	19
D. Receptor-antigen specificity	7	4. Size of cytokine niches	20
1. Cross-reactivity	7	V. Cell fate	20
2. Models of receptor-antigen binding	7	A. Gene regulation and cell differentiation	20
3. Data-driven receptor-antigen binding models	8	1. Th1/Th2 differentiation	21
4. Modeling immunogeneticity	8	2. Other differentiation switches	22
III. Antigen discrimination	9	3. Experimental test of bistability in cell differentiation	22
A. T cells	9	B. Hematopoiesis	24
1. Kinetic proofreading for ligand discrimination	9	1. Timescales	24
2. Adaptive kinetic proofreading	10	C. Inferring the hematopoiesis differentiation tree	25
3. Coupling mechanics and biochemistry: the significance of forces for ligand discrimination	11	D. Cell fate during the immune response	26
B. B cells	13	1. Choice and timing of cell fate under stimulation	26
C. Coarse-graining of molecular details and model reduction	14	2. An aside on cell types	27
IV. Cell-to-cell communication through cytokines	14	3. Inferring cell fate timelines during the immune response	27
		4. Quorum sensing	29
		VI. Repertoires	29
		A. Size of immune repertoires	29
		B. Inference of the stochastic repertoire generation process	31
		C. Thymic selection and central tolerance	32
		D. Diversity and the clone size distribution	34
		E. Repertoire sharing	35
		F. Optimal receptor distribution	36

* Authors are listed alphabetically.

G. Repertoire response to an immune challenge	37	12. Population growth rates	62
VII. Lymphocyte population dynamics	37	D. Ecological models	62
A. Neutral dynamics	37	1. Generalized Lotka–Volterra models	62
B. A note about “neutral processes”	38	2. Susceptible-Infected-Recovered (SIR) Models	63
C. Population dynamics model with external signals	39	3. Solution of stochastic population dynamics with a source	64
D. In host HIV dynamics.	41	4. Solution to forward jump process with opposing drift and source	65
VIII. Affinity maturation	41	5. The Yule process	66
A. Modeling the hypermutation process	42	E. Inference	66
B. Cycles of selection	42	1. Probabilistic inference, maximum likelihood and Bayesian statistics	66
C. Evolution of broadly neutralizing antibodies	43	2. Model selection	67
D. Population genetics approaches to affinity maturation	44	3. Expectation-Maximization	68
1. Evolutionary analysis of repertoire dynamics	44	4. Hidden Markov models	68
2. Models of co-evolution of phenotypic traits	45	5. Information theory	69
IX. Population dynamics of pathogens and hosts	46	6. Maximum entropy models	69
A. Viral fitness models	47	7. Machine learning and Neural networks	69
B. Co-evolution between host and pathogen populations	47	F. High-throughput repertoire sequencing	69
X. Discussion	49	XIII. Acknowledgments	70
XI. Glossary	49	References	70
XII. Methods	50		
A. Physical kinetics	50		
1. Diffusion-limited reaction rate	50		
2. The rates of dissociation between two biomolecules	51		
3. The formation of ligand-receptor pairs: equilibrium and kinetics	51		
B. Gene regulation	52		
1. Basic model	52		
2. Auto-amplification with a single transcription factor	52		
3. Auto-amplification with multiple transcription factors	52		
C. Population dynamics, genetics	53		
1. Deterministic mutation-selection balance	53		
2. Genetic drift	54		
3. Wright-Fisher model	54		
4. Probability of extinction	54		
5. Moran model, continuous limit, and time varying selection	55		
6. Branching processes	56		
7. Coalescence process	57		
8. Site frequency spectra and tree balancing	57		
9. Clonal interference	58		
10. Quantitative traits	59		
11. Lineage reconstruction	60		

I. INTRODUCTION

The role of the immune system is to detect potential pathogens, confirm they really are undesirable pathogens, and destroy them. The goal is in principle well defined. However recognizing molecular friends from foes is not easy, and organisms have evolved many complementary ways of dealing with this problem. Immunologists separate the molecularly non-specific response of the “innate” immune system, which includes everything from scratching to the recognition of protein motifs characteristic of bacteria, and the molecularly specific “adaptive” response, by which specialized cells recognize evolving features of never encountered before pathogens. From another angle we can consider different ways of destroying a pathogen: either swallowing pathogens whole, which is done by cells of the innate immune system called neutrophils and macrophages; or killing our own cells that have been infected by a pathogen or are tumourous — as performed by representatives of the adaptive immune system called killer T-cells. Alternatively, the adaptive immune system produces specialized molecules called antibodies that smother the invader: they attach to pathogens to prevent them from entering cells and multiplying; they bind to bacterial toxins, thereby disarming them; or they bind directly to bacterial cells, flagging them for consumption by macrophages.

That short overview gives us an idea of the many strategies that both pathogens and the host organism have at their disposal (Fig. 1). Pathogenic cells (a bacteria, virus, or a tumour cell) are programmed to divide. The immune system is there to prevent this. To a large

extend, its main challenge is to recognize the unknown. Pathogens are constantly evolving to escape recognition by the immune system. Although the host organism does have a certain list of “warning sign features,” most of which are taken care of by the innate immune system, it is up against a large set of constantly moving targets — as illustrated in our everyday experience by the evolving influenza virus, which requires a new vaccination every year. For this reason, the strategies developed by the immune system are mostly statistical, and require multiple interactions between different types of cells, a lot of checks and balances that leads to a multiplicity of time and length scales (Fig. 2). Despite this complexity, the immune system works remarkably reliably. How do these interactions on many scales dynamically come together in a self-organized way to build a complex sensory system against a high dimensional moving target? This review breaks this high level question into smaller problems and presents some results and concepts contributed by physicists. It also presents a summary of the current methods, experimental and theoretical, used in the field. We do not shy away from the biology of the immune system, but to help the physicist navigate the complexity of immunology, biological details will be introduced as we go along. For an introduction of the immune system for the non-specialist, we refer the reader to the short but excellent book by Lauren Sompayrac [1]. This review does not aim to be exhaustive, but rather focuses on the important physical concepts of immune function, reducing biological complexity to a minimum whenever appropriate.

From an evolutionary perspective, all organisms have some form of protection. Bacteria protect themselves from viruses through specific CRISPR (Clustered Regularly Interspaced Short Palindromic Repeats)-Cas, or unspecific restriction modification systems. The innate immune system is shared by many animals, and is largely similar between us and flies. The adaptive immune system evolved in jawed vertebrates and also has changed little between fish and mammals. Plants also have a well-studied innate system, and have recently been shown to have elements of adaptive immunity. Immunity is therefore a basic element of life. However the details of its implementation and spatial organisation scale with the organism.

Burnet’s clonal selection theory [2], built on previous observations and ideas [3–5], provides a theory of the adaptive immune system in the same sense that Darwin’s theory provides a theory or framework for evolution. Burnet’s theory states that molecules of the pathogens stimulate specific B and T lymphocytes (cells of the adaptive immune system) among a pre-existing ensemble of possible cells, which leads to proliferation of this specifically selected clone. Molecules thus recognized are called antigens. It explains the diversity and specificity of the adaptive immune system, also highlighting its *adaptive nature*. The framework is often summarized in four assumptions: (i) each T and B lymphocyte cell




FIG. 1. **The many scales of the immune system.** The immune system works at many scales from the molecular of receptor-protein interactions, gene regulation, activation of biochemical pathways, to the cellular of cell-to-cell communication directly and through signalling molecules, and organismal level responses using cells of the innate and adaptive immune systems, to the population level where global viral evolution drives the co-evolution of immune systems of different individuals.

has one type of receptor; (ii) receptor-antigen binding is required for cell or receptor proliferation; (iii) offspring of the stimulated cell have the same receptor as their parents; (iv) cells that have receptors that recognize the organism’s self molecules (self-antigens) are removed early in their development. The theory was validated by showing that B-cells always produce one receptor [6] and later by experiments showing immunological tolerance to factors introduced in the embryonic period or immediately after [7, 8]. Burnet’s theory provides an incredibly successful framework to understand the adaptive immune system, yet it is not a quantifiable theory that can be tested against concrete measurements. While what remains to be “filled in” may seem like details, these details still provide a lot of puzzles and uncertainties, also at the conceptual scale. In this review, we present certain examples of quantified ideas that have recently emerged.

We also present some (albeit not all) of the quantitative puzzles, starting from the smallest molecular (sections II and III) and cellular (sections IV and V) scales and moving to the organismal (sections VI, VII and VIII) and population wide (section IX) scales. How do cells discriminate between self and non-self in such an exquisite way integrating processes taking part on many timescales (section III on antigen discrimination)? How do cells communicate and coordinate to orchestrate the immune response as a collective phenomenon (section IV on cytokines)? How do cells adopt specialized pheno-




FIG. 2. **Placing immune processes in length and timescales.** The timescales of immune interactions vary from seconds to years, and the scales from 10^{-10} m for molecular interactions to thousand of kilometers for pathogen evolution


types through cell differentiation to span the range of functions they must fulfil (section V on cell fate)? How does the system cover the space of possible specificities allowing for complete protection against the unknown (section VI on repertoires)? How does the immune system as a whole adapt to the changing environment (section VIII) and how does it influence virus evolution (section IX)?

This review covers a lot of topics, many of them connected. We try to mention these connections, however we have tried to make the sections stand alone and the reader does not have to (and probably should not attempt to) read them all at once, or in the presented order. At the end we include a glossary of the biological terms (section XI) to help navigate the immunological terminology. Detailed appendices (Sec. XII) present general methods that are used in physical immunology as well as other fields.

II. PHYSICAL CHEMISTRY OF LIGAND-RECEPTOR INTERACTION: SPECIFICITY, SENSITIVITY, KINETICS.

Our exploration of the immune system starts at the molecular scale, through the binding of ligands and receptors expressed at the surface of immune cells. Ligands carry information about the pathogenic threat and can be of two types: antigens, i.e. bits of proteins that are recognized by the immune system; and cytokines, which are small molecules secreted by immune cells to communicate with each other about their current state and experience. Binding events between these ligands and their cognate receptors provide the raw signals that cells must interpret to adapt their individual and collective behavior, eventually mounting an immune response in the case of recognition.

All immune decisions start with the classical physical chemistry of ligand-receptor interactions:



This extremely simple reaction leads to strong limits on how fast signals can spread in the immune system (and in biological systems in general), which in turn allows us to discriminate possible regulatory scenarios.

This constraint stems from physical-chemical limits on the parameters of these reactions, typically in the picomolar to millimolar range for the equilibrium dissociation constant $K_D = k_{\text{off}}/k_{\text{on}}$ (we recall that $1nM = 10^{-9} \text{ mol/L} = 10^{-24} \mathcal{N}_{\text{Avogadro}} \mu\text{m}^{-3} \approx 0.6$ molecules per μm^3 , with $\mathcal{N}_{\text{Avogadro}} \approx 6 \cdot 10^{23} \text{ mol}^{-1}$). In this section, we discuss the physical considerations leading to estimates of the biochemical rates driving in these reactions. Such a discussion of the physical chemistry of ligand-receptor interactions is necessary to understand key quantitative aspects of immune activation, as well as discrimination between self and non-self antigens by B & T cells (as we will discuss in Sec. III).

A. Diffusion-limited reaction rates

The basic laws of physical kinetics [9] can be used to obtain an estimate of the diffusion-limited rate of molecular association (see Sec. XII A 1 for a derivation): $k_{\text{on}} \leq k_{\text{diffusion}}$, with

$$k_{\text{diffusion}} = 4\pi(R_{\text{Ligand}} + R_{\text{Receptor}})(D_{\text{Ligand}} + D_{\text{Receptor}}), \quad (2)$$

where R_{Ligand} and R_{Receptor} are the radii of the ligand and of receptor binding pockets (both assumed spherical), while D_{Ligand} and D_{Receptor} are their diffusion coefficients. Since the receptor in Eq. 1 is usually embedded in the lipid bilayer of an immune cell, we can assume it is relatively immobile compared to its ligand L , which diffuses more rapidly in the extracellular medium or in the cytoplasm, $D_{\text{Receptor}} = D_{\text{membrane}} \ll D_{\text{Ligand}} = D_{\text{solution}}$.

The ligand's diffusion coefficient can be estimated from Stokes-Einstein's formula:

$$D_{\text{Ligand}} = D_{\text{solution}} = \frac{k_B T}{6\pi\eta R_{\text{Ligand}}} \approx 300 \mu\text{m}^2 \text{s}^{-1} \quad (3)$$

for a small ligand of size $R_{\text{Ligand}} = 1nm$, diffusing in extracellular medium (whose viscosity is given by that of water, $\eta = 0.7 \text{ mPa} \cdot \text{s}$), at body temperature $T = 310K$ with Boltzmann's constant $k_B = 1.38 \cdot 10^{-23} \text{ J/K}$. Then, from Eq. 2, assuming a small target ($R_{\text{Receptor}} \ll R_{\text{Ligand}}$), we obtain:

$$k_{\text{diffusion}} = 4\pi R_{\text{Ligand}} D_{\text{Ligand}} = \frac{2k_B T}{3\eta} \approx 2 \cdot 10^9 \text{ M}^{-1} \text{ s}^{-1}. \quad (4)$$

Eq. 4 gives a general upper bound for any ligand-receptor association rate, whether it is immunological or not. This limit is conceptually equivalent to the speed of light (although clearly not as fundamental).

Yet, there are examples when ligand-receptor associations between two small proteins seem to "beat" this diffusion limit: $k_{\text{on}} > k_{\text{diffusion}}$ [10]. Such an apparent paradox can be resolved when considering the limiting step for this association. If the ligand interacts weakly with a large object, (e.g. the entire plasma membrane of a cell, DNA coils), inducing directed diffusion in a constrained

space before reaching its specific targets [11, 12], this pre-equilibration step implies that the cross-section of the object our ligand needs to hit can be much larger. The diffusion molecules then can “hop” around these non-specific binding sites to accelerate their search for the specific binding sites. In these cases, the effective R_{Receptor} of collision is the macroscopic scale associated with the large object of weak/non-specific interactions ($5\mu\text{m}$ for cells, $> 0.1\mu\text{m}$ for DNA coils), hence, an apparent increase in the rate of collision according to Eq. 2. A careful calculation (see e.g. [13]) yields the appropriate bound.

The diffusion coefficient will be slowed down for larger macromolecular complexes, and for molecules diffusing inside the cells (there, the apparent viscosity can be increased by 10-fold). Additionally, when considering biomolecules embedded within cell membranes, one must take into consideration the dramatically increased viscosity that leads to slowed-down diffusion for receptor proteins on the surface of immune cells ($D_{\text{solution}} \approx 200\mu\text{m}^2/\text{s}$ vs $D_{\text{membrane}} \approx 5\mu\text{m}^2/\text{s}$ for a typical protein of radius 2nm). For this reason, when considering a soluble protein interacting with proteins embedded within the plasma membrane of immune cells (e.g. extracellular cytokines being captured by a cytokine receptor, intracellular enzymes –such as kinases or phosphatases, that catalyze the phosphorylation/dephosphorylation of proteins– interacting with an activated receptor, etc.), the diffusion of the membrane proteins is so small ($D_{\text{membrane}} \ll D_{\text{solution}}$), than it can be neglected and

$$k_{\text{diffusion}} = 4\pi (R_{\text{Ligand}} + R_{\text{Receptor}}) (D_{\text{solution}}). \quad (5)$$

B. Extrapolating collision rates in solution to association rates in the physiological context

While the rate of collision calculated above is an upper bound to the association rate, a very important limitation must be taken into account: not every molecular collision will lead to their association, and we must estimate the probability of a successful association event. This probability is given by Arrhenius’s law:

$$k_{\text{on}} = k_{\text{diffusion}} e^{-\Delta G_{\text{association}}/(k_B T)}, \quad (6)$$

where $\Delta G_{\text{association}}$ is the free energy barrier (entropic and enthalpic) that molecules must overcome to associate.

The estimation of $\Delta G_{\text{association}}$ is tricky as it requires a deep structural understanding at the atomic level of the entropic, energetic and conformational changes associated with bond formation between two large biomolecules. For the entropic contribution, one rule of thumb is to estimate the numbers of degrees of freedom constrained by the association between biomolecules: simply aligning two biomolecules (in rotation or in translation) incurs an entropic cost of at least $6k_B T$, reducing the probability of association of each collision by

$e^6 \sim 400$. Hence, the basal association rate (before taking into account more subtle molecular constraints) is $\sim 5 \cdot 10^6 M^{-1} s^{-1}$ rather than the $\sim 2 \cdot 10^9 M^{-1} s^{-1}$ estimated in Eq. 4. Additional corrections should be made for each pair of biomolecules (as we will discuss in Sec. IIC 4).

C. Rates and numbers in the physiological context

1. Association rates

The formula in Eq. 2 is particularly useful in the context of quantitative immunology, when one must estimate the kinetics of molecular interactions in different context (solution, intracellular cytoplasm, surface plasma membrane etc.). Most kinetic parameters for biomolecular interactions are measured in solution, as experimentalists routinely purify ligand–receptor pairs and test their interactions in their soluble form (e.g. by calorimetry or by surface plasmon resonance – in the latter case, one molecule must be immobilized). Such measurements can be used to estimate the hard-to-predict activation barrier $\Delta G_{\text{association}}$ for molecular association by inverting Eq. 6, using the measured *in vitro* k_{on} , and the estimated $k_{\text{diffusion}}$. Eq. 2 can be used to rescale the viscosity η and cross-section $R_{\text{Ligand}} + R_{\text{Receptor}}$ of molecular interactions in such a way that one can translate soluble measurements into physiologically relevant parameters.

2. Dissociation rates

As discussed below, immunological interactions span the range of very short lived interactions, $k_{\text{off}} > 10 s^{-1}$ or $\tau_{\text{off}} = (k_{\text{off}})^{-1} < 0.1 s$ for antibody binding to its target in the initial phase of an immune response, to extremely long-lived interactions, $k_{\text{off}} < 10^{-4} s^{-1}$ or $\tau_{\text{off}} > 3$ hours for cytokines interacting with their receptors. These estimates set a huge range of time scales the immune system must deal with, even before taking into consideration delays in cellular responses and how these affect the ligand environment on time scales ranging from hours to days. These considerations form the crux of the matter for quantitative immunology at the cellular scale: while the physical chemistry of ligand-receptor interactions is straightforward, the immune system builds a response of devilish complexity from such elementary interactions.

3. Numbers of receptors per cell

In some situations, such as in T cell antigen recognition, where both the T cell receptor and the antigen exist in a membrane-bound form. All dynamics of interactions must be estimated taking into account the surface concentrations of molecules, with proper adjustment for slower diffusion in the association rates.

One often measures the number of receptors per cell $\#R_{\text{cell}}$, e.g. by assessing the number of cytokine receptors on the surface of cells using quantitative flow cytometry, while the ligand is provided in soluble forms, as is the case with cytokines. In this case, one must estimate the *solution*-level concentrations of available receptors $[R]_{\text{total}}$ in the reaction volume V :

$$[R]_{\text{total}} = \frac{\#R_{\text{cell}}N_{\text{cell}}}{VN_{\text{Avogadro}}}, \quad (7)$$

where N_{cell} is the number of cells in volume V . For example, the effective “concentration” of receptors binding the IL-7 cytokine within a lymph node can be estimated, knowing that each of 10^7 T cells within this $50 \mu\text{L}$ volume express 10^3 receptors:

$$[\text{IL} - 7\text{R}]_{\text{total}} \approx 3 \cdot 10^{-10} M = 300 pM, \quad (8)$$

while $K_D = 10^{-11} M = 10 pM$ for IL-7 binding to its receptor. Hence $[\text{IL} - 7\text{R}]_{\text{total}} \gg K_D$ and any secreted IL-7 will rapidly be captured by the cells within a lymph node. The high density of receptors and cells within lymphoid tissues makes for an interesting regime of competition for soluble ligands, as discussed below (Sec. IV). Such normalization by the Avogadro number, while straightforward, is of critical significance in many immunological configurations, as one must bridge the molecular scale of immune agents (cytokine secretion and consumption) with the functional scale of the system, where competition between cells for soluble ligands takes place.

4. Typical binding rates and some biology

We summed up above the basic and general physical chemistry that drives the molecular association and dissociation of molecules in order to delineate key quantitative parameters in immunological regulation. The precise values of these rates is in fact crucial for the immune system to recognized pathogens.

At its core, the immune system can be considered as a collection of cells, called leukocytes, whose activation registers the presence of “new” molecules — lipid and nucleic acid signatures of viral and bacterial infection for the innate system, or unknown proteins and peptides for the adaptive immune system — and translates into a defensive response — secretion of neutralizing antibodies, killing or phagocytosing of infected targets, etc.

In the innate immune system, non-self recognition is encoded in the structure of the ligand and the receptor. Monocytes, macrophages, dendritic cells, and Mast cells are endowed with genome-encoded receptors called Toll-Like Receptors whose ligands are non-mutable elements of pathogens — peptidoglycans and liposaccharides from the membrane of bacteria, CpG unmethylated dinucleotides derived from viruses. Recognition of non-self in that case is a “simple” lock-and-key proposition whereby pathogenic ligands engage the receptors

and trigger a signaling response that activate innate responses. Hence, for the innate immune system, self vs. non-self discrimination is hard-wired at the molecular level and gets triggered in the early moments of an immune response.

The adaptive immune system offers a more challenging issue. Each B and T cell clone expresses its own unique receptor (and one only), whose assembly is driven by random events (described in detail in Sec. VI and Sec. VIII). The ligands of these random receptors, which are called antigens, are not pre-determined — in fact, they may not even exist at the time of birth of the organism, if for instance they are derived from fast evolving strains of viruses. Understanding how the binding and unbinding rates of such ligand-receptor pairs contribute to self vs. non-self discrimination is one of the core issues in quantitative immunology. Here we give some orders of magnitude of the binding rates to be discriminated.

B and T cell differ fundamentally by the type of antigenic ligands they interact with. T-cell receptors (TCR) interact with their antigen on the surface of antigen-presenting cells, which they scan for anomalies. These antigens are complexes which include a short peptide (around 10 amino acid-long) produced within the antigen-presenting cells by chopping up larger proteins expressed by the cell, either from the genome or from foreign pathogens. Each peptide is loaded onto a large protein called the Major Histocompatibility Complex (MHC), forming a peptide-MHC (pMHC) complex. By contrast, B cell receptors (BCR) bind directly to proteins. The exact position where the binding occurs is called an epitope. Antibodies are soluble versions of the BCR with the same antigenic specificity, and also bind directly to the pathogen proteins to neutralize them.

The affinity of antibodies (produced by B-cells) are highly variable. At the onset of an immune responses, B-cells express and secrete Immunoglobulin M (IgM), which is a pentamerized version of antibodies whose affinity for the target is weak ($K_D > 10 \mu M$): the ability of secreted IgM to bind to pathogens and guide them towards eradication is then driven by the multimeric interaction of IgM with its target. Upon engagement of the adaptive immune response, the interplay between T-cell help and B-cell somatic hypermutation drives the maturation of antibody affinity (see Sec. VIII). At this point, B cells increase the affinity of the antibody they express (down to $K_D < 1 nM$) and switch the class of antibodies from IgM to IgG i.e. from a pentameric version to a monomeric version that does not require as much multimerization of the antibody to bind to their target. Such decrease in K_D and increase in affinity is not driven by better association rates: IgM binds to their target with typical k_{on} around $10^6 M^{-1} s^{-1}$ — similar to IgG. k_{on} is essentially driven by the collision rate and the contribution of the energetic barrier in the association rate is very limited. Instead, the improvement afforded by affinity maturation is driven by a better fit between antibody and antigen, resulting in a lower k_{off} .

In the case of T-cell antigen discrimination, the difference between ligands that will or not activate the immune response is extremely sensitive. Some peptides can elicit a robust activation with a single pMHC molecule, while mutated versions of this peptide differing by just a single amino-acid are unable to trigger T cells even in large quantities ($> 10^6$). The striking feature in T-cell biology is that a single mutation in the peptide of the pMHC only has a marginal effect on k_{off} , while greatly impacting its functional capacity to activate T-cells [14, 15]). Experimentalists can measure the biophysical characteristics of pMHC–TCR interactions in solution, either by surface–plasmon resonance on soluble pMHC interacting with surface–immobilized TCR, by directly observing the surface of cells using single-molecule imaging and fluorescence energy transfer, or by measuring mechanical forces (see Sec. III A 3). Typically, an agonist ligand, defined as a pMHC that will elicit an immune response, binds to their TCR with $\tau_{\text{off}} \sim 1 - 10\text{s}$, while non-agonist ligands bind with timescales that are at least 3-fold shorter (from 0.3 to 3 s), although this varies according to the particular TCR, and also depends on the MHC class (of which there are two, as we will see later). This relatively small difference in binding results in a large fold change ($\geq 10^5$) in response, regardless of the concentrations. Understanding how cells can perform such sensitive discrimination is a major challenge of quantitative immunology, which we will discuss in detail in Sec. III.

D. Receptor-antigen specificity

T- and B-cell receptors do not interact as set of locks and keys matched to each other. Instead, each receptor can bind many different possible ligands, and vice-versa. Here we discuss numbers, data, and models that characterize this many-to-many mapping.

1. Cross-reactivity

Despite the huge diversity of BCR and TCR (reviewed in Sec. VI), the diversity of antigens may be even greater. While such an estimate is difficult for BCR epitopes, a rough estimate of the number of antigens can be obtained for peptide–MHC complexes. For each of the two MHC classes that exist, only a few ($\sim 5\%$) percents of peptides can be presented on any of the 6 MHC genes expressed in an individual. For peptides of length 12, this amounts to $0.05 \times 20^{12} \approx 2 \cdot 10^{14}$ antigens. While this might be a manageable number for humans, who harbor $\sim 10^{12}$ B and T cells, this is not by mice, who have $\sim 10^8$ such cells, and this could be an issue even for humans for longer peptides. This argument led Mason [16] to conclude that each TCR and BCR must be able to recognize more than one antigen, a phenomenon called cross-reactivity or polyspecificity. Cross-reactivity has a less discussed counterpart, which is that each antigen

must be recognized by a large variety of receptors. If there are N antigens, each of which can be bound by k receptors, and R distinct antigen receptors, each of which can recognize ℓ antigens, then one must have $Nk = R\ell$. To fix ideas, the ratio $p = \ell/N = k/R$, the probability that a random antigen and receptor bind together, is thought to be around 10^{-5} [17].

It should be emphasized that finding the sequences of binding pairs of antigens and lymphocyte receptors remains an essentially experimental question, which has received renewed attention lately thanks to the reduced costs of sequencing allowing for high-throughput binding or functional assays [18–21]. However, despite increasing amounts of data on these binding pairs compiled in useful databases [22, 23], and current attempts to leverage these data to make prediction using modern machine learning techniques [24–26], there exists no good predictive model of antigen-receptor specificity. Therefore, the binding models presented in the next paragraph should be viewed as useful toy models for investigating broad properties of cross-reactivity.

2. Models of receptor-antigen binding

The recognition process of antigens by immune receptors is based on molecular interactions between the two proteins: the receptor protein and the antigen, which can be summarized by $K_D = c_0 e^{\Delta G(a,r)/k_B T}$, where $\Delta G(a,r)$ is the interaction free energy of binding between antigen a and receptor r , and c_0 a constant. This energy may be modeled as a string matching problem [27]. In such toy models, each interaction partner is described by a string of length N representing the interacting amino acids. The binding energy between these two specific interaction partners is assumed to depend additively on the interaction between pairs of facing amino-acids:

$$\Delta G(a,r) = \sum_{k=1}^N J_k(a_k, r_k), \quad (9)$$

where the interaction matrix, $J_k(a_k, r_k)$ is an $q \times q$, where q is the size of the space that defines the variability of residues at each position. If we describe each residue as one of the $q = 20$ amino acids, we need to define an amino acid interaction matrix. The Miyazawa–Jernigan matrix [28] was used in such a model to suggest that thymic selection favors moderately interacting residues in TCR [29], or more recently to study the effect of thymic selection on tumor antigenic peptides [30]. The maximum affinity of a receptor to a large set of antigens, which plays a key role in thymic selection (Sec. VI C), is amenable to a statistical mechanistic analysis of extreme value statistics [31, 32], leading to universal features that are robust to the details of the models.

Alternatively, reduced models have been considered, where each residue is defined using a projection that attempts to capture the main biophysical and biochemical

properties in a generalized “shape space” [33]. In the initial string model [34], a_k and r_k were bounded natural integers, and $J_k = a_k \oplus r_k$, where \oplus is the exclusive OR operator acting on each digit of the binary representations of a_k and r_k . This choice was motivated by algorithmic ease rather than biophysical realism. These models were used by Perelson and collaborators to investigate the effect of thymic selection [34, 35] or the immune response [27].

A more drastic reduction is to binarize the antigen: for each epitope position, the amino-acid is defined either as the one present in the viral wild type ($a_k = 1$) or another one (a mutant, $a_k = -1$) [36]. The selective pressure exerted by each receptor, which acts as a “field” on a_k , was drawn as a number r_k from a random continuous distribution, which reduces the model to $J_k(a_k, r_k) = r_k a_k$. In [36], it was additionally assumed that certain positions of the viral epitope were constrained to take the wildtype value, $a_k = 1$ for $k > M$, because of strong conservation at these sites. A similar description in terms of binary strings [37] assumes both the viral epitope and BCR to be binary strings ($a_k = \pm 1$, $r_k = \pm 1$), with a fixed interaction strength, $J_k(a_k, r_k) = \kappa_k r_k a_k$, with again conserved viral positions $k > M$ for which $a_k = 1$ is imposed. Within this description a mismatch between the receptor and the viral “spins” induces an energy penalty. The convention is such that smaller energies imply stronger binding and better recognition. Similar models have been used recently for co-evolution of BCR and HIV [36, 37], the results of which we describe in detail in subsection VIII.

3. Data-driven receptor-antigen binding models

To go beyond the toy models described above, one must experimentally map out the binding energies between specific pairs of antigens and receptors. This can be done in a massively parallel way using high-throughput experiments assaying the binding affinity of many pairs in single experiments. Such an approach was applied to a deep mutational scan experiment reporting the dissociation constant, K_D , of a large number of antibody variants against a fixed antigenic target, fluorescein [19].

The simplest model assumes an additive contribution of each residue to the binding free energy $E(a, r) = \ln(K_D/c_0)$ as in (9), but with $J_k(a_k, r_k) = h_k(r_k)$ fixed to a single value of the antigen a . However, statistical analysis shows that such an additive model is not able to capture all the variability in the binding energy, accounting for less than 2/3 of the variability in double and triple mutants. Epistasis, defined as non-additive effects, accounts for 25–35% of variability in the binding energy between antibodies and the antigen and a large fraction of the epistasis was found to be beneficial [38].

Non-additivity in the binding energy is likely to be a general feature of both TCR and BCR. More sophisti-

cated models including intra-protein or higher-order interactions will be needed to accurately predict receptor-antigen affinity. Ultimately, it would be interesting to reconcile the useful picture of an effective binding shape space with affinity landscapes inferred from data. A major roadblock is that most current experimental techniques only allow for varying one element of the receptor-antigen pair, either testing a library of receptors against a fixed antigen, or a library of antigens against a fixed receptor [18]. Full characterization of the binding landscape would require new methods to test double libraries of antigens and receptors in an ultra-high-throughput manner.

4. Modeling immunogenicity

Rather than modeling the full details of the TCR-antigen interaction, an alternative strategy is to focus on the immunogenic potential of particular antigens, without explicitly modeling the TCR. Such an approach was followed to predict the response of cancer patients to immunotherapy. The prediction is based on the knowledge of the pMHC antigens presented by the tumor cells, called neoantigens [39]. To estimate the ability of a neoantigen a to be recognized by the TCR repertoire, a score was calculated to evaluate its similarity to a database D of antigens known to elicit an immune response:

$$S(a) = \sum_{a' \in D} e^{k s_{a,a'}}, \quad (10)$$

where $s_{a,a'}$ is the alignment score between the neoantigen a and antigens a' from the database, and k an adjustable parameter. The immunogenicity of a is then predicted to be:

$$I(a) = \frac{S(a)}{S_0 + S(a)}. \quad (11)$$

This quantity was combined with the likelihood $A(a)$ that neoantigen a is presented by class-I MHC (predicted by a neural network model [40]) to form a global score $-I(a)A(a)$. With its minus sign, this score reflects the fitness of the tumor cells carrying antigen a . These scores were found to be predictive of patient survival after immunotherapy.

While this similarity-based approach holds great promise for generally predicting the immunogenicity of antigens, further tests are needed to validate the method for broader use. The link between survival and TCR recognition is indirect, and is complicated by the fact that many neoantigens and tumor clones are involved. Direct functional tests of the immune response against libraries of antigens would help consolidate the foundations of this approach.

III. ANTIGEN DISCRIMINATION

We now move to the cellular scale, and describe how the signal propagates in minutes from antigen-receptor binding to biochemical networks, allowing for fine antigen discrimination. We leave to a later section the discussion how this signal is later integrated to make decisions about cell fate and response on timescales spanning hours to days (Sec. V).

A. T cells

1. Kinetic proofreading for ligand discrimination

To model T-cell antigen discrimination, we remain within the quantitative parameters of the “lifetime dogma,” in which the lifetime τ_{off} of TCR-pMHC is the sole determinant of the discrimination process. While this dogma was derived from the biophysical measurements on TCR/pMHC interactions [15], as any dogma it must be revisited regularly for exceptions and refinements. For example, recent studies in the mechanics of TCR triggering have highlighted a new mode of signal initiation, as we shall discuss in Sec. III A 3.

Upon engaging their pMHC ligands, T-cells trigger a cascade of phosphorylation events, which consist of adding phosphoryl groups to conserved residues of TCR-associated chains. The phosphorylations in turn trigger other reactions, ultimately activating transcription factors that modulate gene regulatory responses.

Decision making relative to antigen discrimination can thus be modelled at the phenomenological level by the state of phosphorylation. We assume that, when the TCR complex has accumulated a certain number of phosphoryl groups N_P , it flips a digital switch (to be specific, the activation of the NFAT or ERK molecules [14]) which defines the onset of T-cell activation.

One of the first quantitative models that tackled the exquisite specificity of antigen discrimination during T cell activation was proposed by McKeithan in 1995 [41], building upon the classical kinetic proofreading (KPR) scheme proposed by Hopfield and Ninio in the context of protein translation [42, 43]. Upon binding its ligand, TCR progresses through the different steps of the cascade, controlled by a phosphorylation rate k_P and a de-phosphorylation rate k_{deP} (Figure 3: A). A key aspect of the kinetic proofreading is that, at each step, it is assumed that the complex is de-phosphorylated into unoccupied receptors upon ligand unbinding: this is a reasonable assumption because the CD45 phosphatase — the enzyme that removes phosphoryl groups from the TCR complex — prevents rebinding of TCR by the antigen due to its large size, thus ensuring complete de-phosphorylation before the next binding event. A consequence of this resetting upon unbinding is that antigens forming short-lived complexes with the TCR will

typically only accumulate a small number of phosphoryl groups, below the threshold N_P necessary to activate the immune response, while long-lived complexes can progress more deeply through the phosphorylation cascade.

Neglecting dephosphorylation ($k_{\text{deP}} \ll k_P, k_{\text{off}}$), the number of pMHC–TCR complexes C_i that accumulate i phosphorylations scales with i as [41]:

$$C_i = C_0 \left(\frac{k_P}{k_P + k_{\text{off}}} \right)^i. \quad (12)$$

At steady state, the total number of complexes, $C_{\text{total}} = C_0 \sum_{i=1}^{N_P} C_i$, is given by a second-order equation describing the balance between the rate of binding between free antigens and free receptors, $k_{\text{on}}/(V\mathcal{N}_A)(L - C_{\text{total}})(T - C_{\text{total}})$, where L is the total number of antigens, T the total number of TCR, V the volume, \mathcal{N}_A — the Avogadro number, and the rate of unbinding events, $k_{\text{off}}C_{\text{total}}$. In the limit when TCR are not limiting, $C_{\text{total}} \ll T$, this balance yields the total number of complexes [44]:

$$C_{\text{total}} \sim \frac{(k_{\text{on}}/(V\mathcal{N}_A))TL}{k_{\text{off}} + (k_{\text{on}}/(V\mathcal{N}_A))T}. \quad (13)$$

Further assuming a relatively slow phosphorylation rate, $k_P \ll k_{\text{off}}$, the number of fully phosphorylated complexes C_{N_P} scales as $\propto k_{\text{on}}Lk_{\text{off}}^{-N_P} = Lk_{\text{on}}\tau_{\text{off}}^{N_P}$ as a function of the antigen characteristics — concentration L and binding affinity τ_{off} . This scaling shows how alterations in the pMHC ligand impacting τ_{off} will get amplified into large differences in the amount of phosphorylation accumulating on the TCR:

$$\frac{C_{N_P}(\text{agonist})}{C_{N_P}(\text{non-agonist})} = \left(\frac{\tau_{\text{off}}(\text{agonist})}{\tau_{\text{off}}(\text{non-agonist})} \right)^{N_P}. \quad (14)$$

In other words, kinetic proofreading “reads off” the lifetime of the pMHC–TCR complex and amplifies differences in the output. As with the original KPR scheme of Hopfield and Ninio, this amplification implies energy expenditures caused by the phosphorylation steps. Structurally speaking, the TCR complex contains 20 phosphorylation sites (specifically, tyrosine residues), which can potentially participate in the kinetic proofreading cascade. This means that N_P can be as large as 20, and a two-fold change in pMHC–TCR lifetime could be amplified into a $2^{20} \approx 10^6$ fold difference in the number of phosphorylated receptors C_{N_P} . This mechanism can account for the fact that TCR ligands with minute differences in affinity may elicit very different signals. In addition, it is consistent with the lifetime dogma, in that τ_{off} has a much more determining impact on activation than k_{on} .

However, KPR is insufficient to capture all aspects of T cells’ ability to discriminate between structurally-related ligands. As pointed out by Altan-Bonnet & Germain [45], T cells not only achieve high specificity in

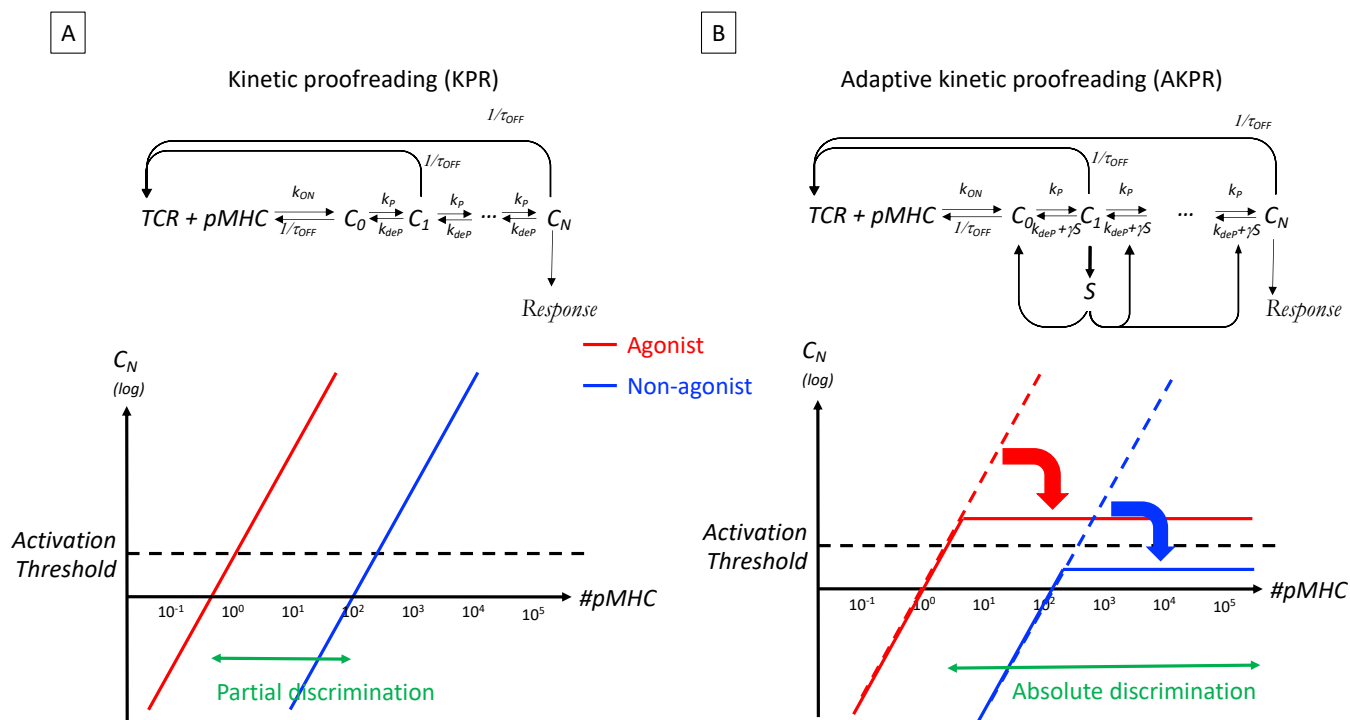


FIG. 3. **Biochemical scheme to reconcile specificity, sensitivity and speed in antigen discrimination.** A) a classical kinetic proofreading (KPR) scheme amplifies differences in TCR phosphorylation (C_N) based on the lifetime τ_{OFF} of the pMHC–TCR complex but gets overwhelmed by a large quantity of pMHC and achieves only partial discrimination of antigens. B) the adaptive kinetic proofreading (AKPR) scheme relies on the activation of a phosphatase (S) to limit spurious activation by a large quantity of non-agonist ligands (see text for details). In the KPR scheme (A) a non-agonist (blue line) can activate a T-cell if it is present in high concentrations. The agonist (red line) activates T-cells also when it is present at lower concentrations. In the AKPR scheme (B) even high non-agonist concentrations cannot lead to T-cell activation, while an agonist still manages to activate T-cells.

ligand discrimination, but they also maintain high sensitivity, as they are able to trigger activation from a single agonist ligand [46, 47]. In addition, they respond very fast, typically within minutes of encountering an antigen presenting cell. These requirements for speed, sensitivity and specificity are hard to achieve all at once. For example, a large number of phosphorylation steps in the KPR scheme allows for high discriminability, but also implies a slower response, as each step must be slow enough to discriminate between agonists and non-agonists. It also affects sensitivity, as more steps imply a lower chance of making it the activation step. These considerations demonstrate how quantitative modeling invalidates a bare KPR scheme, and calls for additional mechanisms.

2. Adaptive kinetic proofreading

To fulfill the conflicting requirements of specificity, sensitivity, and speed, one must expand on the simple kinetic proofreading scheme by adding a mechanism of adaptation to modulate the proofreading steps.

The resulting adaptive kinetic proofreading (AKPR) model [14, 44], which is based on experimental evidence presented by Štefanova *et al.* [48], quantitatively accounts for key features of ligand discrimination by the TCR.

The adaptation module is introduced through a negative feedback mediated by a phosphatase S (which removes phosphoryl groups), which is itself activated by the engaged TCR in state C_1 , so that its steady-state concentration is $S = C_1 / (C_1 + K_S)$, where K_S is a model parameter. The dephosphorylation rate is then enhanced in proportion to the phosphatase concentration, to $k_{deP} + \gamma S$. The equations for steady state have a simple closed form, whose solution is given by the root of a fourth-order polynomial.

Fig. 3 presents a simple graphical argument illustrating how KPR and AKPR schemes perform in their task of discriminating ligands. In the adaptive kinetic proofreading scheme, the negative feedback mediated by S enforces a higher selectivity of antigen discrimination without affecting the high sensitivity of the response.

To better understand how this is possible, consider a large concentration L of non-agonists. Their unspecific engagement with the TCR will cause the first phospho-

rylation of the complex into C_1 , in proportion to their concentration L . This C_1 state activates the phosphatase S , which in turn increases the specificity of the TCR by accelerating its dephosphorylation. When L is large, the amount of phosphatases, and hence dephosphorylation rate, will exactly balance out the amount of non-agonist-bound TCR complexes, preventing these complexes from reaching the activation state C_{NP} (see Figure 3 B). As a result, for non-agonists C_N is essentially flat and stays below the threshold for T cell activation, regardless of L .

The activation threshold for C_{NP} must to be set at very low values to account for the extreme sensitivity of the TCR signalling cascade (a single ligand is sufficient to trigger activation). Hence, one must introduce stochastic equations to appropriately tackle the low number of molecules involved in the TCR decision making, although this does not affect the qualitative features of this particular model.

To conclude, the adaptive kinetic proofreading scheme satisfies the joint requirements of ligand discrimination and sensitivity. It can explain how a single agonist ligand can trigger T cell activation, while a large number ($> 10^5$) of non-agonist ligands cannot. Additionally, it can account for the speed of ligand discrimination by T leukocytes, as only two steps can be sufficient to sort ligands [49]. It is important to emphasize again the importance of a quantitative approach and of modeling to appreciate why achieving speed, sensitivity and specificity in ligand discrimination is indeed such an amazing feat of the immune system.

3. Coupling mechanics and biochemistry: the significance of forces for ligand discrimination

One of the most striking events in the T-cell activation process is the dynamic membrane reorganization into what has been called an immunological synapse. Within minutes of signal initiation, the antigen receptors on the surface of lymphocytes congregate at the center of the interface with their antigen-presenting cells (this has been dubbed a central Supramolecular Assembly Complex or c-SMAC), while transmembrane receptors and phosphatases needed for signal propagation accumulate at the periphery of the contact area (p-SMAC and d-SMAC respectively, standing for peripheral and distal). Such a bullseye pattern has been studied quantitatively using high-resolution and single-cell time-lapsed microscopy. The functional significance of immunological synapses has been identified for the long-term activation of T cells, such as directed killing of antigen-presenting cells or directed secretion of cytokines [51]. As the cell biology of immunological synapse formation was being studied, the question arose as to whether such cellular structure could contribute to self vs non-self discrimination. Qi *et al.* [50] introduced a Ginzburg-Landau model for the mechanical deformation of the T-cell membrane

during synapse formation. The model assumes the presence of several receptor-ligand pairs, indexed by i (e.g. the TCR-pMHC pairs is one these pairs). The free energy \mathcal{F} for the membrane system is estimated as:

$$\mathcal{F} = \sum_i \frac{\lambda_i}{2} \int dxdy C_i(x, y) \left[z(x, y) - z_i^{(0)} \right]^2 + \frac{1}{2} \int dxdy \left[\gamma (\nabla z(x, y))^2 + \kappa (\nabla^2 z(x, y))^2 \right], \quad (15)$$

and the system evolves according to:

$$\frac{\partial C_i}{\partial t} = D_i \left[\nabla^2 C_i + \frac{1}{k_B T} \nabla \cdot \left(C_i \nabla \frac{\delta \mathcal{F}}{\delta C_i} \right) \right] + k_i^{\text{on}}(z) L_i R_i - k_i^{\text{off}} C_i + \zeta_i, \quad (16)$$

and

$$\frac{\partial z}{\partial t} = -M \frac{\delta \mathcal{F}}{\delta z} + \zeta, \quad (17)$$

where $z(x, y)$ is the membrane coordinate for the T-cell membrane, R_i , L_i , and C_i , are the surface densities of receptors, ligands, and receptor-ligand complexes of the i^{th} type, $z_i^{(0)}$ is the length of the $C_i = R_i L_i$ complex bond at rest; and the ζ 's are thermal noises for each process (their distributions are not explicitly stated in the original publications [50]). The free energy \mathcal{F} sums up the elastic energies of each type of complex i (with stiffness λ_i), and the elastic energy of the membrane. k_i^{on} and k_i^{off} are the binding and unbinding rates of ligand-receptor complex formation.

This model is purely mechanical and passive, in the sense that no energy is injected nor is there any biochemical feedback. Yet it captures salient features of the spontaneous formation of immunological synapses. Importantly, synapse formation strongly depends on the lifetime of ligand-receptor complexes, and thus allows for fine discrimination between agonist and non-agonist ligands. The model also emphasizes the functional relevance of size differences in ligand-receptor pairs. For instance, TCR-pMHC complexes have $z_i^{(0)} = 15$ nm, while adhesion complexes such as ICAM-1/LFA-1 are much larger, $z_i^{(0)} = 42$ nm, as it was confirmed experimentally [52, 53].

This model makes several predictions that were validated experimentally. First, it explains how the formation of the characteristic bullseye pattern emerges from the passive sorting of molecules based on their size, upon coupling with the mechanical deformation of the surface membrane. Second, the model predicts an optimum for the lifetime of the pMHC-TCR complex for which the accumulation of antigen in the center of the synapse is maximal. For weak antigens, the membrane energy is dominated by the binding of adhesive complexes, and few antigen-TCR bonds can form because of molecular crowding. For very strong antigens, at intermediate time points (≈ 5 min) or with ligands of intermediate affinity, the membranes of the T-cells adopt an inverted bullseye




FIG. 4. Molecular sorting in the immunological synapse. Upon activation by antigen-presenting cells, T-cells reorganize their membrane proteins to form a bullseye structure (so-called immunological synapse) that stabilizes cell-cell interactions and drives cellular functions. Coupling of mechanical forces derived from the curvature of the membrane and differences in the sizes of pMHC/TCR and integrin (LFA-1/ICAM-1) complexes drives the formation of two "Supramolecular Assembly Complexes" (central: cSMAC; peripheral: pSMAC). At early time points ($t < 2\text{min}$) or for weaker antigens, the systems form an inverted synapse with the tall, rapidly-formed integrin bonds concentrating in the center (red bonds), and pushing the slowly-formed pMHC-TCR bonds to the periphery (blue bonds). For longer timescales and when strong antigen ligands are present, these molecular complexes get sorted into a classical immunological synapse to alleviate the cost of negative curvature in the inverted synapse. A Ginzburg–Landau model (as introduced by Qi et al. [50]) demonstrates that such self-organized sorting could help in antigen discrimination.

pattern with TCR being in the p-SMAC and integrins being in the c-SMAC. Such patterns can be explained when the binding affinity of the integrin binding is larger than the binding energy of the weak ligands. In this case the integrins populate the center of the c-SMAC and make the deformation of the membrane "affordable" energetically (see Figure 4). While this "bell-shaped" activation as a function of antigen affinity has been documented experimentally [54, 55], it has not been consistently observed and there is a sufficient number of exceptions [56, 57] to diminish its significance, especially at the time when clinicians are engineering potent chimeric antigen receptors with unphysiologically-large affinity for their ligands [58, 59]. Third, the state of the membranes of both the T cell and antigen-presenting cell is expected to affect the synapse and thus T-cell activation, as was observed in experiments where the cytoskeleton has been depolymerized [60].

