

Learning and Memory Strategies in Evolving Environments

from the immune system to neural networks

Dissertation

FOR THE AWARD OF THE DEGREE

“Doctor rerum naturalium”

of the University of Göttingen

within the doctoral program Physics of Biological and Complex Systems
of the Georg-August University School of Science (GAUSS)

submitted by

Oskar H. Schnaack

from Munich

Göttingen, 2021

Thesis Committee

PROF. DR. ARMITA NOURMOHAMMAD,
Research Group Statistical Physics of Evolving Systems,
Max-Planck-Institut für Dynamik und Selbstorganisation

PROF. DR. STEFAN KLUMPP,
Institute for Nonlinear Dynamics,
Georg-August-Universität Göttingen

PROF. DR. MATTHIAS KRÜGER,
Institute for Theoretical Physics,
Georg-August-Universität Göttingen

Members of the Examination Board

Referee: PROF. DR. ARMITA NOURMOHAMMAD,
Research Group Statistical Physics of Evolving Systems,
Max-Planck-Institut für Dynamik und Selbstorganisation

2nd Referee: PROF. DR. STEFAN KLUMPP,
Institute for Nonlinear Dynamics,
Georg-August-Universität Göttingen

Further members of the Examination Board:

PROF. DR. MATTHIAS KRÜGER,
Institute for Theoretical Physics,
Georg-August-Universität Göttingen

DR. VIOLA PRIESEMANN,
Research Group Neural Systems Theory,
Max-Planck-Institut für Dynamik und Selbstorganisation

PROF. DR. PETER SOLLICH,
Institute for Theoretical Physics
Georg-August-Universität Göttingen

DR. DAVID ZWICKER,
Research Group Theory of Biological Fluids,
Max-Planck-Institut für Dynamik und Selbstorganisation

Date of oral examination: 10.12.2021

Abstract

Storing memory of molecular encounters is vital for an effective response to recurring external stimuli. Interestingly, memory strategies vary among different biological processes. These strategies range from networks that process input signals and retrieve an associative memory to specialized receptors that bind only to related stimuli. The adaptive immune system uses such a specialized strategy and can provide specific responses against many pathogens. During its response, the immune system retains some cells as memory to act quicker when reinfections with the same or evolved pathogens occur. However, differentiation of memory cells remains one of the least understood cell fate decisions in immunology.

The ability of immune memory to recognize evolved pathogens makes it an ideal starting point to study learning and memory strategies for evolving environments—a topic with applications far beyond immunology.

In this thesis, I present three projects that study different aspects of memory strategies for evolving stimuli. Indeed, we find that specialized memory strategies can follow the evolution of stimuli and reliably recover memory of previous encounters. In contrast, fully connected networks, such as Hopfield networks, fail to reliably recover the memory of evolving stimuli. Thus, pathogen evolution might be the reason that the immune system produces specialized memories. We further find that specialized memory receptors should trade off their maximal binding for cross-reactivity to bind to evolved targets. To produce such receptors, the differentiation into memory cells in the immune system should be highly regulated. Finally, we study update strategies of memory repertoires using an energy-based model. We find that repertoires should have a moderate risk tolerance to fluctuations in performance to adapt to the evolution of targets. Nevertheless, these systems can be very efficient in distinguishing between evolved versions of stored targets and novel random stimuli.

Contents

Glossary and Abbreviations	vi
1 Introduction	1
1.1 Scope of this thesis	1
1.2 The adaptive immune system and immune memory	3
1.3 Models of the immune system, co-evolution and memory	9
1.3.1 Models of immune memory	9
1.3.2 Models of co-evolution	10
1.3.3 Models of working memory	12
2 Methods	15
2.1 Immune recognition	15
2.2 Decision-making	16
2.2.1 Equilibrium decision-making	17
2.2.2 Non-equilibrium decision-making	18
2.3 Hopfield networks	18
2.3.1 Hebbian learning	19
2.3.2 Memory retrieval	20
2.4 The Hopfield energy landscape as repertoire of encountered patterns . .	21
2.5 Binding affinity of a memory repertoire with general affinity function .	22
3 Optimal evolutionary decision-making to store immune memory	23

4	Learning and organization of memory for evolving patterns	47
4.1	Alternative learning rules	65
5	Risk-utility tradeoff shapes memory strategies for evolving patterns	69
6	Discussion and outlook	85
6.1	Optimal memory of the immune system	86
6.2	Memory structure for evolving inputs	87
6.3	Outlook	90
	References	93
	Appendix A Figure supplements for optimal evolutionary decision-making to store immune memory	105
	Appendix B Supplementary figures for learning and organization of memory for evolving patterns	109
	Appendix C Supplementary figures for risk-utility tradeoff shapes memory strategies for evolving patterns	115

Glossary and Abbreviations

Abbreviations

AUC area under the curve

AUROC area under the receiver operating characteristic curve

bnAb broadly neutralizing antibody

FDC follicular dendritic cell

GC germinal center

ROC receiver operating characteristic

Glossary

affinity maturation During affinity maturation, B-cells undergo multiple rounds of mutation and selection in the germinal center to increase the binding affinity of their receptors to the presented antigen.

antibody Antibodies are Y-shaped proteins that can bind to antigens. A B-cell can turn into a plasma cell that secretes large amounts of antibodies. These antibodies share the binding site with the receptors of the B-cell.

antigen Antigens are structures on the surface of pathogens to which antibodies bind.

B-cell B-cells are an integral part of the adaptive immune system. Through affinity maturation, they can increase the binding affinity of their receptors to most presented pathogens. During an immune response, activated B-cells can differentiate into memory B-cells or plasma cells.

B-cell receptor B-cell receptors are Y-shaped proteins with a variable region at the top that binds to antigens.

Hebbian learning An incremental learning rule that includes new information at a learning rate λ while discarding old information at the rate $(1 - \lambda)$.

Hopfield network Hopfield networks are fully connected graphs that serve as an associative memory model. Typically patterns are stored in the energy minima of the network.

memory B-cell B-cells that differentiate into memory are retained for a fast response to reinfections with similar pathogens.

pathogen Pathogens, e.g., viruses or bacteria, are organisms that can cause diseases.

plasma cell Plasma cells secrete large amounts of antibodies during an immune response.

Symbols

α	specificity (only in [1]); in [2, 3] α is the index for patterns
β	deliberation in [1]; corresponds to bias factor in decision-making and inverse temperature in thermodynamics
β_H	inverse Hopfield temperature for equilibration
β_S	inverse temperature of compartmentalization
χ	random pattern
δ	viral evolution in [1]
$J_{i,j}$	connection matrix of a Hopfield network of size $L \times L$
L	length of patterns
λ	learning rate
m^α	weights of the stored states in the memory repertoire
μ	mutation rate of patterns in [2, 3]
N	number of pattern classes stored in the memory
ν	shape-space position of antigen
ν_r	shape-space position of receptor
r	receptor
r_m	memory receptor
σ	stored patterns (only in [2])
t	time
τ	time interval; in [1] deliberation time
Θ	shape parameter of the affinity function

Chapter 1

Introduction

1.1 Scope of this thesis

The starting point and inspiration for the research I conducted over the past few years is the adaptive immune system and, in particular, immune memory. Thus, it only seems right to start my thesis at the same point, allowing me to introduce open questions and motivate the importance of my work.

With the combination of innate and adaptive immunity, our immune system constantly protects us against a multitude of threats. Classically, the innate and the adaptive immune system were seen as two lines of defense against pathogens [4]. In this simplified picture, the innate system quickly mounts a broad response that can clear basic and very common pathogens. In contrast, the adaptive immune system can protect us against more complicated threats such as evolving viruses. This division of tasks is an oversimplified view and the interaction and cooperation between innate and adaptive immune system are ongoing during infections [4, 5]. Nevertheless, it allows us to get a first understanding of the adaptive immune system's role, namely protection against complicated novel and evolving pathogens.

The immune system mounts an adaptive response when encountering such pathogens that the innate response cannot clear. Through a series of mechanisms, the adaptive immune system produces a specific response to the threat. This response takes between 4-10 days [6], during which the pathogen can almost freely grow within the host, often leading to acute infections and illnesses during that time. While fighting new threats might seem like the primary purpose of the adaptive immune system, its ability to produce a long-lasting memory is equally remarkable. In addition to the adaptive immune system producing highly specific and effective responses to pathogens, some intermediate stages of the maturation process in creating effective responses are retained as long-lasting memory. With this memory, the immune system can then mount fast and specific responses during reinfections with the same pathogen. Remarkably, the memory is not only able to neutralize the same pathogen but also mutated variants of that pathogen.

The ability to form an effective memory for future infections is also the basis for vaccination efforts. Vaccines often contain harmless—or at least less harmful—variants of a pathogen to trigger an adaptive immune response that produces memory cells against future infections. Notably, the development of vaccines against various pathogens, such as polio, measles, and rubella, was possible without a comprehensive understanding of the underlying mechanisms [7]. Despite the undoubted success of many vaccines, not all were successful and some, such as the respiratory syncytial vaccine, were even harmful [8]. Moreover, effective vaccines have yet to be developed against some of the most deadly threats, such as HIV and malaria [9, 10]. There is, however, hope that a better understanding of the mechanisms that underlie the generation of protective immune memory will allow the development of safe, effective vaccines for these and other pathogens.

As mentioned above, contributing to the understanding of immune memory was the original goal of this thesis. In particular, I focus on an often overlooked part of the process, namely the ability to recognize evolved pathogens. The first underlying question of this thesis is thus:

What is the optimal immune memory for evolving pathogens?

This is a very broad question, and a complete answer would need to include a full understanding of all biological processes involved in memory production as well as the coupling to the pathogen evolution. This thesis, however, focuses on an abstract theoretical understanding of this question. While it cannot answer this question in all its details, it does provide important insights into understanding immune memory.

Taking a more abstract approach has further advantages since it allows a translation of the results to other fields. In this case, it allows us to study memory in systems of evolving environments in general. This is a problem not only of interest for immune memory but in many other fields such as neuroscience [11], anthropology [12], social learning [13], and machine learning [14]. In particular, this thesis wants to contribute to the following question.

How should memory of evolving environments be structured, and how should it be updated?

To introduce these problems and answer some of the questions, this thesis is organized as follows. First, in the remainder of this chapter, I will give a more detailed introduction to the immune system, focusing on B-cell memory in section 1.2. In section 1.3, I will introduce some models to study immune memory, adaptation to evolving environments, and other models of working memories. Chapter 2 will introduce the mathematical models and methods used for this thesis. Chapter 3 consists of the publication [1], in which we focus on the first question posed and study the optimal specificity of immune memory against evolving pathogens. Chapters 4 and 5 take a broader approach and focus on memory for evolving environments in general. These chapters consist of the preprints [2] and [3], respectively. In [2], we build a generalized Hopfield model that allows us to study differences between specialized memory, such as

immune memory, and fully connected networks. In [3], we use an energy-based model to study the risk of different memory strategies. The discussion in the final chapter 6 connects all the results and puts them in the context of future research.

1.2 The adaptive immune system and immune memory

Naturally, I can only introduce the concepts relevant to my thesis here. I hope that these compellingly motivate physicists to work on problems motivated by the immune system or biological systems in general. For a more general description of the immune system, I recommend the book by Sompayrac [15] and for a detailed description of immune memory the reviews [7, 16–19] that have all inspired this section. To anyone interested in biologically-inspired topics beyond the scope of this thesis, I recommend the review by Altan-Bonnet et al. [20].

Adaptive immunity develops during the lifetime of an organism and its interactions with a multitude of pathogens. Before introducing immune memory formation through a closer look at the adaptive immune system during primary and secondary infections, I will introduce the central actors of the adaptive immune system, B- and T-cells. In the adaptive immune system, we can distinguish between three types of T-cells. (1) Killer T-cells, as the name suggests, can induce the death of infected cells. Additionally, they can recruit other cells during the immune response via cytokines, i.e., signal proteins. Killer T-cells are often referred to as CD8+ T-cells due to their surface protein [21]. (2) CD4+ helper T-cells identify foreign cells and thus potential threats to the host. They can then initiate and influence the immune response via direct signals to B-cells [21]. (3) Regulatory T-cells, which also present the CD4 protein, suppress an immune response against cells recognized as self. These T-cells thus provide an essential mechanism to prevent autoimmune diseases [22]. For the scope of this thesis, the details about T-cell generation, selection, and differentiation are not essential. Instead, it is sufficient to remember the role of helper T-cells during the B-cell selection described below. For a review of T-cell development, I refer to reference [23].

The second key player in the adaptive immune system are B-cells, or more precisely, their receptors. B-cell receptors are the foundation of one of the most effective weapons against pathogens, antibodies. B-cell receptors have a Y-shape, with variable binding sites at the top of the structure, dictating specificity to an antigen. Through genomic rearrangement, mutation, and selection, these binding sites can specialize their binding to unique molecules of the pathogens called antigens. During infections, B-cells can turn into plasma cells that secrete large amounts of antibodies that closely resemble their receptors in both structure and function. These antibodies are also Y-shaped and share the binding sites with the original B-cell receptor. Thus, a B-cell that can effectively bind to an antigen has the potential to produce highly effective antibodies. In addition to the antigenic binding sites, antibodies can interact with other immune cells through the binding site at their base.

When deployed, antibodies have three main ways to neutralize a threat [24]. (1) By directly binding to their targets, antibodies can neutralize pathogens by preventing them from interacting with the organism. This is important for the protection against viruses or toxins. (2) Antibodies use opsonization for larger-size threats such as bacteria. In this phenomenon, they bind to molecules on the pathogen's surface to form a coating that is recognized by phagocytic cells such as macrophages. These cells then proceed to ingest and destroy the pathogen. (3) Antibodies can activate the complement system, a part of the innate immune system. The complement system is a complex network of more than 30 proteins that can elicit highly efficient and tightly regulated inflammatory and cytolytic immune responses [25].

Independent of the three responses mentioned above, it is crucial to have antibodies with a high binding affinity that can respond quickly to a threat. The immune system generates antibodies with high binding affinity by maturing B-cells and their receptors through affinity maturation. Interestingly, B-cells do not interact with pathogens during affinity maturation but rather with follicular dendritic cells (FDCs). FDCs retain the antigens from pathogens and present them to B-cells [26]. Many studies report that FDCs can retain antigens for extended periods on timescales of 12 months or even longer [27–29]. However, the mechanisms behind this ability or the transport of antigens to the FDCs remain unknown [26].

To understand how the immune system initiates affinity maturation and forms a high-affinity response, we will follow an idealized adaptive B-cell response for two consecutive infections. Figure 1.1 shows a schematic that summarizes a primary adaptive immune response, the differentiation into memory, and the secondary memory immune response.

When the immune system encounters a novel pathogen that the innate immune system cannot clear, it initiates an adaptive immune response. A few days after an infection or antigen administration, germinal centers form inside secondary lymphoid organs (such as lymph nodes or the spleen) to allow the affinity maturation of B-cells [6]. However, before entering the germinal center, naïve B-cells undergo the first phase of antigen-driven differentiation. This differentiation also occurs in the lymphoid organs yet outside of the germinal centers. Here, naïve B-cells first encounter FDCs with the presented antigens [30]. If the B-cell receptors can bind to the presented antigens, the B-cells internalize the antigen and present them to helper T-cells [31]. The T-cells can then activate the B-cells to proliferate. It is generally assumed that this activation by T-cells is essential for all following steps. Nonetheless, studies have also shown that a T-cell-independent B-cell response is possible [32–34]. However, these results go beyond the scope of this thesis, and we will focus on the T-cell-dependent response only. In that case, the activation by T-cells can be seen as the first step of binding affinity selection of B-cell receptors since B-cells need to bind to the antigens in order to present them to T-cells [35, 36].

After activation and proliferation, the B-cells appear to have three fates [7]. First, they can differentiate into short-lived plasma cells that secrete a quick yet low-affinity antibody response to control the initial infection. As the name suggests, the lifespan

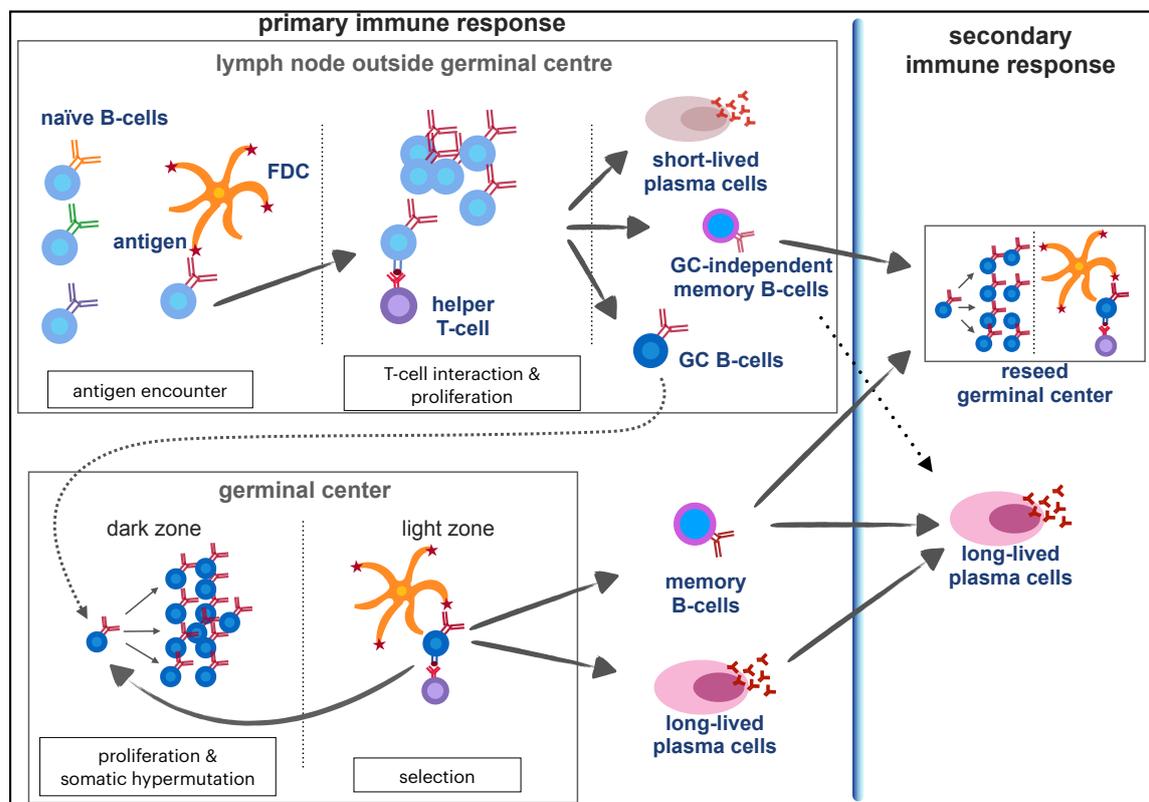


Figure 1.1 Adaptive B-cell immune response and memory differentiation. During the **primary immune response**, B-cells first encounter antigens presented by follicular dendritic cells (FDCs) in the secondary lymphoid organ. B-cells with receptors that bind to the antigen can internalize and present it to helper T-cells for activation. Activated B-cells proliferate and then differentiate into three main types: (1) short-lived plasma cells that provide a quick yet low-affinity antibody response, usually limited to the infection. (2) germinal center (GC)-independent memory B-cells for future infections. (3) GC B-cells that migrate to the germinal center.

The **germinal center** is split into two zones. GC B-cells first enter the dark zone, where they undergo somatic hypermutations and proliferate. They then enter the light zone where they compete for T-cell help. After this selection, GC B-cells have three fates. (1) They can re-enter the dark zone and undergo further rounds of mutation and selection. (2) They can turn into memory B-cells. (3) Typically, high-affinity B-cells turn into long-lived plasma cells that secrete high-affinity antibodies in the late germinal center.

During a **secondary immune response** to a re-encountered pathogen, long-lived plasma cells from the primary response still provide a high-affinity response to homologous (non-mutated) pathogens. Memory B-cells can differentiate into long-lived plasma cells or reseed germinal centers to provide a rapid immune response to evolved pathogens. Typically GC-independent and early GC memory B-cells reseed germinal centers, while late GC memory B-cells differentiate into long-lived plasma cells.

of these plasma cells is generally limited to the infection [34]. Second, other cells differentiate into low-affinity memory B-cells that are kept for future infections [37–39]. These will be discussed in more detail later. Lastly, B-cells can migrate into the germinal center, where they undergo affinity maturation.

Germinal centers are separated into two regions: a dark zone and a light zone [40]. These names are remnants of historical images that show a dark region occupied by many cells and a second lighter region. B-cells that migrate into the germinal center first enter the dark zone. In this zone, B-cells proliferate and somatically hypermutate the variable regions of the receptors with a mutation rate 10^6 times faster than the overall mutation rate of human DNA [41, 42]. This process generates a large diversity of B-cells that originated from B-cells with receptors that bound weakly to the presented antigens.

In the next step, B-cells migrate to the light zone, where they undergo affinity selection. FDCs again present antigens to the B-cells, which they try to internalize for processing and presentation to helper T-cells for further activation. Studies have shown that in the light zone B-cells are significantly more numerous compared to helper T-cells [35, 40] which is the limiting factor in the selection process [6, 40]. Indeed, a recent study suggests that the threshold requirement for B-cells to obtain T-cell help in the germinal center is higher than the threshold for naïve B-cells to obtain T-cell help during initial activation prior to seeding germinal centers [43]. Though a complete description of this light zone process does not yet exist [44], it has been shown experimentally that B-cells with higher binding affinity have a selective advantage to get activated [45–47].

If activated, the B-cells again have one of three fates. First, they can reenter the dark zone for a further round of mutation before again undergoing selection in the light zone. Indeed, B-cells typically undergo many rounds of somatic mutations and selection during affinity maturation and increase the binding affinities of their receptors by up to 10^5 fold [46, 48, 49]. Second, the B-cells can differentiate into memory B-cells as a protection for future infection. The last option is differentiation into long-lived plasma cells. In contrast to the short-lived plasma cells from activated pre-germinal center B-cells, these cells provide a highly specific and high-affinity antibody response that can protect against future infection. While their maximal lifespan is still disputed, they have been observed to live for at least one year [50]. While much progress has been made in understanding affinity maturation, the differentiation in the germinal center is still one of the least understood cell fate decisions in the immune system [6, 40, 46, 51–54].

For a long time, it was assumed that memory B-cells are generated continuously during affinity maturation [55]. Thus, memory B-cells would proliferate without strong affinity-dependent selection [46, 51]. However, studies have shown that memory B-cells often have lower affinity compared to long-lived plasma cells [56], and it has also been found that memory cells show fewer hypermutations than plasma cells [57]. In addition, recent experiments indicate that memory differentiation is highly regulated [58–63], reflecting a preference for memory B-cells at early stages and long-lived plasma cells at

later stages of affinity maturation [59].

To understand why this active regulation might result in the optimal strategy, we look at immune memory response during reinfection. As described above, the immune memory has two kinds of B-cell derived memory available. Akkaya et al. describe the combination of long-lived plasma cells and memory B-cells as two walls of protection [7]. In their analogy, plasma cells with their highly-selected high-affinity antibodies build the first wall. As the plasma cells and antibodies are already circulating in the blood, they provide a rapid response against homologous (unevolved) challenges.

In contrast, the second wall (memory B-cells) can produce a quick and effective response against evolved pathogens whose mutations have escaped the plasma cells. Various studies back up this idea. Purtha et al. used a mouse model to test the immune memory produced for the wild-type West Nile virus [64]. In their study, the original long-lived plasma cells secreted antibodies that only weakly bound to a mutated virus that differed only by one amino acid. Remarkably, memory B-cells that turned into plasma cells during the reexposure secreted antibodies with much stronger binding to the mutated virus. More recently, Leach et al. found that long-lived plasma cells only provide adequate protection against reexposure to a homologous strain of the Narita virus. In contrast, a sufficient immune memory response to evolved viruses requires memory B-cells [65].

The observation that less evolved memory B-cells provide better protection against evolved pathogens introduces cross-reactivity, the ability to generalize, as an additional component to memory production. The results suggest that less evolved and thus lower-affinity B-cells have a larger cross-reactive range against similar or evolved pathogens. Therefore, we can assume a tradeoff between affinity and cross-reactivity during affinity maturation. Different subpopulations of memory B-cells further support this tradeoff. Memory B-cells belong to one of two types, IgM and IgG, depending on their respective constant region [66–69]. During primary infections, IgM B-cells dominate the pre-germinal center region in the lymphoid organs and only some B-cells switch their type to IgG after T-cell activation. IgM B-cells still dominate the early germinal center, associated with fewer mutations and lower affinities. Over time, IgG B-cells outcompete IgM B-cells and IgM B-cells are consequently depleted from the germinal center [70]. Consistent with these results, studies suggest that IgM memory B-cells are generated before or early in the germinal center while IgG B-cells are produced in the later stages of affinity maturation [68].

During a memory response to reinfection, IgG memory B-cells quickly turn into long-lived plasma cells that secrete antibodies of relatively large affinity [71]. Lower affinity IgM memory B-cells, in contrast, are less evolved and can reenter germinal centers to undergo further rounds of affinity maturation. Naturally, reentering a germinal center delays the memory response of IgM memory cells. However, it also allows them to differentiate into IgG memory cells or high-affinity long-lived plasma cells selected for the current threat. While this form of memory response appears similar to a novel naïve response, memory B-cells respond quicker and in higher numbers [72–74]. This allows us to estimate a delay between memory and naïve response of about 1.5 – 5

days [1].

Of course, the tradeoff between affinity and cross-reactivity is not absolute and slight variations are expected. Nevertheless, the observations described above suggest this tradeoff in some form. An apparent exception to this concept seems to be broadly neutralizing antibodies (bnAbs). These antibodies can effectively bind to a wide range of pathogens, seemingly violating the tradeoff. However, bnAbs do not actually bind to many different antigens but are instead optimized for highly conserved regions of pathogens [75, 76], i.e., when a pathogen evolves, the site on which the bnAb binds remains unchanged, allowing recognition of the pathogen. Thus, bnAbs do not contradict the affinity cross-reactivity tradeoff but demonstrate the complexity of the interaction between immune receptors and pathogens. Indeed, many studies are working towards an understanding of bnAbs and much effort goes towards finding vaccination strategies to generate bnAbs against rapidly evolving pathogens such as HIV [77–79]. However, the treatment of bnAbs is beyond the scope of this thesis, and we will instead focus on the interplay between receptors and evolving binding sites.

Beyond bnAbs, the description of immune memory above already omits many biological details about cell signaling and the interaction of all players in the immune response. Nonetheless, it introduces the essential concepts of affinity maturation and B-cell immune memory. However, for the scope of this thesis, it is sufficient to take an even more abstract and simplified point of view and summarize immune memory as follows. During primary infections, the adaptive immune system produces a highly specific response derived from naïve B-cells. Throughout affinity maturation, naïve B-cells receptors evolve from cross-reactive and low-affinity receptors to high-affinity receptors with a small cross-reactive range. While this response can clear the pathogen, it takes about two weeks to produce high-affinity antibodies through affinity maturation [6, 59]. Importantly, this process is not only intended to clear the immediate threat but is also used to produce memory. The memory follows a tradeoff between affinity and cross-reactivity, with early memory showing more cross-reactive and low-affinity properties (IgM memory B-cells). In contrast, memory produced at the final stages of affinity maturation has the largest affinity to the current threat but only a small cross-reactive range (long-lived plasma cells). Memory that is produced at intermediate stages lies in between this tradeoff (IgG memory B-cells). Together, all parts build the immune memory that can quickly respond to reinfections. As expected, cross-reactive memory allows protection against evolved (mutated) pathogens, whereas the most adapted memory can recognize only homologous pathogens.

It is important to note that most of the results on the different types of immune memory and the active regulation of memory production have been published in the last 5 to 10 years, and our understanding is still growing. Thus, while theoretical work has discussed repertoires of cross-reactive immune receptors, it is not surprising that there is little to no theoretical work on optimal memory strategies against evolving threats. These strategies, or from a more general point memory for evolving environments, is the focus of this thesis. The published work [1] presented in chapter 3 uses the tradeoff between cross-reactivity and affinity to discuss optimal memory strategies

against an evolving pathogen. While we study which strategies are optimal when facing various pathogens with different evolutionary rates, this paper does not discuss the formation of a repertoire that simultaneously memorizes multiple pathogens. The repertoire for multiple evolving pathogens is the focus of the following chapters 4 & 5. First, chapter 4 focuses on the optimal storage of evolving patterns. In particular, we focus on the difference between an immune-like specialized repertoire that produces individual B-cells for each different pathogen and other memory strategies such as fully connected networks. Indeed we find that specialized repertoires are favored against a set of evolving targets. Thus, beyond biological constraints, the evolution of pathogens might be the reason for the immune system to encode memory in specialized receptors. Finally, chapter 5 studies how memory repertoires should optimally adapt to evolving targets. In the corresponding paper [1], we study the influence of cross-reactivity, pathogen evolution, and risk on repertoire update rates.

Before we can turn to the methods and results of this thesis, the following chapter will introduce some models that study either immune memory, adaptation to fluctuating or evolving environments, or memory in general.

1.3 Models of the immune system, co-evolution and memory

1.3.1 Models of immune memory

Building models that include the interaction between the immune system and pathogens requires some understanding of the underlying binding between immune receptors and antigens.

Modeling or predicting this binding from genomic sequences, i.e., genotypes of immune receptors and antigens is a challenging problem by itself and goes far beyond the scope of this thesis. To model immune recognition and study binding to evolved pathogens, we rely on the commonly used shape-space introduced by Perelson and Oster [80]. Here we can think of antigens and receptors as points in a high-dimensional space, whose coordinates represent unspecified physicochemical properties [81]. In this space, receptors near each other will recognize similar antigens, and the binding between receptors and antigens depends on their distance. Further, we can interpret random mutations of pathogens and their antigens as a random walk in this space with a step size dependent on the mutation rate. This concept of shape-space is visualized in Fig. 1.2, where the memory receptors have a cross-reactive range in which they can bind to antigens. Through random mutations, the antigen can move away from the neighborhood of the cross-reactive range and can thus avoid recognition during future encounters.

Similar to this thesis, many studies have used a shape-space to describe the interaction between immune receptors and pathogens. In [82], Mayer et al. used the shape-space to introduce a cost function for pathogens given a memory repertoire. They use this cost function to study the optimal resource allocation of a memory

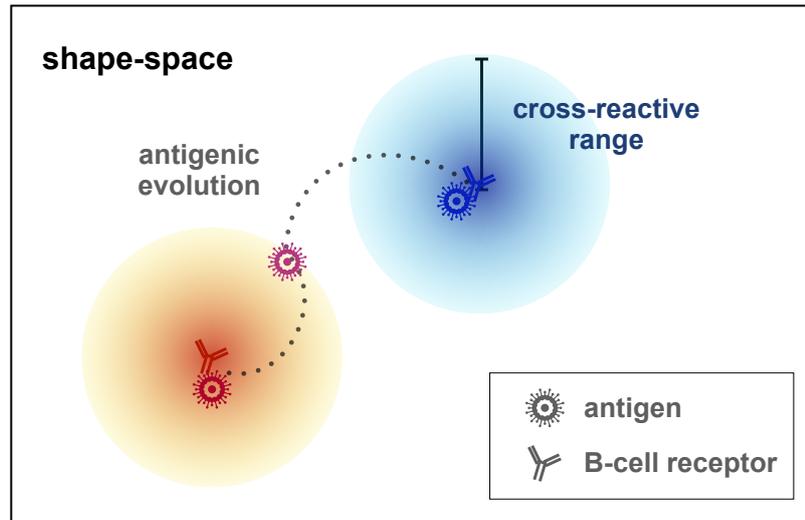


Figure 1.2 Antigen-receptor shape-space. B-cell receptors have a cross-reactive range in which they can bind to antigens. This range is centered around their cognate antigen to which they have the largest affinity. Pathogens evade the immune pressure by evolving their antigens through random mutations. If a pathogen evolves outside of the immune memory’s range, the immune system mounts a new response and forms new memory cells with receptors centered around the evolved antigen.

repertoire that faces a constant set of pathogens. In a later paper, they addressed the reallocation of these resources when the encounter probabilities with the pathogens change [81]. While they do include the effect of different cross-reactivity functions, they constrain their analysis to static pathogens.

Marchi et al. used a different approach and studied the evolution of pathogens in a multi-dimensional shape-space that is constrained by the immune pressure from the host population [83]. Here the immune memories are modeled as regions of cross-reactive binding in which the pathogen cannot evolve. However, the analysis only allows the prediction of different types of viral evolution and does not include an analysis of memory strategies.

Interestingly, most studies ignore the effect of pathogen evolution on the immune strategy. If included at all, pathogen evolution is usually modeled as a fluctuation of the pathogen. Of course for two consecutive encounters, fluctuation and evolution appear identical. However, while fluctuations remain centered around the original pathogen, evolution drives the pathogen away from its original position in shape-space [84]. Thus, the evolution of pathogens plays a crucial role when memory strategies are evaluated over multiple encounters with pathogens.

1.3.2 Models of co-evolution

While the theoretical work on the interaction between immune memory and evolving pathogens has been sparse, much more research exists on evolution or adaptation

in evolving environments. As we will discuss in this section, all models have their justification and field of application. However, they are not suitable to describe the process of memory production and utilization for evolving targets.

Most models rely on some model of fitness. In general, fitness is the measure of evolutionary success, and it is most often interpreted as the long-term growth rate. Biophysicists usually draw an analogy between fitness and negative energy. Similar to energy differences between states, only the fitness differences determine the outcome, with states of higher fitness being favored. A fitness landscape then, similar to an energy landscape, defines the evolutionary success for each state. Thus, a population that produces variance through mutations can explore the space and move towards the fitness peaks through selection. Fisher famously stated his fundamental theorem of natural selection that indeed “the rate of increase in the average fitness of a population is equal to the genetic variance of fitness of that population” [85, 86]. Thus, the evolution of a population can be interpreted as a fitness optimization protocol that moves states of lower fitness towards the peaks in the landscape. Affinity maturation in the germinal centers is an example of this process. During each round of mutations, new states surrounding the original receptor are explored. Then during the selection phase in the light zone, B-cells that bind strongest to the antigen have a fitness advantage as they get more T-cell help. Thus, the population of germinal center B-cells moves up the peaks in the fitness landscape.

To extend studies to changing environments, Mustonen and Lässig suggest dynamic fitness seascapes with evolving fitness peaks [87]. Within these non-equilibrium models, fitness flux, rather than fitness, turns out to be the correct quantity to measure adaptive evolution of genotypes [88]. Models of dynamic fitness seascapes have been used to study adaptive gene expression in *Drosophila* [89] or the host-pathogen coevolution of chronic infections such as HIV [90]. However, the models have not yet been able to capture the effectiveness of states adapted to one landscape and then to a changing landscape at a later point. Instead, these models might be more suited to study how artificial changes to the fitness landscape can direct evolution in a desired direction. References [91, 92] are examples of recent approaches to use control theory to steer evolution. Related work [93] suggests chirp-protocols that rapidly switch between two or more environments that share some fitness peaks. The different environments can, for example, represent different viral strains and the shared peaks are conserved regions. Thus, a protocol that can move the population to a shared peak can produce broadly neutralizing antibodies.

A different approach to studying adaptation to evolving environments is phenotype switching on a population level. Organisms might display different or even change their phenotypes without any change in their genotypic states. For example, butterflies can change their phenotype wing pattern, or mammals vary their fur colors depending on the weather conditions and provide suitable seasonal camouflage [94]. Significantly, these adaptations happen without changes to the genotype. The scenario described here corresponds to a *tracking strategy* that occurs when the environment switches slow enough, the transition is observable by organisms and the entire population can safely

adapt their phenotype [95]. Depending on the transition speed between environments, their observability, and the selectivity of the environment, the long-term optimal solution may favor other strategies such as *bet-hedging strategy* [95]. In this limit, the optimal solution for a population is always having all phenotypes present. Thus, even though some states are suboptimal, the population will not die out because it does not have a fitting solution. This tradeoff between the possible gain in one environment and preparation for other environments illustrates a vital component of these models, namely, focusing on long-term growth, not short-term gains, similar to many other studies [96–101]. This concept is also the basis for the optimizations in this thesis. Here the focus is not on one infection or encounter but instead on long-term outcomes. Beyond this shared method, the results of phenotype switching research do not apply to the memory of evolving environments. Even though the studied environments change, they do not evolve but only switch between a predetermined set of states. However, it would be interesting to see how these models would change if the environment not only switched between states but instead included the evolution of these states.

1.3.3 Models of working memory

This thesis aims to go beyond immune memory and study memory for evolving environments in general. This section will introduce some other forms of memory and possible applications to evolving environments as a first step.

To bridge the gap from immune to other forms of memory, we will start with olfaction. Similar to the immune system, olfactory sensing relies on molecular recognition between olfactory receptors and odor molecules. However, the form of memory production is widely different. In mammals, odor recognition begins with olfactory receptors located in the olfactory bulb in the nasal cavity [102]. In contrast to immune receptors, olfactory receptors can bind to many unrelated molecules. This broad binding allows the olfactory system to recognize orders of magnitude more odors compared to the number of olfactory receptors [103–105]. A given odor is composed of many mono-molecules at different concentrations [106–108]. These molecules are drawn from a space of about 10^4 distinct mono-molecules that react with the olfactory receptors (a total of $\sim 300 - 1000$ in mammals [109–113]). This interaction leads to a distinct distributed signal and spatiotemporal activation pattern in the olfactory bulb [114, 115]. This pattern is then transmitted to the olfactory cortex. Interestingly, only few synapses connect the receptors and cortex, making it one of the most direct links between brain and environment [102]. The olfactory cortex serves as the actual pattern recognition and storage device. Similar to other brain regions, the synaptic connections in the olfactory cortex form when they are co-stimulated by a given odor pattern, thus forming an associative memory [116–121]. This memory can then be retrieved when it encounters the same or similar pattern.

To theoretically study how the associated memory of the olfactory system generates and recovers patterns, artificial neural networks that store auto-associative memory can be used. Neural networks also have activation patterns as their inputs. These patterns (stimuli) then trigger interactions between the network’s encoding nodes, resembling the

co-activation of synaptic connections in a cortex. This way, the ensemble of interactive nodes builds a robust memory of the patterns. The simultaneous stimulation allows pattern recovery when the network is re-exposed to a pattern or even just a fraction of a pattern. One of the models used to study associated memory is the Hopfield neural network [122] with Hebbian learning rule [123]. This energy-based model can encode a set number of patterns as attractors in its energy minima. The number of patterns that Hopfield networks can reliably store and recover grows linearly with the pattern length. At the same time, the model is still analytically solvable. This allows a much better understanding of network dynamics compared to the commonly used black-box machine learning techniques. However, the connection between standard Hopfield networks and real synaptic neural networks has been debated over the past decades. Nevertheless, the Hopfield network provides a simple and solvable coarse-grained model of the synaptic network, relevant for working memory in the olfactory system and the hippocampus [121].

Research on Hopfield networks and other neural networks mostly focuses on disentangling signal from noise. These systems are often highly effective in learning static patterns, even when given noisy inputs [124]. In machine learning, noise inputs are even a key ingredient during the training phase to prevent overfitting [125, 126]. However, research on evolving systems is still sparse, and only some specialized machine learning approaches allow for learning dynamically evolving inputs [14, 127]. Thus, we still lack a general framework for learning evolving patterns relevant to many real-life applications [127].

The papers presented in chapters 4 and 5 aim to contribute to the research of evolving environments. In [2], we show that in contrast to the static case, classical Hopfield models [122] fail to recover evolving patterns efficiently. In particular, we demonstrated that networks need an increased learning rate to follow the dynamics of evolving environments. This heightened learning rate distorts the network's energy landscape so that energy minima become connected by narrow paths (mountain passes). Through these passes, patterns can equilibrate into a wrong attractor and are thus associated with a false memory. Further, we demonstrate that specialized memory compartments can overcome this misclassification. Interestingly, the immune system's B-cell memory is a form of this specialization. We can thus postulate that the evolution of pathogenic patterns may be one of the key reasons why the immune system keeps a specialized memory.

In chapter 5, we use a general energy model, similar to the Hopfield energy function, with Hebbian learning to study updates to a memory repertoire of specific receptors. We introduce a tradeoff between the mean affinity and fluctuations in affinity to investigate different regimes of learning when repertoires face evolving targets. Similar to [2], we find that systems increase their learning rate to follow the evolution of the targets. However, this faster learning also leads to increased fluctuations in performance. Indeed we find that very risk-tolerant systems will use large learning rates and focus the repertoire around the most recently encountered targets. In contrast, risk-averse strategies will effectively shut down the learning when facing evolving targets. We

show that moderate risk tolerance is desirable to build effective repertoires that can distinguish between previously encountered and novel stimuli.

Beyond this biological intuition, both of our papers [2, 3] offer a principled analytical framework to study learning and memory generation in out-of-equilibrium dynamical systems. Thus, this work contributes to the understanding of immune memory and is also a starting point for the needed theoretical framework of machine learning of evolving environments [127].

Chapter 2

Methods

This chapter will introduce the different models I have used to study memory strategies. While I believe that it is helpful to read this section first, it is by no means necessary to understand the articles presented in the next three chapters. All papers are self-contained and introduce all required tools in the materials and methods section [1], or their appendix respectively [2, 3]. Indeed, some overlap and repetition to these sections is unavoidable.

This chapter follows the order of the thesis. First, I will introduce a biophysical model for receptor-antigen binding in section 2.1 and the decision-making theory in section 2.2. Together, these concepts build the foundation for the model of [1] (chapter 3). Section 2.3 introduces the notation of the Hopfield network that we use in [2] (chapter 4). Section 2.4 bridges the gap between the standard description of the Hopfield energy function to a model for the affinity of a memory repertoire. Section 2.5 generalizes this affinity function to the model we use in the final paper of the thesis [3] (chapter 5).

2.1 Immune recognition

In this section, I will introduce a biophysics inspired model for the probability $P_{\text{recog.}}^{(m)}$ of initiating an immune memory response against a pathogen. The first step towards this probability is the binding affinity $E(r_m, \nu)$ of memory receptors r_m to antigens ν that determines their binding during the immune response. To study this binding without going into the physico-chemical details, we can use a common shape-space between receptors and antigens (see Fig. 1.2). In that space, we assume that each receptor r has a cognate antigen ν_r^* against which it has the highest affinity. For memory receptors, this cognate antigen belongs to the pathogen that the receptor was optimized against. As discussed above, receptors typically have some cross-reactive binding to related antigens that allows them to recognize evolved pathogens. In the shared shape-space, we can express the affinity between a receptor r and an arbitrary target antigen ν in terms of the antigenic distance $d_r(\nu) = \|\nu - \nu_r^*\|$ between the receptor's cognate antigen ν_r^* and the target ν : $E(r, \nu) \equiv E(d_r(\nu))$ (see Fig. 1.2).

Beyond the binding affinity $E(r_m, \nu)$, immune-pathogen recognition and thus the initiation of a memory response depends on the encounter rate $\gamma_\nu(t)$ between a memory receptor r_m and the antigen ν at a given time t . This encounter rate $\gamma_\nu(t)$ depends on the abundance of antigens and immune receptors, and hence, can vary during infections within a host. For a short time interval $[t, t + dt]$, we can express the probability that a receptor r_m encounters and binds to an antigen ν by, $\gamma_\nu(t)E(r_m, \nu)dt$. This formulation of the immune-pathogen recognition is similar to the notion used in reference [82]. An immune memory response is triggered through the recognition of an antigen by a circulating memory receptor. However, if no such recognition occurs during the early stage of the infection ($\tau \approx 1.5 - 5$ days [1]), the immune system initiates a naïve response. In [1], we term this time τ the deliberation time, as during this time the memory can be exploited.

As it turns out, it is much easier to calculate the probability that an antigen is recognized through a novel naïve response $P_{\text{recog.}}^{(0)}$, which is equivalent to the probability that the memory receptor cannot recognize the antigen $1 - P_{\text{recog.}}^{(m)}$. To do so, we can split the deliberation time τ into infinitesimal intervals $\lim_{N \rightarrow \infty} \tau/N$ during which we calculate the probability of no memory-pathogen binding. We then multiply these intervals to find

$$\begin{aligned} P_{\text{recog.}}^{(0)}(\nu) &= \lim_{N \rightarrow \infty} \prod_{i=1}^N \left(1 - \gamma_\nu(t_i)E(r_m, \nu) \frac{\tau}{N} \right) \\ &= e^{-\int_0^\tau \gamma_\nu(t)E(r_m, \nu)dt}. \end{aligned} \quad (2.1)$$

This result shows the dependence on the binding affinity well. Suppose the memory receptors r_m have a large affinity $E(r_m, \nu)$ for the current pathogen ν . In that case, the exponent becomes large and no naïve response is needed, i.e., the memory responds to the threat. On the other hand, if the pathogen has evolved too much, the affinity $E(r_m, \nu)$ will be close to zero, and the novel response is mounted with high probability. Indeed, in the case of zero affinity (e.g., for a novel pathogen) the probability of a novel response $P_{\text{recog.}}^{(0)}(\nu) = 1$.

2.2 Decision-making

In our publication [1], we study the adaptive value of immune memory under various evolutionary scenarios, using the framework of decision theory. Simply put, decision-making describes the probabilities of actions as a function of their utilities.

The following section will introduce the relevant concepts of equilibrium and non-equilibrium decision-making. While initially inspired by the work of von Neumann and Morgenstern [128], the concepts are mainly inspired by [129] for equilibrium and [130] for non-equilibrium decision-making. The notation of this section will follow the notation of [1].

2.2.1 Equilibrium decision-making

We assume that the immune system can choose between mounting a memory response or starting a new naïve response. In the language of decision-making theory, a rational decision-maker can choose between two possible actions $a \in \{\text{naïve, memory}\}$ with corresponding utility U_a . In addition, the decision-maker may have a prior preference for each action, denoted by the prior probability distribution $Q_0(a)$. This prior will sway the decision and thus change the outcome probabilities. To tradeoff prior and new information, the constrained decision-maker should choose its actions according to the optimized probability density $Q(a)$ that maximizes the expected utility while satisfying constraints due to the prior assumption [128, 129],

$$Q(a) = \operatorname{argmax}_{Q(a)} \left(\sum_a U_a Q(a) - \frac{1}{\beta} D_{KL}(Q(a)||Q_0(a)) \right). \quad (2.2)$$

Here, $D_{KL}(Q(a)||Q_0(a)) = \sum_a Q(a) \log(Q(a)/Q_0(a))$ is the Kullback-Leibler distance between the rational distribution $Q(a)$ and the prior distribution $Q_0(a)$ and $1/\beta$ is a Lagrange multiplier (bias factor) in the tradeoff between prior knowledge and information about utilities of the outcomes. This tradeoff can either be interpreted as the ability or willingness of decision-makers to process new information and deviate from prior assumptions. Independent of the interpretation, the optimal solution for a rational yet constrained decision follows,

$$Q(a) = \frac{1}{Z} Q_0(a) e^{\beta U_a}, \quad (2.3)$$

where $Z = \sum_a Q_0(a) e^{\beta U_a}$ is a normalization factor. This result allows a good understanding of the bias factor β . When information processing is highly efficient (i.e., the bias factor $1/\beta \rightarrow 0$), the rational decision-maker deterministically chooses the action with the highest utility. On the other hand, if the system cannot process the new information or has a strong prior (i.e., $1/\beta \rightarrow \infty$), the decision-maker hardly changes its opinion and acts according to its prior belief (i.e., $Q(a) = Q_0(a)$). The solution in eq. 2.3 resembles the equilibrium solution to the decision problem with expected free utility

$$\langle U \rangle = \sum_a U_a Q(a). \quad (2.4)$$

This thesis focuses on unbiased systems, i.e., systems where the prior distribution is uniform across actions. In this case, rational decision-making is equivalent to a maximal entropy system with energy constraint [131], resulting in the Boltzmann-like probability of actions $Q(a) \sim \exp[\beta U_a]$. In the thermodynamic analogy, β is the inverse temperature and the utilities U_a correspond to the negative energy of the states. In the analysis of chapter 3, the resulting probabilities for utilizing a memory response

$Q(\text{mem})$ or mounting a naïve response $Q(\text{naïve})$ are thus given by

$$Q(\text{mem}) = 1 - Q(\text{naïve}) = \frac{e^{\beta U_{\text{mem}}}}{e^{\beta U_{\text{mem}}} + e^{\beta U_{\text{naïve}}}}, \quad (2.5)$$

which is a sigmoidal function, dependent on the utility of each action.

In [1], we connect these decision-making probabilities to the biophysical description of the immune response through recognition of an antigen by either of these cell types (2.1). This allows us to derive the utility gain of memory responses compared to a naïve response as well as the bias factor β .

2.2.2 Non-equilibrium decision-making

Analogous to the standard description of thermodynamics, also decision-making theory can be extended to non-equilibrium processes [130]. This extension is relevant when decisions are made at finite time where the decision-maker cannot find the optimal solution or cannot equilibrate to the optimum infinity slowly. Indeed, the response of the immune memory to an evolved pathogen is such a non-equilibrium response as the binding profile of the memory $E(r_m, \nu)$ differs from the optimal solution $E(r_\nu, \nu)$ with the cognate receptor of the evolved pathogen.

In the non-equilibrium decision-making, using a suboptimal strategy $\tilde{Q}(x)$ leads to an extra cost. Analog to the information possessing factor in eq. 2.2, the cost depends on the Kullback-Leibler divergence between the suboptimal strategy $\tilde{Q}(x)$ and the optimal equilibrium strategy $Q_{\text{eq}}(x)$ scaled by the temperature of the decision process [130]. This dissipated utility

$$U_{\text{diss}} = \frac{1}{\beta} D_{KL}(\tilde{Q}(x) || Q_{\text{eq}}(x)), \quad (2.6)$$

similar to the dissipated energy, constitutes the loss of the non-equilibrium process. Together with the equilibrium free utility (see eq. 2.4), we can then calculate the net utility of a non-equilibrium process [130]

$$U_{\text{net}} = \langle U \rangle - U_{\text{diss}}. \quad (2.7)$$

This net utility can be interpreted as the extracted (information-theoretical) utility of the non-equilibrium decision. In the corresponding thermodynamic framework, the net utility equals the extracted work of an out-of-equilibrium process. In these systems, the difference between extracted work and free energy is given by the dissipated heat [132].

2.3 Hopfield networks

In [2] we move beyond our simplified picture of memory as antigenic snapshots, and study how memory actually learns antigenic patterns. To contrast the specialized

receptors of immune memory, we use one of the most widely used solvable models of associative memory, the (Little-)Hopfield network. Some controversy exists about the name as Hopfield based his work [122] on an earlier paper of William Little [133]. Throughout this thesis, I will use the more widely used name Hopfield network. Independent of the name, these networks have received considerable attention as it has been hypothesized that their accessible memory operates according to the same principle as biological associative memory [121, 134, 135].

The storage targets of Hopfield networks are binary patterns (stimuli) σ of length L with entries: $\sigma = (\sigma_1, \dots, \sigma_L)$, with $\sigma_i = \pm 1, \forall i$. To store an associative memory of these patterns, we define a fully connected graph (network) represented by an interaction matrix $J = (J_{i,j})$ of size $L \times L$ (see Fig. 2.1). This interaction defines the Ising-like Hopfield energy function (Hamiltonian) [122]

$$E_J(\sigma) = -\frac{1}{2L} \sum_{i,j} J_{i,j} \sigma_i \sigma_j, \quad (2.8)$$

whose minima (attractors) correspond to the stored states of the system (see Fig. 2.1).

2.3.1 Hebbian learning

To train the network (i.e., store the patterns), we focus on an incremental Hebbian learning process, during which different patterns are presented sequentially and in random order. As a pattern σ^α is presented, the interaction matrix J is updated according to the following Hebbian update rule [136]

$$J_{i,j} \longrightarrow J'_{i,j} = \begin{cases} (1 - \lambda) J_{i,j} + \lambda \sigma_i^\alpha \sigma_j^\alpha, & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (2.9)$$

Here $\lambda \in [0, 1]$ is the *learning rate* that stores the patterns in the energy minima associated with the matrix J . With this particular Hebbian learning rule, Hopfield networks have a *capacity* of $N \approx 0.14 L$ patterns, i.e., they can reliably store and recover up to $N \approx 0.14 L$ patterns¹ [122, 137, 138]. Interestingly, the Hebbian learning rule also works for noisy inputs of patterns [124]. While other learning rules [139–144] can achieve a higher capacity, Hebbian learning is analytically most tractable and we stay well below the Hopfield capacity throughout this thesis. It should be further noted that Hopfield networks have been criticized for the undesirable scaling of the capacity with the length of the patterns (L) while the number of interactions $J_{i,j}$ grows with (L^2) [145]. However, we do not demand optimal efficiency in that regard. Thus, we do not exploit the scaling, instead we focus on the analytic tractability of Hopfield networks.

¹Networks also have minima close to the negated patterns, i.e., to $-\sigma^\alpha$, but these do not contribute the capacity and do not play any role in what follows.

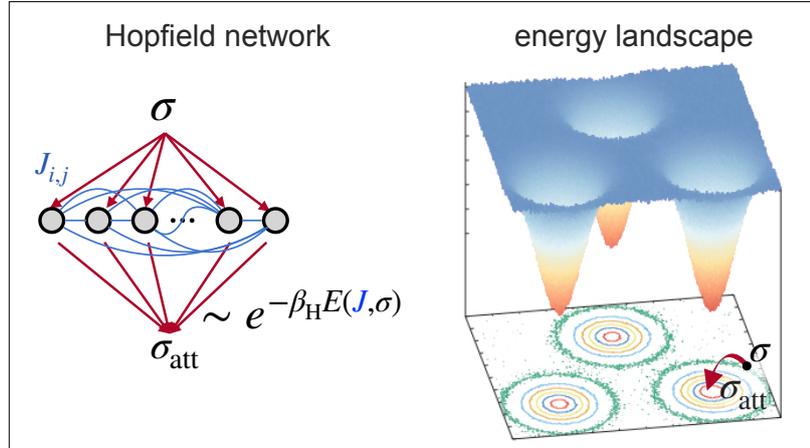


Figure 2.1 Hopfield network and energy landscape. Schematic of Hopfield network (left) as a fully connected graph with coupling matrix $J_{i,j}$ between the nodes. This interaction defines the energy landscape $E(J, \sigma) = -\frac{1}{2L} \sum_{i,j} J_{i,j} \sigma_i \sigma_j$ (eq. 2.8) with equipotential lines are shown in the bottom 2D plane (right). The attractors of the system in the energy minima correspond to stored states. To find an associative memory, a pattern σ with energy $E(J, \sigma)$ equilibrates with inverse temperature β_{H} and falls into the attractor σ_{att} (see eq. 2.10).

2.3.2 Memory retrieval

After a network is sufficiently trained, the patterns can be retrieved even from distorted copies, i.e., copies where a fraction of the entries (spins) σ_i are flipped. To find the associated attractor for a given pattern, the state of that pattern is equilibrated in the energy landscape (eq 2.8). In [2] (see chapter 4) we use a Metropolis algorithm for the equilibration. To do so, we make spin-flips in a presented pattern $\sigma^\alpha \rightarrow \tilde{\sigma}^\alpha$ and accept a spin-flip with probability

$$P(\sigma^\alpha \rightarrow \tilde{\sigma}^\alpha) = \min(1, e^{-\beta_{\text{H}} \Delta E}), \quad (2.10)$$

where $\Delta E = E(J, \tilde{\sigma}^\alpha) - E(J, \sigma^\alpha)$ and β_{H} is the inverse (Hopfield) temperature for pattern retrieval in the network. In the low temperature regime (i.e., high β_{H}), equilibration in networks with working memory drives a presented pattern σ^α towards a similar attractor $\sigma_{\text{att}}^\alpha$, reflecting the memory associated with the corresponding energy minimum (see Fig. 2.1). In the limit of $\beta_{\text{H}} \rightarrow \infty$, this equilibration procedure becomes deterministic and only spin-flips, that reduce the energy, are accepted.

A detailed description of the numerical implementation of the Hopfield network with Hebbian learning as well as the pattern retrieval is given in Appendix A of [2] (chapter 4).

2.4 The Hopfield energy landscape as repertoire of encountered patterns

The treatment of the Hopfield network given above is the classical description found in many textbooks. The formulation of the energy function (eq. 2.8) and learning rule (eq. 2.9) are beneficial when interpreting the system as a fully connected network. We can, however, reformulate the problem to resemble the binding affinity of a memory repertoire. As a first step, we introduce the bra-ket notation where $|\sigma\rangle$ represents the pattern, $\langle\sigma|$ its transpose, and $\langle\sigma|\sigma'\rangle \equiv \sum_i \sigma_i \sigma'_i$ is the scalar product. As it turns out, we can then rewrite the Hebbian learning rule for coupling matrix J (eq. 2.9) for an update with pattern $\sigma^\alpha(t)$ at time t as

$$J(t) \longrightarrow J(t+1) = (1 - \lambda) J(t) + \lambda (|\sigma^\alpha(t)\rangle \langle\sigma^\alpha(t)| - \mathbf{1}), \quad (2.11)$$

where the identity matrix $\mathbf{1}$ ensures $J_{i,i} = 0$. Let us now assume that we start with an empty interaction matrix J at time zero², i.e., $J_{i,j}(t=0) = 0; \forall i, j$. We then present the system with patterns $\sigma(t)$ at each time point. After time T the interaction matrix is then given by

$$\begin{aligned} J(T) &= \lambda \sum_{t=0}^{T-1} (1 - \lambda)^{T-1-t} (|\sigma(t)\rangle \langle\sigma(t)| - \mathbf{1}) \\ &= \lambda \sum_{t=1}^{T-1} (1 - \lambda)^{T-1-t} |\sigma(t)\rangle \langle\sigma(t)| - \lambda \sum_{t=0}^{T-1} (1 - \lambda)^t \mathbf{1}. \end{aligned} \quad (2.12)$$

After a sufficiently large time T the weight in front of $\mathbf{1}$ ($\lambda \sum_{t=0}^{T-1} (1 - \lambda)^t$) will go to 1 and only the weights for the encountered patterns $\sigma(t)$ keep a time dependence. We can further simplify eq. 2.12 by introducing the weights for the encountered patterns $c_t = \lambda(1 - \lambda)^{T-1-t}$. This reformulation allows us to write the energy of any given pattern χ as

$$\begin{aligned} E(J(T), \chi) &= -\frac{1}{2L} \langle\chi|J(T)|\chi\rangle = -\frac{1}{2L} \left(\sum_t c_t \langle\chi|\sigma(t)\rangle \langle\sigma(t)|\chi\rangle - \langle\chi|\mathbf{1}|\chi\rangle \right) \\ &= -\frac{1}{2L} \sum_t c_t |\langle\sigma(t)|\chi\rangle|^2 + \frac{1}{2}. \end{aligned} \quad (2.13)$$

Importantly, this is just a reformulation of the original energy function in a Hopfield network with Hebbian learning. However, the form in eq. 2.13 allows us to reinterpret the energy of any given pattern χ as a weighted sum of squared overlaps with all patterns $\sigma(t)$ previously stored in the system. The sum is then scaled by a model-dependent factor $-\frac{1}{2L}$ and shifted by $\frac{1}{2}$. This shift is chosen such that the expected energy of a random new pattern (i.e., a pattern unrelated to all per-encountered

²Indeed, the Hebbian learning rule ensures that the dependence on the initial state $J(0)$ decays exponentially $\sim (1 - \lambda)^t$ but for simplicity, we will use $J_{i,j}(t=0) = 0$ as an initial state.

patterns) is equal to zero.

Indeed, this interpretation shows a great correspondence with the binding affinity of the immune memory discussed at the beginning of this chapter. The squared distance between the stored patterns $\sigma(t)$ and the tested pattern χ fulfills our requirement of the distance depended binding affinity $E(r_m, \nu)$ between memory receptor r (here $\sigma(t)$) and the presented antigen ν (here χ). The sum over all pre-encountered patterns gives the affinity of the entire memory repertoire (similar to the notation in [82]).

2.5 Binding affinity of a memory repertoire with general affinity function

Using the Hopfield energy as the foundation for a model of memory repertoire limits the analysis to a squared affinity function. In [3] (chapter 5), we extend the model to arbitrary exponents Θ of the overlap between stored states and presented pattern. Here we define the space of all 2^L possible binary patterns ψ^α with $\alpha \in \{1, \dots, 2^L\}$. The memory repertoire allocates normalized weights m^α to each pattern ψ^α . In a Hebbian learning setting, these weights are updated by $m^\alpha(t+1) = (1 - \lambda)m^\alpha(t) + \lambda\delta^{\alpha,\beta}$, when the pattern ψ^β is presented at time t . Analog to the Hopfield energy as shown in eq. 2.13, we define the affinity of the memory repertoire to a given pattern

$$\mathcal{A}(\chi) = \mathcal{A}_0 \sum_{\alpha} m^\alpha |\langle \psi^\alpha | \chi \rangle|^\Theta - \mathcal{A}_{\text{rand}}, \quad (2.14)$$

where \mathcal{A}_0 sets the model's affinity (energy) scale and $\mathcal{A}_{\text{rand}}$ is a shift of the affinity function that sets the expected affinity of novel patterns to zero. For $\mathcal{A}_0 = -\frac{1}{2^L}$, $\mathcal{A}_{\text{rand}} = -\frac{1}{2}$, and $\Theta = 2$ this model is equivalent to the Hopfield energy in eq. 2.13. In [3], we use the general affinity function in eq. 2.14 to study the behavior of memory repertoires in various scenarios analytically. However, comparing these analytic results to simulations becomes exponentially hard with the size of the patterns (L) since simulations would need to evaluate the vector m^α of size 2^L . At this point, the correspondence between the affinity model (eq. 2.14) and the Hopfield energy (eq. 2.13) comes in very handy, as it allows us to efficiently simulate the model (for $\Theta = 2$) by only keeping track of the interaction matrix $J_{i,j}$ that is only of size L^2 .

Chapter 3

Optimal evolutionary decision-making to store immune memory

The following section consists of the article

O. H. Schnaack and A. Nourmohammad. “Optimal evolutionary decision-making to store immune memory”. In: *eLife* 10 (Apr. 2021), e61346. DOI: 10.7554/eLife.61346

reprinted under the Creative Commons Attribution License (CC BY 4.0)¹.

Armita and I designed the scope of this research together. I then built the model and wrote the Julia code for the simulations and data analysis. All authors contributed to the writing of the manuscript, and I produced all figures shown in the paper. The figure supplements referenced in the paper are shown in appendix A.

In this publication, we build a model to evaluate the performance of immune memory strategies against evolving pathogens. We base our model on an analogy between the biophysical probability of mounting a memory immune response and the theory of decision-making. Through this analogy, we derive a utility function for a memory immune response that is independent of biological details.

The strength of our model lies in its ability to describe memory production and usage with two parameters. These parameters are based on two observations about immune memory. The first (α) describes the tradeoff between maximal binding affinity and the range of cross-reactivity observed in memory B-cells. The second parameter β depends on the delay between a memory and novel naïve response.

In the context of this thesis, this publication aims to answer the question of which germinal center B-cells should differentiate into memory. Here we primarily focus on the memory of one evolving pathogen and connect the evolutionary rate of these pathogens to the optimal cross-reactivity of the memory. We also study what memory strategy is best when fighting a multitude of independent pathogens with varying

¹<https://creativecommons.org/licenses/by/4.0/>

evolutionary rates. However, within this framework, we do not include the interaction of different memories in a repertoire.

Optimal evolutionary decision-making to store immune memory

Oskar H Schnaack^{1,2}, Armita Nourmohammad^{1,2,3*}

¹Max Planck Institute for Dynamics and Self-organization, Göttingen, Germany;

²Department of Physics, University of Washington, Seattle, United States; ³Fred Hutchinson Cancer Research Center, Seattle, United States

Abstract The adaptive immune system provides a diverse set of molecules that can mount specific responses against a multitude of pathogens. Memory is a key feature of adaptive immunity, which allows organisms to respond more readily upon re-infections. However, differentiation of memory cells is still one of the least understood cell fate decisions. Here, we introduce a mathematical framework to characterize optimal strategies to store memory to maximize the utility of immune response over an organism's lifetime. We show that memory production should be actively regulated to balance between affinity and cross-reactivity of immune receptors for an effective protection against evolving pathogens. Moreover, we predict that specificity of memory should depend on the organism's lifespan, and shorter lived organisms with fewer pathogenic encounters should store more cross-reactive memory. Our framework provides a baseline to gauge the efficacy of immune memory in light of an organism's coevolutionary history with pathogens.

Introduction

Adaptive immunity in vertebrates develops during the lifetime of an organism to battle a multitude of evolving pathogens. The central actors in our adaptive immune system are diverse B- and T-cells, whose unique surface receptors are generated through genomic rearrangement, mutation, and selection (Janeway et al., 2005). The diversity of receptors allows the immune system to mount specific responses against diverse pathogens. B-cell receptors (BCRs) in particular can specialize through a process of affinity maturation, which is a form of *somatic Darwinian evolution* within an individual to enhance the affinity of BCRs to pathogens. Several rounds of somatic mutation and selection during affinity maturation can increase binding affinities of BCRs up to 10,000 fold (Victoria and Nussenzweig, 2012; Meyer-Hermann et al., 2012).

Beside receptor diversity, immune cells also differentiate and specialize to take on different roles, including plasma B-cells, which are antibody factories, effector T-cells, which can actively battle infections, or memory cells. Memory responses are highly efficient because memory cells can be reactivated faster than naive cells and can mount a more robust response to an infection (McHeyzer-Williams et al., 2000; Tangye et al., 2003; Tangye and Hodgkin, 2004; Moens et al., 2016). Memory generation is a form of cell fate decision in the immune system, which can occur at different stages of an immune response. In B-cells, activated naive cells can differentiate into antibody-secreting long-lived plasma cells, a T-cell-independent un-hypermutated memory cells, or they can initiate a germinal center (Goodnow et al., 2010). B-cells that enter germinal centers differentiate during affinity maturation into high-affinity plasma cells or T-cell-dependent long-lived memory cells that circulate in the blood for antigen surveillance; see schematic **Figure 1**.

The basis for differentiation of B-cells into memory, especially during affinity maturation, is among the least understood in cell fate decision-making in the immune system (Goodnow et al., 2010). A long-standing view was that memory is continuously produced during affinity maturation

*For correspondence:
armita@uw.edu

Competing interest: See
page 18

Funding: See page 18

Received: 22 July 2020

Accepted: 23 April 2021

Published: 28 April 2021

Reviewing editor: Arvind
Murugan, University of Chicago,
United States

© Copyright Schnaack and
Nourmohammad. This article is
distributed under the terms of
the [Creative Commons
Attribution License](#), which
permits unrestricted use and
redistribution provided that the
original author and source are
credited.

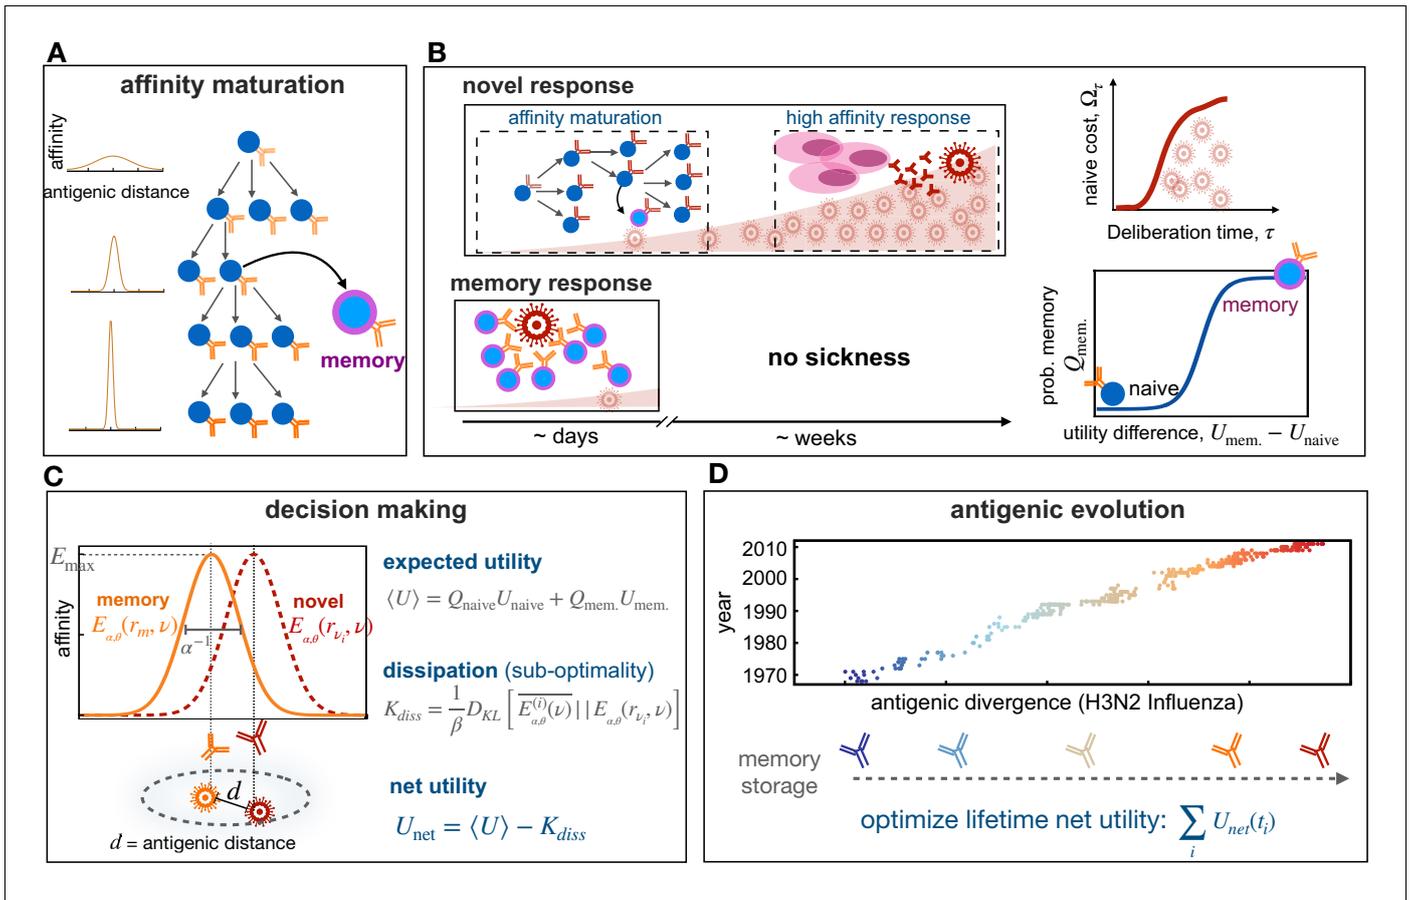


Figure 1. Immune memory or naive response upon infection. (A) Schematic shows affinity maturation in germinal centers(right), where B-cell receptors acquire mutations and undergo selection, resulting in an increase in their affinity to an antigen (from light to dark receptors), indicated by the sharpening of receptors' affinity profiles (on left). (B) Upon infection, the immune system can initiate a novel response (top) or a memory response (bottom). A novel B-cell response could involve affinity maturation to generate memory or high-affinity plasma cells (pink) that can secrete antibodies to battle the pathogen. A novel response can take 1–2 weeks, during which pathogen can replicate within a host and a patient can show symptoms from the disease (top, left). During this time, the proliferation of pathogens within a host incurs a cost associated with a naive response Ω_τ , which is a monotonic function of the deliberation time τ (top, right). If the host carries memory from a previous infection or vaccination (bottom), the immune system can robustly and rapidly activate a memory response to battle the infection. The probability to mount such memory response $Q_{\text{mem.}}$ depends non-linearly on the relative utilities of memory versus naive responses against a given infection $\Delta U = U_{\text{mem.}} - U_{\text{naive}}$ (bottom, right). (C) Affinity profile $E_{\alpha,\theta}(r_m, v) \sim \alpha \exp[-(\alpha d)^\theta]$ of a memory receptor r_m is shown in orange as a function of the distance $d = \|v^* - v\|$ in the antigenic shape space, between the receptor's cognate antigen v^* (orange) and an evolved novel target v_i (red). The affinity of a receptor decays with increasing distance between targets and its cognate antigen. The antigenic range over which a receptor is reactive inversely depends on its specificity α . The shape of the binding profile is tuned by the factor θ , here shown for $\theta = 2$. The expected binding profile $\overline{E_{\alpha,\theta}^{(i)}(v)}$ and the expected utility $\langle U \rangle$ for an immune response are weighted averages of these quantities over memory and naive responses. The Kullback-Leibler distance between the expected profile $\overline{E_{\alpha,\theta}^{(i)}(v)}$ and the profile centered around the infecting antigen $E_{\alpha,\theta}(r_{v_i}, v)$, in units of the deliberation factor β , defines the sub-optimality of a response, that is, dissipation K_{diss} (Equation 1). The net utility U_{net} measures the goodness of a decision to mount a memory vs. naive response against an infection (Equation 2). (D) Antigenic evolution of the H3N2 influenza virus is shown over 40 years along its first (most variable) antigenic dimension (data from Bedford et al., 2014). The decision of an immune system to utilize memory or to mount a novel response (B,C) is determined by the specificity α of receptors and the deliberation factor β . We characterize the optimal immune strategies (α^*, β^*) by maximizing the total net utility of immune responses against pathogens with different antigenic divergences, experienced over the lifetime of an organisms (Equation 3).

(Blink et al., 2005). Memory receptors often have lower affinity compared to plasma cells (Smith et al., 1997), and therefore, if memory B-cells were to be generated continuously it should be able to proliferate without strong affinity-dependent selection (Goodnow et al., 2010; Victora and Nussenzweig, 2012). However, recent experiments indicate that memory

differentiation is highly regulated (Paus et al., 2006; Weisel et al., 2016; Shinnakasu et al., 2016; Recaldin and Fear, 2016; Shinnakasu and Kurosaki, 2017; Viant et al., 2020), reflecting a temporal switch in germinal centers that preferentially produces memory at early stages and plasma at later stages of affinity maturation (Weisel et al., 2016). This active regulation introduces an affinity-dependent cell fate decision, leading to a preferential selection of low-affinity cells to the memory compartment. Low-affinity memory may be at a disadvantage in mounting a protective immune response since immune-pathogen recognition is largely determined by the binding affinity between an immune receptor and antigenic epitopes. On the other hand, immune-pathogen recognition is cross-reactive, which would allow memory receptors to recognize slightly evolved forms of the antigen, in response to which they were originally generated.

We propose that the program for differentiation of immune cells to memory should be viewed in light of the immune system's coevolution with pathogens. We have developed a theoretical framework that incorporates the kinetics and energetics of memory responses as ingredients of memory strategy, which we seek to optimize under various evolutionary scenarios. We propose that the hard-wired affinity-dependent regulatory measures for memory differentiation could be understood as a way to optimize the long-term utility of immune memory against evolving pathogens. Individuals encounter many distinct pathogens with varying evolutionary rates, ranging from relatively conserved pathogens like chickenpox to rapidly evolving viruses like influenza. To battle such a spectrum of evolving pathogens, we propose that an optimal immune system should store a combination of low-affinity memory with high cross-reactivity to counter evolving pathogens, and high-affinity and specific memory to counter the relatively conserved pathogens—a strategy consistent with B-cell memory, which often involves storage of both cross-reactive IgM and high-affinity IgG receptors (Shlomchik, 2018; McHeyzer-Williams et al., 2018). Lastly, we study the impact of organisms' life expectancy on their evolved memory strategies and predict that cross-reactive memory should dominate the immune response in short-lived organisms that encounter only a few pathogens.

Previous work on theoretical modeling of cellular differentiation together with experiments has been instrumental in understanding immune memory generation; for example see reviewed work in Perelson and Weisbuch, 1997; Altan-Bonnet et al., 2020. For example, mechanistic models have indicated the importance of signal integration at the cellular level (Laffleur et al., 2014) and the relevance of stochastic effects at the population level (Hawkins et al., 2007), to explain heterogeneous cell fate decisions for the generation of memory. Our statistical framework aims to characterize high-level features for an optimal memory strategy, without relying on mechanistic details of the underlying process, some of which are at least partially unknown (Bialek, 2012; Nourmohammad et al., 2013). In the case of the immune system, statistical models have provided an intuition for how an immune repertoire should be organized to optimally counter diverse pathogens (Perelson and Oster, 1979; Mayer et al., 2015; Bradde et al., 2020). In a similar fashion, optimal memory strategies identified by our model provide a baseline to gauge the performance of real immune systems in storing and utilizing memory.

Model

The efficacy of an immune response to a pathogen is determined by two key factors: (i) the affinity of immune-pathogen recognition (i.e. energetics) and (ii) the speed of response (i.e. kinetics) to neutralize an infection.

Recognition of a pathogen (or its antigenic epitope) v by an immune receptor r is mediated by the affinity of the molecular interactions $E(r, v)$ between them. We describe cross-reactive immune-pathogen recognition in an immune shape space (Perelson and Oster, 1979), where receptors located near each other in shape space can recognize similar antigens, and in the complementary space, antigens that are close to each other can be recognized by the same immune receptor (Figure 1). We express the binding affinity between a receptor r and an arbitrary target antigen v in terms of the antigenic distance $d_r(v) = \|v - v_r^*\|$ between the receptor's cognate antigen v_r^* and the target v : $E(r, v) \equiv E(d_r(v))$.

Physico-chemical constraints in protein structures can introduce a tradeoff between immune receptors' affinity and cross-reactivity. Although we lack a systematic understanding of these structural constraints, affinity-specificity tradeoffs have been reported repeatedly for B-cells and antibodies (Wedemayer et al., 1997; Frank, 2002; Li et al., 2003; Wu et al., 2017; Mishra and Mariuzza, 2018; Fernández-Quintero et al., 2020). Specifically, while affinity maturation can significantly

increase the binding affinity of a B-cell receptor, it also makes the receptor more rigid and specific to its cognate antigen (Wedemayer et al., 1997; Li et al., 2003; Mishra and Mariuzza, 2018; Fernández-Quintero et al., 2020). Broadly neutralizing antibodies (bNAbs) appear to be an exception to this rule since they have high potency and can react to a broad range of viral strains. However, it should be noted that bNAbs often react to vulnerable regions of a virus where escape mutations are very deleterious, including the CD4 binding site of HIV or the stem proteins in influenza (Mascola and Haynes, 2013; Lee and Wilson, 2015). In other words, the majority of bNAbs are not cross-reactive per se, but they are exceptionally successful in targeting conserved epitopes in otherwise diverse viral strains.

To qualitatively capture this affinity-specificity tradeoff, we use a simple functional form: We assume that the binding affinity of a receptor r to an antigen v depends on the antigenic distance $d_r(v)$ through a kernel with a specificity factor α and a shape factor θ such that, $E(r, v) \equiv E_{\alpha, \theta}(d_r(v)) \sim \alpha \exp[-(\alpha d_r(v))^\theta]$, with $\theta \geq 0$. This affinity function defines a receptor's binding profile over the space of antigens. As specificity α increases (or cross-reactivity $1/\alpha$ decays), the binding affinity profile sharpens and binding becomes more restrictive to antigens closer to the receptor's cognate antigen (Figure 1). Moreover, the absolute strength of binding to the cognate antigen (i.e. a receptor's maximum affinity) increases with specificity α , resulting in a tradeoff between affinity and cross-reactivity. The parameter θ tunes the shape of the receptor's binding profile $E_{\alpha, \theta}(d_r(v))$, resulting in a flat function (i.e. no tradeoff) for $\theta = 0$, a double-sided exponential function for $\theta = 1$, a Gaussian (bell-curve) function for $\theta = 2$, and top-hat functions for $\theta \gg 2$; see Materials and methods.

Upon encountering a pathogen, the adaptive immune system mounts a response by activating the naïve repertoire (i.e. a novel response) and/or by triggering previously stored immune receptors in the memory compartment. A memory receptor often shows a reduced affinity in interacting with an evolved form of the pathogen. Nonetheless, memory plays a central role in protecting against re-infections since even a suboptimal memory can be kinetically more efficient than a naïve response, both in B-cells (Tangye and Hodgkin, 2004) and T-cells (Whitmire et al., 2008; Martin et al., 2012). Specifically, following an infection, memory B-cells initiate cell division about 1–2 days earlier, and they are recruited to proliferate in 2–3 times larger numbers compared to the naïve population (Tangye et al., 2003; Tangye and Hodgkin, 2004; Blanchard-Rohner et al., 2009). Once recruited, however, memory and naïve cells have approximately a similar doubling time of about $t_{1/2} \approx 0.5 - 2$ days (Tangye et al., 2003; Macallan et al., 2005). Taken together, we can define an effective deliberation time $\tau \approx 1.5 - 5$ days for the naïve population to reach an activity level (i.e. a clone size) comparable to the memory; see Materials and methods and Figure 1.

The decision to mount a naïve or a memory response depends on the energetics and the kinetics of the immune machinery, including the cross-reactivity of memory to recognize evolved pathogens and the deliberation time to mount a naïve response upon infection—we refer to these choices as *memory strategies*. We expect that the biochemical machinery involved in making this decision upon an infection has been fine-tuned and selected over evolutionary time scales in order to utilize immune memory and mount an effective response against recurring pathogens. The theory of decision-making (von Neumann and Morgenstern, 1944; Ortega and Braun, 2013) enables us to characterize the response of the immune system as a rational decision-maker that chooses between two possible actions $a \in \{\text{naïve}, \text{memory}\}$ each contributing a utility U_a (Materials and methods). Specifically, the action of a rational decision-maker should follow an optimal distribution Q_a , which maximizes the expected utility while satisfying the constraints in processing new information, for example due to prior preferences (von Neumann and Morgenstern, 1944; Ortega and Braun, 2013). We assume that the immune system has no intrinsic prior for mounting a naïve or a memory response against a given pathogen. In this case, the utility U_a of an action (memory vs. naïve) determines the type of response, and rational decisions follow a maximum entropy distribution $Q_a \sim \exp[\beta U_a]$ (Jaynes, 1957), where β is the efficacy of information processing (see Materials and methods). As β increases, a rational decision-maker more readily chooses the action with the highest utility. The expected utility of the immune response to an infection is equal to the sum of the utilities of a naïve and a memory response, weighted by their respective probabilities: $\langle U \rangle = U_{\text{mem}} Q_{\text{mem}} + U_{\text{naïve}} Q_{\text{naïve}}$. If memory is effective, the utility difference between mounting a memory or a naïve response is

determined by the affinity of the interaction between the responding memory receptor r_m and the infecting antigen v : $U_{\text{mem}} - U_{\text{naive}} = E_{\alpha,\theta}(r_m, v)$; see **Figure 1** and Materials and methods for details.

The time lag (deliberation) between memory and naive response also plays a key role in the decision-making process. On the one hand, if memory is inefficient, long deliberations would allow pathogens to proliferate, incurring a larger cost Ω_τ to a host prior to activation of a novel response; this cost can be interpreted as the negative utility of naïve response $U_{\text{naive}} \equiv -\Omega_\tau$. On the other hand, a long deliberation would allow the immune system to exploit the utility of a usable memory (i.e. process information), even if the available memory has only a slight advantage over a responsive naive receptor (see Materials and methods). Indeed, for a responsive memory, the information processing factor β is equal to accumulated pathogenic load Γ_τ during the deliberation period τ , and thus, we refer to β as the *deliberation factor*.

The expected binding profile of stored memory $\overline{E_{\alpha,\theta}^{(i)}(v)}$ after i^{th} round of re-infection with an antigen v_i can be characterized as the superposition of the binding profiles following a memory or a naive response, weighted by the respective probability of each of these events (**Figure 1** and Materials and methods). Since mounting a sub-optimal memory against evolved variants of a re-infecting pathogen can still be kinetically favorable, the expected profile can deviate from the optimal profile of the cognate receptor centered around the infecting pathogen $E_{\alpha,\theta}(r_{v_i}, v)$ (**Figure 1**). This tradeoff between the kinetics and the energetics of immune response results in a *non-equilibrium decision-making* **Grau-Moya et al., 2018** by the immune system (Materials and methods). In analogy to non-equilibrium thermodynamics, we express this deviation as a dissipative cost of memory response $K_{\text{diss}}(t_i; \alpha, \theta)$ at the i^{th} round of re-infection (time point t_i), which we quantify by the Kullback-Leibler distance between the expected and the optimal binding profiles $D_{\text{KL}}(\overline{E_{\alpha,\theta}^{(i)}(v)} || E_{\alpha,\theta}(r_{v_i}, v))$, in units of the deliberation factor β (**Figure 1**),

$$K_{\text{diss}}(t_i) = \frac{1}{\beta} D_{\text{KL}}(\overline{E_{\alpha,\theta}^{(i)}(v)} || E_{\alpha,\theta}(r_{v_i}, v))$$

$$= \frac{1}{\beta} \sum_{\text{antigens:}v} \overline{E_{\alpha,\theta}^{(i)}(v)} \log \left[\frac{\overline{E_{\alpha,\theta}^{(i)}(v)}}{E_{\alpha,\theta}(r_{v_i}, v)} \right]. \tag{1}$$

An optimal memory strategy should be chosen such that it maximizes the expected utility of the immune response $\langle U \rangle$, while minimizing the dissipation cost due to the non-equilibrium response K_{diss} , over the lifetime of an organism. To infer an optimal strategy, we introduce net utility that accounts for the tradeoff between the expected utility and dissipation at a given round of infection at time point t_i ,

$$U_{\text{net}}(t_i) = \langle U(t_i) \rangle - K_{\text{diss}}(t_i) \tag{2}$$

We infer the optimal memory protocol (i.e. the optimal memory specificity α^* and deliberation factor β^*) by maximizing the total net utility of memory responses throughout the lifetime of an organism (**Figure 1**),

$$(\alpha^*, \beta^*) = \underset{\alpha, \beta}{\text{argmax}} \sum_{i:\text{infections}} U_{\text{net}}(t_i). \tag{3}$$

Results

Efficient immune memory balances specificity and speed. The extent of cross-reactivity and deliberation needed for the memory to react to pathogens should be set by the amount of pathogenic evolution and more specifically, the antigenic divergence $\hat{\delta} \equiv \sqrt{\langle \|v_i - v_{i-1}\|^2 \rangle}$ that a pathogen traces between two infections. An example of such antigenic divergence is shown in Fig. **Figure 1D** for 40 years of H3N2 Influenza evolution along its first (most variable) evolutionary dimension (**Bedford et al., 2014**). We set to find an optimal immune protocol (i.e. specificity α^* and deliberation β^*) by maximizing the net utility U_{net} of an immune system (**Equation 3**) that is trained to counter pathogens with a given antigenic divergence $\hat{\delta}$; see Fig. **Figure 1D** and Materials and methods for details on the optimization procedure.

To battle slowly evolving pathogens ($\hat{\delta} \leq 20\%$) an optimal immune system stores highly specific memory receptors, with a specificity that approaches the upper bound α_{\max} ; see **Figure 2A** and **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**. Importantly, the dependency of optimal specificity on antigenic divergence is insensitive to the cost of deliberation Ω prior to mounting a naive response (**Figure 2A**), the shape factor θ for the specificity profile (**Figure 2—figure supplement 2**), and the specificity threshold α_{\max} (**Figure 2—figure supplement 3**). For relatively conserved pathogens ($\hat{\delta} \simeq 0$), the highly specific memory (with $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\max} \simeq 1$) stored from a previous

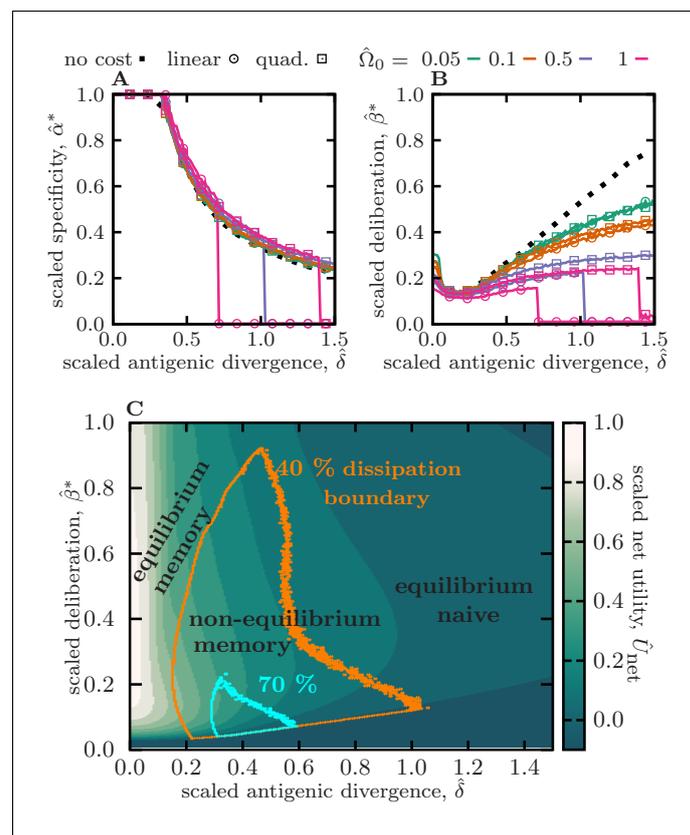


Figure 2. Optimal memory strategies against evolving pathogens. (A) and (B) show the optimal specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\max}$ and deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\max}$, scaled by their respective upper bounds, as a function of the antigenic divergence per infection, scaled by the cross-reactive range (or inverse of maximum specificity) $\hat{\delta} \equiv \delta/(\alpha_{\max}^{-1})$. Colors/markers indicate different naive cost functions for deliberation, including no-cost $\hat{\Omega} \equiv \Omega/E_{\max} = 0$, linear cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$, and quadratic cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}^2$, with varying amplitudes $\hat{\Omega}_0$. (C) The heat map shows the expected rescaled net utility $\hat{U}_{\text{net}} = U_{\text{net}}/E_{\max}$ (Equation 2) per round of infection for an immune system with an optimal specificity $\hat{\alpha}^*$, as a function of rescaled antigenic divergence $\hat{\delta}$ and deliberation factor $\hat{\beta}$. Rescaling by E_{\max} sets the magnitude of net utility to one, for a response to conserved antigens (with $\hat{\delta} = 0$) and in the limit of zero deliberation cost $\Omega \rightarrow 0$. Boundaries indicate different levels of dissipation, with orange and blue encompassing regions of $\geq 40\%$ and $\geq 70\%$ of the maximum dissipation K_{\max} , respectively. The three modes of immune response are indicated based on the magnitude of dissipation and net utility in each region: (i) equilibrium memory, (ii) non-equilibrium memory, and (iii) equilibrium naive. Simulation parameters, (A–C): $\alpha_{\max} = 4$, $\beta_{\max} = 10$, and $\theta = 2$, (C): linear deliberation cost function $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$ with $\hat{\Omega}_0 = 0.1$. Results for other shape parameters θ and specificity thresholds α_{\max} are shown in **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**, respectively.

The online version of this article includes the following figure supplement(s) for figure 2:

Figure supplement 1. Utility, dissipation, and usage of optimal memory.

Figure supplement 2. Optimal memory strategies for different specificity shape factors θ .

Figure supplement 3. Optimal memory strategies for different specificity thresholds α_{\max} .

infection still has high affinity and remains centered and close to the reinfecting pathogens. Therefore, the immune system maintains a moderate level of deliberation to exploit this efficient memory during infections. However, as antigenic divergence grows, specific memory becomes less effective against future infections and therefore, the immune system reduces the deliberation factor to allow a timely novel response, once memory becomes inefficient (**Figure 2B**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**). The magnitude of deliberation decays as the cost of deliberation Ω increases but its overall dependency on antigenic divergence remains comparable for different cost functions (shown in **Figure 2B** for zero cost, and cost functions that grow linearly and quadratically with deliberation factor β). Overall, the net utility of the stored memory in response to slowly evolving pathogens is high (**Figure 2C**, **Figure 2—figure supplement 1**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**), while its dissipation remains small $K_{\text{diss}} \simeq 0$ (**Figure 2C**, **Figure 2—figure supplement 1**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**). Therefore, in analogy to thermodynamics, we term this immune strategy with low dissipation as *equilibrium memory response*; **Figure 2C**.

To battle moderately evolving pathogens (with $\hat{\delta} \simeq 20\% - 60\%$), an optimal immune system stores cross-reactive memory (i.e. with a lower specificity $\hat{\alpha}$) that can recognize moderately evolved form of the primary antigen (**Figure 2A**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**). However, cross-reactive receptors tend to have lower affinities (Wedemayer et al., 1997; Frank, 2002), which could lead to deficient responses against antigens. Importantly, activation of energetically sub-optimal yet cross-reactive memory could be detrimental as it may hinder a stronger novel response without providing protective immunity to the host—a deficiency known as the original antigenic sin (Francis, 1960; Vatti et al., 2017). An optimal immune system can mitigate this problem by using kinetic optimization to tune the deliberation factor β in order to avoid an elongated memory engagement prior to a naive response. This optimization results in a smaller deliberation factor β (i.e. a faster naive response) compared to the scenario with slowly evolving pathogens, yet a long enough deliberation to allow the energetically suboptimal memory to react to an infection, whenever feasible (**Figure 2B**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**). With this kinetic optimization, the immune system can utilize cross-reactive memories through multiple rounds of infection (**Figure 2—figure supplement 1C**), yet with a declining efficiency and net utility as pathogens evolve away from the primary infection (**Figure 2C**, **Figure 2—figure supplement 1**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**). The prominent memory response to moderately evolving pathogens is dissipative with $K_{\text{diss}} \gg 0$ (**Figure 2C**, **Figure 2—figure supplement 1**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**), and in analogy with thermodynamics, we term this dissipative immune strategy as *non-equilibrium memory response*; **Figure 2C**.

For extremely rapidly evolving pathogens ($\hat{\delta} > 60\%$), the immune system would not be able to store an efficient memory to battle future encounters, and hence, each infection would trigger a novel naive response — the reduced net utility of memory and the decay of memory usage in this regime are shown in **Figure 2C**, **Figure 2—figure supplement 1**, **Figure 2—figure supplement 2**, **Figure 2—figure supplement 3**, respectively. Without a protective memory, a novel response is triggered to counter each infection and it matures specifically around the infecting pathogen, resulting in a non-dissipative naive-dominated immune response with $K_{\text{diss}} \simeq 0$, which we term *equilibrium naive response*; **Figure 2C**.

It should be noted that when the cost of deliberation Ω is very high, utilizing memory against pathogens with relatively high evolutionary rates becomes highly unfavorable. In this extreme case, the immune system switches into a state where it invariably mounts a novel response upon an infection (**Figure 2—figure supplement 1C**), and it assures that memory is not utilized by setting the parameters for specificity α and deliberation β to zero (**Figure 2A,B**).

Our analyses in **Figure 2** indicate that a rational decision to become a memory or a plasma cell during an immune response should depend on the affinity of a cell's receptors and it should not be a stochastic choice with a constant rate throughout affinity maturation. Indeed, cell fate decision for B-cells during affinity maturation is highly regulated and dependent on receptors' affinity (Good-Jacobson and Shlomchik, 2010; Kometani et al., 2013; Shinnakasu et al., 2016; Weisel et al., 2016; Shinnakasu and Kurosaki, 2017; Shlomchik et al., 2019). Recent experiments have demonstrated that memory generation is highly correlated with the activity of the transcription factor

Bach2 whose expression level is negatively regulated with the abundance of helper CD4+ T-cells (Kometani et al., 2013; Shinnakasu et al., 2016; Shinnakasu and Kurosaki, 2017). As the affinity of B-cell receptors increases during affinity maturation, more CD4+ T-cells are recruited to germinal centers, resulting in suppression of *Bach2* and a hence, a decline in production of memory cells (Kometani et al., 2013; Shinnakasu et al., 2016; Shinnakasu and Kurosaki, 2017). In other words, our adaptive immune system has encoded a negative feedback mechanism to store memory with intermediate affinity and cross-reactivity to suppress the production of highly specific memory, which is likely to be impotent against evolved pathogens in future infections.

A mixture memory strategy is necessary to counter pathogens with a broad range of evolutionary rates

The decision to trigger an equilibrium or a non-equilibrium memory response depends on the extent of antigenic divergence that an immune system is trained to cope with (Figure 2, Figure 2—figure supplement 1, Figure 2—figure supplement 2, Figure 2—figure supplement 3). Equilibrium memory is highly effective (i.e. it has high net utility) against relatively conserved pathogens, however, it fails to counter evolving pathogens (Figure 2C). On the other hand, cross-reactive non-equilibrium memory is more versatile and can counter a broader range of evolved pathogens but at a cost of reduced net utility in immune response; Figure 2C, Figure 2—figure supplement 1, Figure 2—figure supplement 2, Figure 2—figure supplement 3.

An optimal immune system should have memory strategies to counter pathogens with varying evolutionary rates, ranging from relatively conserved pathogens like chickenpox to rapidly evolving viruses like influenza. We use our optimization protocol to find such memory strategies that maximize the net utility of an immune system that encounters evolving pathogens with (scaled) antigenic divergences uniformly drawn from a broad range of $\hat{\delta} \in [0, 1.6]$; see Materials and methods. This optimization results in a bimodal distribution of optimal specificity for functional memory receptors $P(\alpha)$, with separated peaks corresponding to equilibrium ($\hat{\alpha} \sim 1$) and non-equilibrium ($\hat{\alpha} \sim 0.5$) memory (Figure 3, Figure 3—figure supplement 1). This result suggests that specific and cross-reactive memory

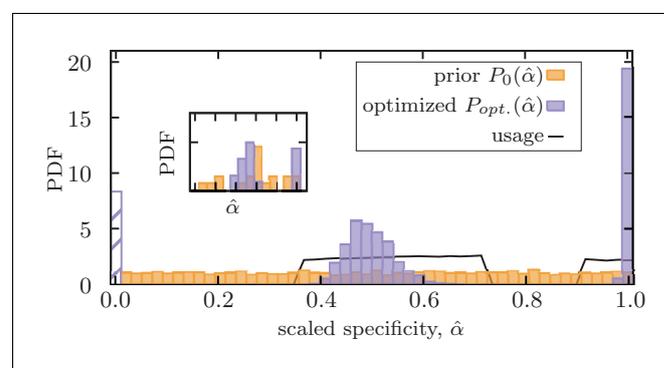


Figure 3. Mixed memory strategy against a mixture of pathogens with a broad range of evolutionary rates. Distribution of scaled optimized specificities $\hat{\alpha}^*$ for functional memory (purple) is shown for an immune system with a fixed deliberation factor $\hat{\beta} = 0.2$. A mixture strategy with a bimodal distribution of specificities $P(\hat{\alpha})$ is established to counter pathogens with a broad range of antigenic divergences. The dashed bar indicates stored memory with specificity $\alpha = 0$, which is not further used in response to infections. The solid line indicates the probability P_{usage} that a stored memory with a given specificity is utilized in future infections (Materials and methods). Optimization is done by maximizing the net utility of immune response averaged over encounters with 1000 independently evolving antigens with (scaled) antigenic divergences drawn uniformly from a range $\hat{\delta} \in (0, 1.6)$ (Materials and methods). The distribution shows the ensemble statistics of functional memory accumulated from 200 independent optimizations, each starting from a flat prior for specificities (orange). The insert shows the optimized mixture strategy for one optimization with 3000 steps. Simulation parameters: $\alpha_{\text{max}} = 4$, $\beta_{\text{max}} = 10$, and $\theta = 2$.

The online version of this article includes the following figure supplement(s) for figure 3:

Figure supplement 1. Mixed memory strategy against pathogens for different deliberation factors $\hat{\beta}$.

strategies are complementary modes of immune response that cannot substitute each other. Moreover, non-equilibrium memory tends to be flexible and moderate values of cross-reactivity $1/\hat{\alpha}$ can counter a range of antigenic divergences, without a need for fine-tuning. Therefore, upon production of memory, an optimal immune system should harvest both specific equilibrium memory and cross-reactive non-equilibrium memory, as it does not have a priori knowledge about the evolutionary rate of the infecting pathogen.

Interestingly, the adaptive immune system stores a mixture of IgM and class-switched IgG isotypes of B-cell memory that show different levels of specificity. IgM memory is an earlier product of affinity maturation with higher cross-reactivity and a lower affinity to antigens, reflecting a non-equilibrium memory that can counter evolving pathogens. On the other hand, memory from class-switched (e.g. IgG) isotype is produced during later stages of affinity maturation and is highly specific to the infecting pathogen, reflecting equilibrium memory that is effective against relatively conserved pathogens (Weisel et al., 2016). Storing a mixture of IgM and class-switched IgG memory is consistent with our recipe for optimal immune strategies to counter pathogens with a broad range of evolutionary rates.

Cross-reactive memory dominates immune response in organisms that encounter fewer pathogens over a shorter lifetime

So far, our analysis has focused on maximizing the net utility of immune response, assuming that organisms encounter many such infections throughout their lifetime. This optimization provides a recipe for optimal immune strategies in response to commonly infecting pathogens. However, the expected frequency of infections is also an important factor that can inform immune strategies. For example, imagine the extreme case that an immune system expects to encounter a pathogen at most only once during an organism's lifetime, for example in short-lived organisms. In this case, there is no benefit in keeping a memory even to counter extremely conserved pathogens, for which memory would be otherwise very beneficial.

To study the impact of infection frequency on immune strategies, we use our optimization procedure to maximize the net utility of immune response, while setting a bound on the number of infections throughout an organism's lifetime (see Materials and methods). Organisms with an *unrealistically* very short lifetime (measured in units of the number of infections) experience only a few infections, and therefore, a small (cumulative) antigenic drift from the primary infection during their lifetime $\hat{\delta}\sqrt{\text{lifetime}} \leq 1$. In this case, it would be sufficient for an optimal immune system to generate specific memory ($\hat{\alpha} \approx 1$), which can mount an effective response with only an intermediate deliberation ($\hat{\beta} \sim 0.4$) upon reinfection (Figure 4A–B), even for pathogens with a moderate evolutionary rate (Figure 4B). Organisms with moderately short lifetime experience evolutionary divergence of re-infecting antigens. In this regime, the immune system stores cross-reactive memory (smaller $\hat{\alpha}$) and uses a larger deliberation factor $\hat{\beta}$ such that this lower-affinity and often off-centered memory can mount an effective response to evolved infections (Figure 4A–B). Since the organism is relatively short-lived, such cross-reactive memory could be sufficient throughout the whole lifetime of the organism, without a need for renewal.

Organisms with long lifetimes, with pathogen encounters that surpassing the threshold c^* , expect higher re-infections with pathogens that are highly diverged from the primary infection. In this case, an optimal immune strategy switches from storing and utilizing cross-reactive memory to generating more specific memory receptors (Figure 4A). This specific memory would not hinder activation of preventive novel responses against evolved pathogens (the problem known as original antigenic sin), resulting in continual renewal of memory during organisms' lifetime. In this regime, the deliberation factor also decreases to facilitate novel responses against antigens that are not readily recognized by memory (Figure 4A–B). The increase in memory specificity from short- to long-lived organisms is more substantial for immune strategies optimized to counter relatively conserved pathogens, that is the specific equilibrium memory (Figure 2C, Figure 4A), compared to the memory against evolving pathogens, that is the cross-reactive non-equilibrium memory (Figure 2C, Figure 4B). The exact value of the transition threshold c^* depends on the expected antigenic divergence $\hat{\delta}$ during pathogenic evolution and the details of the immune machinery, and specifically the cost of deliberation $\Omega(\tau)$ due to an elevated level of pathogenic proliferation prior to a novel response (Figure 4—

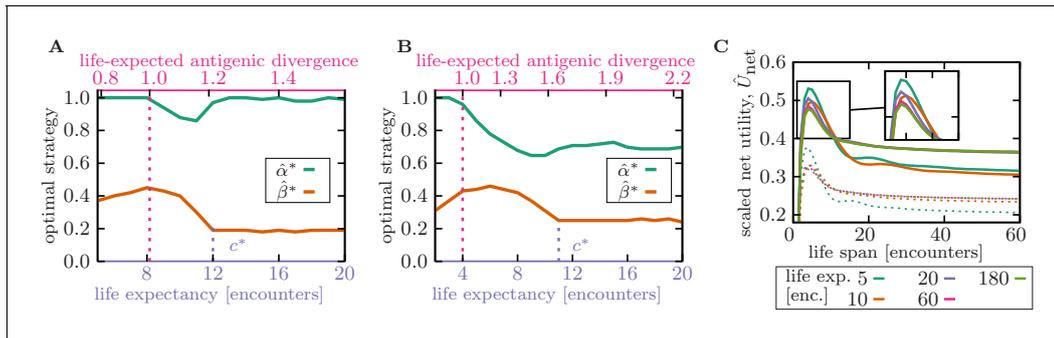


Figure 4. Life expectancy influences the specificity of optimal memory. (A,B) Memory strategies, that is, optimal rescaled specificity $\hat{\alpha}^*$ (green) and deliberation factor $\hat{\beta}^*$ (orange) are shown as a function of the organism's life expectancy (bottom axis) and the corresponding expected antigenic divergence over the organism's life-time $\hat{\delta}\sqrt{\text{lifetime}}$ (top axis). Antigenic divergence (per encounter) of the infecting pathogen is $\hat{\delta} = 0.35$ in (A) and $\hat{\delta} = 0.5$ in (B). Memory is highly specific in organisms with very short lifetimes, during which re-infections with evolved forms of a pathogen are unlikely (i.e. when life-expected antigenic divergence is smaller than 1, indicated by a dotted pink line). Memory becomes more cross-reactive with a smaller deliberation in organisms with (realistic) short lifetimes, up to a transition point c^* (indicated by a dotted purple line), after which specificity increases again. (C) Scaled net utility \hat{U}_{net} is shown as a function of organism's life span, whose immune strategies ($\hat{\alpha}^*$, $\hat{\beta}^*$) are optimized for a specified life expectancy (colors as indicated in the legend). Net utility for memory optimized against pathogens with antigenic divergence $\hat{\delta} = 0.35$ (panel A) and $\hat{\delta} = 0.5$ (panel B) are shown by full and dashed lines, respectively. Life span and life expectancy are measured in units of the number of pathogenic encounters during lifetime. Simulation parameters: linear deliberation cost function $\Omega = \Omega_0\hat{\beta}$ with an amplitude $\hat{\Omega}_0 = 0.1$, $\alpha_{max} = 4$, $\beta_{max} = 10$, and $\theta = 2$.

The online version of this article includes the following figure supplement(s) for figure 4:

Figure supplement 1. Pathogen encounter threshold to transition between cross-reactive and specific memory.

figure supplement 1). However, the qualitative trend for cross-reactivity as a function of the organism's lifetime remain consistent across a range of parameters.