On the other hand, the model also fails to explain more recent observations, such as the ability of T cells to respond to a single agonist ligand while ignoring the

presence of many non-agonist ligands. More recently, its interest has been rekindled by new observations emphasizing the importance of mechanics for ligand discrimination.

Using a biophysical setup to measure forces and lifetimes of individual pMHC-TCR pairs under tension, Zhu and colleagues [61, 62] discovered two types of bonds depending on the peptide sequence: slip bonds and catch bonds. The lifetime of slip bonds decreases monotonously under tension: in other words, they break more easily when pulled. By contrast, the lifetime of catch bonds *increases* under tension. Molecular dynamics simulations [63, 64] suggest that the bonds tension helps expose new residues in pMHC, inducing a better contact with the TCR. At high forces, catch bonds do ultimately break as well, so there exists an optimal force of around 10 pN for which the catch bond lifetime is longest.

This distinction between slip and catch bonds has been suggested to have functional significance in terms of T-cell activation. For instance, it was shown [63] that TCR-pMHC pairs with similar binding constants elicited dif-

ferent activation modes in T cells, with the stimulatory pair forming a catch-bond and the non-stimulatory a slip bond.

How these catch bonds relate to signalling remains to be explained in detail, but this new effect may be critical for predicting binding pairs, and immunogenic neo-antigens of crucial relevance to immunotherapies (see Sec. IID). However, the ability of pMHC to activate the TCR signalling pathway might still be set by the lifetime of the pMHC-TCR complex, but with a “catch”: this lifetime would have to be assessed under the tension exerted by the membrane dynamics, e.g. 10 pN. This would also explain the dependence of activation upon membrane properties such as stiffness. The Ginzburg-Laudau model of Eq. 15 will need to be revisited to account for the existence of catch bonds, e.g. by letting the k_i^{off} of pMHC-TCR pairs depend on the tension $\lambda_i(z - z_i^{(0)})$, possibly affecting how ligands get sorted by mechanical forces in the immunological synapse.

Additional measurements in recent years (in particular, using super-resolution microscopy [65]) emphasize the active role that cytoskeletal rearrangements may play in concentrating TCR in the synapse [66] and the complex interplay between membrane ruffling, receptor sorting, and mechanical tensions. A more complete quantitative model would help identify key limiting steps deciding how strongly T cells get activated in different contexts — different antigen-presenting cells, or different costimulatory contexts.

B. B cells

As mentioned before, the affinity of B-cells with their cognate antigen undergoes a Darwinian selection process that improves their affinities from $K_D = 10 \mu\text{M}$ down to 100pM. Simple counting of the number of occupied receptors in equilibrium could be sufficient to enforce antigen discrimination in B-cells. As for TCR binding to pMHC, the association rate of BCR to antigens is constant for all antigens. But unlike TCR-pMHC binding, which is subjected to an activation barrier, this rate is essentially diffusion limited, $k_{\text{on}} > 10^6 \text{M}^{-1} \text{s}^{-1}$. Strong and weak antigens thus only differ by their binding lifetimes τ_{off} , which varies between 0.1s and 10^4 s, suggesting a possible kinetic proofreading scheme as for TCR [67]. However, since the number of potential phosphorylation sites in the BCR complex is small compared to TCR, such a mechanism may not be as important.

The previously mentioned monomeric version of the BCR, IgG, is itself actually a dimer (so that IgM is a pentamer of dimers), with two binding sites. By analogy with other dimeric receptors on the surface of cells (e.g. Epidermal Growth Factor Receptors or Insulin Receptors [68, 69]), one could assume that this dimerization would help concentrate kinases (the enzymes responsible for phosphorylation) around the receptors, causing them to form clusters. This “cross-linking model” of BCR ac-

tivation is appealing because it is consistent with the formation of microclusters of BCR on the surface of B cells in response to antigens, as observed by super-resolution microscopy.

An alternative model proposed by Reth and colleagues [70] assumes that receptors cluster even in the absence of antigens. When in clusters, BCR inhibit each other and do not signal, but they can assemble and disassemble dynamically. Monomers coming out of these clusters are more prone to signaling, and antigen-binding stabilizes this monomeric state supposedly because of steric hindrances preventing the return into the clustered/inhibited state. This “dissociation-activation” model explains how B cells limit spurious activation of its $\sim 120,000$ surface BCR through their clustering, with signaling made only possible by isolated, antigen-bound monomers. The measurements of Reth and colleagues pose a theoretical challenge in terms of understanding the role dynamic clustering-release-re-clustering of BCR during B cell activation. A full theoretical model of such process and its impact on antigen discrimination remains to be proposed.

One observation that must be taken into account when considering B cell activation is the role of membrane spreading and contraction. Upon initiation of the signaling cascade, B cells rapidly (< 10 min) reorganize their cytoskeleton and membrane, first by spreading to capture as many antigens on the presenting cells as possible, second by contracting to “concentrate” the active receptors. The quantitative model proposed in [71] accounts for binding and unbinding events driving cell spreading and contraction. The model predicts that such a dynamic process can increase dramatically the number of agonist ligands that get captured, compared to a static interface, enhancing the difference between weak and strong ligands, while weaker affinity antigens fail to concentrate. A critical aspect of the model is that, if B cells fail to trigger a sufficient number of BCR by 1 min (a hard cut-off), they terminate the process and shut down their signaling response. Alternatively, the activation response switches to a contraction phase with the surface area decaying according to the phenomenological law $A_{\text{max}} t^{-0.35}$ with A_{max} being the largest area the B cell spreads to, t is the time, and 0.35 is an experimentally-determined exponent. This model is phenomenological as it does not model explicitly the biochemical mechanism driving the spreading and contraction, and makes *ad hoc* assumptions about their behaviour. Yet it illustrates quantitatively how such membrane dynamics can help discriminate ligands.

To conclude, although the issue of antigen discrimination in B cells may not be as stringent as for T cells, it poses interesting quantitative issues for a different parameter range of binding constants ($K_D = 100 \text{ pM} - 10 \mu\text{M}$), over longer timescales ($> 10 \text{ min}$), and using different cellular mechanisms: receptor clustering and membrane dynamics. It is worth recalling that this initial recognition process connects to additional processes over

longer timescales, such as validation by helper T-cells, and dynamics within germinal centers. We will come back to these questions in Sec. VIII.

C. Coarse-graining of molecular details and model reduction

As we go deeper into the molecular details of immune recognition, the number of molecular species, reactions, intermediates, and therefore parameters explodes. This poses a challenge for a number of reasons. First, the effective behaviour of the system as a whole may still be relatively simple, suggesting that simpler phenomenological models may describe these processes equally well. While the variables of such models may be hard to relate directly to molecular entities, they are easier to interpret and allow for better analytical progress and predictions. Second, even assuming that the full complexity of all interactions is needed, large numbers of parameters are likely to lead to overfitting problems, meaning that many parameters or combinations of parameters are underdetermined, undermining the accuracy of predictions. And even when they can be determined, it is not always clear which ones need to be fine-tuned to ensure proper function, or what are the broad design principles presiding over their choices. In Sec. XII E 2, we briefly review the principles of model selection – the classical approach of reducing model complexity from statistics. However, that approach requires to have first defined a hierarchy of models to test, from least to most complex. Besides, model selection relies on goodness of fit as a criterion to evaluate models, while in many cases we may be more interested in capturing the principal features of a biological function, rather than fitting all the data.

In another approach to reducing complexity, François and Hakim [72] developed an method for generating simple molecular networks *in silico* that realize a desired biological function, simply based on a genetic algorithm that selects the “best” solutions. Applied to the problem of absolute ligand discrimination reviewed in Sec. III A, this method infers a class of network motifs, called “adaptive sorting”, that recapitulates known features of T cell activation [49]. In particular, it predicts the emergence of kinetic proofreading and biochemical adaptation. However, the chemical species and reactions of the networks produced by that method may generally not be directly related to the known actors of the phenomenon under study.

To reduce model complexity while keeping close to the details of actual biochemical reactions, one can instead start from a complex biochemical network described by many ordinary differential equations, and “prune” its parameters by setting them, individually or by their combinations (ratio or products), to 0 or infinity [73]. Applying this strategy to a complex model of T cell recognition [74], which contains close to 100 parameters, shows that its behaviour can be boiled down to just three cou-

pled differential equations highlighting its main features of adaptation and discrimination, and reveals the broad design principles that implement these features.

Beyond this particular example, this kind of approach has great potential for helping to make sense of complex biological systems with many entities and interactions, as is often the case in the immune system.

IV. CELL-TO-CELL COMMUNICATION THROUGH CYTOKINES

Cytokine communication is critical to synchronizing the activation of various immune cells and to bridging multiple spatial-temporal scales in immune responses, from local or individual cell activation to global, systemic responses. Cytokines are small glycoproteins that get produced and secreted by all cells, immune or not, with varied dynamics, amplitude and frequency. These cytokines then diffuse and bind to receptors present on the surface of adjacent cells to elicit a signaling response that trigger a gene regulatory response. In short, cells use cytokines to communicate between themselves. In this section, we discuss quantitative models that have been introduced to model how individual leukocytes respond to cytokines (*e.g.* using the JAK-STAT pathway) over multiple time- and length-scales.

A. Cytokine signaling and the JAK-STAT pathway

Cytokine receptors are heterodimers, or more rarely heterotrimers, so that they carry at least two intracellular signaling domains. The JAK-STAT pathway is the dominant pathway engaged by cytokine signaling. Upon binding by its cognate cytokine, the receptor undergoes a conformational change that presents phosphoryl attachment sites on the JAK (Janus Kinase) receptor to face the kinase domain that induces phosphorylation, which in turn leads to activation. Then a protein called STAT (Signal Transducer and Activator of Transcription) interacts with the intracellular domains of the receptors, gets phosphorylated by the activated JAK, and dimerizes upon release from the receptor. Dimers of STAT then translocate into the cell nucleus, bind to the chromatin in specific sites and elicit a transcriptional response. Such signal transduction is one of the simplest biological pathway connecting the extracellular environment and its messenger cytokines to a transcriptional response (Figure 5). Its complexity is encoded in the large number of pairs of cytokines (> 40) and their receptors that cells can express, as well as the 7 forms of STAT and $7^2 = 49$ homo- or hetero-dimers that they can form upon activation [75]. Additionally, the dynamics of cytokine signaling enriches the biology of these pathways, with the existence of positive and negative feedback regulations: cytokine signaling inducing the expression of additional cytokine receptors, cytokine degradation, cy-

tokine receptor endocytosis, expression of negative regulators such as Suppressor of Cytokine Signaling (SOCS). Here we present simple derivations that can help understand the quantitative regulation of cytokine communication within the immune system, which underlies the coordination and orchestration of the immune response.

1. Cytokine binding and signaling at equilibrium

In this section, we discuss the contribution of the field of quantitative immunology to understanding the JAK-STAT pathway as a signal transduction cascade in well-mixed conditions. Spatial heterogeneities in cell-to-cell communication via cytokines will be tackled in section IV B.

Within an individual cell, the biochemical reactions of the JAK-STAT pathway can be treated as well-mixed: the typical concentrations of molecules are fairly high, with ~ 1000 molecules per cell of diameter $10\mu\text{m}$, which translates into concentrations $\sim 100\text{nM}$. We also assume that most reactions are essentially diffusion-limited, as discussed before: $k_{\text{on}} > 10^7 \text{M}^{-1}\text{s}^{-1}$. The signal transduction cascade for JAK-STAT cascade can generically be modelled in a step-wise manner, as sketched in Figure 5.

The kinetics of cytokine binding to its receptor activating JAK into JAK* can be modeled as a simple bimolecular process:

$$\frac{d}{dt} \# \text{Receptor}^* = k_{\text{on}} [\text{Cytokine}] \# \text{Receptor} - k_{\text{off}} \# \text{Receptor}^* \quad (18)$$

where $\# \text{Receptor}^*$ is the number of cytokine-engaged receptors, $[\text{Cytokine}]$ is the cytokine concentration. Because binding is very strong (typically $K_D = k_{\text{off}}/k_{\text{on}} \approx 10\text{pM}$ with $k_{\text{off}}^{-1} \approx 3600\text{s}$ and $k_{\text{on}} \approx 3 \cdot 10^7 \text{M}^{-1}\text{s}^{-1}$), most cytokine molecules bind at a diffusion-limited speed with very strong binding. We assume a large extracellular volume, so that cells do not deplete cytokines as they bind them. Receptors are usually pre-loaded with JAK, waiting for the conformational change associated with cytokine binding, to induce phosphorylation. From Eq. 18 the number of activated JAK (JAK*) at steady-state is:

$$\text{JAK}^* = \text{JAK}_{\text{total}} \frac{[\text{Cytokine}]}{K_D + [\text{Cytokine}]}, \quad (19)$$

where $\text{JAK}_{\text{total}}$ is the total JAK receptor concentration. The kinetics of STAT phosphorylation into pSTAT is described using a classical Michaelis-Menten equation:

$$\frac{d}{dt} [\text{pSTAT}] = k_{\text{cat}} [\text{JAK}^*] \frac{[\text{STAT}]}{[\text{STAT}] + K_m} - k_{\text{dephos}} [\text{pSTAT}], \quad (20)$$


where $[\text{JAK}^*] = \text{JAK}^*/(V\mathcal{N}_A)$ is the concentration of activated JAK within the cytoplasmic volume V , pSTAT is the phosphorylated form of STAT, k_{cat} and k_{dephos} are phosphorylation and dephosphorylation rates, and K_m is

the dissociation constant of STAT to the receptor, which controls the specificity between different receptors and different variants of STAT. Note that, for most signaling networks, the identity of phosphatases that take care of dephosphorylating receptors and kinases remain often undetermined (because of overlapping and pleiotropic activities): their activity is modeled phenomenologically as a single rate reaction.

This system reaches steady-state within minutes [75, 76] such that one can solve $d[\text{pSTAT}]/dt$ in Eq. 20 for pSTAT (with the conservation of matter condition $[\text{pSTAT}] + [\text{STAT}] = S_{\text{total}}$) as a quadratic equation. Combined with Eq. 19, the result gives a direct expression of the response STAT as a function of the input cytokine concentration. This model serves as the basic building block to tackle the dynamic complexity of cytokine responses.

2. Tunability of cytokine responses.

In the previous section we assumed that cytokines, which are formed of several sub-units, were in a pre-assembled configuration. When that is not the case, cytokine binding to the cytokine receptor can be modeled as a two-step scheme.



The cytokine ligand L binds to the cytokine receptor composed of two parts R_1 and R_2 . A simple equilibrium model can be solved, with the dissociation constants $K_i = k_i^{\text{off}}/k_i^{\text{on}}$:

$$K_1 [\text{LR}_1] = [\text{L}][\text{R}_1] = [\text{L}]([\text{R}_1]_{\text{total}} - [\text{LR}_1])$$

$$K_2 [\text{LR}_1\text{R}_2] = [\text{LR}_1][\text{R}_2] = [\text{LR}_1]([\text{R}_2]_{\text{total}} - [\text{LR}_1\text{R}_2])$$

that yields the total concentration of fully-engaged and signaling receptors:

$$[\text{LR}_1\text{R}_2] = \frac{[\text{R}_2]_{\text{total}}}{1 + \frac{K_2}{[\text{R}_1]_{\text{total}}} \left(1 + \frac{K_1}{[\text{L}]}\right)}. \quad (22)$$

This two-step model was shown to be valid for one of the most studied cytokines, IL-2 [77]. IL-2 binds weakly to the abundant α chain of the IL-2 receptor ($\text{R}_1 = \text{IL2R}\alpha$), before being “locked” into a stable and signaling complex by association with the two other chains ($\text{R}_2 = \text{IL2R}\beta\text{-}\gamma_C$) to form the full IL-2–IL2R complex. In that case, $K_1[\text{L}] \ll 1$ such that equation 22 simplifies to:

$$[\text{LR}_1\text{R}_2] = [\text{R}_2]_{\text{total}} \frac{[\text{L}]}{[\text{L}] + \frac{K_1 K_2}{[\text{R}_1]_{\text{total}}}}. \quad (23)$$

This equation reveals key insights for IL-2 and other cytokines [78]: the amplitude of cytokine signaling is proportional the number of $\beta - \gamma$ part of the IL2 receptor

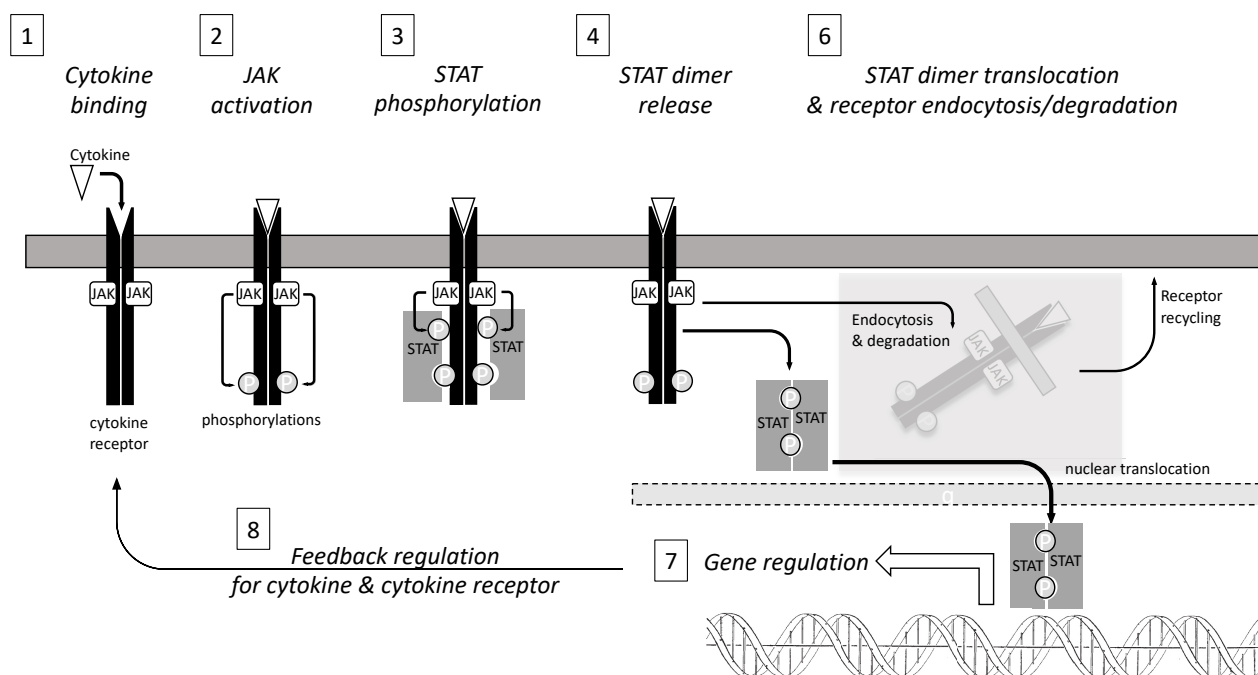


FIG. 5. Model of the regulation of JAK-STAT signaling in response to cytokines. This model in the adiabatic regime can be solved analytically.

(the limiting part), and reaches half of its maximum at $[L]_{50} = K_1 K_2 / [R_1]_{\text{total}}$, inversely proportional to the α chain of the IL-2 receptor (the non limiting-part).

Such tunability of both the amplitude and sensitivity of the dose response of cytokines based on the number of cytokine receptor chains can be critical to achieve immune plasticity, as demonstrated in the context of competition for limited amounts of cytokines. For example, the tug-of-war for IL-2 between effector T-cells and regulatory T-cells is mediated by the exact levels of IL2R α on the surface of these cells. Stimulated *effector* T-cells initiate an immune response, whereas stimulated *regulatory* T cells (Treg) downregulate the response to suppress autoimmune responses (effector response to self-antigens). In general, Treg cells are thought to bypass negative thymic selection despite their strong recognition of self-antigens. These cells then act as pre-activated sentinels that respond to inflammatory cues (such as the upregulation of self-antigens) while not contributing to inflammation itself through cytokine secretion: Tregs constitutively express transcription factor FoxP3, a general downregulator of cytokine production.

Whenever an effector T-cell initiates a spurious response to self-antigens, Tregs can extinguish it by consuming cytokines (*e.g.* IL-2) and by downregulating their inflammatory impact. In this tug of war, whichever cell type, effector or regulatory, expresses more receptors boosts its ability to capture the cytokine and deprives

the other cells of this key cytokine for anti-apoptosis and proliferation in the T-cell compartment. Quantitative models of such competition for IL-2 based on differential expression of IL-2R α receptors have illuminated how self vs. non-self discrimination can emerge from such IL-2 tug-of-war [77, 79, 80].

Some cytokine receptors have an extracellular domain that gets secreted extracellularly, in a soluble form. These can “pre-bind” the cytokine in the extracellular milieu, and deliver it to the complementary chain bound to the cell membrane: such molecular event can have positive or negative effects on cytokine signaling depending on the context. If the pre-formed soluble cytokine/receptor complex binds to the cell to form an incomplete receptor which lacks the intracellular signaling domain of the soluble cytokine receptor, the JAK misses its trans-phosphorylation partner and the cytokine/cytokine receptor fails to signal— “no clap from one hand.” Such soluble complexes act as decoys that antagonize cytokine response (as with viral analogues of IFN receptors, or IL-1 receptors), limiting inflammation to the benefit of viruses [81].

Alternatively, secretion of soluble cytokine receptors can trans-activate cytokine signaling and result in a boost in cytokine response (as with IL-1, IL-2 or IL-6). For example, in the case of IL-6, the soluble portion of the IL-6 receptor (sIL-6R) gets secreted to stabilize IL-6 in the extracellular medium (most cytokines are very small

proteins with short half-lives) and to accelerate the assembly of a complete IL-6/IL-6R/gp30 signaling complex by binding to gp130 dimers on the surface of cells: such activation of the IL-6 signaling response thus does not require the presence of the IL-6 receptor in the membrane of the receiving cells, but only the ubiquitous presence of gp130: this is called cytokine trans-activation [82].

Hence, depending on the exact molecular details, secretion of cytokine receptors in the extracellular environment can trigger or antagonize the cytokine signaling response. The role of quantitative immunology in that context is then to tease out the physico-chemical parameters of such regulation to better understand in which regime inflammatory cues are regulated.

3. Regulation by cytokine consumption

Immune responses need tight regulation to avoid spurious auto-immune activation. This is particularly important for cytokine regulation as overabundant cytokines can be extremely deleterious to the organism as a whole. For example, high concentrations of cytokines induce a “capillary leak syndrome,” whereby tissues lose their barrier against blood serum: this results in septic shocks or viral hemorrhagic fevers. One simple mode of regulation of cytokine signaling by the JAK-STAT pathway is simply to limit the availability of cytokines in the extracellular medium. Most cytokines are very small proteins (with molecular weights in the 10-20kDa range) such that their half-life in the bodily fluid is short. For example, the half-life of IL-2 injected intravenously for the treatment of renal cell carcinoma is at most 3 h in the serum. More significantly, even in the extracellular environment of lymphoid organs, cytokines have a very short half-life due to rapid binding and consumption.

Given the strong binding of cytokine to their receptors (typically $K_D < 100$ pM or $k_{\text{off}} < (3000\text{s})^{-1}$), immune cells can rely on binding to switch off the cytokine signaling response through buffering. This resetting process relies on cytokine consumption by endocytosis of the cytokine with its receptor, trafficking them towards lysosomes, release of the cytokine because of low pH within the lysosomal compartment, and degradation. Note that the signal transduction may persist after receptor endocytosis as long as the cytokine-receptor pair remains engaged [83]. What happens to the cytokine receptor in this process varies based on the cytokine identity. For example, the endocytosed IL-2 receptor dissociates and has its chains sorted towards different compartments: the IL2R γ_C chain goes to the lysosome and gets degraded, while the IL2R α chain goes to early endosomes and gets recycled. This process of endocytosis, sorting, recycling and degradation can be very complex at the molecular level. Yet, a coarse-grained model of this process as a single-biochemical step can be sufficient when modeling IL-2 availability [84, 85]: the typical rate for this step has been measured to be of the order of $(900\text{ s})^{-1}$.

One functional consequence is that cytokines get rapidly consumed while triggering signaling. At the more global level, one can integrate a dynamic equation for production/consumption that accounts for the rise and decay of inflammatory signals in the immune system:

$$\frac{d}{dt}[\text{Cytokine}] = \kappa_{\text{prod}} - \kappa_{\text{consum}}. \quad (24)$$

In the most simple case, $\kappa_{\text{prod}} = N_{\text{prod}}k_{\text{prod}}/(V\mathcal{N}_A)$ is fixed by the number N_{prod} of activated cells at a given time, the rate of secretion k_{prod} per individual cells, and the extracellular volume V . The rate κ_{consum} is determined by the rate of binding of the cytokine to its receptor on the surface of the cells:

$$\kappa_{\text{consum}} = k_{\text{on}}N_{\text{consum}}N_{\text{R}}[\text{Cytokine}]/V\mathcal{N}_A,$$

where N_{consum} is the number of consuming cells and N_{R} the number of receptors per cell. The time dependency of N_{prod} , N_{consum} , κ_{prod} , and N_{R} can be arbitrarily complicated and must be parametrized for each immunological setting under consideration. For example, in the case of IL-2 in the early events of an immune response within a lymph node, $V \approx 50\mu\text{l}$ (free volume), $k_{\text{prod}} = 10\text{s}^{-1}$, $k_{\text{on}} = 3 \cdot 10^7 M^{-1}\text{s}^{-1}$, $N_{\text{consum}} = 10,000$ Treg cells, each endowed with typically $N_{\text{R}} = 3,000$ receptors. For IL-2 to accumulate and reach a significant concentration (typically, $10^{-11}M$ to trigger STAT5 phosphorylation), one must have $\kappa_{\text{prod}} > \kappa_{\text{consum}}$, i.e. $N_{\text{prod}} > 1,000$. This estimate illustrates that there exist thresholds of activation for immune responses, whereby the systems needs a critical mass of activated, cytokine-secreting cells to overcome consumption and drive activation and differentiation [86]. The estimate above does not take into account some of the intricacies of IL-2 regulation: positive feedback in IL-2 secretion, recycling of IL2R α chains, upregulation of IL-2R in T_{reg} cells have been documented. To account for these, Eq. 24 must be solved numerically in more complex settings [84, 85].

4. Other regulations

Cytokine consumption is one major mechanism to limit the duration of availability of cytokines in the extracellular medium, but there exist additional, cell-intrinsic mechanisms that also limit or expand the duration of JAK-STAT signaling in cells, as illustrated by the following examples.

Certain cytokine pathways rely on intracellularly stored pools of receptors that get recruited upon receptor engagement and JAK-STAT signaling. This recruitment further fuels JAK-STAT signaling by avoiding saturation of the initial receptors on the membrane surface. This process was analyzed quantitatively in Ref. [87] to reveal how a certain type of leukocyte called erythroid progenitors can sense a cytokine called Epo over a large range of concentrations (> 1000 -fold range), where classical ligand-receptor binding would predict fast saturation with increasing cytokine concentration.

Another positive feedback in JAK-STAT signaling exists in the IL-2 signaling pathway: The α chain of the IL-2 receptor gets upregulated upon IL-2 signaling, lowering the cytokine concentration at which activation is half-maximum (by virtue of Eq. 23), hence driving further IL-2 signaling [77–79]. Such positive feedback was shown to be critical to committing cells to long-term JAK activation and proliferation [85].

Alternatively, JAK-STAT pathways are also endowed with negative regulators, such as SOCS and Cytokine Inducible SH2-containing (CISH) proteins, which compete with STAT for JAK binding, preventing activation [88]. This can be quantitatively modeled with SOCS or CISH binding competitively with a stronger K_m than STAT onto the receptor (Eq. 20). Varying the exact value of K_m allows SOCS and CISH to negatively regulate JAK activity in two separate regimes of concentrations of cytokines [89]. Such dual feedback regulation of signaling was also shown to extend the range of cytokines that trigger STAT5 phosphorylation while avoiding saturation, and ultimately control cell survival. The modeling of such complex regulation of the JAK-STAT pathway involves adding biochemical steps that can be integrated numerically. The challenge in this context is to acquire a large number of system-specific biochemical parameters. This task has become more amenable as the field progresses and quantitative approaches have been delivering more and more estimates: protein abundances, biophysical parameters, enzymatic rates [78, 90–92].

To conclude, in this section we have reviewed the basic equations governing the regulation of the cytokine-activated JAK-STAT pathway in the immune system. While JAK-STAT signaling can be solved analytically in the adiabatic limit, the long-term dynamics of cytokine accumulation and consumption are complicated by the multitude of feedback regulations triggered in leukocytes. Such rich dynamical complexity will need to be tackled quantitatively and systematically to deliver a more comprehensive understanding of cytokine communication in the immune system [75].

B. Communication across space and time, and cytokine niches

Cytokines allow immune cells to communicate and to modulate their response collectively. In section VA, we will discuss how lymphocyte differentiation can be decided by toggle switches encoded in gene regulatory networks. Most of the positive feedbacks in these toggle switches are in fact associated with a response to cytokines, and are thus intrinsically collective. Tackling the nonlinearities and spatio-temporal aspects of cytokine communication at a more quantitative level is critical to expanding our understanding how antigen recognition leads to such a collective response. We now focus on the dynamics of cytokine secretion and capture in space

and time, over the relatively short timescales (minute to hours) over which cytokines convey information between cells.

1. Cytokine dynamics as diffusion-degradation

The propagation of cytokines within tissues is governed by reaction-diffusion equations which generalize Eq. 24:

$$\frac{\partial c(\vec{r}, t)}{\partial t} = D\nabla^2 c(\vec{r}, t) + \kappa_{\text{prod}}(\vec{r}, t) - \kappa_{\text{consum}}(\vec{r}, t), \quad (25)$$

where $c(\vec{r}, t)$ is the spatio-temporal profile of the concentration of cytokine, D is the diffusion coefficient for the cytokine in the extracellular medium, κ_{prod} and κ_{cons} are volumic rates of cytokine production and consumption respectively. Cytokines are produced by discrete cells, and the production rate $\kappa_{\text{prod}}(\vec{r}, t)$ concentrates on the cell surface. The consumption rate can be estimated by taking into account that cytokines bind tightly to cytokine receptors on the cell surface and get consumed by endocytosis of the engaged receptors. The limiting step for this cytokine consumption is then the binding to receptors, so that

$$\kappa_{\text{consum}} = k_{\text{on}} N_{\text{R}} n_{\text{consum}} c(\vec{r}, t) = k_c c(\vec{r}, t), \quad (26)$$

where n_{consum} is the density of consuming cells. This expression is valid at low concentrations, when cytokine receptors are not saturated and the kinetics of binding remain linear with the concentration of cytokines. It also assumes that consuming cells are uniformly distributed, in a mean-field manner, rather than modelling the precise location of their surfaces. In more general cases, κ_{prod} and κ_{consum} may have additional, nonlinear dependencies on $c(\vec{r}, t)$, for instance because of spatial heterogeneity or receptor saturation. Little can be done analytically in these cases and one must resort to numerical solutions [93]. In the linear regime, Eq. 25 becomes:

$$\frac{\partial c(\vec{r}, t)}{\partial t} = D [\nabla^2 c(\vec{r}, t) - \xi^{-2} c(\vec{r}, t)], \quad (27)$$

with a characteristic length $\xi = \sqrt{D/k_c}$ over which cytokines diffuse before being consumed. In the following section, we focus on the simple homogenous and linear case to derive insights about the quantitative aspects of cytokine communication.

2. Screening by cytokine-consuming cells

Diffusion and consumption of cytokines regulates how the cytokine gradient spreads i.e. the lengthscale for cell communication. It was demonstrated that the short timescale to reach stationary concentration profiles for soluble proteins are explained by the first arrival time of the cytokine ligands, rather than the characteristic diffusion timescale for an individual molecule [94, 95]. This

result implies that the characteristic timescale to reach steady state at a distance r from the secreted cell scales linearly with r rather than with r^2 . The analytical expression for this local relaxation time $\tau(r)$ is [94]:

$$\tau(r) = \frac{1}{2k_c} \left(1 + \frac{r}{\xi} \right). \quad (28)$$

Two relaxation processes are at play in this system: the diffusion timescale ($\tau_d(r) = r^2/2D$) and the reaction timescale (k_c^{-1}). The relaxation timescale is dominated by the reaction time at short distances ($r \ll \xi$) where cytokine consumption dominates. However, at large distances ($r \gg \xi$), both diffusion and consumption are important and the relaxation timescale is the geometric mean of the two, $\tau(r) \sim (\tau_d(r)k_c^{-1})^{1/2} \propto r$. Numerical simulations in the context of immune responses [96] validated this theoretical argument and estimated that it takes only a few minutes for the cytokine gradient to reach its steady state.

The timescale at which immunological regulation by cytokine exchange operates is typically in the range of hours for the activation of gene regulatory elements, to days for the decision to proliferate or to die. The typical timescale at which cells move is also very slow. In the first phase of an adaptive immune responses, antigen-responding T-cells are essentially stuck on their antigen presenting cells. One can thus assume steady state for the distribution profile of the cytokines, $c(\vec{r}, t) = c(\vec{r})$, and solve for $dc/dt = 0$ in Eq. 27. Using spherical symmetry $c(\vec{r}, t) = c(r)$ for the Laplace operator, the steady state profile for the cytokine distribution is

$$c(r) = c(R_{\text{cell}}) \frac{R_{\text{cell}}}{r} e^{(R_{\text{cell}}-r)/\xi}, \quad (29)$$

and the boundary value $c(R_{\text{cell}})$ can be determined by estimating the flux of molecules secreted and consumed at the cell surface:

$$\begin{aligned} \kappa_{\text{prod}} - \kappa_{\text{consum}} &= \oint \vec{j} \cdot d\vec{r} = -4\pi DR^2 \left. \frac{dc(r)}{dr} \right|_{R_{\text{cell}}} \\ &= 4\pi Dc(R_{\text{cell}})R_{\text{cell}} \left(1 + \frac{R_{\text{cell}}}{\xi} \right). \end{aligned} \quad (30)$$

Typically, cells secrete cytokines at a constant rate between $k_{\text{prod}} = 10$ to 1000 molecules per second (e.g. IL-2 and IFN- γ cytokines) and consumption by the secreting cell is given by $k_{\text{consum}} = k_{\text{on}}N_R c(R_{\text{cell}})$. Solving for $c(R_{\text{cell}})$ yields:

$$c(R_{\text{cell}}) = k_{\text{prod}} / (4\pi DR_{\text{cell}}(1 + R_{\text{cell}}/\xi) + k_{\text{on}}N_R) \quad (31)$$

The characteristic scale

$$\xi = \sqrt{D/k_c} = \sqrt{D/(k_{\text{on}}N_R n_{\text{consum}})} \quad (32)$$

is analogous to the screening length of electrostatic interactions and is inversely proportional to the square root of the density of consuming cells. By this analogy,

consuming cells “screen” the diffusion of cytokines from the secreting cells and end up determining the extent of cell-to-cell communication in a dense tissue. In other words, the simple diffusion-consumption of cytokines can account for the heterogeneous accessibility of cytokines in dense tissues and the formation of cytokine “niches,” defined as the portion of space in where cytokines can be sensed.

ξ is controlled by the density of cytokine-consuming cells, and the number N_R of receptors on the surface of consuming cells. In close-packed tissues the density of consuming cells is typically 1 cell every $10\mu\text{m}$, implying $n_{\text{consum}} \sim 10^{-3}\mu\text{m}^{-3}$; for IL-2 D is approximately $100\mu\text{m}^2/s$ (even in close-packed tissues, there is sufficient free space in the extracellular medium for cytokines to diffuse freely as in solution), and each cell consumes typically 10 cytokine molecule per second [96], implying $k_{\text{on}}N_R = 10^{11}M^{-1}s^{-1}$, so that $\xi \sim 25\mu\text{m} = 2.5$ cell diameters. In more physiological conditions, where cells are more dilute, occupying only 10% of the space and cells express lower level of receptors, $n_{\text{consum}} \sim 10^{-2}\mu\text{m}^{-3}$, this length scale rises to $\xi \sim 250\mu\text{m} \sim 25$ cell diameters: there is negligible screening and all the cells within a lymph node or a spleen have access to the cytokine. Note that this 10-fold increase in the screening length ξ implies an increase by 1000-fold in the volume of the cytokine niche and a 1000-fold increase in the number of cells that potentially respond to the cytokine. Such tuneability can be critical to decide which cells (helper T cells or regulatory T cells) win the tug of war for IL-2 and ultimately whether the immune response expands or gets extinguished [77, 79, 80].

3. Probability of autocrine capture

Cells that secrete cytokines often express the associated receptors, and can in principle communicate with themselves. To quantify this effect, we must estimate the non-zero probability P_{auto} that cells capture their own cytokines rather than let them diffuse and interact with neighbouring cells. This phenomenon, called autocrine signaling, is particularly significant when cytokine-secreting cells are surrounded by cytokine-consuming cells and compete for their consumption.

At steady state, there is a balance between cytokine production k_{prod} , cytokine consumption by the secreting cell k_{consum} , and diffusion away from the cell followed by consumption by other cells, as given by Eq. 30. The probability of absorption by the secreting cell is given by:

$$P_{\text{auto}} = \frac{k_{\text{consum}}}{k_{\text{prod}}} = \left[1 + \frac{4\pi DR_{\text{cell}}(1 + R_{\text{cell}}/\xi)}{k_{\text{on}}N_R} \right]^{-1}. \quad (33)$$

We can define the characteristic number of receptors at which this probability reached 50%, $N_R^* = 4\pi DR_{\text{cell}}(1 + R_{\text{cell}}/\xi)/k_{\text{on}} \approx R_{\text{cell}}/\sigma_R$, where $\sigma_R = R_{\text{receptor}}e^{-\Delta G_{\text{assoc}}}$

is the effective cross-section of the receptor-cytokine interaction (Eq. 2 and 6), and where we assumed $\xi \gg R_{\text{cell}}$.

Leukocytes are typically $R_{\text{cell}} = 5$ to $10 \mu\text{m}$ in radius, and $\sigma \sim 0.5\text{nm}$, hence $N_R^* \approx 10^4$. If cells have much fewer than 10^4 receptors (a common situation), $P_{\text{auto}} \ll 1$ and most of the secreted cytokines diffuse away without being captured in an autocrine manner. Alternatively, if cells have more than 10^4 receptors, $P_{\text{auto}} \sim 1$ and secreting cells capture a large fraction of their own cytokines. The functional relevance of autocrine capture remains debated in the field: why would a cell need to respond to the cytokine it secretes. There are also molecular mechanisms limiting its impact. For instance, T cells secreting IL-2 upon antigen recognition express low levels of IL-2 receptors as long as the antigen response lasts to limit autocrine capture [84]). In the tug of war between cell types for cytokine consumption, competition between autocrine and paracrine signaling is crucial, and thus the amount of autocrine capture is significant [80]. Discriminating between autocrine and paracrine signaling is often difficult and it has been suggested that dilution experiments in vitro can be used to distinguish the two modes [97].

4. Size of cytokine niches

Eqs. 29-31 can also be used to estimate the functional size of a cytokine niche established by a single secreting cell. In the limit of no autocrine capture and long dispersal $\xi \gg R_{\text{cell}}$, we have:

$$c(r) = \frac{k_{\text{prod}}}{4\pi D r} e^{-r/\xi} \quad (34)$$

The radius of the niche r_{niche} can be defined as the r at which $c(r)$ falls below the characteristic concentrations K_D needed to trigger cytokine signaling, and is the solution to the implicit equation: $r_{\text{niche}} = \xi \ln(k_{\text{prod}}/4\pi r K_D)$. The logarithmic dependence on the rate of cytokine secretion k_{prod} is significant because that rate may vary greatly depending on the cytokine or depending on the state of differentiation of cells. For example, CD8+ T cells can secrete up to 1000 IFN- γ per second, CD4+ T cells secrete 10 IL-2 per second upon activation from a naive state, and 1000 IL-2 per second as they mature [84] or when get activated from a memory state [98].

In Figure IV B 4, we present the numerical solution for r_{niche} for different secretion rates for cytokines with different screening lengths. Note how variable the volume of the cytokine niche can be, based on realistic parameters of immune responses.

To conclude, in this section we discussed how simple equations for the diffusion and consumption of cytokines in dense immunological tissues can quantitatively account for the tunability of cell-to-cell communication in the immune system. It is particularly striking that such active modulation of cell communication plays a

functional role in modulating immune responses. In the first moments of activation, there are very few cytokine-consuming cells, and the field of cytokines extends to the entirety of the lymphoid organ ($\xi > 100\mu\text{m}$), reaching blood vessels and making communications global. Within 24 hours, many cells migrate closer to cytokine-secreting cells. As they do so, they increase their cytokine consuming capabilities by upregulating their receptors, thereby “screening” the diffusion of cytokines ($\xi \sim 10\mu\text{m}$) and limiting the extent of cell communication to nearest neighbours, within a tight niche around secreting cells.

V. CELL FATE

There is a large variety of types and subtypes of immune cells, making the field of immunology sometimes difficult for the non-specialist. These fates are acquired either at the very beginning during haematopoiesis, the process by which all types of blood cells are generated by differentiation from stem cells, or following an immune response or challenge, during which cells further specialize to better fight infections. We first describe the physical mechanisms of gene regulation that lead to stable distinct cell fates. We then review modeling strategies for understanding and inferring the hematopoiesis differentiation tree. Finally, we discuss the problem of differentiation during an immune response.

A. Gene regulation and cell differentiation

Many cell decision making processes in the immune system are carried out at the level of gene regulation, whereby expression of key transcription factors decide cell fate and subsequent immune effector function. For example, T-cells sense their inflammatory environment as defined by the combinations of cytokines in the extracellular environment, and drive their signaling response towards expressing key transcription factors that in turn decide which cytokines get produced and secreted.

To consider a concrete example of gene regulation, we use the classical example of the differentiation of CD4+ helper cells, a particular type of T cells whose function is to regulate the action of other immune cells through the production of cytokines — the messenger molecules used by immune cells to communicate with each other. In their naïve state, T cells are essentially a blank slate which can later differentiate into more specialized cell states to best match the pathogenic threat. Upon antigen stimulation, helper T cells can orchestrate three types of immune responses. They can elicit a ‘Th1’ response, further unleashing T- and NK cells’ cytotoxic response against infected cells. This is particularly relevant when eradicating an intracellular infection such as a viral infection. Alternatively, helper T cells can elicit a ‘Th2’ response to

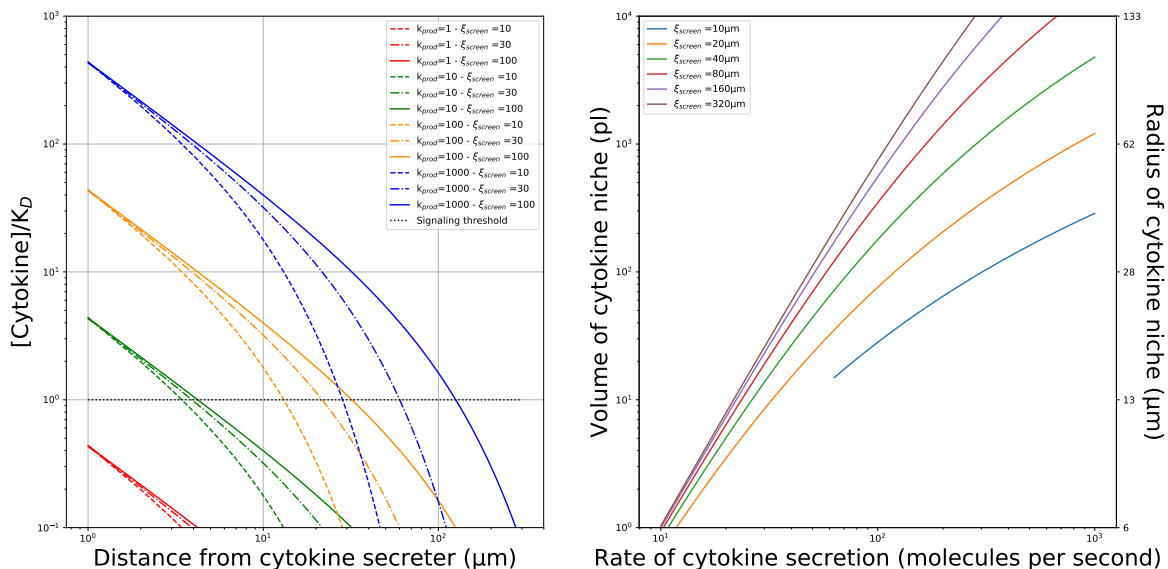


FIG. 6. **Size of cytokine niches for different parameters of secretion/consumption.** (Left) Cytokine profile as a function of the distance from the surface of the secreting cell (normalized by K_D – the critical cytokine concentration to induce STAT phosphorylation), for different rates of secretion k_{prod} (in molecules per second), and different screening length ξ (in μm). (Right) Size of the cytokine niche as a function of the rate of cytokine secretion. For these graphs, $K_D = 3\text{pMol}$ and $R_{cell} = 5\mu\text{m}$ in equation (34).

engage with B cells, drive affinity maturation and class switching in antibody production towards secretion of antibodies that can annihilate extracellular (e.g. bacterial) infections. Finally, T-cells can elicit a ‘Th17’ response that protects mucosal barriers from infection. These three characteristic types of T-helper cell differentiation respond to and amplify distinct cytokine environments. Note that similar decisions are made by other immune cells: for example macrophages differentiate into ‘M1’ (pro-inflammatory) and ‘M2’ (anti-inflammatory) types, which mirror the decisions and inflammatory outputs associated with the Th1 and Th2 helper T-cell types.

Immunologists are interested in dissecting the interplay and feedback between inflammatory environments and immune cell differentiation [99]. In that context, quantitative immunology brings the tools of nonlinear dynamics to explain how sharp and cross-inhibitory decisions can be made by activated T-cells during differentiation. Many immunological systems involve nonlinearities with self-reinforcing feedback loops, whereby the protein of interest acts as a positive feedback that drives further expression. Additionally, most gene regulation involve multimerization of transcription factors onto one gene locus, as in the case with phosphorylated STAT, a transcription factor induced by cytokine signaling. Multimerization implies a nonlinearity because the association rate of an n -mer scales with the n^{th} power of the monomer’s

concentration. These nonlinearities and their functional consequences for cell fate in the immune system can be studied using nonlinear stability analysis. Many transcription factors auto-amplify their production in a direct loop or, more generically, in a more convoluted cytokine-mediated manner. For example, the transcription factor Tbet drives the secretion of a cytokine called interferon γ (IFN- γ) that signals through pSTAT1 to produce more Tbet in the context of Th1 T-cell differentiation. In section XII B 1 of the Methods, we recall the basic formalism of nonlinear analysis, which explains the generic occurrence of bimodal distributions of expression of transcription factors, cytokines and surface markers that immunologists encounter in their single-cell measurements (cytometry or single-cell transcriptomics). We discuss classical models of gene regulation and feedback regulation that account for the molecular programs enforcing sharp cell differentiation in the immune system.

1. Th1/Th2 differentiation

Early experimental evidence demonstrated that CD4+ helper T-cells commit to two distinct and incompatible states of differentiation, Th1 and Th2, upon antigen activation and response to inflammatory cues. A model describing this system is based on two inhibitory loops

(see Sec. XII B 1) whereby transcription factor P_1 inhibits the production of P_2 and vice versa (Figure 7A & B).

$$\begin{aligned}\frac{dR_1}{dt} &= v_1 \frac{K_2^{n_2}}{K_2^{n_2} + P_2^{n_2}} - \gamma_{R_1} R_1 \\ \frac{dP_1}{dt} &= k_1 R_1 - \gamma_{P_1} P_1 \\ \frac{dR_2}{dt} &= v_2 \frac{K_1^{n_1}}{K_1^{n_1} + P_1^{n_1}} - \gamma_{R_2} R_2 \\ \frac{dP_2}{dt} &= k_2 R_2 - \gamma_{P_2} P_2,\end{aligned}$$

where R_1 and R_2 are number of mRNA leading to the production of P_1 and P_2 , γ decay rates, k_i and v_i are production rates, K_i are binding dissociation constants, and n_i are Hill coefficients. At steady state,

$$\begin{aligned}P_1 &\propto \frac{K_2^{n_2}}{K_2^{n_2} + P_2^{n_2}} \\ P_2 &\propto \frac{K_1^{n_1}}{K_1^{n_1} + P_1^{n_1}},\end{aligned}\quad (35)$$

to which some background expression may be added to account for leaky gene regulation.

The nonlinearities, which are controlled by the Hill coefficients $n_i > 1$, originate from the multimeric control of transcription by the transcription factors. A graphical inspection of the nullclines (Eq. 35) for this dynamical model (Figure 7.C) reveals that there exist three fixed points for this gene regulatory system, two stable ones, (P_1^{high}, P_2^{low}) and (P_2^{high}, P_1^{low}) , and one unstable fixed point (P_1^{high}, P_2^{high}) . This simple dynamical system reveals how cross-inhibition of transcription factors yields a classical toggle-switch that yields two incompatible states. Note that such a model predicts hysteresis in T cell differentiation (Figure 7D) as has been observed experimentally [100–102].