The results in **Figure 4** predict that organisms with few pathogenic encounters or a shorter life-span should generate more cross-reactive and lower affinity (i.e. a naive-type) memory receptors. Indeed, consistent with our prediction, analysis of immune repertoire data indicates that sequence features of memory and naïve B-cell receptors tend to be more similar to each other in mouse compared to humans that enjoy a longer life expectancy (**Sethna et al., 2017**). Nonetheless, more comprehensive data on cross-species comparison of immune strategies is needed to test our predictions.

With the increase in human life expectancy, a pressing question is how well our immune system could cope with a larger number of pathogenic challenges that we are now encountering throughout our lifetimes? Aging has many implications for our immune machinery and the history of infections throughout lifetime leaves a complex mark on immune memory that can have long-lasting consequences (**Saule et al., 2006**), which has also been studied through theoretical modeling (**Mayer et al., 2019**). In our framework, we can study one aspect of this problem and ask how an immune strategy optimized to battle a given number of infections would perform if the organism were to live longer or equivalently, to encounter pathogens more frequently. **Figure 4C** shows that cross-reactive memory generated by an immune system optimized to counter few infections (short life expectancy) becomes highly inefficient (i.e., with a lower net utility U_{net}) as the number of encounters increases beyond the organism's expectation (long life span)—an effect that may be in part responsible for the observed decline in the efficacy of our adaptive immunity as we age.

Discussion

Memory is central to our adaptive immunity by providing a robust and preventive response to re-infecting pathogens. In the presence of continually evolving pathogens, immune memory is only beneficial if receptors can recognize evolved antigens by cross-reactivity. However, biophysical constraints can impose a trade-off between affinity and cross-reactivity of antibodies. Specifically, as receptors undergo affinity maturation, their structures become more rigid and less cross-reactive,

while affinity increases (Wedemayer *et al.*, 1997; Frank, 2002; Li *et al.*, 2003; Wu *et al.*, 2017; Mishra and Mariuzza, 2018; Fernández-Quintero *et al.*, 2020). Consistent with recent experiments (Weisel *et al.*, 2016; Shinnakasu *et al.*, 2016; Recaldin and Fear, 2016; Shinnakasu and Kurosaki, 2017; Viant *et al.*, 2020), we show that memory differentiation should be regulated to preferentially produce lower affinity receptors, which can allow cross-reactive recognition of evolved pathogens. To overcome the resulting energetic impediment of these memory receptors, we infer that the immune system should tune the kinetics of the immune response and allocate a longer deliberation time for memory to react before initiating a novel response—a feature that is also in accordance with observations (Tangye *et al.*, 2003; Tangye and Hodgkin, 2004; Blanchard-Rohner *et al.*, 2009). Co-optimizing kinetics and energetics of memory ensures an effective response against evolving pathogens, throughout an organism's lifetime.

Optimal cross-reactive immune memory provides a long-term advantage to an organism, yet it may seem energetically sub-optimal over short time scales (Figure 1). One important consequence of a sub-optimal memory response is known as original antigenic sin, where cross-reactive memory from primary infections could interfere with and suppress a protective novel response (Francis, 1960; Vatti *et al.*, 2017). The viral exposure history and the original antigenic sin may have profound consequences on protective immunity against evolving viruses (Cobey and Hensley, 2017). For example, the 2009 H1N1 pandemic triggered memory responses in individuals with childhood exposures to seasonal H1N1 (Linderman and Hensley, 2016; Li *et al.*, 2013; Hensley, 2014), which in some led to a highly focused antibody response toward the conserved epitopes of H1N1. This focus was a problem when in 2013–2014 the pandemic H1N1 acquired mutations in those epitopes (Linderman and Hensley, 2016), resulting in a disproportionate impact of infection on middle-aged individuals with pre-existing memory (Petrie *et al.*, 2016). This recent example, among others, showcases how immune history and antigenic sin can impact a population's immune response to the a rapidly evolving virus like influenza.

Composition of the immune memory coupled with the exposure history of the host should be taken into account when designing new vaccines (Cobey and Hensley, 2017). For example, current vaccine strategies against influenza use sera isolated from ferrets infected with the virus to measure the antigenic distance of circulating strains against the previous years (Smith *et al.*, 2004). However, these ferrets have no immune history for influenza and the antibodies they produce may be distinct from the immune response in the adult population with prior memory, resulting in incorrect measures of antigenic distances (Hensley, 2014). This problem has been recognized by the World Health Organization and there is now an effort to choose vaccine strains based on human serology.

The impact of immune deficiency related to the original antigenic sin can even be more pronounced due to changes in an organism's life expectancy. Importantly, we show that immune strategies optimized to benefit short-lived organisms produce highly cross-reactive memory (Figure 4). If an organism's life-expectancy increases, which is the case for humans, it would be likely for individuals to encounter evolved forms of a pathogen at antigenic distances larger than expected by their immune systems. In this case, cross-reactive memory, optimized for a shorter lifetime, could still be activated but with lower efficacy, which could suppress a protective novel response, consistent with original antigenic sin. It is therefore important to consider sub-optimality of immune strategies in the face of extensive elongation of the human lifespan as one of the plausible factors responsible for immune deficiencies brought by aging.

One characteristic of memory B-cells, which is currently missing from our model, is their ability to seed secondary germinal centers and undergo further affinity maturation upon reinfection. Evolvability of memory B-cells can allow cross-reactive memory to further specialize against evolved pathogens, without a need to start a germinal center reaction from an un-mutated naive receptor. Interestingly, different experiments suggest that the capacity of memory to re-diversify depends on various factors including the memory isotype (IgM vs. class-switch receptors), the type of antigenic target (viruses vs. others) and the extent of memory maturation (Shlomchik, 2018; McHeyzer-Williams *et al.*, 2018). Therefore, it is interesting to extend our model to study how evolvability of memory can influence its longterm utility to respond to evolving pathogens, and especially viruses.

Evolvability of memory is also relevant for characterizing the dynamics of immune response to chronic viral infections like HIV. Analyses of immune repertoires in HIV patients over multiple years of infection have shown a rapid turnover and somatic evolution of B-cell clonal lineages to counter

the evolution of the virus within hosts (*Nourmohammad et al., 2019*). It would be interesting to see how the constant pressure from the evolving HIV on a host's immune system impacts the dynamics and efficacy of immune memory over time. In addition, understanding the limits of memory re-diversification is instrumental in designing successive vaccination protocols with antigen cocktails to drive extensive affinity maturation of BCR lineages to elicit broadly neutralizing antibodies *Wang et al., 2015; Shaffer et al., 2016; Stephenson et al., 2020*—an approach that is the current hope for universal vaccines against rapidly evolving viruses like HIV.

Although mechanistically distinct from B-cells, T-cells also differentiate into effector and memory in response to infections. The T-cell response does not involve affinity maturation by hypermutations. However, competition among T-cells with varying receptor affinities acts as selection that leads to immuno-dominant responses by the high-affinity clones. Receptor affinity and the subsequent T-cell signaling determine the extent of clonal expansion and differentiation to an effector versus a memory T-cell population (*Kim and Williams, 2010*). Although it is still unresolved as how T-cell signaling determines cell fate decision, the process is known to be highly regulated (*Rutishauser et al., 2009; Roychoudhuri et al., 2016*). Notably, the transcription factor IRF4 selectively promotes expansion and differentiation of high-affinity cytotoxic T-cells into effectors. In contrast, low-affinity T-cells are lost or they could differentiate into early memory (*Man et al., 2013*). There is also accumulating evidence for the circulation of cross-reactive memory T-cells, which often result in protective immunity against evolving forms of a virus (*Greenbaum et al., 2009; Sette and Crotty, 2020*), but could also be detrimental by suppressing novel and specific responses—an effect similar to the original antigenic sin by B-cells (*Selin et al., 2004*). Taken together, there are parallels between differentiation of T-cells and B-cells to memory, and it will be interesting to investigate the advantages of storing cross-reactive (and plausibly low-affinity) T-cell memory as a strategy to counter evolving pathogens.

Materials and methods

All codes for simulations and numerical analysis can be found at: <https://github.com/StatPhysBio/ImmuneMemoryDM> (swh:1:rev:c71f7ab35ebcdd251e4a26fdf9628386fe404e86; *Schnaack, 2021*).

Numerical optimization

Numerical optimization is performed on ensembles of immune systems that encounter evolving pathogens. Recognition of an evolved pathogen at the i^{th} round of infection v_i by a memory that was stored in response to a primary infection v_0 (0^{th} round) depends on the antigenic distance $d_i = \|v_i - v_0\|$. We model pathogenic evolution as diffusion in the antigenic shape space. In this model, the expected antigenic distance between the primary infection v_0 and the evolved antigen v_i can be characterized as, $\langle d_i^2 \rangle \equiv \langle \|v_i - v_0\|^2 \rangle = \zeta^2(t_i - t_0) = i\delta^2$, where ζ is the diffusion coefficient (i.e. the evolutionary rate) and δ is the (averaged) antigenic divergence per round of infection. Importantly, this relationship does not depend on the dimensionality of the antigenic shape space, which in general, is difficult to characterize. We simulate pathogenic evolution relative to a primary infection by drawing the corresponding antigenic distance d_i of the i^{th} round of infection from a normal distribution with mean $\delta\sqrt{i}$ and standard deviation $0.05\delta\sqrt{i}$. The width of this normal distribution characterizes the fluctuations in the mean divergence between infections and reflects how the evolutionary trajectory of a pathogen samples the multi-dimensional shape space surrounding the antigen from the primary infection. Nonetheless, our results are insensitive to the exact choice of this width.

To characterize optimal specificity α^* and deliberation factor β^* (*Figure 2, Figure 3, Figure 4*), we simulate ensembles of immune systems with different immune strategies (α, β) , chosen uniformly from the range $\alpha \in [0, \alpha_{\text{max}}]$ and $\beta \in [0, \beta_{\text{max}}]$, with 500 increments in both parameters. Each immune system experiences successive rounds of infection with an evolving pathogen with a given antigenic divergence δ . During each encounter, the immune system chooses between utilizing an existing memory or initiating a novel response according to *Equation 6*. The net utility of each encounter is calculated according to *Equation 2*. We estimate the expected net utility per encounter over a lifetime of 60 total encounters and repeat this experiment across 10^5 independent ensembles to find the optimal immune strategies (α^*, β^*) with the highest net utility. As shown in *Figure 4*, simulating

up to 60 encounters is sufficient for the inference of optimal strategies in the asymptotic regime (i.e. a long lifetime).

To characterize optimal immune strategies against a mixture of pathogens with distinct levels of antigenic divergences, we define the mixture immune strategy by a set of specificities $\vec{\alpha} = \{\alpha_i\}$ (with, $i = 1, \dots, N_m$), where each α_i is a degree of specificity that a stored memory receptor can potentially have, and N_m is the number of possible specificity strategies that an immune system can choose from. The probability that an immune system with the mixture strategy $\vec{\alpha}$ recognizes a pathogen v through a memory response follows from an extension of = **Equation 6**,

$$\begin{aligned}
 P_{\text{recog.}}^{(m)}(\vec{\alpha}, v) &= 1 - \prod_{\text{specificity:}\alpha_i} \left(1 - P_{\text{recog.}}^{(m)}(r_m^{\alpha_i}, v)\right) \\
 &= 1 - \prod_{\text{specificity:}\alpha_i} e^{-E_{\theta}(r_m^{\alpha_i}, v)\Gamma(\tau)} = 1 - e^{-\sum_{\alpha_i} E_{\theta}(r_m^{\alpha_i}, v)\Gamma(\tau)} \equiv 1 - e^{-\tilde{\beta}\bar{E}_{\theta}(v)}
 \end{aligned}
 \tag{4}$$

where $\bar{E}_{\theta}(v) = \frac{1}{N_m} \sum_{r_m^{\alpha_i}} E_{\theta}(r_m^{\alpha_i}, v)$ is the expected affinity of memory (with distinct specificities) against antigen v in an immune repertoire and $\tilde{\beta} \equiv N_m \beta$ is an effective deliberation factor for all choices of specificity. It should be noted that this effective deliberation factor $\tilde{\beta}$ is an extensive quantity with respect to the number of specificity strategies that an immune system can choose from, and therefore, is comparable across immune systems with different numbers of strategies.

We set out to characterize the mixture strategy as the probability $P_{\beta}(\alpha)$ based on which an immune system with a given effective deliberation factor $\tilde{\beta}$ should store a memory receptor with specificity α , in order to optimally counter infecting pathogens with distinct antigenic divergences, drawn from a distribution $P(\delta)$. We start our optimization by defining a uniform mixture strategy, where the elements of the immune specificity vector $\vec{\alpha} = \{\alpha_i\}$ (of size $N_m = 20$), are drawn uniformly from the range $[0, \alpha_{\text{max}}]$. Each optimization step aims to improve the specificity vector $\vec{\alpha}$ to maximize the net utility (per encounter) of the mixture immune response $U_{\text{net}}(\vec{\alpha}^k)$ against 1000 independently evolving antigens whose (scaled) antigenic divergences are drawn uniformly from the range $\hat{\delta} = [0, \hat{\delta}_{\text{max}}]$. We use stochastic simulations to estimate the net utility of the mixture strategy $U_{\text{net}}(\vec{\alpha}^k)$, whereby the relative affinity of memory receptors (with varying specificities), $E_{\theta}(r_m^{\alpha_i}, v)/\bar{E}_{\theta}(v)$, determines the stochastic rate of their response to the infecting antigen v . The net utility (per encounter) of the immune response against each of the 1000 independently evolving antigens is estimated by averaging over a host's lifetime with 200 rounds of pathogenic encounters. We update the mixture strategy over 3000 steps, using local gradient ascent by sampling 100 points in the space of specificity vectors at each step to maximize net utility,

$$\vec{\alpha}^{k+1} = \vec{\alpha}^k + \epsilon \nabla U_{\text{net}}(\vec{\alpha}^k)
 \tag{5}$$

Here, k indicates the optimization step and $\epsilon = 0.1$ is a hyper-parameter for gradient ascent. We repeat the optimization process starting from 200 independently drawn initial uniform mixture strategies $\vec{\alpha}^0$ to characterize the ensemble of optimal memory strategies $P_{\beta}(\alpha)$ against pathogens with distinct antigenic divergences drawn uniformly from a given range $\hat{\delta} = [0, \hat{\delta}_{\text{max}}]$, as shown in **Figure 3**. We also characterize the probability that a stored memory with a given specificity is utilized against future infections (solid line in **Figure 3**). To do so, we test the optimized ensemble of specificities $P_{\beta}(\alpha)$ against 5000 independent pathogens with antigenic divergences drawn uniformly from the range $\hat{\delta} = [0, \hat{\delta}_{\text{max}}]$. We evaluate the usage of a memory with a given specificity α (solid line in **Figure 3**) as the conditional probability $P_{\beta}(\text{use}\alpha | \text{produce}\alpha)$ for using that memory given that it is produced (i.e. drawn from the distribution $P_{\beta}(\alpha)$).

Model of evolutionary decision-making for adaptive immune response

Kinetics of naive and memory immune response

Upon encountering a pathogen, the adaptive immune system mounts a response by activating the naïve repertoire (i.e. a novel response) and/or by triggering previously stored immune receptors in the memory compartment. A memory receptor often shows a reduced affinity in interacting with an evolved form of the pathogen. Nonetheless, memory plays a central role in protecting against re-infections since even a suboptimal memory can be kinetically more efficient than a naive response,

both in B-cells (Tangye and Hodgkin, 2004) and T-cells (Whitmire et al., 2008; Martin et al., 2012). First, memory cells are fast responders and initiate cell division about $\tau_0 \approx 1 - 2$ days before naive cells (Tangye et al., 2003; Tangye and Hodgkin, 2004; Blanchard-Rohner et al., 2009). Second, the number of memory cells that are recruited to proliferate and differentiate to effector cells is $b \approx 2 - 3$ times larger than the number of naive cells (Tangye et al., 2003; Tangye and Hodgkin, 2004). Once recruited, however, memory and naive cells have approximately a similar doubling time of about $t_{1/2} \approx 0.5 - 2$ days (Tangye et al., 2003; Macallan et al., 2005). Putting these kinetic factors together, we can define an effective deliberation time τ for the naive population to reach an activity level (i.e. a population size) comparable to the memory. Assuming an exponential growth during the early stages of memory and naïve proliferation, the deliberation time can be estimated in terms of the kinetic factors by $\tau = \tau_0 + t_{1/2} \ln b / \ln 2$ and it is within a range of $\tau \approx 1.5 - 5$ days; see Figure 1.

Energetics of immune recognition

We assume that each immune receptor r has a cognate antigen v_r^* against which it has the highest affinity. We express the binding affinity between a receptor r and an arbitrary target antigen v in terms of the antigenic distance $d_r(v) = \|v - v_r^*\|$ between the receptor's cognate antigen v_r^* and the target v : $E(r, v) \equiv E(d_r(v))$. This distance-dependent binding affinity is measured with respect to the affinity of unspecific antigen-receptor interactions, sufficient to trigger a generic naïve response.

Physico-chemical constraints in protein structures can introduce a tradeoff between immune receptors' affinity and cross-reactivity (i.e. ability to equally react to multiple targets). Prior to affinity maturation, the structure of naïve receptors is relatively flexible whereas hypermutations often reconfigure the active sites of a receptor and make them more specific so that they match their target antigens like a lock and key (Wedemayer et al., 1997; Frank, 2002). As a result, the IgM class of antibodies, which are the first line of defense in B-cell response, often have low affinities, yet they are cross-reactive and can recognize mutated forms of the same epitope. On the other hand, the high-affinity IgG class of antibodies, which are the late outcomes of affinity maturation in germinal centers, have higher affinities but bind very specifically to their cognate antigen (Frank, 2002). Broadly neutralizing antibodies (bNAbs) are exceptions to this rule since they often have high potency and can react to a broad range of viral strains. However, bNAbs often react to vulnerable regions of a virus where escape mutations are very deleterious (Mascola and Haynes, 2013). In other words, the majority of bNAbs are not cross-reactive per se, but they are exceptionally successful in targeting conserved epitopes in otherwise diverse viral strains. Nevertheless, an affinity-specificity tradeoff has been reported for a bNAb against the hemagglutinin epitope of influenza (Wu et al., 2017).

We use a simple functional form to qualitatively capture the tradeoff between cross-reactivity and affinity of antigen-receptor binding interactions: We assume that the binding affinity of a receptor r to an antigen v depends on the antigenic distance $d_r(v) = \|v - v_r^*\|$ through a kernel with a specificity factor α and a shape factor θ such that, $E(r, v) \equiv E_{\alpha, \theta}(d_r(v)) \sim \alpha \exp[-(\alpha \|v - v_r^*\|)^\theta]$, with $\theta \geq 0$. The width of this binding profile (i.e. the cross-reactivity) is set by the inverse of the specificity factor $1/\alpha$ (Figure 1), which decays as the height of the function (i.e. the maximum affinity) increases. The parameter θ tunes the shape of the receptor's binding profile $E_{\alpha, \theta}(d_r(v))$, resulting in a flat function (i.e. no tradeoff) for $\theta = 0$, a double-sided exponential function for $\theta = 1$, a Gaussian (bell-curve) function for $\theta = 2$, and top-hat functions for $\theta \gg 2$. Structural constraints and molecular features of protein receptors define a bound on the minimum cross-reactivity or equivalently, a maximum specificity α_{\max} , achievable by a receptor. Using this bound, we define rescaled specificity $\hat{\alpha} \equiv \alpha/\alpha_{\max}$ to characterize the energetics of an immune response in a dimensionless form.

Immune response to evolving pathogens

Upon primary infection (i.e. an encounter with a novel pathogen) naive immune receptors with moderate affinity are activated to develop a specific response through affinity maturation (Figure 1). Since the naive repertoire is diverse enough to contain receptors of moderate affinity against different antigens, we assume that the affinity of responsive naïve receptors, and hence, the strength of a primary immune response to be approximately the same for all pathogens. This simplification becomes less accurate as the immune system ages and the supply of effective receptors become more scarce.

Following a naive response to a primary infection and the subsequent affinity maturation, the immune system stores memory cells with an enhanced affinity to use them against future infections (Janeway et al., 2005; see Figure 1). Therefore, the cognate antigen v_m^* for a given memory receptor r_m is an epitope derived from the primary infection that led to the formation of memory, which we denote by v_0 with a subscript that indicates round of infection. Thus, the binding profile $E_{\alpha,\theta}(r_m, v)$ of the memory receptor r_m is peaked around the primary antigenic epitope $v_m^* = v_0$ (Figure 1). As pathogens evolve globally to escape the immune challenge, drugs, or vaccination, they drift away from the primary antigen in antigenic space. We model this antigenic shift as a diffusion in shape space whereby a reinfecting pathogen at the i^{th} round of infection v_i is on average at a distance $\delta = \sqrt{\langle \|v_i - v_{i-1}\|^2 \rangle}$ from the previous infection v_{i-1} . This antigenic shift is proportional to the rate of pathogen evolution ζ_v and the average time between infections $\Delta t = t_i - t_{i-1}$, such that $\delta \propto \zeta_v \sqrt{\Delta t}$. A cross-reactive memory can mount a response to an evolved antigen, yet with a reduced affinity that decays with antigenic shift; see Figure 1. It should be noted that the minimum level of receptor's cross-reactivity (or maximum specificity) $(\alpha_{\max})^{-1}$ defines a natural scale against which we can measure antigenic divergence δ and hence, form a dimensionless measure of antigenic divergence $\hat{\delta} \equiv \delta / (\alpha_{\max})^{-1}$.

Immune-pathogen recognition depends both on the binding affinity $E_{\alpha,\theta}(r, v)$ and the encounter rate $\gamma_v(t)$ between an immune receptor r and the antigen v at a given time t . The encounter rate $\gamma_v(t)$ depends on the abundance of the antigen and the immune receptor, and hence, can vary during an infection within a host. The probability that a receptor r encounters and binds to an antigen v in a short time interval $[t, t + dt]$ can be expressed by, $\rho(r, v, t)dt = \gamma_v(t)E_{\alpha,\theta}(r, v)dt$; a similar notion of encounter rate has been previously used in Mayer et al., 2016. A memory response in an individual is triggered through the recognition of an antigen by a circulating memory receptor. If no such recognition occurs during the deliberation time $\tau \approx 1.5 - 5$ days, the immune system initiates a naive response. Therefore, the probability that an antigen is recognized through a novel naive response $P_{\text{recog.}}^{(0)}$ can be expressed as the probability of the antigen not being recognized $1 - P_{\text{recog.}}^{(m)}$ by an available memory receptor r_m over the deliberation period τ ,

$$P_{\text{recog.}}^{(0)}(v) = 1 - P_{\text{recog.}}^{(m)}(r_m, v) = e^{-\int_0^\tau \rho(v, t)dt} = e^{-E_{\alpha,\theta}(r_m, v)\Gamma(v, \tau)} \tag{6}$$

where $\Gamma(v, \tau) = \int_0^\tau \gamma_v(t)dt$ is the expected number of pathogenic encounters over the deliberation time τ and depends on the accumulated pathogenic load, as pathogens proliferate in the absence of an effective memory prior to a naive response. Here, we have assumed that the affinity of the memory receptor does not change over the response time, which is a simplification since memory receptor can undergo limited affinity maturation (Shlomchik, 2018; McHeyzer-Williams et al., 2018). To further simplify, we also assume that the accumulated pathogenic load is independent of the type of the pathogen $\Gamma(v, \tau) \equiv \Gamma(\tau)$. As pathogens evolve away from the primary infector, the binding affinity $E_{\alpha,\theta}(r_m, v)$ of the stored memory receptor r_m , and hence, the probability to mount a memory response $P_{\text{recog.}}^{(m)}(r_m, v, \tau)$ decays.

The deliberation time prior to a novel response provides a window for memory to react with an antigen and mount an immune response by initiating an irreversible cascade of downstream events. Although initiation of this pathogenic recognition can be modeled as an equilibrium process, the resulting immune response is a non-equilibrium and an irreversible process, the details of which are not included in our model.

Decision-making to mount a memory or naive response

In the theory of decision-making, a rational decision-maker chooses between two possible actions $a \in \{\text{naive, memory}\}$ each contributing a utility U_a . If the decision-maker has prior preference for each action, which we denote by the prior probability distribution $Q_0(a)$, its decisions could be swayed by this knowledge. As a result, the constrained decision-maker should choose actions according to an optimized probability density $Q(a)$, which maximizes the expected utility while satisfying constraints due to the prior assumption (von Neumann and Morgenstern, 1944; Ortega and Braun, 2013),

$$Q(a) = \operatorname{argmax}_{Q(a)} \left(\sum_a U_a Q(a) - \frac{1}{\beta} D_{KL}(Q(a) || Q_0(a)) \right) \tag{7}$$

Here, $D_{KL}(Q(a) || Q_0(a)) = \sum_a Q(a) \log(Q(a)/Q_0(a))$ is the Kullback-Leibler distance between the rational distribution $Q(a)$ and the prior distribution $Q_0(a)$ and $1/\beta$ is a Lagrange multiplier that constrains the efficacy of a decision-maker to process new information and deviate from its prior assumption. The optimal solution for a rational yet constrained decision follows,

$$Q(a) = \frac{1}{Z} Q_0(a) e^{\beta U_a} \tag{8}$$

where $Z = \sum_a Q_0(a) e^{\beta U_a}$ is a normalization factor. If information processing is highly efficient (i.e. the bias factor $1/\beta \rightarrow 0$) the rational decision-maker deterministically chooses the action with the highest utility. On the other hand, if the prior is strong (i.e. $1/\beta \rightarrow \infty$), the decision-maker hardly changes its opinion and acts according to its prior belief (i.e. $Q(a) = Q_0(a)$). Moreover, if the prior distribution is uniform across actions (i.e. no prior preference), rational decision maximizes the entropy of the system (Jaynes, 1957), resulting in the probability of actions $Q(a) \sim \exp[\beta U_a]$. In our analysis, we consider the case of unbiased maximum entropy solution for decision-making. As a result the probability to utilize memory $Q_{\text{mem.}}$ or naive Q_{naive} follows,

$$Q_{\text{mem.}} = 1 - Q_{\text{naive}} = \frac{e^{\beta U_{\text{mem}}}}{e^{\beta U_{\text{mem}}} + e^{\beta U_{\text{naive}}}} \tag{9}$$

which is a sigmoidal function, dependent on the utility of each action.

A decision to mount a memory or naive response $Q(a)$ based on their respective utilities (Equation 8) should be consistent with the biophysical description of the immune response through recognition of an antigen by either of these cell types (Equation 6). By equating these two descriptions of an immune response (Equation 6, Equation 8), we can specify the utility gain associated with mounting a memory or a naïve response in terms of the biophysics and kinetics of receptor-antigen interactions,

$$\begin{aligned} Q_{\text{mem.}} = P_{\text{recog.}}^{(m)}(r_m, v) &\rightarrow \frac{e^{\beta U_{\text{mem}}}}{e^{\beta U_{\text{mem}}} + e^{\beta U_{\text{naive}}}} = 1 - e^{-E_{\alpha,\theta}(r_m, v)\Gamma(v, \tau)} \\ &\rightarrow \beta(U_{\text{mem.}} - U_{\text{naive}}) = \log \left[e^{E_{\alpha,\theta}(r_m, v)\Gamma(v, \tau)} - 1 \right] \end{aligned} \tag{10}$$

Importantly, in the regime that memory is efficient and being utilized to mount a response (i.e. a low chance for naive recognition: $P_{\text{recog.}}^{(0)} = e^{-E(v)\Gamma(v, \tau)} \ll 1$), the sigmoid form for decision to use memory (Equation 9) is dominated by an exponential factor. Therefore, the utility gain by a memory or a naïve response to an evolved antigen v_i at an antigenic distance $d_i = \|v_i - v_0\|$ from the memory receptor’s cognate antigen $v_{r_m}^* \equiv v_0$ follows (see Materials and methods),

$$\begin{aligned} U_{\text{mem}}(\|v_i - v_0\|; \alpha, \theta) &= U_{\text{naive}} + E_{\alpha,\theta}(r_m, v_i) \\ &= -\Omega(\Gamma_\tau) + E_{\alpha,\theta}(\|v_i - v_0\|) \end{aligned} \tag{11}$$

Here, we introduce the cost for deliberation $\Omega(\Gamma_\tau)$ as the negative utility of the naive response U_{naive} . Deliberation cost $\Omega(\Gamma_\tau)$ is a monotonically increasing function of the cumulative pathogen load Γ_τ and reflects the damage (cost) incurred by pathogens as they proliferate during the deliberation time τ prior to activation of the novel naive response; see Figure 1. It is important to note that the difference in the memory and the naïve utility $\Delta U = U_{\text{mem}} - U_{\text{naive}}$ determines the decision to mount either of these responses.

The same consistency criteria between decision-making (Equation 8) and cellular recognition (Equation 6) indicates that the information processing factor β in Equation 8 should be equal to the accumulated pathogenic load $\Gamma(\tau)$ during the deliberation period τ : $\beta = \Gamma(\tau)$. A longer deliberation, which on one hand leads to the accumulation of pathogens, would allow the immune system to exploit the utility of a usable memory (i.e. process information), even if the memory has only a slight advantage over a responsive naive receptor. As a result, we refer to β as the *deliberation factor*. Moreover, this analogy relates the efficacy of information processing β , which plays the role of

inverse temperature in thermodynamics, and the total accumulated pathogenic load $\Gamma(v, \tau)$, which acts as the sample size for memory receptors as they encounter and accumulate information about pathogens. Interestingly, previous work has drawn a similar correspondence between the inverse temperature in thermodynamics and the effect of sample size on statistical inference **LaMont and Wiggins, 2019**.

The deliberation factor in the immune system should be bounded $\beta \leq \beta_{\max}$ in order for the organism to survive new infections by mounting a novel response that can suppress an exponentially replicating pathogen before it overwhelms the host. Using this bound, we define rescaled deliberation factor $\hat{\beta} \equiv \beta/\beta_{\max} \leq 1$ to characterize the kinetics of an immune response in a dimensionless fashion.

It should be noted that our decision-making formalism assumes that if memory is available, it can be utilized much more efficiently and robustly than a naive response. Therefore, we do not consider scenarios where memory and naive responses are equally involved in countering an infection—a possibility that could play a role in real immune responses. Nonetheless, since such mixed responses are relatively rare, we expect that including them in our model would only result in a slightly different interpretation of the deliberation factor β and should not qualitatively impact our results.

If the immune system decides to mount a memory response against an evolved antigen v_i , the binding profile of memory against the target pathogen remains unchanged and equal to the profile $E_{\alpha,\theta}(r_{v_0}, v)$ against the primary infection v_0 . However, if the immune system mounts a naive response, a new memory receptor r_{v_i} would be generated with a binding profile $E_{\alpha,\theta}(r_{v_i}, v)$, centered around the latest infection v_i . As a result, the expected binding profile $\overline{E_{\alpha,\theta}^{(i)}}(v)$ at the i^{th} round of infection is an interpolation between the profiles associated with memory and naive response, weighted by the likelihood of each decision (**Equation 6**),

$$\overline{E_{\alpha,\theta}^{(i)}}(v) = P_{\text{recog.}}^{(m)}(r_{v_0}, v_i)E_{\alpha,\theta}(r_{v_0}, v) + P_{\text{recog.}}^{(0)}(v_i)E_{\alpha,\theta}(r_{v_i}, v) \tag{12}$$

The expected binding profile at the i^{th} round of infection $\overline{E_{\alpha,\theta}^{(i)}}(v)$ (**Equation 12**) deviates from the optimal profile centered around the infecting pathogen $E_{\alpha,\theta}(r_{v_i}, v)$ (i.e. memory profile stored following a novel response); see **Figure 1**. This deviation arises because an energetically sub-optimal memory response can still be favorable when time is of an essence and the decision has to be made on the fly with short deliberation. This tradeoff between the kinetics and the energetics of immune response results in a *non-equilibrium decision-making* **Grau-Moya et al., 2018** by the immune system. In analogy to non-equilibrium thermodynamics, we express this deviation as a dissipative cost of memory response $K_{\text{diss}}(t_i; \alpha, \theta)$ at the i^{th} round of infection (time point t_i), which we quantify by the Kullback-Leibler distance between the expected and the optimal binding profiles, in units of the deliberation factor β ,

$$\begin{aligned} K_{\text{diss}}(t_i; \alpha, \theta) &= \frac{1}{\beta} D_{\text{KL}}\left(\overline{E_{\alpha,\theta}^{(i)}}(v) \parallel E_{\alpha,\theta}(r_{v_i}, v)\right) \\ &= \frac{1}{\beta} \sum_{\text{antigens: } v} \overline{E_{\alpha,\theta}^{(i)}}(v) \log \left[\frac{\overline{E_{\alpha,\theta}^{(i)}}(v)}{E_{\alpha,\theta}(r_{v_i}, v)} \right] \end{aligned} \tag{13}$$

where we ensure that binding profiles are normalized over the space of antigens. The dissipation K_{diss} measures the sub-optimality (cost) of the mounted response through non-equilibrium decision-making and quantifies deviation from an equilibrium immune response **Grau-Moya et al., 2018**.

An optimal memory strategy should be chosen such that it maximizes the expected utility of the immune response $\langle U \rangle = U_{\text{mem}} P_{\text{recog.}}^{(m)} + U_{\text{naive}} P_{\text{recog.}}^{(0)}$, while minimizing the dissipation cost due to the non-equilibrium response K_{diss} , over the lifetime of an organism. To infer an optimal strategy, we introduce net utility U_{net} that accounts for the tradeoff between the expected utility and dissipation at a given round of infection at time point t_i ,

$$U_{\text{net}}(t_i; \alpha, \beta, \theta) = \langle U_{\alpha,\beta,\theta}(t_i) \rangle - K_{\text{diss}}(t_i; \alpha, \theta) \tag{14}$$

Net utility can be interpreted as the extracted (information theoretical) work of a rational decision-maker that acts in a limited time, and hence, is constantly kept out of equilibrium (**Grau-Moya et al., 2018**). We infer the optimal memory protocol (i.e. the optimal memory specificity α^*

and deliberation factor β^*) by maximizing the total net utility of memory responses throughout the lifetime of an organism,

$$(\alpha^*, \beta^*) = \operatorname{argmax}_{\alpha, \beta} \sum_{i: \text{infections}} U_{\text{net}}(t_i; \alpha, \beta, \theta). \quad (15)$$

While we do not model time limits to memory, we effectively model only one memory at a time. This effect is the consequence of modeling the memory as only being beneficial until a novel immune response is triggered resulting in the storage of an updated memory centered around a more recent antigen (**Figure 1**). After such an update, the old memory is no longer relevant as antigens have drifted away.

In our model, the characteristic time for a novel response (and memory update) is set by the expected antigenic divergence (**Figure 2**). Accordingly, cross-reactivity of memory is optimized so that the organism can mount effective responses against evolved forms of antigens in this window of time. However, if the lifetime of memory were to be shorter than this characteristic time of memory update, we expect the organism to store more specific memory since this memory would be utilized to counter a more limited antigenic evolution before it is lost. In other words, the shorter of either the memory lifetime or the characteristic time for memory updates determines the optimal cross-reactivity for immune memory.

Acknowledgements

We are thankful to Sarah Cobey, Sid Goyal, Lauren McGough, and Josh Plotkin for insightful comments on the manuscript. This work has been supported by the DFG grant (SFB1310) for Predictability in Evolution and the MPRG funding through the Max Planck Society. OHS also acknowledges funding from Georg-August University School of Science (GAUSS) and the Fulbright foundation.

Additional information

Competing interests

Armita Nourmohammad: Reviewing editor, *eLife*. The other author declares that no competing interests exist.

Funding

Funder	Grant reference number	Author
Deutsche Forschungsgemeinschaft	SFB1310	Armita Nourmohammad
Max Planck Society	MPRG funding	Armita Nourmohammad
University of Washington	Royalty Research Fund: A153352	Armita Nourmohammad

The funders had no role in study design, data collection and interpretation, or the decision to submit the work for publication.

Author contributions

Oskar H Schnaack, Conceptualization, Resources, Formal analysis, Validation, Investigation, Methodology, Writing - original draft, Project administration, Writing - review and editing; Armita Nourmohammad, Conceptualization, Resources, Formal analysis, Supervision, Funding acquisition, Validation, Investigation, Visualization, Methodology, Writing - original draft, Project administration, Writing - review and editing

Author ORCIDs

Armita Nourmohammad  <https://orcid.org/0000-0002-6245-3553>

Decision letter and Author responseDecision letter <https://doi.org/10.7554/eLife.61346.sa1>Author response <https://doi.org/10.7554/eLife.61346.sa2>

Additional files**Supplementary files**

- Transparent reporting form

Data availability

All codes for simulations and numerical analysis can be found at: <https://github.com/StatPhysBio/ImmuneMemoryDM> (copy archived at <https://archive.softwareheritage.org/swh:1:rev:c71f7ab35ebcdd251e4a26fdf9628386fe404e86>).

References

- Altan-Bonnet G**, Mora T, Walczak AM. 2020. Quantitative immunology for physicists. *Physics Reports* **849**:1–83. DOI: <https://doi.org/10.1016/j.physrep.2020.01.001>
- Bedford T**, Suchard MA, Lemey P, Dudas G, Gregory V, Hay AJ, McCauley JW, Russell CA, Smith DJ, Rambaut A. 2014. Integrating influenza antigenic dynamics with molecular evolution. *eLife* **3**:e01914. DOI: <https://doi.org/10.7554/eLife.01914>, PMID: 24497547
- Bialek WS**. 2012. *Biophysics : Searching for Principles*. Princeton: Princeton University Press. DOI: <https://doi.org/10.1063/PT.3.2281>
- Blanchard-Rohner G**, Pulickal AS, Jol-van der Zijde CM, Snape MD, Pollard AJ. 2009. Appearance of peripheral blood plasma cells and memory B cells in a primary and secondary immune response in humans. *Blood* **114**: 4998–5002. DOI: <https://doi.org/10.1182/blood-2009-03-211052>, PMID: 19843885
- Blink EJ**, Light A, Kallies A, Nutt SL, Hodgkin PD, Tarlinton DM. 2005. Early appearance of germinal center-derived memory B cells and plasma cells in blood after primary immunization. *The Journal of experimental medicine* **201**:545–554. DOI: <https://doi.org/10.1084/jem.20042060>, PMID: 15710653
- Bradde S**, Nourmohammad A, Goyal S, Balasubramanian V. 2020. The size of the immune repertoire of bacteria. *PNAS* **117**:201903666–201905151. DOI: <https://doi.org/10.1073/pnas.1903666117>, PMID: 32071241
- Cobey S**, Hensley SE. 2017. Immune history and influenza virus susceptibility. *Current opinion in virology* **22**:105–111. DOI: <https://doi.org/10.1016/j.coviro.2016.12.004>, PMID: 28088686
- Fernández-Quintero ML**, Loeffler JR, Bacher LM, Waibl F, Seidler CA, Liedl KR. 2020. Local and Global Rigidity Upon Antibody Affinity Maturation. *Frontiers in molecular biosciences* **7**:182. DOI: <https://doi.org/10.3389/fmolb.2020.00182>, PMID: 32850970
- Francis JRT**. 1960. On the doctrine of original antigenic sin. *Proceedings of the American Philosophical Society* **104**:572–578.
- Frank SA**. 2002. *Immunology and Evolution of Infectious Disease*. Princeton: Princeton University Press. DOI: <https://doi.org/10.2307/j.ctv15r5djw>
- Good-Jacobson KL**, Shlomchik MJ. 2010. Plasticity and heterogeneity in the generation of memory B cells and long-lived plasma cells: the influence of germinal center interactions and dynamics. *Journal of immunology* **185**:3117–3125. DOI: <https://doi.org/10.4049/jimmunol.1001155>, PMID: 20814029
- Goodnow CC**, Vinuesa CG, Randall KL, Mackay F, Brink R. 2010. Control systems and decision making for antibody production. *Nature immunology* **11**:681–688. DOI: <https://doi.org/10.1038/ni.1900>, PMID: 20644574
- Grau-Moya J**, Krüger M, Braun D. 2018. Non-Equilibrium relations for bounded rational Decision-Making in changing environments. *Entropy* **20**:1. DOI: <https://doi.org/10.3390/e20010001>
- Greenbaum JA**, Kotturi MF, Kim Y, Oseroff C, Vaughan K, Salimi N, Vita R, Ponomarenko J, Scheuermann RH, Sette A, Peters B. 2009. Pre-existing immunity against swine-origin H1N1 influenza viruses in the general human population. *PNAS* **106**:20365–20370. DOI: <https://doi.org/10.1073/pnas.0911580106>, PMID: 19918065
- Hawkins ED**, Turner ML, Dowling MR, van Gend C, Hodgkin PD. 2007. A model of immune regulation as a consequence of randomized lymphocyte division and death times. *PNAS* **104**:5032–5037. DOI: <https://doi.org/10.1073/pnas.0700026104>, PMID: 17360353
- Hensley SE**. 2014. Challenges of selecting seasonal influenza vaccine strains for humans with diverse pre-exposure histories. *Current opinion in virology* **8**:85–89. DOI: <https://doi.org/10.1016/j.coviro.2014.07.007>, PMID: 25108824
- Janeway CA**, Travers P, Walport M, Shlomchik M. 2005. *Immunobiology: The Immune System in Health and Disease*. New York: Garland Science. DOI: [https://doi.org/10.1016/S0091-6749\(95\)70025-0](https://doi.org/10.1016/S0091-6749(95)70025-0)
- Jaynes ET**. 1957. Information theory and statistical mechanics. *Physical Review* **106**:620–630. DOI: <https://doi.org/10.1103/PhysRev.106.620>

- Kim C, Williams MA. 2010. Nature and nurture: t-cell receptor-dependent and T-cell receptor-independent differentiation cues in the selection of the memory T-cell pool. *Immunology* **131**:310–317. DOI: <https://doi.org/10.1111/j.1365-2567.2010.03338.x>, PMID: 20738422
- Kometani K, Nakagawa R, Shinnakasu R, Kaji T, Rybouchkin A, Moriyama S, Furukawa K, Koseki H, Takemori T, Kurosaki T. 2013. Repression of the transcription factor Bach2 contributes to predisposition of IgG1 memory B cells toward plasma cell differentiation. *Immunity* **39**:136–147. DOI: <https://doi.org/10.1016/j.immuni.2013.06.011>, PMID: 23850379
- Laffleur B, Bardet SM, Garot A, Brousse M, Baylet A, Cogné M. 2014. Immunoglobulin genes undergo legitimate repair in human B cells not only after cis- but also frequent trans-class switch recombination. *Genes and Immunity* **15**:341–346. DOI: <https://doi.org/10.1038/gene.2014.25>, PMID: 24848929
- LaMont CH, Wiggins PA. 2019. Correspondence between thermodynamics and inference. *Physical Review E* **99**:052140. DOI: <https://doi.org/10.1103/PhysRevE.99.052140>, PMID: 31212576
- Lee PS, Wilson IA. 2015. Structural characterization of viral epitopes recognized by broadly cross-reactive antibodies. *Current topics in microbiology and immunology* **386**:323–341. DOI: https://doi.org/10.1007/82_2014_413, PMID: 25037260
- Li Y, Li H, Yang F, Smith-Gill SJ, Mariuzza RA. 2003. X-ray snapshots of the maturation of an antibody response to a protein antigen. *Nature structural biology* **10**:482–488. DOI: <https://doi.org/10.1038/nsb930>, PMID: 12740607
- Li Y, Myers JL, Bostick DL, Sullivan CB, Madara J, Linderman SL, Liu Q, Carter DM, Wrarmert J, Esposito S, Principi N, Plotkin JB, Ross TM, Ahmed R, Wilson PC, Hensley SE. 2013. Immune history shapes specificity of pandemic H1N1 influenza antibody responses. *The Journal of experimental medicine* **210**:1493–1500. DOI: <https://doi.org/10.1084/jem.20130212>, PMID: 23857983
- Linderman SL, Hensley SE. 2016. Antibodies with ‘Original Antigenic Sin’ Properties Are Valuable Components of Secondary Immune Responses to Influenza Viruses. *PLOS Pathogens* **12**:e1005806. DOI: <https://doi.org/10.1371/journal.ppat.1005806>, PMID: 27537358
- Macallan DC, Wallace DL, Zhang Y, Ghattas H, Asquith B, de Lara C, Worth A, Panayiotakopoulos G, Griffin GE, Tough DF, Beverley PC. 2005. B-cell kinetics in humans: rapid turnover of peripheral blood memory cells. *Blood* **105**:3633–3640. DOI: <https://doi.org/10.1182/blood-2004-09-3740>, PMID: 15644412
- Man K, Miasari M, Shi W, Xin A, Henstridge DC, Preston S, Pellegrini M, Belz GT, Smyth GK, Febbraio MA, Nutt SL, Kallies A. 2013. The transcription factor IRF4 is essential for TCR affinity-mediated metabolic programming and clonal expansion of T cells. *Nature immunology* **14**:1155–1165. DOI: <https://doi.org/10.1038/ni.2710>, PMID: 24056747
- Martin MD, Condotta SA, Harty JT, Badovinac VP. 2012. Population dynamics of naive and memory CD8 T cell responses after antigen stimulations in vivo. *Journal of immunology* **188**:1255–1265. DOI: <https://doi.org/10.4049/jimmunol.1101579>, PMID: 22205031
- Mascola JR, Haynes BF. 2013. HIV-1 neutralizing antibodies: understanding nature’s pathways. *Immunological Reviews* **254**:225–244. DOI: <https://doi.org/10.1111/imr.12075>, PMID: 23772623
- Mayer A, Balasubramanian V, Mora T, Walczak AM. 2015. How a well-adapted immune system is organized. *PNAS* **112**:5950–5955. DOI: <https://doi.org/10.1073/pnas.1421827112>, PMID: 25918407
- Mayer A, Mora T, Rivoire O, Walczak AM. 2016. Diversity of immune strategies explained by adaptation to pathogen statistics. *PNAS* **113**:8630–8635. DOI: <https://doi.org/10.1073/pnas.1600663113>, PMID: 27432970
- Mayer A, Balasubramanian V, Walczak AM, Mora T. 2019. How a well-adapting immune system remembers. *PNAS* **116**:8815–8823. DOI: <https://doi.org/10.1073/pnas.1812810116>, PMID: 30988203
- McHeyzer-Williams LJ, Cool M, McHeyzer-Williams MG. 2000. Antigen-specific B cell memory: expression and replenishment of a novel b220(-) memory b cell compartment. *The Journal of Experimental Medicine* **191**:1149–1166. DOI: <https://doi.org/10.1084/jem.191.7.1149>, PMID: 10748233
- McHeyzer-Williams LJ, Dufaud C, McHeyzer-Williams MG. 2018. Do Memory B Cells Form Secondary Germinal Centers? Impact of Antibody Class and Quality of Memory T-Cell Help at Recall. *Cold Spring Harbor perspectives in biology* **10**:a028878. DOI: <https://doi.org/10.1101/cshperspect.a028878>, PMID: 28320753
- Meyer-Hermann M, Mohr E, Pelletier N, Zhang Y, Victora GD, Toellner KM. 2012. A theory of germinal center B cell selection, division, and exit. *Cell reports* **2**:162–174. DOI: <https://doi.org/10.1016/j.celrep.2012.05.010>, PMID: 22840406
- Mishra AK, Mariuzza RA. 2018. Insights into the Structural Basis of Antibody Affinity Maturation from Next-Generation Sequencing. *Frontiers in immunology* **9**:117. DOI: <https://doi.org/10.3389/fimmu.2018.00117>, PMID: 29449843
- Moens L, Kane A, Tangye SG. 2016. Naïve and memory B cells exhibit distinct biochemical responses following BCR engagement. *Immunology & Cell Biology* **94**:774–786. DOI: <https://doi.org/10.1038/icb.2016.41>, PMID: 27101923
- Nourmohammad A, Held T, Lässig M. 2013. Universality and predictability in molecular quantitative genetics. *Current Opinion in Genetics & Development* **23**:684–693. DOI: <https://doi.org/10.1016/j.gde.2013.11.001>, PMID: 24291213
- Nourmohammad A, Otwinowski J, Łuksza M, Mora T, Walczak AM. 2019. Fierce selection and interference in B-Cell repertoire response to chronic HIV-1. *Molecular Biology and Evolution* **36**:2184–2194. DOI: <https://doi.org/10.1093/molbev/msz143>, PMID: 31209469
- Ortega PA, Braun DA. 2013. Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **469**:20120683. DOI: <https://doi.org/10.1098/rspa.2012.0683>

- Paus D**, Phan TG, Chan TD, Gardam S, Basten A, Brink R. 2006. Antigen recognition strength regulates the choice between extrafollicular plasma cell and germinal center B cell differentiation. *The Journal of Experimental Medicine* **203**:1081–1091. DOI: <https://doi.org/10.1084/jem.20060087>, PMID: 16606676
- Perelson AS**, Oster GF. 1979. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *Journal of Theoretical Biology* **81**:645–670. DOI: [https://doi.org/10.1016/0022-5193\(79\)90275-3](https://doi.org/10.1016/0022-5193(79)90275-3), PMID: 94141
- Perelson AS**, Weisbuch G. 1997. Immunology for physicists. *Reviews of Modern Physics* **69**:1219–1268. DOI: <https://doi.org/10.1103/RevModPhys.69.1219>
- Petrie JG**, Parkhouse K, Ohmit SE, Malosh RE, Monto AS, Hensley SE. 2016. Antibodies against the current influenza A(H1N1) Vaccine strain do not protect some individuals from infection with contemporary circulating influenza A(H1N1) Virus strains. *Journal of Infectious Diseases* **214**:1947–1951. DOI: <https://doi.org/10.1093/infdis/jiw479>, PMID: 27923954
- Recaldin T**, Fear DJ. 2016. Transcription factors regulating B cell fate in the germinal centre. *Clinical & Experimental Immunology* **183**:65–75. DOI: <https://doi.org/10.1111/cei.12702>, PMID: 26352785
- Roychoudhuri R**, Clever D, Li P, Wakabayashi Y, Quinn KM, Klebanoff CA, Ji Y, Sukumar M, Eil RL, Yu Z, Spolski R, Palmer DC, Pan JH, Patel SJ, Macallan DC, Fabozzi G, Shih HY, Kanno Y, Muto A, Zhu J, et al. 2016. BACH2 regulates CD8(+) T cell differentiation by controlling access of AP-1 factors to enhancers. *Nature Immunology* **17**:851–860. DOI: <https://doi.org/10.1038/ni.3441>, PMID: 27158840
- Rutishauser RL**, Martins GA, Kalachikov S, Chandele A, Parish IA, Meffre E, Jacob J, Calame K, Kaech SM. 2009. Transcriptional repressor Blimp-1 promotes CD8(+) T cell terminal differentiation and represses the acquisition of central memory T cell properties. *Immunity* **31**:296–308. DOI: <https://doi.org/10.1016/j.immuni.2009.05.014>, PMID: 19664941
- Saule P**, Trauet J, Dutriez V, Lekeux V, Dessaint JP, Labalette M. 2006. Accumulation of memory T cells from childhood to old age: central and effector memory cells in CD4(+) versus effector memory and terminally differentiated memory cells in CD8(+) compartment. *Mechanisms of Ageing and Development* **127**:274–281. DOI: <https://doi.org/10.1016/j.mad.2005.11.001>, PMID: 16352331
- Schnaack OH**. 2021. ImmuneMemoryDM. *Software Host*. swh:1:rev:c71f7ab35ebcdd251e4a26fdf9628386fe404e86. <https://archive.softwareheritage.org/swh:1:rev:c71f7ab35ebcdd251e4a26fdf9628386fe404e86>
- Selin LK**, Cornberg M, Brehm MA, Kim SK, Calcagno C, Ghersi D, Puzone R, Celada F, Welsh RM. 2004. CD8 memory T cells: cross-reactivity and heterologous immunity. *Seminars in Immunology* **16**:335–347. DOI: <https://doi.org/10.1016/j.smim.2004.08.014>, PMID: 15528078
- Sethna Z**, Elhanati Y, Dudgeon CS, Callan CG, Levine AJ, Mora T, Walczak AM. 2017. Insights into immune system development and function from mouse T-cell repertoires. *PNAS* **114**:2253–2258. DOI: <https://doi.org/10.1073/pnas.1700241114>, PMID: 28196891
- Sette A**, Crotty S. 2020. Pre-existing immunity to SARS-CoV-2: the knowns and unknowns. *Nature Reviews Immunology* **20**:457–458. DOI: <https://doi.org/10.1038/s41577-020-0389-z>, PMID: 32636479
- Shaffer JS**, Moore PL, Kardar M, Chakraborty AK. 2016. Optimal immunization cocktails can promote induction of broadly neutralizing Abs against highly mutable pathogens. *PNAS* **113**:E7039–E7048. DOI: <https://doi.org/10.1073/pnas.1614940113>, PMID: 27791170
- Shinnakasu R**, Inoue T, Kometani K, Moriyama S, Adachi Y, Nakayama M, Takahashi Y, Fukuyama H, Okada T, Kurosaki T. 2016. Regulated selection of germinal-center cells into the memory B cell compartment. *Nature immunology* **17**:861–869. DOI: <https://doi.org/10.1038/ni.3460>, PMID: 27158841
- Shinnakasu R**, Kurosaki T. 2017. Regulation of memory B and plasma cell differentiation. *Current opinion in immunology* **45**:126–131. DOI: <https://doi.org/10.1016/j.coi.2017.03.003>, PMID: 28359033
- Shlomchik MJ**. 2018. Do Memory B Cells Form Secondary Germinal Centers? Yes and No. *Cold Spring Harbor Perspectives in Biology* **10**:a029405. DOI: <https://doi.org/10.1101/cshperspect.a029405>, PMID: 28320754
- Shlomchik MJ**, Luo W, Weisel F. 2019. Linking signaling and selection in the germinal center. *Immunological Reviews* **288**:49–63. DOI: <https://doi.org/10.1111/imr.12744>, PMID: 30874353
- Smith KG**, Light A, Nossal GJ, Tarlinton DM. 1997. The extent of affinity maturation differs between the memory and antibody-forming cell compartments in the primary immune response. *The EMBO Journal* **16**:2996–3006. DOI: <https://doi.org/10.1093/emboj/16.11.2996>, PMID: 9214617
- Smith DJ**, Lapedes AS, de Jong JC, Bestebroer TM, Rimmelzwaan GF, Osterhaus AD, Fouchier RA. 2004. Mapping the antigenic and genetic evolution of influenza virus. *Science* **305**:371–376. DOI: <https://doi.org/10.1126/science.1097211>, PMID: 15218094
- Stephenson KE**, Wagh K, Korber B, Barouch DH. 2020. Vaccines and Broadly Neutralizing Antibodies for HIV-1 Prevention. *Annual Review of Immunology* **38**:673–703. DOI: <https://doi.org/10.1146/annurev-immunol-080219-023629>, PMID: 32340576
- 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. *Journal of Immunology* **170**:686–694. DOI: <https://doi.org/10.4049/jimmunol.170.2.686>, PMID: 12517929
- Tangye SG**, Hodgkin PD. 2004. Divide and conquer: the importance of cell division in regulating B-cell responses. *Immunology* **112**:509–520. DOI: <https://doi.org/10.1111/j.1365-2567.2004.01950.x>, PMID: 15270721
- Vatti A**, Monsalve DM, Pacheco Y, Chang C, Anaya JM, Gershwin ME. 2017. Original antigenic sin: A comprehensive review. *Journal of Autoimmunity* **83**:12–21. DOI: <https://doi.org/10.1016/j.jaut.2017.04.008>, PMID: 28479213

- Viant C**, Weymar GHJ, Escolano A, Chen S, Hartweger H, Cipolla M, Gazumyan A, Nussenzweig MC. 2020. Antibody Affinity Shapes the Choice between Memory and Germinal Center B Cell Fates. *Cell* **183**:1298–1311. DOI: <https://doi.org/10.1016/j.cell.2020.09.063>, PMID: 33125897
- Victoria GD**, Nussenzweig MC. 2012. Germinal centers. *Annual Review of Immunology* **30**:429–457. DOI: <https://doi.org/10.1146/annurev-immunol-020711-075032>, PMID: 22224772
- von Neumann J**, Morgenstern O. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press. DOI: <https://doi.org/10.1038/157172a0>
- Wang S**, Mata-Fink J, Kriegsman B, Hanson M, Irvine DJ, Eisen HN, Burton DR, Wittrup KD, Kardar M, Chakraborty AK. 2015. Manipulating the selection forces during affinity maturation to generate cross-reactive HIV antibodies. *Cell* **160**:785–797. DOI: <https://doi.org/10.1016/j.cell.2015.01.027>, PMID: 25662010
- Wedemayer GJ**, Patten PA, Wang LH, Schultz PG, Stevens RC. 1997. Structural insights into the evolution of an antibody combining site. *Science* **276**:1665–1669. DOI: <https://doi.org/10.1126/science.276.5319.1665>, PMID: 9180069
- Weisel FJ**, Zuccarino-Catania GV, Chikina M, Shlomchik MJ. 2016. A temporal switch in the germinal center determines differential output of memory B and plasma cells. *Immunity* **44**:116–130. DOI: <https://doi.org/10.1016/j.immuni.2015.12.004>, PMID: 26795247
- Whitmire JK**, Eam B, Whitton JL. 2008. Tentative T cells: memory cells are quick to respond, but slow to divide. *PLOS Pathogens* **4**:e1000041. DOI: <https://doi.org/10.1371/journal.ppat.1000041>, PMID: 18404208
- Wu NC**, Grande G, Turner HL, Ward AB, Xie J, Lerner RA, Wilson IA. 2017. In vitro evolution of an influenza broadly neutralizing antibody is modulated by hemagglutinin receptor specificity. *Nature Communications* **8**:15371. DOI: <https://doi.org/10.1038/ncomms15371>, PMID: 28504265

Chapter 4

Learning and organization of memory for evolving patterns

The following chapter consists of the preprint

O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Learning and organization of memory for evolving patterns”. In: *arXiv:2106.02186 [physics]* (June 2021). arXiv: 2106.02186

that was submitted in July 2021 and is currently under revision. The corresponding supplementary figures are shown in appendix B.

This work is part of a collaboration with Luca Peliti from the Santa Mariella Institute. He had the idea to extend the original research of Hopfield networks with a Hebbian learning rule to evolving patterns. All authors then conceptualized the project. I wrote the Julia code for simulations and data analysis, performed all analytical derivations, and produced all figures. All authors contributed to the writing and editing of the manuscript.

In this paper, we study how Hopfield networks need to adapt to follow the evolution of their target patterns. We investigate systems that can either solely solve this problem by adjusting their learning rate or further compartmentalizing their network. We find that fully compartmentalized solutions reach the maximal performance when learning evolving patterns, while all other systems only retrieve the memory of a fraction of the patterns.

In the context of this thesis, this project thus suggests that specialized memory strategies, such as immune memory, are the optimal solution for evolving environments.

Learning and organization of memory for evolving patterns

Oskar H Schnaack

*Max Planck Institute for Dynamics and Self-organization,
Am Faßberg 17, 37077 Göttingen, Germany and
Department of Physics, University of Washington,
3910 15th Ave Northeast, Seattle, WA 98195, USA*

Luca Peliti

Santa Marinella Research Institute, 00058 Santa Marinella, Italy

Armita Nourmohammad*

*Max Planck Institute for Dynamics and Self-organization, Am Faßberg 17, 37077 Göttingen, Germany
Department of Physics, University of Washington,
3910 15th Ave Northeast, Seattle, WA 98195, USA and
Fred Hutchinson Cancer Research Center, 1100 Fairview ave N, Seattle, WA 98109, USA*

(Dated: July 26, 2021)

Storing memory for molecular recognition is an efficient strategy for responding to external stimuli. Biological processes use different strategies to store memory. In the olfactory cortex, synaptic connections form when stimulated by an odor, and establish an associative distributed memory that can be retrieved upon re-exposure to the same odors. In contrast, the immune system encodes specialized memory by diverse receptors that can recognize a multitude of evolving pathogens. Despite the mechanistic differences between memory storage in the olfactory and the immune system, these processes can still be viewed as different information encoding strategies. Here, we develop analytical and numerical techniques for a generalized Hopfield network to probe the utility of distinct memory strategies against both static and dynamic (evolving) patterns. We find that while classical Hopfield networks with distributed memory can efficiently encode a memory of static patterns, they are inadequate against evolving patterns. To follow an evolving pattern, we show that a Hopfield network should use a higher learning rate, which can in turn distort the energy landscape associated with the stored memory attractors. Specifically, we observe the emergence of narrow connecting paths between memory attractors that lead to misclassification of evolving patterns. We demonstrate that compartmentalized networks with specialized subnetworks are the optimal solutions to memory storage for evolving patterns. We postulate that evolution of pathogens may be the reason for the immune system to be encoded in a focused memory, in contrast to the distributed memory used in the olfactory cortex that interacts with mixtures of static odors. Our approach offers a principled framework to study learning and memory retrieval in out-of-equilibrium dynamical systems.

I. INTRODUCTION

Storing memory for molecular recognition is an efficient strategy for sensing and response to external stimuli in biology. Apart from the cortical memory in the nervous system, molecular memory is also an integral part of the immune response, present in a broad range of organisms from the CRISPR-Cas system in bacteria [1–3] to adaptive immunity in vertebrates [4–6]. In all of these systems, an encounter with a pattern is encoded as a memory and is later retrieved and activated in response to a similar stimulus, be it a pathogenic reinfection or a re-exposure to a pheromone. Despite this high-level similarity, the immune system and the synaptic nervous system utilize vastly distinct molecular mechanisms for storage and retrieval of their memory.

Memory storage, and in particular, associative memory in the hippocampus and olfactory cortex has been

a focus of theoretical and computational studies in neuroscience [7–11]. In the case of the olfactory cortex, the input is a combinatorial pattern produced by olfactory receptors which recognize the constituent mono-molecules of a given odor. An odor signal is transmitted from the receptors to the olfactory cortex, which serves as a pattern recognition device and enables an organism to distinguish orders of magnitudes more odors compared to the number of olfactory receptors [12–14]. The synaptic connections in the cortex are formed as they are co-stimulated by a given odor pattern, thus forming a distributed associative memory that can be retrieved in future exposures [7–11, 15].

Immune memory is encoded very differently from associative memory in the nervous system. Immune receptors are extremely diverse and can specifically recognize pathogenic molecules without the need for a distributed and combinatorial encoding. In vertebrates, for example, the adaptive immune system consists of tens of billions of diverse B- and T-cells that can recognize and mount specific responses against the multitude of pathogens [5]. Immune cells activated in response to a pathogen can

* Correspondence should be addressed to Armita Nourmohammad: armita@uw.edu

differentiate into memory cells, which are long lived and can more efficiently respond upon reinfections. Unlike the distributed memory in the olfactory cortex, the receptors encoding immune memory are specialized for a given pathogen class. However, within the same class, they can recognize evolved variants of a primary pathogen, in response to which memory was originally generated [5, 16–21]

Learning and encoding of memory in nervous system has inspired the development of efficient algorithms in machine learning with artificial neural networks [22–25]. In one class of such networks, input patterns trigger interactions between encoding nodes. The ensemble of interactive nodes keeps a robust memory of the pattern since they can be simultaneously stimulated upon re-exposure and thus a stored pattern can be recovered by just knowing part of its content. This mode of memory storage resembles the co-activation of synaptic connections in a cortex. Energy-based models, such as Hopfield neural networks with Hebbian update rules [26], are among such systems, in which memory is stored as the network’s energy minima [27]. These algorithms are effective in disentangling signal from noise, which makes them highly efficient in encoding static inputs with noise. Although some specialized machine learning approaches allow for learning dynamically evolving inputs [22, 23], we still lack a general framework for learning evolving patterns, relevant for many real-life applications [23].

Inspired by the ability of the immune memory in recognizing evolving patterns (pathogens), we propose a flexible model of learning with neural networks that can interpolate between the specialized and the distributed memory strategies used by the immune and the nervous system. We formulate this problem as an optimization task to find a strategy (i.e., learning rate and network structure) that confers the highest accuracy for memory retrieval from the static and the dynamic (evolving) patterns.

In contrast to the static case, we show that a distributed memory in the style of a classical Hopfield model [27] fails to efficiently work for evolving patterns. We show that the optimal learning rate should increase with faster evolution of patterns, so that a network can follow the dynamics of the evolving patterns. This heightened learning rate tends to carve narrow connecting paths (mountain passes) between the memory attractors of a network’s energy landscape, through which patterns can equilibrate in and be associated with a wrong memory. Importantly, we demonstrate that the problem of memory retrieval for continuously evolving patterns is distinct from that of the noisy patterns [28]. Unlike noise, evolution can systematically eliminate shared features among patterns, making it difficult to retrieve a pattern from an associative memory over time. To overcome this misclassification, we demonstrate that specialized memory compartments emerge in a neural network as an optimal solution to efficiently recognize and retrieve a memory of out-of-equilibrium evolving patterns.

Our results suggest that evolution of pathogenic patterns may be one of the key reasons why the immune system encodes a focused (compartmentalized) memory, as opposed to the distributed memory used in the olfactory system, for which molecular mixtures largely present static patterns. Beyond this biological intuition, our model offers a principled analytical framework to study learning and memory generation in out-of-equilibrium dynamical systems.

II. RESULTS

A. Model of working memory for evolving patterns

To probe memory strategies against different types of stimuli, we propose a generalized energy-based model of associative memory, in which a Hopfield-like neural network is able to learn and subsequently recognize binary patterns. This neural network is characterized by an energy landscape and memory is stored as the network’s energy minima. We encode the target of recognition (stimulus) in a binary vector σ (*pattern*) with L entries: $\sigma = (\sigma_1, \dots, \sigma_L)$, with $\sigma_i = \pm 1$, $\forall i$ (Fig. 1A). To store associative memory, we define a fully connected network represented by an interaction matrix $J = (J_{i,j})$ of size $L \times L$, and use a Hopfield-like energy function (Hamiltonian) to describe pattern recognition $E_J(\sigma) = -\frac{1}{2L} \sum_{i,j} J_{i,j} \sigma_i \sigma_j \equiv -\frac{1}{2} \langle \sigma | J | \sigma \rangle$ [27] (Fig. 1C). Here, we used a short-hand notation to denote the normalized pattern vector by $|\sigma\rangle \equiv \frac{1}{\sqrt{L}} \sigma$, its transpose by $\langle \sigma |$, resulting in a normalized scalar product $\langle \sigma | \sigma' \rangle \equiv \frac{1}{L} \sum_i \sigma_i \sigma'_i$, and a matrix product $\langle \sigma | J | \sigma \rangle \equiv \frac{1}{L} \sum_{i,j} \sigma_i J_{i,j} \sigma_j$.

The network undergoes a learning process, during which different patterns are presented sequentially and in random order (Fig. 1B). As a pattern σ^α is presented, the interaction matrix J is updated according to the following Hebbian update rule [29]

$$J_{i,j} \longrightarrow J'_{i,j} = \begin{cases} (1 - \lambda) J_{i,j} + \lambda \sigma_i^\alpha \sigma_j^\alpha, & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Here λ is the *learning rate*. In this model, the memorized patterns are represented by energy minima associated with the matrix J . We consider the case where the number N of different pattern classes is below the *Hopfield capacity* of the network (i.e., $N \lesssim 0.14L$; see refs. [27, 30, 31]).

With the update rule in eq. 1, the network develops energy minima as associative memory close to each of the previously presented pattern classes σ^α ($\alpha \in \{1, \dots, N\}$) (Fig. 1C). Although the network also has minima close to the negated patterns, i.e., to $-\sigma^\alpha$, they do not play any role in what follows. To find an associative memory we let a presented pattern σ^α equilibrate in the energy landscape, whereby we accept spin-flips

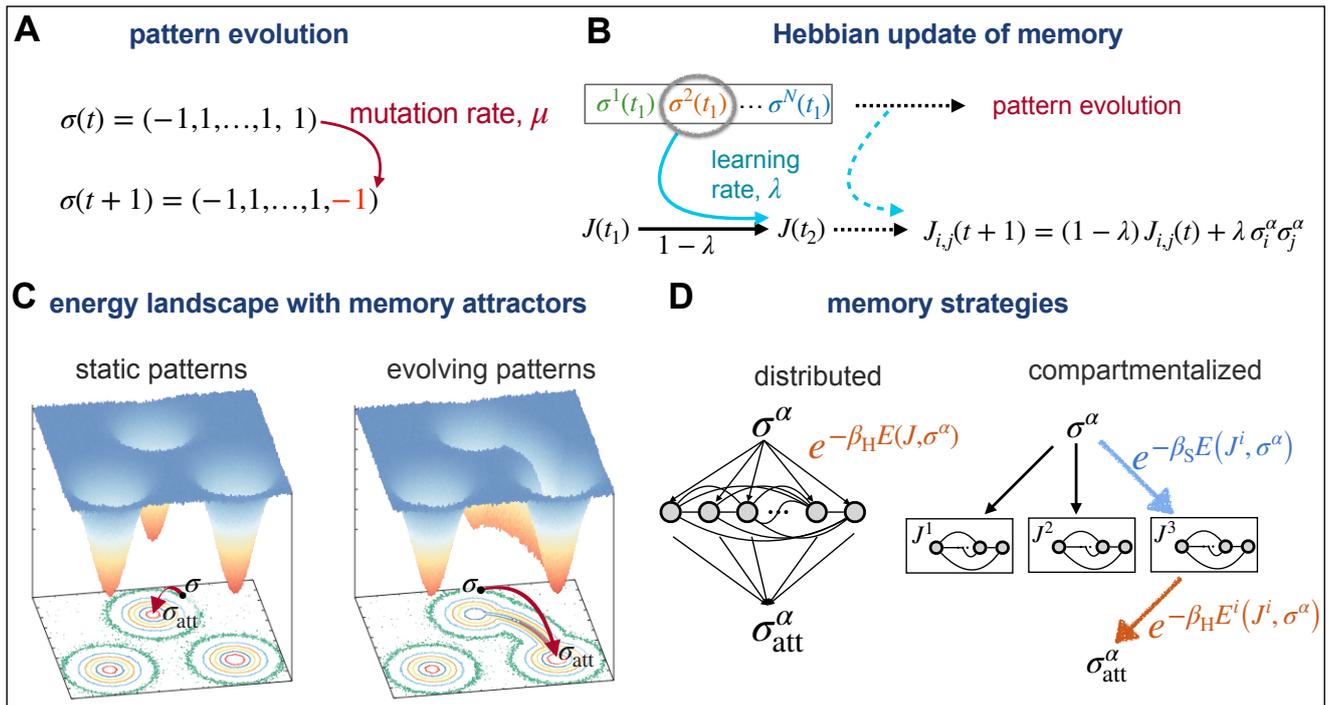


FIG. 1. Model of Working memory for evolving patterns. (A) The targets of recognition are encoded by binary vectors $\{\sigma\}$ of length L . Patterns can evolve over time with a mutation rate μ , denoting the fraction of spin-flips in a pattern per network update event. (B) Hebbian learning rule is shown for network J , which is presented a set of N patterns $\{\sigma^\alpha\}$ (colors) over time. At each step, one pattern σ^α is randomly presented to the network and the network is updated with learning rate λ (eq. 1). (C) The energy landscape for networks with distributed memory with optimal learning rate for static (left) and evolving (right) patterns are shown. The equipotential lines are shown in the bottom 2D plane. The energy minima correspond to memory attractors. For static patterns (left), equilibration in the network's energy landscape drives a pattern towards its associated memory attractor, resulting in an accurate reconstruction of the pattern. For evolving patterns (right), the heightened optimal learning rate of the network results in the emergence of connecting paths (mountain passes) between the energy minima. The equilibration process can drive a pattern through a mountain pass towards a wrong memory attractors, resulting in pattern misclassification. (D) A network with distributed memory (left) is compared to a specialized network with multiple compartments (right). To find an associative memory, a presented pattern σ^α with energy $E(J, \sigma^\alpha)$ in network J equilibrates with inverse temperature β_H in the network's energy landscape and falls into an energy attractor $\sigma_{\text{att}}^\alpha$. Memory retrieval is a two-step process in a compartmentalized network (right): First, the sub-network J^i is chosen with a probability $P_i \sim \exp[-\beta_S E^i(J^i, \sigma^\alpha)]$, where β_S is the inverse temperature for this decision. Second, the pattern equilibrates within the sub-network and falls into an energy attractor $\sigma_{\text{att}}^\alpha$.

$\sigma^\alpha \rightarrow \tilde{\sigma}^\alpha$ with a probability $\min(1, e^{-\beta_H(E_J(\tilde{\sigma}) - E_J(\sigma))})$, where β_H is the inverse equilibration (Hopfield) temperature (Appendix A). In the low temperature regime (i.e., high β_H), equilibration in networks with working memory drives a presented pattern σ^α towards a similar attractor $\sigma_{\text{att}}^\alpha$, reflecting the memory associated with the corresponding energy minimum (Fig. 1C). This similarity is measured by the *overlap* $q^\alpha \equiv \langle \sigma_{\text{att}}^\alpha | \sigma^\alpha \rangle$ and determines the accuracy of the associative memory.

Unlike the classical cases of pattern recognition by Hopfield networks, we assume that patterns can evolve over time with a rate μ that reflects the average number of spin-flips in a given pattern per network's update event (Fig. 1A). Therefore, the expected number of spin-flips in a given pattern between two encounters is $\mu_{\text{eff}} = N\mu$, as two successive encounters of the same pattern are on

average separated by $N - 1$ encounters (and updates) of the network with the other patterns. We work in the regime where the mutation rate μ is small enough such that the evolved patterns stemming from a common ancestor $\sigma^\alpha(t_0)$ at time t_0 (i.e., the members of the *class* α) remain more similar to each other than to members of the other classes (i.e., $\mu N L \ll L/2$).

The special case of static patterns ($\mu_{\text{eff}} = 0$) can reflect distinct odor molecules, for which associative memory is stored in the olfactory cortex. On the other hand, the distinct pattern classes in the dynamic case ($\mu_{\text{eff}} > 0$) can be attributed to different types of evolving pathogens (e.g., influenza, HIV, etc), and the patterns within a class as different variants of a given pathogen. In our model, we will use the mutation rate as an order parameter to characterize the different phases of memory strategies in

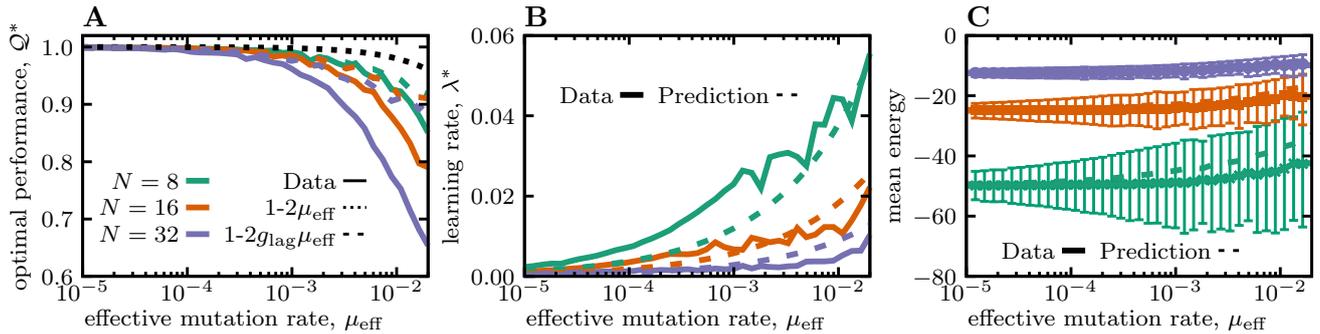


FIG. 2. **Reduced performance of Hopfield networks in retrieving memory of evolving patterns.** (A) The optimal performance of a network $Q^* \equiv Q(\lambda^*)$ (eq. 2) is shown as a function of the effective mutation rate $\mu_{\text{eff}} = N\mu$. The solid lines show the simulation results for networks encountering different number of patterns (colors). The black dotted line shows the naïve expectation for the performance solely based on the evolutionary divergence of the patterns $Q_0 \approx 1 - 2\mu_{\text{eff}}$, and the colored dashed lines show the expected performance after accounting for the memory lag g_{lag} , $Q_{\text{lag}} \approx 1 - 2g_{\text{lag}}\mu_{\text{eff}}$; see Fig. S2 for more details. (B) The optimal learning rate λ^* is shown as a function of effective mutation rate. The solid lines are the numerical estimates and dashed lines show the theoretical predictions (eq. 4). (C) The mean energy obtained by simulations for ordered patterns (solid lines) and the analytical approximation (eq. 3) for randomly ordered patterns (dotted lines) are shown. Error bars show standard error from the independent realizations (Appendix A). The color code for the number of presented patterns is consistent across panels, and the length of patterns is set to $L = 800$.

biological systems.

B. Optimal learning rate for evolving patterns

In the classical Hopfield model ($\mu_{\text{eff}} = 0$) the learning rate λ is set to very small values for the network to efficiently learn the patterns [29]. For evolving patterns, the learning rate should be tuned so the network can efficiently update the memory retained from the prior learning steps. At each encounter, the overlap $q^\alpha(t; \lambda) = \langle \sigma_{\text{att}}^\alpha(t; \lambda) | \sigma^\alpha(t) \rangle$ between a pattern $\sigma^\alpha(t)$ and the corresponding attractor for the associated energy minimum $\sigma_{\text{att}}^\alpha(t; \lambda)$ determines the accuracy of pattern recognition; the parameter λ explicitly indicates the dependency of the network's energy landscape on the learning rate. We declare pattern recognition as successful if the accuracy of reconstruction (overlap) is larger than a set threshold $q^\alpha(t) \geq 0.8$, but our results are insensitive to the exact value of this threshold (Appendix C and Fig. S4). We define a network's performance as the asymptotic accuracy of its associative memory averaged over the ensemble of pattern classes (Fig. 2A),

$$Q(\lambda) \equiv \mathbb{E}[q^\alpha(t; \lambda)] \simeq \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \frac{1}{N} \sum_{\alpha=1}^N \langle \sigma_{\text{att}}^\alpha(t) | \sigma^\alpha(t) \rangle \quad (2)$$

The expectation $\mathbb{E}[\cdot]$ is an empirical average over the ensemble of presented pattern classes over time, which in the stationary state approaches the asymptotic average of the argument. The optimal learning rate is determined by maximizing the network's performance, $\lambda^* = \text{argmax}_\lambda Q(\lambda)$.

The optimal learning rate increases with growing mutation rate so that a network can follow the evolving pat-

terns (Fig. 2B). Although it is difficult to analytically calculate the optimal learning rate, we can use an approximate approach and find the learning rate that minimizes the expected energy of the patterns $\mathbb{E}[E_{\lambda, \rho}(J, \sigma)]$, assuming that patterns are shown to the network at a fixed order (Appendix B). In this case, the expected energy is given by

$$\mathbb{E}[E_{\lambda, \rho}(J, \sigma)] = \frac{L-1}{2} \times \frac{\lambda}{1-\lambda} \times \frac{1}{\rho^{-2N}(1-\lambda)^{-N} - 1} \quad (3)$$

where $\rho^N \equiv (1-2\mu)^N \approx 1-2\mu_{\text{eff}}$ is the upper bound for the overlap between a pattern and its evolved form, when separated by the other $N-1$ patterns that are shown in between. The expected energy grows slowly with increasing mutation rate (i.e., with decreasing overlap q), and the approximation in eq. 3 agrees very well with the numerical estimates for the scenario where patterns are shown in a random order (Fig. 2C). In the regime where memory can still be associated with the evolved patterns ($\mu_{\text{eff}} \ll 0.5$), minimization of the expected energy (eq. 3) results in an optimal learning rate,

$$\lambda^*(\mu) = \sqrt{8\mu/(N-1)}, \quad (4)$$

that scales with the square root of the mutation rate. Notably, this theoretical approximation agrees well with the numerical estimates (Fig. 2B).

C. Reduced accuracy of distributed associative memory against evolving patterns

Despite using an optimized learning rate, a network's accuracy in pattern retrieval $Q(\lambda)$ decays much faster than the naïve expectation solely based on the evolutionary divergence of patterns between two encounters

with a given class (i.e., $\mathcal{Q}_0 = (1 - 2\mu)^N \approx 1 - 2\mu_{\text{eff}}$); see Fig. 2A. It should be noted that a classical Hopfield network with a small learning rate can accurately retrieve the memory of noisy static patterns by relying on the shared features of the patterns within a class [28] (Fig. S1). However, continuous evolution can systematically eliminate the shared features among patterns of a given class, resulting in a significant reduction in the accuracy of memory retrieval for evolving patterns. There are two reasons for such reduced accuracy: (i) the lag in the network’s memory, and (ii) misclassification of presented patterns.

The memory attractors associated with a given pattern class can lag behind and only reflect the older patterns presented prior to the most recent encounter of the network with the specified class. We characterize this lag g_{lag} by identifying a previous version of the pattern that has the maximum overlap with the network’s energy landscape at a given time t : $g_{\text{lag}} = \text{argmax}_{g \geq 0} \mathbb{E} [\langle \sigma(t - gN) | J(t) | \sigma(t - gN) \rangle]$ (Appendix B.2). g_{lag} measures time in units of N (i.e., the effective separation time of pattern of the same class). An increase in the optimal learning rate reduces the time lag and enables the network to follow the evolving patterns more closely (Fig. S2). The accuracy of the memory subject to such a lag decays as $\mathcal{Q}_{\text{lag}} = \rho^{g_{\text{lag}}N} \approx 1 - 2g_{\text{lag}}\mu_{\text{eff}}$, which is faster than the naïve expectation (i.e., $1 - 2\mu_{\text{eff}}$); see Fig. 2A. This memory lag explains the loss of performance for patterns that are still reconstructed by the network’s memory attractors (i.e., those with $q^\alpha > 0.8$; Fig. S2A). However, the overall performance of the network $\mathcal{Q}(\lambda)$ remains lower than the expectation obtained by taking into account this time lag (Fig. 2A)—a discrepancy that leads us to the second root of reduction in accuracy, i.e., pattern misclassification.

As the learning rate increases, the structure of the network’s energy landscape changes. In particular, we see that with large learning rates a few narrow paths emerge between the memory attractors of networks (Fig. 1C). As a result, the equilibration process for pattern retrieval can drive a presented pattern through the connecting paths towards a wrong memory attractor (i.e., one with a small overlap $\langle \sigma_{\text{att}} | \sigma \rangle$), which leads to pattern misclassification (Fig. S3A and Fig. S4A,C). These paths are narrow as there are only a few possible spin-flips (mutations) that can drive a pattern from one valley to another during equilibration (Fig. S4B,D and Fig. S5A,C). In other words, a large learning rate carves narrow mountain passes in the network’s energy landscape (Fig. 1C), resulting in a growing fraction of patterns to be misclassified. Interestingly, pattern misclassification occurs even in the absence of mutations for networks with an increased learning rate (Fig. S3A). This suggests that mutations only indirectly contribute to the misclassification of memory, as they necessitate a larger learning rate for the networks to optimally operate, which in turn results in the emergence of mountain passes in the energy landscape.

To understand the memory misclassification, particularly for patterns with moderately low (i.e., non-random) energy (Fig. 2C), we use spectral decomposition to characterize the relative positioning of patterns in the energy landscape (see Appendix C). The vector representing each pattern $|\sigma\rangle$ can be expressed in terms of the network’s eigenvectors $\{\Phi^i\}$, $|\sigma\rangle = \sum_i m_i |\Phi^i\rangle$, where the overlap $m_i \equiv \langle \Phi^i | \sigma \rangle$ is the i^{th} component of the pattern in the network’s coordinate system. During equilibration, we flip individual spins in a pattern and accept the flips based on their contribution to the recognition energy. We can view these spin-flips as rotations of the pattern in the space spanned by the eigenvectors of the network. Stability of a pattern depends on whether these rotations could carry the pattern from its original subspace over to an alternative region associated with a different energy minimum.

There are two key factors that modulate the stability of a pattern in a network. The dimensionality of the subspace in which a pattern resides, i.e., support of a pattern by the network’s eigenvectors, is one of the key determining factors for pattern stability. We quantify the support of a pattern σ using the participation ratio $\pi(\sigma) = (\sum_i m_i^2)^2 / \sum_i m_i^4$ [32, 33] that counts the number of eigenvectors that substantially overlap with the pattern. A small support $\pi(\sigma) \approx 1$ indicates that the pattern is spanned by only a few eigenvectors and is restricted to a small sub-space, whereas a large support indicates that the pattern is orthogonal to only a few eigenvectors. As the learning rate increases, patterns lie in lower dimensional sub-spaces supported by only a few eigenvectors (Fig. S5B,D). This effect is exacerbated by the fact the energy gap between the eigenstates of the network also broaden with increasing learning rate (Fig. S6). The combination of a smaller support for patterns and a larger energy gap in networks with increased learning rate leads to the destabilization of patterns by enabling the spin-flips during equilibration to drive a pattern from one subspace to another, through the mountain passes carved within the landscape; see Appendix C and Fig. S7 for the exact analytical criteria for pattern stability.

D. Compartmentalized learning and memory storage

Hopfield-like networks can store accurate associative memory for static patterns. However, these networks fail to perform and store retrievable associative memory for evolving patterns (e.g., pathogens), even when learning is done at an optimal rate (Fig. 2). To overcome this difficulty, we propose to store memory in compartmentalized networks, with C sub-networks of size L_c (i.e., the number of nodes in a sub-network). Each compartment (sub-network) can store a few of the total of N pattern classes without interference from the other compartments (Fig. 1D).

Recognition of a pattern σ in a compartmentalized network involves a two step process (Fig. 1D): First, we choose a sub-network J^i associated with compartment i with a probability $P_i = \exp[-\beta_S E(J^i, \sigma)]/\mathcal{N}$, where β_S is the inverse temperature for this decision and \mathcal{N} is the normalizing factor. Once the compartment is chosen, we follow the recipe for pattern retrieval in the energy landscape of the associated sub-network, whereby a pattern equilibrates into a memory attractor.

On average, each compartment stores a memory for $N_c = N/C$ pattern classes. To keep the networks with different number of compartments C comparable, we scale the size of each compartment L_c to keep $C \times L_c = \text{const.}$, which keeps the (Hopfield) capacity of the network $\alpha = N_c/L_c$ invariant under compartmentalization. Moreover, the mutation rate experienced by each sub-network scales with the number of compartments $\mu_c = C\mu$, which keeps the effective mutation rate $\mu_{\text{eff}} = N_c\mu_c$ invariant under compartmentalization. As a result, the optimal learning rate (eq. 4) scales with the number of compartments C as, $\lambda_c^* = \sqrt{8\mu_c/(N_c - 1)} \approx C\lambda_1^*$. However, since updates are restricted to sub-networks of size L_c at a time, the expected amount of updates within a network $L_c\lambda_c$ remains invariant under compartmentalization. Lastly, since the change in energy by a single spin-flip scales as $\Delta E \sim 1/L_c$, we introduce the scaled Hopfield temperature $\beta_{H_c} \equiv C\beta_H$ to make the equilibration process comparable across networks with different number of compartments. No such scaling is necessary for β_S .

By restricting the networks to satisfy the aforementioned scaling relationships, we are left with two independent variables, i.e., (i) the number of compartments C , and (ii) the learning rate λ_c , which define a memory strategy $\{C, \lambda_c\}$. A memory strategy can be then optimized to achieve a maximum accuracy for retrieving an associative memory for evolving patterns with a given effective mutation rate μ_{eff} .

E. Phases of learning and memory production

Pattern retrieval can be stochastic due to the noise in choosing the right compartment from the C sub-networks (tuned by the inverse temperature β_S), or the noise in equilibrating into the right memory attractor in the energy landscape of the chosen sub-network (tuned by the Hopfield inverse temperature β_{H_c}). We use mutual information to quantify the accuracy of pattern-compartment association, where larger values indicate a more accurate association; see Appendix A and Fig. 4. The optimal performance Q^* determines the overall accuracy of memory retrieval, which depends on both finding the right compartment and equilibrating into the correct memory attractor. The amplitudes of intra-versus inter-compartment stochasticity determine the optimal strategy $\{C^*, \lambda_c^*\}$ used for learning and retrieval of patterns with a specified mutation rate. Varying the

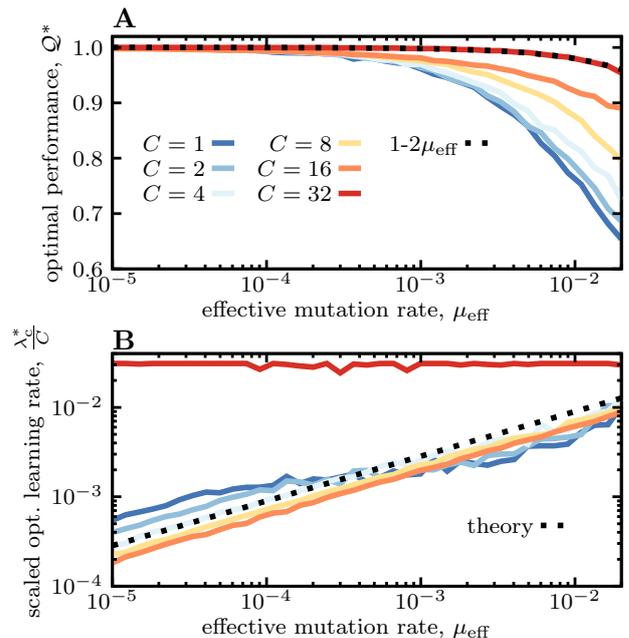


FIG. 3. **Compartmentalized memory storage.** (A) The optimal performance is shown as a function of the effective mutation rate (similar to Fig. 2A) for networks with different number of compartments C (colors), ranging from a network with distributed memory $C = 1$ (blue) to a 1-to-1 compartmentalized network $C = N$ (red). (B) The optimal (scaled) learning rate λ_c/C is shown as a function of the effective mutation rate for networks with different number of compartments (colors according to (A)). Full lines show the numerical estimates and the dashed line is from the analytical approximation, $\lambda_c^* = \sqrt{8\mu_c/(N_c - 1)} \approx C\lambda_1^*$. The scaled learning rates collapse on the analytical approximation for all networks except for the 1-to-1 compartmentalized network (red), where the maximal learning rate $\lambda \approx 1$ is used and each compartment is fully updated upon an encounter with a new version of a pattern. The number of presented patterns is set to $N = 32$. We keep $L \times C = \text{const.}$, with $L = 800$ used for the network with $C = 1$.

corresponding inverse temperatures (β_{H_c}, β_S) results in three distinct phases of optimal memory storage.

i. Small intra- and inter-compartment noise ($\beta_{H_c} \gg 1$, $\beta_S \gg 1$). In this regime, neither the compartment choice nor the pattern retrieval within a compartment are subject to strong noise. As a result, networks are functional with working memory and the optimal strategies can achieve the highest overall performance. For small mutation rates, we see that all networks perform equally well and can achieve almost perfect performance, irrespective of their number of compartments (Figs. 3A, 4A,B). As the mutation rate increases, networks with a larger number of compartments show a more favorable performance, and the 1-to-1 specialized network, in which each pattern is stored in a separate compartment (i.e., $N = C$), reaches the

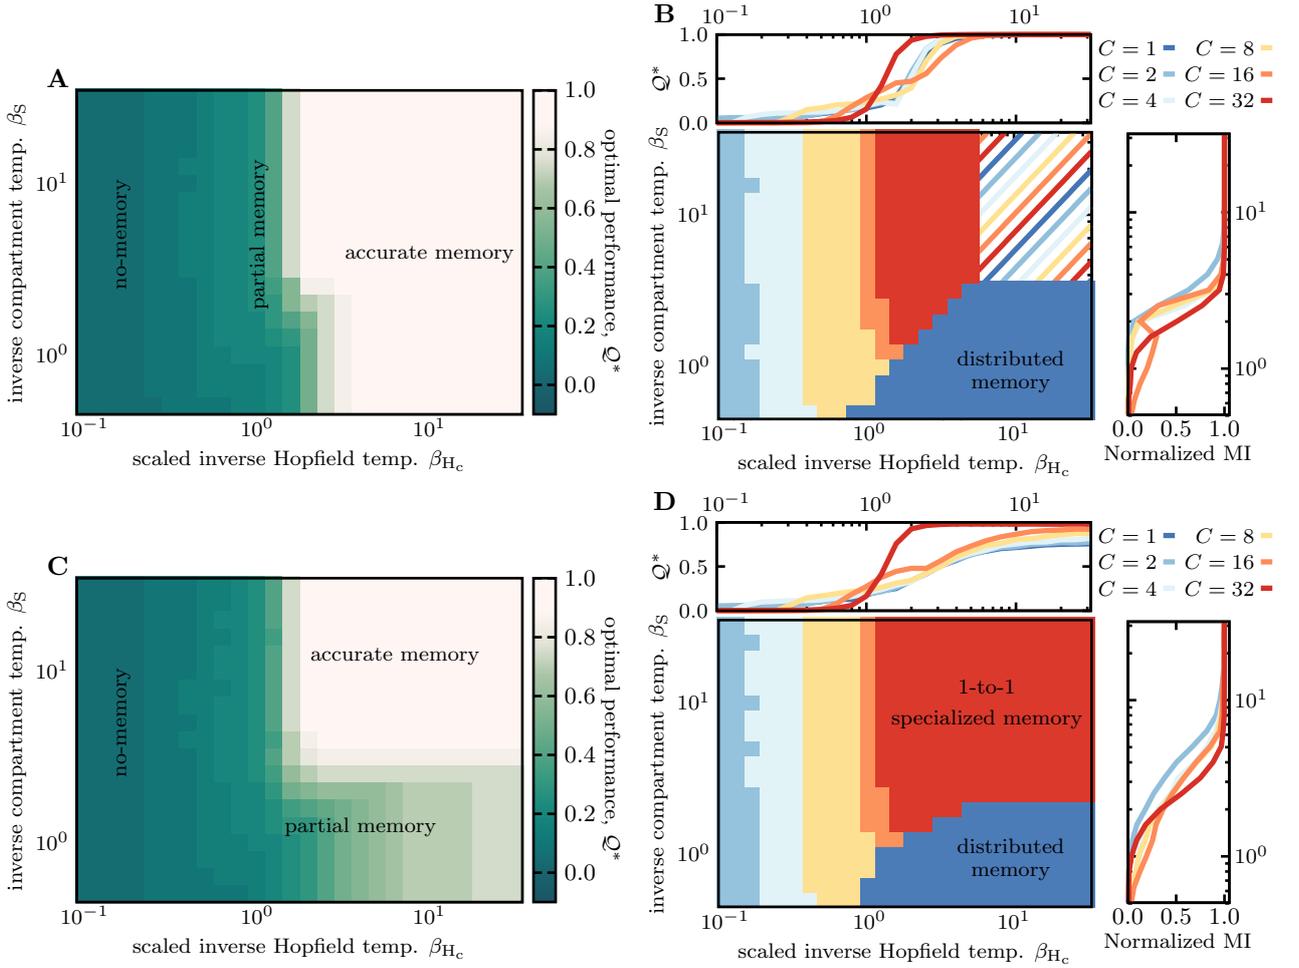


FIG. 4. Phases of learning and memory production. Different optimal memory strategies are shown. **(A)** The heatmap shows the optimal memory performance Q^* as a function of the (scaled) Hopfield inverse temperature $\beta_{H_c} = \beta_H \cdot C$ and the inverse temperature associated with compartmentalization β_S for networks learning and retrieving a memory of static patterns ($\mu = 0$); colors indicated in the color bar. The optimal performance is achieved by using the optimal strategy (i.e., learning rate λ_c^* and the number of compartments c^*) for networks at each value of β_{H_c} and β_S . The three phases of accurate, partial, and no-memory are indicated. **(B)** The heatmap shows the memory strategies for the optimal number of compartments (colors as in the legend) corresponding to the memory performance shown in (A). We limit the optimization to the possible number of compartments indicated in the legend to keep N/C an integer. The dashed region corresponds to the case where all strategies perform equally well. Regions of distributed memory ($C = 1$) and the 1-to-1 specialized memory ($C = N$) are indicated. The top panel shows the optimal performance Q^* of different strategies as a function of the Hopfield inverse temperature β_{H_c} . The right panel shows the mutual information $MI(\Sigma, \mathcal{C})$ between the set of pattern classes $\Sigma \equiv \{\sigma^\alpha\}$ and the set of all compartments \mathcal{C} normalized by the entropy of the compartments $H(\mathcal{C})$ as a function of the inverse temperature β_S ; see Appendix A.3. This normalized mutual information quantifies the ability of the system to assign a pattern to the correct compartment. **(C-D)** Similar to (A-B) but for evolving patterns with the effective mutation rate $\mu_{\text{eff}} = 0.01$. The number of presented patterns is set to $N = 32$ (all panels). Similar to Fig. 3 we keep $L \times C = \text{const.}$, with $L = 800$ used for networks with $C = 1$.

optimal performance $1 - 2\mu_{\text{eff}}$ (Figs. 3A, 4C,D). As predicted by the theory, the optimal learning rate for compartmentalized networks scales with the mutation rate as $\lambda_c^* \sim \mu_c^{1/2}$, except for the 1-to-1 network in which $\lambda_c^* \rightarrow 1$ and sub-networks are steadily updated upon an encounter with a pattern (Fig. 3B). This rapid update is expected since there is no interference between the stored memories in the 1-to-1 network, and a steady

update can keep the stored memory in each sub-network close to its associated pattern class without disrupting the other energy minima.

ii. Small intra- and large inter-compartment noise ($\beta_{H_c} \gg 1$, $\beta_S \ll 1$). In this regime there is low noise for equilibration within a compartment but a high level of noise in choosing the right compartment. The optimal strategy in this regime is to store patterns

in a single network with a distributed memory, since identifying the correct compartment is difficult due to noise (Fig. 4B,D). For static patterns this strategy corresponds to the classical Hopfield model with a high accuracy (Figs. 2A, 4A,B). On the other hand, for evolving patterns this strategy results in a partial memory (Fig. 4C,D) due to the reduced accuracy of the distributed associative memory, as shown in Fig. 2A. Interestingly, the transition between the optimal strategy with highly specific (compartmentalized) memory for evolving patterns in the first regime and the generalized (distributed) memory in this regime is very sharp (Fig. 4D). This sharp transition suggests that depending on the noise in choosing the compartments β_S , an optimal strategy either stores memory in a 1-to-1 specialized fashion ($C = N$) or in a distributed generalized fashion ($C = 1$), but no intermediate solution (i.e., quasi-specialized memory with $1 < C < N$) is desirable.

iii. Large intra-compartment noise ($\beta_{H_c} < 1$). In this regime there is a high level of noise in equilibration within a network and memory cannot be reliably retrieved (Fig. 4A,C), regardless of the compartmentalization temperature β_S . However, the impact of the equilibration noise β_{H_c} on the accuracy of memory retrieval depends on the degree of compartmentalization. For the 1-to-1 specialized network ($C = N$), the transition between the high and the low accuracy is smooth and occurs at $\beta_{H_c} = 1$, below which no memory attractor can be found. As we increase the equilibration noise (decrease β_{H_c}), the networks with distributed memory ($C < N$) show two-step transitions, with a plateau in the range of $1/N_c \lesssim \beta_{H_c} \lesssim 1$. Similar to the 1-to-1 network, the first transition at $\beta_{H_c} \approx 1$ results in the reduced accuracy of the networks' memory retrieval. At this transition point, the networks' learning rate λ_c approaches its maximum value 1 (Fig. S8), which implies that the memory is stored (and steadily updated) for only $C < N$ patterns (i.e., one pattern per sub-network). Due to the invariance of the networks' mean energy under compartmentalization, the depth of the energy minima associated with the stored memory in each sub-network scales as N/C , resulting in deeper and more stable energy minima in networks with smaller number of compartments C . Therefore, as the noise increases (i.e., β_{H_c} decreases), we observe a gradient in transition from partial retrieval to a no-memory state at $\beta_H \approx 1/N_c$, with the most compartmentalized network (larger C) transitioning the fastest, reflecting the shallowness of their energy minima.

Taken together, the optimal strategy leading to working memory depends on whether a network is trained to learn and retrieve dynamic (evolving) patterns or static patterns. Specifically, we see that the 1-to-1 specialized network is the unique optimal solution for storing working memory for evolving patterns, whereas the distributed generalized memory (i.e., the classical Hopfield network) performs equally well in learning and retrieval of memory for static patterns. The contrast between

these memory strategies can shed light on the distinct molecular mechanisms utilized by different biological systems to store memory.

III. DISCUSSION

Storing and retrieving memory from prior molecular interactions is an efficient scheme to sense and respond to external stimuli. Here, we introduced a flexible energy-based neural network model that can adopt different memory strategies, including distributed memory, similar to the classical Hopfield network, or compartmentalized memory. The learning rate and the number of compartments in a network define a memory strategy, and we probed the efficacy of different strategies for static and dynamic patterns. We found that Hopfield-like networks with distributed memory are highly accurate in storing associative memory for static patterns, even when patterns are noisy. However, these networks fail to reliably store retrievable associative memory for evolving patterns, even when learning is done at an optimal rate.

To achieve a high accuracy, we showed that a retrievable memory for evolving patterns should be compartmentalized, where each pattern class is stored in a separate sub-network. In addition, we found a sharp transition between the different phases of working memory (i.e., compartmentalized and distributed memory), suggesting that intermediate solutions (i.e., quasi-specialized memory) are sub-optimal against evolving patterns.

The contrast between these memory strategies is reflective of the distinct molecular mechanisms used for memory storage in the adaptive immune system and in the olfactory cortex. In particular, the memory of odor complexes, which can be assumed as static, is stored in a distributed fashion in the olfactory cortex [7–11, 15]. On the other hand, the adaptive immune system, which encounters evolving pathogens, allocates distinct immune cells (i.e., compartments) to store a memory for different types of pathogens (e.g. different variants of influenza or HIV)—a strategy that resembles that of the 1-to-1 specialized networks [5, 16–21]. Our results suggest that pathogenic evolution may be one of the reasons for the immune system to encode a specialized memory, as opposed to the distributed memory used in the olfactory system.

The increase in the optimal learning rate in anticipation of patterns' evolution significantly changes the structure of the energy landscape for associative memory. In particular, we found the emergence of narrow connectors (mountain passes) between the memory attractors of a network, which destabilize the equilibration process and significantly reduce the accuracy of memory retrieval. Indeed, tuning the learning rate as a hyper-parameter is one of the challenges of current machine learning algorithms with deep neural networks (DNNs) [24, 25]. The goal is to navigate the tradeoff between the speed (i.e., rate of convergence) and accuracy without overshooting

during optimization. It will be interesting to see how the insights developed in this work can inform rational approaches to choose an optimal learning rate in optimization tasks with DNNs.

Machine learning algorithms with DNNs [24] and modified Hopfield networks [34–37] are able to accurately classify hierarchically correlated patterns, where different objects can be organized into an ultrametric tree based on some specified relations of similarity. For example, faces of cats and dogs have the oval shape in common but they branch out in the ultrametric tree according to the organism-specific features, e.g., whiskers in a cat, and the cat branch can then further diversify based on the breed-specific features. A classification algorithm can use these hierarchical relations to find features common among members of a given sub-type (cats) that can distinguish them from another sub-type (dogs). Although evolving patterns within each class are correlated, the random evolutionary dynamics of these patterns does not build a hierarchical structure where a pattern class branches in two sub-classes that share a common ancestral root. Therefore, the optimal memory strategies found here for evolving patterns are distinct from those of the hierarchically correlated patterns. It will be interesting to see how our approaches can be implemented in DNNs to classify dynamic and evolving patterns.

ACKNOWLEDGEMENT

This work has been supported by the DFG grant (SFB1310) for Predictability in Evolution and the MPRG funding through the Max Planck Society. O.H.S also acknowledges funding from Georg-August University School of Science (GAUSS) and the Fulbright foundation.

Appendix A: Computational procedures

1. Initialization of the network

A network J (with elements J_{ij}) is presented with N random (orthogonal) patterns $|\sigma^\alpha\rangle$ (with $\alpha = 1, \dots, N$), with entries $\sigma_i^\alpha \in \{-1, 1\}$, reflecting the N pattern classes. For a network with C compartments (with $1 \leq C \leq N$), we initialize each sub-network J^s at time t_0 as $J_{i,j}^s(t_0) = \frac{1}{N/C} \sum_{\alpha \in \mathcal{A}_s} \sigma_i^\alpha \sigma_j^\alpha$ and $J_{ii}^s(t_0) = 0$; here, \mathcal{A}_s is a set of N/C randomly chosen (without replacement) patterns initially assigned to the compartment (sub-network) s . We then let the network undergo an initial learning process. At each step an arbitrary pattern σ^ν is presented to the network and a sub-network J^s is chosen for an update with a probability

$$P_s = \frac{\exp[-\beta_S E(J^s(t), \sigma^\nu(t))]}{\sum_{r=1}^C \exp[-\beta_S E(J^r(t), \sigma^\nu(t))]}, \quad (\text{A1})$$

where the energy is defined as,

$$\begin{aligned} E(J^s(t), \sigma^\nu(t)) &= \frac{-1}{2L} \sum_{i,j} J_{i,j}^s(t) \sigma_i^\nu(t) \sigma_j^\nu(t) \\ &\equiv \frac{-1}{2} \langle \sigma^\nu(t) | J^s(t) | \sigma^\nu(t) \rangle \end{aligned} \quad (\text{A2})$$

and β_S is the inverse temperature associated with choosing the right compartment. We then update the selected sub-network J^s , using the Hebbian update rule,

$$J_{i,j}^s(t+1) = \begin{cases} (1-\lambda) J_{i,j}^s(t) + \lambda \sigma_i^\nu \sigma_j^\nu, & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A3})$$

For dynamic patterns, the presented patterns undergo evolution with mutation rate μ , which reflects the average fraction of flipped spins in a given pattern per network's update event (Fig. 1). For noisy patterns, the actual patterns remain unchanged, but the network is presented with noisy versions of these patterns. Here, the noise amplitude reflects the average fraction of flipped spins presented and actual patterns.

Our goal is to study the memory retrieval problem in a network that has reached its steady state. The state of a network $J(t_n)$ at the time step n can be traced back to the initial state $J(t_0)$ as,

$$J(t_n) = (1-\lambda)^n J(t_0) + \lambda \sum_{i=1}^n (1-\lambda)^{n-i} |\sigma(t_i)\rangle \langle \sigma(t_i)| \quad (\text{A4})$$

The contribution of the initial state $J(t_0)$ to the state of the network at time t_n decays as $(1-\lambda)^n$ (eq. A4). Therefore, we choose the number of steps to reach the steady state as $n_{\text{stat.}} = \max \left[10N, 2C \text{ ceil} \left(\frac{\log 10^{-5}}{\log(1-\lambda)} \right) \right]$. This criteria ensures that $(1-\lambda)^{n_{\text{stat.}}} \leq 10^{-5}$ and the memory of the initial state $J(t_0)$ is removed from the network $J(t)$. We will then use this updated network to collect the statistics for memory retrieval. To report a specific quantity from the network (e.g., the energy), we pool the $n_{\text{stat.}}$ samples collected from each of the 50 realizations.