2. Other differentiation switches

A similar formalism has been applied to account for the differentiation of CD4+ T-cells in Th17 and regulatory T-cells (Treg). Th17 cells are pro-inflammatory and critical to eradicate infections in mucosal tissues; regulatory T cells are anti-inflammatory and limit overzealous immune responses, such as auto-immune disorders caused by a strong response against self antigens, or runaway immune responses such as septic shock. These two states of differentiation for CD4+ T cells can be accessed upon exposure to the cytokine TGF β . However, that cytokine has been shown to be both pro- or anti-inflammatory depending on the context. Tyson & coworkers [103] proposed a model to account for such bimodality in TGF- β action. In that model, the concentration of each protein species P_i generically follows:

$$\frac{dX_i}{dt} = \gamma_i \left(X_i^{eq} \frac{1}{1 + e^{-x}} - X_i \right) \quad (36)$$

where $\omega_i = \omega_i^0 + \sum_j \omega_{j \rightarrow i} P_j$, P_i^{eq} is the characteristic concentration of P_i in differentiated cells at steady state, γ_i are degradation rates, ω_i^0 is the constitutive level of expression of regulation of X_i , $\omega_{j \rightarrow i}$ is the (possibly negative) influence of X_j on X_i , and σ_i is a parameter modulating the steepness of the regulation. Note that proteins can regulate their own production, $\omega_{i \rightarrow i} \neq 0$.

There are three main species involved in Th17/Treg differentiation: P_1 =FoxP3, P_2 =ROR γ t (the transcription factors controlling Treg and Th17 differentiation, respectively), and P_3 =TGF β (a cytokine that regulate T cell differentiation), which acts as an externally fixed stimulus. Their interactions are represented in Figure 8A. For instance ROR γ t⁺ is activated by TGF β ($\omega_{3 \rightarrow 2} > 0$), as well as by itself ($\omega_{2 \rightarrow 2} > 0$), and but is repressed by ROR γ t ($\omega_{1 \rightarrow 2} < 0$).

Graphically, one can represent the nullclines ($dX_i/dt = 0$) for both transcription factors as shown in Figure 8. There are three stable fixed points where the nullclines meet: ROR γ t⁺FoxP3⁻ for the Th17 cells, FoxP3⁺ROR γ t⁻ for the Treg cells and ROR γ t⁺FoxP3⁺ for a mixed state whose existence was later confirmed experimentally [103]. This formalism is interesting because it handles coupling and feedback phenomenologically, through the effective regulation parameters $\omega_{i \rightarrow j}$, rather than through detailed mechanisms of how these regulations are implemented.

3. Experimental test of bistability in cell differentiation

The elegant models presented in the two previous subsections (Th1/Th2 and Th17/Treg differentiation) account qualitatively for the observation that many activation conditions lead T-cells to commit to clear, separate and stable cell fates. The crux of these models is the existence of toggle switches based on cross-inhibition of transcription factors and, optionally, self-reinforcing positive feedback loops for transcription factors. Similar models have been proposed in the context of hematopoietic differentiation [104, 105], and they provide a quantitative framework to account for the bimodal distribution of transcription factors in immune cells. How do these models compare with experiments?

A theoretical inspection of the Th1/Th2 model of T-cell differentiation by Antebi *et al.* [102] highlighted how the bistable solution of the model presented in Figure 7 does not generically produce a bistable solution. A simple scan of parameters within the physiological range for induction of Tbet and GATA-3, the candidate P_1 and P_2 in Eq. 35, demonstrated that the ratios of rates for their self-reinforcement and inhibition most commonly predicted a single stable fixed point, and only rarely more than two. In addition, the induction of transcription factors is known to be inherently stochastic, because of intrinsic noise in gene induction [106–109] or extrinsic noise in the levels of expression of signalling molecules in T cells [110]. Overall, the “simplest” model

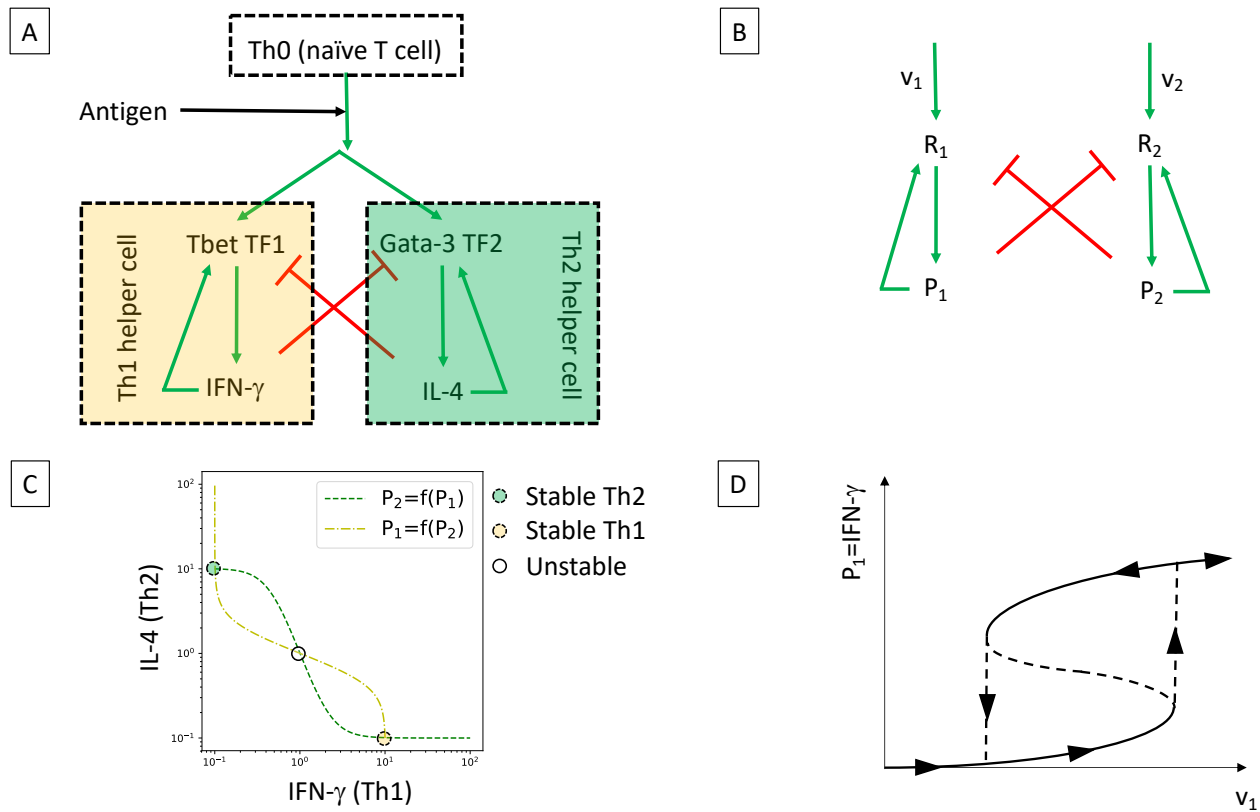


FIG. 7. Modeling Th1/Th2 differentiation. A. Cartoon model of the cross-regulation of Tbet and GATA-3 transcription factors (TF) whose expression determines T cell fate in terms of Th1/Th2 differentiation. B. Diagram of the interactions of TF deciding cell fates. C. Nullcline for the regulation of TF (P_1 and P_2): three fixed points can emerge in the dynamics of the system, two stable, one unstable. D. Dynamics of expression of P_1 when the tonic rate of expression v_1 (driven by antigen response) is increased: note the hysteresis in the levels of P_1 whereby a bimodal distribution can be anticipated for intermediate values of v_1 .

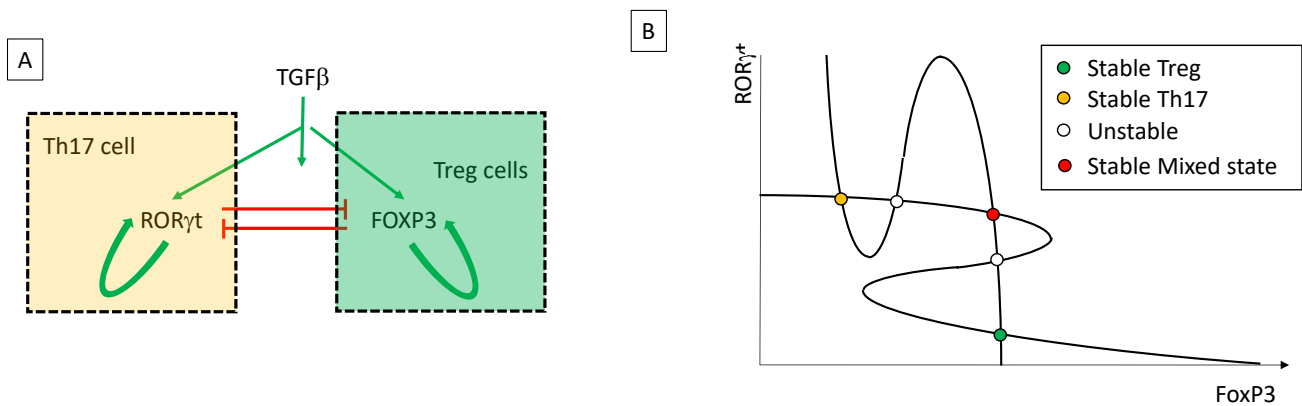


FIG. 8. Modeling the trimodal distribution of T-cell fate upon exposure to $TGF\beta$. A. Cartoon of the genetic network associated with the regulation of $ROR\gamma t$ and $FoxP3$ upon $TGF\beta$. B. Nullclines for the dynamics of regulation of $ROR\gamma t$ and $FoxP3$. Each intersubsection of the nullclines corresponds to a fixed point; graphical considerations can be used to deduce the stability of each fixed point.

of a toggle switch in T-cell differentiation, though elegant in its inception, is not consistent with observa-

tions. This case is typical in quantitative immunology, where initial insight is often derived from genetic manipulations with severe consequences as in the case of the Th1/Th2 system where knocking out Tbet leads to a reduction/maintenance of induction of Th1/Th2 cytokine secretion, respectively [111]. These knockout experiments must be balanced and reinterpreted in non-perturbative settings, e.g. by monitoring the spontaneous induction of Tbet in differentiating T-cells. Similar reassessments are happening in many models of leukocyte differentiation, such as in the M1/M2 model of macrophage differentiation [112], or PU-1-controlled differentiation in the myeloid compartment [113]). Future work will require better parametrization an innovative modeling approach that can embrace the combinatorial and dynamic complexity of cytokine communications and gene regulation in the immune system.

Such “murkiness” in immunology must be embraced as it opens up functional possibilities. For example, Peine *et al.* [114] tracked the formation of Tbet⁺GATA-3⁺ mixed phenotypes, i.e. with characteristics of both Th1 and Th2, so in apparent contradiction with bistability. Their presence was demonstrated to limit the deleterious impact of all-out, non-mixed Th1 or Th2 inflammation. Thus, a better understanding of noise and stability in immune cell differentiation will be key to understanding virtuous and pathologic inflammation.

B. Hematopoiesis

1. Timescales

Where do immune cells come from? A human being has $\sim 5 \cdot 10^{13}$ cells in the body and more than half of them are made by hematopoietic stem cells (HSC) found in the bone marrow. There are an estimated order of magnitude $\sim 10^{16}$ cell divisions per human, which gives $\sim 10^6$ cell divisions per second [115, 116], and most of them are linked to the HSC cells. HSC number $\sim 10^4$ in mice (there are no reliable numbers for humans), many of them already made in the fetus. Through a series of differentiation and phenotypic commitment events (HSC \rightarrow Short Term HSC (ST-HSC) \rightarrow Multipotent progenitor (MPP), these cells give rise to different kinds of cells found in the blood, including red blood cells and all immune cells (see Fig. 9). The first branching decision, whose precise timing is currently being questioned, is about becoming a myeloid progenitor, that will give rise to red blood cells, mast cells, thrombocytes (so not immune cells) but also macrophages, granulocytes (neutrophils, basophils and eosinophils) and dendritic cells – so cells of the immune system that eat up cells and proteins non-specifically. They function as cells of the innate immune system, simply eliminating bacteria and other pathogens, but some of them (e.g. dendritic cells) also play an important role in adaptive immunity as antigen presenting cells (APC). The remaining branch of the differentiation tree leads to




FIG. 9. **Hematopoiesis.** A cartoon representation of the outline of the hematopoiesis process that leads to the formation of immune cells. The cells that are discussed in some detail in this review are marked in red.

lymphoids that include cells of the adaptive immune system: B and T-cells as well as natural killer (NK) cells. The decision about becoming a myeloid or a lymphocyte is widely assumed not be an autonomous decision but is believed to be influenced by external signal (see below for a detailed discussion).

General quantitative questions about differentiation apply to the differentiation of cells in the immune system. For example, given that in mice there are of the order of 10^4 HSC, what kind of dynamics results in $\sim 10^7$ very short lived granulocytes (of the order of a day) while at the same time producing $\sim 10^{11}$ lymphocytes that can live even for years? How long does it take to produce these cells? What kind of differentiation process can produce this kind of diversity, and how is it regulated to produce the right numbers of cells? A lot of information about these differentiation processes has been gained from two types of experiments. The first involves exposing cells at a given upstream stage with died or radioactive markers (i.e. bromium, deuterium) that get taken up — and then diluted — upon cell division. Analysis of the decay curves is informative about cellular lifetimes. The second type of experiment — an adoptive transfer experiment — is experimentally harder, since it involves transplanting new marked cells into an animal (typically mouse) and then tracking them. To study cells during the early stages of hematopoiesis the mouse must first be cleared of its natural cells, and recent results suggest that the dynamics of differentiation after transplantation may be very different from regular dynamics [117]. Lymphocyte divisions can be studied without killing the host’s own immune system. The principle of these experiments is simple — after some time δt the marked cell will appear in a given compartment.

At the population level, given that cells at stage i of differentiation can proliferate with rate λ_i , differentiate with rate α_i or die with rate δ_i , the dynamics of a given cell type n_R in the differentiation process follows [117, 118]

$$\frac{dn_R}{dt} = \alpha_u n_u - (\alpha_R + \delta_R - \lambda_R) n_R, \quad (37)$$

where n_u describes the upstream (pre-differentiation) cell

type. Cells can also migrate, which introduces a spatial component to the equations. In general, all these rates are functions of the concentrations of the different populations n_i , resulting in non-linear equations. $\alpha_R + \delta_R - \lambda_R = k_R$ define a collective timescale of decay of the n_R population. This highlights the identifiability problem that without specially planned out experiments guided by theory, it may be hard to tease apart the timescales of biological interest.

Commonly in this kind of analysis it is assumed that the cellular differentiation of the adult individual is in steady state. We will return to this assumption when we learn the timescales for the process. Within this assumption the steady state ratio of cells in the compartment of interest, compared to the upstream compartment is $n_R/n_u = \alpha_u/(\alpha_R - \beta_R)$, with $\beta_R = \lambda_R - \delta_R$. This kind of approach has been used to learn proliferation rates for T-cells [119]. Stochastic versions of such models have also been considered in detail [119–121], showing that theories of age-structured cell populations [120] are equivalent to branching processes [119] (see section XII C 6). For hemopoietic stem cells, a HSC in the bone marrow divides once every 50 days ($\lambda_R = 1/110\text{days}$) and a HSC becomes a granulocytes after ~ 1 year (and then lives for a day). This is known from so-called fate-mapping experiments, consisting of tracking labeled cells, which is easier for granulocytes since they are known not to proliferate. It is also known that granulocytes have no other cell source, unlike other cell types. In this case the label from the HSC appears in a granulocyte after a year. Since a lifetime of a mouse is ~ 2 years, a 1 year timescale to produce a new cell may put the steady state assumption into question. The fetal dynamics is very different, the steady state assumption cannot be used, and a new granulocyte is produced in a few weeks [117].

Technically, measuring cell counts directly is not reliable because of sampling bias. Instead, one should measure cell ratios, $f_i = n^i/n^{\text{total}}$, as was shown in the previous example. In steady state, the fraction of cells of a given type evolves according to $df_i/dt = k_i(f_{i-1} - f_i)$, with $k_i = \alpha_i - \beta_i$. By looking at subsequent compartments it is possible to disentangle differentiation and proliferation. This gives an estimate of how *self-renewing* the compartment is (β_i), and how much comes from upstream differentiation (α_i). For a perfect stem cell we expect $k_i \rightarrow 0$. Real stem cells (e.g. HSC) are essentially self-renewing with τ (in mice) on the order of the lifetime of the mouse. Other differentiation compartments are clearly transitional, with $k_i^{-1} \sim 1$ day. HSC divide once every 110 days but they also differentiate once every 110 days, so they seem to have a lifetime of (they replace themselves every) 110 days (this remains to be verified with more direct means). Short Term HSC (ST-HSC, see Fig. 9) die at a rate of 1 per month and Multipotent progenitor (MPP) of 1 per day. This produces a strong self-renewal gradient from HSC to MPP. At the same time the proliferation rates of these cells go up from HSC to MPP with a similar gradient. Some have speculated that

these inverse trends of proliferation and self-renewal protect the organism from malignancy by making sure the right amounts of cells are produced. Myeloid lymphomas — a condition where the organisms produces too many myeloid cells — have been linked to disturbances in this balance and the condition is only visible in aged individuals when the overproduced cells have accumulated.

Most HSC become myeloid cells, with only one in 100 cells becoming a lymphoid cell. Lymphoid cells are very long lived and the thymus can function for long time without input from new HSC. In general, due to the slow timescales of the whole process, the system has a lot of inertia. If all HSC are killed in a young adult mouse, there are no visible phenotypic effects for ~ 5 months, unless the hemopoietic pathway itself is stressed [122].

There are *a priori* no biological reasons for downstream cells in the hematopoietic pathway not to influence upstream cells, however this has not yet been directly observed. Probably certain feedback mechanisms are at play.

C. Inferring the hematopoiesis differentiation tree

Analyzing experiments from mice with progenitor cells carrying unique barcodes with a maximum likelihood approach (see section XII E 1) coupled to a stochastic model, Perié et al [123], uncovered a more complex picture of myeloid and lymphoid differentiation. They considered the lymphoid primed multipotent progenitors (LMPPs) differentiation process during which each LMPP is called an MDB, since it has the potential to become a myeloid (M), dendritic (D) or B cell. They built a stochastic branching process model of cell differentiation (see section XII C 6), where in each step of the decision tree a cell can lose one or two of its potentialities: an MDB can become an MD, MB, DB or directly an M, D, B. Similarly an MD can become an M or a D etc. They contrasted this with a traditional differentiation model, where the MB phenotype is forbidden, and an MDB cannot lose two potentials at once. Using mice with progenitor cells carrying unique barcodes, they build lineage trees, traced each offspring cell to its parent and counted how many offspring of each type (N_i for $i \in \{M, D, B\}$) were produced by each MDB cell. The generating function $G(z_M, z_D, z_B) = \sum_{m,d,b} z_M^m z_D^d z_B^b P(m, d, b)$ for the probability of a cell to give rise to a given combination of offspring cells $P(N_M, N_D, N_B)$ can be calculated from a convolution of simple branching process generating functions [123]):

$$G(z_M, z_D, z_B) = H(p_{\text{MDB}}, G_{\text{MDB}}(z_M, z_D, z_B)), \quad (38)$$

where $H(p, z)$ is given by the generating function, in Eq. 130 in section XII C 6, and

$$G_{MDB}(z_M, z_D, z_B) = \sum_{i \in \{M, D, B\}} p_{MDB \rightarrow i} z_i + \sum_{i, j \in \{M, D, B\}, i \neq j} p_{MDB \rightarrow ij} H(p_{ij}, G_{ij}(z_j, z_j)), \quad (39)$$

$p_{MDB}, p_{MB}, p_{MD}, p_{BD}$ are the probabilities of not differentiating from (staying in) the MDB, MB, MD, BD states and $p_{MDB \rightarrow ij}$ describes the transition probabilities from the MDB state to the $ij \in \{MD, MB, DB\}$ states. The remaining functions are defined recursively as in Eq. 39, e.g. for $ij = MD$: $G_{MD}(z_M, z_D) = p_{MD \rightarrow M} z_M + p_{MD \rightarrow D} z_D$ etc.

The marginals of this distribution give the probability that a given barcode appears in each of the 7 possible cell types. Following the equal loss of potential model (ELP) [124, 125], the rate of loss of the cells' ability to make different cell types was assumed to be constant and independent of current cell type, which gives three rates $\alpha_M, \alpha_B, \alpha_D$, that parametrize the transition probabilities between the cell types $p_{i \rightarrow j}$ (for all allowed combinations of $i, j \in MDB, MD, MB, DB, M, D, B$) and reduces the number of parameters. The parameters are determined by carrying out a maximum likelihood fit $\vec{p}_{\max} = \text{argmax}_{\vec{p}} (\sum_i n_i \pi_i(\vec{p}))$, where n_i is the number of observed barcodes in state i and $\pi_i(\vec{p})$ is the model probability (from Eq. 38) that an ancestral barcoded cell gives rise to offsprings in states i , where $i \in MDB, MD, MB, DB, M, D, B$ calculated from the generating function: $\pi_{MDB}(\vec{p}) = \sum_{m, d, b \geq 1} P(m, d, b)$, $\pi_{MD}(\vec{p}) = \sum_{m, d \geq 1} P(m, d, b = 0)$, $\pi_M(\vec{p}) = \sum_{m \geq 1} P(m, d = 0, b = 0)$ etc.

This approach showed that the classical model of sequential loss forbidding the MB state is not consistent with data. In fact, the MB state is very unlikely but is essential to explain the observed barcode distribution. Simulation of the dynamics of these models showed that the differentiation process between LMPPs and the committed cell types requires ~ 20 rounds of differentiation, which translates into about 2 weeks. Recent barcoding experiments combined with similar inference methods that looked at earlier stages of differentiation from HSC cells [126] showed that most HSC cells give rise to many cell fates.

Similar techniques of branching processes coupled to likelihood inference methods were initially used by Yates et al [121], to analyse fluorescent dye experiments of differentiating T-cells and estimate cell division and death rates.

D. Cell fate during the immune response

1. Choice and timing of cell fate under stimulation

Lymphocytes, both T and B-cells, start growing and dividing upon stimulation during an infection. When they divide they also acquire new specialized functions, make many decisions at the individual cell level, e.g. whether to divide further, whether to switch class for immunoglobulins (Ig) molecules, whether to become memory cells for both T-cells and B-cells, or to become plasmablasts for B-cells. As their number increases, they eventually make another decision to stop dividing and start dying. At the population level we observe a decrease of the population size, after the large expansion due to proliferation, and a return to pre-infection cell counts. Since each cell can in principle adopt a different decision path, at the population level we see a large combinatoric diversity of cell states: from cells that have divided a small number of types and become memory cells, to cells that have divided many types to leave no offspring.

This diversity is even more staggering for B-cells, since they undergo a phenomena called class (or isotype) switching recombination (CSR), where the variable region (discussed in Sec. VI) of the antibody does not change but the constant region of the receptor is modified. This process does not change the affinity of the antibody for the antigen, but it means the antibody cannot interact with signaling molecules (see sec. III B). B-cells do this by expressing a specific gene located in the heavy chain locus, organized in a specific order. During CSR, segments of genes in this locus are removed, and the remaining DNA is recombined to encode a different isotype. Since the double stranded breaks occur at conserved nucleotide motifs, the identity of the expressed isotype-encoding genes are conserved between cells and individuals. Since the non-expressed isotype-encoding genes are deleted from the locus, in general the order of cycling through the classes is conserved, although inter-chromosomal translocation from the other allele can add deleted loci [127].

How are these decisions made by cells? One view is that cells integrate signals from the environment, which act as cues to trigger decisions. While this is certainly true [128, 129], and experimentalists can induce certain cell fates in vitro using cytokine cocktails (see sec. IV), an alternative (yet not contradictory) idea was proposed in the "cyton" model [130]. The basic idea behind this model is that each cell in the population makes a stochastic decision, choosing from one of the accessible cell fates. To do this it uses an intrinsic pre-programmed timer, telling it when it should divide and die. However since the timer in each cell picks a time from a distribution of times, each cell will go through a different decision scenario. Putting together the different decisions each cell can make this leads to a heterogeneity of cell fates in one population. Yet, for each trait (division times, becoming a memory cell), the population level distribution is reproducible.

The model can be and has been adapted to different cells [130–133] but let us present it on the simple example

of three possible options: to divide, to die or do nothing even if stimulated. Given an experimentally characterized distribution of division time $\phi_i(t)$ and death times $\psi_i(t)$ for each division round i , and assuming that a fraction F_i of cells will divide when stimulated in division round i , the number of cells dividing or dying per unit time is a given round $i = 1, \dots, m$ of divisions is:

$$n_i^{\text{div}}(t) = 2pF_i \int_0^t dt' n_{i-1}^{\text{div}}(t') \left(1 - \int_0^{t-t'} dt'' \psi_i(t'') \right) \phi_i(t-t'),$$

$$n_i^{\text{die}}(t) = 2 \int_0^t dt' n_{i-1}^{\text{div}}(t') \left(1 - F_i \int_0^{t-t'} dt'' \phi_i(t'') \right) \psi_i(t-t'),$$

where the terms in parantheses account for cells that would have divided (died) at time t but had previously died (divided). The total number of cells in each division round is then calculated by summing gains and losses in each round:

$$N_0(t) = N - \int_0^t dt' (n_0^{\text{div}}(t') + n_0^{\text{die}}(t')), \quad (40)$$

$$N_i(t) = \int_0^t dt' (2n_{i-1}^{\text{div}}(t') - n_i^{\text{div}}(t') - n_i^{\text{die}}(t')),$$

for $i = 1, \dots, m$. These integrals can be performed analytically for exponential distributions but since the experimental data are better described by log-normal distributions [134, 135], the integrals are performed numerically.

The cyton model and its generalizations have been shown to be able to fit the number of cells in experiments with and without stimulation [130]. Of course, one could note that the cyton model takes the experimentally observed distribution of division and death times and simply calculates the number of surviving cells. Yet the power of the model lies in the paradigm shift that assumes there are rules and, while the immune system functions stochastically and relies on heterogeneity, we can nevertheless predict and understand its behaviour. The idea of a programmed stochastic system, as opposed to a black box integration of cues, is extremely powerful.

Using data from transgenic mice infected with the influenza virus, Marchingo et al [131] showed that T-cell receptors and costimulatory signals impose an intrinsic division fate on each cell, and cells then count through generations, as defined by the cyton model, before returning to the pre-stimulation state. This initial heritable priming can be later modified by dose dependent cytokine signaling, which is also integrated by the cells and influences their the response. As a result different combinations of costimulatory signal and cytokines can generate signals of similar magnitude. These experiments also show that the cyton model is not in contradiction with signaling-based decision making.

2. An aside on cell types

Cell types in immunology are defined using surface markers. Cells communicate by binding and unbinding

signaling molecules (cytokines), for which they need specialized receptors. A cell that commits to participate in a given communication channel, expresses a surface receptor and the combination of surface receptors gives the spying experimenter an idea of how this cell is bound to behave when triggered – this is what we call a cell type. This phenomenological approach has been very successful in the history of immunology and provides a way to classify different cells and give us some idea about their function and properties. In practice, FACS (Fluorescence-activated cell sorting) sorting experiments segregate cells according into ones that have a high concentration and a low concentration of a given surface marker on their surface. By doing this in many dimensions, one can zero in on a very specific cell type. Technically, this leaves the problem of deciding where the boundary between high and low is – in practice cells are heterogenous and the experiment produces a distribution of marker concentrations, which, if we are lucky, is bimodal. This problem of finding the separatrices between cell types is called *gating*, and is an art implemented in analysis softwares. Advances in machine learning will alleviate the need for expert-guided gated and make the definition of cell types more automatic and easy-to-validate. Being aware and using the heterogeneity of the population data is often useful. Looking at the population distribution has made people realize that often there is continuum of phenotypes and a binary approach is not valid. This has been amplified by recent high throughput cell sorting experiments by mass cytometry –so-called CYTOF [136], which make it clear that cell types are more continuous and dynamic than traditionally assumed. With that in mind, we can now try to learn something about how cells acquire and switch identity.

3. Inferring cell fate timelines during the immune response

The problem of inferring the timeline of cell differentiation during T-cell response can also be treated using the cyton theory [130–133]. Here we present another more data-driven approach to the problem described in Buchholz, Flossdorf et al [137]. We will concentrate on the example of CD8 T-cells. A clone, here defined as all the TCR that respond to the same antigen, starts with ~ 10 – 100 long-lived naive cells (lifetime of about 80 days in mice). During an infection this clone grows to $\sim 10^7$ short-lived effector cells. After infection clearance, there remains 10^3 memory cells with an intermediate lifetime (about 15 days in mice). These memory cells can then respond more robustly in a subsequent infection. Memory and effector cells also proliferate within their classes to maintain their pool. This is done through small antigenic signals they constantly receive for memory, and strong antigenic signals for effector cells.

The above textbook description provides a simplified picture, but is unable to put measurable numbers behind each of the described populations. More importantly,

this picture was painted by looking by eye at bulk population data summing over many cells. Can inference and data discrimination approaches help us infer the differentiation pathways? Population-level adoptive transfer experiments can be used to infer the parameters of known models, but it is much harder to infer the topology of the network, especially from single snapshots. Flossdorf et al [137, 138] analyzed single cell adoptive transfer (single cell fate mapping) experiments to discriminate between different differentiation tree topologies and find the underlying model parameters. They considered a model with four cell types: naive (N), effector (TEF), effector memory precursor (TEMp) and central effector memory precursor (TCMp) and wrote down the most general network diagram, in which every non-naive cell type could differentiate into any other cell type, all of which also could proliferate and further differentiate. Naive cells could become any of the other cells, but no other cell could become a naive cell. This led to 302 possible models, defined by their network topologies, as well as their parameters. The precise definitions of the particular cell types in the experimental data was done using surface markers (see discussion above) can be found in the original paper. The immunological meaning of these cell types is under debate – but we are mainly interested here in the power of the theoretical method.

The considered models where simple stochastic dynamics for each cell type that considered differentiation ($\vec{\alpha}$) and proliferation ($\vec{\beta}$) reactions. As mentioned above, a lot of heterogeneity was observed within naive cells and in vitro experiments show heterogeneity in the reaction rates, which had implications on the model. Differentiation rates were chosen from underlying distributions: an exponential and gamma distribution of differentiation rates were considered. The data resolution did not allow to discriminate between the two, with an exponential giving a sufficiently good fit. The proliferation rates were taken to be fixed, since longitudinal correlations of proliferation events were previously shown to decay rapidly [139]. The considered models ignored the heterogeneity in the phenotype of naive cells, but since the heterogeneity in the number of naive cells would render model selection impossible, the experiment was tuned to start with one naive cell ($N_{\text{naive}} = 1$).

The model consisted of a Master equation for the evolution of the joint probability of the different cell types:

$$\partial_t P(\vec{N}) = F \left(P(\vec{N}), \vec{\alpha}, \vec{\beta}, n_{\text{naive}}, \vec{r} \right), \quad (41)$$

where is F a linear operator describing the proliferation and differentiation dynamics of the different cells types, $\vec{N} = (N_{\text{TEMp}}, N_{\text{TCMp}}, N_{\text{TEF}})$ are the numbers of cells of each type, and \vec{r} is the death rate vector of all cell types. The analytically calculated first and second moments of these equations are fit to the variances ($CV_i(t)$) and covariances ($\Sigma_{ij}(t)$) and mean concentrations ($\langle N \rangle_i(t)$) for each cell type i from cell fate mapping data using a mean squared error minimisation. The

model selection is then finalized by comparing the AIC (see section XII E 1) of the best fit for each topology. The necessity to use correlations in the inference procedure, and not just population averages, was demonstrated on synthetic data generated with known parameters and topologies. For instance, both a linear ($A \rightarrow B \rightarrow C$) and tree ($C \leftarrow A \rightarrow B$) topology can give rise to the exact same evolution of the mean number of cells as a function of time, for well chosen parameters. However covariances between these numbers give clearly different signatures, allowing us to distinguish the different topologies. These data-based, model-selection approaches are similar to those developed for stem cell hemapoiesis [123], described in section V C.

As a result, from the 302 possible models, two strongly stand out as potential candidates to describe the cell differentiation of CD8 T-cells responding to a listeria expressing chicken ovalbumin (to which these T-cells were specific) in mice. These two models are:

$$\begin{aligned} \text{naive} &\rightarrow \text{TCMp} \rightarrow \text{TEMp} \rightarrow \text{TEF}, & (42) \\ \text{naive} &\rightarrow \text{TCMp} \rightarrow \text{TEMp} \rightarrow \text{TEF} \ \& \ \text{naive} \xrightarrow{10\%} \text{TEMp}, \end{aligned}$$

where each cell type except for the naive one proliferates. The models have well defined parameters that can easily be identified. This bacterial response is an acute infection and generates exponential cell growth, and the inferred proliferation rates form a gradient from TEF \rightarrow naive cells with naive cells proliferating the least. Both of these models correctly predict the time dependent dynamics of the mean concentrations of the different cell types. Even if learned on single-cell progeny data at only one time point, the first linear diversification model is able to predict the phenotypic composition of the expanding population at earlier time points. These kinds of approaches shed light on the possible differentiation dynamics and allow us to rule out possible scenarios. To give just one example, for this system the cell differentiation is largely ($\sim 90\%$) symmetric: the two offspring of a cell are usually the same, although not the same as the mother, as opposed to asymmetric divisions where the two daughter cells have different cell types which happen about $\sim 10\%$ of times. Also, it turns out that the differentiation and proliferation rates must depend on time to explain the data. These time-dependent rates are the same in a response to both bacterial (listeria) and viral challenges, suggesting universal response dynamics.

Using similar methods to those proposed by Flossdorf et al [137, 138], Miles et al [140] showed that the model is consistent with the data of Buchholz [137] et al, but applied to the data of Kinjyo et al [141] suggests that memory precursors are produced before effector cells. This shows that this type of inference is robust but the biological interpretation of the results depends both on the experimental conditions and model assumptions.

4. Quorum sensing

Recent experiments have revisited the question of naive T-cells differentiation into different memory cell types upon antigen stimulation, by considering collective effects through cell-cell communication. Polonsky et al [86] tracked live cells tagged with antibody markers over time in microwells with different initial cell numbers. By continuously tracking individual cell differentiation states and proliferation, they are able to overcome the limitations of bulk experiments where local cell densities are hard to control. Live-cell imaging showed that the decision to differentiate into progenitor central memory (pTCM) is collective: in equal medium conditions, cells surrounded by more cells had a larger differentiation rate. That rate depended solely on the instantaneous number of cells in the well, as shown by data collapse of the differentiation rate onto a single curve as a function of the number of cells, for various conditions (different media, different initial number of cells in the well). The curve shows that cells are more likely to differentiate into pTCM if there are more than $N_c \sim 30$ cells in the cluster. The process can be described by a stochastic differentiation model, in which cells divide and die, and differentiate with a rate R . However, to explain the data $R(N)$ must be made to depend on the number of cells N . Such a stochastic model also captures the observed well-to-well variability. The collective decision could be linked to the IL-2 and IL-6 cytokine communication: blocking IL-2 reduced the maximum of the universal differentiation rate curve, while blocking IL-6 increased N_c without significantly altering the plateau value. Interestingly, IL-2 and IL-6 receptors were found to cluster on cell surfaces, with receptor patches directed towards neighbouring cells. This experiment shows that the percentage of cells that differentiate into memory cells and effector cells depends on the instantaneous local T-cell density, independently of the additional post-stimulation influence of T-cell receptor signalling strength, or the effect of antigen presenting cells.

Polonsky et al [86] give a threshold in terms of absolute number of cells in a micro-well needed for collective decision making, $N_c \sim 30$. However, that number must depend on the size of the well, if cells do communicate through cytokine diffusion. As we have discussed in section IV B in the context of CD4+ T-cell differentiation into Tregs and T-helper cells, signal propagate to short distances as they are taken up by neighboring cells [96]. For pTCM differentiation, this is facilitated by the IL-2 and IL-6 orientation of the receptors towards secreting cells. But at large distances, diffusion starts playing a role. Active clustering seems to be a mechanism that puts cells in the non-diffusive regime.

The scale of cytokine communication has also been shown to depend on the type of signal: an anti-inflammatory IL-10 signal produced by dendritic cells is long-ranged, while the pro-inflammatory TNF α signal produced by the same cells is short ranged. These

cells are also known to cluster in cultures and in-vivo. A two step model for the integration of these two opposing signals was proposed to explain the nonlinear response of dendritic cells when integrating pro- and anti-inflammatory information [97]: an initial intermediate molecule integrates both signals in an indiscriminate way until a threshold value, providing a bottleneck. This way either the pro- or anti-inflammatory signal can stimulate the system. But the output of the this bottleneck is later down-regulated by only the anti-inflammatory signal. This combination of collective decision making with a modulated bottleneck safeguard was proposed as a way to control possibly excessive immune responses.

Collective decision making, also called “quorum sensing”, was previously proposed theoretically as a way for immune cells to help solve the self non-self discrimination problem [32]. According to that theory, which we discuss in more detail in Sec. VIC, communication allows cells to make a census-based decision by integrating the signal read out by many clones in order to correct mistake made by individual cells. This is a cellular implementation of an error-correcting code, similarly to kinetic proof reading discussed in section III A 1. The context of that proposal is different from that of [86], where the T-cell population is monoclonal, while the idea of error correction relies on taking a census of different T cell clones that make independent decisions. This difference in clonality does not change the general similarity in the nature of the collective decision. The experimental results are likely to hold for different T-cell clones with similar affinities for the stimulating antigen. The T-cell response is multiclonal (see section VI G) and any quorum-sensing mechanism is likely to involve multiple distinct clones. Understanding the details of how the memory and effector cells pools are controlled in a multiclonal setting remains an open question.

VI. REPERTOIRES

A. Size of immune repertoires

For the adaptive immune system to protect us against all the different pathogens we may encounter, including ones that may not exist when we are born, we need a large set of different immune receptors.

There are of the order of $4 \cdot 10^{11}$ T cells circulating in the human body [142], and of the same order of B cells [143], each expressing a single type of receptor to a first approximation [144, 145]. However, T and B cells divide, meaning that many cells can carry the exact same receptor, defining a “clonotype”. The repertoire size, or number of clonotypes, is therefore smaller than the number of cells. Early work [146] based on a subset of the repertoire gave an estimate of $\sim 10^6$ unique β chains in one human individual.

However, these estimates were indirect, and have been updated recently thanks to the advent of repertoire

high-throughput sequencing (RepSeq) techniques in 2009 [147–149], discussed in a number of recent reviews [150–155]. This method, which we briefly discuss in section XIIF, allows one to obtain the sequences of a non-exhaustive but fairly substantial portion of all receptors in a biological sample. As different technologies have been emerging, the need for standardisation of RepSeq data reporting has appeared, and the Adaptive Immune Receptor Repertoire (AIRR) community [156] has been organizing meetings and producing web-based and journal publications to promote data reporting unification.

RepSeq experiments on human blood samples, combined with statistical extrapolation estimators, report unique TCR β numbers ranging from 4×10^6 [148] (using Fisher’s Poisson abundance model [157]) to the order of 10^8 [158] (using the Chao2 estimator [159]). For BCR heavy chains, estimates obtained with the Poisson abundance model yields $1\text{--}2 \cdot 10^9$ [160]. One should take all these empirical estimate with great caution, as they make unverified, and probably wrong, assumptions about the shape of the clonotype abundance distribution when doing the extrapolation from small blood samples to the entire organism (more on that in Sec. VID). Theoretical estimates based on population dynamics models assuming that naive cells divide little or not at all logically give much higher estimates, closer to the number of cells itself, $10^{10} - 10^{11}$ [161, 162].

How large is large enough? De Boer and Perelson [17] proposed that the minimal repertoire size is ultimately determined by self-tolerance. We revisit their argument in the light of modern estimates, with simplified notations, starting with the easier case of B-cells. B-cells interact directly with antigens, unlike T-cells that interact with short peptide fragments presented on the multi-histocompatibility complex (MHC). Call $p \ll 1$ the probability that a randomly chosen receptor recognizes a randomly chosen epitope. During the self-tolerance selection process, the initially generated repertoire of size R_0 is reduced to R , $R = fR_0$, where f is fraction of receptors that survive negative selection, i.e. that do not recognize any self-epitope, so that

$$f = (1 - p)^n \approx e^{-pn}, \quad (43)$$

where $n \sim 10^5\text{--}10^6$ is the number of self-epitopes (in humans). The probability that an epitope escapes the immune system is then given by

$$p_e = (1 - p)^R \approx e^{-pR}. \quad (44)$$

This gives the repertoire size $R = -(1/p) \ln p_e$ as well as the pre-selection repertoire size

$$R_0 = -(1/p)e^{pn} \ln p_e. \quad (45)$$

Assuming that evolution has optimized the recognition probability p so as to minimize R_0 , solving for $\partial R_0 / \partial p$ an optimal $p^* = 1/n$, and an optimal repertoire size

$$R^* = R_0 f^* = -n \ln p_e, \quad (46)$$

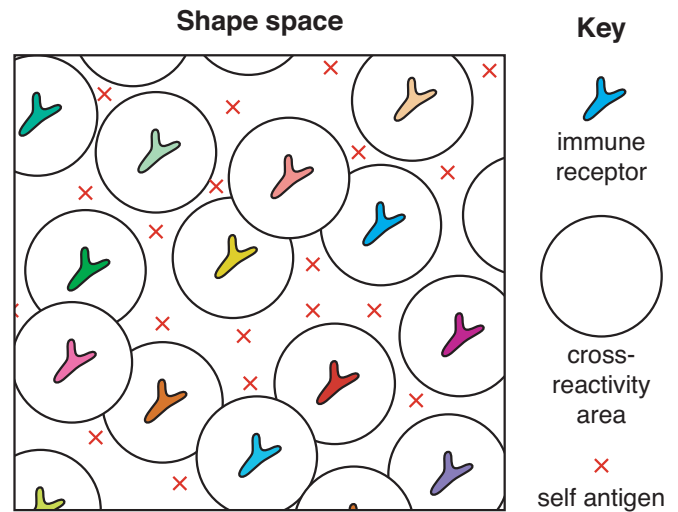


FIG. 10. **The minimal repertoire size is set by the number of self-antigens through negative selection.** A two-dimensional cartoon of the recognition “shape space”, an abstract space which summarizes the main bio-chemical properties of antigen-receptor binding into a unique space [33]. Antigen-receptor pairs that are close are likely to bind and trigger an immune response, while distant pairs do not interact. Each expressed receptor of the repertoire covers a ball of cross-reactivity where antigens are recognized by that receptor. Negative selection ensures that self-antigens (red crosses) are not included in any of the cross-reactivity balls of the repertoire. To cover shape space efficiently with that constraint, the repertoire size (number of balls) must scale with the number of self-antigens to avoid.

with an optimal selected fraction $f^* = \exp(-1) \approx 0.37$.

This estimate suggests that the number of self-epitopes is the main determinant of the minimal repertoire size, with which it scales linearly. Another prediction is that the recognition probability should be inversely proportional to the number of self-epitopes. We can intuitively understand this result by thinking about the repertoire as a covering problem: viewing the set of epitopes recognized by a single receptor as a ball in an abstract phenotypic space, how many such balls does one need to cover the entire space of foreign epitopes, while avoiding a finite number of self-epitopes? The simple argument given above, as well as the schematic of Fig. 10, tells us that the volume of each ball, p , should be inversely proportional to the number of self-epitopes, n , implying in turn that their number (R) should scale with n .

Empirical estimates suggest $p \sim 10^{-5}$ and $n \sim 10^5 - 10^6$ [17], consistent with the theoretical prediction $p \sim 1/n$. The prediction of $f \approx 37\%$ of cells passing negative selection is consistent with recent experimental estimates $25\% - 45\%$ [163]. The predicted B cell receptor diversity crucially depends on p_e , which is hard to estimate. De Boer and Perelson estimated that 99% of antigens, presenting each 10 epitopes, were recognized by at least one antibody of the immune system. This implies $p_e^{10} \sim 10^{-2}$, and $R^* \sim 5 \cdot 10^4 - 5 \cdot 10^5$, which

is much smaller than current experimental estimates of BCR diversity ($\sim 10^9$).

These estimates become a bit more involved for TCR since epitopes are presented by the Multi-Histocompatibility Complexes (MHCs), which translates into an additional level of sampling. As already mentioned in previous sections, there are 2 types of MHCs relevant to adaptive immunity: MHC type I complexes present peptides from inside the cell to CD8+ (also called killer) T-cells, whose goal is to kill infected or cancerous cells; MHC type II complexes present peptides from outside the cell to CD4+ (helper) T-cells whose role is to stimulate B-cells during affinity maturation process (see section below).

Each epitope, or peptide derived from the antigen of interest (self or foreign) must both be presented by one of the $m = 6$ MHC genes for each MHC type, and be recognized by one of the R TCR expressed in the body. This implies that the probability of escape reads:

$$p_e = (1 - q + q(1 - p)^R)^m \approx \exp[-mq(1 - e^{-pR})], \quad (47)$$

where $q \ll 1$ is the fraction of peptides that can be presented by a given MHC molecule. Eq. 43 is still valid, with n being the total number of possible MHC-self-peptide complexes. The repertoire size is then given by $R_0 = -(e^{pm}/p) \ln[1 + (\ln p_e/mq)]$, and the optimal size is obtained through $\partial R_0/\partial p$:

$$R^* = -n \ln \left(1 + \frac{\ln p_e}{mq} \right) \quad (48)$$

with $p^* = 1/n$ and $f^* = \exp(-1)$. Thus, the scaling and selection probability are predicted to be the same as for B cells.

The probability of presenting a given peptide on MHC class I complexes is relatively well known: deep learning algorithms have been very successful in building algorithms that predict which peptide is likely to be presented. These methods are implemented in software packages such as netMHC [40, 164], which account for known biological features to guide the learning. The same prediction task has been implemented for MHC class II [165], but the prediction is much harder, partly because unlike MHC class I presented peptides, which have a relatively fixed length of 9-11 amino acids, MHC class II presented peptides have varying length, from 13 to 25. More data than is currently available is needed to successfully train the neural networks, and to improve prediction. The models predict that about 1% of peptides are presented by any given particular MHC allele, i.e. $q = 0.01$.

The number of MHC-self-peptide complexes is given by $n = qmn_p$, where n_p is the number of peptides in the human proteome. Each individual has 6 MHC alleles for each class, i.e. $m = 6$. n_p is roughly the number of proteins, $\sim 3 \cdot 10^4$, times their average length, ~ 400 , i.e. $n_p \sim 10^7$, and $n \sim 6 \cdot 10^5$. The recognition probability, p , has been estimated to be in the range 10^{-4} - 10^{-6} [166,

167]. These numbers are consistent with the theory $p \sim 1/n$. The fraction of negatively selected cells, predicted to be $f^* \approx 37\%$ by the theory, is estimated to fall in the range 20-50% [168]. As for B cells, predicting the size of the repertoire is hard without a good estimate of p_e , but current estimates of $R = 10^8$ - 10^{10} are probably much larger than predicted by (48) regardless of the estimate of p_e .

In summary, these theoretical estimates still seem to be relevant today for the relationship between recognition probability, p , and number of self-epitopes, n , both for T and B cells. Recent estimates of TCR and BCR diversity, however, are much larger than predicted by these simple theories. However appealing, a major limitation of these estimates and predictions is that Eq. 43 assumes that each receptor is tested against all possible self-epitopes, while in practice this is impossible due to the limited duration of lymphocyte maturation. We will come back to this point in Section VIC devoted to modeling thymic selection.

B. Inference of the stochastic repertoire generation process

Antigen receptors are proteins, which must be encoded as genes in the DNA. Humans have of the order of 10^4 protein-coding genes. Directly encoding the whole diversity of immune receptor genes ($\sim 10^8 - 10^{10}$) in each genome would make it impossible for the DNA to fit in the nucleus. The immune system has solved that problem by stochastically creating receptors in each cell, in a process called V(D)J recombination that combines combinatorics and randomness to generate diversity. B- and T- cell receptors are made of two chains, light and heavy for BCR, and α and β for TCR. The genome encodes a certain number of gene templates, called V, D and J for the heavy and β chains, and V and J for the light and α chains. Upon creation of a chain, DNA is edited and one of each of these gene templates is chosen per receptor. The combinatorics of templates typically results in $\sim 10^3$ different receptors. To obtain the (much larger) observed diversity, nucleotides are randomly inserted in a non-templated way and deleted at the junctions between the V and D and D and J genes (or V and J genes for α and light chains). This process of generating “junctional diversity” in fact accounts for most of the diversity of the repertoire [169, 170].

Characterizing the above described process of receptor generation in quantitative detail has been made possible thanks to the development of RepSeq methods (see Sec. XII F). This can be done using out-of-frame sequences (with a frameshift due to the random number of additions and deletions at the junctions), which are nonproductive and hence a raw product of recombination, as they are free of selection effects. These sequences survive selection because they are in the same cell as an in-frame sequence. Since each cell has two chromosomes,

out-of-frame sequences come from cells where the first rearranged chromosome resulted in an out-of-frame sequence and the second rearranged chromosome resulted in an in-frame sequence. Due to the random insertions and deletions of nucleotides, it is impossible to reliably determine how a given receptor was formed (V, D, J assignments, as well as distinguishing inserted from templated nucleotides) from its mere sequence. Instead one can consider a list of scenarios (which includes V, D and J gene choice plus a number of insertions and deletions at each of the junctions) and sum over them weighted by their probabilities as determined self-consistently using a probabilistic model. Concretely, in the case of the simpler α or light chains, the probability of a given recombination scenario r is:

$$P_{\text{rearr}}(r) = P(V, J)P(\text{del}V|V)P(\text{del}J|J)P(\text{ins}VJ), \quad (49)$$

where $\text{del}V$ and $\text{del}J$ denote the number of deletions at the V and J ends, and $\text{ins}VJ$ is the list of inserted nucleotides. A similar expression can be written for the β or heavy chains by adding the D gene and the related deletions and insertions. The model is the most general factorizable distribution that is consistent with the data and the known biological constraints (such as the relative positioning of the genes in the genome). The overall generation probability of a given sequence s is obtained by summing of the probabilities of all scenarios that could have given rise to this sequence:

$$P_{\text{gen}}(s) = \sum_{r \rightarrow s} P_{\text{rearr}}(r). \quad (50)$$

The learning of the model parameters, encoded in the probability distributions in the factorized form (49), is performed by maximizing the likelihood of the sequences, using the Expectation-Maximization algorithm to deal with the sum over the hidden variable r (see Sec. XII E 1 for details about Maximum Likelihood and Expectation-Maximization).