2. Pattern retrieval from associative memory

Once the trained network approaches a stationary state, we collect the statistics of the stored memory.

To find a memory attractor $\sigma_{\text{att}}^\alpha$ for a given pattern σ^α we use a Metropolis algorithm in the energy landscape $E(J^s, \sigma^\alpha)$ (eq. A2). To do so, we make spin-flips in a presented pattern $\sigma^\alpha \rightarrow \tilde{\sigma}^\alpha$ and accept a spin-flip with probability

$$P(\sigma^\alpha \rightarrow \tilde{\sigma}^\alpha) = \min(1, e^{-\beta_H \Delta E}), \quad (\text{A5})$$

where $\Delta E = E(J^s, \tilde{\sigma}^\alpha) - E(J^s, \sigma^\alpha)$ and β_H is the inverse (Hopfield) temperature for pattern retrieval in the network (see Fig. 1). We repeat this step for 2×10^6 steps,

which is sufficient to find a minimum of the landscape (see Fig. S4).

For systems with more than one compartment C , we first choose a compartment according to eq. A1, and then perform the Metropolis algorithm within the associated compartment.

After finding the energy minima, we update the systems for $n'_{\text{stat.}} = \max[2 \cdot 10^3, n_{\text{stat.}}]$ steps. At each step we present patterns as described above and collect the statistics of the recognition energy $E(J^s(t), \sigma^\alpha(t))$ between a presented pattern σ^α and the memory compartment $J^s(t)$, assigned according to eq. A1. These measurements are used to describe the energy statistics (Figs. 2,S3) of the patterns and the accuracy of pattern-compartment association (Fig. 4B,D). After the $n'_{\text{stat.}}$ steps, we again use the Metropolis algorithm to find the memory attractors associated with the presented patterns. We repeat this analysis for 50 independent realizations of the initializing pattern classes $\{\sigma^\alpha(t_0)\}$, for each set of parameters $\{L, N, C, \lambda, \mu, \beta_S, \beta_H\}$.

When calculating the mean performance \mathcal{Q} of a strategy (see Figs. 2,3,4,S8), we set the overlap between attractor and pattern $q^\alpha = |\langle \sigma_{\text{att}}^\alpha | \sigma^\alpha \rangle|$ equal to zero when patterns are not recognized ($q^\alpha < 0.8$). As a result, systems can only achieve a non-zero performance if they recognize some of the patterns. This choice eliminates the finite size effect of a random overlap $\sim 1/\sqrt{L}$ between an attractor and a pattern (see Fig. S4). This correction is especially important when comparing systems with different sub-network sizes ($L_c \equiv L/C$) in the $\beta_H < 1$ regime (Figs. 4,S8), where random overlaps for small L_c could otherwise result in a larger mean performance compared to larger systems that correctly reconstruct a fraction of the patterns.

3. Accuracy of pattern-compartment association

We use the mutual information $\text{MI}(\Sigma, \mathcal{C})$ between the set of pattern classes $\Sigma \equiv \{\sigma^\alpha\}$ and the set of all compartments \mathcal{C} to quantify the accuracy in associating a presented pattern with the correct compartment,

$$\begin{aligned} \text{MI}(\Sigma, \mathcal{C}) &= H(\mathcal{C}) - H(\mathcal{C}|\Sigma) \\ &= - \sum_{c \in \mathcal{C}} P(c) \log P(c) \\ &\quad - \left[- \sum_{\sigma^\alpha \in \Sigma} P(\sigma^\alpha) \sum_{c \in \mathcal{C}} P(c|\sigma^\alpha) \log P(c|\sigma^\alpha) \right]. \end{aligned} \tag{A6}$$

Here $H(\mathcal{C})$ and $H(\mathcal{C}|\Sigma)$ are the entropy of the compartments, and the conditional entropy of the compartments given the presented patterns, respectively. If chosen randomly, the entropy associated with choosing a compartment is $H_{\text{random}}(\mathcal{C}) = \log C$. The mutual information (eq. A6) measures the reduction in the entropy of compartments due to the association between the patterns and the compartments, measured by the conditional entropy $H(\mathcal{C}|\Sigma)$. Figure 4B,D shows the mutual information $\text{MI}(\Sigma, \mathcal{C})$ scaled by its upper bound $H(\mathcal{C})$, in order to compare networks with a different number of compartments.

Appendix B: Estimating energy and optimal learning rate for working memory

1. Approximate solution for optimal learning rate

The optimal learning rate is determined by maximizing the network's performance $\mathcal{Q}(\lambda)$ (eq. 2) against evolving patterns with a specified mutation rate:

$$\lambda^* = \underset{\lambda}{\text{argmax}} \mathcal{Q}(\lambda) \tag{B1}$$

We can numerically estimate the optimal learning rate as defined eq. B1; see Figs. 2,3. However, the exact analytical evaluation of the optimal learning rate is difficult and we use an approximate approach and find the learning rate that minimizes the expected energy of the patterns in the stationary state $\mathbb{E}[E_{\lambda, \rho}(J, \sigma)]$, assuming that patterns are shown to the network at a fixed order. Here, the subscripts explicitly indicate the learning rate of the network λ , and the evolutionary overlap of the pattern ρ . To evaluate an analytical approximation for the energy, we first evaluate the state of the network $J(t)$ at time step t , given all the prior encounters of the networks with patterns shown at a fixed order.

$$\frac{1}{L} (J(t) + \mathbf{1}) = \lambda \sum_{j=1}^{\infty} (1-\lambda)^{(j-1)} |\sigma(t-j)\rangle \langle \sigma(t-j)| \quad (\text{B2})$$

$$= \lambda \sum_{i=1}^{\infty} (1-\lambda)^{(i-1)N} \underbrace{\sum_{\alpha=1}^N (1-\lambda)^{\alpha-1} |\sigma^\alpha(t-\alpha-(i-1)N)\rangle \langle \sigma^\alpha(t-\alpha-(i-1)N)|}_{\substack{\text{sum over } N \text{ pattern classes} \\ \text{sum over time (generations, } i)}} \quad (\text{B3})$$

$$= \lambda \underbrace{\sum_{\alpha=1}^N \sum_{i=0}^{\infty} (1-\lambda)^{(\alpha-1)+iN} |\sigma^\alpha(t-\alpha-iN)\rangle \langle \sigma^\alpha(t-\alpha-iN)|}_{\substack{\text{sum over time} \\ \text{sum over patterns}}} \quad (\text{B4})$$

Here, we referred to the (normalized) pattern vector from the class α presented to the network at time step t by $|\sigma^\alpha(t)\rangle \equiv \frac{1}{\sqrt{L}} \sigma^\alpha(t)$. Without loss of generality, we assumed that the last pattern presented to the network at time step $t-1$ is from the first pattern class $|\sigma^1(t-1)\rangle$, which enabled us to split the sum in eq. B2 into two separate summations over pattern classes and N time-

steps generations (eq. B3). Adding the identity matrix $\mathbf{1}$ on the left-hand side of eq. B2 assures that the diagonal elements vanish, as defined in eq. A3.

The mean energy of the patterns, which in our setup is the energy of the pattern from the N^{th} class at time t , follows

$$\begin{aligned} \mathbb{E} [E_{\lambda,\rho}(J, \sigma)] &= \mathbb{E} \left[-\frac{1}{2} \langle \sigma^N(t) | J(t) | \sigma^N(t) \rangle \right] \\ &= \mathbb{E} \left[-\frac{L-1}{2} \lambda \sum_{\alpha=1}^N \sum_{i=0}^{\infty} (1-\lambda)^{(\alpha-1)+iN} \langle \sigma^N(t) | \sigma^\alpha(t-\alpha-iN) \rangle \langle \sigma^\alpha(t-\alpha-iN) | \sigma^N(t) \rangle \right]. \end{aligned} \quad (\text{B5})$$

Since the pattern families are orthogonal to each other, we can express the over-

lap between patterns at different times as $\langle \sigma^\alpha(t_1) | \sigma^\nu(t_2) \rangle = \delta_{\alpha,\nu} (1-2\mu)^{|t_2-t_1|} \equiv \delta_{\alpha,\nu} \rho^{|t_2-t_1|}$, and simplify the energy function in eq. B5,

$$\begin{aligned} \mathbb{E} [E_{\lambda,\rho}(J, \sigma)] &= -\frac{L-1}{2} \lambda \sum_{i=0}^{\infty} (1-\lambda)^{(N-1)+iN} \rho^{2(N+iN)} \\ &= -\frac{L-1}{2} \lambda (1-\lambda)^{(N-1)} \rho^{2N} \sum_{i=0}^{\infty} ((1-\lambda)^N \rho^{2N})^i \\ &= -\frac{L-1}{2} \lambda \frac{(1-\lambda)^{(N-1)} \rho^{2N}}{1-(1-\lambda)^N \rho^{2N}}. \end{aligned} \quad (\text{B6})$$

Since accurate pattern retrieval depends on the depth of the energy valley for the associative memory, we will use the expected energy of the patterns as a proxy for the performance of the network. We can find the approxi-

mate optimal learning rate that minimized the expected energy by setting $\partial \mathbb{E} [E_{\lambda,\rho}(J, \sigma)] / \partial \lambda = 0$, which results

in

$$\begin{aligned}
(1 - 2\mu)^{2N} &= (1 - \lambda^*)^{-N} (1 - N\lambda^*) \\
\implies 1 - 4N\mu + \mathcal{O}(\mu^2) &= 1 + \frac{1}{2}(N - N^2)(\lambda^*)^2 + \mathcal{O}(\lambda^3); \\
\implies \lambda^*(\mu) &\simeq \sqrt{8\mu/(N - 1)}.
\end{aligned} \tag{B7}$$

where we used the fact that both the mutation rate μ and the learning rate λ are small, and therefore, expanded eq. B7 up to the leading orders in these parameters.

In addition, eq. B7 establishes an upper bound for the learning rate: $\lambda < \frac{1}{N}$. Therefore, our expansion in mutation rate (eq. B7) is only valid for $8\mu < \frac{1}{N}$, or equivalently for $\mu_{\text{eff}} = N\mu < 12.5\%$; the rates used in our analyses lie far below these upper bounds.

2. Lag of memory against evolving patterns

The memory attractors associated with a given pattern class can lag behind and only reflect the older patterns presented prior to the most recent encounter of the network with the specified class. As a result, the upper bound for the performance of a network $\mathcal{Q}_{\text{lag}} = \rho^{g_{\text{lag}} N} \approx$

$1 - 2g_{\text{lag}}\mu_{\text{eff}}$ is determined by both the evolutionary divergence of patterns between two encounters μ_{eff} and number of generations g_{lag} by which the stored memory lags behind; we measure g_{lag} in units of generations; one generation is defined as the average time between a network's encounter with the same pattern class i.e., N . We characterize this lag g_{lag} by identifying the past pattern (at time $t - g_{\text{lag}}N$) that has the maximum overlap with the network's energy landscape at given time t :

$$\begin{aligned}
g_{\text{lag}} &= \operatorname{argmax}_{g \geq 0} \mathbb{E} [\langle \sigma(t - gN) | J(t) | \sigma(t - gN) \rangle] \\
&\equiv \operatorname{argmin}_{g \geq 0} \mathbb{E} [E_{\text{lag}}(g)]
\end{aligned} \tag{B8}$$

where we introduced the expected lagged energy $\mathbb{E} [E_{\text{lag}}(g)]$. Here, the vector $|\sigma(t)\rangle$ refers to the pattern σ presented to the network at time t , which can be from any of the pattern classes. Because of the translational symmetry in time in the stationary state, the lagged energy can also be expressed in terms of the overlap between a pattern at time t and the network at a later time $t + gN$. We evaluate the lagged energy by substituting the expression for the network's state $J(t + gN)$ from eq. B2, which entails,

$$\frac{2}{L-1} \mathbb{E} [E_{\text{lag}}(g)] = -\frac{1}{L-1} \mathbb{E} [\langle \sigma(t) | J(t + gN) | \sigma(t) \rangle] \tag{B9}$$

$$= -\mathbb{E} \left[\frac{1}{L-1} (1 - \lambda)^{Ng} \langle \sigma(t) | J(t) | \sigma(t) \rangle + \lambda \sum_{j=0}^{gN-1} (1 - \lambda)^{gN-1-j} \langle \sigma(t) | \sigma(t + j) \rangle^2 \right] \tag{B10}$$

$$= \frac{2}{L-1} (1 - \lambda)^{Ng} \mathbb{E} [E_{\lambda, \rho}(J, \sigma)] - \lambda \sum_{i=0}^{g-1} \sum_{\alpha=0}^{N-1} (1 - \lambda)^{gN-1-Ni-\alpha} \langle \sigma^N(t) | \sigma^{N-\alpha}(t + Ni + \alpha) \rangle^2 \tag{B11}$$

$$= \frac{2}{L-1} (1 - \lambda)^{Ng} \mathbb{E} [E_{\lambda, \rho}(J, \sigma)] - \lambda \sum_{i=0}^{g-1} (1 - \lambda)^{gN-1-Ni} \rho^{2Ni} \tag{B12}$$

$$= -\lambda \left(\frac{(1 - \lambda)^{N(g+1)} \rho^{2N}}{1 - (1 - \lambda)^N \rho^{2N}} + \frac{(1 - \lambda)^{N(g+1)-1} - (1 - \lambda)^{N-1} \rho^{2Ng}}{(1 - \lambda)^N - \rho^{2N}} \right). \tag{B13}$$

Here, we used the expression of the network's matrix J in eq. A4 to arrive at eq. B10, and then followed the procedure introduced in eq. B3 to arrived at the double-summation in eq. B11. We then used the equation for pattern overlap $\langle \sigma^\alpha(t_1) | \sigma^\nu(t_2) \rangle = \delta_{\alpha, \nu} \rho^{|t_2 - t_1|}$ to reduce the sum in eq. B12 and arrived at the result in eq. B13 by evaluating the geometric sum and substituting the empirical average for the energy $\mathbb{E} [E_{\lambda, \rho}(J, \sigma)]$ from eq. B6.

We probe this lagged memory by looking at the performance \mathcal{Q} for patterns that are correctly associated with their memory attractors (i.e., those with $\langle \sigma_{\text{att}} | \sigma \rangle > 0.8$). As shown in Fig. S2, for a broad parameter regime,

the mean performance for these correctly associated patterns agrees well with the theoretical expectation $\mathcal{Q}_{\text{lag}} = \rho^{g_{\text{lag}} N}$, which is lower than the naive expectation \mathcal{Q}_0 .

Appendix C: Structure of the energy landscape for working memory

1. Formation of mountain passes in the energy landscape of memory for evolving patterns

As shown in Fig. 1, large learning rates in networks with memory for evolving patterns result in the emergence of narrow connecting paths between the minima of the energy landscape. We refer to these narrow connecting paths as *mountain passes*. In pattern retrieval, the Monte Carlo search can drive a pattern out of one energy minimum into another minimum and potentially lead to pattern misclassification.

We use two features of the energy landscape to probe the emergence of the mountain passes.

First, we show that if a pattern is misclassified, it has fallen into a memory attractor associated with another pattern class and not a spuriously made energy minima. To do so, we compare the overlap of the attractor with the original pattern $|\langle \sigma_{\text{att}}^\alpha | \sigma^\alpha \rangle|$ (i.e., the reconstruction performance of the patterns) with the maximal overlap of the attractor with all other patterns $\max_{\nu \neq \alpha} |\langle \sigma_{\text{att}}^\alpha | \sigma^\nu \rangle|$. Indeed, as shown in Fig. S4A for evolving patterns, the memory attractors associated with 99.4% of the originally stored patterns have either a large overlap with the correct pattern or with one of the other previously presented pattern. 71.3% of the patterns are correctly classified (stable patterns in sector **I** in Fig. S4A), whereas 28.1% of them are associated with a secondary energy minima after equilibration (unstable patterns in sector **II** in Fig. S4A). A very small fraction of patterns ($< 1\%$) fall into local minima given by the linear combinations of the presented patterns (sector **IV** in Fig. S4A). These minima are well-known in the classical Hopfield model [38, 39]. Moreover, we see that equilibration of a random pattern (i.e., a pattern orthogonal to all the presented classes) in the energy landscape leads to memory attractors for one of the originally presented pattern classes. The majority of these random patterns lie in sector **II** of Fig. S4A, i.e., they have a small overlap with the network since they are orthogonal to the originally presented pattern classes, and they fall into one of the existing memory attractors after equilibration.

Second, we characterize the possible paths for a pattern to move from one valley to another during equilibration, using Monte Carlo algorithm with the Metropolis acceptance probability,

$$\rho(\sigma \rightarrow \sigma') = \min \left(1, e^{-\beta(E(J, \sigma') - E(J, \sigma))} \right). \quad (\text{C1})$$

We estimate the number of beneficial spin-flips (i.e., open paths) that decrease the energy of a pattern at the start of equilibration (Fig. S4B). The average number of open paths is smaller for stable patterns compared to the unstable patterns that are be misclassified during retrieval (Fig. S4B). However, the distributions for the number of open paths largely overlap for stable and un-

stable patterns. Therefore, the local energy landscape of stable and unstable patterns are quite similar and it is difficult to discriminate between them solely based on the local gradients in the landscape. Fig. S5A shows that the average number of beneficial spin-flips grows with the mutation rate of the patterns but this number is comparable for stable and unstable patterns. Moreover, the unstable stored patterns (blue) have far fewer open paths available to them during equilibration compared to random patterns (red) that are presented to the network for the first time (Figs. S4B & S5A). Notably, on average half of the spin-flips reduce the energy of for random patterns, irrespective of the mutation rate. This indicates that even though previously presented pattern classes are statistically distinct from random patterns, they can still become unstable, especially in networks which are presented with evolving patterns.

It should be noted that the evolution of the patterns only indirectly contributes to the misclassification of memory, as it necessitates a larger learning rate for the networks to optimally operate, which in turn results in the emergence of mountain passes. To clearly demonstrate this effect, Figs. S4C,D, and S5D shows the misclassification behavior for a network trained to store memory for static pattern, while using a larger learning rate that is optimized for evolving patterns. Indeed, we see that pattern misclassification in this case is consistent with the existence of mountain passes in the network's energy landscape.

2. Spectral decomposition of the energy landscape

We use spectral decomposition of the energy landscape to characterize the relative positioning and the stability of patterns in the landscape. As shown in Figs. S4, S5, destabilization of patterns due to equilibration over mountain passes occurs in networks with high learning rates, even for static patterns. Therefore, we focus on how the learning rate impacts the spectral decomposition of the energy landscape in networks presented with static patterns. This simplification will enable us to analytically probe the structure of the energy landscape, which we will compare with numerical results for evolving patterns.

We can represent the network J (of size $L \times L$) that store a memory of N static patterns with N non-trivial eigenvectors $|\Phi^i\rangle$ with corresponding eigenvalues Γ_i , and $N - L$ degenerate eigenvectors, $|\Psi^k\rangle$ with corresponding trivial eigenvalues $\gamma_k = \gamma = -1$:

$$J = \sum_{i=1}^N \Gamma_i |\Phi^i\rangle \langle \Phi^i| + \sum_{k=1}^{L-N} \gamma_k |\Psi^k\rangle \langle \Psi^k|. \quad (\text{C2})$$

The non-trivial eigenvectors span the space of the presented patterns, for which the recognition energy can be

expressed by

$$E(J, \sigma^\alpha) = -\frac{1}{2} \sum_{i=1}^N \Gamma_i \langle \sigma^\alpha | \Phi^i \rangle \langle \Phi^i | \sigma^\alpha \rangle. \quad (\text{C3})$$

An arbitrary configuration χ in general can have components orthogonal to the N eigenvectors $|\Phi^i\rangle$, as it points to a vertex of the hypercube, and should be expressed in terms of all the eigenvectors $\{\Phi^1, \dots, \Phi^N, \Psi^1, \dots, \Psi^{L-N}\}$:

$$E(J, \chi) = -\frac{1}{2} \left(\underbrace{\sum_{i=1}^N \Gamma_i \langle \chi | \Phi^i \rangle \langle \Phi^i | \chi \rangle}_{\text{stored patterns}} + \underbrace{\sum_{k=1}^{L-N} \gamma \langle \chi | \Psi^k \rangle \langle \Psi^k | \chi \rangle}_{\text{trivial space}} \right). \quad (\text{C4})$$

Any spin-flip in a pattern (e.g., during equilibration) can be understood as a rotation in the eigenspace of the network (eq. C4). As a first step in characterizing these rotations we remind ourselves of the identity

$$|\chi\rangle = \sum_{i=1}^N \langle \Phi^i | \chi \rangle |\Phi^i\rangle + \sum_{k=1}^{L-N} \langle \Psi^k | \chi \rangle |\Psi^k\rangle, \quad (\text{C5})$$

with the normalization condition

$$\sum_{i=1}^N (\langle \Phi^i | \chi \rangle)^2 + \sum_{k=1}^{L-N} (\langle \Psi^k | \chi \rangle)^2 = 1. \quad (\text{C6})$$

In addition, since the diagonal elements of the network are set to $J_{ii} = 0$ (eq. A3), the eigenvalues should sum to zero, or alternatively,

$$\sum_{i=1}^N \Gamma_i = - \sum_{k=1}^{L-N} \gamma_k = L - N. \quad (\text{C7})$$

To assess the stability of a pattern σ^ν , we compare its recognition energy $E(J, \sigma^\nu)$ with the energy of the rotated pattern after a spin-flip $E(J, \tilde{\sigma}^\nu)$. To do so, we first consider a simple scenario, where we assume that the pattern σ^ν has a large overlap with one dominant non-trivial eigenvector Φ^A (i.e., $\langle \sigma^\nu | \Phi^A \rangle^2 = m^2 \approx 1$). The other components of the pattern can be expressed in terms of the remaining $N - 1$ non-trivial eigenvectors with mean squared overlap $\frac{1-m^2}{N-1}$. The expansion of the recognition energy for the presented pattern is restricted to the N non-trivial directions (eq. C4), resulting in

$$\begin{aligned} E(J, \sigma^\nu) &= -\frac{1}{2} \left(m^2 \Gamma_A + \sum_{i \neq A} \frac{1-m^2}{N-1} \Gamma_i \right) \\ &= -\frac{1}{2} \left(m^2 \Gamma_A + (1-m^2) \tilde{\Gamma} \right), \end{aligned} \quad (\text{C8})$$

where $\tilde{\Gamma} = \frac{1}{N-1} \sum_{i \neq A} \Gamma_i = \frac{N\tilde{\Gamma} - \Gamma_A}{N-1}$ is the mean eigenvalue for the non-dominant directions.

A spin-flip ($|\sigma^\nu\rangle \rightarrow |\tilde{\sigma}^\nu\rangle$) can rotate the pattern out of the dominant direction Φ^A and reduce the squared overlap by ϵ^2 . The rotated pattern $|\tilde{\sigma}^\nu\rangle$ in general lies in the L -dimensional space and is not restricted to the N -dimensional (non-trivial) subspace. We first take a *mean-field* approach in describing the rotation of the pattern after a spin-flip. Because of the normalization condition (eq. C6), the loss in the overlap with the dominant direction should result in an *average* increase in the overlap with the other $L - 1$ eigenvectors by $\frac{\epsilon^2}{L-1}$. The energy of the rotated pattern after a spin-flip $E(J, \tilde{\sigma}^\nu)$ can be expressed in terms of all the L eigenvectors (eq. C4),

$$\begin{aligned} E(J, \tilde{\sigma}^\nu) &= -\frac{1}{2} \left[(m^2 - \epsilon^2) \Gamma_A + \sum_{i \neq A} \left(\frac{1-m^2}{N-1} + \frac{\epsilon^2}{L-1} \right) \Gamma_i + \sum_k \frac{\epsilon^2}{L-1} \gamma_k \right] \\ &= E(J, \sigma^\nu) + \frac{\epsilon^2}{2} \left[\Gamma_A - \frac{1}{L-1} \left(\sum_{i \neq A} \Gamma_i + \sum_k \gamma_k \right) \right] \end{aligned} \quad (\text{C9})$$

$$= E(J, \sigma^\nu) + \frac{\epsilon^2}{2} \Gamma_A \left(1 + \frac{1}{L-1} \right). \quad (\text{C10})$$

where in eq. C10 we used the fact that the eigenvalues should sum up to zero. On average, a spin-flip $|\sigma^\nu\rangle \rightarrow |\tilde{\sigma}^\nu\rangle$ increases the recognition energy by $E(J, \tilde{\sigma}^\nu) - E(J, \sigma^\nu) = \frac{\epsilon^2}{2} \Gamma_A [1 + \mathcal{O}(L^{-1})]$. This is

consistent with the results shown in Figs. S4B,D and Figs. S5A,D, which indicate that the majority of the spin-flips keep a pattern in the original energy minimum and only a few of the spin-flips drive a pattern out of the

original attractor.

In the analysis above, we assumed that the reduction in overlap with the dominant eigenvector ϵ^2 is absorbed equally by all the other eigenvectors (i.e., the mean-field approach). In this case, the change in energy is equally distributed across the positive and the negative eigenvalues (Γ 's and γ 's in eq. C9), resulting in an overall increase in the energy due to the reduced overlap with the dominant direction $|\Phi^A\rangle$. The destabilizing spin-flips are associated with atypical changes that rotate a pattern onto a secondary non-trivial direction $|\Phi^B\rangle$ (with positive eigenvalue Γ_B), as a result of which the total energy could be reduced. To better characterize the con-

ditions under which patterns become unstable, we will introduce a perturbation to the mean-field approach used in eq. C10. We will assume that a spin-flip results in a rotation with a dominant component along a secondary non-trivial direction $|\Phi^B\rangle$. Specifically, we will assume the reduced overlap ϵ^2 between the original pattern $|\sigma^\nu\rangle$ and the dominant direction $|\Phi^A\rangle$ is distributed in an imbalanced fashion between the other eigenvectors, with a fraction p projected onto a new (non-trivial) direction $|\Phi^B\rangle$, while all the other $L - 2$ directions span the remaining $(1 - p)\epsilon^2$. In this case, the energy of the rotated pattern is given by

$$\begin{aligned} E(J, \tilde{\sigma}^\nu) &= -\frac{1}{2} \left[(m^2 - \epsilon^2)\Gamma_A + \left(\frac{1 - m^2}{N - 1} + p\epsilon^2 \right) \Gamma_B + \sum_{i \neq A, B} \left(\frac{1 - m^2}{N - 1} + \frac{(1 - p)\epsilon^2}{L - 2} \right) \Gamma_i + \sum_k \frac{(1 - p)\epsilon^2}{L - 2} \gamma_k \right] \\ &= E(J, \sigma^\nu) + \frac{\epsilon^2}{2} [\Gamma_A - p\Gamma_B + \mathcal{O}(L^{-1})]. \end{aligned} \quad (\text{C11})$$

Therefore, a spin-flip is beneficial if $\Gamma_A < p\Gamma_B$. To further concretize this condition, we will estimate the typical loss ϵ^2 and gain $p\epsilon^2$ in the squared overlap between the pattern and its dominating directions due to rotation by a single spin-flip.

Let us consider a rotation $|\sigma^\nu\rangle \rightarrow |\tilde{\sigma}^\nu\rangle$ by a flip in the n -th spin of the original pattern $|\sigma^\nu\rangle$. This spin flip reduces the original overlap of the pattern $m = \langle \sigma^\nu | \Phi^A \rangle$ with the dominant direction $|\Phi^A\rangle$ by the amount $\frac{2}{\sqrt{L}} \Phi_n^A$, where Φ_n^A is the n -th entry of the eigenvector $|\Phi^A\rangle$. Since the original overlap is large (i.e., $m \simeq 1$), all entries of the dominant eigenvector are approximately $\Phi_i^A \simeq 1/\sqrt{L}$, $\forall i$, resulting in a reduced overlap of the rotated pattern $\langle \tilde{\sigma}^\nu | \Phi^A \rangle = m - \frac{2}{L}$. Therefore, the loss in the squared

overlap ϵ^2 by a spin flip is given by

$$\begin{aligned} \epsilon^2 &= \langle \sigma^\nu | \Phi^j \rangle^2 - \langle \tilde{\sigma}^\nu | \Phi^j \rangle^2 = m^2 - \left(m^2 - 4\frac{m}{L} + 4\frac{1}{L^2} \right) \\ &= 4\frac{m}{L} + \mathcal{O}\left(\frac{1}{L^2}\right). \end{aligned} \quad (\text{C12})$$

Similarly, we can derive the gain in the squared overlap $p\epsilon^2$ between the pattern $|\sigma^\nu\rangle$ and the new dominant direction $|\Phi^B\rangle$ after a spin-flip. Except for the direction $|\Phi^A\rangle$, the expected squared overlap between the original pattern (prior to the spin flip) and any of the non-trivial eigenvectors including $|\Phi^B\rangle$ is $\langle \sigma^\nu | \Phi^B \rangle^2 = \frac{1 - m^2}{N - 1}$. The flip in the n -th spin of the original pattern increases the overlap of the rotated pattern with the new dominant direction $|\Phi^B\rangle$ by $2\frac{\Phi_n^B}{\sqrt{L}}$, where Φ_n^B should be of the order of $\sqrt{1/L}$. Therefore, the largest gain in overlap due to a spin-flip is given by

$$\begin{aligned} p\epsilon^2 &= \langle \tilde{\sigma}^\nu | \Phi^B \rangle^2 - \langle \sigma^\nu | \Phi^B \rangle^2 \simeq \left(\frac{1 - m^2}{N - 1} + 4\sqrt{\frac{1 - m^2}{N - 1}} \frac{\Phi_n^B}{\sqrt{L}} + 4\frac{(\Phi_n^B)^2}{L} \right) - \frac{1 - m^2}{N - 1} \\ &= \sqrt{\frac{1 - m^2}{N - 1}} \frac{\Phi_n^B}{\sqrt{L}} + \mathcal{O}\left(\frac{1}{L^2}\right). \end{aligned} \quad (\text{C13})$$

By using the results from eqs. C12 and C13, we can express the condition for beneficial spin-flips to drive a pattern over the carved mountain passes during equilibration

(eq. C11),

$$\epsilon^2 \Gamma_A < \epsilon^2 p \Gamma_B, \quad \rightarrow \quad \frac{\Gamma_A}{\Gamma_B} < \sqrt{\frac{1 - m^2}{m^2}} \frac{1}{\sqrt{\alpha}} \Phi_n^B, \quad (\text{C14})$$

where $\alpha = N/L$. This result suggests that the stability of a pattern depends on how the ratio of the eigenvalues associated with the dominant projections of the pattern before and after the spin-flip Γ_A/Γ_B compare to the overlap m of the original pattern with the dominant eigenvector Φ^A and the change due to the spin-flip Φ_n^B .

So far, we have constrained our analysis to patterns that have a dominant contribution to only one eigenvector Φ^A . To extend our analysis to patterns which are instead constrained to a small sub-space \mathcal{A} spanned by non-trivial eigenvalues, we define the squared pattern overlap with the subspace $m_{\mathcal{A}}^2 = \sum_{\alpha \in \mathcal{A}} \langle \sigma^\nu | \Phi^\alpha \rangle^2$ and a weighted averaged eigenvalue in the subspace $\Gamma_{\mathcal{A}} = \sum_{\alpha \in \mathcal{A}} \langle \sigma^\nu | \Phi^\alpha \rangle^2 \Gamma_\alpha$. As a result, the difference in the energy of a pattern before and after a spin-flip (eq. C11) can be extended to $E(J, \sigma^\nu) - E(J, \tilde{\sigma}^\nu) = \frac{\epsilon^2}{2} [\Gamma_{\mathcal{A}} - p\Gamma_B + \mathcal{O}(L^{-1})]$. Similarly, the stability condition in eq. C14 can be extended to $\frac{\Gamma_{\mathcal{A}}}{\Gamma_B} < \sqrt{\frac{1-m_{\mathcal{A}}^2}{m_{\mathcal{A}}^2}} \frac{1}{\sqrt{\alpha}} \Phi_n^B$.

Patterns that are constrained to larger subspaces are more stable, since the weighted averaged eigenvalue for their containing subspace $\Gamma_{\mathcal{A}}$ is closer to the mean of all eigenvalues $\bar{\Gamma} = 1 - N/L$ (law of large numbers). Therefore, in these cases a much larger eigenvalue gap (or a broader eigenvalue spectrum) is necessary to satisfy the condition for pattern instability.

Fig. S7 compares the loss in energy with the original dominant direction $\epsilon^2\Gamma_{\mathcal{A}}$ to the maximal gain in any of the other directions $\epsilon^2p\Gamma_B$ to test the pattern stability criteria presented in eq. C14. To do so, we identify a spin flip n in a secondary direction B that confers the maximal energy gain: $\epsilon^2p\Gamma_B \approx \max_{n,B} \sqrt{\frac{1-m_{\mathcal{A}}^2}{N-1}} \frac{\Phi_n^B}{\sqrt{L}} \Gamma_B$. In Fig. S7A,C we specifically focus on the subset of patterns that show a large (squared) overlap with the one

dominant direction (i.e., $m > 0.85$). Given that evolving patterns are not constraint to the $\{\Phi\}$ (non-trivial) sub-space, we find a smaller fraction of these patterns to fulfill the condition for such large overlap m (Fig. S7A), compared to the static patterns (Fig. S7C). Nonetheless, we see that the criteria in eq. C14 can be used to predict the stability of patterns in a network for both static and evolving patterns; note that here we use the same learning rate for both the static and evolving patterns.

We then relax the overlap condition by including all patterns that have a large overlap with a subspace \mathcal{A} , spanned by up to 10 eigenvectors (i.e., $m_{\mathcal{A}}^2 = \sum_{\alpha \in \mathcal{A}} \langle \sigma | \Phi^\alpha \rangle^2 > 0.85$). For these larger subspaces the transition between stable and unstable patterns is no longer exactly given by eq. C14. However, the two contributions $\epsilon^2\Gamma_{\mathcal{A}}$ and $\epsilon^2p\Gamma_B$ still clearly separate the patterns into stable and unstable classes for both static and evolving patterns (Figs. S7B,D). The softening of this condition is expected as in this regime we can no longer assume that a single spin-flip can reduce the overlap with all the eigenvectors in the original subspace. As a result, the effective loss in overlap become smaller than ϵ^2 and patterns become unstable below the dotted line in Fig. S7B,D.

As the learning rate increases, the gap between the eigenvalues $\Gamma_B/\Gamma_{\mathcal{A}}$ (Fig. S6) become larger. At the same time, patterns become more constrained to smaller subspaces (Fig. S5C,D). As a result of these two effects, more patterns satisfy the instability criteria in eq. C14. These patterns are misclassified as they fall into a wrong energy minimum by equilibrating through the mountain passes carved in the energy landscape of networks with large learning rates.

-
- [1] S. J. Labrie, J. E. Samson, and S. Moineau, Bacteriophage resistance mechanisms, *Nature Rev. Microbiol.* **8**, 317 (2010).
 - [2] R. Barrangou and L. A. Marraffini, CRISPR-Cas systems: Prokaryotes upgrade to adaptive immunity, *Molecular cell* **54**, 234 (2014).
 - [3] S. Bradde, A. Nourmohammad, S. Goyal, and V. Balasubramanian, The size of the immune repertoire of bacteria, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5144 (2020).
 - [4] A. S. Perelson and G. Weisbuch, Immunology for physicists, *Rev Mod Phys* **69**, 1219 (1997).
 - [5] C. Janeway, P. Travers, M. Walport, and M. Schlomchik, *Immunobiology*, 5th ed., The Immune System in Health and Disease (Garland Science, New York, 2001).
 - [6] G. Altan-Bonnet, T. Mora, and A. M. Walczak, Quantitative immunology for physicists, *Physics Reports* **849**, 1 (2020).
 - [7] L. B. Haberly and J. M. Bower, Olfactory cortex: Model circuit for study of associative memory?, *Trends Neurosci.* **12**, 258 (1989).
 - [8] P. Brennan, H. Kaba, and E. B. Keverne, Olfactory recognition: A simple memory system, *Science* **250**, 1223 (1990).
 - [9] R. Granger and G. Lynch, Higher olfactory processes: Perceptual learning and memory, *Curr. Opin. Neurobiol.* **1**, 209 (1991).
 - [10] L. B. Haberly, Parallel-distributed processing in olfactory cortex: New insights from morphological and physiological analysis of neuronal circuitry, *Chemical senses* **26**, 551 (2001).
 - [11] D. A. Wilson, A. R. Best, and R. M. Sullivan, Plasticity in the olfactory system: Lessons for the neurobiology of memory, *Neuroscientist* **10**, 513 (2004).
 - [12] C. Bushdid, M. O. Magnasco, L. B. Vosshall, and A. Keller, Humans Can Discriminate More than 1 Trillion Olfactory Stimuli, *Science* **343**, 1370 (2014).
 - [13] R. C. Gerkin and J. B. Castro, The number of olfactory stimuli that humans can discriminate is still unknown, *eLife* **4**, e08127 (2015).
 - [14] E. J. Mayhew, C. J. Arayata, R. C. Gerkin, B. K. Lee, J. M. Magill, L. L. Snyder, K. A. Little, C. W. Yu, and J. D. Mainland, *Drawing the Borders of Olfactory Space*, preprint (Neuroscience, 2020).
 - [15] A. Lansner, Associative memory models: From the cell-assembly theory to biophysically detailed cortex simulations, *Trends Neurosci.* **32**, 178 (2009).
 - [16] A. Mayer, V. Balasubramanian, T. Mora, and A. M. Walczak, How a well-adapted immune system is organized., *Proc. Natl. Acad. Sci. U.S.A.* **112**, 5950 (2015).
 - [17] R. Shinnakasu, T. Inoue, K. Kometani, S. Moriyama,

- Y. Adachi, M. Nakayama, Y. Takahashi, H. Fukuyama, T. Okada, and T. Kurosaki, Regulated selection of germinal-center cells into the memory B cell compartment, *Nat. Immunol.* **17**, 861 (2016).
- [18] R. Shinnakasu and T. Kurosaki, Regulation of memory B and plasma cell differentiation, *Curr. Opin. Immunol.* **45**, 126 (2017).
- [19] A. Mayer, V. Balasubramanian, A. M. Walczak, and T. Mora, How a well-adapting immune system remembers, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 8815 (2019).
- [20] O. H. Schnaack and A. Nourmohammad, Optimal evolutionary decision-making to store immune memory, *eLife* **10**, e61346 (2021).
- [21] C. Viant, G. H. J. Weymar, A. Escolano, S. Chen, H. Hartweger, M. Cipolla, A. Gazumyan, and M. C. Nussenzweig, Antibody affinity shapes the choice between memory and germinal center B cell fates, *Cell* **183**, 1298 (2020).
- [22] R. Polikar and C. Alippi, Guest Editorial Learning in Nonstationary and Evolving Environments, *IEEE Trans. Neural Netw. Learning Syst.* **25**, 9 (2014).
- [23] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar, Learning in Nonstationary Environments: A Survey, *IEEE Comput. Intell. Mag.* **10**, 12 (2015).
- [24] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) <http://www.deeplearningbook.org>.
- [25] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to Machine Learning for physicists, *Physics Reports* **810**, 1 (2019), arXiv: 1803.08823.
- [26] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (Wiley, New York, 1949).
- [27] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).
- [28] J. F. Fontanari and R. Meir, Learning noisy patterns in a Hopfield network, *Phys. Rev. A* **40**, 2806 (1989).
- [29] M. Mezard, J. P. Nadal, and G. Toulouse, Solvable models of working memories, *J. Physique* **47**, 1457 (1986).
- [30] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [31] R. McEliece, E. Posner, E. Rodemich, and S. Venkatesh, The capacity of the hopfield associative memory, *IEEE Transactions on Information Theory* **33**, 461 (1987).
- [32] J.-P. Bouchaud and M. Mezard, Universality classes for extreme-value statistics, *J. Phys. A Math. Gen.* **30**, 7997 (1997).
- [33] B. Derrida, From random walks to spin glasses, *Physica D: Nonlinear Phenomena* **107**, 186 (1997).
- [34] N. Parga and M. Virasoro, The ultrametric organization of memories in a neural network, *J. Phys. France* **47**, 1857 (1986).
- [35] M. A. Virasoro, Ultrametricity, Hopfield Model and all that, in *Disordered Systems and Biological Organization* (Springer, Berlin, Heidelberg, 1986) pp. 197–204.
- [36] H. Gutfreund, Neural networks with hierarchically correlated patterns, *Phys. Rev. A* **37**, 570 (1988).
- [37] M. V. Tsodyks, Hierarchical associative memory in neural networks with low activity level, *Mod. Phys. Lett. B* **04**, 259 (1990).
- [38] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Spin-glass models of neural networks, *Physical Review A* **32**, 1007 (1985).
- [39] J. F. Fontanari, Generalization in a hopfield network, *J. Phys. France* **51**, 2421 (1990).

4.1 Alternative learning rules

In the preprint [2], we focus on the standard Hebbian learning rule (eq. 2.9), both for the entire network and for individual compartments. Hopfield networks are among the most studied models for learning. Thus, it is no surprise that many other learning rules have been developed for these networks [144].

Here, we will discuss a number of these rules and assess their impact on our results. We focus on local and incremental learning rules, in which the updates of the weights only depend on one pattern at a time. Other learning rules, such as the Kraut-Mézard class [140], use the information from all patterns during each round of update, which makes them unrealistic for biological systems. With this condition in mind, we will introduce and discuss the consequences of the Storkey learning rule [143], the gradient descent learning rule [144], and the sparse Hebbian learning rule.

Storkey learning

Introduced by Storkey in 1997, this rule had the goal of increasing the capacity of Hopfield networks [143]. Indeed, Hopfield networks that are trained with this rule have a capacity of $L/\sqrt{2 \ln L}$ [143], which is significantly larger than the capacity of $L/(2 \ln L)$, reached with the conventional Hebbian learning [138].

This rule was originally designed for a learning phase in which all patterns are known. In that case, the update to the interaction matrix $J_{i,j}$ is given by

$$\Delta J_{i,j} = \begin{cases} \frac{1}{N} (\sigma_i - f_{i,j}) (\sigma_j - f_{ji}), & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (4.1)$$

Where $f_{i,j}$ is the local field on the spin σ_i except for the contribution from σ_j

$$f_{i,j} = \frac{1}{L-2} \left(\sum_k J_{i,k} \sigma_k - J_{i,i} \sigma_i - J_{i,j} \sigma_j \right). \quad (4.2)$$

To use this learning rule for consecutive encounters with (evolving) patterns, we add a learning rate λ to eq. 4.1 and then use the rule

$$\Delta J_{i,j} = \begin{cases} \lambda (\sigma_i - f_{i,j}) (\sigma_j - f_{ji}), & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (4.3)$$

Gradient descent

The discussion of gradient descent in the Hopfield model follows from ref. [144].

Gradient descent is the foundation of many optimization problems. With respect to the Hopfield model, we want to construct energy minima associated with the stored memory that are as deep as possible. For any given pattern σ this is achieved if

$W\sigma = \sigma$, where $W_{i,j} = \frac{1}{L-1}J_{i,j}$ is a normalized interaction matrix. In other words, we want to minimize the distance $D^{(p)}(\sigma, W\sigma)$ between the pattern σ and its projection $W\sigma$. Here, $D^{(p)}(\cdot, \cdot)$ is the distance measure for a general L^p -norm, which for the L^2 -norm follows,

$$D^{(2)}(\sigma, W\sigma) = \sum_i \left(\sigma_i - \sum_j W_{i,j}\sigma_j \right)^2. \quad (4.4)$$

The derivative of the L^2 distance with respect to the element $W_{i,j}$ is given by,

$$\frac{d}{dW_{i,j}} D^{(2)}(\sigma, W\sigma) = -2 \left(\sigma_i - \sum_k W_{i,k}\sigma_k \right) \sigma_j. \quad (4.5)$$

Therefore, we can define the gradient descent learning rule towards the energy minimum with learning rate (step size) λ that is consistent with minimization of the L^2 distance as,

$$\Delta J_{i,j} = \begin{cases} \lambda (\sigma_i - \sum_k W_{i,k}\sigma_k) \sigma_j, & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (4.6)$$

Interestingly, when using an L^1 -norm distance, the gradient descent learning rule is equivalent to the original Hebbian learning rule (eq. 2.9). This can be shown by first evaluating the derivative of the distance for the L^1 -norm with respect to the element $W_{i,j}$ as,

$$\begin{aligned} \frac{d}{dW_{i,j}} D^{(1)}(\sigma, W\sigma) &= -\frac{(\sigma_i - \sum_k W_{i,k}\sigma_k)}{\sqrt{(\sigma_i - \sum_k W_{i,k}\sigma_k)^2}} \sigma_j \\ &= -\text{sign} \left(\sigma_i - \sum_k W_{i,k}\sigma_k \right) \sigma_j \\ &= -\sigma_i \sigma_j \end{aligned} \quad (4.7)$$

where we used the fact that $|\sigma_i| \geq |\sum_k W_{i,k}\sigma_k|$, which results in $\text{sign}(\sigma_i - \sum_k W_{i,k}\sigma_k) = \text{sign}(\sigma_i) = \sigma_i$. Therefore, for the L^1 -norm distance, we recover the original Hebbian learning rule with $\Delta J_{i,j} = \lambda \sigma_i \sigma_j$.

Sparse Hebbian learning

Machine learning algorithms often enforce sparsity to regularize neural networks to avoid overfitting [126]. While a direct translation of such regularization to the Hopfield model is non-trivial, we can enforce sparsity on the interaction matrix $J_{i,j}$. To achieve a sparsity of X % in the interaction matrix $J_{i,j}$, we will use the standard Hebbian learning rule (eq. 2.9) and set entries $J_{i,j}$ with absolute values smaller than $(100 - X)$ % of all entries to zero.

Performance of networks with alternative learning rules

To characterize the impact of learning rules on our results, we perform the same optimization procedure as for the standard Hebbian learning in [2]. It should be noted that these alternative learning protocols are substantially more complex than the standard Hebbian learning, which limits our simulations to networks of maximum size $L = 100$ and $N = 8$ (in contrast to $L = 800$ and $N = 32$ in [2]). As we increase the ratio of N/L , we observe stronger finite-size effects. Still, we stay far below the capacity of the network and have no reason to expect any qualitative changes in the outcomes for networks of larger size.

In Fig. 4.1A we compare the performance of the networks trained with the standard Hebbian learning, to that of the alternative models, i.e., the Storkey, the gradient descent, and sparse Hebbian learning rules, in recognizing patterns evolving with a range of effective mutation rates μ_{eff} . For small sparsity (10%), the networks trained with sparse Hebbian learning perform similar to those with the standard Hebbian learning. However, when sparsity is large (50%), the sparse networks appear to lose the exact position of the minima. As a result, the system's performance systematically decays even for very slowly evolving patterns (small μ_{eff}). The Storkey and the gradient descent learning rules perform slightly better than the Hebbian learning for evolving patterns. This slight increase in performance is most likely a consequence of the increased capacity. A similar effect is seen in Fig. 2, as emptier (larger) networks perform better at a fixed effective mutation rate. However, the gain due to capacity is negligible compared to the reduction in performance of all networks with increasing mutation rate.

Aside from their similar performances, the distortion of the energy landscape for networks following the alternative learning rules are also comparable to that of the standard Hebbian learning. Specifically, we see that with all the alternative learning rules the misclassified patterns fall into attractors associated with one of the other pattern classes (Fig. 4.1B), consistent with the results for the standard Hebbian learning in Fig. S4A. Moreover, alternative learning rules also give rise to network structures in which the average number of open paths (i.e., number of beneficial spin-flips during equilibration) is smaller for stable (correctly classified) patterns compared to the unstable (misclassified) patterns, and both are smaller than for random patterns (Fig. 4.1C). This result is also similar to that of the standard Hebbian learning in Fig. S4B. It should be noted that the larger fluctuations seen in Figs. 4.1B,C compared to the Figs. S4A,B are due to the smaller system size used for simulations with the alternative learning rules.

In conclusion, the fraction of the correctly reconstructed patterns decay with increasing evolutionary rates due to the emergence of narrow passes in a network's energy landscape, irrespective of the choice of the learning rule. While we cannot exclude that other learning rules might achieve better performances, the results suggest that local (in time) learning rules which act on all network weights cannot learn multiple evolving patterns with maximal performance. Therefore, it is likely that the specialized

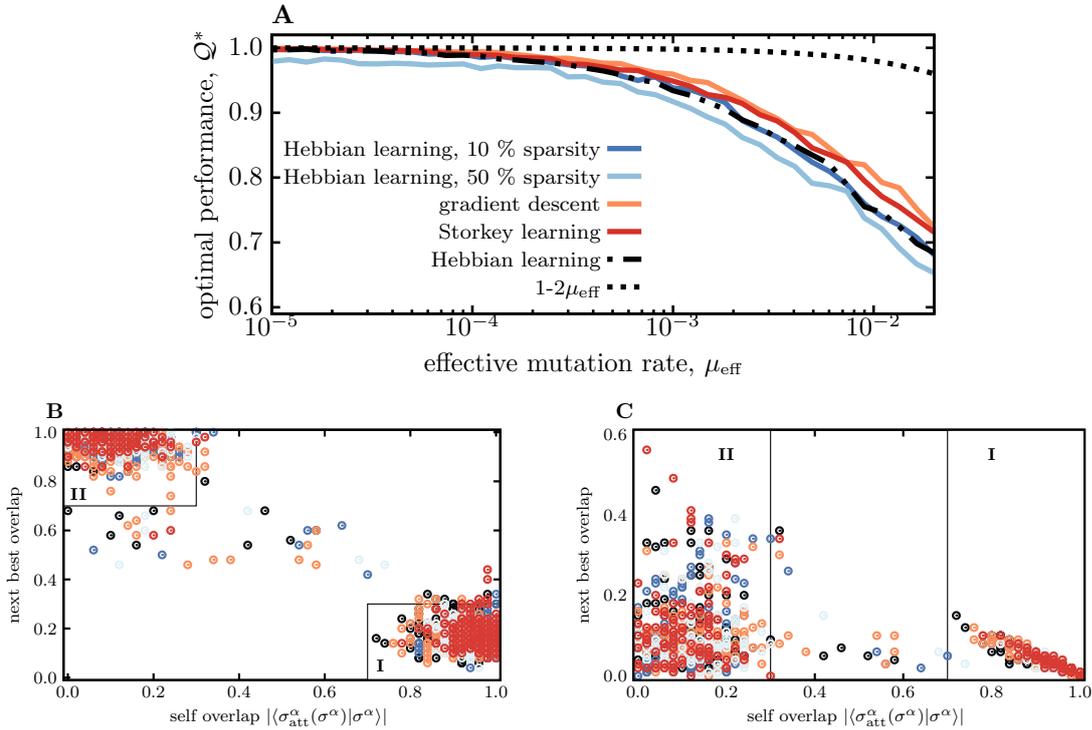


Figure 4.1 Alternative learning rules to train a Hopfield network. (A) The optimal performance of networks $Q^* = Q(\lambda^*)$ is shown as a function of the effective mutation rate $\mu_{\text{eff}} = N\mu$. The solid lines show the simulation results for networks with different learning rules (colors), indicated in the legend and introduced in Appendix D (colors). The dashed line shows the optimal performance for networks trained with the standard Hebbian learning (similar to Fig. 2A). The dotted line shows the naïve expectation for the performance solely based on the evolutionary divergence of the patterns ($1 - 2\mu_{\text{eff}}$). (B) The overlap between a presented pattern σ^α and the memory associated with the same pattern class $\sigma_{\text{att}}^\alpha(\sigma^\alpha)$ is shown against the overlap of the pattern with the next best memory attractor associated with any of the other presented pattern classes $\max_{\nu \neq \alpha} |\langle \sigma_{\text{att}}^\alpha | \sigma^\nu \rangle|$. Different colors indicate the alternative learning rules, according to the color code in (A). Similar to Fig. S4A, patterns are either correctly reconstructed (sector I), or they fall into an attractor associated with one of the other stored patterns (sector II). (C) The number of beneficial spin-flips for a presented pattern at the beginning of equilibration (i.e., the number of open equilibration paths) is shown against patterns' self-overlap (horizontal axis similar to (A)). For stable patterns (sector I), the number of open paths is anti-correlated with the overlap between the attractor and the presented pattern, consistent with the standard Hebbian learning in Fig. S4B. For unstable patterns (sector II), the number of open paths is on average larger than that of the stable patterns but smaller than for random patterns, shown in Fig. S4B. Simulation parameters: $L = 100$, $N = 8$; $\mu_{\text{eff}} = 0.02$ (for B and C).

1-to-1 memory to remain the only strategy that effectively learns and recovers evolving patterns.

Chapter 5

Risk-utility tradeoff shapes memory strategies for evolving patterns

The following section consists of the preprint

O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Risk-utility tradeoff shapes memory strategies for evolving patterns”. In: *arXiv:2110.15008 [physics]* (Oct. 2021). arXiv: 2110.15008.

The corresponding supplementary figures are shown in appendix C.

This paper is the second part of the collaboration with Luca Peliti. However, it should not be seen as an extension of the first project [2]. Instead, it is a different approach to the problem that uses a related method.

Similar to the first project, we study different memory strategies for a set of evolving patterns. In contrast to [2], we do not use a network structure that stores an associative memory. Instead, our system resembles a repertoire of specialized memories. In this setup, we investigate how the repertoire should be updated to follow the evolution of patterns.

In the context of this thesis, this project builds on the result that memory for evolving environments should be specialized [2]. As a result, we assume that different specialized memories (compartments) can reconstruct the stored states. Here, we study what the optimal resource allocation to these different compartments is.

Similar to [2], all authors did the conceptualization for the project. I performed the analytic and numerical analysis and produced all figures. All authors contributed to the writing and editing of the manuscript.

Risk-utility tradeoff shapes memory strategies for evolving patterns

Oskar H Schnaack,^{1,2} Luca Peliti,³ and Armita Nourmohammad^{1,2,4,*}

¹*Max Planck Institute for Dynamics and Self-organization, Am Fassberg 17, 37077 Göttingen, Germany*

²*Department of Physics, University of Washington,
3910 15th Ave Northeast, Seattle, WA 98195, USA*

³*Santa Marinella Research Institute, 00058 Santa Marinella, Italy*

⁴*Fred Hutchinson Cancer Research Center, 1100 Fairview ave N, Seattle, WA 98109, USA*

(Dated: October 28, 2021)

Keeping a memory of evolving stimuli is ubiquitous in biology, an example of which is immune memory for evolving pathogens. However, learning and memory storage for dynamic patterns still pose challenges in machine learning. Here, we introduce an analytical energy-based framework to address this problem. By accounting for the tradeoff between utility in keeping a high-affinity memory and the risk in forgetting some of the diverse stimuli, we show that a moderate tolerance for risk enables a repertoire to robustly classify evolving patterns, without much fine-tuning. Our approach offers a general guideline for learning and memory storage in systems interacting with diverse and evolving stimuli.

I. INTRODUCTION

Biological systems, ranging from the brain to the immune system, store memory of molecular interactions to efficiently recognize and respond to stimuli. Memory encoding in biological networks has also inspired a growing host of algorithms for learning and memory storage in image and pattern recognition by artificial neural networks [1–3]. A critical step in these algorithms is to find regularities in data to associate related patterns with each other. As such, these learning algorithms often assume that the set of training data comes from a stationary distribution that represents the regularities necessary for pattern recognition in data well.

Memory recognition, however, is not limited to static patterns and can be desirable when classifying evolving stimuli that drive the system out of equilibrium. One such example is the adaptive immune system in which memory can effectively recognize evolved variants of previously encountered pathogens [4–8]. In a recent work, we have demonstrated that distributed learning strategies, which are desirable for pattern recognition in the stationary setup, can fail to reliably learn and classify dynamically evolving patterns [9]. Specifically, we showed that to follow evolving patterns, an energy-based Hopfield-like neural network [10] should use a higher learning rate, which in turn, can distort the energy landscape associated with the stored memory attractors, leading to pattern misclassification. To remedy this problem, we proposed compartmentalized networks as the optimal solutions to memory storage for evolving patterns [9].

Irrespective of the network structure, an increase in learning rate is necessary for a network to follow, recognize, and store effective memory of evolving patterns [9]. Increasing learning rate leads to a risky strategy, as the

memory repertoire begins to reflect only the most recently encountered pattern while effectively destroying the memory of the prior encounters. Here, we present an analytical approach to explore how the tradeoff between utility and risk can determine learning strategies of a repertoire in keeping a memory of evolving patterns. We show that a moderate risk tolerance enables a repertoire to store an effective and robust memory. Our approach puts forward a guideline for optimal learning and memory storage for systems interacting with multiple evolving pattern classes without much fine-tuning.

II. MODEL

To probe memory strategies, we define a repertoire \mathcal{M} that can store stimuli (patterns) ψ and later utilize them to recognize newly presented stimuli. We consider the space of possible binary patterns of length L , $\{\psi^\alpha\}$ with $\alpha \in 1, \dots, 2^L$ enumerating all unique patterns, with entries $\psi_i^\alpha = \pm 1$, for $i \in 1, \dots, L$. The memory repertoire \mathcal{M} associates normalized weights m^α ($\sum m^\alpha = 1$) to all patterns it encounters. The non-zero weights reflect the relative importance of the stored memory from different stimuli, and patterns that are not encountered are associated with a zero weight. After an encounter with a given stimulus ψ^β at time t , all weights are updated according to a Hebbian learning rule $m^\alpha(t) = (1 - \lambda)m^\alpha(t - 1) + \delta_{\alpha,\beta}\lambda$, where λ is the learning rate [11]. Thus, the weight $m^\alpha(t + \tau)$ associated with pattern α , which was previously encountered at time t , decays in an approximate exponential form over time,

$$m^\alpha(t + \tau) = \lambda(1 - \lambda)^\tau \approx \lambda e^{-\lambda\tau}. \quad (1)$$

To determine the memory performance for a given stimulus χ , we define the affinity $\mathcal{A}(\mathcal{M}, \chi)$ as the overlap between the patterns stored in the memory repertoire \mathcal{M}

* Correspondence should be addressed to Armita Nourmohammad: armita@uw.edu

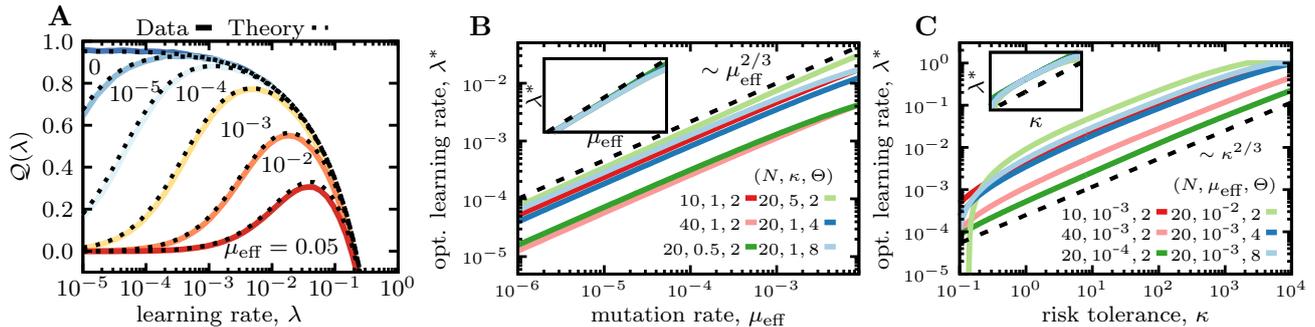


FIG. 1. **Optimal memory strategy for evolving patterns.** (A) The objective function $\mathcal{Q}(\lambda)$ (Eq. 3) is shown as a function of the learning rate, for $\kappa = 1$ and $\Theta = 2$. Analytical approximations (dashed lines) are compared to numerical estimates (full lines), for different mutation rates μ_{eff} (colors). (B, C) The optimal learning rate (Eq. 5) is shown as a function of the mutation rate μ_{eff} (B) and the risk tolerance κ (C), for different sets of parameters indicated in the legend. Insets show collapsed plots after rescaling according to Eq. 5.

and the pattern χ ,

$$\mathcal{A}(\mathcal{M}, \chi) = \mathcal{A}_0 \sum_{\psi^\alpha \in \mathcal{M}} m^\alpha |\langle \psi^\alpha | \chi \rangle|^\Theta - \mathcal{A}_{\text{rand}}, \quad (2)$$

Here, we used a short-hand notation to denote the normalized pattern vector by $|\psi\rangle \equiv \psi/\sqrt{L}$, its transpose by $\langle\psi|$, and a normalized scalar product $\langle\psi|\chi\rangle \equiv (\sum_i \psi_i \chi_i)/L$. We use the shape parameter Θ to modulate the dependency of the affinity function on pattern overlap. The scaling parameter \mathcal{A}_0 is model dependent and sets the unit of affinity in a given system, yet its precise value does not impact our analysis. The offset $\mathcal{A}_{\text{rand}}$ is chosen such that the expected affinity of random patterns remains zero. The choice of affinity as a measure of performance is inspired by memory retrieval in biological systems, where recognition is mediated through biophysical interactions.

We characterize the response of a memory repertoire to N independently evolving pattern classes Ψ^c (with $c \in 1, \dots, N$). Each class Ψ^c denotes a set $\{\psi\}_c$ of patterns generated over time through evolution. Patterns within each class evolve (mutate) by random spin-flips of rate μ . At each time step t , a pattern from a randomly chose class is presented to the memory repertoire, resulting in an effective observed mutation rate of $\mu_{\text{eff}} = N\mu$ per expected encounter with the same pattern class. To simplify, we use $\Psi^c(t)$ to denote the evolved pattern from class c that is generated and presented to the repertoire at time t . Since pattern classes are orthogonal to each other (up to finite-size effects), the expected overlap between patterns presented at different times follows $\langle\langle\Psi^c(t)|\Psi^{c'}(t+\tau)\rangle\rangle = \delta_{c,c'}\rho^\tau + \mathcal{O}(1/\sqrt{L})$, where $\rho = (1 - 2\mu)$ measures the similarity between evolved patterns.

Interestingly, for $\Theta = 2$ this model is equivalent to the energy function of the classical Hopfield network with Hebbian learning rule [9, 10, 12, 13] (Appendix B). This correspondence enables us to simulate the memory repertoire efficiently and to test analytic predictions with

numerical experiments; see Appendix C for numerical method.

III. RESULTS

A. Optimal learning for evolving patterns with risk-return tradeoff

We seek to find an optimal strategy to set the learning rate such that the stored memory in a repertoire can be reliably retrieved for patterns evolving at a specified rate μ . Learning and updates of memory repertoires over time impact both the expected affinity and the variance of the affinity across patterns. A high learning rate can sustain a high affinity in a repertoire for the most recently presented patterns. However, keeping up with the latest trend can be risky as the variance of the affinity across patterns can increase, with older patterns suffering most from this tradeoff. To account for this effect, we optimize an objective function $\mathcal{Q}(\lambda; \mu)$ that balances the risk-utility tradeoff by maximizing the mean (utility) $\langle\mathcal{A}\rangle$ and minimizing the standard deviation (risk) $\sigma_{\mathcal{A}}$ of the affinity across patterns,

$$\mathcal{Q}(\lambda; \mu) = \langle\mathcal{A}\rangle - \frac{1}{\kappa} \sigma_{\mathcal{A}}. \quad (3)$$

Here, κ measures the risk tolerance of the repertoire. Such a risk-utility analysis was initially introduced in economics [14, 15], but has since been used to characterize tradeoffs in biological and evolutionary processes [16–22].

We can analytically evaluate both the mean and the variance of the affinity by using explicit expansion of the encounter history or by evaluating the cumulant-generating function for the distribution of affinities, valid in the large- N limit (i.e, many patterns); see Appendix A and Fig. S1. The expression for the n^{th} cumulant of the

affinity follows,

$$c_n = a_0^n \frac{\lambda^n \rho^{\Theta n}}{N(1 - (1 - \lambda)^n \rho^{\Theta n})} + \mathcal{O}(N^{-2}) \quad (4)$$

where $a_0 = \mathcal{A}_0 - \mathcal{A}_{\text{rand}}$, and the mean and the variance are given by $n = 1$ and $n = 2$, respectively. For static patterns ($\mu = 0 \leftrightarrow \rho = 1$), the mean affinity becomes independent of the learning rate and reaches the naive expectation a_0/N . For evolving patterns, a repertoire with a maximal learning rate $\lambda = 1$ can achieve the highest mean affinity $a_0 \rho^{\Theta}/N$ by storing a memory of the most recent pattern with the maximal affinity a_0 , while treating the other patterns as random.

For a broad range of evolutionary rates, the optimum of the objective function $\mathcal{Q}(\lambda; \mu)$ is achieved for intermediate values of learning rate λ , where the mean and variance of affinity across patterns are balanced; see Fig. 1A for analytical and numerical results. The optimal learning rate $\lambda^*(\mu) = \max_{\lambda} \mathcal{Q}(\lambda; \mu)$ that maximizes the objective function scales with different model parameters as,

$$\lambda^* = \frac{2}{N} (2\kappa \Theta \mu_{\text{eff}})^{2/3} + \mathcal{O}\left(\frac{\mu_{\text{eff}}}{N}\right). \quad (5)$$

The analytical scaling relation in Eq. 5 is in excellent agreement with numerical simulations (Fig. 1B,C). As the evolutionary rate of patterns μ_{eff} increases, the optimal learning rate grows so that the repertoire closely follows the patterns' evolution (Fig. 1B). As the shape parameter Θ increases, the affinity function becomes more peaked around the recently stored patterns (Eq. 2), resulting in an increase in the optimal learning rate to keep the memory focused on the more recent (less evolved) encounters. The learning rate scales inversely with the number of patterns N for the repertoire to evenly distribute the resources (allocate memory). Repertoires with higher risk tolerance κ use larger learning rates (Fig. 1C) and store a risky but a high-affinity memory against recent encounters. As repertoire becomes more risk-averse (small κ), it stores a more equitable memory across patterns but at a loss for affinity. In the limit of no risk tolerance ($\kappa \rightarrow 0$), the repertoire stops learning ($\lambda = 0$) and adopts a risk-free but impractical strategy where the memory has zero affinity for all patterns. This tradeoff can be depicted by a Pareto front in the affinity-risk space, along which one cannot increase the mean affinity without increasing the risk or vice versa. Fig. 2 shows these Pareto fronts parametrized by scaled mean affinity and risk, for different mutation rates (colors) and by varying the risk tolerance κ along each line. The combinations of risk and affinity values that lie below the Pareto front are inaccessible, and those that lie above are sub-optimal solutions for a memory repertoire.

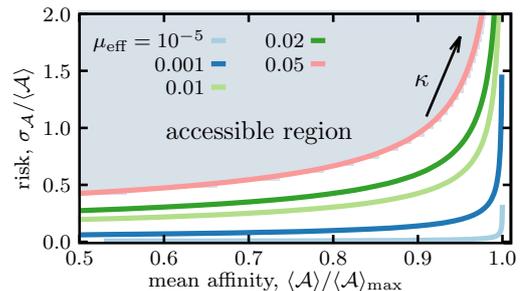


FIG. 2. **Pareto front for risk-utility tradeoff of memory for evolving patterns.** Optimized objective function $\mathcal{Q}(\lambda^*)$ (Eq. 3) describes a Pareto front in the scaled affinity-risk plane by varying the risk tolerance κ along each line, for different mutation rates (colors). To present the Pareto front in a dimensionless space, the mean affinity is scaled by its maximal value $\langle \mathcal{A} \rangle_{\text{max}} = \frac{a_0}{N}$, and the standard deviation as risk is scaled by the mean affinity. Only the combinations of risk and affinity values that lie above the Pareto front (gray shading) are accessible to a memory repertoire. Other parameters: $N = 200$ and $\Theta = 2$. Results for other shape parameters are shown in Fig. S2.

B. Discrimination of random and stored patterns

One key goal of memory usage is to recognize and accurately classify presented patterns with a memory stored from prior encounters. A misclassification could have dire consequences. For example, in the case of immune system, if memory response is not triggered by a secondary infection (i.e., *false negative*), the host would pay a cost by enduring sickness and having to mount a novel response and re-store a new memory. *False positive* responses are also costly, as they can be associated with autoimmunity if mounted against self-antigens [6], or they can interfere with novel responses without preventing the disease, e.g., in the case of original antigenic sin against viruses like influenza [23].

Memory strategies optimized to operate under different risk tolerance κ can yield varying levels of pattern misclassification. We characterize the discrimination accuracy of a repertoire by quantifying the rate by which it recognizes evolved patterns associated with previously stored memory (true positive), or randomly generated patterns without any prior encounter history (false positive). To do so, we need to characterize the distribution of affinities for patterns with prior encounter history, and for novel (random) patterns.

The recognition affinity of a memory repertoire for recurring patterns depends on the history of pattern encounters and on the learning rate λ . For example, in the case where patterns from a given class Ψ^c are presented at time points $[t_1, \dots, t_n]$, the affinity $\mathcal{A}(\Psi^c(T))$ for an evolved pattern from the same class shown at time

$T > t_n$ can be expressed as

$$\mathcal{A}(\Psi^c(T)) = a_0 \sum_{i=1}^n \lambda e^{-\lambda(T-t_{i-1})} \rho^{\Theta(T-t_i)} \quad (6)$$

Here, the factor $e^{-\lambda(T-t_{i-1})}$ accounts for the exponential decay for the affinity of a memory stored at time t_i from its maximum level λ , due to updates in the repertoire (Eq. 1). $\rho^{\Theta(T-t_i)}$ accounts for the decay in the overlap between the memory stored at time t_i and the presented patterns in the future, due to evolution.

Patterns are presented to the repertoire in random order, and the time $\tau_i = t_i - t_{i-1}$ between consecutive encounters with the same class is exponentially distributed with a mean of N steps (i.e., number of classes), $p(\tau) \approx e^{-N\tau}/N$. The distribution of affinities $P_\Psi(\mathcal{A})$ for patterns of a given class can be derived by convolving the affinity function in Eq. 6 with the exponential waiting time distribution for pattern history. Although the exact form of this affinity distribution is difficult to evaluate, we expect it to be from an exponential family. Indeed, the Gamma distribution is a good approximation to the distribution of the affinities (Fig. S3). We quantify the accuracy of this approximation by the Kullback-Leibler distance $D_{\text{KL}}(P_\Psi(\mathcal{A})\|\Gamma_{\mathcal{A}})$ between the true affinity distribution $P_\Psi(\mathcal{A})$ and a Gamma distribution $\Gamma_{\mathcal{A}}$ with the same mean and variance. By using an Edgeworth expansion [24, 25] with the cumulants of the affinity distribution in Eq. 4, we can show that in the limit of small learning and mutation rates the Kullback-Leibler distance between the two distributions is small, with $D_{\text{KL}}(P_\Psi(\mathcal{A})\|\Gamma_{\mathcal{A}}) = (25/27)(2\Theta\mu_{\text{eff}}/\kappa)^{2/3} + \mathcal{O}(\Theta\mu_{\text{eff}}) \ll 1$ (Appendix A 4). This result can also be intuitively understood, since the Gamma distribution arises in processes for which the waiting times between events are relevant.

The distribution of affinities for random patterns (i.e., not belonging to any of the presented classes) $P_0(\mathcal{A})$ can be similarly characterized. The law of large numbers suggest that the overlap between unrelated patterns should be normally distributed with mean zero and variance $1/(4L)$. The overlap between the memory repertoire and a random pattern is the sum of the normally distributed random overlaps weighed by the weights $\{m^\alpha\}$, associated with the stored patterns (Appendix A 5 and Fig. S4). Thus, the distribution of affinities for random patterns $P_0(\mathcal{A})$ is well-approximated by a Gamma distribution with mean $\propto L^{-\Theta/2}$ and variance $\propto L^{-\Theta}$. This distribution should be contrasted to that of the affinities for patterns with prior encounter histories $P_\Psi(\mathcal{A})$, the statistics of which primarily depends on the number of patterns and the learning rate, with mean $\propto 1/N$, and variance $\propto \lambda/N$.

The sensitivity (fraction of true positives to all patterns associated with memory) and specificity (fraction of false negatives to all random patterns) of a repertoire in classifying patterns depends on an affinity cut-off for distinguishing between familiar and random patterns. A

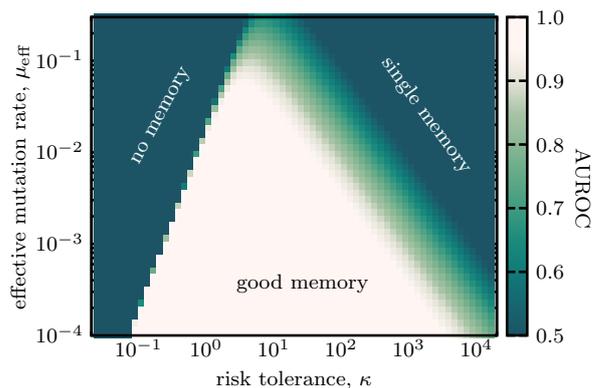


FIG. 3. Three phases of memory discrimination. The phase diagram shows the discrimination ability of repertoires (AUROC) between familiar patterns with prior encounter history and random patterns, based on their respective affinities. AUROC is estimated using numerical approximations to the affinity distributions $P_\Psi(\mathcal{A})$ and $P_0(\mathcal{A})$ for familiar and random patterns (Appendix A). Each point in the phase diagram shows AUROC for a repertoire optimized with a given risk tolerance κ for patterns with a specified mutation rate μ_{eff} . Other parameters: $L = 200$, $N = 40$, and $\Theta = 2$. Results for other shape parameters Θ are shown in Fig. S6. Full simulation results are shown in Fig. S7.

receiver operating characteristic (ROC) curve shows the relationship between sensitivity and specificity of a classification in a repertoire for every possible affinity cut-off. The area under the ROC curve (AUROC), which measures the discriminative ability of the repertoire, depends on the evolutionary rate μ_{eff} of patterns it encounters, its risk tolerance κ , and the parameter Θ that determines the shape of the affinity function (Eq. 4). These parameters also determine the optimal learning rate λ^* (Eq. 5), which sets the memory strategy of a repertoire in a given evolutionary setup.

The phase diagram for the discrimination ability of a repertoire defines three distinct regions determined by a combination of risk tolerance κ and the evolutionary rate of patterns μ_{eff} , for a specified shape parameter Θ (Fig. 3): (i) A triangular region at the center of the phase diagram with $\text{AUROC} \simeq 1$ indicates the range of parameters for which the repertoire can efficiently discriminate between familiar and novel patterns. Within this region, as the evolutionary rate of patterns increases, the range for risk tolerance that enables a repertoire to efficiently discriminate between familiar and random patterns narrows, and the repertoire has to fine-tune its learning rate to match the faster evolution of the patterns (Fig. S5). (ii) When risk tolerance κ is large, the repertoire updates rapidly and only keeps a memory of the most recent encounters, resulting in an inefficient discrimination between familiar and novel patterns with $\text{AUROC} \simeq 0.5$. (iii) For risk avert strategies (small κ), the memory effectively shuts down and the repertoire cannot anymore

discriminate between random and familiar patterns. It should be noted that a change in the affinity shape parameter Θ can shift the exact boundaries between these phases, but it does not impact the overall structure of the phase diagram (Fig. S6).

Taken together, a moderate risk tolerance ($\kappa\Theta = \mathcal{O}(1)$) enables a repertoire to store an effective and a robust memory and to operate without fine tuning. We further confirm these results with simulations for $\Theta = 2$ in Fig. S7.

IV. DISCUSSION

Devising a learning strategy to store functional memory for evolving stimuli is an open problem with potential applications in many fields. Inspired by immune memory that reliably recognizes evolving threats, we present an analytical energy-based memory model against evolving patterns that captures the risk-utility tradeoff in memory repertoires.

We found that risk-tolerant repertoires adopt faster learning rates and keep a memory of the recent patterns with high affinity, at the risk of disregarding older patterns. On the other hand, risk-averse repertoires effectively shut down their learning to minimize the variance in their recognition affinity for different patterns, at the cost of having a low affinity for all patterns. A moderate risk tolerance enables a repertoire to achieve a desirable balance between risk and affinity (utility) to robustly classify evolving patterns and associate them with prior memory. This risk-utility tradeoff defines Pareto front for accessible memory strategies by a repertoire (Fig. 2).

One key role of memory storage is to reliably discriminate between familiar and novel stimuli. The discrimination ability of an optimal memory repertoire depends on its risk tolerance and the evolutionary rate of the patterns that it encounters. Interestingly, as the evolutionary rate of the presented patterns increases, the range of risk tolerance that allows a repertoire to retain a functional memory narrows and approaches a regime where risk and utility are equally valued. The classification criteria in our analysis are solely based on the affinity of the memory repertoire to a presented pattern. This stands in contrast to energy-based Hopfield-like networks that achieve recognition by retrieving associative memory stored in the networks' energy minima [10].

We have used the fluctuations (standard deviation) of the memory's affinity over the ensemble of presented patterns to measure the risk of misclassification by the stored storage. While variance might be more suitable in other settings [26, 27], using standard deviation as a measure of risk keeps the risk tolerance κ dimensionless and comparable across different systems. Nonetheless, we expect the Pareto front's overall structure for the risk-utility tradeoff and the phase diagram for discrimination ability of the repertoires to remain qualitatively intact, irrespective of the exact choices made for the risk function.

Although our model is inspired by immune memory repertoires, the introduced approach is general enough that can be applied to other models of memory. In particular, incorporating risk and utility in training of artificial neural networks can guide these machine learning efforts to store efficient of memory for evolving signals [28].

ACKNOWLEDGEMENT

This work has been supported by the NSF CAREER award (grant No: 2045054), DFG grant (SFB1310) for Predictability in Evolution and the MPRG funding through the Max Planck Society. O.H.S also acknowledges funding from Georg-August University School of Science (GAUSS) and the Fulbright foundation.

Appendix A: Statistics of the affinity distribution

1. The average affinity of a repertoire for presented patterns

In the main text (Eq. 6) we express the affinity of a pattern class at time T in terms of the prior encounters with the patterns of the same class at times $[t_1, \dots, t_n]$ with $t_{n-1} < t_n < T$. To compute the statistics (mean and variance) of a repertoire's affinity, we first express the affinity function in terms of the times $\hat{t}_i = T - t_{n-i+1}$, passed since the i^{th} encounter, when counting backward in time. This allows us to write the affinity as

$$\mathcal{A}(\Psi^c(T)) = a_0 \lambda \sum_{i=1}^n e^{-\lambda \hat{t}_i - 1} \rho^{2\hat{t}_i}. \quad (\text{S1})$$

Similar to forward times t_i , the reverse times \hat{t}_i are separated by exponentially distributed independent waiting times $\hat{\tau}_i = \hat{t}_i - \hat{t}_{i-1}$,

$$p(\hat{\tau}) = \frac{(N-1)^{\hat{\tau}-1}}{N^{\hat{\tau}}} \approx \frac{1}{N} e^{-N\hat{\tau}}. \quad (\text{S2})$$

Given the relationship $\sum_{j=1}^i \hat{\tau}_j = \hat{t}_i$, we can express the affinity of the patterns in Eq. S1 in terms of the statistically independent $\hat{\tau}_i$ as,

$$\mathcal{A}(\Psi) = a_0 \lambda \sum_{g=1}^n (1-\lambda)^{\sum_{g'=1}^g \hat{\tau}_{g'} - 1} \rho^{\Theta \sum_{g'=1}^g \hat{\tau}_{g'}}. \quad (\text{S3})$$

Here, we can interpret each term in the sum as the contribution of a memory from the corresponding prior generation to the affinity at the current time. These contributions decay due to the updates (learning) in the repertoire with rate $(1-\lambda)$ and due to evolution of the patterns with rate ρ . The expected affinity contribution $\langle \mathcal{A}_g \rangle$ from the g^{th} generation follows,

$$\begin{aligned} \langle \mathcal{A}_g \rangle &= a_0 \lambda \sum_{\hat{\tau}_1} \cdots \sum_{\hat{\tau}_g} \left(\prod_{i=1}^g p(\hat{\tau}_i) \right) (1-\lambda)^{\sum_{i=1}^g \hat{\tau}_i - 1} \rho^{\Theta \sum_{i=1}^g \hat{\tau}_i} \\ &= a_0 \lambda (1-\lambda)^{-1} \prod_{i=1}^g \left(\sum_{\hat{\tau}_i} p(\hat{\tau}_i) (1-\lambda)^{\hat{\tau}_i} \rho^{\Theta \hat{\tau}_i} \right) \\ &= a_0 \lambda (1-\lambda)^{-1} \left(\frac{(1-\lambda)\rho^{\Theta}}{N + (1-\lambda)\rho^{\Theta} - N(1-\lambda)\rho^{\Theta}} \right)^g \end{aligned} \quad (\text{S4})$$

where we used the fact that the time windows $\hat{\tau}_i$ are independent from each other. The expected affinity follows from adding up the contributions from all generations,

$$\langle \mathcal{A} \rangle = \sum_{g=1}^{\infty} \langle \mathcal{A}_g \rangle = a_0 \frac{\lambda \rho^{\Theta}}{N(1 - (1-\lambda)\rho^{\Theta})}. \quad (\text{S5})$$

As mentioned in the main text, this result immediately shows that the system reaches the maximal mean affinity of $a_0/N\rho^{\Theta}$ for the maximal learning rate $\lambda = 1$. Moreover, the mean affinity becomes independent of the learning rate when the patterns are static ($\rho = 1$).

2. The variance of repertoire's affinity across presented patterns

To calculate the variance of the affinity (Eq. S4 in the main text) we need to account for the covariance $\text{cov}(\mathcal{A}_{g_1}, \mathcal{A}_{g_2})$ between the contributions from different generations. Because the covariance is symmetric, we can write the variance of the affinity as

$$\text{var}(\mathcal{A}) = \sum_{g=1}^{\infty} \text{var}(\mathcal{A}_g) + 2 \sum_{g_1=1}^{\infty} \sum_{g_2=g_1+1}^{\infty} \text{cov}(\mathcal{A}_{g_1}, \mathcal{A}_{g_2}). \quad (\text{S6})$$

First we calculate the variance of the individual terms \mathcal{A}_g . Following the notation introduced in Eq. S4, we find

$$\begin{aligned}\mathbf{var}(\mathcal{A}_g) &= \mathbf{var} \left(a_0 \lambda (1 - \lambda)^{-1} \prod_{i=1}^g ((1 - \lambda)^{\hat{\tau}_i} \rho^{\Theta \hat{\tau}_i}) \right) \\ &= \mathbf{var} \left(a_0 \lambda (1 - \lambda)^{-1} \prod_{i=1}^{g-1} ((1 - \lambda)^{\hat{\tau}_i} \rho^{\Theta \hat{\tau}_i}) \left((1 - \lambda)^{\hat{\tau}_g} \rho^{\Theta \hat{\tau}_g} \right) \right)\end{aligned}\quad (\text{S7})$$

$$\begin{aligned}&= \langle (1 - \lambda)^{\hat{\tau}_g} \rho^{\Theta \hat{\tau}_g} \rangle^2 \mathbf{var} \left(a_0 \lambda (1 - \lambda)^{-1} \prod_{i=1}^{g-1} ((1 - \lambda)^{\hat{\tau}_i} \rho^{\Theta \hat{\tau}_i}) \right) \\ &\quad + \left\langle a_0 \lambda (1 - \lambda)^{-1} \prod_{i=1}^{g-1} ((1 - \lambda)^{\hat{\tau}_i} \rho^{\Theta \hat{\tau}_i}) \right\rangle^2 \mathbf{var} \left((1 - \lambda)^{\hat{\tau}_g} \rho^{\Theta \hat{\tau}_g} \right)\end{aligned}\quad (\text{S8})$$

$$= (a_0 \lambda (1 - \lambda)^{-1})^{-2} \left(\langle \mathcal{A}_1 \rangle^2 \mathbf{var}(\mathcal{A}_{g-1}) + \langle \mathcal{A}_{g-1} \rangle^2 \mathbf{var}(\mathcal{A}_1) \right)\quad (\text{S9})$$

$$= g (a_0 \lambda (1 - \lambda)^{-1})^{-(2g-2)} \langle \mathcal{A}_1 \rangle^{2g-2} \mathbf{var}(\mathcal{A}_1).\quad (\text{S10})$$

Because the time intervals $\hat{\tau}_i$ are independent from each other, we could use error propagation to get from Eq. S7 to Eq. S8. Moreover, since all time intervals follow the same distribution (Eq. S2), we could insert the statistics of \mathcal{A}_g with the correct normalization $(a_0 \lambda (1 - \lambda)^{-1})^{-2}$ to arrive at Eq. S9. Finally, we performed $g - 1$ further iterations to get a result that only depends on the statistics of the first contribution \mathcal{A}_1 in Eq. S10. Thus, it is sufficient to calculate the variance of only one generation,

$$\mathbf{var}(\mathcal{A}_1) = a_0^2 \rho^{2\Theta} \left(\frac{1}{N - (1 - \lambda)^2 (N - 1) \rho^{2\Theta}} - \frac{1}{(N - (1 - \lambda)(N - 1) \rho^\Theta)^2} \right).\quad (\text{S11})$$

To evaluate the variance of the affinity (Eq. S6), we still need to calculate the covariance between the contributions from different generations. Similar to Eq. S7, we use,

$$\mathcal{A}_g = \mathcal{A}_{g-i} \cdot \prod_{j=g-i}^g ((1 - \lambda)^{\hat{\tau}_j} \rho^{\Theta \hat{\tau}_j}).\quad (\text{S12})$$

which entails the following relationship for the covariance between \mathcal{A}_g and \mathcal{A}_{g-i} ,

$$\begin{aligned}\mathbf{cov}(\mathcal{A}_g, \mathcal{A}_{g-i}) &= \mathbf{cov} \left(\mathcal{A}_{g-i} \cdot \prod_{j=g-i}^g ((1 - \lambda)^{\hat{\tau}_j} \rho^{\Theta \hat{\tau}_j}), \mathcal{A}_{g-i} \right) = \mathbf{var}(\mathcal{A}_{g-i}) \left\langle \prod_{j=g-i}^g ((1 - \lambda)^{\hat{\tau}_j} \rho^{\Theta \hat{\tau}_j}) \right\rangle \\ &= (a_0 \lambda (1 - \lambda)^{-i})^{-1} \mathbf{var}(\mathcal{A}_{g-i}) \langle \mathcal{A}_1 \rangle^i.\end{aligned}\quad (\text{S13})$$

Here, we again used the fact that all time intervals are independent and that they follow the same distribution in Eq. S2. Using the expressions for variance of individual terms (Eq. S10) and the covariance between them (Eq. S13), we can characterize the variance of the affinity (Eq. S6) as,

$$\mathbf{var}(\mathcal{A}) = a_0^2 \left[\frac{\lambda(N - 1) \rho^{2\Theta} (N - (1 - \lambda)(N - 1) \rho^\Theta)^2 (\lambda N - (1 - \lambda) \rho^\Theta (\lambda(N + 2) - 2))}{N^2 (1 - (1 - \lambda) \rho^\Theta) (N - (1 - \lambda)(N - 2) \rho^\Theta)^2 (N - (1 - \lambda)^2 (N - 1) \rho^{2\Theta})} \right].\quad (\text{S14})$$

Note that the expression in Eq. S14 is accurate up to terms of order $a_0^2 \mathcal{O}(N^{-1} L^{-\Theta})$, which arise from the (negligible) overlap between patterns of difference classes due to the finite size of the patterns; see Appendix A 5.

3. Cumulant-generating function for the affinity distribution

To characterize the distribution of the affinities, we can rely on the corresponding cumulant-generating function. The cumulant-generating function $R_X(q)$ of a random variable X is defined as the logarithm of the moment-generating function

$$R_X(q) = \log \langle e^{qX} \rangle.\quad (\text{S15})$$

If we define a new random variable $Y = \sum_i X_i$ as the sum of independent random variables, its cumulant-generating function is given by $\mathcal{R}_Y(q) = \sum_i \mathcal{R}_{X_i}(q)$. To use this relation for the affinity function, we need to rewrite the affinity of a pattern class (Eq. S1) as a sum of independent random variables. In Eq. S1 we calculate the affinity as the sum over the contributions of past encounters. In that view, the times of the past encounters \hat{t}_i are random variables while the contributions to the affinity at those times are deterministic.

We now change the point of view and write the affinity function as sum of contributions from all time, instead of only the past encounters (Eq. S1),

$$\mathcal{A}(\Psi^c) = \sum_{\hat{t}=1} \mathcal{A}^c(\Psi^{c'}(\hat{t})) \quad (\text{S16})$$

Here, \hat{t} is the reverse time with the origin at the final (current) time point, as defined in Eq. S1. In this case, the encounter of a network with pattern $\Psi^{c'}(\hat{t})$ from class c' at (reverse) time \hat{t} is accounted for by the contribution $\mathcal{A}^c(\Psi^{c'}(\hat{t}))$ to the network's affinity against the pattern Ψ^c . We then find

$$\mathcal{A}^c(\Psi^{c'}(\hat{t})) = \begin{cases} 0 & \text{if } c \neq c' \\ a_0 \lambda (1 - \lambda)^{\hat{t}-1} \rho^{\Theta \hat{t}} =: E(\hat{t}) & \text{if } c = c' \end{cases} \quad (\text{S17})$$

which implies that a repertoire's affinity against a pattern from a given class is only determined by the prior history of the repertoire's encounters with patterns of the same class, and the memory of these prior encounters decay over time according to Eq. S17.

The repertoire encounters a pattern of a specific class at a given time point with probably $1/N$. As a result, the distribution of affinity contributions from a given time point $\mathcal{A}^c(\Psi^{c'}(\hat{t}))$ can be expressed as,

$$P[\mathcal{A}^c(\Psi^{c'}(\hat{t}))] = \left(1 - \frac{1}{N}\right) \delta\left(\mathcal{A}^c(\Psi^{c'}(\hat{t})) - 0\right) + \frac{1}{N} \delta\left(\mathcal{A}^c(\Psi^{c'}(\hat{t})) - E(\hat{t})\right), \quad (\text{S18})$$

where $E(\hat{t})$ is defined in Eq. S17. As a result, the expectation value $\langle \mathcal{A}(\Psi^c(T)) \rangle$ for the affinity of the repertoire against a pattern from class c (Eq. S16) can be evaluated as,

$$\langle \mathcal{A}(\Psi^c) \rangle = \sum_{\hat{t}=1} \int d\mathcal{A}_\hat{t}^{c,c'} \mathcal{A}_\hat{t}^{c,c'} P[\mathcal{A}_\hat{t}^{c,c'}]. \quad (\text{S19})$$

where we used the shorthand notation $\mathcal{A}_\hat{t}^{c,c'} \equiv \mathcal{A}^c(\Psi^{c'}(\hat{t}))$. Fig. S3 shows an agreement between simulations and the expected affinities estimated with this procedure.

We can now express the cumulant-generating function of the affinity distribution $P[\mathcal{A}^c(\Psi^{c'}(\hat{t}))]$ as the sum of the cumulant-generating functions of the independent terms $\mathcal{A}^c(\Psi^{c'}(\hat{t}))$. With the definition in Eq. S15 we can evaluate the cumulant-generating functions $R_\hat{t}(q, \hat{t})$ for the contributions of each time point $\mathcal{A}^c(\Psi^{c'}(\hat{t}))$ to the affinity function,

$$\begin{aligned} R_\hat{t}(q, \hat{t}) &= \log \langle e^{q \mathcal{A}_\hat{t}^{c,c'}} \rangle = \log \left(\left(1 - \frac{1}{N}\right) + \frac{1}{N} e^{q E(\hat{t})} \right) \\ &\approx -\frac{1}{N} + \frac{1}{N} e^{q E(\hat{t})} + \mathcal{O}(N^{-2}). \end{aligned} \quad (\text{S20})$$

The cumulant-generating function $\mathcal{R}(q)$ of the affinity distribution $P[\mathcal{A}]$ can be expressed as the sum of the cumulant-generating functions of the independent contributions from each time point (Eq. S20), which entail,

$$\begin{aligned} \mathcal{R}(q) &= \sum_{\hat{t}=1} R_\hat{t}(q, \hat{t}) = \sum_{\hat{t}=1} \left[-\frac{1}{N} + \frac{1}{N} e^{q E(\hat{t})} \right] + \mathcal{O}(N^{-2}) \\ &= \sum_{\hat{t}=1} \left[-\frac{1}{N} + \frac{1}{N} \sum_{n=0} \frac{q^n E(\hat{t})^n}{n!} \right] + \mathcal{O}(N^{-2}) \\ &= \frac{1}{N} \sum_{n=1} \frac{q^n}{n!} a_0^n \sum_{\hat{t}=0} \lambda^n \rho^{\Theta n} ((1 - \lambda) \rho^\Theta)^{n \hat{t}} + \mathcal{O}(N^{-2}) \\ &= \frac{1}{N} \sum_{n=1} \frac{q^n}{n!} \times \frac{a_0^n \lambda^n \rho^{n \Theta}}{1 - (1 - \lambda)^n \rho^{n \Theta}} + \mathcal{O}(N^{-2}). \end{aligned} \quad (\text{S21})$$

where we have substituted the exponential function with an infinite sum, used the expression in Eq. S17 for $E(\hat{t})$, and performed the resulting geometric sum.

We can now evaluate the n^{th} cumulant c_n of the affinity function as,

$$c_n = \left. \frac{d^n \mathcal{R}(q)}{dq^n} \right|_{q=0} = a_0^n \frac{\lambda^n \rho^{n\Theta}}{N(1 - (1 - \lambda)^n \rho^{\Theta n})} + \mathcal{O}(N^{-2}). \quad (\text{S22})$$

While the first cumulant is equal to the mean affinity in Eq. S5, obtained from the direct calculation of the moments, the second cumulant differs from the variance in Eq. S14. These differences arise due to the expansion of the logarithm in Eq. S20, which is only valid for large N . However, both results describe well the behavior of the variance in the regime that we are interested in (Fig. S1), and therefore, it is warranted to use the simplified form in Eq. 2 for our analysis.

To characterize the scaling relation between the cumulants c_n and the model parameters, we can expand Eq. S22 for the case where the learning rate is close to its optimal value $\lambda^* \sim \mu_{\text{eff}}^{2/3}$ (Eq. 5), which entails,

$$c_n = a_0^n \frac{\lambda^{n-1}}{nN} + \mathcal{O}\left(\frac{\lambda^{n-2} \mu_{\text{eff}}}{nN^2}\right), \quad (\text{S23})$$

and thus, to the leading order, the mean affinity scales as $c_1 \sim a_0/N$ and the variance scales as $c_2 \sim a_0^2 \lambda / (2N)$.

4. Approximating the affinity distribution with a Gamma distribution

As mentioned in the main text, the interpretation of the affinity as a sum of samples from an exponential function separated by exponentially distributed times motivates us to model the affinity distribution as a Γ -distribution. To corroborate this choice, we evaluate the Kullback-Leibler divergence (D_{KL}) between the distribution of patterns affinities in the repertoire $P_{\Psi}(\mathcal{A})$ and a Gamma distribution with matching mean and variance. Since we do not have an analytical expression for $P_{\Psi}(\mathcal{A})$ we use an Edgeworth approximation [25] to evaluate the Kullback-Leibler divergence between the two distribution, by relying on the cumulants of $P_{\Psi}(\mathcal{A})$ (Eq. S22).

In brief, Edgeworth series expands a probability density function around a normal distribution in terms of its cumulants, and it provides a true asymptotic expansion with controlled error [25]. To use the Edgeworth series, we transform the data to have mean zero and variance one, resulting in a modified probability density function for affinities $\hat{P}_{\Psi}(\mathcal{A})$. The second leading order approximation to $\hat{P}_{\Psi}(\mathcal{A})$ is given by

$$\hat{P}_{\Psi}(\mathcal{A}) \approx \phi(\mathcal{A}) \left(1 + \frac{1}{3!} \hat{c}_3 He_3(\mathcal{A}) + \frac{1}{4!} \hat{c}_4 He_4(\mathcal{A}) + \frac{10}{6!} \hat{c}_3^2 He_6(\mathcal{A}) \right) := \phi(\mathcal{A}) (1 + u_{P_{\Psi}}), \quad (\text{S24})$$

where $\phi(\mathcal{A})$ is a standard normal distribution, He_n is a Hermite polynomial of order n , and $\hat{c}_n = c_n / c_2^{n/2}$ is the n^{th} normalized cumulant [24, 25]. As seen in Fig. S3, the approximation in Eq. S24 describes the distribution and especially the bulk of the affinity distribution very well.

To evaluate the Kullback-Leibler divergence $D_{\text{KL}}(\hat{P}_{\Psi} || \hat{\Gamma}_{\Psi})$ between the modified affinity distribution $\hat{P}_{\Psi}(\mathcal{A})$ and a Gamma distribution $\hat{\Gamma}$ with matching mean and variance, we use an Edgeworth expansion for both of these distributions,

$$\begin{aligned} D_{\text{KL}}(\hat{P}_{\Psi}(\mathcal{A}) || \hat{\Gamma}_{\Psi}(\mathcal{A})) &= \int_{\mathcal{A}} \hat{P}_{\Psi}(\mathcal{A}) \log \frac{\hat{P}_{\Psi}(\mathcal{A})}{\hat{\Gamma}_{\Psi}(\mathcal{A})} = \int_{\mathcal{A}} \phi(\mathcal{A}) (1 + u_{P_{\Psi}}) \log \frac{\phi(\mathcal{A}) (1 + u_{P_{\Psi}})}{\phi(\mathcal{A}) (1 + u_{\Gamma_{\Psi}})} \\ &= \int_{\mathcal{A}} \phi(\mathcal{A}) \left[\left(u_{P_{\Psi}} + \frac{1}{2} u_{P_{\Psi}}^2 \right) - \left(u_{\Gamma_{\Psi}} + u_{P_{\Psi}} u_{\Gamma_{\Psi}} - \frac{1}{2} u_{\Gamma_{\Psi}}^2 \right) \right] + \mathcal{O}(u^3) \end{aligned} \quad (\text{S25})$$

where $u_{\Gamma_{\Psi}}$ arises from the Edgeworth expansion of the Gamma distribution, in analogy to Eq. S24. A similar approach has previously been used to approximate the Kullback-Leibler divergence between two distributions [29, 30].

To evaluate the integral in Eq. S25, we use the orthogonality of the Hermit polynomials, i.e.,

$$\int_x \phi(x) He_n(x) He_m(x) = n! \delta_{n,m} \quad (\text{S26})$$

with $He_0 = 1$. As a result, all the linear terms in Eq. S25 vanish and only the squared terms with equal polynomial orders He_n contribute. We thus obtain,

$$\begin{aligned} D_{\text{KL}}(\hat{P}_{\Psi}(\mathcal{A})||\hat{\Gamma}_{\Psi}(\mathcal{A})) &= \frac{1}{2} \left(\frac{1}{3!} \hat{c}_3^2 + \frac{1}{4!} \hat{c}_4^2 + \frac{100}{6!} \hat{c}_3^4 \right) + \frac{1}{2} \left(\frac{1}{3!} \hat{\gamma}_3^2 + \frac{1}{4!} \hat{\gamma}_4^2 + \frac{100}{6!} \hat{\gamma}_3^4 \right) \\ &\quad - \left(\frac{1}{3!} \hat{c}_3 \hat{\gamma}_3 + \frac{1}{4!} \hat{c}_4 \hat{\gamma}_4 + \frac{100}{6!} \hat{c}_3^2 \hat{\gamma}_3^2 \right) \\ &= \frac{1}{2} \left[\frac{1}{3!} (\hat{c}_3 - \hat{\gamma}_3)^2 + \frac{1}{4!} (\hat{c}_4 - \hat{\gamma}_4)^2 + \frac{100}{6!} (\hat{c}_3^2 - \hat{\gamma}_3^2)^2 \right]. \end{aligned} \quad (\text{S27})$$

where $\hat{\gamma}_n$ is the n^{th} cumulant of the modified Gamma distribution $\hat{\Gamma}_{\Psi}$.

By substituting the cumulant of the affinity distribution (Eq. S22) and that of the Gamma distribution, we arrive at an approximation for the Kullback-Leibler distance between the two distributions,

$$D_{\text{KL}}(P_{\Psi}(\mathcal{A})||\Gamma_{\Psi}(\mathcal{A})) = \frac{25}{27} (2\Theta\mu_{\text{eff}}\kappa)^{2/3} + \mathcal{O}(\Theta\mu_{\text{eff}}) \ll 1. \quad (\text{S28})$$

indicating that the Gamma distribution is a good approximation for the true distribution of the affinities $P_{\Psi}(\mathcal{A})$ for small mutation rate μ_{eff} as long as $\Theta\kappa$ is of order 1. However, for risk tolerant strategies (large κ), this approximation fails. In this regime the repertoire only learns one pattern effectively, resulting in a bi-modal distribution of the affinities with one mode reflecting the low-affinity recognition and the other, the higher affinity of the latest stored pattern. This bimodal distribution cannot be approximated by a Gamma distribution.

5. Affinity of random patterns

The affinity of a random pattern χ (i.e., patterns unrelated to the previously encountered classes Ψ^c) is determined by summing over the overlaps $|\langle \psi^{\alpha} | \chi \rangle|^{\Theta}$ between the random pattern χ and all previously stored patterns in the memory repertoire $\psi^{\alpha} \in \mathcal{M}$, weighed by their contributions to memory m^{α} . It should be noted that the mean affinity of random patterns is set to be zero and therefore, the affinity shift $\mathcal{A}_{\text{rand}}$, defined in Eq. 2, can be evaluated as,

$$\mathcal{A}_{\text{rand}} = \left\langle \mathcal{A}_0 \sum_{\psi^{\alpha} \in \mathcal{M}} m^{\alpha} |\langle \psi^{\alpha} | \chi \rangle|^{\Theta} \right\rangle \quad (\text{S29})$$

$$= \mathcal{A}_0 \langle |\langle \psi | \chi \rangle|^{\Theta} \rangle = \mathcal{A}_0 \frac{2^{\Theta/2} \Gamma\left[\frac{1+\Theta}{2}\right]}{\sqrt{\pi L^{\Theta}}}, \quad (\text{S30})$$

where we use the fact that random pattern χ is independent of the stored patterns ψ^{α} , and that on average their overlap $\langle \psi | \chi \rangle$ is a normally distributed random variable with mean zero and variance $\sigma_{\text{rand}}^2 = \frac{1}{4L}$, according to the law of large numbers.

To evaluate the variance of random patterns, we first introduce a basis that spans all the stored patterns in a repertoire ψ^{α} that have non-zero weights m^{α} at a given time point t . One choice for this basis is to use the N directions corresponding to the pattern classes $\Psi^c(t)$ at time t . However, these N vectors do not fully span all the previously presented patterns, due to the evolutionary divergence of these patterns over time. To account for this remaining subspace, we introduce N' auxiliary directions $\Phi^{c'}(t)$ (with $c' = 1, \dots, N'$). In principle, the space encompassing all possible the patterns is L -dimensional. However, the space of all stored patterns is typically more restricted (i.e., $N + N' < L$) due to the relatively fast updates in the repertoire such that it only keeps a memory of patterns that are similar to the current states of the presented classes $\Psi^c(t)$. Using the set of basis vectors $\{\Psi^c(t), \Phi^{c'}(t)\}$, we can express any stored pattern in the repertoire as,

$$\langle \psi^{\alpha} | = \sum_{c=1}^N \langle \psi^{\alpha} | \Psi^c(t) \rangle \langle \Psi^c(t) | + \sum_{c'=1}^{N'} \langle \psi^{\alpha} | \Phi^{c'}(t) \rangle \langle \Phi^{c'}(t) |. \quad (\text{S31})$$

As a result, the overlaps between a random pattern χ and all the stored patterns in the repertoire in Eq. S29 can be

expressed as,

$$\sum_{\psi^\alpha \in \mathcal{M}} m^\alpha |\langle \psi^\alpha | \chi \rangle|^\Theta = \sum_{\psi^\alpha \in \mathcal{M}} m^\alpha \left| \sum_{c=1}^N \langle \Psi^c(t) | \chi \rangle \langle \psi^\alpha | \Psi^c(t) \rangle + \sum_{c'=1}^{N'} \langle \Phi^{c'}(t) | \chi \rangle \langle \psi^\alpha | \Phi^{c'}(t) \rangle \right|^\Theta \quad (\text{S32})$$

$$\approx \sum_{c=1}^N \left(\sum_{\psi^\alpha \in \mathcal{M}} m^\alpha |\langle \psi^\alpha | \Psi^c(t) \rangle \langle \Psi^c(t) | \chi \rangle|^\Theta \right) + \sum_{c'=1}^{N'} \left(\sum_{\psi^\alpha \in \mathcal{M}} m^\alpha |\langle \psi^\alpha | \Phi^{c'}(t) \rangle \langle \Phi^{c'}(t) | \chi \rangle|^\Theta \right) \quad (\text{S33})$$

$$\equiv \sum_{c=1}^N M_c |\langle \Psi^c(t) | \chi \rangle|^\Theta + \sum_{c'=1}^{N'} M'_{c'} |\langle \Phi^{c'}(t) | \chi \rangle|^\Theta. \quad (\text{S34})$$

To arrive at Eq. S33, we assumed that the stored patterns ψ^α have non-vanishing overlaps with only one of the bases in the set $\{\Psi^c, \Phi^{c'}\}$. As a result, in the expansion of the expression to the power Θ in Eq. S32, all the cross terms associated with different bases vanish, and the expression can be simply written as the sum of independent terms. The final form in Eq. S34 expresses the overlap of a random pattern with the repertoire as a sum over the overlaps with the bases $\{\Psi^c(t), \Phi^{c'}(t)\}$, with effective weights $M_c(t) = \sum_\alpha m^\alpha |\langle \psi^\alpha | \Psi^c(t) \rangle|^\Theta$, and $M'_{c'}(t) = \sum_\alpha m^\alpha |\langle \psi^\alpha | \Phi^{c'}(t) \rangle|^\Theta$. The effective weight M_c associated with the presented pattern classes $\Psi^c(t)$ follows,

$$M_c(t) = \sum_\alpha m^\alpha |\langle \psi^\alpha | \Psi^c(t) \rangle|^\Theta = \frac{\mathcal{A}(\Psi^c(t)) + \mathcal{A}_{\text{rand}}}{\mathcal{A}_0} \approx \frac{\mathcal{A}(\Psi^c(t))}{a_0}, \quad (\text{S35})$$

where we used the approximation $a_0 = \mathcal{A}_0 - \mathcal{A}_{\text{rand}} = \mathcal{A}_0 + \mathcal{O}(L^{-\Theta/2})$.