This inference procedure has been applied to variety of immune receptors and species, TCR β [169] and α [171] chains, BCR heavy [170] and light chains [172] in humans, as well as TCR β in mouse [173] and BCR heavy chains in trout [174]. It is implemented in the IGoR software [175] which can be used to learn additional models from other species and locus combinations. IGoR can also be used to generate synthetic sequences from the P_{gen} distribution, and to estimate $P_{\text{gen}}(s)$ of an arbitrary nucleotide sequence s . Another method, OLGA [176], was designed to estimate the probability of amino-acid sequences, using dynamic programming to deal with the enormous sums involving the enumeration of all possible nucleotide variants. Other methods relying on hidden Markov models (see Sec. XII E 1), were also proposed to handle the high dimensionality of hidden variables [177–180].

Repertoires generated *in silico* using the learned model confirm that the inference is able to call the correct recombination scenario in at best $\sim 25\%$ of cases, making

the deterministic annotation of sequences impossible and necessitating a probabilistic inference approach. A perhaps counterintuitive consequence is that the model parameters can be inferred with arbitrary accuracy despite our inability to annotate any sequence reliably.

The results of the inference show that the probability of generation of receptors is incredibly reproducible between individuals of the same species. The distribution is dominated by insertions and deletions, whereas gene choice contributes relatively little to the overall probability. Among the features of V(D)J recombination, the gene choice distribution varies the most between unrelated individuals [169], even when corrected for single nucleotide polymorphisms (SNP) in V, D and J alleles [181]. As far as can be determined, the generations of the α and β chains are largely independent, and the overall generation probability of a receptor is well approximated by the product of the generation of its two chains [145]. Although no similar analysis has been performed for BCR, it is likely that the formation of its two chains are also independent.

Although these recombination models ignore selection effects, to be discussed in the section below, it is a good predictor of the abundance of particular TCR β in a human population [176], and it has been used to detect signatures of immune responses as deviations from this baseline distribution [182, 183].

C. Thymic selection and central tolerance

After the receptors are generated, they undergo an initial selection step, known as thymic selection in T-cells. A similar process called central tolerance occurs in B-cells maturation [163], but let us focus on T-cell for concreteness. For a more comprehensive survey of models of thymic selection, we refer readers to the review by A. Yates [168], and focus here on recent development involving quorum sensing and data-driven models of sequence specific selection.

T-cells mature in the thymus, an organ that contains only proteins that are native to the host organism, termed self-proteins. The newly generated receptors are expressed on cell surfaces and their binding properties are tested against the peptides from the self-proteins, presented by the MHC complexes. If the receptor fails to bind any self-peptide, even weakly, it will probably fail to bind any protein and the cell carrying these receptors is discarded – a process called positive selection which removes $\sim 80\%$ of immature cells. Conversely, if a receptor binds any one self-protein too strongly, it is also discarded as a result of negative selection, since it is likely to bind self-proteins later on and trigger autoimmune diseases. Negative selection was at the core of the argument for the optimal recognition probability and repertoire size presented in Sec. VI A, and it is estimated to remove 50% – 80% of cells. The detailed process of T cell selection, complete with the timing of the α and




FIG. 11. **Thymic selection.** A cartoon representation of the basic processes that take the cells from the bone marrow to the thymic cortex and medulla and finally lead to the generation of functional naive CD4+ and CD8+ T-cells that are exported to the periphery. The duration of the processes is not drawn to scale – it is meant solely as an indication of temporal order and overlap.

β chain recombination events, is summarized in Fig. 11. During positive selection, cells with TCR that recognize peptides presented by the MHC class I commit to becoming CD8+ cells, whose main function is to kill infected cells, while cells with MHC class II TCR specificity become CD4+ cells, whose main function is to help B cells in their affinity maturation, and to regulate the immune response. There are about 4 times as many CD4+ cells as CD8+ cells. A small subclass of CD4+ T cells, called regulatory T cells (Tregs), actually suppress the immune response, and play an essential role to prevent auto-immunity [184]. Tregs are selected for higher affinities to self-epitopes than regular T cells, which allows them to selectively suppress immune responses to self-antigens, although the picture seem to be quite complex (see [185]).

To be able to ensure self tolerance, each receptor should in principle be tested against all possible self-epitope, which would take an impractically long time. In practice, experiments suggest that each T cell may encounter around 500 antigen presenting cells [186], while theoretical arguments estimate the number of presented self peptides around a few thousands [32] (these numbers are not inconsistent because each antigen-presenting cell may present many different peptides in single encounter with a T cell). These numbers are much lower than the estimated diversity of presentable self peptide-MHC complexes, $5 \cdot 10^5 - 5 \cdot 10^6$ (see Sec. VIA). Calling x the fraction of presented self peptides during negative selection, the probability that none of the selected TCR are self-reactive is $(1 - p)^{RNx}$, which quickly goes to 0 as R becomes even moderately large, even if p scales as $1/N$.

A recently proposed solution to this problem is to assume that a minimal number t of TCR must recognize a peptide to trigger an immune response [32]. By virtue of the law of large numbers, the number of T cells responding to a self-peptide is distributed as a Gaussian with mean and variance $= Rp(1 - x)$, while the number of T cells responding to a foreign peptide is of mean and variance $= Rp$. Good discrimination is achieved when

the two distributions are well separated, giving the condition $Rpx^2 \gg 1$, with an optimal discrimination threshold $t^* = Rp(1 - x/2)$ in the limit $x \ll 1$. This condition is barely satisfied by $x = 10^{-2}$, $R = 10^9$ and $p = 10^{-5}$, for which the optimal threshold for the number of TCR involved is $t^* \sim 10^4$. This argument suggests another, stricter lower bound on the diversity R of the T cell repertoire that is necessary to make such a collective decision. It also implies a “quorum sensing” mechanism by which responding T cells have a way of estimating how many other cells are involved in order to commit to a response. Recent work suggests that such quorum sensing does occur locally, probably using cytokine signaling [86].

To model thymic selection in more quantitative detail, a common strategy has been to use additive models of binding free energy similar to (9) [29, 30, 35]. In these models, the only receptors that survive are those whose maximal affinity E^* to any self peptide falls within a range (E_p, E_n) corresponding to the positive and negative selection thresholds. This condition can be mapped onto an extreme value statistics problem, allowing for a statistical mechanics treatment [31]. Under this framework, it was shown that the sequence composition of negatively selected TCR was biased towards weakly binding residues [29], and the theory was subsequently applied to explain clinical data on “elite controllers” of HIV expressing a particular type of MHC class I molecule, HLA-B27 [187]. A similar theory was used to study the sensitivity of TCR that target tumor cells [30]. The parameters of these models are not inferred empirically, but are instead picked from popular but unrealistic interaction potentials between amino acids such as the Miyazawa-Jernigan matrix [28]. Yet many conclusions of these studies are relatively insensitive to the details of the interaction matrices in the extreme value statistics regime.

Statistical inference methods based on immune repertoire sequencing can also be used to estimate the probability that a particular receptor passes selection based on its sequence. We can define a selection factor corresponding to the ratio of probabilities to find a sequence in a selected repertoire, $P_{\text{sel}}(s)$, relative to its probability in the unselected repertoire, as given by the recombination model $P_{\text{gen}}(s)$: $Q(s) = P_{\text{sel}}(s)/P_{\text{gen}}(s)$. In practice, it is impossible to evaluate $P_{\text{sel}}(s)$ directly, as the number of possible sequences to be considered is too large, spanning more than 20 orders of magnitude in generation probabilities. However, simplifying assumptions can be made on the form of $Q(s)$. Specifically, Elhanati et al [188] considered:

$$Q(s) = q(L)q(V, J) \prod_{i=1}^L q_{i|L}(a_i), \quad (51)$$

where (a_1, a_2, \dots, a_L) is the amino-acid sequence of the Complementarity Determining Region 3 (CDR3, which forms a loop important for antigen recognition, and covers the most variable part of the receptor ranging from the end of the V to the beginning of the J segments).

Selection is assumed to act independently on the length of the CDR3 region $q(L)$, the V gene, $q(V)$, and the J gene, $q(J)$. The D gene is not taken into account separately, but selection on amino acids in the CDR3 is considered explicitly. The parameters of the model are inferred by maximizing the likelihood using Expectation Maximization (Sec. XII E 1). This approach has been applied to the initial selection process of β -chain sequences [188], α -chain sequences [171], and BCR heavy chain [170] and light chain [172] sequences. For many of these cases, the $Q(s)$ selection factors did not differ substantially from individual to individual. Instead, they reflected an global selection for general biophysical and biochemical features of amino acids (suggesting positive selection for proper protein function), rather than individual-specific removals in the repertoire caused by negative selection. Similar observations were made for BCR by direct comparison of pre-mature and mature repertoires [189]. Modeling negative selection is much harder, as it requires to learn the “holes” that selection pokes into the repertoire, whose post-selection landscape looks like a dense golf course. A simple way to model negative selection is to remove a random fraction $1 - q$ of the sequences from the selected repertoire, boosting the likelihood of surviving ones by a factor $1/q$ [190]. In practice, in simulations one can use a hashing function, which associates a quasi-random number to each sequence, and select sequences whose hashing number falls below a chosen threshold, so that selection is reproducible, yet agnostic to the features of the sequences.

Note that models of the form (51) are not restricted to thymic selection, and can be used to describe the sequence-wide selection pressure during any process, by comparing the initial distribution before selection to the final distribution after selection, as was for instance done to characterize responsive receptors to yellow fever vaccination [183]. For instance, different T cell subsets (e.g. CD4+ and CD8+ cells), different BCR isotypes, in different organs or environments, are probably characterized by distinct selective pressure, each of which could be modeled using similar $Q(s)$ factors.

D. Diversity and the clone size distribution

The distributions of TCR and BCR span a very large space, and are highly skewed. It is often useful to summarize these distributions by a single summary statistics that quantifies their diversity. Diversity measures in the context of immune repertoires have been thoroughly discussed in a recent review [162]. Diversity is important for understanding how well our immune repertoire prepares us against a wide range of pathogenic threats. As we will see, diversity measures are deeply linked with the distribution of sizes of immune clones (number of cells with the same receptor), which contains important information about the dynamics of immune cells in response to environmental challenges (see Sec. VII). In addition,

it determines the amount of overlap one expects between repertoires of distinct individual, and underlies the concept of “public repertoire” shared by all individuals (see Sec. VI E). All these aspects can be formalized within the common language of statistical mechanics through the definition of a density of state, which we introduce below.

The most general family of diversity measures for a distribution $p(s)$ is given by Renyi entropies:

$$H_\beta = \frac{1}{1 - \beta} \ln \left[\sum_s p(s)^\beta \right]. \quad (52)$$

This definition reduces to the Shannon entropy for $\beta = 1$, $H_1 = -\sum_s p(s) \ln p(s)$ (see Sec. XII E 5), to the total number of receptors for $\beta = 0$, $R = \exp(H_0)$, and to the probability of drawing the same receptor twice with replacement for $\beta = 2$, $\exp(-H_2) = \sum_s p(s)^2$, also equal to the inverse of the the Simpson index. More generally, the family H_β recapitulates the entire clone size distribution, defined in its cumulative form as the number $G(E)$ of receptor sequences with $-\ln p(s) < E$, through a Laplace transform:

$$H_\beta = \frac{1}{1 - \beta} \ln \left[\int dG(E) e^{-\beta E} \right]. \quad (53)$$

Note that $G(E)$ is formally equivalent to a cumulative density of states in statistical physics. $E = -\ln p$ is sometimes called the “surprise” and is also formally similar to an “energy” by analogy with Boltzmann’s law, $p \sim \exp(-E)$, making β analogous to an inverse temperature. The mapping between $G(E)$ and H_β is formally analogous to the equivalence of the canonical and micro-canonical ensembles, and is amenable to a statistical physics analysis [191]. The lower the temperature (the higher the β), the more the Rényi entropy concentrates on low-energy, high probability sequences.

$G(E)$ can also be interpreted as the rank (ordered from most probable to least probable) of sequences with probability $p = e^{-E}$. The distribution of clone sizes is often presented in terms of a clone-size frequency rank distribution, where clones are ranked ordered by their sizes, and their normalized frequency is plotted as a decreasing function of its rank. $G(E)$ precisely encodes this rank-frequency relation.

It is important to distinguish between the potential diversity, corresponding to a theoretical distribution $P_{\text{gen}}(s)$ or $P_{\text{sel}}(s)$, and the realized diversity, corresponding to the actual distribution of a finite number of receptors in a particular sample, $p(s) = n(s)/N_{\text{tot}}$, where $n(s)$ is the number of molecules or cells with receptor sequence s , and $N_{\text{tot}} = \sum_s p(s)$.

Let us start with the potential diversity derived from the probability distribution $P_{\text{gen}}(s)$ (see Sec. VI B). For human TCR β , the total number of possible sequences, e^{H_0} , is infinite for all practical purposes ($> 10^{39}$ [162]), as it is larger than the total number of TCR β receptors having ever been produced by a human being. For this reason, it makes more sense to compare Shannon entropies.




FIG. 12. **Distribution of clone sizes follow power laws.** Cumulative distribution $G(p)$ of clone frequencies p of unfractioed TCR β sequences sampled from the blood of 6 human donors [171]. The dashed line represents slope -1 .

These can be computed using (53), where $G(E)$ is evaluated from the probability density distribution of E among randomly drawn sequences according to $p(s)$, $\rho(E) = (dG/dE)e^{-E}$, so that $G(E) = \int_0^E dE' \rho(E') \exp(-E')$.

Entropy estimates for both nucleotide and amino acid sequences are reported in [176]. Human nucleotide TCR β diversity is $H_1 \sim 44$ bits (bits refer to $\ln(2)$ units) for nucleotides and 30 bits for amino-acids, TCR α diversity is ~ 30 bits for nucleotides and 25 bits for amino acids. The BCR heavy chain entropy is 67 bits for nucleotides and 53 for amino acids. These numbers are very large. If the distribution were uniform, that would correspond to $2^{55} \approx 3.6 \cdot 10^{16}$ distinct amino acid TCR $\alpha\beta$ sequences. Because the distribution is non-uniform, the pooled repertoire of all humans having ever lived do not exhaust the potential diversity. As noted before, most this diversity is due to random insertions and deletions.

Thymic selection and central tolerance effectively reduce these potential diversities in $P_{\text{sel}}(s)$, by around 9 bits for TCR β [188], 4 bits for TCR α [171], and 12 bits for BCR heavy chain [170], but the corresponding diversity number remain very high. This diversity loss does not mean that e.g. $2^{-13} \approx 10^{-4}$ of TCR must be discarded; the probability of thymic selection $Q(s)$ is in fact correlated with generation probability $P_{\text{gen}}(s)$, meaning that diversity is reduced by selecting already likely sequences.

Let us now turn to the realized diversity in a sample or an individual. This diversity is best described by the whole clone-size distribution encoded by $G(E)$. Based on recent high-throughput sequencing experiments, the distributions are often long-tailed, spanning several orders of magnitude (see Fig. 12) [162, 192–194]. These tails seem to be mostly due to the memory fraction of the repertoire [195]. Initial worries that these long tails may

arise from noise is the Polymerase Chain Reaction, which amplifies differences exponentially [196], were put to rest thanks to the introduction of unique molecular barcodes associated to each initial mRNA molecule [154, 197, 198]. Specifically, the clone size distribution follows a power-law:

$$G(E) \sim e^{\hat{\beta}E} \sim \frac{1}{p^{\hat{\beta}}}, \quad (54)$$

with $\beta \leq 1$ but close to 1. The probability distribution of clone frequencies,

$$\rho(p) = -\frac{dG}{dp} \sim \frac{1}{p^{1+\hat{\beta}}}, \quad (55)$$

is also a power law of exponent slightly below 2. Another way to look at the distribution is to consider the inverse relationship,

$$p \sim \frac{1}{G^{1/\hat{\beta}}}, \quad (56)$$

which also follows a power law, called Zipf’s law in this particular context when its exponent $\hat{\beta}^{-1}$ is close to 1. Since G is the rank of the clone (from most frequent to least frequent), this relation is called the rank-frequency relationship.

Because of these long tails and limited sampling, very few of the diversity measures can be estimated reliably. One way to overcome these issues is to build a statistical model that is then parametrised by the experimental data. In Sec. VIA we already mentioned Fisher’s Poisson abundance model and the Chao estimators, which use this idea. More sophisticated methods have been proposed for immune repertoire, e.g. using Expectation Maximization (see Sec. XII E3) [199] or using rarefaction curves [200], but all these methods are susceptible to huge errors when the clone size distribution follows a power law [201]. We expect the reported numbers to grossly underestimate the true diversity of both BCR and TCR. To make progress, models of lymphocyte population dynamics should be leveraged to reliably extrapolate the tails of clone-size distributions. We will discuss some of these models in Sec. VII.

E. Repertoire sharing

Receptor sharing is at the heart of the public vs private debate in immunology. Public receptors are those shared between individuals in a given cohort, while private are ones that are seen only in individuals. If one studies a group of individuals with a specific condition (e.g. cytomegalovirus or CMV [204] or Ankylosing Spondylitis [205]), receptors that are shared between these people can be expected to be linked to a disease. This justifies the search and investigation of “public” sequences. Both humans and mice do share a non-negligible numbers of receptors in functional repertoires, regardless of their

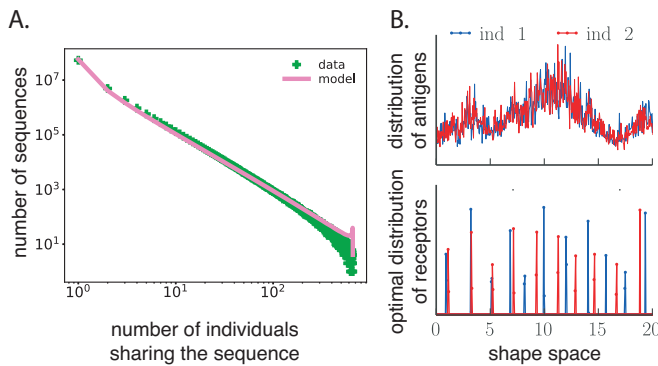


FIG. 13. Overlap of immune repertoires between individuals. A. The amount of shared lymphocyte receptors (TCR β chain) between individuals is consistent with pure chance [190] (data from [202]). B. To explain that observation, a simple model of optimal immune coverage of pathogens [203] predicts that even very similar immune environments experienced by two individuals (top) can lead to very different (and peaked) distributions of protections (bottom).

lifestyles, environmental factors and family ties. However, receptors are also expected to be shared by chance, because some sequences are likely to recombine independently in different individuals [188, 206–208]. One can try to correct for this convergent recombination by using an outgroup — a group of individuals that are as similar to the condition-specific group but do not have this condition — and identify sequences shared more in the studied group than in the control group, as was done for CMV [204]. Alternatively, one can use the knowledge of immune repertoire diversity encoded in $P_{\text{gen}}(s)$ to estimate our null expectation.

All sharing properties are recapitulated by the cumulative density of states $G(E)$. Given n independent repertoires of sizes N_1, \dots, N_n , the expected number of sequences shared by exactly m repertoires, M_m , is given by the following generating function [190]:

$$g(x) = \sum_{m=0}^n M_m x^m \quad (57)$$

$$= \int_0^{+\infty} dG(E) \prod_{i=1}^n \left[e^{-N_i e^{-E}} + \left(1 - e^{-N_i e^{-E}}\right) x \right]$$

For example, for $n = 2$, in the limit of small samples, $N_{1,2} e^{-E} \ll 1$, the expected number of shared sequences is

$$M_2 \approx N_1 N_2 \int dG(E) e^{-2E} \quad (58)$$

$$\approx N_1 N_2 \sum_s p(s)^2 = N_1 N_2 e^{-H_2},$$

reducing to simple biased birthday problem related to Simpson’s diversity and thus Rényi entropy of order 2. More generally, sharing between m small samples is governed by $\int dG(E) e^{-mE}$ and to Rényi entropies of order

m . The ensemble of sequences shared between m samples thus corresponds to sampling the generation probability at temperature $1/m$, focusing on more and more likely sequences as m increases.

This calculation shows that in such large ensembles of cells as the B or T cell repertoire, which contain upward of 10^8 unique receptors, we expect the most common receptors with a high probability to be independently generated multiple times in different people. The number of shared TCR β sequences among sampled repertoires from a cohort of 658 human donors [204] was well predicted by the model in Eq. 57. By itself, the generation probability P_{gen} was enough to build a classifier that can determine whether a given sequence will be shared by a minimal number of people, with a $> 95\%$ accuracy for $m > 10$. This analysis also brings our attention to the definition of public repertoires. A receptor that is shared between $m = 2$ and $m = 1000$ people has a very different notion of publicness. For this reason it is better to talk about degrees of publicness. In addition, sharing and publicness depends very sensitively on the collected sample size N_i . A reasonable choice for a single definition of receptor publicness could be that the sequence must be present on average once in each individual, implying $P_{\text{gen}}(s) > 1/R$.

Interestingly, the theory can also identify receptors that are shared for other reasons than sheer chance. These receptors have a high propensity to be shared, but a relatively low generation probability. The discrepancy between chance prediction and observed sharing was used to discover shared T cell clones in identical twins [171], to find candidate TCR sequences associated with conditions such as CMV or diabetes without a control group [183], or to assess the public BCR response to vaccination in trout [174].

Identifying common, even unusual receptors does not necessarily mean they are responding to the same antigen or pathogen, and even when it does it does not tell us what the target antigen is. As discussed in Sec. IID, immune response is a complex phenomenon that involves the binding different epitopes to immune receptors, followed by signal propagation, cell commitment, and cell-cell communication through messenger molecules. Yet tools based on the statistical expectations of an unbiased repertoire allow us to identify deviations from this baseline, and propose candidate responding receptor sequences to be tested in other experiments.

F. Optimal receptor distribution

In Sec. VIA we discussed theoretical arguments for an optimal repertoire size, but recent data suggests that actual repertoire sizes are much bigger than these theoretical predictions. However, as stressed in Sec. VID, the distribution of receptors is itself highly skewed, suggesting that the relative abundance of receptor types, rather than their absolute number, is an important factor of

repertoire design. Can we find an optimization principle over the composition of the repertoire? Take a distribution Q_a of antigens to be recognized, where a lives in a theoretical high-dimensional “shape space” encompassing both antigens and receptors, and in which proximity reflects receptor-ligand affinity [33, 209], through a recognition or “cross-reactivity” matrix f_{ar} between antigen a and receptor r which encodes the probability that a random encounter between the two results in an immune response. For a given distribution of receptors, P_r . The expected cost of an infection linked to a is a decreasing function $c(\cdot)$ of the probability that a chance encounter with a receptor is successful, $\sum_r f_{ar}P_r$, so that the expected cost for a random infection reads [210]:

$$C(\{P_r\}, \{Q_a\}) = \sum_a Q_a c \left(\sum_r f_{ar} P_r \right). \quad (59)$$

Minimizing that cost with respect to the receptor distribution $\{P_r\}$ gives an interesting multi-peak structure (Fig. 13B). Because of the degeneracy of recognition allowed by cross-reactivity, similarly well protecting repertoires in different individuals may exist, even if they sample essentially the same pathogenic environment (Figure 13B). As a result, we should not be surprised by great differences at the phenotypic level, let alone at the sequence level, when looking at repertoires responding to specific threats. In other words, even if convergent selection is at play, each immune repertoire may find a different molecular solution to it, which may explain the observation that most repertoire sharing seem to occur by chance.

G. Repertoire response to an immune challenge

RepSeq technologies allow one to track changes in the repertoire in time and in response to environmental changes, infections, or vaccinations. In many animal models such as mice or fish, it is generally not possible to sample the repertoire in the same individuals. Instead, the repertoires of distinct but isogenic individuals are sampled at the different times before and after an immune challenge. By construction, this strategy restricts the analysis to public features of the repertoire response. To identify common sequences features of repertoires that responded to a particular challenge, machine learning techniques based on low-level features [211] or support vector machines [212] have been used in mice to classify responding repertoires. Enrichment of particular V classes and sequences features have also been reported in the BCR repertoires of trout following immunization [174, 213].

In humans, the BCR response to influenza vaccination [197, 214], BCR and TCR response to varicella-zoster vaccination [215, 216], and the TCR response to yellow fever vaccination (a model for an acute infection) [217, 218] have been studied at the repertoire level using RepSeq. These studies are mostly descriptive, but

they pave the way towards a better characterization of the specificity, reproducibility, and publicness of the immune response at the repertoire level. A common observation is that the response is mostly private, although responding clones are more shared than would expected by chance. Responding sequences also tend to be clustered in sequence space [218]. This clustering gives a criterion for identifying responding sequences in a single repertoire snapshot, based on the density of similar sequences (differing by at most one amino acid) in the repertoire, relative to this density expected by the recombination model, $\sum_{s', |s-s'| \leq 1} P_{\text{gen}}(s')$ [183].

VII. LYMPHOCYTE POPULATION DYNAMICS

T and B cells are organized in clones of cells that express the same immune receptor. These clones grow and decay as individual cells divide or die either spontaneously or in response to external signals — e.g. an antigen recognition event triggering cell proliferation, or cell-cell communication through messenger molecules called cytokines. The correct way to approach and to model these complex lymphocyte dynamics, as well as their implication for experimental observables such as clone size distributions or clone expansion, is still a largely open question. There is a vast body of literature on modeling lymphocyte dynamics for broad subpopulations with no clonal information, using models ordinary differential equations. Ref. [120] provides a useful entry point into that literature. Here we will focus on stochastic models population dynamics, where the identities and sizes of individual lymphocyte clones are tracked. This approach is in line with recent developments in high-throughput repertoire sequencing (Sec. VI), which allows for such a fine-grain description of population dynamics, and for which much work remains to be done, both experimentally and theoretically. While progress on this topic is still in its infancy, it promises to give us better insight into the collective decision-making and dynamics of populations of immune cells, and to extract important *in vivo* parameters from data, at a moment when most studies are based on cultured cells or mouse models.

We also briefly review the application of differential equation models to HIV dynamics, in what constitutes perhaps the most successful application of modeling to translational immunology.

A. Neutral dynamics

Let us first consider a simple stochastic model of repertoire evolution in the absence of any antigenic proliferation. Such a model is relevant for the naive pool, which is often believed to be unaffected by external stimuli (although this is also debated, see below), but can also serve as a “neutral” baseline against which to compare other dynamics or longitudinal data. While some of the results

discussed in this section were presented in [194], here we expand on them and add a few original results, notably on the establishment of the steady state, and on estimates of the total number of clones.

In this simple model, new clones of size C come out of the thymus (for T cells) or bone marrow (for B cells) with rate θ_C , so that the total cell output $\theta = \sum_C C\theta_C$. Then each cell can divide with rate ν , and die with rate $\mu > \nu$.

The total number of cells $C_{\text{tot}}(t)$ follows the simple differential equation:

$$\frac{dC_{\text{tot}}}{dt} = \theta - (\mu - \nu)C. \quad (60)$$

This simple equation is at the basis of many studies aiming at quantifying lymphocyte population dynamics, using experiments with isotope labeling [219, 220] or other markers of cell divisions [221]. When thymic output, division and death rates are constant, Eq. 60 is solved by (assuming no cell at $t = 0$):

$$C_{\text{tot}}(t) = \frac{\theta}{\mu - \nu} \left(1 - e^{-t(\mu - \nu)}\right). \quad (61)$$

At steady state, the total number of lymphocytes $C_{\text{tot}} = \theta/(\mu - \nu)$ reflects the balance between thymic output, cell division and death. It is believed that the division and death rates adjust themselves so as to keep C_{tot} constant, a process called homeostasis. For T cells, $\theta \sim 10^6/\text{day}$ for mice, and $\theta \sim 10^8/\text{day}$ for humans, although that number varies with age. Taking $C_{\text{tot}} \sim 10^{11}$, this gives an effective decay time of $(\mu - \nu)^{-1} \sim 1,000$ days, which is consistent with lifetime estimates of T cells in human using deuterium water [220].

However, these numbers are not informative about the repertoire structure and its clone size distribution. These calculations are also unable to disentangle division and death, lumped in the single parameter $\mu - \nu$. The expected number of clones of size C , $N_C(t)$, is governed by the following equations [193]:

$$\begin{aligned} \frac{dN_C}{dt} = & \nu((C - 1)N_{C-1} - CN_C) \\ & + \mu((C + 1)N_{C+1} - CN_C) + \theta_C \end{aligned} \quad (62)$$

At steady state, $dN_C/dt = 0$, the solution for $C > \max\{C : \theta_C > 0\}$ is a power-law of exponent 1 with an exponential cutoff:

$$N_C \sim \frac{1}{C} \left(\frac{\nu}{\mu}\right)^C. \quad (63)$$

When birth and death are balanced, $\nu \sim \mu$, the exponential cutoff disappears. Comparing to the power law Eq. 55 in the distribution of frequencies $\rho(p)$ with the correspondance $N_C = N_{\text{tot}}\rho(p = C/C_{\text{tot}})$, this model would predict a exponent $\hat{\beta} = 0$, which is not supported by repertoire data on unsorted or memory cells. This suggests that the dynamics of the memory repertoire cannot

be simply neutral as modeled by Eq. 62, but must be affected by selection, such as expansions events following antigen recognition, as we will see in the next section.

The dynamics of Eq. 62 can in fact be solved analytically using generating functions and the method of characteristics [222]. Assuming for simplicity that clones have initial size 1, $\theta_C = \theta\delta_{C,1}$, we obtain (see Sec. XII D 3):

$$N_C(t) = \frac{\theta}{\nu} \frac{1}{C} \left(\frac{\nu}{\mu}\right)^C \left(\frac{1 - e^{-t(\mu - \nu)}}{1 - (\nu/\mu)e^{-t(\mu - \nu)}}\right)^C. \quad (64)$$

This relation is still a power law with an exponential cut-off. The power law exponent remains 1 at all times, but the cut-off gets larger and larger as the system reaches steady state. The total number of clones, $N_{\text{tot}}(t)$, then evolves according to:

$$N_{\text{tot}}(t) = \frac{\theta}{\nu} \ln \left(\frac{\mu - \nu e^{-t(\mu - \nu)}}{\mu - \nu}\right), \quad (65)$$

and the total number of cells, $C_{\text{tot}}(t) = \sum_C CN_C(t)$ follows Eq. 61. At steady-state, the total number of clones, $N_{\text{tot}} = (\theta/\nu) \ln[\mu/(\mu - \nu)]$ depends on both the division and death rates, and not just their difference. In the limit of small division rate, $\nu \ll \mu$, which is often assumed for naive cells, we get $N_{\text{tot}} = \theta/\mu \approx C_{\text{tot}}$. Conversely, in the limit $\mu - \nu \ll \mu$, one has $N_{\text{tot}} \sim C_{\text{tot}}\epsilon \ln(1/\epsilon)$ with $\epsilon = (\mu - \nu)/\mu$, meaning that the number of clones may be arbitrarily smaller than the number of cells. Starting with initial clone sizes of k_0 does not substantially affect this picture, with $N_{\text{tot}} \sim (H_{k_0}/k_0)C_{\text{tot}}$ (where $H_n = \sum_{i=1}^n 1/i$) in the limit $\nu \ll \mu$, and with unchanged scaling $N_{\text{tot}} \sim \epsilon \ln(1/\epsilon)C_{\text{tot}}$ in the limit $\mu - \nu \ll \mu$ (see Sec. XII D 3).

These models can be refined by accounting for convergent recombination, by which clones sizes can also be increased by thymic exports that have the exact same sequence, which happens with probability $P_{\text{gen}}(s)$ for a particular sequence s . With this correction, Zheng and collaborators [223] found fair agreement between the results of a neutral model with a source Eq. 62 and the bulk distribution of clone frequencies observed in the naive repertoire of mice. However, they also found many outliers, *i.e.* large naive clones that cannot be explained by the neutral assumption. A major difficulty of such a comparison is that, as we already mentioned in Sec. VI D, the clone frequency distribution of the full repertoire is heavy-tailed. Any small contamination of non-naive cells into the studied repertoire is likely to introduce spurious outliers in the distribution, confounding the analysis.

B. A note about “neutral processes”

In this review we discuss four different types of processes that are referred to as neutral in different domains: (i) the neutral model of population genetics (see section XII C 8), (ii) the neutral clone size distribution

Neutral model	N_C	Comment	Eq.
Lymphocyte dynamics	$\sim (\mu/\nu)^C/C$	Source	Eq. 63
Population genetics	$\sim 1/C$	Fixed pop. size	Eq. 141
Yule process	$\sim 1/C^{1+\rho}$	Expanding pop.	Eq. 207

TABLE I. Summary of different neutral models discussed in the review, with their clone size distributions. $\nu < \mu$ refer to the division and death rates respectively, and $\rho = 1/(1 - \alpha)$, where α the probability of mutation upon division in the Yule process.

model (see section VII A), (iii) the Yule-Simon process of speciation (see section XII D 5) and (iv) the neutral ecological model (see section XII D 1). Since these models are often confused with one another, sometimes simply for semantic reasons, otherwise because many of them give power law distributions for some observable, we thought it is pedagogical to discuss how they are different and summarize the exponents they give and what kind of distributions they predict.

Many of these models predict a power law in the distribution of clone or species frequencies, but details differ. Table VII B gives a summary of these distributions and differences between the model assumptions.

Neutral models of ecology come in different flavours [224], but the simplest one is essentially equivalent to Eq. 62, and the corresponding clone size distribution, called the species abundance distribution in that context, has long been known as Fisher’s logseries [157], although Fisher derived it through purely statistical means using unjustified assumptions.

The neutral model of population genetics is similar to the neutral model of Eq. 63, with a major difference: the population size C_{tot} is fixed. Most often, novel types arise from mutations arising in existing types. In population genetics the “site-frequency spectrum” corresponds to the distribution of allele frequencies, and is thus the equivalent of the clone size distribution N_C . In the neutral model, the site-frequency spectrum is a straight power law of exponent -1 . The difference with the result of Eq. 63, beyond the absence of an exponential cutoff, is that mutants that reach the (fixed) population size C_{tot} become the wildtype, and are thus removed from the dynamics of mutants, while in lymphocyte dynamics no single clone ever takes over the whole population.

The Yule process is a model of an ever expanding population with mutations. It assumes that individuals of all types divide with rate μ , with a mutation probability α that generates new types. The crucial difference with Eq. 63, beyond the absence of death, is that the rate of novelty increases linearly with the population size. Also, the population expands exponentially, and the clone size distribution only reaches a quasi steady state, after renormalizing the total population size, C_{tot} . The renormalized size distribution asymptotically follows a power law of exponent -2 in the limit $\alpha \ll 1$. Although this exponent is tantalizing close to those observed in data,

the assumption that new types arise with a rate proportional to the population size is incompatible with the biology of T cells, as new clones originate from the thymus, and not from already circulating cells. On the other hand, the Yule process may be appropriate for modeling population-wide affinity maturation of B cells, as their receptors hypermutate upon expansion.

C. Population dynamics model with external signals

To explain the power laws observed in the clone size distribution from sequenced repertoires, one needs to go beyond neutral models and introduce external signals affecting the division and death rates — or fitness — of cells as a function of their phenotypic state or of their immune receptor. Several models of clonal dynamics have been proposed to describe the growth and decline of clones in various populations (see section XII D 1 for a general discussion of competition models). We will not describe all of these models, but try and give the reader an idea of their general features. The key ingredient is that cells divide in response to antigenic stimulation, which depends on the concentration or frequency of available antigens that are susceptible to be recognized. Cells in a clone expressing receptor r die with constant rate μ as in Eq. 62, but they divide with clone-specific rate

$$\nu_r(\{Q_a\}, \{C_{r'}\}) = \nu \sum_a Q_a f_{r,a} A \left(\sum_{r'} C_{r'} f_{r',a} \right), \quad (66)$$

where $f_{r,a}$ is the cross-reactivity function between an antigen and a receptor, Q_a the frequency of antigen a , and $A(x)$ is a decreasing function describing the availability of a given antigen as it is being bound by other receptors than r ; for example, $A(x) = (1 + \epsilon)/(1 + \epsilon x)^2$ with ϵ setting the strength of competition. In this model, receptors interact only indirectly through this competition factor.

As a historical note, models with *direct* interactions between receptors, namely antibodies binding to each other, were once proposed as a way to generate interesting dynamics on so-called idiotypic networks [225]. While idiotypic networks were often used to explain phenomena in the 1980’s, and became very popular with physicists because of their link to network theory and the physics of disordered systems [209], it is not currently considered a dominant paradigm in immunology (and already was not by the time Ref. [209] was written, as discussed therein), as its predictions can be explained by clonal selection theory alone without invoking additional elements.

In the linear noise approximation, the size of the clone of cells expressing receptor r is described by a stochastic differential equation:

$$\frac{dC_r}{dt} = [\nu_r(\{Q_a\}, \{C_{r'}\}) - \mu] C_r + \xi_r(t), \quad (67)$$

with $\langle \xi_i(t)\xi_i(t') \rangle = (\nu_r(\{Q_a\}, \{C_{r'}\}) + \mu) C_r \delta(t-t')$ with Itô's convention. This general form of the model can describe a wide variety of situations. It can be used as a model for the naive T cell repertoire, where Q_a is the distribution of self-antigens for which T cells compete for survival cues. Such models of competitive exclusion [226] have been studied and used to explain biological facts for both T- [227] and B cells [228] repertoires, such as the fact that a larger diversity is beneficial in terms of overall repertoire fitness. Interestingly, the same dynamical equation emerges as a way to approach an optimal repertoire [210], as defined in Sec. VIF.

Eq. 67 can also be used to model lymphocyte population dynamics in response to transient antigenic stimulation. Antigens appear with some rate s , and decay as they are being cleared by the immune system as $Q_a(t) = Q_a^0 e^{-\lambda(t-t_a)}$, where t_a is the time at which antigen a appeared in the population. In this context, the antigen is not limiting, and is fully available, $A(x) = 1$. When λ is very large compared to μ and ν , each clone experiences a series of short spikes of expansion of magnitude $\exp(\nu\langle Q^0 \rangle/\lambda)$ with rate sp , where p is the probability that $f_{ar} = 1$ for a random choice of a and r (and $f_{ar} = 0$ with probability $1 - p$, in an all-or-nothing approximation). With the change of variable $x = \ln C$, each expansion event causes an average jump $\Delta x = \nu\langle Q^0 \rangle/\lambda$.

In Ref. [193], expansion events were assumed to be both small and frequent compared to other time scales, which led to an effective diffusion equation for x . Here, as an original result of this review, we present the solution to the general process. A more detailed solution using Laplace transforms is presented in Sec. XIID 4). The density of clones with logarithmic size x , $\rho(x, t) = (dC/dx)N_{C=e^x}(t) = e^x N_{e^x}(t)$, follows the simple jump process with constant negative drift:

$$\frac{\partial \rho}{\partial t} = \mu \frac{\partial \rho}{\partial x} + sp[\rho(x - \Delta x) - \rho(x)] + \tilde{\theta}(x). \quad (68)$$

where $\tilde{\theta}(x) = e^x \theta(C = e^x)$ is the density rate of new clones entering the population with initial logarithmic size x . Assuming a power law Ansatz for the clone size distribution,

$$N_C \propto \frac{1}{C^{1+\alpha}}, \quad (69)$$

translates into exponential decay in x , $\rho(x) = \rho_0 e^{-\alpha x}$, giving the consistency equation for α :

$$\mu\alpha + sp[1 - \exp(\alpha\Delta x)] = 0. \quad (70)$$

The total numbers of cells and clones can be calculated at steady state (see Eqs. 204,205 in Sec. XIID 4), and read in the case of fixed introduction size $\theta_C = \theta\delta_{C,C_0}$:

$$C_{\text{tot}} = \frac{\theta \ln(C_0)}{sp\Delta x - \mu}, \quad N_{\text{tot}} = \frac{\theta(C_0 - 1)}{sp(e^{\Delta x} - 1) - \mu}. \quad (71)$$

If we assume many small expansion events, $\Delta x \ll 1$, expanding the exponential at second order allows us to recover the power-law exponent of [193]:

$$\alpha = \frac{\mu - sp\Delta x}{sp\Delta x^2/2}, \quad (72)$$

where the numerator is the net decay rate of the clone size obtained from the difference between death and expansion, while the denominator corresponds to an effective diffusion coefficient stemming from the random arrival of expansion events.

The result of Eq. 69 suggests that a wide range of models with clone-dependent, random expansion or selection events produce power laws with arbitrary exponents, in agreement with the data on memory or unfractured lymphocyte repertoires. To distinguish between different types of dynamics, more details of the dynamics of individual infection events, as well as of the kinetics of unstimulated clones should be studied in terms of their impact on the clone size distribution.

The hypothesis that the expansion factor Δx is constant, or at least does not depend on the clone size, is questionable. For instance, in secondary immune responses, the expanded cells already have a memory phenotype, and these cells tend to expand less than naive cells upon a primary infection. In addition, the amount of inflammatory signals that promote expansion should in general depend on the number of cells involved in the response. From a design perspective, it would seem beneficial not to expand clones that are already big, as the organism is already well protected against the pathogens that these clones are specific to. The optimal expansion strategy upon each pathogen encounter can be calculated within the simplified framework described in Sec. VIF. Using Eq. 59 with a logarithmic cost $c(x) = -\ln(x)$ and a uniquely specific Kernel $f_{ar} = \delta_{a,r}$, combined with Bayesian prediction theory, the optimal dynamics can be shown to approximately follow [229]:

$$\frac{dC_r}{dt} = \chi \sum_{i \in E_r} \delta(t - t_i) - \tau_m^{-1}(C_r - \chi\theta_a), \quad (73)$$

where $\tau_m = 2\tau/(\chi^{-1}C_{\text{tot}} - 1)$ (with τ the effective timescale of the pathogen dynamics), E_r is the set of infection events i occurring at time t_i in response to which lymphocytes carrying receptor r expand, and χ is a scaling parameter setting the total number of lymphocytes. The main difference between these and previously considered dynamics is that the growth of C_r is not exponential, meaning that regulation or homeostasis mechanisms must exist to tune the magnitude of the expansion as a function of the size of the antigen-specific repertoire subset.

So far we have assumed that selection manifests itself at the level of clones: all cells in the same clone has the same average growth rate. However, selection could be cell dependent, which is the case if cells respond differently to non-antigenic stimulation signals, such as cy-

tokines. Growing experimental evidence shows the heterogeneous response to cytokines, due to differential expression of cytokine receptors, signalling molecules and their diffusion. Models with a cell-dependent fitness do not give power law behaviour, but do produce long tails [193]. With the current experimental cut-offs, it may not be possible to discriminate between clone-specific and cell-specific fitness models using repertoire data only. Further studies with longitudinal tracking of clone sizes may help settle these questions.

D. In host HIV dynamics.

Perelson and collaborators [120, 230] considered a modified SIS model of an HIV infection. This is one of the most influential examples of how computational models influenced medicine. It was used to predict the effects of anti-viral drugs on HIV: reverse transcriptase (RT) which blocks the ability of HIV to infect a cell, and protease inhibitors (PI) that result in the production of non-infectious viruses. Target cells T , that correspond mainly to CD4+ T-cells, become infected (I) at rate β , are born with rate α and die with rate μ . Infected cells die at rate ν , but they also help the virus (V) reproduce and produce p new virions that are cleared with rate c .

$$\frac{dT}{dt} = \alpha - \mu T - \beta IS, \quad \frac{dI}{dt} = \beta VT - \nu I, \quad \frac{dV}{dt} = pI - cV.$$

To study the effect of an antiviral drug, the virus equations are modified to account for two viral species: non-infectious viruses (V_{NI}) and infectious viruses (V_I):

$$\frac{dV_I}{dt} = (1 - \epsilon_{PI})pI - cV_I, \quad \frac{dV_{NI}}{dt} = \epsilon_{PI}pI - cV_{NI},$$

where $\beta = 1 - \epsilon_{RT}$ and $0 < \epsilon_{PI} < 1$ and $0 < \epsilon_{RT} < 1$ are efficacies of RT and PI drugs. The total viral populations is held constant $V_{NI} + V_I = V = \text{const}$. For $\epsilon_{PI} = 1$ and $\epsilon_{RT} = 1$, assuming $T = \text{const}$, the total viral load will decay according to [120, 230]:

$$V(t) = V(t=0)e^{-ct} + \frac{cV(t=0)}{c-\nu} \times \left[\frac{cV(t=0)}{c-\nu} \{e^{-\nu t} - e^{-ct}\} - \nu t e^{-ct} \right]. \quad (74)$$

This decay was fit to data from an HIV-infected patient on anti-viral therapy by adjusting the c and ν parameters using least-squares regression, showing very good agreement [231]. Explaining viral dynamics in the case of combination therapy (tri-cocktails) requires introducing long-lived infected cells that act as a secondary viral source [232]. Fitting these more complex, yet still extremely simple models to data, allowed Perelson and collaborators to estimate the half-lives of the different species in Eq. 74. This in turn made it possible to effectively administer a tri-therapy cocktail treatment in intervals that made it very hard for the virus to escape. This extremely simple calculation saved lives.




FIG. 14. **Affinity maturation in germinal centers.** Naive B-cells get recruited into germinal centers where they become activated and acquire somatic hypermutations upon proliferation in the dark zone. This produces cells with different BCRs. If the BCR manages to get expressed on the cell surface, it moves to the light zone to undergo selection for binding to antigen presented on follicular dendritic cells (FDC). Cells that bind strongly internalize the antigen and present them to helper T-cells. They undergo a second selection step where they compete to help T-cell signals. The successful cells can undergo another round of affinity maturation by returning into the dark zone to somatically hypermutate, or they can leave the germinal center as a memory B-cell or an antibody secreting plasma cell.

VIII. AFFINITY MATURATION

Upon antigenic stimulation, B-cell receptors acquire somatic hypermutations (SHM) that can help them explore the binding landscape. Based on high-throughput repertoire sequencing data, combinations of clustering and tree building methods [233, 234] have been proposed to characterize lineages. Since affinity maturation is a fascinating example of a Darwinian evolutionary process, a lot of effort is going into extracting the details of the evolutionary processes from data, along with more theoretical efforts. We very briefly review tree building, clustering and lineage reconstruction approaches that are available in different software in section XII C 11. A recent review more thoroughly summarizes these attempts [235], arguing for methods departing from the traditional assumptions of population genetics, and tailor-made for the hypersomatic mutation scenario. BCR repertoire sequencing provides large amounts of data that allow software packages like Immcantation [234, 236–240], Partis [179] (see section XII C 11) SPURF [241] or IGoR [175, 233] to learn hypermutation models (as described below), and understand the evolutionary process. As mentioned in section XII C 11 the problem of lineage reconstruction is quite general to many areas in immunology. Incorporating the observation that more abundant clones are likely to have more offspring allows for reconstructing B-cell lineage ancestries from germinal center imaging experiments [242].

Affinity maturation is a sped-up Darwinian evolution process. Upon recognizing an antigen even with low affinity [243], naive B-cells migrate to germinal centers (GC) in lymph nodes and about 3 days after antigen injection the germinal center reaction starts [244], initially dividing without hypermutations (see section VD1) for about 7 days and reaching a population size of ~ 1500 cells [243]. About a week after the infection starts, they first enter the dark zone region of the germinal center, where the Activation-induced cytidine deaminase (AID) enzyme introduces hypermutations as the cells continue to divide.

A. Modeling the hypermutation process

Repertoire sequencing analysis of unselected mutations [170, 175, 234] has shown that AID acts highly non-uniformly, with hypermutation hotspots and strong context dependent motifs. The difficulty in learning these models stems from eliminating selection that acted on memory sequences. A 5-mer model, called S5F, was learned from synonymous mutations (that are in principle free of selection) from long heavy chain sequences [234]. Only half of possible 5-mers were observed from data with synonymous mutations, the remaining ones were inferred by averaging over related observed pentamers. Alternatively, out-of-frame memory sequences were used to learn heavy V and J segment di- and tri-nucleotide context dependent hypermutation models from low throughput sequencing experiments [245].

The S5F model learns the probability $P_{\text{mut}}(\vec{\pi})$ of mutating the central position of a 5-mer motif $\vec{\pi} = (\pi_{-2}, \dots, \pi_2)$ by counting occurrences of synonymous mutations in this motif centered at position i along the sequences $\vec{\sigma}^j$ of a large dataset. Because only synonymous mutations must be taken into account, *ad hoc* heuristic normalization rules were used. In practice, mutation rates are given by:

$$P_{\text{mut}}(\vec{\pi}) = \mu \frac{m_{\vec{\pi}}}{\sum_{\vec{\pi}'} m_{\vec{\pi}'}} \quad (75)$$

with μ an adjustable baseline mutation rate (not learned within S5F), and

$$m_{\vec{\pi}} = \frac{1}{\sum_{j, \pi'} C_{\vec{\pi}'}^j} \sum_j \frac{C_{\vec{\pi}}^j / B_{\vec{\pi}}^j}{\sum_{\vec{\pi}'} C_{\vec{\pi}'}^j / B_{\vec{\pi}'}^j} \sum_{\vec{\pi}'} C_{\vec{\pi}'}^j \quad (76)$$

is the expected number of mutations of each 5-mer in a sequence. This number is obtained by normalizing for being synonymous in each sequence, and then averaging over sequences weighted by their number of synonymous mutations. In that expression, $C_{\vec{\pi}}^j$ is the observed number of synonymous mutations in sequence $\vec{\sigma}^j$ in motif $\vec{\pi}$, and $B_{\vec{\pi}}^j$ is a background normalization:

$$C_{\vec{\pi}}^j = \sum_i I_{\vec{\pi}^j}(i, \vec{\pi}, \sigma_i^j), \quad B_{\vec{\pi}}^j = \sum_{i, b} s_b^{\vec{\pi}} I_{\vec{\pi}^j}(i, \vec{\pi}, b) \quad (77)$$

where: $\vec{\pi}^j$ is the germline (unmutated) ancestor of $\vec{\sigma}^j$; $I_{\vec{\pi}^j}(i, \vec{\pi}, b)$ is an indicator function that is equal to 1 if $\vec{\pi}^j$ matches $\vec{\pi}$ over the position range $(i-2, \dots, i+2)$, and $\tau_i^j \rightarrow b$ is a synonymous mutation; finally $s_b^{\vec{\pi}} = N_b^{\vec{\pi}} / \sum_{b' \neq \pi_0} N_b^{\vec{\pi}}$ and $s_{\pi_0}^{\vec{\pi}} = 0$, where $N_b^{\vec{\pi}}$ is the total number of synonymous mutations to b observed in 5-mer motifs $\vec{\pi}$ across all sequences in the training dataset.