When the shape parameter is $\Theta = 2$, the effective weights define a normalized set (i.e., $\sum_{c=1}^N M_c + \sum_{c'=1}^{N'} M'_{c'} = 1$). However, for sharper affinity functions ($\Theta > 2$), the relationship between the effective weights is bounded from above by $\sum_{c=1}^N M_c + \sum_{c'=1}^{N'} M'_{c'} \leq 1$, and for broader affinity functions ($\Theta < 2$), the weights are bounded from below as, $1 \leq \sum_{c=1}^N M_c + \sum_{c'=1}^{N'} M'_{c'}$. Here, we discuss the case with $\Theta > 2$, which entails

$$\sum_{c'=1}^{N'} M'_{c'}(t) \leq 1 - \sum_{c=1}^N M_c(t) \approx 1 - N \frac{\langle \mathcal{A} \rangle}{a_0}. \quad (\text{S36})$$

This upper bound implies that when the patterns are well represented in the repertoire (i.e., $\langle \mathcal{A} \rangle \approx \mathcal{A}_{\text{max}} = a_0/N$; see Eq. S5), the weights of the auxiliary bases shrink to zero, i.e., $N' = 0$. This result is in line with the eigen-decomposition analysis of the generalized Hopfield network in Appendix C of ref. [9].

For the cases that $M'_{c'} > 0$, we will consider a mean field approximation, where all the auxiliary directions are equally important (i.e., $M'_{c'} = M', \forall c'$) and that on average the memory stored in the auxiliary directions has a comparable weight to that of the bases spanned by the presented patterns, (i.e., $\mathbb{E}_c[M_c] = \mathbb{E}_{c'}[M'_{c'}] = M'$, where $\mathbb{E}_\star[\cdot]$ denotes expectation over the argument in the subscript). These approximations together with Eq. S36 define an upper bound for the number of auxiliary bases,

$$N' \leq \frac{a_0}{\langle \mathcal{A} \rangle} \left(1 - N \frac{\langle \mathcal{A} \rangle}{a_0} \right). \quad (\text{S37})$$

Using these relationships, we can now evaluate the variance of the affinities across random patterns χ , as,

$$\mathbf{var}_{\text{rand}} = \mathbf{var} \left[\mathcal{A}_0 \sum_{\psi^\alpha \in \mathcal{M}} m^\alpha |\langle \psi^\alpha | \chi \rangle|^\Theta - \mathcal{A}_{\text{rand}} \right]_\chi \quad (\text{S38})$$

$$= \mathcal{A}_0^2 \mathbf{var} \left[\sum_{c=1}^N M_c |\langle \Psi^c(t) | \chi \rangle|^\Theta + \sum_{c'=1}^{N'} M'_{c'} |\langle \Phi^{c'}(t) | \chi \rangle|^\Theta \right] \quad (\text{S39})$$

$$= \mathcal{A}_0^2 \sum_{c=1}^N (M_c)^2 \mathbf{var} \left(|\langle \Psi^c | \chi \rangle|^\Theta \right) + \mathcal{A}_0^2 \sum_{c'=1}^{N'} (M'_{c'})^2 \mathbf{var} \left(|\langle \Phi^{c'} | \chi \rangle|^\Theta \right) \quad (\text{S40})$$

$$= \mathcal{A}_0^2 \left(\sum_{c=1}^N (M_c)^2 + \sum_{c'=1}^{N'} (M'_{c'})^2 \right) \mathbf{var} \left(|\langle \psi | \chi \rangle|^\Theta \right) \quad (\text{S41})$$

$$\leq \mathcal{A}_0^2 \left(\frac{N}{a_0^2} (\mathbf{var}(\mathcal{A}) + \langle \mathcal{A} \rangle^2) + \frac{1}{a_0} \langle \mathcal{A} \rangle \left(1 - \frac{N}{a_0} \langle \mathcal{A} \rangle \right) \right) \times \frac{2^\Theta \left(\sqrt{\pi} \Gamma \left[\frac{1}{2} + \Theta \right] - \Gamma \left[\frac{1+\Theta}{2} \right]^2 \right)}{\pi L^\Theta}, \quad (\text{S42})$$

where, we first used Eq. S30 to express the affinity function in Eq. S38 in terms of the bases in Eq. S39. Given that the projections of a random pattern along different bases are statistically independent from each other, we expressed the variance of the sum of these contributions in Eq. S39 as the sum of their variances in Eq. S40. Next we used the fact that the overlaps of a random pattern with all the bases (i.e., $\langle \Psi^c | \chi \rangle$, $\langle \Phi^{c'} | \chi \rangle$, $\forall c, c'$) are Gaussian random variates with mean zero and variance $L/4$. Thus, the variance of these overlaps to the power Θ are equal for all bases, which we simply expressed as $\mathbf{var} \left[|\langle \psi | \chi \rangle|^\Theta \right]$ in Eq. S41. We then evaluated the term $\mathbf{var} \left[|\langle \psi | \chi \rangle|^\Theta \right]$ from the underlying Gaussian distribution, and used the definition of the effective weight M_c in Eq. S35 and the inequality in Eq. S37 to substitute $\sum_{c=1}^N (M_c)^2 + \sum_{c'=1}^{N'} (M'_{c'})^2$ and arrive at the upper bound of the variance for the affinity for random patterns in Eq. S42. The results in Eq. S42 match very well with simulations for the affinity function with the shape parameter $\Theta = 2$ (Fig S4).

Appendix B: Mapping between immune memory repertoires and the Hopfield model

The Hopfield network is among the most frequently used models for associative memory [10]. A classical Hopfield network describes a fully connected graph with interaction matrix J_{ij} . A binary pattern ψ of length L presented to the network J is assigned an energy $E(\psi, J)$,

$$E(\psi, J) = -\frac{1}{2L} \sum_{i,j} J_{i,j} \psi_i \psi_j. \quad (\text{S1})$$

Hopfield networks can learn and store an associative memory of the presented patterns as the minima of the energy landscape [10]. One way to construct such network is by Hebbian learning, whereby the network is updated with a learning rate λ upon an encounter with a pattern $\psi(t)$ at time t ,

$$J_{i,j}(t+1) = \begin{cases} (1-\lambda) J_{i,j}(t) + \lambda \psi_i(t) \psi_j(t), & \text{if } i \neq j; \\ 0, & \text{otherwise.} \end{cases} \quad (\text{S2})$$

In the supplementary information of ref. [9], we show that this learning rule can be expressed as

$$J(t+1) = (1-\lambda)J(t) + \lambda(L|\psi(t)\rangle \langle \psi(t)| - \mathbb{1}). \quad (\text{S3})$$

Using this expression for the network $J(T)$ at time T with an encounter history with patterns $\psi(t)$ (for $t \leq T$), we

can evaluate the energy $E(\chi, J(T))$ of an arbitrary pattern χ presented to the network as,

$$\begin{aligned}
E(\chi, J(T)) &= -\frac{1}{2} \langle \chi | J(T) | \chi \rangle \\
&= -\frac{L}{2} \sum_{t=1}^{T-1} \lambda(1-\lambda)^{T-1-t} \langle \chi | \psi(t) \rangle \langle \psi(t) | \chi \rangle + \frac{1}{2} \sum_{t=1}^{T-1} \lambda(1-\lambda)^{T-1-t} \langle \chi | \mathbf{1} | \chi \rangle \\
&= -\frac{L}{2} \sum_{t=1}^{T-1} \lambda(1-\lambda)^{T-1-t} |\langle \psi(t) | \chi \rangle|^2 + \frac{1}{2} + \mathcal{O}((1-\lambda)^T) \\
&= -\frac{L}{2} \sum_{t=1}^{T-1} m_t |\langle \psi(t) | \chi \rangle|^2 + \frac{1}{2} + \mathcal{O}((1-\lambda)^T)
\end{aligned} \tag{S4}$$

with the time-dependent weights, $m_t = \lambda(1-\lambda)^{T-1-t}$. Note that the correction terms vanish in the limit of large time T and that the weights sum up to one, i.e., $\sum m_t = 1$. In this limit, the energy function $E(\chi, J)$ corresponds to the affinity function given in Eq. 2 of the main text, for the choice of the shape parameter $\Theta = 2$, the energy scale $A_0 = -L/2$, and the random energy $\mathcal{A}_{\text{rand}} = -1/2$.

Appendix C: Numerical methods

As discussed in Section B, for the shape parameter $\Theta = 2$ the affinity function in Eq. 2 maps onto the energy function of a Hopfield network with Hebbian learning [9, 10, 12], which is easily tractable with numerical techniques. Specifically, using a Hopfield model for simulating the repertoire problem with $\Theta = 2$ has the advantage that we only need to keep track of the interaction matrix $J_{i,j}$ of size L^2 as opposed to all the 2^L memory weights m^α . This dramatic reduction in complexity makes the numerical simulations for $\Theta = 2$ highly efficient.

For the simulations, we use the same approach as in [9]. We initialize the interaction matrix J of size $L \times L$ with all entries set to zero $J_{i,j} = 0$ and choose N independent random patterns of size L .

Before collecting any data, we first update the network until it reaches a quasi-stationary state so that it no longer depends on the initial condition. Since the stored memory of a pattern within the network decays as $(1-\lambda)^s$ with the number the number update steps s since the original encounter, we update the network following the initialization by $s_{\text{stat.}} = \text{ceil}\left(\frac{\log 10^{-5}}{\log(1-\lambda)}\right)$ steps for the network to reach a quasi-stationary state. This criterion ensures that $(1-\lambda)^{s_{\text{stat.}}} \leq 10^{-5}$ and the memory of the initial state is removed. Moreover, this criterion implies that in the quasi-stationary state the memory weights become normalized i.e., $1 - \sum_\alpha m^\alpha(s_{\text{stat.}}) < 10^{-5}$, after starting from a no memory initial state of $\sum_\alpha m^\alpha(0) = 0$. During each step of this preparation phase, we also evolve the patterns, whereby we flip the spins of patterns within each class at rate μ . We then randomly choose one of the patterns to present to the network and update the network according to the learning rule in Eq. S2.

After reaching the quasi-stationary state of memory, we collect data from the network over 10^4 steps. During this process, the patterns evolve with rate μ and are presented at a random order to the network. We record the affinity of each presented pattern before updating the network. At each step we also record the affinity of a randomly generated pattern that belongs to none of the previously encountered pattern classes. For each learning rate λ , we repeat this process for 50 independent initial sets of pattern classes $\{\Psi^c\}$. Overall, we perform a total of 5×10^5 affinity measurements on both the previously encountered and the random patterns.

-
- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) <http://www.deeplearningbook.org>.
 - [2] A. Soltoggio, K. O. Stanley, and S. Risi, Born to learn: The inspiration, progress, and future of evolved plastic artificial neural networks, *Neural Networks* **108**, 48 (2018).
 - [3] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to Machine Learning for physicists, *Physics Reports* **810**, 1 (2019), arXiv: 1803.08823.
 - [4] C. Janeway, P. Travers, M. Walport, and M. Schlomchik, *Immunobiology*, 5th ed., The Immune System in Health and Disease (Garland Science, New York, 2001).
 - [5] R. Barrangou and L. A. Marraffini, CRISPR-Cas systems: Prokaryotes upgrade to adaptive immunity, *Molecular cell* **54**, 234 (2014).
 - [6] G. Altan-Bonnet, T. Mora, and A. M. Walczak, Quantitative immunology for physicists, *Physics Reports* **849**, 1 (2020).
 - [7] S. Bradde, A. Nourmohammad, S. Goyal, and V. Balasubramanian, The size of the immune repertoire of bacteria, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5144 (2020).

- [8] O. H. Schnaack and A. Nourmohammad, Optimal evolutionary decision-making to store immune memory, *eLife* **10**, e61346 (2021).
- [9] O. H. Schnaack, L. Peliti, and A. Nourmohammad, Learning and organization of memory for evolving patterns, arXiv:2106.02186 [physics] (2021), arXiv: 2106.02186.
- [10] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).
- [11] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (Wiley, New York, 1949).
- [12] M. Mezard, J. P. Nadal, and G. Toulouse, Solvable models of working memories, *J. Physique* **47**, 1457 (1986).
- [13] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [14] R. E. Steuer, *Multiple criteria optimization: theory, computation, and application*, Wiley series in probability and mathematical statistics (Wiley, New York, 1986).
- [15] A. Sen, Markets and Freedoms: Achievements and Limitations of the Market Mechanism in Promoting Individual Freedoms, *Oxf. Econ. Pap.* **45**, 519 (1993).
- [16] O. Shoval, H. Sheftel, G. Shinar, Y. Hart, O. Ramote, A. Mayo, E. Dekel, K. Kavanagh, and U. Alon, Evolutionary Trade-Offs, Pareto Optimality, and the Geometry of Phenotype Space, *Science* **336**, 1157 (2012).
- [17] R. Schuetz, N. Zamboni, M. Zampieri, M. Heinemann, and U. Sauer, Multidimensional Optimality of Microbial Metabolism, *Science* **336**, 601 (2012).
- [18] Y. Hart, H. Sheftel, J. Hausser, P. Szekely, N. B. Ben-Moshe, Y. Korem, A. Tendler, A. E. Mayo, and U. Alon, Inferring biological tasks using Pareto analysis of high-dimensional data, *Nat Methods* **12**, 233 (2015).
- [19] P. Szekely, Y. Korem, U. Moran, A. Mayo, and U. Alon, The Mass-Longevity Triangle: Pareto Optimality and the Geometry of Life-History Trait Space, *PLoS Comput Biol* **11**, e1004524 (2015).
- [20] L. F. Seoane and R. Solé, Phase transitions in Pareto optimal complex networks, *Phys. Rev. E* **92**, 032807 (2015).
- [21] A. Tendler, A. Mayo, and U. Alon, Evolutionary tradeoffs, Pareto optimality and the morphology of ammonite shells, *BMC Syst Biol* **9**, 12 (2015).
- [22] L. Koçillari, P. Fariselli, A. Trovato, F. Seno, and A. Maritan, Signature of Pareto optimization in the Escherichia coli proteome, *Sci Rep* **8**, 9141 (2018).
- [23] S. Cobey and S. E. Hensley, Immune history and influenza virus susceptibility, *Current Opinion in Virology* **22**, 105 (2017).
- [24] P. Hall, *The Bootstrap and Edgeworth Expansion*, Springer Series in Statistics (Springer New York, New York, NY, 1992).
- [25] S. Blinnikov and R. Moessner, Expansions for nearly Gaussian distributions, *Astron. Astrophys. Suppl. Ser.* **130**, 193 (1998).
- [26] M. Schweizer, Mean-Variance Hedging for General Claims, *Ann. Appl. Probab.* **2**, 10.1214/aoap/1177005776 (1992).
- [27] K. B. Toft, On the Mean-Variance Tradeoff in Option Replication with Transactions Costs, *The Journal of Financial and Quantitative Analysis* **31**, 233 (1996).
- [28] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar, Learning in Nonstationary Environments: A Survey, *IEEE Comput. Intell. Mag.* **10**, 12 (2015).
- [29] J.-J. Lin, N. Saito, and R. A. Levine, Edgeworth approximation of the Kullback-Leibler distance towards problems in image analysis, University of California, Davis, Tech. Rep (1999).
- [30] J. Inglada and G. Mercier, A New Statistical Similarity Measure for Change Detection in Multitemporal SAR Images and Its Extension to Multiscale Change Analysis, *IEEE Trans. Geosci. Remote Sensing* **45**, 1432 (2007).

Chapter 6

Discussion and outlook

Storing and retrieving memory from prior molecular interactions is an efficient scheme to sense and respond to external stimuli. So far, the standard literature has focused on learning strategies for constant stimuli, and though many studies have included noisy inputs, most have neglected evolving stimuli. One field where learning and storing memory of evolving signals is essential is immunology. A vital feature of the adaptive immune system is its ability to store memory of pathogens for future infections. Importantly, the immune system can also utilize this memory against evolved forms of pathogens. While understanding immune memory for evolving pathogens is an important and interesting problem on its own, understanding general learning strategies for evolving stimuli is an open problem with potential applications far beyond immunology.

In this thesis, I have presented three projects that study different aspects of learning and memory of evolving targets. The key findings of these projects can be summarized as follows. A single unit of memory encountering only one evolving target should have some cross-reactivity even if its cross-reactivity is traded off against having maximal affinity (utility) for some encounters. Further, it can be beneficial to delay a novel response and first try to utilize the memory even if that increases the cost of the novel response. When multiple evolving targets are encountered, fully connected networks such as the Hopfield network fail to encode and retrieve the stored state of all targets correctly. One strategy for maximal recognition performance is using a compartmentalized memory with a specialized 1-to-1 memory for each target, akin to what is used by the adaptive immune system. Moreover, repertoires of multiple specialized memories need to increase their learning rate with the evolutionary rate of the targets. This increased learning rate leads to broader distributions of memory performance. Thus, although these memory repertoires need to have moderate levels of risk tolerance to fluctuations in performance, they can be very efficient in distinguishing between evolved versions of stored targets and novel random stimuli.

6.1 Optimal memory of the immune system

In the first publication of this thesis ([1] see chapter 3), we studied the optimal choice of memory B-cells. As naïve B-cells undergo affinity maturation, their receptors become more specialized and less cross-reactive [146–150]. To characterize this phenomenon and incorporate these biophysical constraints, we constructed a model which includes a tradeoff between receptors' (antibodies') cross-reactivity and maximal affinity. We found that B-cell differentiation into the memory compartment should be highly regulated to produce lower affinity receptors, allowing cross-reactive recognition of evolved pathogens. This prediction is consistent with recent experiments [59–63]. We also found that it is beneficial to allocate more deliberation time when a novel antigen is encountered so that lower affinity memory can react before initiating a novel response. This delayed novel response has been reported in numerous studies [72–74].

To study different modes of the immune memory response, we co-optimized the kinetics and energetics of memory for an effective response against evolving pathogens over the organism's lifetime. We found three distinct regimes of memory production. (1) For non- or slowly evolving pathogens, the optimal solution uses maximally specialized high-affinity receptors. (2) When pathogens evolve at moderate speed, the best solution uses a non-equilibrium response with a less adapted, more cross-reactive memory. While this is the best solution in this setting, it will always be suboptimal because the response was initially adapted for a different version of the pathogen. (3) When the pathogen evolution becomes too fast, the system should not use memory but instead rely only on a naïve response.

While all these observations might appear to be intuitively clear, our work was the first to propose a utility of an immune memory response to evolved pathogens. Indeed, a recent (published after [1]) preprint suggests a different model for this problem [101]. While they use a different function to model the cross-reactive affinity function between receptors and antigens, their parameters to describe the immune system are simply different interpretations of our model. First, they use receptors with constant affinity and cross-reactive range, and instead introduce a tradeoff between a monoclonal response (all receptors are centered at the pathogen) and a polyclonal response (receptors are distributed further around the pathogen). Because they study a repertoire with constant size, the effect of their tradeoff is equivalent to ours, a tradeoff between the maximal affinity and cross-reactivity of receptors. The second parameter they use is the cost of a naïve response¹. Here the connection to our deliberation parameter (β) is even clearer as we chose our naïve cost to be a monotonic function of β . With this clear mapping to our model, it is no surprise that [101] finds the same three regimes described above, while their model relies on more assumptions of the actual biological process.

A further advantage of our model is that it features a stochastic choice between memory and naïve response. This non-deterministic decision enables the usage of sub-

¹Indeed, they use the naïve diversity as the quantity they discuss. But as pointed out in the paper, the cost used for the naïve response is the inverse of the diversity.

optimal memory in instances where a naïve response would be beneficial. A comparable biological process is known as original antigenic sin, where cross-reactive memory from primary infections interferes and suppresses a protective novel response [151]. Thus, the viral exposure history may profoundly affect protective immunity against evolving viruses [152].

Immune deficiency as a result of encounter history, such as in antigenic sin, can be more pronounced when the life expectancy of organisms changes. Our work shows that the immune strategy of short-lived organisms should favor highly cross-reactive memory. If an organism's life expectancy increases, individuals will likely encounter more evolved forms of a pathogen with antigenic distances larger than expected by their immune systems. Consistent with antigenic sin, cross-reactive memory optimized for a shorter lifetime, can still be activated but with lower efficacy, thus suppressing a protective novel response. Therefore, the elongation of the human lifespan is one of the plausible factors responsible for immune deficiencies brought by aging.

When the immune system first encounters a pathogen, it cannot predict the pathogen's evolutionary rate. Thus, the memory strategy it uses (i.e., which B-cells differentiate into memory B-cells) needs to be general for all pathogens it encounters. To study this constraint, we optimized which distribution of memory strategies is best when facing pathogens with different evolutionary rates. Here we found that the optimal solution uses a mixture of high-affinity and lower affinity memories. This mixture might be interpreted as the combination of low-affinity IgM and high-affinity IgG memory B-cells. However, these memory B-cells also have different roles in the memory response that go beyond our model. Nevertheless, it is interesting to observe that the optimal strategy for multiple independent threats consists of a regulated mixture between two strategies instead of constantly producing memory.

6.2 Memory structure for evolving inputs

The second publication of this thesis ([2] in chapter 4) studied different structures that store and retrieve memory. We introduced a flexible energy-based neural network model that can adopt different memory strategies. These strategies range from the classical fully connected Hopfield network to a compartmentalized memory that stores a specialized 1-to-1 memory for each state. The contrast between these strategies reflects distinct molecular mechanisms used for memory storage in the adaptive immune system and neural structures such as the olfactory cortex. In particular, the olfactory system uses distributed memory in the olfactory cortex to recognize odor signals from the receptors in the olfactory bulb [116–121]. On the other hand, the adaptive immune system adopts a 1-to-1 strategy and generates specialized immune cells (i.e., compartments) for each pathogen [1, 60, 62, 63, 81, 82].

We studied different memory strategies, defined by the learning rate and the number of compartments in the network, and we probed their efficiency for static and dynamic patterns. We recovered the well-known result that Hopfield-like networks with distributed memory are highly accurate in storing associative memory for static or

noisy patterns. Thus for these patterns, specialized 1-to-1 memories have no advantage over distributed memories. However, we found that fully connected networks fail to reliably store retrievable associative memory for multiple evolving patterns. This observation holds even when the networks learn at an optimal rate. We showed that networks increase their learning rate to follow the patterns' evolution. This increase in the optimal learning rate drastically changes the structure of the energy landscape for associative memory. In particular, we found the emergence of narrow connections (mountain passes) between the memory attractors. These connections destabilized the attractors of some stored states, and the corresponding patterns fall into an incorrect attractor during the equilibration process, leading to a significant reduction in retrieval accuracy.

We further showed that full compartmentalization of the network, where each pattern class is stored in a separate sub-network, is one solution to achieve maximal accuracy. Indeed, in our compartmentalization model, the 1-to-1 strategy is the only solution that achieves the maximal accuracy for evolving patterns. We also found that intermediate solutions (i.e., quasi-specialized memory) do not perform significantly better than the fully connected strategy, as long as the total memory capacity is kept constant.

Our study focuses on a Hebbian learning rule, but we additionally tested three other incremental rules (Hebbian learning with a sparse matrix, gradient descent, and Storkey learning). While all these results suggest that distributed networks cannot efficiently learn evolving patterns, we cannot exclude that some specialized learning rules can achieve a performance comparable to the 1-to-1 strategy. Particularly, learning rules that disconnect the attractors of the energy function will be able to recognize evolving patterns effectively. However, it is questionable whether these rules still resemble a neural network or whether they represent a different form of our fully compartmentalized solution. Further, the ongoing search for artificial neural network structures that can learn evolving environments [127] suggests that optimal learning of evolving stimuli requires a specialized 1-to-1 strategy.

Our results have significant implications for the learning of evolving stimuli. If indeed only a 1-to-1 strategy can achieve the maximal performance, the size of the memory would need to grow linearly with the number of states it stores. This scaling contrasts current machine learning techniques that aim to reduce the input space while still recognizing many signals. The resolution of this problem needs a nuanced discussion.

First, we need to remind ourselves that the fully connected networks we used do not lose the ability to retrieve the memory for all patterns. Instead, only a fraction of stored states is incorrectly recognized. Thus, the optimal memory strategy might follow a tradeoff between accepted error in recognition and resources allocated to the memory. The results of this tradeoff will depend on the constraints of the model that is used. Indeed, systems that can tolerate some degree of error will likely prefer connected network structures. Interestingly, the quasi-specialized strategies do not outperform the fully connected model of a generalized Hopfield memory with constant capacity.

Thus, we would only expect either fully connected or 1-to-1 solutions in our model. However, models with other constraints or different cost allocations might very well find partially compartmentalized solutions to be optimal.

Independent of the used model, we predict that the 1-to-1 strategy is the only solution that reaches the maximal performance. Thus systems that cannot tolerate any error need to use this strategy. The immune memory might be seen as such a system because false recognition of a pathogen can have severe consequences such as death. Our results thus suggest that pathogenic evolution may be one of the reasons for the immune system to encode a specialized memory, whereas distributed memory is used in other systems.

The memory's task also informs on the optimal storage strategy. As discussed, the increased learning rate leads to narrow connections between the attractors in the energy landscape. However, the stored states still lie in these attractors, and their energy is easily distinguishable from random states. Thus, fully connected networks are sufficient if the system only needs to distinguish between previously encountered patterns and novel patterns. This distinction also has some overlap to the problem studied in the final paper of this thesis ([3] see chapter 5).

Such a task can also influence the level of compartmentalization of the solution. For example, if the task is to distinguish between many (evolving) harmful and non-harmful signals, two compartments would be sufficient. One compartment would store harmful inputs (e.g., lions and sharks), while the other compartment stores harmless inputs (e.g., bunnies and goldfish). This system might confuse lions with sharks, but it would have no problem distinguishing between sharks and goldfish, which arguably is much more critical².

In summary, we can say that the memory storage strategies can depend on many model-dependent factors, such as a tradeoff between performance and cost or the actual memory task. While tuning these factors will allow maneuvering in the solution space, our model determines the standalone performance of the memory without model-dependent constraints.

The third and final paper of this thesis ([3] see chapter 5) also studies the structure of memory for evolving targets but takes a different approach and focuses on the memory repertoire. The affinity function we used can be seen as the sum over the affinities of all states in the repertoire with normalized weights. The affinity function thus implies a 1-to-1 memory where each encoded state is stored individually. Importantly, we did not model a retrieval of the memory in a Hopfield sense, where we equilibrate the state in the affinity landscape. Instead, we assumed that the memory is 1-to-1 and the stored states with the highest binding to a presented pattern will trigger a memory response. Indeed the mapping to the Hopfield energy has a pure simulation purpose, and the two projects represent different approaches to modeling memory. Nevertheless, the affinity-energy correspondence allows us to translate the classification results to

²I am aware that the evolution of all these animals happens on time scales that are far beyond the human lifetime. Thus, they appear static to us and are easy to learn by one single fully connected network. However, this should be seen as an illustrative example, not a task humans face.

the fully connected Hopfield model.

In [3] we studied how a repertoire of specific memories needs to be updated when it attempts to store multiple evolving states. The 1-to-1 strategy in [2] is a special state of such a repertoire with N specialized memories with equal weights. In [3] we instead studied how a repertoire can update these weights when it incorporates new information. Indeed, the model in [3] is equivalent to adding a new compartment every time we encounter a new state.

With this model, we studied the distribution of affinities when facing evolving targets. To describe different memory strategies (defined by the learning rate), we introduced a tradeoff between the mean affinity and the fluctuations of the affinities. These fluctuations can lead to a misclassification between stored and random patterns. Therefore, we used the width of the affinity distribution (standard deviation) as the risk measure of a strategy.

Interestingly, systems can learn with maximal mean affinity and risk-free when patterns are static (i.e., do not evolve). In this case, systems will learn infinitely slow to distribute the resources of the memory perfectly. However, once patterns start evolving, memory repertoires need to adapt to the pattern evolution and learn faster, a behavior similar to [2].

We found that, with a risk of discarding older patterns, risk-tolerant repertoires adopt faster learning rates and keep a memory of the recent patterns with high affinity. On the other hand, risk-averse repertoires effectively shut down their learning to minimize the fluctuations in their recognition affinity for different patterns, taking on the cost of having a low affinity for all patterns. Ideally, systems adopt a moderate risk tolerance that enables the repertoire to achieve a desirable balance between risk and affinity. In this regime, systems can robustly recognize previously encountered patterns that evolved and distinguish them from random unseen patterns.

This classification also plays a vital role in immune memory. As discussed above, it is not sufficient to only have a highly effective response to previously encountered pathogens. Instead, the immune memory needs to be specific enough and not mount a response against novel pathogens to avoid immune defects such as antigenic sin. Thus, the classification in our model provides an understanding of the update rates immune repertoires should use.

As our model focuses on evolving patterns, the learning and update behavior in our paper is different from the results of Mayer et al., who studied the learning for memory repertoires that face non-evolving pathogens with changing encounter probabilities [81]. Naturally, future work should investigate the combination of both models.

6.3 Outlook

The projects presented in this thesis introduce frameworks to study different aspects of learning and memory for evolving targets. With these models, we can investigate and predict optimal memory strategies for the immune system but also for general memory systems. As described above, the results of this thesis already significantly

enrich the understanding of memory for evolving targets. However, this thesis is still just a starting point and much more research is needed to extend the results to more systems.

The first natural extension of this thesis is the generalization of the framework introduced in [1]. So far, our model only includes the decision between mounting a naïve or memory response. Further, we assumed that the memory receptors remain constant during reinfections. However, as discussed in the introduction, an immune memory response has different components. These range from long-lived plasma cells to IgG memory B-cells that turn into plasma cells to IgM memory B-cells that can seed secondary germinal centers and undergo further affinity maturation upon reinfection. While we found that an optimal solution should use a combination of high- and low-affinity memory, we used the same time scale for all memory responses. It would be interesting to model the memory response as a multilayer decision process in the spirit of the two walls of protection [7]. In this process, memory B-cells could only be activated if long-lived plasma cells do not neutralize the threat.

A different extension could include the evolution of memory cells and its influence on their long-term utility for evolving pathogens. This evolvability of memory is relevant for characterizing immune response dynamics to chronic viral infections like HIV. The analysis of immune repertoires in HIV patients over multiple years shows a rapid evolution of B-cell lineages to counter the evolution of the virus within hosts [153].

Including the possibility of reseeding germinal centers and thus the adaptation of memory to evolved forms would also allow us to study the coevolution between evolving pathogens and memory in more detail. This inclusion bridges the gap between [1] and the treatment of a memory repertoire that adapts to changing targets in [3]. Indeed, it would be interesting to see how the adaptation of stored memory and production of new memory can be combined and how this might affect the memory performance.

Beyond immune memory, it would be interesting to study strategies for phenotype switching similar to [95] but include the evolution of environments. Because the phenotypes are connected through the genotype, a change (adaptation) in one phenotype might significantly influence other phenotypes. This behavior resembles the fully connected Hopfield network we studied in [2]. Thus, it would be interesting to see if, similar to the attractors in the network, some phenotypes become unstable when the system tries to adapt to changing environments.

Interestingly, Xue et al. trained a neural network to find optimal strategies depending on switching rates and environmental noise levels [95]. Using a similar approach when environments can evolve might be helpful to study the limits of artificial neural networks further. So beyond studying phenotype strategies, it would be interesting to investigate which network structures are able to find these strategies.

Indeed, understanding the learning of evolving environments remains an open question in machine learning. Even though some specialized algorithms have been successfully implemented [14, 127], a general theory for the learning of evolving signals does not yet exist [127]. With the treatment of Hopfield networks in [2] and learning strategies of other energy-based models [3], this thesis offers a starting point in the

search for such a theory. Here the results regarding the increased learning rate for evolving signals might be particularly interesting as tuning of this hyperparameter is one of the challenges of current machine learning algorithms with deep neural networks [125, 126].

References

- [1] O. H. Schnaack and A. Nourmohammad. “Optimal evolutionary decision-making to store immune memory”. In: *eLife* 10 (Apr. 2021), e61346. DOI: 10.7554/eLife.61346.
- [2] O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Learning and organization of memory for evolving patterns”. In: *arXiv:2106.02186 [physics]* (June 2021). arXiv: 2106.02186.
- [3] O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Risk-utility tradeoff shapes memory strategies for evolving patterns”. In: *arXiv:2110.15008 [physics]* (Oct. 2021). arXiv: 2110.15008.
- [4] S. E. Turvey and D. H. Broide. “Innate immunity”. In: *Journal of Allergy and Clinical Immunology* 125.2 (Feb. 2010), S24–S32. DOI: 10.1016/j.jaci.2009.07.016.
- [5] A. Iwasaki and R. Medzhitov. “Control of adaptive immunity by the innate immune system”. In: *Nat Immunol* 16.4 (Apr. 2015), pp. 343–353. DOI: 10.1038/ni.3123.
- [6] N. S. De Silva and U. Klein. “Dynamics of B cells in germinal centres”. In: *Nat Rev Immunol* 15.3 (Mar. 2015), pp. 137–148. DOI: 10.1038/nri3804.
- [7] M. Akkaya, K. Kwak, and S. K. Pierce. “B cell memory: building two walls of protection against pathogens”. In: *Nat Rev Immunol* 20.4 (Apr. 2020), pp. 229–238. DOI: 10.1038/s41577-019-0244-2.
- [8] A. Z. Kapikian, R. H. Mitchell, R. M. Chanock, R. A. Shvedoff, and C. E. Stewart. “An epidemiologic study of altered clinical reactivity to respiratory syncytial (RS) virus infection in children previously vaccinated with an inactivated RS virus vaccine”. In: *American Journal of Epidemiology* 89.4 (Apr. 1969), pp. 405–421. DOI: 10.1093/oxfordjournals.aje.a120954.
- [9] L. M. Walker and D. R. Burton. “Passive immunotherapy of viral infections: ‘super-antibodies’ enter the fray”. In: *Nat Rev Immunol* 18.5 (May 2018), pp. 297–308. DOI: 10.1038/nri.2017.148.
- [10] I. A. Cockburn and R. A. Seder. “Malaria prevention: from immunological concepts to effective vaccines and protective antibodies”. In: *Nat Immunol* 19.11 (Nov. 2018), pp. 1199–1211. DOI: 10.1038/s41590-018-0228-6.
- [11] J. A. Bullinaria. “Evolution of learning strategies in changing environments”. In: *Cognitive Systems Research* 52 (Dec. 2018), pp. 429–449. DOI: 10.1016/j.cogsys.2018.07.024.

- [12] R. McElreath and P. Strimling. “When Natural Selection Favors Imitation of Parents”. In: *Current Anthropology* 49.2 (Apr. 2008), pp. 307–316. DOI: 10.1086/524364.
- [13] M. Ehn and K. Laland. “Adaptive strategies for cumulative cultural learning”. In: *Journal of Theoretical Biology* 301 (May 2012), pp. 103–111. DOI: 10.1016/j.jtbi.2012.02.004.
- [14] R. Polikar and C. Alippi. “Guest Editorial Learning in Nonstationary and Evolving Environments”. In: *IEEE Trans. Neural Netw. Learning Syst.* 25.1 (Jan. 2014), pp. 9–11. DOI: 10.1109/TNNLS.2013.2283547.
- [15] L. M. Sompayrac. *How the Immune System Works*. 5. Chichester, West Sussex, UK ; Ames, Iowa: Wiley-Blackwell, 2015.
- [16] T. Kurosaki, K. Kometani, and W. Ise. “Memory B cells”. In: *Nat Rev Immunol* 15.3 (Mar. 2015), pp. 149–159. DOI: 10.1038/nri3802.
- [17] T. G. Phan and S. G. Tangye. “Memory B cells: total recall”. In: *Current Opinion in Immunology* 45 (Apr. 2017), pp. 132–140. DOI: 10.1016/j.coi.2017.03.005.
- [18] F. Weisel and M. Shlomchik. “Memory B Cells of Mice and Humans”. In: *Annu. Rev. Immunol.* 35.1 (Apr. 2017), pp. 255–284. DOI: 10.1146/annurev-immunol-041015-055531.
- [19] J. G. Cyster and C. D. Allen. “B Cell Responses: Cell Interaction Dynamics and Decisions”. In: *Cell* 177.3 (Apr. 2019), pp. 524–540. DOI: 10.1016/j.cell.2019.03.016.
- [20] G. Altan-Bonnet, T. Mora, and A. M. Walczak. “Quantitative immunology for physicists”. In: *Physics Reports* 849 (Mar. 2020), pp. 1–83. DOI: 10.1016/j.physrep.2020.01.001.
- [21] T. K. Starr, S. C. Jameson, and K. A. Hogquist. “Positive and negative selection of T cells”. In: *Annu. Rev. Immunol.* 21.1 (Apr. 2003), pp. 139–176. DOI: 10.1146/annurev.immunol.21.120601.141107.
- [22] K. Wing and S. Sakaguchi. “Regulatory T cells exert checks and balances on self tolerance and autoimmunity”. In: *Nat Immunol* 11.1 (Jan. 2010), pp. 7–13. DOI: 10.1038/ni.1818.
- [23] A. E. Moran and K. A. Hogquist. “T-cell receptor affinity in thymic development: TCR affinity in thymic development”. In: *Immunology* 135.4 (Apr. 2012), pp. 261–267. DOI: 10.1111/j.1365-2567.2011.03547.x.
- [24] L. L. Lu, T. J. Suscovich, S. M. Fortune, and G. Alter. “Beyond binding: antibody effector functions in infectious diseases”. In: *Nat Rev Immunol* 18.1 (Jan. 2018), pp. 46–61. DOI: 10.1038/nri.2017.106.
- [25] J. R. Dunkelberger and W.-C. Song. “Complement and its role in innate and adaptive immune responses”. In: *Cell Res* 20.1 (Jan. 2010), pp. 34–50. DOI: 10.1038/cr.2009.139.
- [26] B. A. Heesters, R. C. Myers, and M. C. Carroll. “Follicular dendritic cells: dynamic antigen libraries”. In: *Nat Rev Immunol* 14.7 (July 2014), pp. 495–504. DOI: 10.1038/nri3689.

-
- [27] T. E. Mandels, R. P. Phippsi, A. Abbot, and J. G. Tew. “The Follicular Dendritic Cell: Long Term Antigen Retention During Immunity”. In: *Immunol Rev* 53.1 (Dec. 1980), pp. 29–59. DOI: 10.1111/j.1600-065X.1980.tb01039.x.
- [28] C. D. Allen and J. G. Cyster. “Follicular dendritic cell networks of primary follicles and germinal centers: Phenotype and function”. In: *Seminars in Immunology* 20.1 (Feb. 2008), pp. 14–25. DOI: 10.1016/j.smim.2007.12.001.
- [29] M. E. M. El Shikh and C. Pitzalis. “Follicular dendritic cells in health and disease”. In: *Front. Immun.* 3 (2012). DOI: 10.3389/fimmu.2012.00292.
- [30] T. G. Phan, E. E. Gray, and J. G. Cyster. “The microanatomy of B cell activation”. In: *Current Opinion in Immunology* 21.3 (June 2009), pp. 258–265. DOI: 10.1016/j.coi.2009.05.006.
- [31] P. Garside. “Visualization of Specific B and T Lymphocyte Interactions in the Lymph Node”. In: *Science* 281.5373 (July 1998), pp. 96–99. DOI: 10.1126/science.281.5373.96.
- [32] T. V. Obukhanych and M. C. Nussenzweig. “T-independent type II immune responses generate memory B cells”. In: *Journal of Experimental Medicine* 203.2 (Feb. 2006), pp. 305–310. DOI: 10.1084/jem.20052036.
- [33] T. Defrance, M. Taillardet, and L. Genestier. “T cell-independent B cell memory”. In: *Current Opinion in Immunology* 23.3 (June 2011), pp. 330–336. DOI: 10.1016/j.coi.2011.03.004.
- [34] A. Bortnick and D. Allman. “What Is and What Should Always Have Been: Long-Lived Plasma Cells Induced by T Cell-Independent Antigens”. In: *J.I.* 190.12 (June 2013), pp. 5913–5918. DOI: 10.4049/jimmunol.1300161.
- [35] C. D. C. Allen, T. Okada, H. L. Tang, and J. G. Cyster. “Imaging of Germinal Center Selection Events During Affinity Maturation”. In: *Science* 315.5811 (Jan. 2007), pp. 528–531. DOI: 10.1126/science.1136736.
- [36] T. A. Schwickert, G. D. Victora, D. R. Fooksman, A. O. Kamphorst, M. R. Mugnier, A. D. Gitlin, M. L. Dustin, and M. C. Nussenzweig. “A dynamic T cell-limited checkpoint regulates affinity-dependent B cell entry into the germinal center”. In: *Journal of Experimental Medicine* 208.6 (June 2011), pp. 1243–1252. DOI: 10.1084/jem.20102477.
- [37] S. M. Anderson, M. M. Tomayko, A. Ahuja, A. M. Haberman, and M. J. Shlomchik. “New markers for murine memory B cells that define mutated and unmutated subsets”. In: *Journal of Experimental Medicine* 204.9 (Sept. 2007), pp. 2103–2114. DOI: 10.1084/jem.20062571.
- [38] J. J. Taylor, K. A. Pape, and M. K. Jenkins. “A germinal center-independent pathway generates unswitched memory B cells early in the primary response”. In: *Journal of Experimental Medicine* 209.3 (Mar. 2012), pp. 597–606. DOI: 10.1084/jem.20111696.
- [39] T. Kaji et al. “Distinct cellular pathways select germline-encoded and somatically mutated antibodies into immunological memory”. In: *Journal of Experimental Medicine* 209.11 (Oct. 2012), pp. 2079–2097. DOI: 10.1084/jem.20120127.

- [40] C. D. Allen, T. Okada, and J. G. Cyster. “Germinal-Center Organization and Cellular Dynamics”. In: *Immunity* 27.2 (Aug. 2007), pp. 190–202. DOI: 10.1016/j.immuni.2007.07.009.
- [41] D. McKean, K. Huppi, M. Bell, L. Staudt, W. Gerhard, and M. Weigert. “Generation of antibody diversity in the immune response of BALB/c mice to influenza virus hemagglutinin.” In: *Proceedings of the National Academy of Sciences* 81.10 (May 1984), pp. 3180–3184. DOI: 10.1073/pnas.81.10.3180.
- [42] S. H. Kleinstein, Y. Louzoun, and M. J. Shlomchik. “Estimating Hypermutation Rates from Clonal Tree Data”. In: *J Immunol* 171.9 (Nov. 2003), pp. 4639–4649. DOI: 10.4049/jimmunol.171.9.4639.
- [43] K. Kwak et al. “Intrinsic properties of human germinal center B cells set antigen affinity thresholds”. In: *Sci. Immunol.* 3.29 (Nov. 2018), eaau6598. DOI: 10.1126/sciimmunol.aau6598.
- [44] M. Meyer-Hermann. “A molecular theory of germinal center B cell selection and division”. In: *Cell Reports* 36.8 (Aug. 2021), p. 109552. DOI: 10.1016/j.celrep.2021.109552.
- [45] I. C. M. MacLennan. “Germinal Centers”. In: *Annu. Rev. Immunol.* 12.1 (Apr. 1994). Publisher: Annual Reviews, pp. 117–139. DOI: 10.1146/annurev.iy.12.040194.001001.
- [46] G. D. Victora and M. C. Nussenzweig. “Germinal Centers”. In: *Annu. Rev. Immunol.* 30.1 (Apr. 2012), pp. 429–457. DOI: 10.1146/annurev-immunol-020711-075032.
- [47] R. Nakagawa and D. P. Calado. “Positive Selection in the Light Zone of Germinal Centers”. In: *Front. Immunol.* 12 (Mar. 2021), p. 661678. DOI: 10.3389/fimmu.2021.661678.
- [48] M. Meyer-Hermann, E. Mohr, N. Pelletier, Y. Zhang, G. D. Victora, and K.-M. Toellner. “A Theory of Germinal Center B Cell Selection, Division, and Exit”. In: *Cell Reports* 2.1 (July 2012), pp. 162–174. DOI: 10.1016/j.celrep.2012.05.010.
- [49] S. Wang, J. Mata-Fink, B. Kriegsman, M. Hanson, D. J. Irvine, H. N. Eisen, D. R. Burton, K. D. Wittrup, M. Kardar, and A. K. Chakraborty. “Manipulating the Selection Forces during Affinity Maturation to Generate Cross-Reactive HIV Antibodies”. In: *Cell* 160.4 (Feb. 2015), pp. 785–797. DOI: 10.1016/j.cell.2015.01.027.
- [50] R. Manz. “Survival of long-lived plasma cells is independent of antigen [In Process Citation]”. In: *International Immunology* 10.11 (Nov. 1998), pp. 1703–1711. DOI: 10.1093/intimm/10.11.1703.
- [51] C. C. Goodnow, C. G. Vinuesa, K. L. Randall, F. Mackay, and R. Brink. “Control systems and decision making for antibody production.” In: *Nat. Immunol.* 11.8 (Aug. 2010), pp. 681–688.
- [52] M. J. Shlomchik and F. Weisel. “Germinal center selection and the development of memory B and plasma cells: Germinal center differentiation and selection”. In: *Immunological Reviews* 247.1 (May 2012), pp. 52–63. DOI: 10.1111/j.1600-065X.2012.01124.x.

-
- [53] D. Zotos and D. M. Tarlinton. “Determining germinal centre B cell fate”. In: *Trends in Immunology* 33.6 (June 2012), pp. 281–288. DOI: 10.1016/j.it.2012.04.003.
- [54] S. L. Nutt, P. D. Hodgkin, D. M. Tarlinton, and L. M. Corcoran. “The generation of antibody-secreting plasma cells”. In: *Nat Rev Immunol* 15.3 (Mar. 2015), pp. 160–171. DOI: 10.1038/nri3795.
- [55] E. J. Blink, A. Light, A. Kallies, S. L. Nutt, P. D. Hodgkin, and D. M. Tarlinton. “Early appearance of germinal center-derived memory B cells and plasma cells in blood after primary immunization.” In: *The Journal of experimental medicine* 201.4 (Feb. 2005), pp. 545–554.
- [56] K. G. Smith, A. Light, G. J. Nossal, and D. M. Tarlinton. “The extent of affinity maturation differs between the memory and antibody-forming cell compartments in the primary immune response.” In: *EMBO J.* 16.11 (June 1997), pp. 2996–3006.
- [57] T. G. Phan, D. Paus, T. D. Chan, M. L. Turner, S. L. Nutt, A. Basten, and R. Brink. “High affinity germinal center B cells are actively selected into the plasma cell compartment”. In: *Journal of Experimental Medicine* 203.11 (Oct. 2006), pp. 2419–2424. DOI: 10.1084/jem.20061254.
- [58] D. Paus, T. G. Phan, T. D. Chan, S. Gardam, A. Basten, and R. Brink. “Antigen recognition strength regulates the choice between extrafollicular plasma cell and germinal center B cell differentiation.” In: *The Journal of experimental medicine* 203.4 (Apr. 2006), pp. 1081–1091.
- [59] F. J. Weisel, G. V. Zuccarino-Catania, M. Chikina, and M. J. Shlomchik. “A Temporal Switch in the Germinal Center Determines Differential Output of Memory B and Plasma Cells”. In: *Immunity* 44.1 (Jan. 2016), pp. 116–130. DOI: 10.1016/j.immuni.2015.12.004.
- [60] R. Shinnakasu, T. Inoue, K. Kometani, S. Moriyama, Y. Adachi, M. Nakayama, Y. Takahashi, H. Fukuyama, T. Okada, and T. Kurosaki. “Regulated selection of germinal-center cells into the memory B cell compartment”. In: *Nat. Immunol.* 17.7 (July 2016), pp. 861–869.
- [61] T. Recaladin and D. J. Fear. “Transcription factors regulating B cell fate in the germinal centre.” In: *Clin. Exp. Immunol.* 183.1 (Jan. 2016), pp. 65–75.
- [62] R. Shinnakasu and T. Kurosaki. “Regulation of memory B and plasma cell differentiation”. In: *Curr. Opin. Immunol.* 45 (Apr. 2017), pp. 126–131.
- [63] C. Viant, G. H. J. Weymar, A. Escolano, S. Chen, H. Hartweger, M. Cipolla, A. Gazumyan, and M. C. Nussenzweig. “Antibody affinity shapes the choice between memory and germinal center B cell fates”. In: *Cell* 183.5 (Nov. 2020), 1298–1311.e11.
- [64] W. E. Purtha, T. F. Tedder, S. Johnson, D. Bhattacharya, and M. S. Diamond. “Memory B cells, but not long-lived plasma cells, possess antigen specificities for viral escape mutants”. In: *Journal of Experimental Medicine* 208.13 (Dec. 2011), pp. 2599–2606. DOI: 10.1084/jem.20110740.

- [65] S. Leach, R. Shinnakasu, Y. Adachi, M. Momota, C. Makino-Okamura, T. Yamamoto, K. J. Ishii, H. Fukuyama, Y. Takahashi, and T. Kurosaki. “Requirement for memory B-cell activation in protection from heterologous influenza virus reinfection”. In: *International Immunology* 31.12 (Nov. 2019), pp. 771–779. DOI: 10.1093/intimm/dxz049.
- [66] U. Klein, R. Küppers, and K. Rajewsky. “Evidence for a Large Compartment of IgM-Expressing Memory B Cells in Humans”. In: *Blood* 89.4 (Feb. 1997), pp. 1288–1298. DOI: 10.1182/blood.V89.4.1288.
- [67] I. Dogan, B. Bertocci, V. Vilmont, F. Delbos, J. Mégret, S. Storck, C.-A. Reynaud, and J.-C. Weill. “Multiple layers of B cell memory with different effector functions”. In: *Nat Immunol* 10.12 (Dec. 2009), pp. 1292–1299. DOI: 10.1038/ni.1814.
- [68] K. A. Pape, J. J. Taylor, R. W. Maul, P. J. Gearhart, and M. K. Jenkins. “Different B Cell Populations Mediate Early and Late Memory During an Endogenous Immune Response”. In: *Science* 331.6021 (Mar. 2011), pp. 1203–1207. DOI: 10.1126/science.1201730.
- [69] A.-K. E. Palm and C. Henry. “Remembrance of Things Past: Long-Term B Cell Memory After Infection and Vaccination”. In: *Front. Immunol.* 10 (July 2019), p. 1787. DOI: 10.3389/fimmu.2019.01787.
- [70] C. Sundling et al. “Positive selection of IgG+ over IgM+ B cells in the germinal center reaction”. In: *Immunity* 54.5 (May 2021), 988–1001.e5. DOI: 10.1016/j.immuni.2021.03.013.
- [71] K. Kometani, R. Nakagawa, R. Shinnakasu, T. Kaji, A. Rybouchkin, S. Moriyama, K. Furukawa, H. Koseki, T. Takemori, and T. Kurosaki. “Repression of the Transcription Factor Bach2 Contributes to Predisposition of IgG1 Memory B Cells toward Plasma Cell Differentiation”. In: *Immunity* 39.1 (July 2013), pp. 136–147. DOI: 10.1016/j.immuni.2013.06.011.
- [72] S. G. Tangye, D. T. Avery, E. K. Deenick, and P. D. Hodgkin. “Intrinsic Differences in the Proliferation of Naive and Memory Human B Cells as a Mechanism for Enhanced Secondary Immune Responses”. In: *J Immunol* 170.2 (Jan. 2003), pp. 686–694. DOI: 10.4049/jimmunol.170.2.686.
- [73] S. G. Tangye and P. D. Hodgkin. “Divide and conquer: the importance of cell division in regulating B-cell responses”. In: *Immunology* 112.4 (Aug. 2004), pp. 509–520. DOI: 10.1111/j.1365-2567.2004.01950.x.
- [74] G. Blanchard-Rohner, A. S. Pulickal, C. M. Jol-van der Zijde, M. D. Snape, and A. J. Pollard. “Appearance of peripheral blood plasma cells and memory B cells in a primary and secondary immune response in humans”. In: *Blood* 114.24 (Dec. 2009), pp. 4998–5002. DOI: 10.1182/blood-2009-03-211052.
- [75] D. R. Burton, P. Poignard, R. L. Stanfield, and I. A. Wilson. “Broadly Neutralizing Antibodies Present New Prospects to Counter Highly Antigenically Diverse Viruses”. In: *Science* 337.6091 (July 2012), pp. 183–186. DOI: 10.1126/science.1225416.
- [76] P. D. Kwong and J. R. Mascola. “Human Antibodies that Neutralize HIV-1: Identification, Structures, and B Cell Ontogenies”. In: *Immunity* 37.3 (Sept. 2012), pp. 412–425. DOI: 10.1016/j.immuni.2012.08.012.