The S5F model provide a profile of hypermutation hot and cold spots in the absence of selection, which is widely used for analyzing BCR data. In particular, it can used as a baseline for quantifying selection from synonymous mutations [246, 247]. Note that a 5-mer model was also learned from VH and JH genes sequences in rearrangements engineered not to be productive, in both the heavy and light chains of mouse, without having to use the trick of using synonymous mutations [239].

To avoid for over-fitting, one can further assume that the impact of the motif on mutability is additive. This strategy was applied to non-productive heavy-chain sequences in humans [175]. Mutability of $(2m+1)$ -mer motifs $\vec{\pi}$ is given by:

$$P_{\text{mut}} = \frac{\mu \exp(\sum_{i=-m}^m e(\pi_i))}{1 + \mu \exp(\sum_{i=-m}^m e(\pi_i))}, \quad (78)$$

where (π_{-m}, \dots, π_m) is the $2m+1$ -mer sequence context around the mutation i , and $e(\pi_i)$ are the inferred elements of the Position-Weight Matrix (PWM), learned within the IGoR framework [175] (see Sec. VI B) using the Expectation-Maximization algorithm (see Sec. XII E 1).

The actual hypermutation process is yet more complicated than captured by all these models. It has been pointed out that because of the context dependence of mutations, the order in which they appear matters. Since considering all possible histories is computationally not feasible, Feng et al proposed to Gibbs sample the orderings to learn a hypermutation model that was validated on simulated data [248]. More importantly, hypermutations are not independent of each other, and tend to cluster along the sequence [175]. Hypermutations result from a lesion of DNA followed by error-prone DNA repair over extended regions (~ 20 nucleotide) of the sequence [249]. These events may induce several simultaneous mutations at close-by positions along the sequence. It also suggests that context can influence the mutation rate over fairly large distances, since the erroneous repair can occur far from the original damage causing the lesion. Finally, context models cannot explain all of the variance of hypermutation rates, suggesting that other factors, such as length dependence, may be at play.

B. Cycles of selection

After acquiring hypermutations B-cells pass to the light zone of the germinal center where they undergo the selection step of Darwinian evolution. In this step

a B-cell receptor must bind the antigen presented on a follicular dendritic cell (FDC) strongly enough for the B-cell to internalize the antigen and present them to helper T-cells. B-cells compete for binding of their pMHC to these helper T-cells that give them an essential survival signal. Therefore each B-cell passes a two-fold selection step: selection for recognizing an antigen, and selection for competitive binding to a helper T-cell. Binding to T-cells is specific and depends on the peptide presented on the MHC. B-cells that have a higher affinity for the antigen, have a larger internalization rate and a larger probability to successfully solicit T-cell help. Helper T-cell stimulation is essential for survival. One unverified hypothesis is that helper T-cells that have previously been trained in the thymus against self-proteins, now discriminate against hypermutations that lead to self-reactive B-cells. B-cells that do not receive a survival signal die.

Theoretical work [250, 251] that was later verified by experiments has shown [252] show that a single cycle of passing through the light and dark zones is not enough to generate the numbers of observed high affinity B-cells. This leads to a hypothesis called recycling where $\sim 90\%$ of B-cells go back to the dark zone. At the end of the affinity maturation process, B-cells increase their affinity to the antigen by ~ 1000 fold [251, 253] (or even ~ 10000 in rabbits [254]) and the B-cell population is always oligoclonal [255], with a large diversity of different B-cell receptors [256]. Interestingly, a germinal center is binary: either it produces many high-affinity cells, or it fails to produce high-affinity B-cells [257].

The details of the heterogeneity of GC have been discovered through so-called *rainbow* experiments [255, 258], in which individual B-cells were permanently tagged during BCR acquisition with one of 10 possible combinations of four different colored fluorescent proteins. This technique combined with imaging allowed for multicolor fate mapping of a cell's progeny since all cells that come from the same ancestral cell have the same color. By injecting cells in germinal centers with different dyes and tracking their movement in mice, Victora et al [255] have shown that in one lymphnode there are many germinal centers active at the same time, producing different diversities of cells. Most GCs probed at both 6 and 15 days showed a lot of different colored clusters, with lower estimates of ~ 50 clones per GC, going up to hundreds. By delaying tag formation until the cells are in the GC, selection in GCs was shown to keep multicolored clusters, showing that diversity of lineages is always maintained even if a certain dominance of one specific lineage (typically less than 40% of all cells belong to the dominant lineage) is observed at later times. Notably the heterogeneity was huge: some GC had one lineage making up to 80% of cells, others as little as $\sim 20\%$ two weeks after the infection. The rate of diversity loss was also very heterogeneous: some GCs converged to a single lineage within a few days, and others took over two weeks, or never converged. Even in adoptive transfer experiments (see section V), where high affinity clones were intro-

duced and the mouse was challenged with either bacteria or virus antigens, the GCs showed great heterogeneity regardless of the type of antigen - from nearly monoclonal to ones with an effectively neutral distribution of lineage diversity. The GCs expressing one lineage can be traced back to one single ancestral BCR (with possible SHM) that expands in a clonal *burst* over a short period of time, leading to a loss of diversity. Interestingly, the affinity of BCRs coming from GCs with varying levels of lineage diversity is similar. The oligoclonality of antibody repertoires is also confirmed by RepSeq experiments that show that antibodies responding to a specific antigen found at physiologically relevant concentrations can have between 3 – 147 distinct CDR3s [259].

C. Evolution of broadly neutralizing antibodies

Many pathogens such as viruses (influenza, VIH) come in a variety of strains. Typically, antibodies mature to recognize just one type of strain, by targeting easily accessible epitopes on the surface of its proteins. However, viruses can often easily escape immunity afforded by these antibodies by mutating these epitopes, with little fitness cost. Epitopes that are more conserved (i.e. in which mutations carry a significant fitness cost to the virus) would be better targets, but they are usually harder to access by antibodies, as viruses have evolved to hide those conserved regions. To mature antibodies that bind strongly to more than one strain (called broadly neutralizing antibodies — BnAbs), the immune system needs to be trained with different antigens from different strains. In the context of a vaccination strategy, a crucial question is in what antigens to present, and in what order.

To address this question in the context of HIV, Wang et al [260] simulated different temporal immunization schemes, by considering three antigens (wildtype HIV antigen and 2 mutants) and modelling their interaction with antibodies using a BCR-antigen binding model such as described in section IID. The virus was modeled as consisting of a conserved part, in which mutations are deleterious, and a variable part. The three considered immunization schedules were: all three variants together (scheme 1), WT with mutant 1 at the same time, followed by mutant 2 (scheme 2), WT followed by mutant 1 followed by mutant 2 (scheme 3). Running these simulations many times, they looked at the antibodies that come out of these *in silico* germinal centers and scanned them against a standard panel of antigens that are used in affinity maturation experiments. The breadth of each antibody is defined by how many different test antigens it recognizes. Scheme 1 does not produce any broadly neutralizing antibodies (bnAbs). In fact, it hardly produces any germinal centers that have any output as B cells die quickly during the process. The frustration of not being able to discriminate between conserved and non-conserved residues results to a situation where the

antibodies are not really selected since a BCR is likely seeing very different antigen in each round of selection. Scheme 2 does produce bnAbs, but with very low probability. This is likely what happens during a normal infection because you get infected with a single strain that later diversifies. In this case, the affinity maturation process in the first round does produce binders to the conserved residues, breaking the frustration. Finally, scheme 3 produces bnAbs with high probability ($\sim 69\%$). The sequential application allows the evolving BCR to focus on the only non-moving part of the target — the conserved residues. These predictions have been tested experimentally, showing that mice that were sequentially immunized focus their immune response.

There are effectively two parameters that govern the B-cell germinal center distribution: the probability of surviving selection P_s and the probability of mutation μ . With probability $P_s(1 - \mu)$ the existing B-cells will expand, with probability $P_s\mu$ it will mutate, and with probability $1 - P_s$ it will die. Wang et al [260] suggest this system is optimally frustrated and the probability of survival puts the system in a special regime: too stringent selection (small P_s will kill all cells), while too lenient selection (large P_s) will fill the germinal centers with existing cells, since μ is small. So only intermediate levels of P_s lead to mutant cells surviving, allowing for many rounds of selection that results in acquiring many mutations that give large breadth. Support for this intermediate level selection comes from simulations that show that intermediate antigen numbers result in the largest number of mutations in a BCR. However experimental validation is still lacking.

Repeated antigen exposure helps increase the affinity of cells through selection, as in Darwinian evolution. In addition, new naive cells are recruited, adding novelty in the selection process other than through mutation. Murugan et al [261] compared the impact of selecting naive clones to that of producing new mutants by SHM. Looking at the response of BCR to the malaria parasite in humans coupled with string models of affinity maturation, they showed that antigens with low complexity are much more efficient in generating good binders by means of SHM affinity maturation compared to high complexity antigens, which rely on recruitment of new naive cells. Both routes (SHM and selecting naive cells) were observed experimentally. Influenza vaccine studies show that, upon secondary immunization (booster vaccines), $\sim 7\%$ of BCR contained no SHM [262] (although these experiments showed strong clonal dominance with $\sim 10\%$ of clones representing $\sim 90\%$ of sequences [259]). The model of Murugan et al [261] predicts that long exposure to small amounts of antigen leads to selecting for antibodies with a lot of SHM, whereas short exposure to high antigen concentrations results in selection on existing naive cell variation, which is confirmed in experiments. These results, which hold for relatively short timescales (\sim few months), are consistent with the argument of Wang et al [260] in the case of chronic in-

fections. Interestingly, both Murugan et al [261] and Neu et al [262] experimentally found that the majority of SHM do not necessarily improve antigen binding, and that both routes can lead to equally high affinity mutants, which supports the heterogeneity observed by Tas et al [255] in the brainbow experiments. Yet, it is still not clear why some individuals fail to produce any good binding antibodies to the very strong antigenic challenge of malaria.

Additionally, the BCR repertoire changes with age, with both positive and negative selection acting differently in older people [263], producing longer CDR3s that are more promiscuous and likely to bind self-proteins. The rate of observed BCRs with SHM also changes from infancy to adulthood [264]. In influenza vaccine studies [262], people born before 1977 show less inter clonal diversity but achieve the same affinity as people born after 1977. Adaptation via the SHM route is less likely in older people, and they are more likely to use their larger existing memory pools than younger people whose antibodies acquire a lot of SHM. As a result, the antibodies of the older group target a conserved part (the stalk) of the influenza protein as do BnAbs, whereas younger people's antibodies target the head (aiming for a strain-specific antibody). Younger people can produce BnAbs when hit with a strong challenge (vaccine combined with a pandemic), as predicted by the analysis of Wang et al [260], but the response to subsequent infections is still specific. In summary, these results highlight the complexity of BCR maturation, which involves both naive cells (with no SHM) and highly mutated cells. The details of the response depend on the dose concentration and timing, past infections, age and antigen complexity.

D. Population genetics approaches to affinity maturation

1. Evolutionary analysis of repertoire dynamics

Recent RepSeq data was used to quantify the evolutionary regime of in-host HIV evolution [265] and in response to influenza vaccines [266]. The flu vaccine is a model of an acute infection, which occurs on short timescales. BCRs from 9 timepoint blood samples, before and after the infection over the course of ~ 2 weeks, resulted in very skewed lineage trees and a U-shaped site-frequency spectrum (SFS). SFS corresponds to the distribution of frequencies of mutations in a population, and its shape is often informative about the underlying evolutionary process (see section XII C 8). U-shaped SFS, as observed in BCR lineages after vaccination, are characteristic of strong selection. However, 3 out of 5 of the strongly selected BCRs recognized neither the flu vaccine epitope nor the full virus, suggesting bystander evolution, by which non-responsive clonotypes expand in response to a multitude of signals. This could be a mechanism for upkeeping memory clones between infections with the

same or similar pathogens.

Skewed trees and U-shaped SFS were also observed in the BCR lineages of HIV-infected individuals [265]. However, the same observations were made in out-of-frame sequences of healthy individuals. This is not surprising because the repertoires of healthy people are under constant selection, and out-of-frame sequences evolve in cells with functional receptors, and their evolution reflects the hitchhiking and selection on the whole cell. However, a finer analysis exploiting the dynamic trends of synonymous versus non-synonymous mutations is able to distinguish selection due to the chronic HIV infection from overall selection patterns. More interestingly, it finds that the CDR3 region of the BCR of untreated HIV carriers evolves according to a regime known as “clonal interference” (see section XII C 9). In that regime, several new beneficial mutations arise at similar times and compete with each other. This competition slows down adaptation, as only one the beneficial mutations can survive. This analysis is based on estimating the probability that a (beneficial) mutation first rises to a threshold frequency x , but then is driven to extinction by a competing mutation, $H(x, x_i) = G(0|x)G(x|x_i)$, where $G(x|y)$, called a propagator is the probability of ever reaching x starting at y , and x_i is the initial frequency of the mutant. Both $H(x, x_i)$ and $G(x|x_i)$ can be estimated from the data and for various models of evolution. Neutral models cannot explain empirical observation, and neither does a simple model of selection with a fitness advantage. Instead, data can be fit by a model of varying selection (see section XII C 5), where the fitness advantage fluctuates as competing mutants come and go.

Analysis of phylogenetic trees build from RepSeq data from HIV patients combined with S5F hypermutation models [240] also led to identifying the constraints on BCR adaptation [267]. Ancestral sequences in lineages are more likely to mutate their CDR3s than the framework (FWR). The propensity of a given residue to mutate was decreased more in framework (FWR) regions than in CDR3s, but in both cases a decrease in mutability was much more likely than an increase. Although most of the constraints on the residue to mutate in CDR3 that are under strong positive selection to increase their binding affinity, come from nonsynonymous mutations, in all regions up to 21% of loss in residue mutability was caused by synonymous mutations, which are the result of neutral evolution. These results also point to a slow down in adaptation from clonal interference, as expected in the clonal interference regime.

2. Models of co-evolution of phenotypic traits

The fitness function for the co-evolution between HIV and BCR depends on the ability of the BCRs to recognize the virus, which is a (possibly nonlinear) function of the binding affinity. By focusing on a single antibody lineage and a viral population, Nourmohammad et al [268]

derived an effective stochastic equation for the rescaled mean binding energy ϵ between antibodies and viral epitopes:

$$\frac{d\epsilon}{dt} = -2[\theta_A + \theta_V(N_A/N_V)]\epsilon + s_A\sigma_A - s_V\sigma_V + \sqrt{\sigma_A + N_A/N_V}\sigma_V\xi_\epsilon, \quad (79)$$

where σ_i are the variances of the binding energy across the viral and antibody populations, θ_i , s_i and N_i are the mutation, selection coefficients and population sizes of antibodies ($i = A$) and viruses ($i = V$); ξ_ϵ is Gaussian noise and times is measured in units of antibody coalescence time. This equation was derived by assuming a linear relationship between binding energy and fitness, and using a additive string model of binding energy as a function of genotype (see section IID 2). The approximate method to get from the evolutionary dynamics, which are defined on the genotype, to an effective equation for the phenotype (binding energy), is described in a simpler context in XII C 10.

Intuitively, in general mutations on both antibodies and viruses reduce recognition. On the other hand, antibody diversity (σ_A) increases binding energy by selecting the best binders, while virus diversity (σ_V) decreases it by selecting the best escapers on the viral side of co-evolution. The same formalism can be applied to the evolution of BnAbs, which target epitopes on the conserved regions of the virus. In that context, the equations simplify to $\theta_V = 0$ and $\sigma_V = 0$, because the viral epitope is constant.

However, tracking the binding energy is not alone a good signature of co-evolution. Time-shifted statistics, such as the time delayed viral fitness (see Fig. 15), can be used to identify strongly co-evolving populations. The time delayed viral fitness is defined by the interaction between the viral population at time t , with distribution $y^\gamma(t)$ of genotypes γ , and a population of antibodies α taken at a later time $t + \tau$ distributed according to $x^\alpha(t + \tau)$: $F_{V,\tau}(t) \sim -s_V(t) \sum_{\alpha\gamma} E_{\alpha\gamma} y^\gamma(t) x^\alpha(t + \tau)$, where $E_{\alpha\gamma}$ is the binding energy from the string model between antibody α and viral epitope γ . Note that same-time fitness $F_{V,0}$ is equivalent to $-\epsilon$ up to a selectivity constant.

If the selective effects on the phenotype are comparable in the two populations, $s_A\theta_A \approx s_V\theta_V$, the current virus has highest fitness against antibodies from the near past since it has acquired mutations to escape them, and smaller fitness against antibodies from the future, which have caught up with these escape mutations (Fig. 15). For long times, mutations randomize the genome and $F_\tau(t \rightarrow \infty) = 0$. If the selection pressure of viruses is larger than that of the antibody $s_V\theta_V \gg s_A\theta_A$, the virus is effectively evolving against a neutral antibody population. It still has a high fitness against previously seen antibodies, and genome randomization decreases its fitness advantage against future antibodies, but without a penalty.

The left slope of the time-shifted viral fitness quantifies the adaptation of the viral population to the exist-




FIG. 15. **Time delayed viral fitness.** Adapted from Nourmohammad et al [268]. Viruses have a positive fitness against antibodies from the past, and a negative fitness against antibodies from the future. The slope of the current viral fitness coming from past antibodies is equal to the fitness flux $\phi_V(t)$ – a measure of adaptation to existing antibody populations, and the slope of the current viral fitness coming from future antibodies is equal to the transfer flux $\mathcal{T}_{A \rightarrow V}(t)$ that describes the pressure on the viral population from the adapting antibody population. The blue line describes a co-adapting population, $s_A \theta_A \approx s_V \theta_V$, the red line describes a virus evolving effectively without antibody selection pressure, $s_V \theta_V \gg s_A \theta_A$ and the green line is a regime of much stronger antibody adaptation, $s_V \theta_V \ll s_A \theta_A$.

ing antibody population and defines the a “fitness flux” $\partial_\tau F_{V, \tau}(t - \tau)|_{\tau=0^-} = \phi_V(t) = \sum_\gamma \partial_{y^\gamma} F_V(t)(dy^\gamma(t)/dt)$. The right slope measures the pressure on the viral population from the adapting antibody population, which defines the transfer flux of fitness from antibodies to viruses $\partial_\tau F_{V, \tau}|_{\tau=0^+} = \mathcal{T}_{A \rightarrow V}(t) = \sum_\alpha \partial_{x^\alpha} F_V(t)(dx^\alpha(t)/dt)$. At stationarity both fluxes sum to zero, and the derivative is continuous, as in Fig. 15.

The characteristic S-curve shown for $s_A \theta_A \approx s_V \theta_V$ is indicative of two competing populations. Neutralization measurements [269–272] show that viruses are more resistant to past antibodies and more susceptible to future antibodies, which results in the S-curved time delayed viral fitness [268]. Interestingly, Blanquart and Gandon [269] also showed that antibodies that evolved in one HIV positive patient are better at targeting the virus found in another person than this person’s current antibodies. This result remains unexplained.

While this analysis only involves a single antibody lineage, in germinal centers multiple lineages can compete with each other. Extending their analysis based on

the dynamics of quantitative traits, Nourmohammad et al [268] considered the competition of different lineages during affinity maturation. The change in the frequency of a given antibody lineage (each composed of many genotypes α) of size N_A^C in the population $\rho^C = \sum_{\alpha \in C} x^\alpha$ is driven by its mean fitness F_A^C compared to the mean fitness of antibody lineages $F_A = \sum_C F_A^C \rho^C$:

$$\frac{d\rho^C}{dt} = (F_A^C - F_A) \rho^C + \sqrt{\frac{\rho^C(1 - \rho^C)}{N_A}} \xi_C, \quad (80)$$

where ξ_C is a Gaussian white noise. The lineage fitness is an average of the genotype A^α fitnesses $f_C^\alpha(t)$ that make up the lineage C, weighted by their frequencies in that lineage $x_C^\alpha(t)$, $F_A^C = \sum_\alpha f_C^\alpha(t)x_C^\alpha(t)$. The mean fitness of a lineage depends on the genotypes within this lineage and also on the frequency of virus y^γ through $f_C^\alpha(t)$.

The probability of fixation P_{fix}^C of lineage C is not only a function of the mean fitness advantage and population size, as in Eq. 114 (Sec. XII C 4), but also of the ability of the population within the lineage to adapt. Its value is given by:

$$P_{\text{fix}}^C/P_{\text{fix}}^0 \approx 1 + \langle N_A (F_A^C(t=0) - F_A(t=0)) \rangle + \frac{N_A^2}{3} \langle \phi_A^C(t=0) - \phi_A(t=0) \rangle - N_A N_V \langle |\mathcal{T}_{V \rightarrow A}^C(t=0)| - |\mathcal{T}_{V \rightarrow A}(t=0)| \rangle, \quad (81)$$

where $P_{\text{fix}}^0 = \rho^C(t=0)$ is the neutral fixation probability, and where ϕ_A^C and $\mathcal{T}_{V \rightarrow A}^C$ are the antibody counterpart of ϕ_V and $\mathcal{T}_{A \rightarrow V}$ for a specific lineage, and ϕ_A and $\mathcal{T}_{V \rightarrow A}$ their average over lineages. The first corrective term is standard and comes from the mean fitness advantage. The second corrective term accounts for the adaptation of each lineage in response to selection of antibodies for viral recognition within lineages. The third corrective term corresponds to the adaptation or escape capability of the viral population and its effect of the lineage.

For typical antibodies targeting a variable region of the virus, a broad viral diversity is detrimental as it allows the viral population to escape immunity (large $\mathcal{T}_{V \rightarrow A}^C$). For BnAbs, which bind to constant regions of the virus, this effect disappears. In fact the effect is the opposite: higher diversity increases the probability of BnAb fixation compared to a non-BnAbs antibody, whereas at low viral diversity the probability of fixation of both a non-BnAb and BnAbs is the same. This argument was proposed to explain why BnAbs show up later in the infection when the viral diversity is large.

IX. POPULATION DYNAMICS OF PATHOGENS AND HOSTS

A. Viral fitness models

The immune system of a population of hosts defines a fitness landscape F in which viruses evolve. Models have been developed to describe how this immune pressure shapes the evolution of both flu [273] and HIV [274–276] at the population level. Viruses are usually taken to evolve according to traditional population genetics models.

This formalism can be used to draw short term predictions about the future fate of existing strains, e.g. of circulating influenza over the course of a year. Viral populations are subject to stochastic Wright-Fisher dynamics, but once they reach large enough numbers they grow exponentially following the fitness function (see section XII C 3). In that deterministic stage, the size of a given viral strain α , $X_\alpha(t)$ grows (or decreases) according to the fitness it experiences:

$$X_\alpha(t + \Delta t) = X_\alpha(t) \exp(F_\alpha \Delta t). \quad (82)$$

Alternatively, one can consider the longer-term prevalence of particular strains of very diverse viruses such as HIV across many hosts, and use that to infer the fitness landscape of the virus. To do this, we can assume that the population reaches an “equilibrium” as the result of many mutations and fixation events across time in the entire population of viruses and hosts. The probability of finding a given viral strain is then assumed to follow Boltzmann’s law:

$$p_\alpha \approx \exp(\beta F_\alpha), \quad (83)$$

where β plays the role of an inverse temperature setting the tension between entropy from mutations, or genetic drift, and the fitness advantage. One possibility is to make use of the equilibrium assumption and learn the fitness Hamiltonian directly from existing viral data. This has been done for the Gag envelope protein of HIV [274] using maximum entropy approaches (see section XII E 1), such that the probability of seeing a given HIV strain α defined by its amino acid sequence $\vec{\sigma}$, is $p_\alpha = p(\vec{\sigma}) = Z^{-1} \exp(\sum h_i \sigma_i + \sum_{i,j} J_{i,j} \sigma_i \sigma_j)$, where h_i define single-site fields that constrain the probability of observing the wildtype amino acid at a given site i , and J_{ij} are interaction terms between amino acids (assuming a binary representation of the virus with $\sigma_i = 1$ denoting the consensus amino acid and $\sigma_i = 0$ a mutation at time i). Such models have been inferred from multiple sequence alignment of the protein, and then used to evolve *in silico* viral proteins in a population of host individuals [274, 276], using quasi-species equations introduced by Eigen [277, 278], to make predictions about HIV evolution.

Instead of learning the fitness directly from the sequence data, another strategy is to derive the fitness of each strain by decomposing it into two components

$$F_\alpha(t) = F_\alpha^{\text{ag}}(t) + F_\alpha^{\text{stability}}, \quad (84)$$

where the $F_\alpha^{\text{ag}}(t)$ is due to the antigenic component of the immune pressure exerted by the on the virus, and $F_\alpha^{\text{stability}}$ encodes the fact that mutations that change the ability of the protein to fold should be detrimental to fitness. The fitness due to folding depends on the free energy of folding G , assuming a two state thermodynamic model of folding [279]:

$$F_\alpha^{\text{stability}} \sim (1 + \exp[(G_\alpha - \bar{G}/G_0]))^{-1}, \quad (85)$$

where \bar{G} and G_0 set the energy scale and are learned from data. Similarly, the antigenic component of the fitness is taken to be a sigmoidal function of the binding affinity H :

$$F_\alpha^{\text{ag}} \sim (1 + \exp[(H_\alpha - \bar{H}/H_0]))^{-1}, \quad (86)$$

where again \bar{G} and G_0 set the energy scale and are learned from data. These constants are learned from data assuming that non-epitope mutations in the virus decrease protein stability, whereas epitope mutations decrease binding affinity. This approach was applied to influenza by Luksza and Lässig [273]. The antigenic component F^{ag} was inferred from ferret blood titers, in which sera of blood from ferrets are challenged with different viral strains and their antibody response measured. To gain precision, the growth prediction of Eq. 82 were considered at the level of clades rather than individual strains by summing over all strains in each clade. In the next section we will see how this model could be used to predict the upcoming dominant strain of influenza.

B. Co-evolution between host and pathogen populations

The adaptive immune system of hosts also evolves under the selective pressure of the antigenic environment, by expanding immune receptors that led to successful recognition. The antigens and the immune systems of individuals in a population engage on an arms race, where one forces the evolution of the other. The interaction between the immune system and different viruses takes on different forms. While HIV strains undergo very dynamic in-host evolution [270] experiencing many mutations and strong selection accompanied by clonal interference [280], the influenza virus evolves more slowly, and at scale of the whole population of hosts, with few mutations in the same host. We have discussed HIV evolutionary models in section VIII. Here we turn to the effects of population level co-evolution on immune repertoires, characteristic of viruses such as influenza.

This is currently an emerging field. Most of the treatments of immune systems so far have been coarse-grained and reduced to the effective selective pressure they exert on pathogens. This pressure is seen on antigenic maps, which place pathogens in a common space according to the similarity of the immune response to them (measured

in the sera of ferrets). In practice, antigenic maps are produced using dimensionality reduction algorithms to reduce the response to a two dimensional manifold. These methods have been extremely useful in tracking influenza evolution [281] and show that single point mutations can result in a completely new response (new faraway cluster), while some multiple mutations do not change the type of response (same cluster).

However, the link between the molecular interaction between immune receptor (both BCR and TCR, as introduced in section IID 2), and the these phenotypic maps remains to be explored. An outstanding question is what evolutionary constraints these molecular details impose. For example, is the order of mutations important, and are all the mutations independent? Deep mutational scanning experimental techniques exist now to both find the best binding proteins in multiple rounds of selection experiments, and to map out the spectrum of possible solutions.

Recently researchers have very successfully predicted short-term flu evolution [273, 282]. There are two successful types of models in this class. The first uses the diversity of existing strains and extrapolates the tree branches that have recently expanded to predict the dominant strains in the near future [282]. This approach makes use of the coalescence framework described in section XII C 7 and since it relies mainly on evolutionary properties, can easily be extended to other globally evolving viruses. The main idea is to statistically infer the characteristics of the evolutionary process from the existing influenza trees, with no reference to the immune system. The second [273] is based on molecular information about the influenza antigen [281, 283] and identifying successful mutations. In practice these are stochastic models that assess the probability of future strains. This method incorporates an effective treatment of the immune system, based on the response of ferret serum to flu strains, as explained in the previous section. In practice this treatment ranks influenza strains by the strength of serum response, but has no information about the molecular basis of the response, potential overlap for similar strains and its evolution. It then builds a fitness model, where the future frequency of a given viral strain depends on its structural stability and ability to escape the immune system. The details of these interaction models are learned from data using advanced statistical inference techniques. The second method relies much more on biophysical details and is therefore in principle more adapted to influenza, although the approach has been extended to predict the evolution of tumorous clones [39] as a function of their immunogenicity score (see Sec. IID 4). Despite their differences both approaches are similar in style: they rely on recent evolutionary traces (from the last couple of months) to predict the dominant strain up to a year in advance. From the evolutionary perspective both models encode the idea of clonal interference between viral strains (see section XII C 9).

Another theoretical approach to the viral-immune

co-evolution problem is based on Susceptible-Infected-Recovered (SIR) Models. These models have been used for a long time to look at epidemic spreading within populations [120, 284]. Many of these approaches are based on simulating sets of nonlinear equations. Originally these models were used to study slowly evolving viruses, such as measles [284]. However recently, they have been tied with viral fitness models that account for the evolution of the virus. We already discussed one type of these models, for in-host evolution of HIV in section VII. Influenza, unlike HIV that evolves within the host organism, evolves mainly at the level of a population. Recent models [285, 286] of influenza evolution combine traditional SIR approaches with data-derived knowledge about flu strain evolution. Specifically, in Yan et al. [286], the model is of the form given by Eqs. 170-172 for each antigenic strain a , with the Susceptible equation solved explicitly, $S_\alpha \approx \exp(-\sum_{\alpha'} K_{\alpha\alpha'} R_{\alpha'})$, where S_α and R_α are the fractions of individuals who are susceptible and recovered from (and thus immune to) strain α , respectively. $K_{\alpha\alpha'} = e^{-|\alpha-\alpha'|/d}$ defines a cross-reactivity Kernel of range d , where distances correspond to the number of mutations in an infinite-genome model. This equation means that individual who are outside of the cross-reactivity range of the recovered individuals are susceptible to be infected by a given strain.

Within this setup, a mutation in the virus introduces a new strain α' that increases its distance to all existing strains by one. In a perturbative limit, this mutation increases fitness by a fixed and constant amount $s \propto d^{-1}$. This situation can be mapped onto a fitness wave model where many beneficial mutations compete against each other [287] (see section XII C 9). In the fitness wave description, the viral population has a distribution of fitnesses and reproduces according to that fitness. The dynamics are dominated by events happening at the stochastic nose of the distribution, i.e. strains with the the largest fitness f_m . The beneficial mutations that govern the future fate of the population all occur at that stochastic nose, with rate $r \sim f_m / \ln(f_m/\mu)$, where μ is the mutation rate.

The cumulative effect of these stochastic beneficial mutations at the nose can be described by an effective Langevin equation:

$$\frac{df_m}{dt} = \frac{sf_m}{\ln(f_m/\mu)} - I_{\text{tot}} + s\xi(t), \quad (87)$$

where $\xi(t)$ is a Gaussian white noise of amplitude $\langle \xi(t)\xi(t') \rangle = r\delta(t-t')$, and where the negative fitness term $-I_{\text{tot}}$ comes from the population of immune systems catching up with the viral strains in the bulk of the distribution, decreasing the fitnesses of all existing strains. This Langevin description makes it possible to estimate the rate of “speciation”, whereby two mutant strains at the nose escape the influence of the current immune systems in different way, creating two independently evolving antigenic niches which separate once their genetic distance reaches at least $\sim d$, $r_{\text{sp}} \sim re^{-d/q}$, where

$q = f_n/s$ is the typical number of beneficial mutations that strains at the fitness nose accumulate relative to the rest of the population.

These arguments can be used to explain the deep splits in flu strain topologies that have been observed in data and simulated using similar types of equations [288, 289]. In other words, the traveling wave picture also explains how influenza can constantly escape the immune system without continuous accumulation of genetic diversity.

Rouzine and Rozhnova [285] applied a similar mapping onto a fitness wave description, but in a one-dimensional antigenic space. They used the predictions of their continuous traveling wave framework to estimate an antigenic mutation rate of $\mu \sim 3 \cdot 10^{-5}$ per transmission event and predict the cross-immunity distance of $d \sim 15$ nucleotide substitutions that agrees well with independent estimates.

While these models give general insights and scaling laws to understand viral-immune co-evolution, they ignore the specific molecular details and mechanisms of immunity. As more data becomes available about the specific interactions between immune cells and viral strains, it may be possible to combine these approaches with data-driven models of immune repertoires and viral genomes, and to use them to make specific predictions about the fate of particular viral strains and their relation to immune repertoires of hosts.

X. DISCUSSION

Despite the length of this review, there are many subjects we did not touch upon that definitely fall into the global topic of quantitative immunology. Some of these, like the whole area of transcriptomics [290] applied to immunology, are new and while the methodology is quantitative, the current experiments are only just starting to give quantitative models of the immune system that can be linked with a physical understanding (although things are moving so fast that by the time this review is published, this sentence may be obsolete). We also did not go into the methodology of certain analysis or theoretical approaches that are widely used, such as machine learning [291] or stochastic gene expression and biochemical regulation [292, 293], because detailed reviews for the physics audience exist for these topics. We refer the curious reader to these papers for the necessary background. Lastly, we only briefly mention some amazing experimental advances, such as imaging [252] since they are currently in the process of being used to verify quantitative models. We hope we have managed to give the idea of a vibrant and multi-direction field. We also note that we made presentation choices, which were not easy, because many of the presented topics are linked to other ones. These links will surely become better explored in the coming years as solid experimental quantification and validation of theoretical models becomes the norm.

XI. GLOSSARY

Immunology can be painful at time for physicists because of its “jargon”. To help with accessibility, we try not to focus on specific molecules, but sometimes we need to. So here we give a brief overview of the main players and rules of the game for the immune system, and also other molecules that are mentioned in the text. For an introduction to immunology, we strongly recommend reading “Your Amazing Immune System” courtesy of the European Federation Immunological Society <http://www.oegai.org/oegai/2-PDF/AmazingImmuneSystem.pdf> [294]. After this first glance, we recommend the short but illuminating book “How the immune system works” by physicist Lauren Sompayrac [1].

- Adaptive Immune Response – a response of the organism to specific pathogens that changes in the lifetime of each individual. It is based on lymphocytes (T and B-cells) recognizing pathogens, proliferating and then keeping a subset of cells called memory cells that are adapted to previously encountered pathogens. Recognition involves interactions with many components of the innate immune system (e. g. antigen presenting cells, cytokines) and well as with different cells of the adaptive immune system. Adaptive immune systems first appeared in jawed vertebrates.
- Innate Immune Response – the innate immune response involves non-specific types of defense, from physical and chemical barriers (e.g. skin, clotting, scratching) to cells and molecules that recognize non-specific pathogenic patterns, such as proteins of the bacterial cell envelope (LPS - lipopolysaccharides). The innate immune system is evolutionary older than the adaptive one and is found in plants, insects, funghi, as well as vertebrates.
- Cytokines - small protein secreted by leukocytes to enforce cell-to-cell communications in the immune system. Examples of specific cytokines mentioned in the text: IL-2 is a key anti-apoptotic cytokine for T cells as well as an activation cytokine for regulatory T cells; IL-7 is an hematopoietic growth (anti-apoptosis, pro-proliferation) factor for lymphocytes.
- Interleukins (IL): cytokines historically thought to be produced by leukocytes.
- Leukocytes: the white blood cells of the immune system. These include lymphocytes (B, NK or T cells), granulocytes, monocytes and macrophages (see Fig. 9).
- TCR: T cell receptor, the main receptor on the surface of T cells that recognizes pMHC as a ligand for T cell activation.

- BCR: B cell receptor, the main receptor on the surface of B cells that recognizes membrane-bound molecules as ligands for B cell activation.
- Antigen: biomolecules that trigger an adaptive immune responses. Antigens can be proteins that are recognized by antibodies (B-cell mediated responses) or short peptides that are loaded onto MHC (T-cell mediated responses).
- Epitope: part or whole of the antigen that the immune receptor bind to. For B cells, part of the protein that the B cell receptor bind to. For T cells, this corresponds to the peptide loaded by the MHC.
- APC (antigen presenting cells) – dendritic cells and macrophages. Surveilling cells that internalize molecules and cells from tissues and present them on their MHC type II to T-cells.
- MHC: major histocompatibility complex; its main function is to "present" short peptides, as antigen for T cells. It comes in three types, but two of them are more relevant for the purposes of this review: type I and type II. Type I MHC are expressed on most cells in the body, presenting random bits of protein fragments found in that cell – this informs the surveilling cells if the MHC presenting cell is healthy or not. Type II MHC are expressed only on specialized cells (antigen presenting cells) and carry information about the cellular environment (whether there is an infection in a given tissue).
- CD8(+) (killer) T-cell – a type of T-cell identified by its CD8 marker involved in interactions with MHC type I presenting cells. These T-cells trigger apoptosis (kill by forcing the cells to kill themselves) the infected cells.
- CD4(+) (helper) T-cell – a type of T-cell identified by its CD4 marker involved in interactions with MHC type II presenting cells that present peptides from antigens they engulfed. These T-cells produce cytokines and help orchestrate the response of other immune cells (e.g. B-cells, killer T-cells).
- pMHC: a complex of a short peptide (p) and MHC that constitutes a ligand for TCR. Depending on the nature of the embedded peptide, a pMHC can be agonistic (triggering an immune response), antagonistic (extinguishing an immune response) or null. Such hierarchy often lines up with self/non-self discrimination.
- Somatic (hyper)mutation: genetic alteration acquired by a cell that occurs in body cells (somaplast) and that can be passed to the progeny of the mutated cell in the course of cell division; one of the driver for antibody maturation in B cells, and for repertoire generation in T cells.
- Germline mutation: genetic alteration that occur in the gamete-producing cells (sperm and eggs)
- VDJ recombination – a DNA editing process that creates T- and B-cell receptors (see section VIB for details).
- Ig class – B-cells can express different types of constant regions of their receptors coupled to the same variable region. Throughout their lifetime, they can change (albeit in a specific order) which gene they express – this is called class switching. The expressed constant region gene determines the so-called class of the B-cell and antibody. The different genes are aligned in the immunoglobulin locus, and after the B-cell moves on the next class, previous classes get deleted. For example the μ and δ genes are first in the heavy chain locus and they lead to the expression of IgM and IgD chains, both of which are expressed on naive cells. The order of the remaining classes for the heavy locus is IgG3, IgG1, IgA1, IgG2, IgG4, IgE, IgA2. Class switching is regulated by cytokines, through regulation of gene regulation (see sections IV B and V A).

XII. METHODS

This section is devoted to introducing or expanding the more technical details of methods relevant to quantitative immunology. These methods and concepts may be of different kinds. Some pertain directly to the biophysics of immunology and the molecular details of binding between receptors and cognate ligands. Some may be classical tools from other fields, from statistical learning to population genetics, but which carry some conceptual similarity with approaches from statistical physics. The presented methods are almost always cited in the context of recent work in physical or quantitative immunology. While they are often not established concepts of immunology, we believe that they might play an increasing role in its future development, with more and more imports from other fields.

A. Physical kinetics

1. Diffusion-limited reaction rate

To compute the association rates of 2 of two molecules (Ligand and Receptor), let's place ourselves in the reference frame of the receptor, so that the receptor is immobile, but the ligand diffuses with coefficient $D = D_{\text{Ligand}} + D_{\text{Receptor}}$. The spatial distribution of the ligand concentration $C_{\text{Ligand}}(\vec{r}, t)$ at steady state obeys Fick's law:

$$\nabla^2 C(\vec{r}) = 0, \quad (88)$$

with boundary conditions $C(\vec{r}, t) \xrightarrow{r \rightarrow \infty} C_{\text{Ligand}}$ and $C(r = R) = 0$, where $R = R_{\text{Receptor}} + R_{\text{Ligand}}$ is the sum of the receptor and ligand radii (each modeled by a sphere). Because of spherical symmetry, $C(\vec{r}) = C(r)$, we can consider only the radial part of the Laplacian in spherical coordinates

$$\frac{\partial}{\partial r} \left(r^2 \frac{\partial C}{\partial r} \right) = 0,$$

and the steady-state solution is:

$$C(r) = C_{\text{Ligand}} \left(1 - \frac{R}{r} \right).$$

The flux of ligands j_{Ligand} onto the receptor sphere is

$$j_{\text{Ligand}}(r = R) = D \nabla C = \frac{DC_{\infty} R}{r^2} = \frac{DC_{\infty}}{R},$$

and integrating over the sphere surface one can compute total flux:

$$\begin{aligned} \Phi_{\text{collision}} &= \oint \vec{j}_{\text{Ligand}}(r = R) d\vec{S} \\ &= 4\pi R^2 \frac{DC_{\text{Ligand}}}{R} = 4\pi DC_{\text{Ligand}} R. \end{aligned} \quad (89)$$

The total number of collisions per unit time is $4\pi (R_{\text{Ligand}} + R_{\text{Receptor}}) (D_{\text{Ligand}} + D_{\text{Receptor}}) C_{\text{Ligand}}$ and the rate of diffusion-limited interactions between receptors and ligands $k_{\text{collision}} = \Phi_{\text{collision}}/C_{\text{ligand}}$ is:


$$k_{\text{collision}} = 4\pi (R_{\text{Ligand}} + R_{\text{Receptor}}) (D_{\text{Ligand}} + D_{\text{Receptor}}). \quad (90)$$

2. The rates of dissociation between two biomolecules

The physical chemistry of dissociation of two biomolecules can be modelled in two steps: a slow one that corresponds the breakage of chemical bonds between molecules, and a fast one that corresponds to the rushing in of the solvent. Estimating the rate of dissociation k_{off} thus corresponds to estimating the slow step, which is truly driven by the quantum physics of bond breakage. The rate is then encapsulated by the frequency of bond vibration ($k_B T / \hbar \approx 10^{12} \text{ s}^{-1}$) multiplied by the probability of successfully dissociating the molecular pair ($\exp(-\Delta G_{\text{off}} / (k_B T))$). This bond-breakage free energy sums up the contribution of all the bonds (electrostatic interactions, hydrophobic bonds, Van der Waals forces, hydrogen-bonds etc.) holding a molecular complex together: it could be computed from purely quantum considerations, independently of the context in which the pair is considered (solvated in the intercellular medium, embedded in the plasma membrane, in vacuum etc.). Yet, this remains a tricky proposition as small errors in estimating the bond energies will yield to exponentially-inaccurate estimates of k_{off} .

3. The formation of ligand-receptor pairs: equilibrium and kinetics

Once association and dissociation rates are available, estimating the kinetics of formation of the ligand-receptor pairs is done by considering the reaction:



and integrating the equation:

$$\begin{aligned} \frac{d[\text{Complex}]}{dt} &= k_{\text{on}}[\text{Ligand}][\text{Receptor}] \\ &\quad - k_{\text{off}}[\text{Complex}]. \end{aligned} \quad (92)$$

We have three species whose kinetics we are estimating, and two conservation laws (for the ligand and for the receptor), hence the reaction coordinate for this reaction is one-dimensional and can be solved analytically. We call the complex concentration $x = [\text{Complex}]$

$$\frac{dx}{dt} = k_{\text{on}}(L_{\text{total}} - x)(R_{\text{total}} - x) - k_{\text{off}}x, \quad (93)$$

where L_{total} is the total ligand concentration, free and complexed, and R_{total} is the total receptor concentration. In steady state (for $t \rightarrow \infty$) we have $dx/dt = 0$ and

$$\frac{(L_{\text{total}} - x)(R_{\text{total}} - x)}{x} = \frac{k_{\text{off}}}{k_{\text{on}}} = K_D, \quad (94)$$

which has two formal solutions

$$\begin{aligned} x_{\pm} &= \frac{1}{2} \left(L_{\text{total}} + R_{\text{total}} + K_D \right. \\ &\quad \left. \pm \sqrt{(L_{\text{total}} + R_{\text{total}} + K_D)^2 - 4L_{\text{total}}R_{\text{total}}} \right) \end{aligned} \quad (95)$$

of which only x_- is physical and satisfies $x < L_{\text{total}}, R_{\text{total}}$.

The kinetics of relaxation

$$\frac{dx}{dt} = k_{\text{on}}(x - x_-)(x - x_+), \quad (96)$$

can also be integrated:

$$x(t) = \frac{x_- - ax_+ e^{-k_{\text{reaction}} t}}{1 - a e^{-k_{\text{reaction}} t}}, \quad (97)$$

where the characteristic rate for the reaction is: $k_{\text{reaction}} = k_{\text{on}} \sqrt{(L_{\text{total}} + R_{\text{total}} + K_D)^2 - 4L_{\text{total}}R_{\text{total}}}$, and a is a constant depending on the initial condition.

In the limit of a few receptors $R_{\text{total}} \ll K_D, L_{\text{total}}$, the expression simplifies to $x = L_{\text{total}}R_{\text{total}}/(L_{\text{total}} + K_D)$ and $k_{\text{reaction}} = k_{\text{on}}L_{\text{tot}} + k_{\text{off}}$. Three regimes can be considered:

- When the ligand exists in high concentration, $L \gg K_D$ and $x = R_{\text{total}}$, the system relaxes and fluctuates with a characteristic rate $k_{\text{reaction}} = k_{\text{on}}L_{\text{total}}$.

- When the ligand is sparse and in limited amount, $L_{\text{total}} \ll K_D$, few complexes form, $x \approx L_{\text{total}}R_{\text{total}}/K_D$, and the system relaxes and fluctuates with a characteristic rate k_{off} .
- When $[\text{Ligand}] \approx K_D$, half of the receptors are occupied and half are free of ligands, $x = R_{\text{total}}/2$; this is the mid-point of the dose response when one assess receptor occupancy against increasing concentrations of ligands.

B. Gene regulation

1. Basic model

In the simplest quantitative model of gene regulation, we can write the deterministic dynamics of gene transcription and translation as:

$$\begin{aligned} \frac{d}{dt}[\text{mRNA}] &= p_{\text{mRNA}} - \gamma_{\text{mRNA}}[\text{mRNA}] \\ \frac{d}{dt}[\text{protein}] &= p_{\text{protein}}[\text{mRNA}] - \gamma_{\text{protein}}[\text{protein}], \end{aligned} \quad (98)$$

where p_i are the production rates and γ_i the degradation rates of the molecules, for $i = \text{protein, mRNA}$. Since transcription and degradation rates for mRNA are often larger than translation and degradation rates for the protein, timescale separation leads to

$$\frac{d}{dt}[\text{protein}] = p_{\text{protein}} \frac{p_{\text{mRNA}}}{\gamma_{\text{mRNA}}} - \gamma_{\text{protein}}[\text{protein}]. \quad (99)$$

At steady state:

$$[\text{protein}] = \frac{p_{\text{mRNA}} p_{\text{protein}}}{\gamma_{\text{mRNA}} \gamma_{\text{protein}}}. \quad (100)$$

This straightforward expression can become arbitrarily complicated. Both the production and degradation rates of mRNA and proteins of interest can be complex functions of the signalling response, and post-translational modifications and regulated degradation can also happen.

In simplest case, these rates can be approximated as constants, when adiabatic conditions apply such that signalling responses and cytokine production and consumption occur on timescales much shorter or much longer than gene regulation. Including positive feedback loops produces more interesting dynamical behaviour, especially when the reinforcement is through multimerized transcription factors that introduce additional nonlinearities.

2. Auto-amplification with a single transcription factor

A simple feedback loop corresponds to the case when the rate of transcription of a protein species present in

the cell at concentration $[\text{protein}] = X$ is a dose response of the X itself (e.g. when X is a transcription factor that binds to the promoter region of the gene for X). Then the dynamic equation for X is:

$$\frac{d[X]}{dt} = k \frac{[X]}{[X] + K} - \gamma[X], \quad (101)$$

where k is the maximum production rate, K the concentration of protein X at half-maximum expression, and γ is a degradation rate (we have used the separation of time scales to eliminate the mRNA stage, as in the previous paragraph). This equation has one fixed point $X = 0$ for $k/\gamma < 0$, and $X_{\text{equilibrium}} = k/\gamma - K$, for $k/\gamma > K$. A simple stability analysis around that non trivial fixed point, taking $[X] = X_{\text{equilibrium}} + \epsilon$,

$$\frac{d\epsilon}{dt} = \frac{\gamma^2}{k} \left(K - \frac{k}{\gamma} \right) \epsilon, \quad (102)$$

shows negative restoring force ($K - k/\gamma < 0$), meaning that all fluctuations get quenched and the fixed point is stable.

3. Auto-amplification with multiple transcription factors

The complexity and relevance of the auto-amplification gene regulatory circuit becomes more relevant when the expression of gene X is regulated by two transcription factors TF_1 and TF_2 , i.e. both transcription factors need to bind to the promoter region of gene X to elicit its transcription. We can compute the state of the promoter using classical tools of statistical mechanics. The promoter for gene X can exist in four possible states with the related probabilities: unoccupied promoter (p_0), TF_1 only (p_1), TF_2 only (p_2), and both (p_{12}), with

$$p_0 = \frac{1}{Z}, \quad p_1 = \frac{[TF_1]}{K_1}, \quad p_2 = \frac{[TF_2]}{K_2}, \quad p_{12} = \frac{[TF_1][TF_2]}{K'_1 K'_2}, \quad (103)$$

with the partition function:

$$Z = 1 + \frac{[TF_1]}{K_1} + \frac{[TF_2]}{K_2} + \frac{[TF_1][TF_2]}{K'_1 K'_2}, \quad (104)$$

where $[TF_1]$ and $[TF_2]$ denote concentrations of the two transcription factors and $K_1 = c_0 \exp(\Delta G_1/RT)$, $K_2 = c_0 \exp(\Delta G_2/RT)$, $K'_1 = c_0 \exp(\Delta G'_1/RT)$ and $K'_2 = \exp(\Delta G'_2/RT)$ their equilibrium binding constants related to the free energies of binding individually (non-primed) and cooperatively (primed) to the promoter binding sites.

When the cooperative binding of transcription factors binding to the promoter region is usually thermodynamically favored, $K'_1 \ll K_1$ and $K'_2 \ll K_2$, the partition function simplifies to $Z \approx 1 + [TF_1][TF_2]/(K'_1 K'_2)$.

A common case is when $[TF_1] = [TF_2] = [TF]$ and gene X is transcribed upon homodimerization of a transcription factor in its promoter region (then $K'_i = K'$).




FIG. 16. **Bistable auto regulation of a gene.** Shown are the production and degradation rates of as a function of the protein concentration X , Eq. 106. The intersections of the curves define three fixed points: two stable ones (0 and X_+), and one unstable one (X_-).

The probabilities of the promoter to be unoccupied (transcribed at a basal level for auto-activation), $p_{\text{off}} = p_0$, and to be occupied (transcribed at an enhanced level for auto-activation), $p_{\text{on}} = p_{12}$ are:

$$p_{\text{on}} = \frac{[TF]^2}{[TF]^2 + K^2}, \quad p_{\text{off}} = 1 - p_{\text{on}} \quad (105)$$

This form of regulation, which follows a Hill function (with Hill coefficient $h = 2$), introduces a nonlinearity in the production of X that is of critical relevance in cell differentiation. For auto-activation by homodimers, the dynamic equation for the regulation of X is:

$$\frac{d[X]}{dt} = k \frac{[X]^2}{[X]^2 + K^2} - \gamma[X], \quad (106)$$

where k and γ are production and degradation rates for X respectively.

As in the non cooperative case, this dynamical system has a single fixed point if $k/\gamma < K$: $[X] = 0$. But for $k/\gamma > K$, it has three fixed points at 0 , X_- and X_+ . The stability of these solutions can be deduced graphically (see section V A for details): X_- is an unstable fixed point and 0 and X_+ are stable fixed points. Hence, this simple system of gene regulation will generate two types of cells: cells that do not express X , and cells that express a high level of X . We see many examples of such bimodal distributions in expression of transcription factors, cytokines and surface markers in cells of the immune system (see section V A).

C. Population dynamics, genetics

To study the somatic evolution of immune clones and cell types it is useful to summarize some basic results

from population genetics. The material in this section can be found in many evolution and population textbooks [295, 296]. To describe the evolution of a mutating population, one typically considers a background of genetically identical individuals (or cells, in the context of lymphocyte population dynamics). Each one of these individuals can acquire a mutation with rate μ , which starts a new mutant subpopulation. Each individual in both the ancestral and the mutant subpopulations can die or reproduce in each generation, and mutants carrying non-neutral amino acid substitutions have a different (higher for beneficial and lower for deleterious mutations) growth rates. Since it is more likely that a mutation will confer a disadvantage than an advantage, most mutants are deleterious. But as we shall see, even beneficial mutations are not necessarily destined to succeed and ultimately the fate of most mutants is to die. Out of the lucky few that establish subpopulations of significant frequency, some will completely take over the population — they will *fix*. In this section we calculate the probabilities for fates of mutant subpopulations in the simplified case, compared to somatic evolution of immune repertoires, when there are no sources of new clones other than mutations.

1. Deterministic mutation-selection balance

Let us start with a situation where two subpopulations exist, the ancestral clone of size n_2 individuals and a mutant clone of size n_1 individuals. The subpopulations grow (which accounts for both reproduction and death) with rates γ_1 and γ_2 , respectively, where $g_i = r_i - f(n_1, n_2)$ is the balance of growth rate r_i and a death rate $f(n_1, n_2)$. Individuals also mutate from the ancestral population to the mutant one with rate μ_1 and mutate back with rate μ_2 :

$$\frac{dn_1}{dt} = \mu_2 n_2 - \mu_1 n_1 + g_1(n_1, n_2) n_1 \quad (107)$$

$$\frac{dn_2}{dt} = \mu_1 n_1 - \mu_2 n_2 + g_2(n_1, n_2) n_2. \quad (108)$$

We also assume, as is often the case in population genetics models that the population size is constant, $N = n_1 + n_2 = \text{const}$, which sets $f(n_1, n_2) = (r_1 n_1 + r_2 n_2)/N$. The constant population size constraint means that knowing n the size of both subpopulations is completely determined by the fraction of individuals in the ancestral population $x = n_1/n$, which follows:

$$\frac{dx}{dt} = \mu_2 - (\mu_1 + \mu_2)x + sx(1-x), \quad (109)$$

where the selection coefficient $s = r_1 - r_2$ describes how much faster (or slower if $s < 0$) the ancestral population grows compared to the mutant subpopulation.

At steady state, in the absence of mutations $\mu_1 = \mu_2 = 0$, either the ancestral subpopulation fixes, $x^* = 1$ for $s > 0$, or the mutant one $x^* = 0$ for $s < 0$. In the

absence of selection but with mutations, the relative ratio of the mutation rates determines the fraction of each subpopulation, $x^* = \mu_1/(\mu_1 + \mu_2)$. With both mutation and selection a mutation-selection balance is established (for $s < 0$) at $x^* = (1/2)(1 + 2\mu/|s| - \sqrt{4(\mu/s)^2 + 1})$ for $\mu_1 = \mu_2 = \mu$. When mutations are rare compared to the fitness advantage, $|s| \gg \mu$, this balance simplifies to $x^* \approx \mu/|s|$.

2. Genetic drift

In order to observe a situation described by the deterministic equations in Eq. 109, the mutant population needs to grow to a sizeable fraction of the population. However, every mutant appears first in only one individual and undergoes a subsequent random walk of reproduction and death, which means that the most likely fate is for it go immediately extinct. This makes the effect of small number noise coming from finite population sizes, called genetic drift in population genetics, relevant. We can explore this effect considering a population that produces only neutral mutants (meaning all mutants grow at the same rate as the ancestor), and keeping the population size fixed to N individuals. If we focus on one individual at some initial time, and follow its offspring, at very long times only two outcome are possible. Either its offspring have taken over the whole population, or the lineage has gone completely extinct. Since we start with N individuals, the probability of taking over is $1/N$. The argument generalizes to a subpopulation of size n : the probability that one of its members has taken over the population is n/N .

3. Wright-Fisher model

While we cannot calculate deterministically the fate of any particular mutant individual, we can calculate the probability of the evolution of the ancestral and mutant fractions, assuming two subpopulations and a constant population size as we did in subsection XII C 1. Two models, the Wright-Fisher and the Moran models, each with slightly different setups, describe the neutral evolution of populations. The Wright-Fisher model assumes discrete, non-overlapping generations: at each generation, individuals from the previous generation are cleared, a new sample of N individuals is drawn, each new individual picking an ancestor from the previous generation with probability $1/N$. Following a subpopulation of size n in the parent generation, the probability that their offspring comprise m individuals is given by the transition probability:

$$P_{n \rightarrow m} = \binom{N}{m} \left(\frac{n}{N}\right)^m \left(1 - \frac{n}{N}\right)^{N-m}. \quad (110)$$

The probability distribution for the size of that subpopulation evolves according to the recursion: $P_{t+1}(m) =$

$\sum_n P_{n \rightarrow m} P_t(n)$, where t labels generations.

Within this model the mean frequency of the subpopulation does not change with the time $\sum_n n P_n(t) = \langle n(t) \rangle = \langle n(t+1) \rangle$. Since the binomial distribution for large N becomes Gaussian, the evolutionary trajectory of the number of individuals in each subpopulation size is well described by a random walk: $n(t+1) \approx n(t) + \eta(t)\sqrt{n(t)(1-n(t)/N)}$, with $\eta(t)$ normally distributed.

The Wright-Fisher model can be generalized to include selection. The probability to pick a given member i of the parent generation depends on its selective advantage or disadvantage s_i , as $(1/N)(1 + s_i)/(1 + \bar{s})$, with $\bar{s} = (1/N)\sum_i s_i$.

In the case of two subpopulations, wildtype (of size $n_2 = N - n$) and mutant (of size $n_1 = n$), with selective (dis-)advantage s for the mutant, Eq. is modified to: $P_{n \rightarrow m} = \binom{N}{m} p^m (1-p)^{N-m}$, with $p = (n/N)(1+s)/(1+\bar{s})$ and $\bar{s} = 1 + sn/N$. For large populations, this leads to a biased random walk:

$$n(t+1) \approx n(t) + sn(t) \left(1 - \frac{n(t)}{N}\right) + \eta(t) \sqrt{n(t) \left(1 - \frac{n(t)}{N}\right)}, \quad (111)$$

whose deterministic part reproduces Eq. 109 in the absence of mutations.

4. Probability of extinction

We can calculate the probability of extinction of the mutant subpopulation in the setup of two populations we have described until now. This limit corresponds to a small mutation rate: each subpopulation can fix or go extinct before a new mutation appears. While many immunological situations (e.g. affinity maturation) may not be in that regime, it is an important result to know. The probability that a subpopulation of size n goes extinct, $q(n)$, requires knowing the probability that all their offspring, calculated within the Wright-Fisher model, will go extinct:

$$q(n) = \sum_{m=0}^N \binom{N}{m} p^m (1-p)^{N-m} q(i), \quad (112)$$

given a selective advantage s for the population. For large N the binomial distribution is peaked, and we can Taylor expand $q(i)$ around n , $q(i) \approx q(n) + (i-n)q'(n) + \frac{1}{2}(i-n)^2 q''(n)$, which results in:

$$0 = sn \left(1 - \frac{n}{N}\right) q'(n) + \frac{1}{2} n \left(1 - \frac{n}{N}\right) q''(n). \quad (113)$$

Solving for the extinction probability with boundary conditions $q(n=0) = 1$ (an extinct subpopulation remains extinct) and $q(n=N) = 0$ (a fixed population cannot go extinct) yields $q(n) = 1 - \frac{1-e^{-sn}}{1-e^{-sN}}$. The probability of

fixation is then:

$$P_{\text{fix}} = 1 - q(n) = \frac{1 - e^{-sn}}{1 - e^{-sN}}. \quad (114)$$

If $s = 0$ we recover the result in subsection XII C 2.

In the limit of strong positive selection ($Ns \gg 1$, but $s \ll 1$) the probability that a founder mutant ($n = 1$) fixes is proportional to the selection strength $P_{\text{fix}} \approx s$, if the initial population is one individual, or more generally $P_{\text{fix}} \approx ks$, if the initial population is small $k \ll 1/s$ individuals. On the other hand $P_{\text{fix}} \approx 1$ for large founder populations $k \gg 1/s$. This defines a threshold of $1/s$ individuals to ensure the survival of the subpopulation. For weak selection pressures $|Ns| \ll 1$, the probability of fixation is equal to the frequency of individuals in the population and is independent of s , $P_{\text{fix}} \approx n/N$, making this regime effectively neutral. Lastly, for strongly deleterious mutations, $sN \ll -1$, the probability of fixation is non-zero $P_{\text{fix}} \approx e^{s(N-n)}$, but is exponentially small.

5. Moran model, continuous limit, and time varying selection

The above results can also be obtained considering a diffusion process in the number of individuals in the subpopulation. To do this, we use the Moran model, in which at each time step we choose one individual to die and one to reproduce (it can be the same individual), ensuring the population size is kept fixed. This implies overlapping generations, and the typical generation time is of order N time steps.

Assume a wildtype population of size $N - n$ and a mutant population of size n . The probability of having n at time step i is given by:

$$p_n(i) = p_{n+1}(i-1) \frac{n+1}{N} \frac{N-(n+1)}{N} + p_{n-1}(i-1) \frac{n-1}{N} \frac{N-(n-1)}{N} + p_n(i-1) \left[\left(\frac{n}{N} \right)^2 + \left(\frac{N-n}{N} \right)^2 \right], \quad (115)$$

where the first term describes killing one dominant allele and reproducing a mutant, the second one reproducing a dominant allele and killing a mutant, and the last term describes the two possibilities of killing and reproducing the same allele. In the limit of large population sizes, rescaling time by the typical generation N , $t = i/N$, so that $p_n(i) - p_n(i-1) = (1/N)\partial_t p_n(t)$, and Taylor expanding in $x = n/N$:

$$\frac{\partial p(x, t)}{\partial t} = \frac{1}{2N} \frac{\partial^2}{\partial x^2} [x(1-x)p(x, t)]. \quad (116)$$

Mutation and selection can be added following a similar

derivation and contribute in a mean “drift”:

$$\begin{aligned} \frac{\partial p(x, t)}{\partial t} &= \frac{1}{2N} \frac{\partial^2}{\partial x^2} [x(1-x)p(x, t)] \\ &- \frac{\partial}{\partial x} [(sx(1-x) + \mu_2 - (\mu_1 + \mu_2)x)p(x, t)]. \end{aligned} \quad (117)$$

This Fokker-Planck equation is the stochastic version of 109, with effective diffusion coefficient $x(1-x)/2N$. This expression is also consistent with the random walk approximation of the Wright-Fisher model (111), meaning that the two models are equivalent in that continuous limit.

Eq. 117 has a general solution in terms of Gegenbauer polynomials. At steady state, assuming $\mu_1 = \mu_2 = \mu$ for simplicity it takes the form $p(x) = Z^{-1}[x(1-x)]^{N\mu-1}e^{Nsx}$, where Z is a normalization factor. We can reinterpret it as a Boltzmann distribution, $p(x) = Z^{-1}p_0(x)e^{Nsx}$, where each mutant individual gives an energy gain $-s$, and $p_0(x) = [x(1-x)]^{N\mu-1}$ is the neutral distribution in the absence of selection.

Since Eq. 117 is a one dimensional diffusion equation, the probability of fixation is calculated from the backward equation for reaching the absorbing barrier at $x = 1$. This calculation gives the same result as Eq. 114.

To model the fluctuations of the environment, we can consider a time varying selection pressure $s(t) = s_0 + \sigma(t)$ with mean s_0 and random white noise fluctuations $\langle \sigma(t)\sigma(t') \rangle = \delta(t-t')\Omega$ around this mean [297]. The probability $p(y, x, t)$ that the wildtype, starting at initial frequency y at time $t = 0$, reaches frequency at least x by time t is given by the backwards equation:

$$\partial_t p(y, x, t) = v(y)\partial_y p(y, x, t) + D(y)\partial_y^2 p(y, x, t), \quad (118)$$

where

$$v(y) = s_0 y(1-y) + \Omega y(1-y)(1-2y), \quad (119)$$

and where the diffusion term cumulates genetic drift and fitness fluctuations:

$$D(y) = \frac{1}{2N} y(1-y) + \Omega y^2(1-y^2). \quad (120)$$

With the boundary conditions $p(y \geq x, x, t) = 1$ and $p(y = 0, x, t) = 0$, the probability of ever reaching frequency x is:

$$p(y, x, t \rightarrow \infty) = \frac{1 - \left| \frac{1-y/\alpha_+}{1-y/\alpha_-} \right|^\lambda}{1 - \left| \frac{1-x/\alpha_+}{1-x/\alpha_-} \right|^\lambda}, \quad (121)$$

where $\alpha_\pm = \frac{1}{2} \left[1 \pm \sqrt{1 + 2/N\Omega} \right]$ and $\lambda = s_0/(\Omega\sqrt{1 + 2/N\Omega})$. For $x = 1$ we recover the fixation probability reported in Takahata et al [297], and for $\Omega \rightarrow 0$ we recover the result for constant fitness (114), $P_{\text{fix}} = (1 - e^{-s_0 N y})/(1 - e^{-s_0 N})$.

For more than two genotypes [298], Eq. 117 generalizes to an equation for a vector $\vec{x} = (x^1, \dots, x^{K-1})$ of linearly independent genotype frequencies ($x^K = 1 - \sum_{\alpha}^{K-1} x^{\alpha}$):

$$\frac{\partial p(\vec{x}, t)}{\partial t} = \sum_{\alpha, \beta} \left\{ -\frac{\partial}{\partial x^{\alpha}} [m^{\alpha}(\vec{x}) + C^{\alpha\beta}(\vec{x})s_{\beta}(\vec{x})] + \frac{1}{2N} \frac{\partial^2}{\partial x^{\alpha} \partial x^{\beta}} C^{\alpha\beta}(\vec{x}) \right\} p(\vec{x}, t), \quad (122)$$

where the covariance matrices are:

$$C^{\alpha\beta}(\vec{x}) = \begin{cases} -x^{\alpha}x^{\beta} & \text{if } \alpha \neq \beta \\ x^{\alpha}(1 - x^{\alpha}) & \text{if } \alpha = \beta \end{cases} \quad (123)$$

the mutation coefficients $m^{\alpha}(\vec{x}) = \sum_{\beta} (\mu_{\beta \rightarrow \alpha} x^{\beta} - \mu_{\alpha \rightarrow \beta} x^{\alpha})$ and the selection coefficient is the relative growth rate of genotype β compared to a reference genotype $s_{\alpha} = f^{\beta} - f^{\text{ref}}$.

6. Branching processes

Branching processes are useful for tracking the fate of the offspring of a individual through time. We will first introduce it in the context of the Moran model, and then present its more standard applications. Assuming birth (with rate $1 + s$) and death (with rate 1) as the only possibly processes, we can track the evolution of the probability of having 1 individual at time 0 and n at time t . This probability satisfies a recursion that can be obtained by considering the possible events occurring in the first time step, between times 0 and $dt = 1/N$ [287, 299]:

$$\begin{aligned} p(1, n, t) &= \frac{1}{N} \left(1 - \frac{1}{N}\right) \delta_{n,0} \\ &+ \frac{1}{N} \left(1 - \frac{1}{N}\right) (1 + s)p(2, n, t - dt) + \\ &+ \left[1 - \frac{1}{N} \left(1 - \frac{1}{N}\right) (2 + s)\right] p(1, n, t - dt), \end{aligned} \quad (124)$$

The first term corresponds to the lineage going extinct between times 0 and dt — and remaining extinct until t . The second term corresponds to a division of the initial individual between times $t = 0$ and dt , and the lineage then reaching size n from size 2 in the remaining time, $t - dt$. The last term corresponds to no change at all.

The next step, which characterizes the branching process approach, is to assume that the outcome of a lineage can be deduced from the outcome of each of the descendants of the first division, taken independently: $p(2, n, t) \approx \sum_{m=0}^n p(1, m, t)p(1, n - m, t)$, which assumes that once one individual gives birth to two, the births in these lineages happen independently. This approximation is valid as long as $n \ll N$. Then the recursion becomes in the continuous time limit:

$$\frac{\partial p(1, n)}{\partial t} = \delta_{n,0} + (1 + s) \sum_{m=0}^n p(1, m)p(1, n - m) - (2 + s)p(1, n), \quad (125)$$

Defining the generation function $G(z, t) = \sum_{n=0}^{\infty} p(1, n, t)z^n$ we obtain:

$$\frac{\partial G(z)}{\partial t} = 1 + (1 + s)G^2(z) - (2 + s)G(z), \quad (126)$$

which is solved with boundary conditions $G(z = 1) = 1$ and $G(z, t_0) = z$:

$$G(z, t) = \frac{(z - 1)(1 - e^{st}) + zs}{(z - 1)(1 - (1 + s)e^{st}) + zs}. \quad (127)$$

For $s = 0$ (no selection), we obtain $G(z) = 1 - \frac{1}{1 + (1 - z)^{-1}}$, whose series expansion yields $p(1, n = 0, t) = t/(1 + t)$ and $p(1, n, t) = t^{n-1}/(1 + t)^{n+1}$ for $n \geq 1$.

For arbitrary s the mean number of individuals at time t is $\langle n \rangle = \partial_z G(z)|_{z=1} = e^{st}$. The probability of going extinct at time t is

$$G(z = 0) = p(1, n = 0, t) = \frac{e^{st} - 1}{(1 + s)e^{st} - 1}, \quad (128)$$

which goes to $1 - s$ for $s > 0$ at $t \rightarrow \infty$. We thus recover the result that the fixation probability goes to s , as obtained previously in the strong positive selection limit. In that limit, by the time the mutant escapes genetic drift, for $n \gg 1/s$, its population size is still small compared to the total population size $n \ll N$, making the branching process approximation appropriate.

While we have illustrated branching processes in the context of a Moran model of evolution, branching processes are ubiquitous and used in a variety of contexts. One of the simplest branching process is defined in discrete time, where at each step a individual can divide with probability p , or not divide with probability $1 - p$ [121, 299, 300]. A recursion relation for the generating function of the total number of individuals n in a lineage at time t (past and present), $G(z, t) = \sum_n p(n, t)z^n$, can be written following similar arguments as above:

$$G(z, t) = zpG^2(z, t - 1) + (1 - p)z. \quad (129)$$

This recursion equation is solved for $t \rightarrow \infty$ by:

$$G(z, t \rightarrow \infty) = \frac{1 - \sqrt{1 - z^2 4p(1 - p)}}{2pz}, \quad (130)$$

which has a critical point as $p \rightarrow 1/2$ [300], when the average number of offspring equals 1. Rewriting the generating function as a series we recover the probability of the total number of individuals (which must be even)

$$P_{2k} = \frac{1}{2} \frac{\Gamma(k - 1/2)}{\Gamma(1/2)\Gamma(k + 1)} \frac{(4p(1 - p))^k}{2p}, \quad (131)$$

For $p < 1/2$, P_k is a decaying exponential as expected, and for $p = 1/2$ it decays as a power-law $P_k \sim k^{-3/2}$. Note that this class of critical branching processes has been used to explain power laws in the distribution of activity in neural networks [301].

7. Coalescence process

As seen when calculating fixation probabilities, it is often useful to think about evolutionary processes backwards in time. This is the basic idea behind a coalescence approach, which can be formalised using branching processes in the context of the Wright-Fisher model. Here we give just some basic intuition about how thinking back in time about the history of coalescing sequences in lineages can be useful when studying affinity maturation processes, or tracing phenotypic lineages. We will present the coalescence process in a neutral evolutionary framework. As before the neutral framework provides us with a null model in the case affinity maturation where selection is important, but it may also be useful in immunological phenotyping. The coalescence approach does not concern itself with mutations, but simply tracks genealogies. Mutations can later be added to an existing genealogy (tree).

We consider two individuals and ask how long ago they shared a common ancestor. If one individual has a given parent, the probability that the second cell has the same parent, given there are N cells, is $1/N$. The probability that they do not have the same parent is $1 - 1/N$. Following this reasoning, the probability that they have the same parent t generations ago, but not during the $t - 1$ generations is

$$P[T_2 = t] = [1 - 1/N]^{t-1} 1/N \approx 1/N e^{-t/N}, \quad (132)$$

where we have expanded for large N and T_2 stands for time to mean recent common ancestor (MRCA). The mean time for two cells to coalesce is simply the mean expectation time of this distribution

$$\langle T_2 \rangle = \int_0^\infty dt t / N e^{-t/N} = N. \quad (133)$$

The mean time for two cells to coalesce is equal to the population size, in units of generation time. More generally, the probability that k cells do *not* share the same parent is $P_{\text{diff}} = [1 - 1/N][1 - 2/N] \dots [1 - (k - 1)/N] \approx 1 - \binom{k}{2} N^{-1}$. The probability that *at least* two cells have the same parent (or coalescence in our backwards picture) is then $P_c = 1 - P_{\text{diff}} = \binom{k}{2} N^{-1}$. If $N \gg k$ then the probability that more than two cells share the same parent in a single generation can be neglected and we will assume that in each generation only two cells will share a parent. The distribution of times until the first coalescence is:

$$P[\text{1st coalescence at time } t] = [1 - P_c]^{t-1} P_c \quad (134) \\ \approx P_c e^{-P_c t} = \binom{k}{2} \frac{1}{N} e^{-\binom{k}{2} t/N}.$$

After a coalescent event, there are $k - 1$ individuals left, and the process can be repeated until the whole genealogical tree is reconstructed. The coalescence probability completely determines the statistics of the topology of

the tree and the branch lengths in a neutral process. Once the branch lengths (times between each coalescence events) have been determined, neutral mutations are distributed randomly along the branch with some rate, so that their number on each branch follows a Poisson distribution. Specifically, the number of mutations π between two individuals, also called pairwise heterozygosity, is given by:

$$P(\pi|t) = e^{-2\mu t} \frac{[2\mu t]^\pi}{\pi!} \quad (135)$$

for coalescence time t and per-generation mutation rate μ (the total time is $2t$ because it adds the two branches of length t from the common ancestor). Of course, all of this breaks down in the presence of selection.

Thus, the coalescence probability determines the genetic diversity within a population. The distribution of π in the population can be obtained by integrating over the coalescence time (132),

$$P(\pi) = \int_0^\infty dt \left[\frac{1}{N} e^{-t/N} \right] \left[e^{-2\mu t} \frac{[2\mu t]^\pi}{\pi!} \right] = \frac{1}{1 + \theta} \left[\frac{\theta}{1 + \theta} \right]^\pi, \quad (136)$$

where $\theta = 2\mu N = \langle \pi \rangle$. Thus, in principle, in the regime of neutrality, characterizing the distribution of the mean mutational distance should allow us to read off the mutation rate.

Departures from the Poisson distribution is one of the signatures of selection. In the presence of selection, the distribution of branch lengths depends on the details of the type of selection we are studying. If we do think about affinity maturation (where selection plays an important role), we see that cells undergo bursts of selection in the germinal centers, followed by periods outside of the germinal centers. Two cells sampled at the same time from the blood may therefore have very different recent histories and using their mutation distance to infer a mean mutation rate would be misleading.

8. Site frequency spectra and tree balancing

Another statistics that is easy to calculate within the neutral model, and can thus be used as a null model to compare to data, is the site frequency spectrum (SFS). The SFS is the number of individuals in the population (in our case cells) that have a mutation at a given position in the aligned sequence with respect to the dominant base pair in the most recent common ancestor, and presents it as a histogram: $(n_1, n_2, \dots, n_{N-1})$, where n_1 is the number of mutations present in a single cell, n_2 is the number number of mutations present in two cells, etc. Sites that are not polymorphic (all have the same base pair) do not contribute to the spectrum as they coincide with the most recent common ancestor of the whole population. However a mutation at a site that is close to fixation will contribute to the spectrum, although this mutation now dominates the population. In absence of

information about the most recent common ancestor, in population genetics the mutation is often called with respect to the outgroup. In the case of B-cell receptors, a good estimate can be inferred from the best alignment to the V, D, and J germline sequences.

We can estimate the SFS for a neutrally evolving population of constant size N within the Moran model. In the continuous time limit, the mean number of mutations shared by k individuals evolves according to (for $k > 1$):

$$\frac{dn_k}{dt} = \frac{(k+1)(N-k-1)}{N}n_{k+1} + \frac{(k-1)(N-k+1)}{N}n_{k-1} - 2\frac{k(N-k)}{N}n_k = J_{k-1} - J_k, \quad (137)$$

with

$$J_k = [k(N-k)/N]n_k - [(k+1)(N-k-1)/N]n_{k+1} \quad (138)$$

the “current” of mutations across subpopulation sizes. The first term corresponds to death event occurring in the subpopulation of size $k+1$ carrying the mutation of interest, and the second term to birth events in the subpopulation of size $k-1$. New mutation always starts with one individual, so that

$$\frac{dn_1}{dt} = \frac{2(N-2)}{N}n_2 - 2\frac{(N-1)}{N}n_1 + \mu, \quad (139)$$

where μ is the mutation rate. At steady state the current is constant and equal to $J_k = J_1 = \mu/N$, because each new mutation has a probability $1/N$ to fix, resulting in a current μ/N of mutations traveling from from size 1 to N . This implies:

$$n_k = \frac{(k-1)(N-k+1)}{k(N-k)}n_{k-1} - \frac{\mu}{k(N-k)}, \quad (140)$$

which is solved by

$$n_k = \mu/k. \quad (141)$$

The SFS can also be calculated approximately in models with selection, but still with fixed population sizes, so we do not recall these results here [302].

Another feature that can be used to identify selection through departure from the neutral model is how balanced lineage trees are. Intuitively, a neutral process does not favor adding a new mutation to any of the branches, hence the resulting trees should be symmetric and balanced. A process with selection will preferentially grow the favoured parts of the tree, resulting in some long branches, and other short ones. In the neutral model, the number of leaves n in a sublineage at generation t , follows the distribution $p_t(n)$, which satisfies the recursion:

$$p_t(n) = \left(1 - \frac{n}{t-1}\right)p_{t-1}(n) + \frac{n-1}{t-1}p_{t-1}(n-1). \quad (142)$$

The sublineage of interest can add one leaf by reproducing, with probability proportional to its size n , (second term), or not (first term). Solving by recursion gives a uniform branch length distribution $p_t(n) = 1/(n-1)$.

9. Clonal interference

Another concept from population genetics that has been shown to be relevant in BCR affinity maturation is clonal interference, which we discussed in more detail in section VIII. Here we recall the back-of-the-envelope arguments of Desai and Fisher [287] to show on a more classical example why clonal interference slows down the rate of adaptation.

Given a population size N and mutation rate μ , a mutation occurs in any individual with rate $N\mu$, so the time between mutations is $(N\mu)^{-1}$. Most of these mutations go extinct due to genetic drift. In the strong selection regime, a mutation survives genetic drift with probability s . Thus, when mutations are rare, the average time for a new mutation to occur and fix is $(N\mu s)^{-1}$, and the rate of adaptation is given by $v_{\text{rare}} = N\mu s^2$.

However fixation itself may take time, and the above picture breaks down when mutations are no longer rare, i.e. when the typical time between successful mutation becomes smaller or of the same magnitude as the time for it to fix. After the subpopulation carrying the mutation overcomes genetic drift which results in an initial population size of $1/s$, we can consider that the subpopulation grows deterministically with rate s , $n(t) \sim (1/s)e^{st}$. The mutation will fix when $n(t) = N$, which gives the fixation time $t_{\text{fix}} = (1/s) \ln Ns$. Mutations are no longer rare when $(N\mu s)^{-1} \sim 1/s \ln Ns$, or $N\mu \sim (\ln Ns)^{-1}$.

In that regime, many mutants can co-exist at the same time, and new mutations can appear in existing mutants before these have time to fix. Mutants existing at the same time will have different growth rates, depending on the number of mutations they have acquired, even if we assume for simplicity that all mutations give the same fitness advantage. The whole population can be described by a fitness distribution and we can notice that the fate of a mutation in the bulk of this distribution (a subpopulation that was created some time ago) is different than at its nose (a recently created subpopulation). The nose subpopulations are small and susceptible to genetic drift. The bulk subpopulations have more individuals and grow deterministically, but are subject to nonlinear effects of competition between individuals. We can consider these two subpopulations separately and then “stitch” the two solutions.

The individuals at the nose of the fitness distribution have a fitness advantage, let's call it qs , meaning that they have q mutations, each conferring an advantage s , compared to the bulk of the distribution. The time needed for an even fitter individual (with fitness $(q+1)s$) to appear is given by τ that satisfies $\int_0^\tau dt \mu n_{\text{nose}}(t) qs = 1$, where $n_{\text{nose}}(t) = 1/(qs)e^{qst}$ is the number of individuals at the nose of the fitness distribution, following the same arguments as for the rare mutation case. In the limit of $e^{qst} \gg 1$, one gets:

$$\tau = 1/(qs) \ln(qs/\mu). \quad (143)$$

Every time that a new fitter class is added to the distri-

bution, the mean of the distribution also increases. After $q\tau$, the old nose of the distribution will become the mean (note that the individuals that are less fit than the mean go extinct deterministically). The populations in fitness class q grow as qs , so in class $q-1$ they grow as $(q-1)s$, etc. A lineage originating at the nose grows to dominate the population $q\tau$ later, while progressively losing its fitness advantage. By the time it dominates the population, its relative fitness advantage is zero. On average, it will have grown at rate $qs/2$. This gives the consistency equation $(1/qs)e^{qst/2} = N$, with $t = q\tau$. This gives a second estimate of the establishment time:

$$\tau = \frac{2}{q^2 s} \ln Nqs. \quad (144)$$

Equating Eq. 143 and Eq. 144, we find the average rate at which the population grows:

$$q = \frac{2 \ln Nqs}{\ln(Nq/\mu)}. \quad (145)$$

We can solve this implicit equation approximately assuming that q is not too big and neglecting the $\ln q$ terms:

$$q = \frac{2 \ln Ns}{\ln(N/\mu)}. \quad (146)$$

The rate of adaptation in the regime of multiple competing mutations is then:

$$v_{\text{CI}} = \frac{s}{\tau} = \frac{s^2 \ln Ns}{\ln s/\mu}. \quad (147)$$

This rate scales with the logarithm of the population size, $v_{\text{CI}} \sim \ln N$, much slower than in the rare mutation regime $v_{\text{rare}} \sim N$. The clonal interference regime is not only a regime in which there are many competing mutations, but a regime where new mutations arise on the background of still relatively low frequency mutations, forming competing lineages. The distinction between competition of different clones and clonal interference is especially important in affinity maturation, where the competition in germinal centers could be between very different clones (which is not clonal interference), or between clones with similar histories, in which case it is clonal interference.

10. Quantitative traits

While models of population genetics are often defined in the space of genotypes or alleles, what is often measured in the resulting phenotype, which is a (possibly nonlinear) function of the genotype. For instance, for lymphocyte receptors, the relevant phenotype may be defined as the binding affinity to epitopes of interest.

A projection of a description in phenotypic space, where mutations occur independently at all loci i along the genome, results in an effective description in phenotypic space, which is useful because selection occurs at

the phenotypic level and we often do not have access to the genotype-phenotype map. Given a fraction x^α of the population carrying allele α , the population mean and variance of a phenotype E is:

$$\Gamma = \langle E \rangle = \frac{1}{N} \sum_{\alpha=1}^N E^\alpha x^\alpha \approx \int dE E w(E) \quad (148)$$

$$\Delta = \langle E^2 \rangle - \langle E \rangle^2 \approx \int dE (E - \Gamma)^2 w(E),$$

where N is the population size, and $w(E)$ is the distribution of this trait in the population.

One can derive an effective equation for the joint evolution of the mean phenotype and its variance [298]:

$$\begin{aligned} \frac{\partial}{\partial t} Q(\Gamma, \Delta, t) = & \left[-\frac{\partial}{\partial \Gamma} \left(\frac{d\Gamma^{\text{mut}}}{dt} + \frac{d\Gamma^{\text{sel}}}{dt} \right) \right. \\ & - \frac{\partial}{\partial \Delta} \left(\frac{d\Delta^{\text{mut}}}{dt} + \frac{d\Delta^{\text{sel}}}{dt} \right) \\ & \left. + \frac{1}{2N} \left(\frac{\partial^2}{\partial \Gamma^2} C^{\Gamma\Gamma} + \frac{\partial^2}{\partial \Delta^2} C^{\Delta\Delta} \right) \right] Q(\Gamma, \Delta, t), \end{aligned} \quad (149)$$

where we calculate the drift and diffusion terms below.

First lets focus on the diffusion terms, which correspond to genetic drift are independent of selection:

$$\begin{aligned} C^{\Gamma\Gamma} &= \sum_{\alpha, \beta} \frac{\partial \Gamma}{\partial x^\alpha} \frac{\partial \Gamma}{\partial x^\beta} C^{\alpha\beta} \\ &= \sum_{\alpha} E_\alpha^2 x^\alpha (1 - x^\alpha) - \sum_{\alpha \neq \beta} E_\alpha E_\beta x^\alpha x^\beta \\ &= \langle E^2 \rangle - \langle E \rangle^2 = \Delta, \end{aligned} \quad (150)$$

where we have used the definitions of the covariance matrices in high dimensional space from Eq. 123. Similarly

$$C^{\Delta\Delta} = \sum_{\alpha, \beta} \frac{\partial \Delta}{\partial x^\alpha} \frac{\partial \Delta}{\partial x^\beta} C^{\alpha\beta} = \langle (E - \Gamma)^4 \rangle - \Delta^2 \approx 2\Delta^2, \quad (151)$$

where in the last step we have assumed $w(E)$ is approximately Gaussian. A similar calculation shows that $C^{\Delta\Gamma} = \langle (E - \Gamma)^3 \rangle \approx 0$ with the Gaussian assumption.

To compute the drift terms, which arise from selective effects, we need to specify how fitness and phenotype are related. For simplicity, let us assume that fitness of each individual is a quadratic function of the phenotype (or expand fitness close to its peak at E^*), $f(E) = -c_0(E - E^*)^2$, where c_0 is a prefactor that measures the width of the fitness peak. The mean fitness is then:

$$\begin{aligned} F(\Gamma, \Delta) &= \sum_{\alpha} f^\alpha x^\alpha = -c_0 \sum_{\alpha} (E^\alpha - E^*)^2 x^\alpha \\ &= -c_0 [\Delta + (\Gamma - E^*)^2]. \end{aligned} \quad (152)$$

The change in the mean phenotype Γ due to selection is:

$$\begin{aligned} \frac{d\Gamma^{\text{sel}}}{dt} &= \frac{d}{dt} \int dE E w(E) = \int dE E \left(\frac{dw(E)}{dt} \right) \quad (153) \\ &= \int dE E (f(E) - F) w(E) \\ &= -c_0 ((E - E^*)^2 - \Delta - (\Gamma - E^*)^2) w(E) \\ &= -2c_0 \Delta (\Gamma - E^*). \end{aligned}$$

Since (by analogy with Eq. 122) $\frac{d\Gamma^{\text{sel}}}{dt} = C^{\Gamma} s_{\Gamma}$, using Eq. 150 we can verify that the selection coefficient is $s_{\Gamma} = \frac{\partial}{\partial \Gamma} F(\Gamma, \Delta) = -2c_0 (\Gamma - E^*)$.

In principle we can calculate the change in the mean variance due to selection in the same way $\frac{d\Delta^{\text{sel}}}{dt} = \frac{d}{dt} \int dE (E - \Gamma)^2 w(E)$ assuming a peaked phenotype distribution, but in practice its easier to use the analogy with Eq. 122 and Eq. 151:

$$\frac{d\Delta^{\text{sel}}}{dt} = \frac{\partial}{\partial \Gamma} F(\Gamma, \Delta) \approx -2c_0 \Delta^2. \quad (154)$$

To consider the changes of Γ and Δ due to mutations we have to write down a more detailed genotype model since mutations act on base pairs. If each locus can take a WT value ($\sigma_i = 0$) or mutant value ($\sigma_i = 1$), each linearly additive phenotype (for simplicity) E (for example binding energy) can be written as $E^\alpha = \sum_i^L E_i^\alpha \sigma_i$, where E_i^α is a given sites contribution to the overall phenotype. We have $\Gamma = (1/N) \sum_{i,\alpha} E_i^\alpha \sigma_i^\alpha x^\alpha$ and:

$$\begin{aligned} \frac{d\Gamma^{\text{mut}}}{dt} &= \frac{1}{N} \sum_{i,\alpha} \frac{\partial \Gamma}{\partial \sigma_i^\alpha} \frac{d\sigma_i^\alpha}{dt} = \quad (155) \\ &= - \sum_{\alpha,i} x^\alpha E_i^\alpha \mu (2\sigma_i^\alpha - 1) = -2\mu (\Gamma - \Gamma_0), \end{aligned}$$

where $\Gamma_0 = (1/2) \sum_i E_i$ is a phenotype average.

Similarly, the variance defined in Eq. 148 is also a sum over sites, and the change in the variance due to mutations is:

$$\frac{d\Delta^{\text{mut}}}{dt} = \sum_{i,\alpha} \frac{\partial \Delta}{\partial \sigma_i^\alpha} \frac{d\sigma_i^\alpha}{dt} \approx -4\mu (\Delta - E_0^2) \quad (156)$$

where $E_0^2 = (1/4) \sum_i E_i^2$.

Now we have all the elements, the effective diffusion equation (149) in phenotypic space reads:

$$\begin{aligned} \frac{\partial Q(\Gamma, \Delta, t)}{\partial t} &= \frac{1}{2N} \left[\frac{\partial^2}{\partial \Gamma^2} \Delta + \frac{\partial^2}{\partial \Delta^2} (2\Delta^2) \right] Q(\Gamma, \Delta, t) \\ &+ \left[2c_0 \Delta \frac{\partial}{\partial \Gamma} (\Gamma - E^*) + 2c_0 \frac{\partial}{\partial \Delta} \Delta^2 \right] Q(\Gamma, \Delta, t) \\ &+ \left[2\mu \frac{\partial}{\partial \Gamma} (\Gamma - \Gamma_0) + 4\mu \frac{\partial}{\partial \Delta} (\Delta^2 - E_0^2) \right] Q(\Gamma, \Delta, t). \quad (157) \end{aligned}$$

The stochastic equation for the evolution of the mean phenotype can then be read off as:

$$\frac{d\Gamma}{dt} = -2\mu (\Gamma - \Gamma_0) + \Delta \partial_{\Gamma} F + \sqrt{\frac{\Delta}{N}} \xi_{\Gamma}, \quad (158)$$

where the first term describes mutations, the second term selection and $\partial_{\Gamma} F = -2c_0 (\Gamma - E^*(t))$ and ξ_{Γ} is a normalized Gaussian white noise. $E^*(t)$ can be a time dependent moving fitness maximum. An analogous equation holds for Δ . Assuming that Δ changes on much faster timescales than Γ , it can be replaced by its mean, and the steady state solution has a Boltzmann form $Q_{\text{eq}}(t) = Z^{-1} Q_0(\Gamma) e^{-2NF(\Gamma)}$, where as before $Q_0(\Gamma) \sim \exp \left[-\frac{1}{2} \frac{(\Gamma - \Gamma_0)^2}{\langle \Delta \rangle / (4N)} \right]$ is the distribution with no selection.

In a time dependent environment the peak of the distribution also changes with time according to a prescribed model. The population tries to track the fitness peak, without really ever reaching it. From the physical point of view, the system is maintained out of equilibrium. In a changing fitness landscape, where the population history is described by a sequence of phenotypic trait measurement ($\Gamma_0, \dots, \Gamma_M$) over time (t_0, \dots, t_M) the fitness flux of a population history Φ [303] describes a cumulative selective effect of phenotypic trait changes:

$$\Phi = \sum_{i=1}^M \delta \Gamma_i \partial_{\Gamma_i} F(\Gamma_i, t_i) \neq F(\Gamma_M, t_M) - F(\Gamma_0, t_0). \quad (159)$$

The lack of equality holds in general because F also has an explicit time dependency. One can show that this discrepancy is formally equivalent to dissipation in thermal systems, and can be related to the entropy production:

$$\langle 2N\Phi \rangle = \text{KL}(\mathcal{P}|\mathcal{P}^{\mathcal{T}}) + \text{boundary terms}, \quad (160)$$

where $\mathcal{P}^{\mathcal{T}}$ is the probability of the forward trajectory and \mathcal{P} of the backward trajectory and the Kullback-Leibler divergence is defined in Eq. 227. This dissipation has been called ‘‘fitness flux’’ in the context of population genetics.

These kinds of phenotypic trait models have been used as starting points for studying co-evolution as we describe in section VIII. The main difference is that the fitness objective E^* is itself a function of the composition of the population with which the initial population evolves. That other population is also subject to selection and drift and evolves stochastically, giving rise to coupled equations [268].

11. Lineage reconstruction

Lineage reconstruction is a necessary first step in both BCR evolutionary analysis and phenotypic tracking. While many software methods exist to reconstruct lineages for evolutionary problems, they are not always well adapted for immunological data. Nevertheless they are often used, and more adapted methods are usually built upon classical ones, so here we provide a general overview of these existing approaches.

The first problem in reconstructing lineages requires taking the sequence data and identifying clusters of se-

quences that share a common ancestor, and therefore belong to the same lineage. A classical strategy to cluster datapoints is single-linkage clustering, which builds hierarchical clusters by iteratively merging pairs together [304]. The Partis software [180] uses a likelihood ratio test to decide if a give set of sequences $\sigma_1, \sigma_2, \dots, \sigma_N$ can be grouped as descending from one ancestor or not. Specifically Partis uses a Hidden Markov Model (HMM) based method to annotate each nucleotide in a specific BCR sequence as coming from a given a set of hidden states corresponding to V, D or J genomic template, or from a set of exponentially distributed non-templated insertions. The sum over paths determines the probability $P(\sigma_i)$ for each sequence. The same procedure repeated for a set of N sequences, results in the total probability of a common scenario for these sequences $P(\sigma_1, \dots, \sigma_N)$. If the likelihood ratio $P(\sigma_1, \dots, \sigma_N) / \prod_{i=1}^N P(\sigma_i) > \text{threshold}$, one concludes that the sequences originate from a common recombination. Partis starts with pairs of sequences, keeping the cluster with the largest likelihood ratio, and then adding new members to the cluster based on the same test. In practice, to speed up the algorithm Partis does not test all possible pairs but creates initial subsets of data with low Hamming distances. Joining a cluster is irreversible, which can lead to errors in clustering.

Once a lineage is defined, annotation softwares such as Partis [179] or IGoR [175] can be used to propose the naive root of the tree.

This problem alone is simpler than finding the whole tree genealogy, known as the topology of the tree or the branching pattern, because part of the sequence is templated by the V, D, and J genes (see Sec. VIB). There are two main classes of methods for tree topology reconstruction: maximum parsimony and maximum likelihood. Let us first explain why it is practically impossible to exhaustively sample all tree topologies, and then explain the differences in the two approaches.

We will use the example of binary trees — trees in which each node gives rise to only two branches — since essentially all loopless trees can be cast into a binary form by adding branches of zero length. Starting from the simplest unrooted tree of two leaves connected by one branch, and adding new leaves one at a time, it is simple to convince oneself that in each step we add 1 leaf, 1 internal node and 2 edges, such that a tree with N leaves has $N - 2$ internal nodes and $2N - 3$ edges (or branches). Adding the N th leaf ($N - 1 \rightarrow N$) adds $2N - 5$ topologies, such that the number of unrooted tree topologies with N leaves T_N , is $T_N = (2N - 5)T_{N-1}$, which can be recursively solved to give $T_N = (2N - 5)!!$. For rooted trees, the number of rooted topologies with N leaves is given by T_{N+1} in terms of the number of non-rooted tree topologies, since it just requires adding a root to the existing unrooted tree topologies.

The goal of tree reconstruction is to find the topology that is consistent with the data. For a tree with $N = 10$ leaves, that means exploring $T_N \sim 10^6$ trees and for $N = 50$, $T_N \sim 10^{76}$ trees, which is prohibitively large. BCR

clusters often have hundreds of leaves.

There are two ways of determining distance between two node sequences along an edge of a tree: one is to calculate the Hamming distance by simply counting the number of single nucleotide differences between two sequences h_j , the other is to use the estimated time between two mutations t_j . The number of mutations in a given time t_j is distributed according to a Poisson distribution of mean μt_j : $P(h_j) = \frac{\mu t_j^{h_j}}{h_j!} e^{-\mu t_j}$. The branch length refers to the time between two nodes. The tree length is defined as the total number of substitutions, $h_{\text{tot}} = \sum_{\text{branches}} h_j$, and is not equal to the sum of branch lengths. Maximum parsimony methods use Hamming distance to define distances on trees, which requires knowing the identity of all the internal nodes and the tree topology, whereas maximum likelihood sums over possibilities for the sequences on the internal nodes, but requires a mutational model.

The maximum parsimony approach is easy to formulate: for a given topology, it assigns internal nodes so as to minimize h_{tot} . It then selects topologies with minimal h_{tot} . The parsimony scores differ from 0 (when all the leaves are the same) to NL (when all the leaves are different, which essentially means this position carries no information). There are $(2N - 5)!!$ topologies to scan so many trees have the same h_{tot} , including at its minimum. Therefore there is often a large family of most parsimonious trees. Apart from the practical problems, the assumptions behind the evolutionary model in this approach are not clear [305]. In practice, the reconstructed trees match the true structure if there is little variation in the internal nodes. The main advantage is that it embodies the basic intuition for reconstructing the phylogeny.

In the maximum likelihood method, given an overall lineage age T and tree topology, the likelihood is given by a specific mutation model by varying the branch lengths t_j . Within the independent site assumption, the likelihood factorizes over sites and the log-likelihood $\ell = \sum_{r=1}^L \ln \ell_r$ independently traces the evolution of each site r . Thus the following equations apply to a single site, but generalize readily to many sites. The probability of mutating the base pair from x_i into x_j between nodes i and j , is defined as $P_{x_i, x_j}(t_j)$. For instance, this probability can take the following form, although many others are possible:

$$P_{x_i, x_j}(t_j) = e^{-\mu t_j} \delta_{x_i, x_j} + (1 - e^{-\mu t_j}) \pi_{x_j}, \quad (161)$$

where π_x (with $\sum_x \pi_x = 1$) describes the probability that a mutation results in base pair x , and is equal to $1/4$ for the unbiased 4-base pair model, and t_j is the branch length (time between nodes i and j). The first term corresponds to no mutation, and the second to a mutation (including one into the same base pair). This formulation guarantees detailed balance: $\pi_x P_{xy}(t) = \pi_y P_{yx}(t)$ for all t .

We can now build the likelihood recursively, assuming a fixed root identity 0, with sequences denotes as x_0 that

are not fixed. First consider a tree made of a root and two branches of length t_1 and t_2 leading to leaves 1 and 2, whose identity x_1 and x_2 is known. The likelihood of this tree is:

$$L = \sum_{x_0} \pi_{x_0} P_{x_0 x_1}(t_1) P_{x_0 x_2}(t_2). \quad (162)$$

We can now add an internal node 4 to the tree, descending from 0 and ancestor of 1 and 2, and a leaf 3 descending from 0:

$$\begin{aligned} L &= \sum_{x_0} \sum_{x_4} \pi_{x_0} P_{x_0 x_4}(t_4) P_{x_0 x_3}(t_3) P_{x_4 x_1}(t_1) P_{x_4 x_2}(t_2) = \\ &= \sum_{x_0} \pi_{x_0} P_{x_0 x_3}(t_3) \sum_{x_4} P_{x_0 x_4}(t_4) L_4(x_4), \end{aligned} \quad (163)$$

where $L_4(x_0) = P_{x_4 x_1}(t_1) P_{x_4 x_2}$. This recursive form allows for efficient calculation of likelihood by successively joining roots of trees into a new root. Formally, joining root j and k into common parent i :

$$L_i(x_i) = \left[\sum_{x_j} P_{x_i x_j}(t_j) L_j(x_j) \right] \left[\sum_{x_k} P_{x_i x_k}(t_k) L_k(x_k) \right], \quad (164)$$

Leaves are initialized to $L_j(x_j) = 1$, and the final results is given by:

$$L = \sum_{x_0} \pi_{x_0} L_0(x_0). \quad (165)$$

Given a fixed topology, we can now maximize the likelihood L over the set of intermediate branch lengths $\{t_j\}$, which results in a score for each topology. The topology \mathcal{T} with the best ranking score, is the most likely topology (similarly to the maximum parsimony approach).

Existing maximum likelihood methods use the independent site assumption to decrease computational time. Many methods also assume homogeneous mutation rates across sites and the reversibility of the evolutionary process. The GTRGAMMA substitution model of the RAxML software [306] uses gamma-distributed mutation rates for different base pairs, and IgPhyML [307] encodes a non-reversible mutation model that effectively accounts for the context dependence of mutations (although in a site-independent setup).

Davidsen and Matsen [308] compared maximum likelihood and maximum parsimony methods for BCR lineage reconstruction. They concluded that an improved and informed maximum parsimony method outperforms classic maximum likelihood.

12. Population growth rates

Another way to study the evolution of populations is to study dynamics that maximize the long term growth rate

of the population given a fixed environment [309, 310]:

$$\Lambda(\text{dynamics, environment}) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0} \ln(N_t), \quad (166)$$

where N_t is the population size at a given time, whose evolution is driven by the dynamics. One needs to specify classes of evolutionary dynamics to consider, and the class of interactions it has with the environment. The framework is clearly very general, and reduces to a constrained optimization problem. The main conceptual difference with population genetics models studied in this section is that populations cannot go extinct, and all possible variants are represented. Populations are in a regime where fast growth governs evolution. Alternative approaches could consider optimal strategies for avoiding extinction [311], by minimizing the probability of extinction, P_{ext} , whose form depends on the problem considered.

D. Ecological models

Population of lymphocytes interact with other through signaling, competition for resources, cytokines, or antigens. For this reason, concepts and mathematical models from ecology are often useful to describe their dynamics.

1. Generalized Lotka–Volterra models

Ecological generalized Lotka–Volterra models [312–314] describe the co-habitation of multiple species in the same environment and account for their interactions, both direct (one species needs another to reproduce) and indirect (through competition for external resources). In general these models have the form:

$$\frac{dx_i}{dt} = x_i [\alpha(\vec{x}, \vec{c}) - \beta(\vec{x}, \vec{c})], \quad (167)$$

where x_i is the frequency of a given species, c_i are external factors such as nutrients $\alpha(\vec{x}, \vec{c})$ is the growth rate and $\beta(\vec{x}, \vec{c})$ the death rate with $\vec{x} = (x_1, x_2, \dots)$ and $\vec{c} = (c_1, c_2, \dots)$. The form of the dependence of the growth rate on the different species and external factors determines the non-linearity of the problem. This is usually encoded by an interaction matrix between species. Models where the growth and death rates are the same for all species are called neutral. Often one introduces a carrying capacity which is a form of non-linearity that describes competition for resources.

This formulation of the model is very general and the solution depends on the specific assumptions. In general these models are solved numerically, and the results depend on the numerical values of the parameters that are often not known with great certainty. These models are often used to ask questions about the coexistence of different species, as well as speciation itself — why

are ecological environments with different species stable? Since we often do not have experimental information to parametrize the interaction matrix, random matrix models have been successfully used to effectively describe the interactions [315]. In this approach one uses the fact that within a family of random matrices (i.e. matrices whose elements are chosen from the same distribution) the eigenvalues of these matrices are the same. Near the fixed point, the stability of the system is explored by linearizing the system of equations in Eq. 167, and the eigenvalues of this linearized matrix determine the stability of the ecosystem:

$$\frac{dx_i}{dt} = -x_i + \epsilon \sum_j K_{ij} x_j(t), \quad (168)$$

where ϵ is the interaction strength and \mathbf{K} is a random interaction matrix. As a result the stability of the system does not depend on the realization of the interaction matrix, just on its statistics. For the matrix to be stable, all the eigenvalues of \mathbf{K} , λ_i have to satisfy $\epsilon \lambda_i \leq 1$, which is fulfilled in the largest eigenvalue satisfies $\lambda_{\max} < 1/\epsilon$. In the case of Gaussian random matrices, the properties are determined by their first two moments and a strong transition occurs for $N \rightarrow \infty$ where the system is stable for $\epsilon < 1/\sqrt{2}$ and otherwise unstable. This result holds only if the connectivity scales with the size of the system.

Incidentally, random matrix theory also has been used when looking at covariances in sequence variation (see section XII E 6). The general idea is that beyond the first couple of eigenvalues of the covariance matrix of amino acid variability (which often stems from phylogenetic bias), eigenvalues are well approximated by a random matrix. For $R \times C$ Gaussian random matrices, where the elements K_{ij} are chosen from a Gaussian distribution with mean 0 and variance σ , the density of eigenvalues of the covariance matrix $Y = \langle KK^T \rangle / R$ is given by the Marchenko-Pastur distribution [316]:

$$\rho(x) = \frac{1}{2\pi\sigma^2} \frac{\sqrt{(\theta_+ - x)(x - \theta_-)}}{\theta x}, \quad (169)$$

where $\theta = C/R$ and $\theta_{\pm} = \sigma^2(1 \pm \sqrt{\theta})$.

2. Susceptible-Infected-Recovered (SIR) Models

Susceptible-Infected-Recovered (SIR) models are used to study the spread of epidemics at the population level. An SIR model considers three possible kinds of individuals: ones that are susceptible to the infection (S), infected (I) and recovered (R) and therefore usually immune. The simplest SIR model assumes that the birth-death of the host individuals happens on slower timescales than the spread of the epidemic itself. This model is a good description of measles and other infections by slowly evol-

ing pathogens.

$$\frac{dS}{dt} = -\beta IS, \quad (170)$$

$$\frac{dI}{dt} = \beta IS - \nu I = xI, \quad (171)$$

$$\frac{dR}{dt} = \nu R, \quad (172)$$

where β is an effective infection rate, ν is the rate of recovery, and $\nu(\beta S/\nu - 1) = x$ defines an effective growth rate, or fitness of the host population. Within this formulation no one can die. Since $\partial_t S + \partial_t I + \partial_t R = 0$, the total population size is constant, $N = S + I + R = \text{const}$. The typical recovery time is $T_r = \nu^{-1}$ and the typical timescale for a infected-susceptible individual interaction is $T_{\text{inf}} = \beta^{-1}$, meaning that an infected individual manages to infect $R_0 = T_r/T_{\text{inf}} = \beta/\nu$ individuals before recovering. R_0 is also called the infection radius or basic reproduction number.

Using Eqs. 170-172, we find $S = S(t = 0)e^{-R_0(R-R(t=0))}$. At infinite time ($t \rightarrow \infty$, $I = 0$), the number of recovered individuals reads:

$$R_{\infty} = N - S(t = 0)e^{-R_0(R_{\infty} - R(t=0))}, \quad (173)$$

which means that at an end of an epidemic there are still susceptible individuals. From Eq. 171 we find that $S(t = 0)R_0 > 1$, leads to $\partial_t I > 0$ and an exponentially growing infection (i.e. epidemic outbreak), whereas $S(t = 0)R_0 < 1$ leads to $\partial_t I < 0$ and quenching of the infection. Since these equations are deterministic, the initial condition determines the infection spread for all times. The effective number of susceptible individuals that can propagate the infection determines its future.

SIR equations can be extended to account for host birth (with rate α) and death (with rate μ), that lead to equations of the form:

$$\frac{dS}{dt} = \alpha - \mu S - \beta IS, \quad (174)$$

$$\frac{dI}{dt} = \beta IS - \nu I - \mu I = xI, \quad (175)$$

$$\frac{dR}{dt} = \nu R - \mu R, \quad (176)$$

with two stable steady state solution: pathogen free, $(S, I, R) = (\alpha/\mu, 0, 0)$; and permanent infection, $(S, I, R) = ((\nu + \mu)/\beta, (\mu/\beta)(R_0 - 1), (\gamma/\beta)(R_0 - 1))$, where $R_0 = \beta\alpha/(\mu(\mu + \nu))$ is the infection radius. For $R_0 \leq 1$, the system evolves towards the pathogen-free solution, while for $R_0 \geq 1$ the permanently infected solution is reached.

The basic SIR model can be made more complicated in many different ways. One other variation called the SIS (Susceptible - Infected - Susceptible) model describes a situation where an infected individual does not have long lasting immunity. This model is inspired by influenza, where the virus mutates fast enough that a past infection does not necessarily guarantee immunity (although

current models of influenza account for more detailed descriptions of cross-reactivity – see section IX). The model just has two states:

$$\frac{dS}{dt} = -\beta IS + \nu I, \quad (177)$$

$$\frac{dI}{dt} = \beta IS - \nu I = xI. \quad (178)$$

Using $S = N - I$, we obtain a single differential equation:

$$\frac{dI}{dt} = (\beta N - \nu)I - \beta I^2, \quad (179)$$

which has two fixed points $I = 0$ for $R_0 = (\beta/\nu)N < 1$ and $I = (\beta N - \nu)/\beta$ for $R_0 > 1$. The second fixed results from the solution:

$$I(t) = \frac{(\beta N - \nu)}{\beta} \frac{1}{1 + \left(\frac{\beta N - \nu}{\beta I_0} - 1\right) e^{(\nu - \beta N)(t - t_0)}}. \quad (180)$$

We note that the SIR equations (Eqs. 170-172) differ in its assumptions from the two species Lotka-Volterra equations (see section XII D 1) between and a host and a pathogen. Specifically, SIR equations do not assume that everyone can get infected, while Lotka-Volterra equations do. However, generalized Lotka-Volterra equations can easily be modified to account for different subclasses of hosts. Lotka-Volterra equations on the other hand do explicitly model the viral population, and allow for host-pathogen oscillations that SIR models do not, and they do not assume a constant population size. Notably, the SIS equations (Eq. 177) are a two-species realization of Lotka-Volterra dynamics, with the additional assumption for constant population size.

In section VII we describe equations that model the in-host evolution during HIV infections. They are very similar in spirit to SIS equations which explicitly consider viral dynamics.

3. Solution of stochastic population dynamics with a source

Here we present the calculations leading the results presented in Sec. VII, some of which are original to this review. To solve Eq. 62, we define the generating functions:

$$G(x) = \sum_{C=1}^{+\infty} N_C x^C, \quad (181)$$

$$\Theta(x) = \sum_{C=1}^{+\infty} \theta_C x^C. \quad (182)$$

The time evolution of the population is then governed by:

$$\frac{dG(x)}{dt} = \Theta(x) - \mu N_1 + (\nu x^2 + \mu - (\mu + \nu)x)G'(x). \quad (183)$$

At steady state, we have:

$$G'(x) = \frac{\Theta(1) - \Theta(x)}{(\nu - \mu x)(1 - x)}, \quad (184)$$

where we have used the balance between birth and death of clones $dG(1)/dt = \Theta(1) - \nu N_1 = 0$. Using $(\Theta(1) - \Theta(x))/(1 - x) = \sum_{C>0} \theta_C \sum_{i=0}^{C-1} x^i$, and expanding $1/(\nu - \mu x) = (1/\nu) \sum_{j=0}^{\infty} (\mu/\nu)^j x^j$, we obtain:

$$N_C = \frac{1}{\mu - \nu} \frac{1}{C} \left\{ \left(\frac{\nu}{\mu}\right)^C \sum_{k=1}^{C-1} \theta_k \left[\left(\frac{\mu}{\nu}\right)^k - 1\right] + \sum_{k=C}^{\infty} \theta_k \left[1 - \left(\frac{\nu}{\mu}\right)^C\right] \right\}, \quad (185)$$

which reduces to

$$N_C = \frac{\theta}{\nu} \frac{1}{C} \left(\frac{\nu}{\mu}\right)^C \quad (186)$$

in the simple case $\theta_C = \theta \delta_{C,1}$. The total number of clones reads:

$$N_{\text{tot}} = G(1) = \frac{1}{\mu - \nu} \sum_k \theta_k \left\{ \sum_{C=1}^k \frac{1 - \left(\frac{\nu}{\mu}\right)^C}{C} + \left[\left(\frac{\mu}{\nu}\right)^k - 1\right] \sum_{C=k+1}^{\infty} \left(\frac{\nu}{\mu}\right)^C \frac{1}{C} \right\}, \quad (187)$$

and

$$N_{\text{tot}} = (\theta/\nu) \ln[1/(1 - \nu/\mu)] \quad (188)$$

for $\theta_C = \theta \delta_{C,1}$. In the limit of rare divisions $\nu \ll \mu$, only clone sizes below the initial clone size, $C \leq k$, contribute, so that

$$N_{\text{tot}} = \frac{\sum_k \theta_k H_k}{\mu} = C_{\text{tot}} \frac{\sum_k H_k \theta_k}{\sum_k k \theta_k}, \quad (189)$$

where $H_k = \sum_{i=1}^k 1/i$ are harmonic numbers. In the opposite limit of division balancing death, $\mu - \nu \ll \mu$, the long tail of the second sum dominates:

$$N_{\text{tot}} = \frac{\ln(1/\epsilon) \sum_k k \theta_k}{\mu} = C_{\text{tot}} \epsilon \ln(1/\epsilon), \quad (190)$$

with $\epsilon = (\mu - \nu)/\mu$.

Eq. 183 can be solved out of steady state with the method of characteristics. Defining $F(x) = G(x) - G(1)$ and making the change of variable $y = \ln[(x - \mu/\nu)/(x - 1)]/(\mu - \nu)$, this equation reduces to:

$$\left(\frac{\partial}{\partial t} - \frac{\partial}{\partial y}\right) F = \Theta(x) - \Theta(1). \quad (191)$$

This equation can be solved in absence of a source term ($\Theta = 0$): $F_0(x, t) = A(t + y)$, with $A(y) = F_0(x, 0)$, yielding:

$$F_0(x, t) = F_0 \left(\frac{\mu/\nu - (x - \mu/\nu)/(x - 1)e^{t(\mu-\nu)}}{1 - (x - \mu/\nu)/(x - 1)e^{t(\mu-\nu)}}, 0 \right) \quad (192)$$

Starting with a single clone of size s at $t = 0$, $G_0(x, 0) = x^s$ and $F_0(x, 0) = x^s - 1$, the coefficients of the solution $G_0(x, t)$ to the homogeneous equation (Eq. 191 with $\Theta = 0$) give Green's function, defined as the probability that the clone starting at size s at $t = 0$ has size C at a later time t , $G_0(x, t) = \sum_C P(C, t|s, 0)x^C$. Expanding (192) in x , one obtains:

$$P(C, t|s, 0) = g(C, s, t) = \frac{r^C}{(1 - rz)^{C+s}} \times \sum_{n=1}^{\min(C, s)} \binom{s}{n} \binom{C-1}{n-1} r^n (1 - r^{-1})^{2n} z^n (1 - z)^{s+C-2n}, \quad (193)$$

with the shorthands $z = e^{-t(\mu-\nu)}$, $r = \nu/\mu$. The combinatorial factor in the sum comes from counting the number of ways there are to choose $s - n$ factors of order 0 in x , $\binom{s}{n}$, and n factors of order (j_1, \dots, j_n) so that $\sum_{i=1}^n j_i = C$, $\binom{C-1}{n-1}$, in a series expansion of the form $[a + b((cx) + (cx)^2 + (cx)^3 + \dots)]^s = [a + bcx/(1 - cx)]^s$ to which (192) can be reduced.

We can now use Green's function to calculate the dynamics of the distribution with an arbitrary initial condition:

$$N_C(t) = \sum_k N_k(0)g(C, k, t) + \int_0^t dt' \sum_k \theta_k(t')g(C, k, t-t'). \quad (194)$$

In the case of an empty immune system at $t = 0$, $N_C(0) = 0$, and constant thymic output θ_k , we get:

$$N_C(t) = \frac{r^C}{\mu - \nu} \sum_k \theta_k \sum_{n=1}^{\min(C, k)} \binom{k}{n} \binom{C-1}{n-1} \frac{r^n (1 - r^{-1})^{2n}}{n} \times [z^n F_1(n, 2n - k - C, C + k, 1 + n, z, rz)]_{e^{-t(\mu-\nu)}}, \quad (195)$$

where F_1 is the Appell hypergeometric function of two variables, and where we made the change of variable $z = e^{-(t-t')(\mu-\nu)}$ to carry out the integral in (194). This equation simplifies greatly for $\theta_k = \theta \delta_{k,1}$:

$$N_C(t) = \frac{1}{\nu} \frac{r^C}{C} \left(\frac{1 - e^{-t(\mu-\nu)}}{1 - r e^{-t(\mu-\nu)}} \right)^C. \quad (196)$$

4. Solution to forward jump process with opposing drift and source

Consider the process described by (68) generalized to an arbitrary distribution of jumps:

$$\frac{\partial \rho}{\partial t} = \mu \frac{\partial \rho}{\partial x} + \int_0^\infty dy J(y) [\rho(x - y) - \rho(x)] + \tilde{\theta}(x), \quad (197)$$

where absorbing boundary condition at $x = 0$. where $J(y)dy$ is the rate (per unit time) of jumps of size between y and $y + dy$. This equation can be solved exactly by considering the Laplace transform of ρ :

$$\hat{\rho}(k, t) = \int_0^\infty dx e^{-kx} \rho(x, t), \quad (198)$$

which satisfies:

$$\frac{\partial \hat{\rho}}{\partial t} = \mu(\rho(0, t) + k\hat{\rho}(k, t)) + (\hat{J}(k) - \hat{J}(0))\hat{\rho}(k, t) + \hat{\theta}(k), \quad (199)$$

where $\hat{J}(k)$ and $\hat{\rho}(k)$ are the Laplace transforms for J and $\tilde{\theta}$. The steady state solution is:

$$\hat{\rho}(k) = \frac{\hat{\theta}(k) - \mu\rho(0)}{\hat{J}(k) - \hat{J}(0) + \mu k}. \quad (200)$$

To get a physical solution, a pole at $k = 0$ cannot exist, so that the total number of clones $\hat{\rho}(0)$ is well defined. This is satisfied for $\mu\rho(0) = \hat{\theta}(0)$, which is simply the condition that the rate of birth of new clones, $\hat{\theta}(0)$ be balanced by the rate of death, $\mu\rho(0)$. Then solution becomes

$$\hat{\rho}(k) = \frac{\hat{\theta}(k) - \hat{\theta}(0)}{\hat{J}(k) - \hat{J}(0) + \mu k}. \quad (201)$$

To understand the behaviour of $\rho(x)$ at large x , we have to examine poles for negative k , $k = -\alpha$, which satisfy the condition:

$$\int_0^\infty dx J(x)(e^{\alpha x} - 1) = \mu\alpha. \quad (202)$$

The left-hand side has derivative $\bar{\nu} \equiv \int dx J(x)x$ with respect to k_0 , which is the average growth rate of a clone. To guarantee that all clones eventually go extinct, that number must be smaller than μ . Therefore (202) has only one solution $\alpha > 0$. At large x the behaviour is dominated by that pole, $\hat{\rho}(k) \sim 1/(k+\alpha)$, yielding $\rho(x) \sim e^{-\alpha x}$ and:

$$N_C \sim \frac{1}{C^{1+\alpha}}. \quad (203)$$

The total numbers of clones and cells read:

$$N_{\text{tot}} = \hat{\rho}(0) = \frac{\int_0^\infty dx \theta(x)x}{\int_0^\infty dx J(x)x - \mu}, \quad (204)$$

$$C_{\text{tot}} = \hat{\rho}(-1) = \frac{\int_0^\infty dx \theta(x)(e^x - 1)}{\int_0^\infty dx J(x)(e^x - 1) - \mu}. \quad (205)$$

In the specific case of fixed-size jump, $J(x) = sp\delta(x - \Delta x)$, we get back (70) and (71). Eq. 202 simplifies in the limit of many very small jumps, corresponding to $J_{\text{tot}} \equiv \int_0^\infty dx J(x) \rightarrow \infty$ while keeping the average effect of jumps, $\bar{\nu} < \mu$, as well as its second moment, $\Gamma = \int_0^\infty dx J(x)x^2$, finite, meaning that the average jump size $\langle \Delta x \rangle = \nu/J_{\text{tot}}$ goes to zero. Expanding (202) at small x , we get:

$$\alpha = 2 \frac{\mu - \bar{\nu}}{\Gamma}. \quad (206)$$

which gives back (72) in the case of fixed-size jumps.

5. The Yule process

The Yule-Simon [317, 318] process is another type of “neutral” process that is characterized by a distribution with power long tails, which in its tail is reminiscent of Zipf’s law (see section VII B). The model was initially formulated to describe the distribution of biological genera. It was later adapted in network science and is known under the name of preferential attachment, or the rich-get-richer model. In general it is a realization of a birth-death process, where the birth coefficient is proportional to the number of individuals. In the model, new individuals are added at each time step. With probability α , the new individual gets a mutation (or undergoes a speciation event if we think about species), and with probability $1 - \alpha$ the new individual’s type is chosen among all the possible ancestors, i.e. proportionally to the abundance of each type.

The probability that the rank of an individual is k results in a power law distribution at large C :

$$\frac{N_C}{C_{\text{tot}}} = \frac{\alpha}{2 - \alpha} \frac{\Gamma(C)\Gamma(\rho + 1)}{\Gamma(C + \rho + 1)} \sim \frac{1}{C^{\rho+1}}, \quad (207)$$

where $\Gamma(x)$ denotes a gamma function, and $\rho = 1/(1 - \alpha)$.

A precise derivation of (207) can be found in [318], but the power-law behaviour can be understood intuitively in a continuous approximation. Assume that the rate of division is μ , so that there are $C_{\text{tot}} = e^{\mu t}$ individual at time t , starting from a single individual at time $t = 0$. The rate of new emergence of new mutants is $\alpha\mu e^{\mu t}$, so that at time t there are $n(t) = \int_0^t dt' \alpha\mu e^{\mu t'} = \alpha(e^{\mu t} - 1)$ new types. The key point is that the abundance of a mutant is an exponentially increasing function of its age. A mutant that arose at t' has size $C = e^{\mu(1-\alpha)(t-t')}$ at time t . $n(t')$ can thus be viewed as the rank of the mutant, ordered by increasing frequency. The abundance of the mutant arising at time t' , of rank $n = n(t')$ is, at time t :

$$C = \frac{e^{\mu(1-\alpha)t}}{(1 + n/\alpha)^{(1-\alpha)}} \sim \frac{1}{n^{1-\alpha}}, \quad (208)$$

where the approximation is only valid at large n . We recognize Zipf’s law (56), which implies a power law in the clone size distribution according to (55):

$$N_C \sim \frac{1}{C^{1+1/(1-\alpha)}} = \frac{1}{C^{\rho+1}}. \quad (209)$$

E. Inference

Here we introduce basic concepts and methods for constructing models and inferring their parameters from data. These approaches are becoming increasingly important as large datasets are being produced by high-throughput methods, from imaging to sequencing. Some of these methods have been used in only a few applications in computational immunology, but we anticipate that their usage will become widespread in future studies.

1. Probabilistic inference, maximum likelihood and Bayesian statistics

Often, empirical data originates from a stochastic process. The first source of variability, which is present in almost any measurement, is experimental noise, which is random. In many cases related to biological systems, from gene expression to cell signaling and random recombination of DNA, the underlying processes are intrinsically stochastic, and so are the models that must describe them.

Let us assume that a model can be encapsulated into a probability distribution for observations $P(\mathbf{x}|\theta)$, where \mathbf{x} is the empirical data, and θ the set of parameters defining the model. In many cases, the data can be decomposed into a set of independent datapoints, $\mathbf{x} = (x_1, \dots, x_N)$, so that the probability distribution factorizes over these observations:

$$P(\mathbf{x}|\theta) = \prod_{i=1}^N P(x_i|\theta). \quad (210)$$

This probability is called the likelihood (of the data given the model). Because of its often multiplicative form, it is common to consider instead its logarithm, the log-likelihood $\mathcal{L}(x|\theta)$, which is additive, or its normalized variant $\ell(x|\theta) = (1/N) \sum_{i=1}^N \ln P(x_i|\theta)$. One can think of x_i as a sequence, for instance of an immune receptor or of a pathogen; a fluorescence signal coming from flow cytometry or single-cell microscopy imaging; an abundance of RNA transcripts for gene expression, or of sequencing reads of immune receptors when estimating the abundance of lymphocyte clones. The models may be statistical models for the occurrence of a particular data point, or stochastic dynamical models of the process.

A popular way to infer the parameters $\theta = (\theta_1, \dots, \theta_K)$ of a model is to find those that maximize the likelihood of

the data:

$$\theta^*(\mathbf{x}) = \arg \max_{\theta} P(\mathbf{x}|\theta). \quad (211)$$

This estimate of the parameters is called the maximum likelihood estimator, and it can be shown to be unbiased and optimal in the limit of large numbers of observation, in the sense that it saturates the Cramer-Rao bound. The Cramer-Rao bound gives a lower bound on the variance of unbiased estimators of a fixed and deterministic parameter. In practice, this implies that the fluctuations of the estimator around its true value, $\theta^*(\mathbf{x}) - \theta$ are given by a multivariate Gaussian distribution of mean 0 and co-variance I^{-1}/N , with:

$$I_{ab}(\theta) = \sum_{\mathbf{x}} P(\mathbf{x}|\theta) \frac{\partial \ln P(\mathbf{x}|\theta)}{\partial \theta_a \partial \theta_b} \quad (212)$$

is called the Fisher information matrix.

When one has information about the process and its parameters independently of the data, through a prior belief in the distribution of θ , $P_{\text{prior}}(\theta)$, this information can be combined with the data using Bayes rules:

$$P(\theta|\mathbf{x}) = \frac{P(\mathbf{x}|\theta)P_{\text{prior}}(\theta)}{P(\mathbf{x})}. \quad (213)$$

This expression defines the posterior distribution of parameters, given the data. For instance, some parameters may have been measured with some uncertainty, which can be modelled using a Gaussian distribution $P_{\text{prior}}(\theta_i) = (1/\sqrt{2\pi\sigma_i^2}) \exp[-(\theta_i - \theta_{i,\text{meas}})^2/\sigma_i^2]$, or there exists natural physiological range for them. A uniform prior, $P_{\text{prior}} = \text{const}$, amounts to considering the likelihood alone. Also note that in the limit of large datasets, $N \gg 1$, the log-likelihood dominates over the logarithm of the prior, which becomes negligible. In that case, the posterior over possible parameters, $P(\theta|\mathbf{x})$, is given by a Gaussian distribution over θ with mean θ^* and covariance given by the same Fisher information matrix I^{-1}/N . This means that the Bayesian fluctuations, which correspond to our uncertainty about the true values of the parameters, match the fluctuations of the error we make when picking the maximum-likelihood estimate.

Using Eq. 213, one can either consider the full range of acceptable parameter values from the posterior, by e.g. sampling from it. Alternatively, one can take its maximum, $\theta^* = \arg \max_{\theta} P(\theta|\mathbf{x})$, called the maximum a posteriori estimator. In the limit of large N , that estimator is equivalent to the maximum likelihood estimator.

We now turn to two simple examples of inference problems for which the maximum likelihood estimator recovers an intuitive answer. First consider linear regression. Suppose we have data points (x_1, \dots, x_N) taken at times (t_1, \dots, t_N) , which we want to fit with a linear model, $x = \theta_1 + \theta_2 t$. Implicitly, we must assume that the data is noisy, otherwise a fit would not be necessary. Assuming

Gaussian noise, the stochastic model reads:

$$\begin{aligned} P(\mathbf{x}|\theta) &= \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x_i - \theta_1 - \theta_2 t_i)^2/2\sigma^2} \\ &= \frac{1}{(2\pi\sigma^2)^{N/2}} \exp \left[-\frac{1}{2\sigma^2} \sum_{i=1}^N (x_i - \theta_1 - \theta_2 t_i)^2 \right] \end{aligned} \quad (214)$$

where σ is the amplitude of the noise. Examining this expression, one can see maximizing the likelihood corresponds to minimizing the mean squared error, $\sum_{i=1}^N (x_i - \theta_1 - \theta_2 t_i)^2$, which is the standard way to do linear fit. This argument generalizes to any fit of a parametrized function, assuming Gaussian distributed noise.

The second example is estimating frequencies from a finite set of outcomes. For concreteness, let us assume that we want to describe the probability of finding a particular nucleotide, $x = \text{A, C, G, T}$, at a particular location in a genomic sequence. The parameters of the model are the frequencies of A, C, and G, denoted $\theta_1, \theta_2, \theta_3$, while the frequency of T is $1 - \theta_1 - \theta_2 - \theta_3$. The probability of observing a certain set of observations $\mathbf{x} = (x_1, \dots, x_N)$ is given by:

$$P(\mathbf{x}|\theta) = \theta_1^{n_A} \theta_2^{n_C} \theta_3^{n_G} (1 - \theta_1 - \theta_2 - \theta_3)^{n_T} \quad (215)$$

where $n_A = \sum_{i=1}^N \delta_{x_i, A}$ and similarly for the other nucleotides. Maximizing the likelihood with respect to the parameters yields the intuitive counting estimator: $\theta_1^* = n_A/(n_A + n_C + n_G + n_T)$, and similarly for the other nucleotides. Note that this argument holds for sequences of nucleotides rather than single nucleotides, provided that the underlying probabilistic model assumes that the choices of nucleotides at each position along the sequence are independent of each other.

2. Model selection

The maximum likelihood and Bayesian rules outlined above assume that the class of models, if not its parameters, are known, but this is not always the case. For instance, we may have competing hypotheses between different models, or we may consider models of increasing complexity. For instance, in the example of linear regression, we may want to know whether the data is not better explained by a quadratic model $x = \theta_1 + \theta_2 t + \theta_3 t^2$. Comparing the likelihood directly always favors models with more parameters, which can in turn lead to overly complex models with the risk of overfitting the data. We need a way to compare models with different structures and numbers of parameters.

The general approach to compare two model A and B, described by probabilities $P(\mathbf{x}|\theta_A, A)$ and $P(\mathbf{x}|\theta_B, B)$, is to evaluate the overall probability of the data given each

model, by integrating the over the parameters:

$$P(\mathbf{x}|A) = \int d^{K_A} \theta_A P_A(\mathbf{x}|\theta_A) P_{\text{prior}}(\theta_A|A) \quad (216)$$

$$P(\mathbf{x}|B) = \int d^{K_B} \theta_B P_B(\mathbf{x}|\theta_B) P_{\text{prior}}(\theta_B|B). \quad (217)$$

The relative probabilities of each model is then evaluated by Bayes's rule:

$$P(A|\mathbf{x}) = \frac{P(\mathbf{x}|A)P_{\text{prior}}(A)}{P(\mathbf{x}|A)P_{\text{prior}}(A) + P(\mathbf{x}|B)P_{\text{prior}}(B)}. \quad (218)$$

In the limit of large N , $P_A(\mathbf{x}|\theta_A)$ (and the same for B) is very peaked around its maximum likelihood value, with Gaussian fluctuations around that peak given by the Fisher Information matrix I_A :

$$P_A(\mathbf{x}|\theta_A) \approx P(\mathbf{x}|\theta_A^*) e^{-N(\theta_A - \theta_A^*) I_A (\theta_A - \theta_A^*) / 2}. \quad (219)$$

By the saddle-point approximation we obtain:

$$P(\mathbf{x}|A) \approx \sqrt{(2\pi/N)^{K_A} / \det(I_A)} P_A(\mathbf{x}|\theta_A^*) P_{\text{prior}}(\theta_A^*|A). \quad (220)$$

If we replace this expression into Eq. 216, and focus on two terms that do not depend on our prior assumptions, we see two terms: the maximized likelihood for each model $P_A(\mathbf{x}|\theta_A^*)$, which quantifies the quality of the fit, and ‘‘parsimony’’ term $\sqrt{(2\pi/N)^{K_A} / \det(I_A)}$, which quantifies the volume of parameters that are consistent with each model. As evident for this expression of the parsimony, the more parameters, the smaller that volume will be.

The ratio of probabilities in (216) between A and B is sometimes called Bayes's factor, while the ratio of the parsimony terms is called Occam's factor. If we expand the logarithm of (216) in $o(\ln(N))$ using (220), we obtain, up to a -2 factor, the so called Bayesian Information Criterion (BIC):

$$-2 \ln P(A|\mathbf{x}) \approx K_A \ln(N) - 2N\ell(\mathbf{x}|\theta_A^*) \equiv \text{BIC}(A). \quad (221)$$

The BIC is a popular score for comparing models, and can be intuitively understood as a correction to the likelihood by a term that penalizes large numbers of parameters. In the large N limit, Eq. 218 tells us that the model with the smallest BIC is more likely to be correct. Another popular score is the Akaike information criterion (AIC), defined as $2K - 2N\ell(\mathbf{x}|\theta^*)$, which puts a smaller penalty on the number of parameters.

Applications of model selection criterion have been numerous in computational immunology, see e.g. [78, 97, 200, 233, 319].

3. Expectation-Maximization

Some models are better defined in terms of variables that are not accessible to the observer. If we call those

hidden variables \mathbf{h} , and \mathbf{x} the visible variables, then the likelihood of the data can be expressed as the sum over the hidden variables:

$$P(\mathbf{x}|\theta) = \sum_{\mathbf{h}} P(\mathbf{x}, \mathbf{h}|\theta) = \prod_{i=1}^N \sum_{h_i} P(x_i, h_i|\theta). \quad (222)$$

Expectation-Maximization (EM) consists of maximizing $P(\mathbf{x}|\theta)$ iteratively. Starting with a guess θ^t , one can write a pseudo log-likelihood:

$$\tilde{\mathcal{L}}(\mathbf{x}|\theta) = \sum_{i=1}^N \sum_{h_i} P(h_i|x_i, \theta^t) \ln P(x_i, h_i|\theta), \quad (223)$$

which is essentially an average of the log-likelihood of the full model (i.e. including the hidden variables), weighted by the posterior distribution of the hidden variable h_i under the current model θ_t . In an iteration step, one sets $\theta_{t+1} = \arg \max_{\theta} \tilde{\mathcal{L}}(\mathbf{x}|\theta)$. As this scheme is iterated, θ_t converges to a maximum of the true likelihood $P(\mathbf{x}|\theta)$.

Expectation-maximization has notably been used in the inference of generation model for antigen-specific lymphocyte receptors (TCR and BCR), in Refs. [169, 170, 178, 188].

4. Hidden Markov models

Hidden Markov models (HMM) are models with a hidden variable h which is the trajectory of a Markov process, $h = (h_x)_{x=1, \dots, L}$:

$$P(h) = p_1(h_1) \prod_{j=2}^L w(h_j|h_{j-1}). \quad (224)$$

The visible variable x , is also a vector of size L , whose elements x_j are called ‘emissions’ and are drawn according to the value of the hidden variable according to the emission probabilities $P(x|h) = \prod_{j=1}^L e_j(x_j|h_j)$. A key feature of HMM is that high-dimensional sums over $h = (h_1, \dots, h_L)$ can be performed recursively using technique which is equivalent to the transfer matrix technique in statistical physics. For instance, the forward algorithm allows one to calculate $P(x)$, which requires to sum over all possible trajectories of the hidden variable. This is done by defining $z_j(h_j) = P(x_1, \dots, x_j, h_j)$ and using the recursive equation:

$$z_j(h_j) = \sum_{h_{j-1}} e_j(x_j|h_j) w_j(h_j|h_{j-1}) z_{j-1}(h_{j-1}). \quad (225)$$

The probability $P(x) = P(x_1, \dots, x_L)$ is then obtained as $\sum_{h_L} z_L(h_L)$. A similar trick allows one to calculate any marginal of the model, e.g. $P(x_j)$ or correlation functions. Combining these recursions with EM defines a powerful algorithm for learning the parameters of the model, $\theta = \{w_j(\cdot|\cdot), e_j(\cdot|\cdot)\}$, called the Baum-Welch algorithm.

HMM have been applied to model BCR and TCR annotation to the germline in various software, e.g. SODA [177], Partis [179], and repgenhmm [320].

5. Information theory

The **entropy** of a probability distribution $P(x)$ is defined as:

$$S[P] = - \sum_x P(x) \ln P(x). \quad (226)$$

It quantifies the randomness of the distribution, and is maximum for a uniform distribution (and zero for perfectly peaked one). For this reason it is often used as a measure of diversity. Unlike all other diversity measures, it is additive, meaning that the entropy of a joint distribution of two independent variables ($P(x), Q(y)$) is given by $S[P] + S[Q]$.

The entropy is the building block for several information-theoretic measures. The **Kullback-Leibler divergence** or relative entropy, defined as

$$\text{KL}(P||Q) = \sum_x P(x) \ln \frac{P(x)}{Q(x)}, \quad (227)$$

is used as a distance measure between probability distributions (although not a metric in the mathematical sense because of its asymmetry). Another popular measure is the **mutual information** between two variables x and y :

$$I(x, y) = S[P(x)] - S[P(x|y)] = \sum_{x,y} P(x, y) \ln \frac{P(x, y)}{P(x)P(y)}. \quad (228)$$

It quantifies how much the knowledge of y reduces the entropy of x , and vice-versa since it is symmetric in x and y . It is often used as a non-parametric measure of correlations between two variables.

6. Maximum entropy models

When the underlying mechanisms that give rise to the data are not known, it can be useful to define phenomenological models based on the observables that are deemed important. A convenient way to do this is to infer maximum entropy models [321], which are probabilistic models $P(x)$ of maximum entropy $S[P]$ subject to the constraint that they agree with the data on a choice of key mean observables $\mathcal{O}_a(x)$: $\sum_x P(x)\mathcal{O}_a(x) = (1/N) \sum_{i=1}^N \mathcal{O}_a(x_i)$. It can be shown that the distribution takes an exponential form:

$$P(x) = \frac{1}{Z} \exp \left(\sum_a \theta_a \mathcal{O}_a(x) \right), \quad (229)$$

where θ_a are Lagrange multipliers that must be adjusted to satisfy the constraints, and Z is a normalization constant. It turns out that maximizing the likelihood of a dataset (x_1, \dots, x_N) under Eq. 229 over the model parameters (θ_a) is equivalent to satisfying the constraint over the mean observables. The inference procedure is computationally hard, and is usually performed using a combination of Monte Carlo sampling methods and gradient descent, or mean-field techniques [322]. Often, the observables are taken to be marginals of variables of interest, as well as pairwise of higher-order correlations between them. The resulting models then fall into classes of inverse statistical physics models, such as disordered Ising or Potts models [323].

7. Machine learning and Neural networks

Modern machine learning techniques are increasingly popular in computational immunology, and are anticipated to become even more popular in the near future. We refer the reader to Ref. [324] for a review of machine learning methods aimed at physicists.

The best established use of deep neural networks in immunology is for predicting peptide-MHC binding through a set of tools called netMHC [40, 164]. More traditional machine learning approaches have been used to compare and characterize immune repertoires [211, 212].

Current efforts aim at predicting TCR-antigen binding using machine learning techniques [24–26], although the amount of data necessary to obtain truly predictive models is probably still insufficient at this stage.

F. High-throughput repertoire sequencing

The variable region of an antigen receptor chain is about 400 bp long. The sequencing challenge is to capture the whole region in one read. Modern methods of high-throughput genomic or metagenomic sequencing use shotgun sequencing which breaks the genome of interest into short fragments and after sequencing pastes the reads together by putting together overlapping reads with the help of a reference guiding template. Since Repertoire Sequencing (RepSeq) focusses on highly variable regions which includes non-templated insertions and deletions, the sequence must be acquired in one read. Read lengths required for TCR sequences are typically shorter ($\sim 100 - 150$ bp) since short fragments of V and J genes are enough to distinguish different V genes with the help of known genomic templates. However, since B-cells carry a large number of hypermutations in the V gene, sequencing of longer fragments (~ 600 bp) that encompass explicitly most of the V gene is necessary.

Different methods have been developed for sequencing DNA and mRNA. In general, the protocols start with isolating the mRNA or DNA of the cells of interest (BCR, TCR, in subsets of cells of interest sorted using

FACS). The mRNA product is then reverse transcribed onto cDNA and barcoded, sometimes with unique molecular identifier (UMI), and amplified by Polymerase Chain Reaction (PCR) before sequencing. Different techniques exist for PCR amplification that either add primer specific sites or multiple primers. Rapid amplification of cDNA ends (RACE) is one of the most common techniques, which is based on adding a linker with a primer binding site to a conserved region, on the 5' end of the V gene where no constant template is available (unlike on the 3' end). In DNA sequencing multiple primers that target specific regions of the DNA, introducing a primer specific amplification bias that can be controlled for by spiking in known sequences.

Protocols based on sequencing directly DNA have the advantage that the experiments is free of mRNA expression bias. However, currently mRNA based technologies are able to reliably report sequence counts thanks to the use of UMI barcoding techniques that are still being developed for DNA sequencing. The difficulty with using UMIs in the DNA protocol lies in introducing the barcode before the original sequence is amplified. In the mRNA protocols the barcode is introduced during the initial reverse transcription of mRNA into cDNA for every RNA molecule. The barcoded sequence is then purified and the whole product is amplified by PCR. For DNA an additional ligation step is needed, in which the DNA is cut close to the region of interest (keeping in mind that the sequenced lengths is short). Ligation is not an efficient reaction. This approach is simple when many gene copies are available in a sample, but this is not the case in immune repertoire sequencing. Alternatively, the barcode could be added in the first PCR cycle. The product then needs to be well purified making sure no barcode carrying primers are left, at the same time making sure that the barcoded sequences is not lost due to its low fraction compared to non-barcoded sequences, which remains technically tricky.

The above procedures sequence the chains of TCR or BCR repertoire in bulk, leaving no possibility to figure out the pairing of alpha and beta, or light and heavy chains in the same cell. Naturally, this limits the discussions of repertoire diversity, and more importantly attempts to link sequence to a functional phenotype. Recently these limitations have been overcome by the development of single cell reseq sequencing, where the PCR reaction is performed either in wells [325] or in droplets [326–328]. Alternatively, a computational ap-

proach that exploits the statistics of co-occurrence of the rare events partitioning the chains separately into wells is used to identify pairs of TCR sequences that come from the same cell from bulk DNA sequencing experiments [329]. The results of these analysis show that the pairing of alpha-beta chains is largely independent for TCR [145, 328], so the results of bulk analysis for diversity hold. However the affinity and functional properties of the BCR and TCR need to be estimated at the level of the whole receptor.

Lastly, novel high-throughput single-cell barcoded mRNA technologies [223] are promising to change the field of immune repertoire sequencing, providing paired chain reads in large numbers. These technologies are very similar to those used in 5'RACE with UMI. Droplet based platforms allow in principle for large numbers of cells (up to 80 000 per chip) to be analyzed at the same time, although these numbers are still much smaller than the 10^6 cells in bulk experiments. Given that clones have very low frequencies in bulk experiments (possibly down to a single cell out of 10^{11}), sampling of the order of 10^4 does not guarantee reproducible experiments. Additionally the cost remains high, at about 1 USD per cell. At the moment of writing, the first analyses based on this technology are still underway.

XIII. ACKNOWLEDGMENTS

The authors thank all past and current collaborators for the many useful discussions. TM and AMW thanks Meriem Bensouda Koraichi, Barbara Bravi, Victor Chardès, Thomas Dupic, Cosimo Lupo, Carlos Olivares, Jacopo Marchi, Maria Ruiz Ortega and Natanael Spisak for their comments on the manuscript and Mikhail Pogorelyy and Anastasia Minervina for discussions. GAB thanks Sooraj Achar, Emanuel Salazar Cavazos and Van Truong for their comments on the manuscript. This work was partially supported by the European Research Council Consolidator Grant n. 724208 [TM & AMW] and the intramural research program of the National Cancer Institute [GAB]. We also thank the KITP, Les Houches EdP and IES Cargèse institutes for hospitality during workshops on the topics reported in this review.

-
- [1] Sompayrac L (1999) *How the immune system works* (Blackwell Scientific Publications).
 - [2] Burnet FM (1957) A Modification of Jerne's Theory of Antibody Production using the Concept of Clonal Selection. *The Australian Journal of Science* 20:67–69.
 - [3] Ehrlich P (1900) On Immunity with Special Reference to Cell Life. *Roy. Soc. Proc.* 66:424–448.
 - [4] Jerne NK (1955) The natural-selection theory of antibody formation. *Proc Natl Acad Sci U S A* 41:849–857.
 - [5] Hodgkin PD, Heath WR, Baxter AG (2007) The clonal selection theory : 50 years since the revolution. *Nature immunology* 66:424–448.
 - [6] Nossal GJV, Lederberg J (2019) Antibody Production by Single Cells. 1. *The Journal of Immunology*

- 182:1231–1232;
- [7] Billingham RE, Brent L, Medawar PB (1953) Activity Acquired Tolerance ' of Foreign Cells. *The Journal of Immunology* 172:603–606.
- [8] Billingham RE, Brent L, Medawar PB (1956) Quantitative Studies on Tissue Transplantation Immunity. III. Actively Acquired Tolerance. *Philosophical Transactions of the Royal Society B: Biological Sciences* 239:378–414.
- [9] Lifschitz E, Pitaevskii L (1981) *Physical Kinetics* (Butterworth-Heinemann).
- [10] Hager GL, McNally JG, Misteli T (2009) Transcription dynamics. *Mol Cell* 35:741–53.
- [11] Halford SE, Marko JF (2004) How do site-specific dna-binding proteins find their targets? *Nucleic Acids Res* 32:3040–52.
- [12] Gorman J, Greene EC (2008) Visualizing one-dimensional diffusion of proteins along dna. *Nat Struct Mol Biol* 15:768–74.
- [13] Slutsky M, Mirny LA (2004) Kinetics of Protein-DNA Interaction : Facilitated Target Location in Sequence-Dependent Potential. *Biophysical Journal* 87:4021–4035.
- [14] Altan-Bonnet G, Germain RN (2005) Modeling t cell antigen discrimination based on feedback control of digital erk responses. *PLoS Biol* 3:e356.
- [15] Feinerman O, Germain RN, Altan-Bonnet G (2008) Quantitative challenges in understanding ligand discrimination by alphabeta t cells. *Mol Immunol* 45:619–31.
- [16] Mason D (1998) A very high level of crossreactivity is an essential feature of the T- cell receptor.
- [17] de Boer RJ, Perelson ASA (1993) How diverse should the immune system be. *Proc. R. Soc. B Biol. Sci.* 252:171.
- [18] Birnbaum ME, et al. (2014) Deconstructing the peptide-MHC specificity of t cell recognition. *Cell* 157:1073–1087.
- [19] Adams R, Kinney JB, Mora T, Walczak AM (2016) Measuring the sequence-affinity landscape of antibodies with massively parallel titration curves., Technical report.
- [20] Glanville J, et al. (2017) Identifying specificity groups in the T cell receptor repertoire. *Nature* advance on:94–98.
- [21] Dash P, et al. (2017) Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature* advance on:89–93.
- [22] Shugay M, et al. (2018) VDJdb: A curated database of T-cell receptor sequences with known antigen specificity. *Nucleic Acids Res.* 46:D419–D427.
- [23] Tickotsky N, Sagiv T, Prilusky J, Shifrut E, Friedman N (2017) McPAS-TCR: A manually curated catalogue of pathology-associated T cell receptor sequences. *Bioinformatics* 33:2924–2929.
- [24] Jurtz VI, et al. (2018) NetTCR: sequence-based prediction of TCR binding to peptide-MHC complexes using convolutional neural networks. *bioRxiv* p 433706.
- [25] Sidhom JW, Larman HB, Pardoll DM, Baras AS (2018) DeepTCR: a deep learning framework for revealing structural concepts within TCR Repertoire. *bioRxiv* p 464107.
- [26] Jokinen E, Heinonen M, Huuhtanen J, Mustjoki S, Harri L (2019) TCRGP : Determining epitope specificity of T cell receptors. pp 4–12.
- [27] Chao DL, Davenport MP, Forrest S, Perelson AS (2005) The effects of thymic selection on the range of T cell cross-reactivity. *Eur. J. Immunol.* 35:3452–3459.
- [28] Miyazawa S, Jernigan RL (1996) Residue-Residue Potentials with a Favorable Contact Pair Term and an Unfavorable High Packing Density Term, for Simulation and Threading - 1-s2.0-S002228369690114X-main.pdf. *J Mol Biol* pp 623–644.
- [29] Kosmrlj A, Jha AK, Huseby ES, Kardar M, Chakraborty AK (2008) How the thymus designs antigen-specific and self-tolerant T cell receptor sequences. *Proc. Natl. Acad. Sci. U. S. A.* 105:16671–6.
- [30] George JT, Kessler DA, Levine H (2017) Effects of thymic selection on T cell recognition of foreign and tumor antigenic peptides. *Proc. Natl. Acad. Sci.* p 201708573.
- [31] Košmrlj A, Chakraborty AK, Kardar M, Shakhnovich EI (2009) Thymic selection of T-cell receptors as an extreme value problem. *Phys. Rev. Lett.* 103:068103.
- [32] Butler TC, Kardar M, Chakraborty AK (2013) Quorum sensing allows T cells to discriminate between self and nonself. *Proc. Natl. Acad. Sci. U. S. A.* 110:11833–8.
- [33] Perelson AS, Oster GF (1979) Theoretical studies of clonal selection minimal antibody repertoire size and reliability of self non self discrimination. *J. Theor. Biol.* 81:645–670.
- [34] Detours V, Mehr R, Perelson aS (1999) A quantitative theory of affinity-driven T cell repertoire selection. *J. Theor. Biol.* 200:389–403.
- [35] Detours V, Mehr R, Perelson a (2000) Deriving quantitative constraints on T cell selection from data on the mature T cell repertoire. *J Immunol* 164:121–128.
- [36] Wang S, et al. (2015) Manipulating the selection forces during affinity maturation to generate cross-reactive HIV antibodies. *Cell* 160:785–797.
- [37] Nourmohammad A, Otwinowski J, Plotkin JB (2015) Host-pathogen co-evolution and the emergence of broadly neutralizing antibodies in chronic infections. pp 19–22.
- [38] Adams RM, Kinney JB, Walczak AM, Mora T (2019) Epistasis in a Fitness Landscape Defined by Antibody-Antigen Binding Free Energy. *Cell Syst.* pp 1–8.
- [39] Luksza M, et al. (2017) A neoantigen fitness model predicts tumour response to checkpoint blockade immunotherapy. *Nature* 551:517–520.
- [40] Andreatta M, Nielsen M (2015) Gapped sequence alignment using artificial neural networks: Application to the MHC class i system. *Bioinformatics* 32:511–517.
- [41] McKeithan TW (1995) Kinetic proofreading in T-cell receptor signal transduction. *Proceedings of the National Academy of Sciences* 92:5042–5046.
- [42] Hopfield JJ (1974) Kinetic proofreading: a new mechanism for reducing errors in biosynthetic processes requiring high specificity. *Proc Natl Acad Sci U S A* 71:4135–9.
- [43] Ninio J (1975) Kinetic amplification of enzyme discrimination. *Biochimie* 57:587–95.
- [44] François P, Voisinne G, Siggia ED, Altan-bonnet G, Vergassola M (2013) Phenotypic model for early T-cell activation displaying sensitivity , specificity , and antagonism. *Proc Natl Acad Sci U S A* pp E888–E897.
- [45] Germain RN (2005) Modeling T Cell Antigen Discrimination Based on Feedback Control of Digital ERK Responses. *PLoS biology* 3.

- [46] Sykulev Y, Joo M, Vturina I, Tsomides TJ, Eisen HN (1996) Evidence that a single peptide-mhc complex on a target cell can elicit a cytolytic t cell response. *Immunity* 4:565–571.
- [47] Irvine DJ, Purbhoo MA, Krogsgaard M, Davis MM (2002) Direct observation of ligand recognition by t cells. *Nature* 419:845–849.
- [48] Stefanova I, et al. (2003) Tcr ligand discrimination is enforced by competing erk positive and shp-1 negative feedback pathways. *Nat Immunol* 4:248–54.
- [49] Lalanne JB, François P (2013) Principles of adaptive sorting revealed by in silico evolution. *Phys. Rev. Lett.* 110:218102.
- [50] Qi SY, Groves JT, Chakraborty AK (2001) Synaptic pattern formation during cellular recognition. *Proc Natl Acad Sci U S A* 98:6548–53.
- [51] Huse M, Lillemeier BF, Kuhns MS, Chen DS, Davis MM (2006) T cells use two directionally distinct pathways for cytokine secretion. *Nat Immunol* 7:247–55.
- [52] Choudhuri K, Wiseman D, Brown MH, Gould K, van der Merwe PA (2005) T-cell receptor triggering is critically dependent on the dimensions of its peptide-mhc ligand. *Nature* 436:578–82.
- [53] Choudhuri K, van der Merwe PA (2007) Molecular mechanisms involved in t cell receptor triggering. *Semin Immunol* 19:255–61.
- [54] Valitutti S, Dessing M, Aktories K, Gallati H, Lanzavecchia A (1995) Sustained signaling leading to t cell activation results from prolonged t cell receptor occupancy. role of t cell actin cytoskeleton. *J Exp Med* 181:577–84.
- [55] Kalergis AM, et al. (2001) Efficient t cell activation requires an optimal dwell-time of interaction between the tcr and the pmhc complex. *Nat Immunol* 2:229–34.
- [56] Holler PD, et al. (2000) In vitro evolution of a t cell receptor with high affinity for peptide/mhc. *Proc Natl Acad Sci U S A* 97:5387–92.
- [57] Lever M, et al. (2016) Architecture of a minimal signaling pathway explains the t-cell response to a 1 million-fold variation in antigen affinity and dose. *Proc Natl Acad Sci U S A* 113:E6630–E6638.
- [58] Posey, A. D. J, et al. (2016) Engineered car t cells targeting the cancer-associated tn-glycoform of the membrane mucin mucl control adenocarcinoma. *Immunity* 44:1444–54.
- [59] Schmitt TM, et al. (2017) Generation of higher affinity t cell receptors by antigen-driven differentiation of progenitor t cells in vitro. *Nat Biotechnol* 35:1188–1195.
- [60] Valitutti S, Muller S, Cella M, Padovan E, Lanzavecchia A (1995) Serial triggering of many t-cell receptors by a few peptide-mhc complexes. *Nature* 375:148–51.
- [61] Zhu C, Jiang N, Huang J, Zarnitsyna VI, Evavold BD (2013) Insights from in situ analysis of tcr-pmhc recognition: response of an interaction network. *Immunol Rev* 251:49–64.
- [62] Liu B, Chen W, Evavold BD, Zhu C (2014) Accumulation of dynamic catch bonds between tcr and agonist peptide-mhc triggers t cell signaling. *Cell* 157:357–368.
- [63] Sibener LV, et al. (2018) Isolation of a structural mechanism for uncoupling t cell receptor signaling from peptide-mhc binding. *Cell* 174:672–687 e27.
- [64] Wu P, et al. (2019) Mechano-regulation of peptide-mhc class i conformations determines tcr antigen recognition. *Mol Cell* 73:1015–1027 e7.
- [65] Cai E, et al. (2017) Visualizing dynamic microvillar search and stabilization during ligand detection by t cells. *Science* 356.
- [66] Dustin ML, Chakraborty AK, Shaw AS (2010) Understanding the structure and function of the immunological synapse. *Cold Spring Harb Perspect Biol* 2:a002311.
- [67] Tsourkas PK, Liu W, Das SC, Pierce SK, Raychaudhuri S (2012) Discrimination of membrane antigen affinity by b cells requires dominance of kinetic proofreading over serial engagement. *Cell Mol Immunol* 9:62–74.
- [68] Blinov ML, Faeder JR, Goldstein B, Hlavacek WS (2006) A network model of early events in epidermal growth factor receptor signaling that accounts for combinatorial complexity. *Biosystems* 83:136–51.
- [69] Lemmon MA, Schlessinger J (2010) Cell signaling by receptor tyrosine kinases. *Cell* 141:1117–34.
- [70] Yang J, Reth M (2010) The dissociation activation model of b cell antigen receptor triggering. *FEBS Lett* 584:4872–7.
- [71] Fleire SJ, et al. (2006) B cell ligand discrimination through a spreading and contraction response. *Science* 312:738–41.
- [72] François P, Hakim V (2004) Design of genetic networks with specified functions by evolution in silico. *Proc. Natl. Acad. Sci. U. S. A.* 101:580–585.
- [73] Proulx-Giraldeau F, Rademaker TJ, François P (2017) Untangling the Hairball: Fitness-Based Asymptotic Reduction of Biological Networks. *Biophys. J.* 113:1893–1906.
- [74] Lipniacki T, Hat B, Faeder JR, Hlavacek WS (2008) Stochastic effects and bistability in T cell receptor signaling. *J. Theor. Biol.* 254:110–122.
- [75] Altan-Bonnet G, Mukherjee R (2019) Cytokine-mediated communication: a quantitative appraisal of immune complexity. *Nat Rev Immunol* 19:205–217.
- [76] Vogel RM, Erez A, Altan-Bonnet G (2016) Dichotomy of cellular inhibition by small-molecule inhibitors revealed by single-cell analysis. *Nat Commun* 7:12428.
- [77] Feinerman O, et al. (2010) Single-cell quantification of il-2 response by effector and regulatory t cells reveals critical plasticity in immune response. *Mol Syst Biol* 6:437.
- [78] Cotari JW, Voisinne G, Dar OE, Karabacak V, Altan-Bonnet G (2013) Cell-to-cell variability analysis dissects the plasticity of signaling of common γ chain cytokines in T cells. *Sci. Signal.* 6:ra17.
- [79] Busse D, et al. (2010) Competing feedback loops shape il-2 signaling between helper and regulatory t lymphocytes in cellular microenvironments. *Proc Natl Acad Sci U S A* 107:3058–63.
- [80] Hofer T, Krichevsky O, Altan-Bonnet G (2012) Competition for il-2 between regulatory and effector t cells to chisel immune responses. *Front Immunol* 3:268.
- [81] Levine SJ (2004) Mechanisms of soluble cytokine receptor generation. *J Immunol* 173:5343–8.
- [82] Rose-John S, Neurath MF (2004) Il-6 trans-signaling: the heat is on. *Immunity* 20:2–4.
- [83] Cendrowski J, Maminska A, Miaczynska M (2016) Endocytic regulation of cytokine receptor signaling. *Cytokine Growth Factor Rev* 32:63–73.
- [84] Tkach KE, et al. (2014) T cells translate individual, quantal activation into collective, analog cytokine responses via time-integrated feedbacks. *Elife* 3:e01944.

- [85] Voisinne G, et al. (2015) T cells integrate local and global cues to discriminate between structurally similar antigens. *Cell Rep* 11:1208–19.
- [86] Polonsky M, et al. (2018) Induction of CD4 T cell memory by local cellular collectivity. *Science* 360.
- [87] Becker V, et al. (2010) Covering a broad dynamic range: information processing at the erythropoietin receptor. *Science* 328:1404–8.
- [88] Raue A, Kreutz C, Maiwald T, Klingmüller U, Timmer J (2011) Addressing parameter identifiability by model-based experimentation. *IET Syst Biol* 5:120–30.
- [89] Bachmann J, et al. (2011) Division of labor by dual feedback regulators controls jak2/stat5 signaling over broad ligand range. *Mol Syst Biol* 7:516.
- [90] Karr JR, et al. (2015) Summary of the dream8 parameter estimation challenge: Toward parameter identification for whole-cell models. *PLoS Comput Biol* 11:e1004096.
- [91] Shi T, et al. (2016) Conservation of protein abundance patterns reveals the regulatory architecture of the egfr-mapk pathway. *Sci Signal* 9:rs6.
- [92] Mitchell S, Roy K, Zangle TA, Hoffmann A (2018) Non-genetic origins of cell-to-cell variability in b lymphocyte proliferation. *Proc Natl Acad Sci U S A* 115:E2888–E2897.
- [93] Thurley K, Gerecht D, Friedmann E, Hofer T (2015) Three-dimensional gradients of cytokine signaling between t cells. *PLoS Comput Biol* 11:e1004206.
- [94] Berezhkovskii AM, Sample C, Shvartsman SY (2010) How long does it take to establish a morphogen gradient? *Biophys J* 99:L59–61.
- [95] Kolomeisky A (2011) Formation of a morphogen gradient: Acceleration by degradation. *J. Phys. Chem. Lett.* 2:1502–1505.
- [96] Oyler-Yaniv A, et al. (2017) A Tunable Diffusion-Consumption Mechanism of Cytokine Propagation Enables Plasticity in Cell-to-Cell Communication in the Immune System. *Immunity* 46:609–620.
- [97] Marcou Q, et al. (2018) A model for the integration of conflicting exogenous and endogenous signals by dendritic cells. *Phys. Biol.* 15.
- [98] Huang J, et al. (2013) A single peptide-major histocompatibility complex ligand triggers digital cytokine secretion in cd4(+) t cells. *Immunity* 39:846–57.
- [99] Zhu J, Yamane H, Paul WE (2010) Differentiation of effector cd4 t cell populations (*). *Annu Rev Immunol* 28:445–89.
- [100] Chaouat G, et al. (2004) Th1/th2 paradigm in pregnancy: paradigm lost? cytokines in pregnancy/early abortion: reexamining the th1/th2 paradigm. *Int Arch Allergy Immunol* 134:93–119.
- [101] Yates A, Callard R, Stark J (2004) Combining cytokine signalling with t-bet and gata-3 regulation in th1 and th2 differentiation: a model for cellular decision-making. *J Theor Biol* 231:181–96.
- [102] Antebi YE, et al. (2013) Mapping differentiation under mixed culture conditions reveals a tunable continuum of t cell fates. *PLoS Biol* 11:e1001616.
- [103] Hong T, Xing J, Li L, Tyson JJ (2011) A mathematical model for the reciprocal differentiation of t helper 17 cells and induced regulatory t cells. *PLoS Comput Biol* 7:e1002122.
- [104] Laslo P, et al. (2006) Multilineage transcriptional priming and determination of alternate hematopoietic cell fates. *Cell* 126:755–66.
- [105] Kueh HY, Champhekar A, Nutt SL, Elowitz MB, Rothenberg EV (2013) Positive feedback between pu.1 and the cell cycle controls myeloid differentiation. *Science* 341:670–3.
- [106] Walczak AM, Sasai M, Wolynes PG (2005) Self-consistent proteomic field theory of stochastic gene switches. *Biophysical journal* 88:828–50.
- [107] Walczak AM, Onuchic JN, Wolynes PG (2005) Absolute rate theories of epigenetic stability. *Proceedings of the National Academy of Sciences of the United States of America* 102:18926–31.
- [108] Friedman N, Cai L, Xie XS (2006) Linking stochastic dynamics to population distribution: an analytical framework of gene expression. *Phys Rev Lett* 97:168302.
- [109] Elowitz MB, Levine AJ, Siggia ED, Swain PS (2002) Stochastic gene expression in a single cell. *Science* 297:1183–6.
- [110] Feinerman O, Veiga J, Dorfman JR, Germain RN, Altan-Bonnet G (2008) Variability and robustness in t cell activation from regulated heterogeneity in protein levels. *Science* 321:1081–4.
- [111] Szabo SJ, et al. (2000) A novel transcription factor, t-bet, directs th1 lineage commitment. *Cell* 100:655–69.
- [112] Palma A, Jarrar AS, Tieri P, Cesareni G, Castiglione F (2018) Gene regulatory network modeling of macrophage differentiation corroborates the continuum hypothesis of polarization states. *Front Physiol* 9:1659.
- [113] Kueh HY, et al. (2016) Asynchronous combinatorial action of four regulatory factors activates bcl11b for t cell commitment. *Nat Immunol* 17:956–65.
- [114] Peine M, et al. (2013) Stable t-bet(+)/gata-3(+) th1/th2 hybrid cells arise in vivo, can develop directly from naive precursors, and limit immunopathologic inflammation. *PLoS Biol* 11:e1001633.
- [115] Moran U, Phillips R, Milo R (2010) SnapShot: key numbers in biology. *Cell* 141:1262–1262.e1.
- [116] Milo R, Phillips R (2015) *Cell Biology by the Numbers* (Garland Science), p 2015.
- [117] Busch K, et al. (2015) Fundamental properties of unperturbed haematopoiesis from stem cells in vivo. *Nature* 518:542.
- [118] Höfer T, Barile M, Flossdorf M (2016) Stem-cell dynamics and lineage topology from in vivo fate mapping in the hematopoietic system. *Current opinion in biotechnology* 39:150–156.
- [119] Zilman A, Ganusov VV, Perelson AS (2010) Stochastic models of lymphocyte proliferation and death. *PLoS One* 5:1–14.
- [120] Perelson AS (2002) Modelling Viral and Immune System Dynamics. *Nat. Rev. Immunol.* 2:28–36.
- [121] Yates A, et al. (2007) Reconstruction of cell population dynamics using CFSE. *BMC Bioinformatics* 20:1–20.
- [122] Schoedel KB, et al. (2016) The bulk of the hematopoietic stem cell population is dispensable for murine steady-state and stress hematopoiesis. *Blood* 128:2285–2297.
- [123] Perie L, et al. (2014) Determining Lineage Pathways from Cellular Barcoding Experiments. *Cell Reports* 6:617–624.
- [124] Ogawa M, Porter PN, Nakahata T (1983) Renewal And Commitment To Differentiation Of Hemopoietic Stem Cells (An Interpretive Review). *Blood* 5:823–829.

- [125] Tsuji K, Nakahata T (1989) Stochastic Model for Multipotent Hemopoietic Progenitor Differentiation. *Journal of cellular physiology* 653:647–653.
- [126] Pei W, et al. (2017) Polylox barcoding reveals haematopoietic stem cell fates realized in vivo. *Nature Publishing Group* 548:456–460.
- [127] Laffleur B, et al. (2014) Immunoglobulin genes undergo legitimate repair in human B cells not only after cis - but also frequent trans -class switch recombination. *Genes and Immunity* pp 341–346.
- [128] Murphy K, Travers P, Walport M (2007) *Janeway's Immunology* (Garland Science), 7 edition edition.
- [129] Murphy KM, et al. (2000) Signaling AND Transcription in T Helper development. *Annu. Rev. Immunol.* 18:451–494.
- [130] Hawkins ED, Turner ML, Dowling MR, Gend CV, Hodgkin PD (2007) A model of immune regulation as a consequence of randomized lymphocyte division and death times. *Proceedings of the National Academy of Sciences* 104:5032–5037.
- [131] Marchingo JM, et al. (2014) Antigen affinity, costimulation, and cytokine inputs sum linearly to. *Science* 346:1123–1128.
- [132] Marchingo JM, et al. (2016) T-cell stimuli independently sum to regulate an inherited clonal division fate. *Nature Communications* 7:1–12.
- [133] Duffy KR, Hodgkin PD (2012) Intracellular competition for fates in the immune system. *Trends in Cell Biology* 22:457–464.
- [134] Deenick EK, Gett AV, Hodgkin PD (2003) Stochastic Model of T Cell Proliferation: A Calculus Revealing IL-2 Regulation of Precursor Frequencies, Cell Cycle Time, and Survival. *The Journal of Immunology* 170:4963–4972.
- [135] Tangye SG, Avery DT, Deenick EK, Hodgkin PD (2003) Intrinsic Differences in the Proliferation of Naive and Memory Human B Cells as a Mechanism for Enhanced Secondary Immune Responses. *The Journal of Immunology* 170:686–694.
- [136] Bendall SC, Nolan GP, Roederer M, Chattopadhyay PK (2012) A deep profiler 's guide to cytometry. *Trends in Immunology* 33:323–332.
- [137] Buchholz VR, et al. (2013) Disparate Individual Fates Compose Robust CD8+ T Cell Immunity. *Science* 340:630–636.
- [138] Flossdorf M, Rössler J, Buchholz VR, Busch DH, Höfer T (2015) CD8 + T cell diversification by asymmetric cell division. *Nature immunology* 16:891–893.
- [139] Dowling MR, et al. (2014) Stretched cell cycle model for proliferating lymphocytes. *Proc Natl Acad Sci U S A* 111:6377–6382.
- [140] Miles AS, Hodgkin PD, Duffy KR (2018) Inferring differentiation order in adaptive immune responses from population level data. *Springer Volume on T-cell immunology*.
- [141] Kinjyo I, et al. (2015) Real-time tracking of cell cycle progression during CD8+ effector and memory T-cell differentiation. *Nature Communications* 6:1–13.
- [142] Jenkins MK, Chu HH, McLachlan JB, Moon JJ (2009) On the composition of the preimmune repertoire of T cells specific for Peptide-major histocompatibility complex ligands. *Annu. Rev. Immunol.* 28:275–294.
- [143] Glanville J, et al. (2009) Precise determination of the diversity of a combinatorial antibody library gives insight into the human immunoglobulin repertoire. *Proc. Natl. Acad. Sci.* 106:20216–20221.
- [144] Han A, Glanville J, Hansmann L, Davis MM (2014) Linking T-cell receptor sequence to functional phenotype at the single-cell level. *Nat. Biotechnol.* 32:684–692.
- [145] Dupic T, Marcou Q, Walczak AM, Mora T (2018) Genesis of the $\alpha\beta$ T-cell receptor. *arXiv:1806.11030*.
- [146] Arstila TP, et al. (1999) A direct estimate of the human alphabeta T cell receptor diversity. *Science* 286:958–961.
- [147] Weinstein JA, Jiang N, White RA, Fisher DS, Quake SR (2009) High-throughput sequencing of the zebrafish antibody repertoire. *Science (80-.)*. 324:807–810.
- [148] Robins HS, et al. (2009) Comprehensive assessment of T-cell receptor beta-chain diversity in alphabeta T cells. *Blood* 114:4099–4107.
- [149] Boyd SD, et al. (2009) Measurement and clinical monitoring of human lymphocyte clonality by massively parallel {VDJ} pyrosequencing. *Sci Transl Med* 1:12ra23.
- [150] Benichou J, Ben-Hamo R, Louzoun Y, Efroni S (2012) Rep-Seq: Uncovering the immunological repertoire through next-generation sequencing.
- [151] Six A, et al. (2013) The past, present and future of immune repertoire biology - the rise of next-generation repertoire analysis. *Front. Immunol.* 4:413.
- [152] Robins H (2013) Immunosequencing: applications of immune repertoire deep sequencing. *Curr. Opin. Immunol.* 25:646–652.
- [153] Georgiou G, et al. (2014) The promise and challenge of high-throughput sequencing of the antibody repertoire. *Nat. Biotechnol.* 32:158–68.
- [154] Heather JM, Ismail M, Oakes T, Chain B (2017) High-throughput sequencing of the T-cell receptor repertoire: pitfalls and opportunities. *Brief. Bioinform.* p bbw138.
- [155] Minervina A, Pogorely M, Mamedov I (2019) TCR and BCR repertoire profiling in adaptive immunity. *Transpl. Int.* pp 0–2.
- [156] Rubelt F, et al. (2017) Adaptive Immune Receptor Repertoire Community recommendations for sharing immune-repertoire sequencing data. *Nature immunology* 18:1274–1278.
- [157] Fisher R, Steven Corbet A, Williams C (2016) The Relation Between the Number of Species and the Number of Individuals in a Random Sample of an Animal Population. *J. Anim. Ecol.* 12:42–58.
- [158] Qi Q, et al. (2014) Diversity and clonal selection in the human T-cell repertoire. *Proc. Natl. Acad. Sci. U. S. A.* 111:13139–44.
- [159] Chao A, Bunge J (2002) Estimating the number of species in a stochastic abundance model. *Biometrics* 58:531–539.
- [160] DeWitt WS, et al. (2016) A Public Database of Memory and Naive B-Cell Receptor Sequences. *PLoS One* 11:e0160853.
- [161] Lythe G, Callard RE, Hoare R, Molina-París C (2015) How many TCR clonotypes does a body maintain? *J. Theor. Biol.* 389:214–224.
- [162] Mora T, Walczak A (2018) in *Syst. Immunol.*, eds Das JD, Jayaprakash C (CRC Press), pp 185–199.
- [163] Wardemann H, et al. (2003) Predominant Autoantibody Production by Early Human B Cell Precursors. *Science (80-.)*. 301:1374–1377.

- [164] Lauemøller SL, et al. (2003) Sensitive quantitative predictions of peptide-MHC binding by a Query by Committee ' artificial neural network approach. *Tissue Antigens* 62:378–384.
- [165] Jensen KK, et al. (2018) Improved methods for predicting peptide binding affinity to MHC class II molecules. *Immunology* 154:394–406.
- [166] Moon JJ, et al. (2007) Naive CD4+ T Cell Frequency Varies for Different Epitopes and Predicts Repertoire Diversity and Response Magnitude. *Immunity* 27:203–213.
- [167] Jenkins MK, Moon JJ (2012) The role of naive {T} cell precursor frequency and recruitment in dictating immune response magnitude. *J. Immunol.* 188:4135–4140.
- [168] Yates AJ (2014) Theories and quantification of thymic selection. *Front. Immunol.* 5:13.
- [169] Murugan A, Mora T, Walczak AM, Callan CG (2012) Statistical inference of the generation probability of T-cell receptors from sequence repertoires. *Proc. Natl. Acad. Sci.* 109:16161–16166.
- [170] Elhanati Y, et al. (2015) Inferring processes underlying B-cell repertoire diversity. *Philos Trans R Soc Lond, B, Biol Sci* 370:20140243.
- [171] Pogorelyy MV, et al. (2017) Method for identification of condition-associated public antigen receptor sequences. *arXiv:1709.09703*.
- [172] Toledano A, et al. (2018) Evidence for shaping of light chain repertoire by structural selection. *Front. Immunol.* 9:1–9.
- [173] Sethna Z, et al. (2017) Insights into immune system development and function from mouse T-cell repertoires. *Proc. Natl. Acad. Sci.* 114:2253–2258.
- [174] Magadan S, et al. (2018) Origin of Public Memory B Cell Clones in Fish After Antiviral Vaccination. *Front. Immunol.* 9:2115.
- [175] Marcou Q, Mora T, Walczak AM (2018) High-throughput immune repertoire analysis with IGoR. *Nat. Commun.* 9:561.
- [176] Sethna Z, Elhanati Y, Callan CG, Walczak AM, Mora T (2019) OLGA: fast computation of generation probabilities of B- and T-cell receptor amino acid sequences and motifs. *Bioinformatics* btz035.
- [177] Munshaw S, Kepler TB (2010) SoDA2: a Hidden Markov Model approach for identification of immunoglobulin rearrangements. *Bioinformatics* 26:867–72.
- [178] Elhanati Y, Marcou Q, Mora T, Walczak AM (2015) repgenHMM: a dynamic programming tool to infer the rules of immune receptor generation from sequence data. *bioRxiv* In press:1–9.
- [179] Ralph DK, Matsen FA (2016) Consistency of VDJ Rearrangement and Substitution Parameters Enables Accurate B Cell Receptor Sequence Annotation. *PLoS Comput. Biol.* 12:1–25.
- [180] Ralph DK, Matsen FA (2016) Likelihood-Based Inference of B Cell Clonal Families. *PLoS Comput. Biol.* 12:1–28.
- [181] Wang Y, et al. (2011) Genomic screening by 454 pyrosequencing identifies a new human IGHV gene and sixteen other new IGHV allelic variants. *Immunogenetics* 63:259–265.
- [182] Pogorelyy MV, et al. (2018) Method for identification of condition-associated public antigen receptor sequences. *Elife* 7:1–13.
- [183] Pogorelyy MV, et al. (2018) Detecting T-cell receptors involved in immune responses from single repertoire snapshots. *arXiv:1807.08833*.
- [184] Wing K, Sakaguchi S (2010) Regulatory T cells exert checks and balances on self tolerance and autoimmunity. *Nat. Immunol.* 11:7–13.
- [185] Bains I, van Santen HM, Seddon B, Yates AJ (2013) Models of Self-Peptide Sampling by Developing T Cells Identify Candidate Mechanisms of Thymic Selection. *PLoS Comput. Biol.* 9.
- [186] Le Borgne M, et al. (2009) The impact of negative selection on thymocyte migration in the medulla. *Nat. Immunol.* 10:823–830.
- [187] Kosmrlj A, et al. (2010) Effects of thymic selection of the T-cell repertoire on HLA class I-associated control of HIV infection. *Nature* 465:350–354.
- [188] Elhanati Y, Murugan A, Callan CG, Mora T, Walczak AM (2014) Quantifying selection in immune receptor repertoires. *Proc. Natl. Acad. Sci.* 111:9875–9880.
- [189] Kaplinsky J, et al. (2014) Antibody repertoire deep sequencing reveals antigen-independent selection in maturing B cells. *Proc. Natl. Acad. Sci.* 111:E2622–9.
- [190] Elhanati Y, Sethna Z, Callan CG, Mora T, Walczak AM (2018) Predicting the spectrum of TCR repertoire sharing with a data-driven model of recombination. *Immunol Rev* 284:167–179.
- [191] Mora T, Walczak AM (2016) Renyi entropy, abundance distribution and the equivalence of ensembles. *Phys Rev E* 95:052418.
- [192] Mora T, Walczak AM, Bialek W, Callan CG (2010) Maximum entropy models for antibody diversity. *Proc. Natl. Acad. Sci.* 107:5405–5410.
- [193] Desponds J, Mora T, Aleksandra W (2016) Fluctuating fitness shapes the clone size distribution of immune repertoires. *Proc Natl Acad Sci USA* 113:274.
- [194] Desponds J, Mayer A, Mora T, Walczak AM (2017) Population dynamics of immune repertoires. pp 1–9.
- [195] Greef PCD, et al. (2019) The naive T-cell receptor repertoire has an extremely broad distribution of clone sizes. *bioRxiv*.
- [196] Best K, Oakes T, Heather JM, Shawe-Taylor J, Chain B (2015) Computational analysis of stochastic heterogeneity in PCR amplification efficiency revealed by single molecule barcoding. *Sci. Rep.* 5:14629.
- [197] Vollmers C, Sit RV, Weinstein Ja, Dekker CL, Quake SR (2013) Genetic measurement of memory B-cell recall using antibody repertoire sequencing. *Proc. Natl. Acad. Sci.* 110:13463–8.
- [198] Shugay M, et al. (2014) Towards error-free profiling of immune repertoires. *Nat. Methods* 11:653–5.
- [199] Kaplinsky J, Arnaout R (2016) Robust estimates of overall immune-repertoire diversity from high-throughput measurements on samples. *Nat. Commun.* 7:11881.
- [200] Laydon DJ, Bangham CRM, Asquith B, Crm B (2015) Estimating T-cell repertoire diversity: limitations of classical estimators and a new approach. *Philos Trans R Soc Lond, B, Biol Sci* 370:20140291.
- [201] Haegeman B, et al. (2013) Robust estimation of microbial diversity in theory and in practice. *ISME J.* 7:1092–1101.
- [202] Emerson RO, et al. (2017) Immunosequencing identifies signatures of cytomegalovirus exposure history and

- HLA-mediated effects on the T cell repertoire. *Nat. Genet.* 49:659–665.
- [203] Mayer A, Balasubramanian V, Mora T, Walczak AM (2015) How a well-adapted immune system is organized. *Proceedings of the National Academy of Sciences USA* pp 1–15.
- [204] Emerson RO, et al. (2017) Immunosequencing identifies signatures of cytomegalovirus exposure history and HLA-mediated effects on the T cell repertoire. *Nat. Genet.* 49:659–665.
- [205] Faham M, et al. (2017) Discovery of T Cell Receptor β Motifs Specific to HLAB27Positive Ankylosing Spondylitis by Deep Repertoire Sequence Analysis. *Arthritis Rheumatol.* 69:774–784.
- [206] Venturi V, et al. (2006) Sharing of T cell receptors in antigen-specific responses is driven by convergent recombination. *Proc. Natl. Acad. Sci.* 103:18691–18696.
- [207] Venturi V, et al. (2011) A mechanism for TCR sharing between T cell subsets and individuals revealed by pyrosequencing. *J. Immunol.* 186:4285–4294.
- [208] Madi A, et al. (2014) T-cell receptor repertoires share a restricted set of public and abundant CDR3 sequences that are associated with self-related immunity. *Genome Res.* 24:1603–12.
- [209] Perelson A, Weisbuch G (1997) Immunology for physicists. *Reviews of Modern Physics* 69:1219–1268.
- [210] Mayer A, Balasubramanian V, Mora T, Walczak AM (2015) How a well-adapted immune system is organized. *Proc. Natl. Acad. Sci.* 112:5950–5955.
- [211] Thomas N, et al. (2014) Tracking global changes induced in the CD4 T-cell receptor repertoire by immunization with a complex antigen using short stretches of CDR3 protein sequence. *Bioinformatics* 30:3181–3188.
- [212] Cinelli M, et al. (2017) Feature selection using a one dimensional naïve Bayes’ classifier increases the accuracy of support vector machine classification of CDR3 repertoires. *Bioinformatics* 33:btw771.
- [213] Castro R, et al. (2013) Teleost Fish Mount Complex Clonal IgM and IgT Responses in Spleen upon Systemic Viral Infection. *PLoS Pathog.* 9.
- [214] Laserson U, et al. (2014) High-resolution antibody dynamics of vaccine-induced immune responses. *Proc. Natl. Acad. Sci.* 111:4928–4933.
- [215] Wang C, et al. (2015) B-cell repertoire responses to varicella-zoster vaccination in human identical twins. *Proc. Natl. Acad. Sci.* 112:500–505.
- [216] Qi Q, et al. (2016) Diversification of the antigen-specific T cell receptor repertoire after varicella zoster vaccination. *Sci. Transl. Med.* 8:332ra46–332ra46.
- [217] DeWitt WS, et al. (2015) Dynamics of the Cytotoxic T Cell Response to a Model of Acute Viral Infection. *J. Virol.* 249:JVI.03474–14.
- [218] Pogorelyy MV, et al. (2018) Precise tracking of vaccine-responding T-cell clones reveals convergent and personalized response in identical twins. *arXiv:1804.04485*.
- [219] Borghans JAM, De Boer RJ (2007) Quantification of T-cell dynamics: From telomeres to DNA labeling. *Immunol. Rev.* 216:35–47.
- [220] De Boer RJ, Perelson AS (2013) Quantifying T lymphocyte turnover. *J. Theor. Biol.* 327:45–87.
- [221] Bains I, Antia R, Callard R, Yates AJ (2009) Quantifying the development of the peripheral naive CD4+ T-cell pool in humans. *Blood* 113:5480–7.
- [222] Murugan A (2012) private communication.
- [223] Zheng GX, et al. (2017) Massively parallel digital transcriptional profiling of single cells. *Nat. Commun.* 8:1–12.
- [224] Hubbell SP (2001) *The Unified Neutral Theory of Biodiversity and Biogeography*.
- [225] Jerne N (1974) Towards a network theory of the immune system. *Ann Immunol (Paris)* 125C:373–89.
- [226] De Boer RJ, Perelson AS (1994) T cell repertoires and competitive exclusion.
- [227] De Boer RJ, Perelson AS (1997) Competitive control of the self-renewing T cell repertoire. *Int. Immunol.* 9:779–790.
- [228] De Boer RJ, Freitas aa, Perelson aS (2001) Resource competition determines selection of B cell repertoires. *J. Theor. Biol.* 212:333–343.
- [229] Mayer A, Balasubramanian V, Walczak AM, Mora T (2018) How a well-adapting immune system remembers. *arXiv:1806.05753v2*.
- [230] Perelson AS, Ribeiro RM (2013) Modeling the within-host dynamics of HIV infection. *BMC Biology* 11:96.
- [231] Perelson AS, et al. (1996) HIV-1 Dynamics in Vivo : Virion Clearance Rate , Infected Cell Life-Span , and Viral Generation Time. *Science* 271:1582–1586.
- [232] Perelson AS, et al. (1997) Decay characteristics of HIV-1-infected compartments during combination therapy. *Nature* 38718:188–191.
- [233] Mccoy CO, et al. (2015) Quantifying evolutionary constraints on B-cell affinity maturation. *Philos Trans R Soc Lond, B, Biol Sci* 370:20140244.
- [234] Yaari G, Kleinstein SH (2015) Practical guidelines for B-cell receptor repertoire sequencing analysis. *Genome Med.* 7:121.
- [235] Hoehn KB, Fowler A, Lunter G, Pybus OG (2016) The Diversity and Molecular Evolution of B-Cell Receptors during Infection. *Molecular Biology and Evolution* 33:1147–1157.
- [236] Gupta NT, et al. (2015) Change-O: A toolkit for analyzing large-scale B cell immunoglobulin repertoire sequencing data. *Bioinformatics* 31:3356–3358.
- [237] Vander Heiden JA, et al. (2014) PRESTO: A toolkit for processing high-throughput sequencing raw reads of lymphocyte receptor repertoires. *Bioinformatics* 30:1930–1932.
- [238] Yaari G, Benichou JIC, Heiden JAV, Kleinstein SH, Louzoun Y (2015) The mutation patterns in B-cell immunoglobulin receptors reflect the influence of selection acting at multiple time-scales. *Philos Trans R Soc Lond, B, Biol Sci* 370:20140242.
- [239] Cui A, et al. (2016) A Model of Somatic Hypermutation Targeting in Mice Based on High-Throughput Ig Sequencing Data. *Journal of immunological methods*.
- [240] Yaari G, et al. (2013) Models of somatic hypermutation targeting and substitution based on synonymous mutations from high-throughput immunoglobulin sequencing data. *Frontiers in immunology* 4:1–11.
- [241] Dhar A, Davidsen K, Matsen FA, Minin VN (2018) Predicting B Cell Receptor Substitution Profiles Using Public Repertoire Data. *PLoS Genet.*
- [242] DeWitt WS, Mesin L, Victora GD, Minin VN, Matsen FA (2018) Using genotype abundance to improve phylogenetic inference. *Mol. Biol. Evol.* 35:1253–1265.
- [243] Jacob J, Kassir R, Kelsoe G (1991) In Situ Studies of the Primary Immune Response to (4-hydroxy-3-nitrophenyl) acetyl . I . The Architecture and Dynamics

- of Responding Cell Populations. *Journal of Experimental Medicine* 173.
- [244] Nieuwenhuis P, Opstelten D (1984) Functional Anatomy of Germinal Centers. *American Journal of Anatomy* 435:421–435.
- [245] Shapiro GS, Aviszus K, Ikle D, Wysocki LJ (1999) Predicting regional mutability in antibody v genes based solely on di- and trinucleotide sequence composition. *The Journal of Immunology* 163:259–268.
- [246] Uduman M, et al. (2011) Detecting selection in immunoglobulin sequences. *Nucleic Acids Res.* 39:W499–W504.
- [247] Yaari G, et al. (2013) Models of somatic hypermutation targeting and substitution based on synonymous mutations from high-throughput immunoglobulin sequencing data. *Front. Immunol.* 4:358.
- [248] Feng J, Shaw DA, Minin VN, Simon N, Matsen IV FA (2017) Survival analysis of DNA mutation motifs with penalized proportional hazards. *arxiv* 1711.04057.
- [249] Unniraman S, Schatz DG (2007) Strand-Biased Spreading of Mutations During Somatic Hypermutation. *Science (80-.)*. 317:1227–1230.
- [250] Oprea M, Perelson AS (1997) Somatic Mutation Leads to Efficient Afinity Maturation When Centrocytes Recycle Back to Centroblasts. *Journal of immunology (Baltimore, Md. : 1950)* 158:5155–62.
- [251] Oprea M, van Nimwegen E, Perelson AS (2000) Dynamics of One-pass Germinal Center Models : Implications for Affinity Maturation. *Bulletin of Mathematical Biology* pp 121–153.
- [252] Victora GD, Nussenzweig MC (2012) Germinal centers. *Annu. Rev. Immunol.* 30:429–457.
- [253] Berek C, Milstein C (1987) Mutation drift and repertoire shift in the maturation of the immune response. *Immunol. Rev.* 96:23–41.
- [254] Eisen HN, Siskind GW (1964) Variations in Affinities of Antibodies during the Immune Response. *Biochemistry* 3:996–1008.
- [255] Tas JMJ, et al. (2016) Visualizing antibody affinity maturation in germinal centers. *Science* 3439:1–11.
- [256] Kuraoka M, et al. (2016) Complex Antigens Drive Permissive Clonal Selection in Germinal Centers. *Immunity* 44:542–552.
- [257] Victora GD, Nussenzweig MC (2012) Germinal Centers. *Annu. Rev. Immunol.* 2012. 30:429–57.
- [258] Victora GD, et al. (2010) Germinal center dynamics revealed by multiphoton microscopy with a photoactivatable fluorescent reporter. *Cell* 143:592–605.
- [259] Lee J, et al. (2016) Molecular-level analysis of the serum antibody repertoire in young adults before and after seasonal influenza vaccination. *Nature medicine*.
- [260] Wang S, Burton D, Kardar M, Chakraborty A (2014) Guiding the evolution to catch the virus: An in silico study of affinity maturation against rapidly mutating antigen. *Bulletin of the American Physical Society* Volume 59, Number 1.
- [261] Murugan R, et al. (2018) Clonal selection drives protective memory B cell responses in controlled human malaria infection. *Science Immunology* pp 1–10.
- [262] Neu KE, et al. (2019) Spec-seq unveils transcriptional subpopulations of antibody-secreting cells following influenza vaccination. *The Journal of Clinical Investigation* 129:93–105.
- [263] Dunn-walters DK, Hare JSO (2018) in *Handbook of Immunosenescence*, ed T. Fulop et al. (eds.) (Springer International Publishing AG 2018).
- [264] Wendel BS, et al. (2017) Accurate immune repertoire sequencing reveals malaria infection driven antibody lineage diversification in young children. *Nature Communications* 8:531.
- [265] Nourmohammad A, Otwinowski J, Luksza M, Mora T, Walczak AM (2018) Clonal competition in B-cell repertoires during chronic HIV-1 infection. *bioarxiv*.
- [266] Horns F, Vollmers C, Dekker CL, Quake SR (2019) Signatures of selection in the human antibody repertoire : Selective sweeps , competing subclones , and neutral drift. *Proceedings of the National Academy of Sciences* 116.
- [267] Vieira MC, Zinder D, Cobey S (2018) Selection and Neutral Mutations Drive Pervasive Mutability Losses in Long-Lived Anti-HIV B-Cell Lineages. *Molecular Biology and Evolution* 35:1135–1146.
- [268] Nourmohammad A, Otwinowski J, Plotkin JB (2016) Host-Pathogen Coevolution and the Emergence of Broadly Neutralizing Antibodies in Chronic Infections. *PLoS genetics* 12:e1006171.
- [269] Blanquart F, Gandon S (2013) Time-shift experiments and patterns of adaptation across time and space. *Ecology Letters* pp 31–38.
- [270] Richman DD, Wrinn T, Little SJ, Petropoulos CJ (2003) Rapid evolution of the neutralizing antibody response to HIV type 1 infection. *Proceedings of the National Academy of Sciences* 100:4144–4149.
- [271] Frost SDW, et al. (2005) Neutralizing antibody responses drive the evolution of human immunodeficiency virus type 1 envelope during recent HIV infection. *Proceedings of the National Academy of Sciences* 102:18514–18519.
- [272] Moore PL, et al. (2009) Limited Neutralizing Antibody Specificities Drive Neutralization Escape in Early HIV-1 Subtype C Infection. *PLoS Pathogens* 5.
- [273] Luksza M, Lässig M (2014) A predictive fitness model for influenza. *Nature* 507:57–61.
- [274] Ferguson AL, et al. (2013) Translating HIV Sequences into Quantitative Fitness Landscapes Predicts Viral Vulnerabilities for Rational Immunogen Design. *Immunity* 38:606–617.
- [275] Shekhar K, et al. (2013) Spin models inferred from patient-derived viral sequence data faithfully describe HIV fitness landscapes. *Physical Review E* 88:62705.
- [276] Barton JP, et al. (2016) Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable. *Nature Communications* 7:1–10.
- [277] Eigen M (1971) Selforganization of Matter and the Evolution of Biological Macromolecules. *Die Naturwissenschaften*.
- [278] Leuthäusser I, Leuthäusser I (1998) An exact correspondence between Eigen ’ s evolution model and a two-dimensional Ising system. *Journal of Chemical Physics* 108:1985–1987.
- [279] Phillips R, Kondev J, Theriot J, Garcia H (2012) *Physical Biology of the Cell* p 2012.
- [280] Zanini F, Brodin J, Thebo L, Lanz C, Albert J (2015) Population genomics of inpatient HIV-1 evolution. *eLife* pp 1–26.
- [281] Smith DJ, Lapedes AS, Jong JCD (2004) Mapping the antigenic and genetic evolution of influenza virus. *Sci-*

- ence 305:371–377.
- [282] Neher RA, Russell CA, Shraiman BI (2014) Predicting evolution from the shape of genealogical trees. *eLife* 3:1–18.
- [283] Bao Y, et al. (2008) The influenza virus resource at the National Center for Biotechnology Information. *Journal of virology* 82:596–601.
- [284] Grenfell BT (2008) *Measles: Nonlinearity and Stochasticity in an Epidemic Metapopulation* (Oxford University Press).
- [285] Rouzine IM, Rozhnova G (2018) Antigenic evolution of viruses in host populations. *PLoS Pathogens* pp 1–16.
- [286] Yan L, Neher R, Shraiman BI (2018) Phylogenetics of rapidly adapting pathogens: extinction and speciation of a Red Queen. pp 1–15.
- [287] Desai MM, Fisher DS (2007) Beneficial mutation selection balance and the effect of linkage on positive selection. *Genetics* 176:1759–98.
- [288] Bedford T, Rambaut A, Pascual M (2012) Canalization of the evolutionary trajectory of the human influenza virus. *BMC Biology* 10:38.
- [289] Marchi J, Lässig M, Mora T, Walczak AM (2019) Multi-lineage evolution in viral populations driven by host immune systems. *arXiv:1906.07444*.
- [290] Regev A, et al. (2017) The Human Cell Atlas. *eLife* 6:e27041.
- [291] Mehta P, et al. (2018) A high-bias, low-variance introduction to Machine Learning for physicists.
- [292] Walczak A, Mugler A, Wiggins C (2012) in *Computational Modeling of Signaling Networks* pp 273–322.
- [293] Tkačik G, Walczak AM (2011) Information transmission in genetic regulatory networks: a review. *Journal of physics. Condensed matter : an Institute of Physics journal* 23:153102.
- [294] for Immunology TJS, Ishikawa T (2009) *Your Amazing Immune System*.
- [295] Hartl DL (2011) *Genetics: Analysis of Genes and Genomes* ed Learning J&B.
- [296] Gillespie, John H (2004) *Population Genetics : A Concise Guide* (Johns Hopkins University Press), p 2004.
- [297] Takahata N, Ishii K, Matsuda H (1975) Effect of temporal fluctuation of selection coefficient on gene frequency in a population. *Proc. Natl. Acad. Sci. U.S.A.* 72:4541–4545.
- [298] Nourmohammad A, Schiffels S, Lässig M (2013) Evolution of molecular phenotypes under stabilizing selection. *J. Stat. Mech. Theory Exp.* 2013.
- [299] Allen LJS (2003) *An Introduction to Stochastic Processes With Biology Applications*. (Prentice-Hall, New York).
- [300] Zapperi S, Ba K, Stanley HE (1995) Self-Organized Branching Processes: Mean-Field Theory for Avalanche. *Phys. Rev. Lett.* 75:4071–4074.
- [301] Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J. Neurosci.* 23:11167–11177.
- [302] Neher RA (2013) Genetic Draft, Selective Interference, and Population Genetics of Rapid Adaptation. *Annual Review Of Ecology Evolution And Systematics* 44:195–215.
- [303] Mustonen V, Lässig M (2010) Fitness flux and ubiquity of adaptive evolution. *Proceedings of the National Academy of Sciences of the United States of America* 107:4248–4253.
- [304] Everitt B (2011) *Cluster analysis* (Wiley, Chichester, West Sussex, U.K).
- [305] Yang Z (2006) *Computational Molecular Evolution* (Oxford University Press).
- [306] Stamatakis A (2014) RAxML version 8 : a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics (Oxford, England)* 30:1312–1313.
- [307] Hoehn KB, Lunter G, Pybus OG (2017) A Phylogenetic Codon Substitution Model for Antibody Lineages. *Genetics* 206:417–427.
- [308] Davidsen K, Matsen IV FA (2018) Benchmarking Tree and Ancestral Sequence Inference for B Cell Receptor Sequences. *Frontiers in immunology* 9:1–13.
- [309] Kussell E, Leibler S (2005) Phenotypic diversity, population growth, and information in fluctuating environments. *Science* 309:2075–8.
- [310] Rivoire O, Leibler S (2011) The Value of Information for Populations in Varying Environments. *Journal of Statistical Physics* 142:1124–1166.
- [311] Bradde S, Mora T, Walczak AM (2019) Cost and benefits of CRISPR spacer acquisition. *Philosophical Transactions of the Royal Society B: Biological Sciences* pp 1–10.
- [312] Lotka AJ (1920) Analytical note on certain rhythmic relations in organic systems. *Proc. Nat. Acad* 6:410–415.
- [313] Volterra V (1956) *Opere matematiche. Memorie e Note.* ed dei Lincei AN (Accademia Nazionale dei Lincei, Rome), pp xxxiii + 604.
- [314] Edelstein-Keshet L (2005) *Mathematical Models in Biology* (SIAM Philadelphia).
- [315] May RM (2001) *Stability and complexity in model ecosystems* (Princeton University Press).
- [316] Marchenko V, Pastur L (1967) Distribution of eigenvalues for some sets of random matrices. *Mat. Sb.* 72:507–536.
- [317] Yule UG (1925) A mathematical theory of evolution, based on the conclusions of Dr J.C. Willis, F.R.S. *Phil Trans R Soc Lond, Ser B* CCXIII.
- [318] Simon HA (1955) On a Class of Skew Distribution Functions. *Biometrika Trust* 42:425–440.
- [319] Kepler TB, et al. (2014) Reconstructing a B-cell Clonal Lineage. II. Mutation, Selection, and Affinity Maturation. *Front. Immunol.* 5.
- [320] Elhanati Y, Marcou Q, Mora T, Walczak AM (2016) RepgenHMM: A dynamic programming tool to infer the rules of immune receptor generation from sequence data. *Bioinformatics* 32:1943–1951.
- [321] Pressé S, et al. (2013) Principles of maximum entropy and maximum caliber in statistical physics. *Rev. Mod. Phys.* 85:1115–1141.
- [322] Barton JP, De Leonardis E, Coucke A, Cocco S (2016) ACE: Adaptive cluster expansion for maximum entropy graphical model inference. *Bioinformatics* 32:3089–3097.
- [323] Nguyen HC, Zecchina R, Berg J (2017) Inverse statistical problems: from the inverse Ising problem to data science. *Adv. Phys.* 66:197–261.
- [324] Mehta P, et al. (2018) A high-bias, low-variance introduction to Machine Learning for physicists.
- [325] Dekosky BJ, et al. (2014) In-depth determination and analysis of the human paired heavy- and light-chain antibody repertoire. *Nat. Med.* 21:1–8.
- [326] Klein AM, et al. (2015) Droplet Barcoding for Single-Cell Transcriptomics Applied to Embryonic Stem

- Cells Resource Droplet Barcoding for Single-Cell Transcriptomics Applied to Embryonic Stem Cells. *Cell* 161:1187–1201.
- [327] Mcdaniel JR, DeKosky BJ, Tanno H, Ellington AD, Georgiou G (2016) Ultra-high-throughput sequencing of the immune receptor repertoire from millions of lymphocytes. *Nat. Protoc.* 11:429–442.
- [328] Grigaityte K, et al. (2017) Single-cell sequencing reveals $\alpha\beta$ chain pairing shapes the T cell repertoire. *bioRxiv:213462*.
- [329] Howie B, et al. (2015) High-throughput pairing of T cell receptor a and b sequences. *Sci. Transl. Med.* 7:301ra131.