-
- [77] P. D. Kwong, J. R. Mascola, and G. J. Nabel. “Broadly neutralizing antibodies and the search for an HIV-1 vaccine: the end of the beginning”. In: *Nat Rev Immunol* 13.9 (Sept. 2013), pp. 693–701. DOI: 10.1038/nri3516.
- [78] D. R. Burton and L. Hangartner. “Broadly Neutralizing Antibodies to HIV and Their Role in Vaccine Design”. In: *Annu. Rev. Immunol.* 34.1 (May 2016), pp. 635–659. DOI: 10.1146/annurev-immunol-041015-055515.
- [79] D. Sok and D. R. Burton. “Recent progress in broadly neutralizing antibodies to HIV”. In: *Nat Immunol* 19.11 (Nov. 2018), pp. 1179–1188. DOI: 10.1038/s41590-018-0235-7.
- [80] A. S. Perelson and G. F. Oster. “Theoretical studies of clonal selection: Minimal antibody repertoire size and reliability of self-non-self discrimination”. In: *Journal of Theoretical Biology* 81.4 (Dec. 1979), pp. 645–670. DOI: 10.1016/0022-5193(79)90275-3.
- [81] A. Mayer, V. Balasubramanian, A. M. Walczak, and T. Mora. “How a well-adapting immune system remembers”. In: *PNAS* 116.18 (Apr. 2019), pp. 8815–8823. DOI: 10.1073/pnas.1812810116.
- [82] A. Mayer, V. Balasubramanian, T. Mora, and A. M. Walczak. “How a well-adapted immune system is organized”. In: *Proc Natl Acad Sci USA* 112.19 (May 2015), pp. 5950–5955. DOI: 10.1073/pnas.1421827112.
- [83] J. Marchi, M. Lässig, T. Mora, and A. M. Walczak. “Multi-Lineage Evolution in Viral Populations Driven by Host Immune Systems”. In: *Pathogens* 8.3 (July 2019), p. 115. DOI: 10.3390/pathogens8030115.
- [84] B. T. Grenfell. “Unifying the Epidemiological and Evolutionary Dynamics of Pathogens”. In: *Science* 303.5656 (Jan. 2004), pp. 327–332. DOI: 10.1126/science.1090727.
- [85] R. A. Fisher. *The genetical theory of natural selection*. Oxford: Clarendon Press, 1930. DOI: 10.5962/bhl.title.27468.
- [86] R. A. Fisher. “Average excess and average effect of a gene substitution”. In: *Annals of Eugenics* 11.1 (Jan. 1941), pp. 53–63. DOI: 10.1111/j.1469-1809.1941.tb02272.x.
- [87] V. Mustonen and M. Lässig. “From fitness landscapes to seascapes: non-equilibrium dynamics of selection and adaptation”. In: *Trends in Genetics* 25.3 (Mar. 2009), pp. 111–119. DOI: 10.1016/j.tig.2009.01.002.
- [88] V. Mustonen and M. Lässig. “Fitness flux and ubiquity of adaptive evolution”. In: *Proceedings of the National Academy of Sciences* 107.9 (Mar. 2010), pp. 4248–4253. DOI: 10.1073/pnas.0907953107.
- [89] A. Nourmohammad, J. Rambeau, T. Held, V. Kovacova, J. Berg, and M. Lässig. “Adaptive Evolution of Gene Expression in *Drosophila*”. In: *Cell Reports* 20.6 (Aug. 2017), pp. 1385–1395. DOI: 10.1016/j.celrep.2017.07.033.
- [90] A. Nourmohammad, J. Otwinowski, and J. B. Plotkin. “Host-Pathogen Co-evolution and the Emergence of Broadly Neutralizing Antibodies in Chronic Infections”. In: *PLoS Genet* 12.7 (July 2016). Ed. by S. Cobey, e1006171. DOI: 10.1371/journal.pgen.1006171.

- [91] A. Nourmohammad and C. Eksin. “Optimal Evolutionary Control for Artificial Selection on Molecular Phenotypes”. In: *Phys. Rev. X* 11.1 (Mar. 2021), p. 011044. DOI: 10.1103/PhysRevX.11.011044.
- [92] M. Lässig and V. Mustonen. “Eco-evolutionary control of pathogens”. In: *Proc Natl Acad Sci USA* 117.33 (Aug. 2020), pp. 19694–19704. DOI: 10.1073/pnas.1920263117.
- [93] V. Sachdeva, K. Husain, J. Sheng, S. Wang, and A. Murugan. “Tuning environmental timescales to evolve and maintain generalists”. In: *Proc Natl Acad Sci USA* 117.23 (June 2020), pp. 12693–12699. DOI: 10.1073/pnas.1914586117.
- [94] A. P. Moczek, S. Sultan, S. Foster, C. Ledón-Rettig, I. Dworkin, H. F. Nijhout, E. Abouheif, and D. W. Pfennig. “The role of developmental plasticity in evolutionary innovation”. In: *Proc. R. Soc. B.* 278.1719 (Sept. 2011), pp. 2705–2713. DOI: 10.1098/rspb.2011.0971.
- [95] B. Xue, P. Sartori, and S. Leibler. “Environment-to-phenotype mapping and adaptation strategies in varying environments”. In: *Proc Natl Acad Sci USA* 116.28 (July 2019), pp. 13847–13855. DOI: 10.1073/pnas.1903232116.
- [96] O. Rivoire and S. Leibler. “The Value of Information for Populations in Varying Environments”. In: *J Stat Phys* 142.6 (Apr. 2011), pp. 1124–1166. DOI: 10.1007/s10955-011-0166-2.
- [97] O. Rivoire and S. Leibler. “A model for the generation and transmission of variations in evolution”. In: *Proceedings of the National Academy of Sciences* 111.19 (May 2014), E1940–E1949. DOI: 10.1073/pnas.1323901111.
- [98] O. Rivoire. “Informations in Models of Evolutionary Dynamics”. In: *J Stat Phys* 162.5 (Mar. 2016), pp. 1324–1352. DOI: 10.1007/s10955-015-1381-z.
- [99] A. Mayer, T. Mora, O. Rivoire, and A. M. Walczak. “Diversity of immune strategies explained by adaptation to pathogen statistics”. In: *Proc Natl Acad Sci USA* 113.31 (Aug. 2016), pp. 8630–8635. DOI: 10.1073/pnas.1600663113.
- [100] A. Mayer, T. Mora, O. Rivoire, and A. M. Walczak. “Transitions in optimal adaptive strategies for populations in fluctuating environments”. In: *Phys. Rev. E* 96.3 (Sept. 2017), p. 032412. DOI: 10.1103/PhysRevE.96.032412.
- [101] V. Chardès, M. Vergassola, A. M. Walczak, and T. Mora. *Affinity maturation for an optimal balance between long-term immune coverage and short-term resource constraints*. preprint. Immunology, July 2021. DOI: 10.1101/2021.07.26.453765.
- [102] P. J. Moberg, R. Agrin, R. E. Gur, R. C. Gur, B. I. Turetsky, and R. L. Doty. “Olfactory Dysfunction in Schizophrenia: A Qualitative and Quantitative Review”. In: *Neuropsychopharmacology* 21.3 (Sept. 1999), pp. 325–340. DOI: 10.1016/S0893-133X(99)00019-6.
- [103] C. Bushdid, M. O. Magnasco, L. B. Vosshall, and A. Keller. “Humans Can Discriminate More than 1 Trillion Olfactory Stimuli”. In: *Science* 343.6177 (Mar. 2014), pp. 1370–1372. DOI: 10.1126/science.1249168.
- [104] R. C. Gerkin and J. B. Castro. “The number of olfactory stimuli that humans can discriminate is still unknown”. In: *eLife* 4 (July 2015), e08127. DOI: 10.7554/eLife.08127.

-
- [105] E. J. Mayhew, C. J. Arayata, R. C. Gerkin, B. K. Lee, J. M. Magill, L. L. Snyder, K. A. Little, C. W. Yu, and J. D. Mainland. *Drawing the Borders of Olfactory Space*. preprint. Neuroscience, Dec. 2020. DOI: 10.1101/2020.12.04.412254.
- [106] R. A. Raguso. “Wake Up and Smell the Roses: The Ecology and Evolution of Floral Scent”. In: *Annu. Rev. Ecol. Evol. Syst.* 39.1 (Dec. 2008), pp. 549–569. DOI: 10.1146/annurev.ecolsys.38.091206.095601.
- [107] A. Dunkel, M. Steinhaus, M. Kotthoff, B. Nowak, D. Krautwurst, P. Schieberle, and T. Hofmann. “Nature’s Chemical Signatures in Human Olfaction: A Food-borne Perspective for Future Biotechnology”. In: *Angew. Chem. Int. Ed.* 53.28 (July 2014), pp. 7124–7143. DOI: 10.1002/anie.201309508.
- [108] I. Beyaert and M. Hilker. “Plant odour plumes as mediators of plant-insect interactions: Plant odour plumes”. In: *Biol Rev* 89.1 (Feb. 2014), pp. 68–81. DOI: 10.1111/brv.12043.
- [109] G. Glusman, I. Yanai, I. Rubin, and D. Lancet. “The Complete Human Olfactory Subgenome”. In: *Genome Research* 11.5 (May 2001), pp. 685–702. DOI: 10.1101/gr.171001.
- [110] C. I. Bargmann. “Comparative chemosensation from receptors to ecology”. In: *Nature* 444.7117 (Nov. 2006), pp. 295–301. DOI: 10.1038/nature05402.
- [111] K. Touhara and L. B. Vosshall. “Sensing Odorants and Pheromones with Chemosensory Receptors”. In: *Annu. Rev. Physiol.* 71.1 (Mar. 2009), pp. 307–332. DOI: 10.1146/annurev.physiol.010908.163209.
- [112] C.-Y. Su, K. Menuz, and J. R. Carlson. “Olfactory Perception: Receptors, Cells, and Circuits”. In: *Cell* 139.1 (Oct. 2009), pp. 45–59. DOI: 10.1016/j.cell.2009.09.015.
- [113] C. Verbeurgt, F. Wilkin, M. Tarabichi, F. Gregoire, J. E. Dumont, and P. Chatelain. “Profiling of Olfactory Receptor Gene Expression in Whole Human Olfactory Mucosa”. In: *PLoS ONE* 9.5 (May 2014). Ed. by R. D. Newcomb, e96333. DOI: 10.1371/journal.pone.0096333.
- [114] G. M. Shepherd and C. A. Greer. “Olfactory bulb”. In: *The synaptic organization of the brain, 4th ed.* New York, NY, US: Oxford University Press, 1998, pp. 159–203.
- [115] O. Gschwend, N. M. Abraham, S. Lagier, F. Begnaud, I. Rodriguez, and A. Carleton. “Neuronal pattern separation in the olfactory bulb improves odor discrimination learning”. In: *Nat Neurosci* 18.10 (Oct. 2015), pp. 1474–1482. DOI: 10.1038/nn.4089.
- [116] L. B. Haberly and J. M. Bower. “Olfactory cortex: Model circuit for study of associative memory?” In: *Trends Neurosci.* 12.7 (July 1989), pp. 258–264.
- [117] P. Brennan, H. Kaba, and E. B. Keverne. “Olfactory recognition: A simple memory system”. In: *Science* 250.4985 (Nov. 1990), pp. 1223–1226.
- [118] R. Granger and G. Lynch. “Higher olfactory processes: Perceptual learning and memory”. In: *Curr. Opin. Neurobiol.* 1.2 (Aug. 1991), pp. 209–214.
- [119] L. B. Haberly. “Parallel-distributed processing in olfactory cortex: New insights from morphological and physiological analysis of neuronal circuitry”. In: *Chemical senses* 26.5 (2001). Publisher: Oxford University Press, pp. 551–576.

- [120] D. A. Wilson, A. R. Best, and R. M. Sullivan. “Plasticity in the olfactory system: Lessons for the neurobiology of memory”. In: *Neuroscientist* 10.6 (Dec. 2004), pp. 513–524.
- [121] A. Lansner. “Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations”. In: *Trends in Neurosciences* 32.3 (Mar. 2009), pp. 178–186. DOI: 10.1016/j.tins.2008.12.002.
- [122] J. J. Hopfield. “Neural networks and physical systems with emergent collective computational abilities”. In: *Proc. Natl. Acad. Sci. U.S.A.* 79.8 (Apr. 1982), pp. 2554–2558.
- [123] D. O. Hebb. *The organization of behavior: A neuropsychological theory*. New York: Wiley, 1949.
- [124] J. F. Fontanari and R. Meir. “Learning noisy patterns in a Hopfield network”. In: *Phys. Rev. A* 40.5 (Sept. 1989), pp. 2806–2809. DOI: 10.1103/PhysRevA.40.2806.
- [125] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT Press, 2016.
- [126] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab. “A high-bias, low-variance introduction to Machine Learning for physicists”. In: *Physics Reports* 810 (May 2019), pp. 1–124. DOI: 10.1016/j.physrep.2019.03.001.
- [127] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar. “Learning in Nonstationary Environments: A Survey”. In: *IEEE Comput. Intell. Mag.* 10.4 (Nov. 2015), pp. 12–25. DOI: 10.1109/MCI.2015.2471196.
- [128] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton: Princeton University Press, 1944.
- [129] P. A. Ortega and D. A. Braun. “Thermodynamics as a theory of decision-making with information-processing costs”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 469.2153 (May 2013), p. 20120683. DOI: 10.1098/rspa.2012.0683.
- [130] J. Grau-Moya, M. Krüger, and D. Braun. “Non-Equilibrium Relations for Bounded Rational Decision-Making in Changing Environments”. In: *Entropy* 20.1 (Dec. 2017), p. 1. DOI: 10.3390/e20010001.
- [131] E. T. Jaynes. “Information theory and statistical mechanics”. In: *Physical Review* 106.4 (May 1957), pp. 620–630.
- [132] C. Jarzynski. “Equalities and Inequalities: Irreversibility and the Second Law of Thermodynamics at the Nanoscale”. In: *Annu. Rev. Condens. Matter Phys.* 2.1 (Mar. 2011), pp. 329–351. DOI: 10.1146/annurev-conmatphys-062910-140506.
- [133] W. Little. “The existence of persistent states in the brain”. In: *Mathematical Biosciences* 19.1-2 (Feb. 1974), pp. 101–120. DOI: 10.1016/0025-5564(74)90031-5.
- [134] S. Recanatesi, M. Katkov, S. Romani, and M. Tsodyks. “Neural Network Model of Memory Retrieval”. In: *Front. Comput. Neurosci.* 9 (Dec. 2015). DOI: 10.3389/fncom.2015.00149.
- [135] U. Pereira and N. Brunel. “Attractor Dynamics in Networks with Learning Rules Inferred from In Vivo Data”. In: *Neuron* 99.1 (July 2018), 227–238.e4. DOI: 10.1016/j.neuron.2018.05.038.

-
- [136] M. Mézard, J. Nadal, and G. Toulouse. “Solvable models of working memories”. In: *J. Phys. France* 47.9 (1986), pp. 1457–1462. DOI: 10.1051/jphys:019860047090145700.
- [137] D. J. Amit, H. Gutfreund, and H. Sompolinsky. “Storing Infinite Numbers of Patterns in a Spin-Glass Model of Neural Networks”. In: *Phys. Rev. Lett.* 55.14 (Sept. 1985), pp. 1530–1533. DOI: 10.1103/PhysRevLett.55.1530.
- [138] R. McEliece, E. Posner, E. Rodemich, and S. Venkatesh. “The capacity of the Hopfield associative memory”. In: *IEEE Transactions on Information Theory* 33.4 (1987), pp. 461–482. DOI: 10.1109/TIT.1987.1057328.
- [139] L. Personnaz, I. Guyon, and G. Dreyfus. “Collective computational properties of neural networks: New learning mechanisms”. In: *Phys. Rev. A* 34.5 (Nov. 1986), pp. 4217–4228. DOI: 10.1103/PhysRevA.34.4217.
- [140] W. Krauth and M. Mezard. “Learning algorithms with optimal stability in neural networks”. In: *J. Phys. A: Math. Gen.* 20.11 (Aug. 1987), pp. L745–L752. DOI: 10.1088/0305-4470/20/11/013.
- [141] S. Diederich and M. Opper. “Learning of correlated patterns in spin-glass networks by local learning rules”. In: *Phys. Rev. Lett.* 58.9 (Mar. 1987), pp. 949–952. DOI: 10.1103/PhysRevLett.58.949.
- [142] E. Gardner. “The space of interactions in neural network models”. In: *J. Phys. A: Math. Gen.* 21.1 (Jan. 1988), pp. 257–270. DOI: 10.1088/0305-4470/21/1/030.
- [143] A. Storkey. “Increasing the capacity of a hopfield network without sacrificing functionality”. In: *Artificial Neural Networks — ICANN’97*. Ed. by G. Goos, J. Hartmanis, J. van Leeuwen, W. Gerstner, A. Germond, M. Hasler, and J.-D. Nicoud. Vol. 1327. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997, pp. 451–456. DOI: 10.1007/BFb0020196.
- [144] P. Tolmachev and J. H. Manton. “New Insights on Learning Rules for Hopfield Networks: Memory and Objective Function Minimisation”. In: *2020 International Joint Conference on Neural Networks (IJCNN)*. Glasgow, United Kingdom: IEEE, July 2020, pp. 1–8. DOI: 10.1109/IJCNN48605.2020.9207405.
- [145] D. Krotov and J. J. Hopfield. “Dense associative memory for pattern recognition”. In: *Advances in neural information processing systems*. Ed. by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett. Vol. 29. Curran Associates, Inc., 2016.
- [146] G. J. Wedemayer, P. A. Patten, L. H. Wang, P. G. Schultz, and R. C. Stevens. “Structural Insights into the Evolution of an Antibody Combining Site”. In: *Science* 276.5319 (June 1997), pp. 1665–1669. DOI: 10.1126/science.276.5319.1665.
- [147] Y. Li, H. Li, F. Yang, S. J. Smith-Gill, and R. A. Mariuzza. “X-ray snapshots of the maturation of an antibody response to a protein antigen”. In: *Nat Struct Mol Biol* 10.6 (June 2003), pp. 482–488. DOI: 10.1038/nsb930.
- [148] N. C. Wu, G. Grande, H. L. Turner, A. B. Ward, J. Xie, R. A. Lerner, and I. A. Wilson. “In vitro evolution of an influenza broadly neutralizing antibody is modulated by hemagglutinin receptor specificity”. In: *Nat Commun* 8.1 (Aug. 2017), p. 15371. DOI: 10.1038/ncomms15371.

- [149] A. K. Mishra and R. A. Mariuzza. “Insights into the Structural Basis of Antibody Affinity Maturation from Next-Generation Sequencing”. In: *Front. Immunol.* 9 (Feb. 2018), p. 117. DOI: 10.3389/fimmu.2018.00117.
- [150] M. L. Fernández-Quintero, J. R. Loeffler, L. M. Bacher, F. Waibl, C. A. Seidler, and K. R. Liedl. “Local and Global Rigidification Upon Antibody Affinity Maturation”. In: *Front. Mol. Biosci.* 7 (Aug. 2020), p. 182. DOI: 10.3389/fmolb.2020.00182.
- [151] A. Vatti, D. M. Monsalve, Y. Pacheco, C. Chang, J.-M. Anaya, and M. E. Gershwin. “Original antigenic sin: A comprehensive review”. In: *Journal of Autoimmunity* 83 (Sept. 2017), pp. 12–21. DOI: 10.1016/j.jaut.2017.04.008.
- [152] S. Cobey and S. E. Hensley. “Immune history and influenza virus susceptibility”. In: *Current Opinion in Virology* 22 (Feb. 2017), pp. 105–111. DOI: 10.1016/j.coviro.2016.12.004.
- [153] A. Nourmohammad, J. Otwinowski, M. Łuksza, T. Mora, and A. M. Walczak. “Fierce Selection and Interference in B-Cell Repertoire Response to Chronic HIV-1”. In: *Molecular Biology and Evolution* 36.10 (Oct. 2019). Ed. by T. Leitner, pp. 2184–2194. DOI: 10.1093/molbev/msz143.
- [154] D. J. Amit, H. Gutfreund, and H. Sompolinsky. “Spin-glass models of neural networks”. In: *Phys. Rev. A* 32.2 (Aug. 1985), pp. 1007–1018. DOI: 10.1103/PhysRevA.32.1007.
- [155] J. Fontanari. “Generalization in a Hopfield network”. In: *J. Phys. France* 51.21 (1990), pp. 2421–2430. DOI: 10.1051/jphys:0199000510210242100.

Appendix A

Figure supplements for optimal evolutionary decision-making to store immune memory

The following chapter consists of the figure supplements of the article

O. H. Schnaack and A. Nourmohammad. “Optimal evolutionary decision-making to store immune memory”. In: *eLife* 10 (Apr. 2021), e61346. DOI: 10.7554/eLife.61346

in chapter 3 reprinted under the Creative Commons Attribution License (CC BY 4.0)

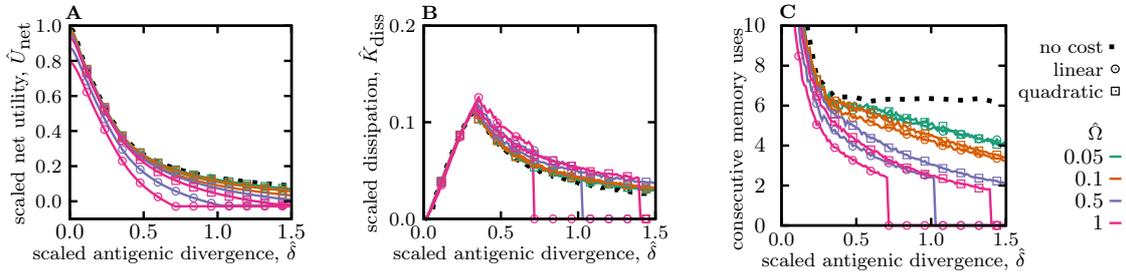


Figure A.1 Figure 2—figure supplement 1. **Utility, dissipation, and usage of optimal memory.** (A) and (B) show the scaled net utility $\hat{U}_{\text{net}} \equiv U_{\text{net}}/E_{\text{max}}$ (Equation 1) and dissipation $\hat{K}_{\text{diss}} \equiv K_{\text{diss}}/E_{\text{max}}$ (Equation 2) per round of infection as a function of the antigenic divergence $\hat{\delta}$. Rescaling by E_{max} sets the magnitude of net utility for a response to conserved antigens (with $\hat{\delta} = 0$), and in the limit of zero deliberation cost $\hat{\Omega} \rightarrow 0$, to 1; see Fig. 2 in the main text for comparison. (C) The expected number of rounds that a memory receptor is utilized prior to a novel response in an optimal system is shown to decay as the antigenic divergence $\hat{\delta}$ increases. The results are evaluated for immune systems with optimized strategies ($\hat{\alpha}^*$, $\hat{\beta}^*$) against pathogens with a given scaled antigenic divergence $\hat{\delta}$; the corresponding strategies are shown in Fig. 2. Colors / markers indicate different naïve cost functions for deliberation, including no-cost $\hat{\Omega} \equiv \Omega/E_{\text{max}} = 0$, linear cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$, and quadratic cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}^2$, with varying amplitudes $\hat{\Omega}_0$. Simulation parameters: $\alpha_{\text{max}} = 4$, $\beta_{\text{max}} = 10$, and $\theta = 2$.

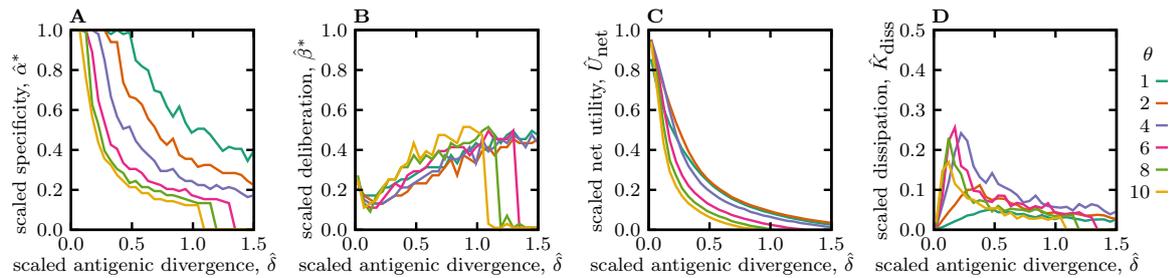


Figure A.2 Figure 2—figure supplement 2. **Optimal memory strategies for different specificity shape factors θ .** (A) Scaled specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\text{max}}$, (B) scaled deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\text{max}}$, (C) scaled net utility $\hat{U}_{\text{net}} \equiv U_{\text{net}}/E_{\text{max}}$, and (D) scaled dissipation are shown as a function of the scaled antigenic divergence per infection $\hat{\delta} = \delta/(\alpha_{\text{max}}^{-1})$ (similar to Fig. 2). Colors indicate different shape factors θ of the specificity function, ranging from a double-sided exponential ($\theta = 1$), to Gaussian for $\theta = 2$ (as in Fig. 2), and top-hat functions $\theta > 2$. The dependence of memory strategies on antigenic divergence is qualitatively insensitive to the shape factor of the specificity function. Simulation parameters: linear deliberation cost function $\Omega = \hat{\Omega}_0\hat{\beta}$ with $\hat{\Omega}_0 = 0.1$, $\alpha_{\text{max}} = 4$, and $\beta_{\text{max}} = 10$.

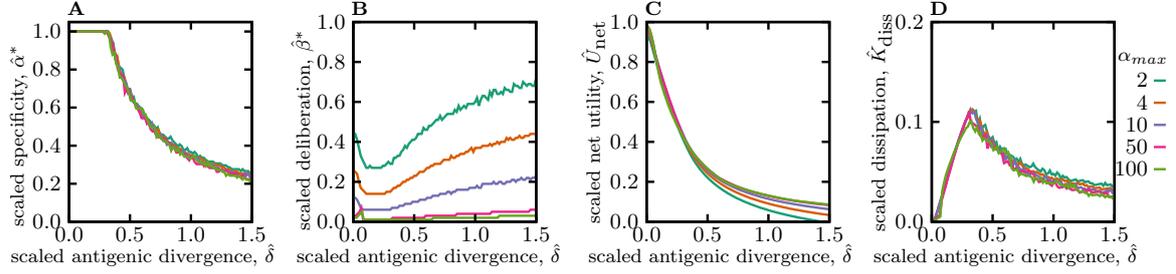


Figure A.3 Figure 2—figure supplement 3. **Optimal memory strategies for different specificity thresholds α_{\max} .** (A) Scaled specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\max}$, (B) scaled deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\max}$, (C) scaled net utility $\hat{U}_{\text{net}} \equiv U_{\text{net}}/E_{\max}$, and (D) scaled dissipation \hat{K}_{diss} are shown as a function of the scaled antigenic divergence per infection $\hat{\delta} = \delta/(\alpha_{\max}^{-1})$ (similar to Fig. 2). Colors indicate different specificity thresholds α_{\max} . Memory strategies are qualitatively insensitive to the specificity threshold. Simulation parameters: linear deliberation cost function $\Omega = \hat{\Omega}_0 \hat{\beta}$ with $\hat{\Omega}_0 = 0.1$ and $\beta_{\max} = 10$.

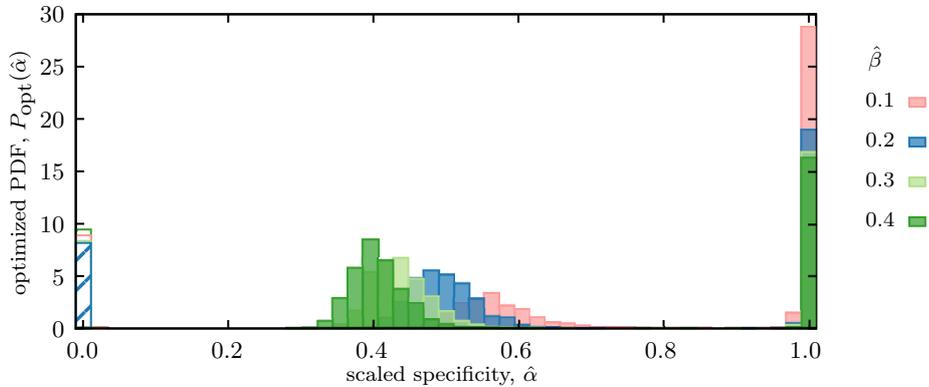


Figure A.4 Figure 3—figure supplement 1. **Mixed memory strategy against pathogens for different deliberation factors $\hat{\beta}$.** Distribution of scaled optimized specificities $\hat{\alpha}^*$ of functional memories is shown for an immune system with a fixed deliberation factor $\hat{\beta} = 0.2$, in which a mixture strategy with a bimodal distribution of specificities $P(\hat{\alpha})$ is established to counter pathogens with a broad range of antigenic divergences, drawn uniformly from a range $\hat{\delta} \in (0, 1.6)$ (similar to Fig. 3). The dashed bars indicate stored memory with specificity $\alpha = 0$, which is not further used in response to infections. Colors indicate different deliberation factors. Simulation parameters: $\alpha_{\max} = 4$, and $\beta_{\max} = 10$.

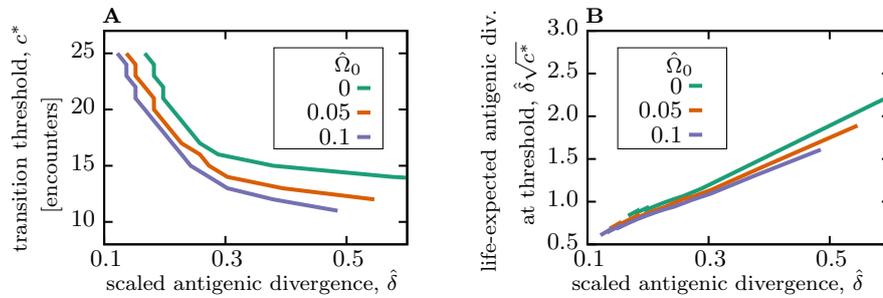


Figure A.5 Figure 4—figure supplement 1. **Pathogen encounter threshold to transition between cross-reactive and specific memory.** (A) The encounter threshold c^* , shown in Fig. 4A,B, decays as a function of the antigenic divergence (per encounter) $\hat{\delta}$ and the amplitude of the naive cost $\hat{\Omega}_0$ (colors). (B) The expected antigenic divergence for the duration of c^* (threshold) encounters $\hat{\delta}\sqrt{c^*}$ is shown as a function of antigenic divergence (per encounter) $\hat{\delta}$. Simulation parameters: linear deliberation cost function $\Omega = \Omega_0\hat{\beta}$, $\alpha_{\max} = 4$, $\beta_{\max} = 10$, and $\theta = 2$.

Appendix B

Supplementary figures for learning and organization of memory for evolving patterns

The following chapter consists of the figure supplements of the preprint

O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Learning and organization of memory for evolving patterns”. In: *arXiv:2106.02186 [physics]* (June 2021). arXiv: 2106.02186

in chapter 4.

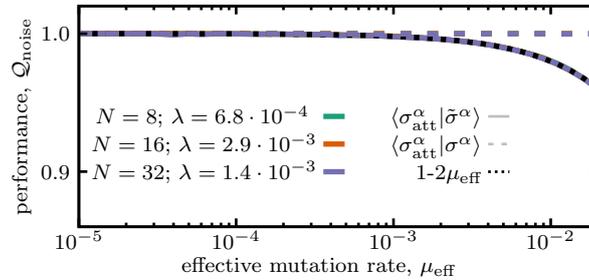


Figure B.1 S1 Performance of networks trained with noisy inputs. A network is trained with inputs $\tilde{\sigma}^\alpha$ which are the noisy version of the true static patterns σ^α . The noisy inputs are generated by randomly flipping spins of a static pattern with a rate according to the effective mutation rate, indicated on the x-axis. The performance of the network is shown for different number of patterns (colors), which collapse on top of each other. As the noise increases, the overlap between the attractor $\sigma_{\text{att}}^\alpha$ and the noisy version of the patterns $\tilde{\sigma}^\alpha$ (full line) decays and falls below the expected overlap for the true static patterns σ^α at one (dashed line). The networks are trained with learning rate λ , which corresponds to the optimal learning rate for nearly static patterns with $\mu_{\text{eff}} = 1.8 \cdot 10^{-5}$. Other simulation parameters: $L = 800$.

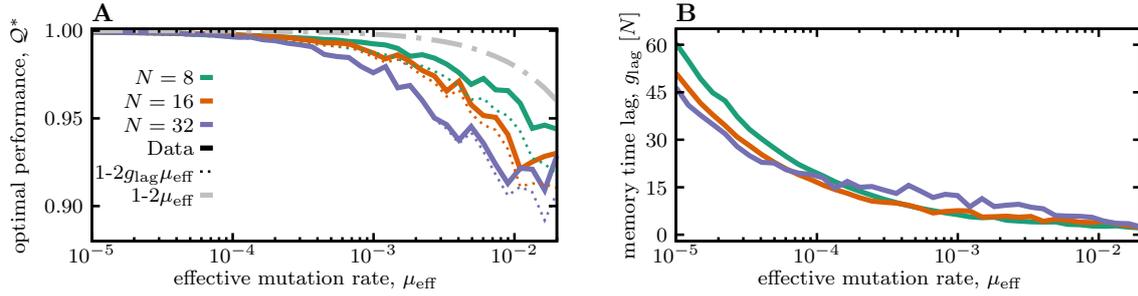


Figure B.2 S2 Reduced performance of Hopfield networks due to memory delay. (A) The optimal performance Q^* for patterns that are correctly associated with their memory attractors (i.e., they have an overlap $q(\sigma) = \langle \sigma_{\text{att}} | \sigma \rangle > 0.8$) is shown as a function of the effective mutation rate μ_{eff} . The solid lines show the simulation results for networks encountering a different number of patterns N (colors). The gray dashed line shows the naïve expectation for the performance ($Q_0 = 1 - 2\mu_{\text{eff}}$), and the colored dashed lines show the expected performance after accounting for the memory lag $Q_{\text{lag}} = 1 - 2g_{\text{lag}}\mu_{\text{eff}}$. (B) The lag time g_{lag} for memory is shown in units of generations $[N]$ as a function of the effective mutation rate for networks encountering a different number of patterns (colors similar to (A)). The networks are trained with a learning rate $\lambda^*(\mu)$ optimized for the mutation rate specified on the x-axis. Other simulation parameters: $L = 800$.

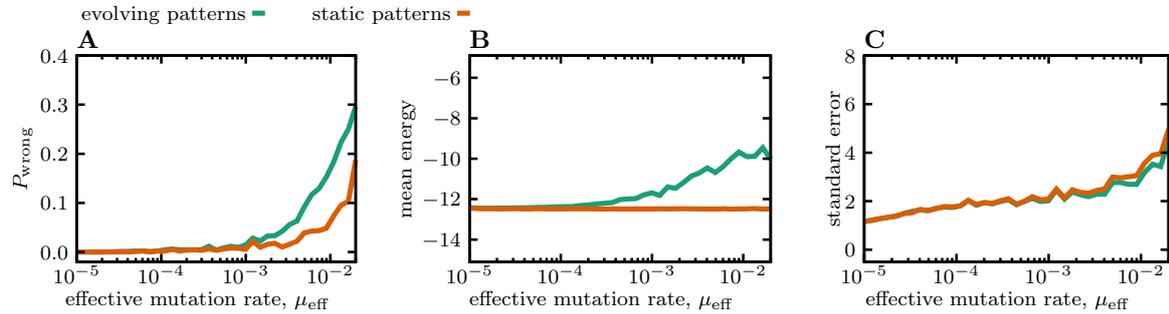


Figure B.3 S3 Statistics of static and evolving patterns for networks with different learning rates. We compare the statistics of evolving (green) and static (orange) patterns in networks trained with a learning rate $\lambda^*(\mu)$ optimized for the mutation rate specified on each panel's x-axis; see Fig. 2B for dependency of the optimal learning rate on mutation rate. The reported statistics are (A) Fraction P_{wrong} of misclassified patterns (i.e., patterns with a small overlap $q(\sigma) = \langle \sigma_{\text{att}} | \sigma \rangle < 0.8$), (B) the mean energy of the patterns, and (C) the standard error of the energy of the patterns in the network. Simulation parameters: $L = 800$ and $N = 32$.

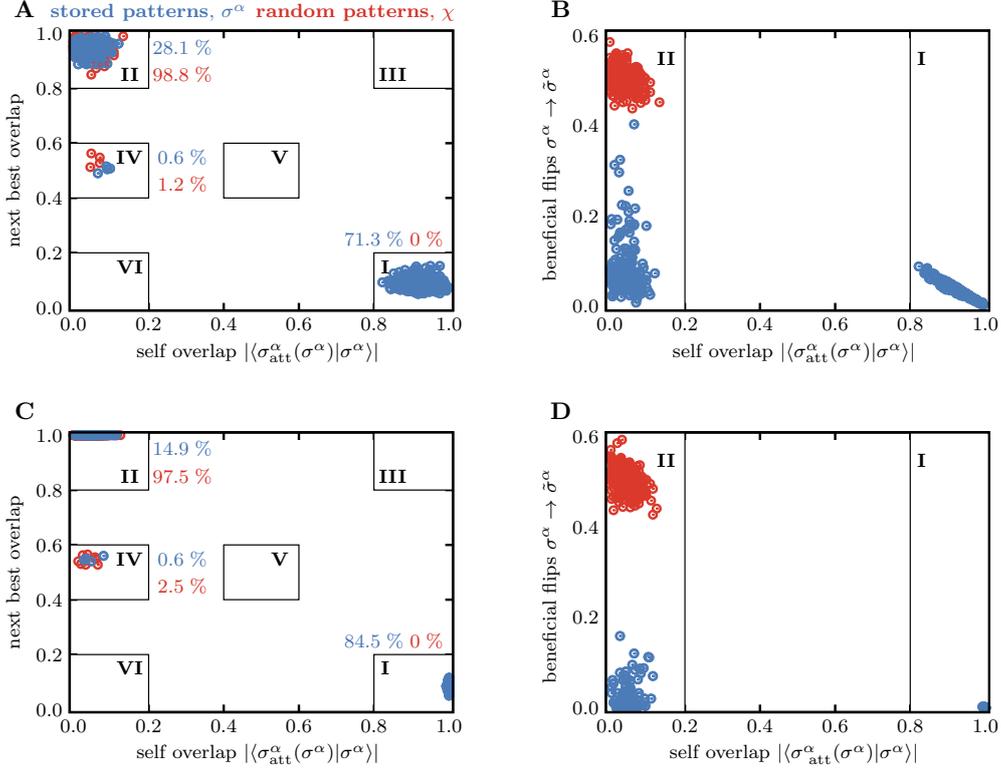


Figure B.4 S4 Attractors and equilibration paths in networks. Overlap of patterns with the networks' attractors are shown for both the patterns σ^α associated with one of the classes that were previously presented to the network during training (blue) and the random patterns χ that are on average orthogonal to the previously presented classes. **(A)** The overlap between a presented pattern σ^α and the memory associated with the same pattern class $\sigma_{\text{att}}^\alpha(\sigma^\alpha)$ is shown against the overlap of the pattern with the next best memory attractor associated with any of the other presented pattern classes $\max_{\nu \neq \alpha} |\langle \sigma_{\text{att}}^\alpha | \sigma^\nu \rangle|$. Fractions of the previously presented patterns and the random patterns that fall into different sectors of the plot are indicated in blue and red, respectively. Sector **I** corresponds to patterns that fall into the correct energy attractors (i.e., $\langle \sigma_{\text{att}}^\alpha | \sigma^\alpha \rangle \approx 1$). In the limit of large self-overlap, the maximal overlap to any other pattern family is close to zero, and thus, no patterns are found in sector **III**. Patterns with a small self-overlap could fall into three different sectors: Sector **II** corresponds to misclassified patterns that fall into a valley associated with a different class ($\max_{\nu \neq \alpha} |\langle \sigma_{\text{att}}^\alpha | \sigma^\nu \rangle| \approx 1$). Patterns in sectors **IV** and **V** fall into local valleys between the minima of two pattern families. This *mixture states* are well known in the classical Hopfield model [154, 155]. Sector **VI** indicates patterns that fall into an attractor in the network that does not correspond to any of the previously presented classes. The fact that neither the previously presented patterns nor the random patterns fall into this sector suggests that the network indeed only stores memory of the presented patterns and is not in the glassy regime. **(B)** The number of beneficial spin-flips for presented pattern at the beginning of equilibration (i.e., the number of open equilibration paths) is shown against patterns' self-overlap (x-axis in (A)). For stable patterns (sector **I**) the number of open paths is anti-correlated with the overlap between the attractor and the presented pattern. For unstable patterns (sector **II**), the number of open paths is on average larger than that of the stable patterns. However, there are fewer paths available to the previously presented patterns compared to the random patterns. In (A,B) patterns evolve with rate $\mu_{\text{eff}} = 0.01$ and the network's learning rate is optimized accordingly. The sharp transition between sector occupations indicates that our results are insensitive to the classification threshold for self-overlap (currently set to $q^\alpha > 0.8$), i.e. any threshold value between sectors **I** and **II** would result in the same classification of patterns. **(C,D)** Similar to (A,B) but for static patterns in a network with a similar learning rate to (A,B). Simulation parameters: $L = 800$ and $N = 32$.

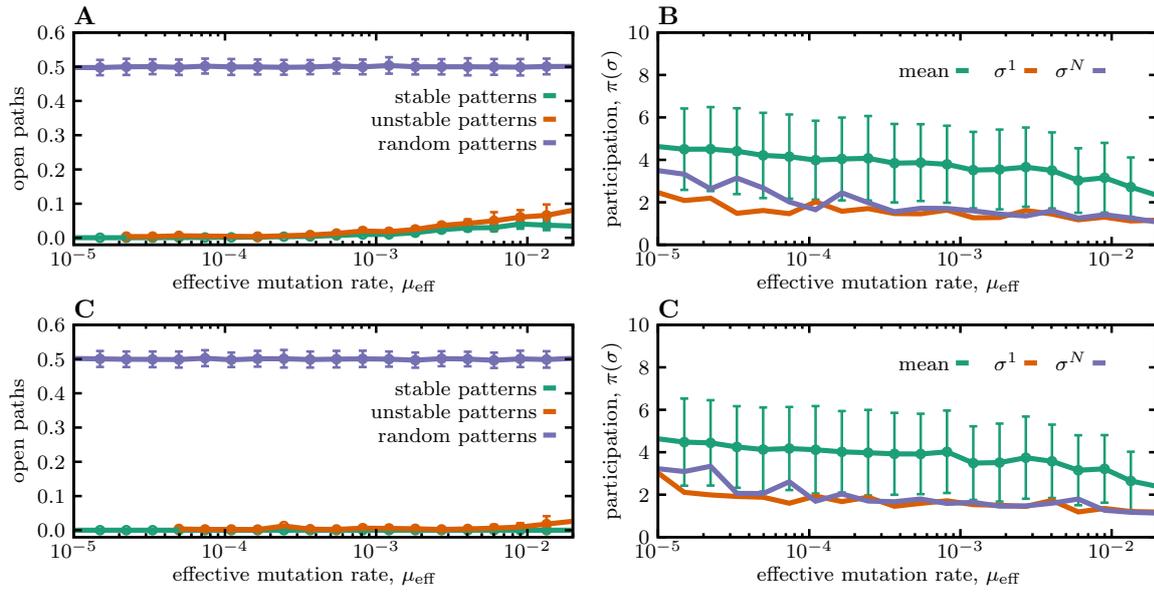


Figure B.5 S5 Open equilibration paths and participation ratio. (A) The mean number of open paths (i.e., the beneficial spin-flips at the beginning of equilibration) is shown for stable, unstable, and random patterns (colors) as a function of the effective mutation rate μ_{eff} in networks trained with the optimal learning rate $\lambda^*(\mu)$. (B) The participation ratio $\pi(\sigma^j) = \frac{(\sum_i m_{i,j}^2)^2}{\sum_i m_{i,j}^4}$, with $m_{i,j} = \langle \Phi^i | \sigma^j \rangle$ is shown for the pattern σ^1 with the lowest energy (orange), the 1 pattern σ^N with the highest energy (purple). The mean participation ratio averaged over all patterns is shown in green. (C,D) Similar to (A,B) but for static patterns ($\mu = 0$). The learning rate of the network in this case is tuned to be optimal for the mutation rate specified on the x-axis. Simulation parameters: $L = 800$ and $N = 32$.

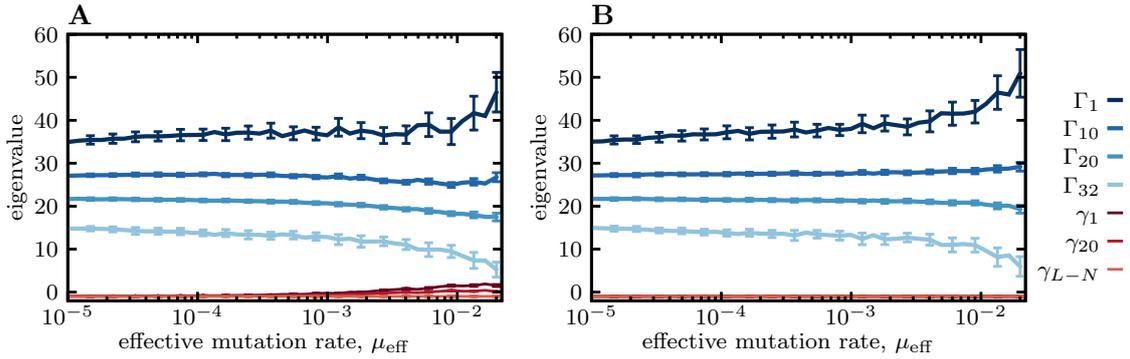


Figure B.6 S6 Eigenvalues of networks with memory against dynamic and static patterns. (A) The first Γ_1 , the 10-th (Γ_{10}), the 20-th (Γ_{20}), and the last ($\Gamma_{N=32}$) non-trivial eigenvalues of a network of size $L = 800$ presented with $N = 32$ patterns is shown as a function the patterns' effective mutation rate (different shades of blue). In each case, the network is trained with the optimal learning rate $\lambda^*(\mu)$. The trivial eigenvalues are shown in different shades of red, with their rank indicated in the legend. For small μ_{eff} all trivial eigenvalues match the prediction $\gamma_k = -1$, which implies that the network updates fast enough to keep the patterns within the N -dimensional sub-space. For larger mutation rates, some of the trivial eigenvalues deviate from -1 , indicating that evolving patterns start spanning in a larger sub-space. Moreover, as the mutation rate (or learning rate) increases the gap between the non-trivial eigenvalues broadens. (B) Similar to (A) but for static patterns in networks trained with a learning rate $\lambda^*(\mu)$ optimized for the mutation rate specified on the panel's x-axis. In contrast to (A) all trivial eigenvalues remain equal to -1 independent of the learning rate, implying that the static patterns remained within the non-trivial N -dimensional sub-space. Similar to (A) the gap between the nontrivial eigenvalues broadens with increasing learning rate.

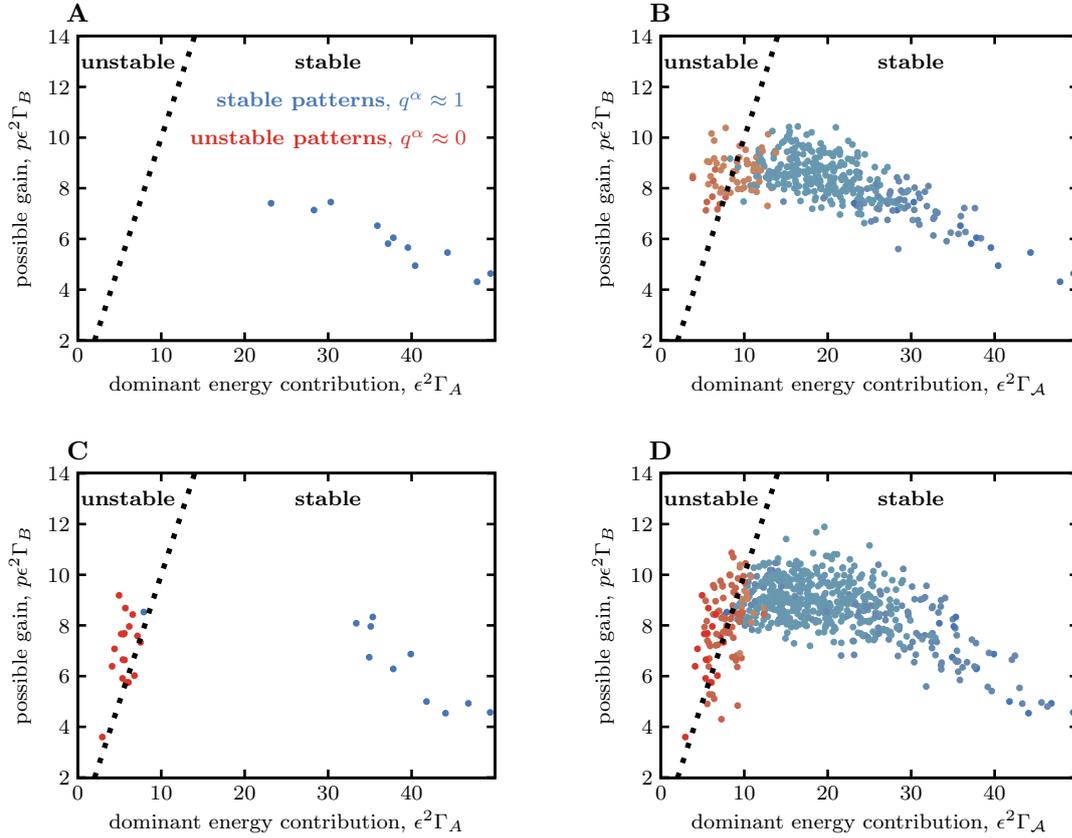


Figure B.7 S7 Stability condition for patterns during equilibration. The stability condition in eq. C14 (dotted line) is used to classify stable (blue) and unstable (red) patterns for **(A)** the patterns that have a squared overlap with one dominant eigenvector $m^2 = \langle \Phi^A | \sigma^\nu \rangle^2 > 0.85$, and **(B)** the patterns that are constrained to a small sub-space \mathcal{A} spanned by up to 10 nontrivial eigenvectors; in this case, $m_{\mathcal{A}}^2 = \sum_{a \in \mathcal{A}} \langle \Phi^a | \sigma^\nu \rangle^2 > 0.85$. The shading indicate the number of eigenvectors needed to represent a pattern from dark (one) to light (ten). **(C, D)** Similar to (A, B) but for static patterns in networks trained with the same learning rate as in (A, B). In general more static patterns reach the threshold of $m > 0.85$ as these patterns remain constrained to the N-dimensional subspace spanned by the non-trivial eigenvectors $\{\Phi_i\}$. Simulation parameters: $N = 32$, $L = 800$, $\mu_{\text{eff}} = 0.02$, and networks are trained with the optimal learning rate $\lambda^*(\mu)$.

Appendix C

Supplementary figures for risk-utility tradeoff shapes memory strategies for evolving patterns

The following chapter consists of the figure supplements of the preprint

O. H. Schnaack, L. Peliti, and A. Nourmohammad. “Risk-utility tradeoff shapes memory strategies for evolving patterns”. In: *arXiv:2110.15008 [physics]* (Oct. 2021). arXiv: 2110.15008

in chapter 5.

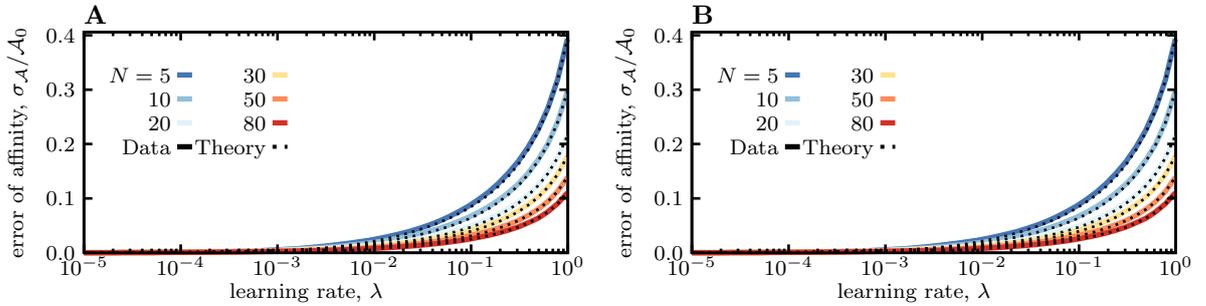


Figure C.1 S1 Standard deviation of affinities for pre-encountered patterns. Solid lines shown the numerical estimates for the standard deviation of pattern affinities ($\sigma_{\mathcal{A}}$) divided by the scale of the affinity (\mathcal{A}_0) for different numbers of patterns N (colors), using Hopfield model to simulate of repertoires. Dotted lines show the analytical estimates, using (A) the full solution in Eq. S14, and (B) the approximation in Eqs. 4 & S22. Simulation parameters: $L = 100$, $\mu_{\text{eff}} = 0.01$, $\Theta = 2$.

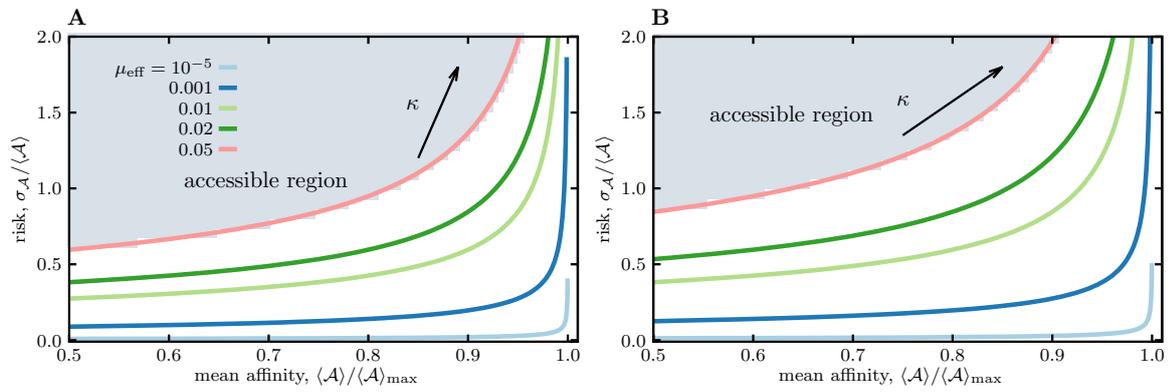


Figure C.2 S2 Pareto front for risk-utility tradeoff of memory for evolving patterns. Similar to Fig. 2 in the main text, the risk-affinity Pareto front for the optimized objective function $\mathcal{Q}(\lambda^*)$ (Eq. 3) is shown for different mutation rates (colors) by varying the risk tolerance κ along each line, for the shape parameters **(A)** $\Theta = 4$, and **(B)** $\Theta = 8$. In both cases, $N = 200$.

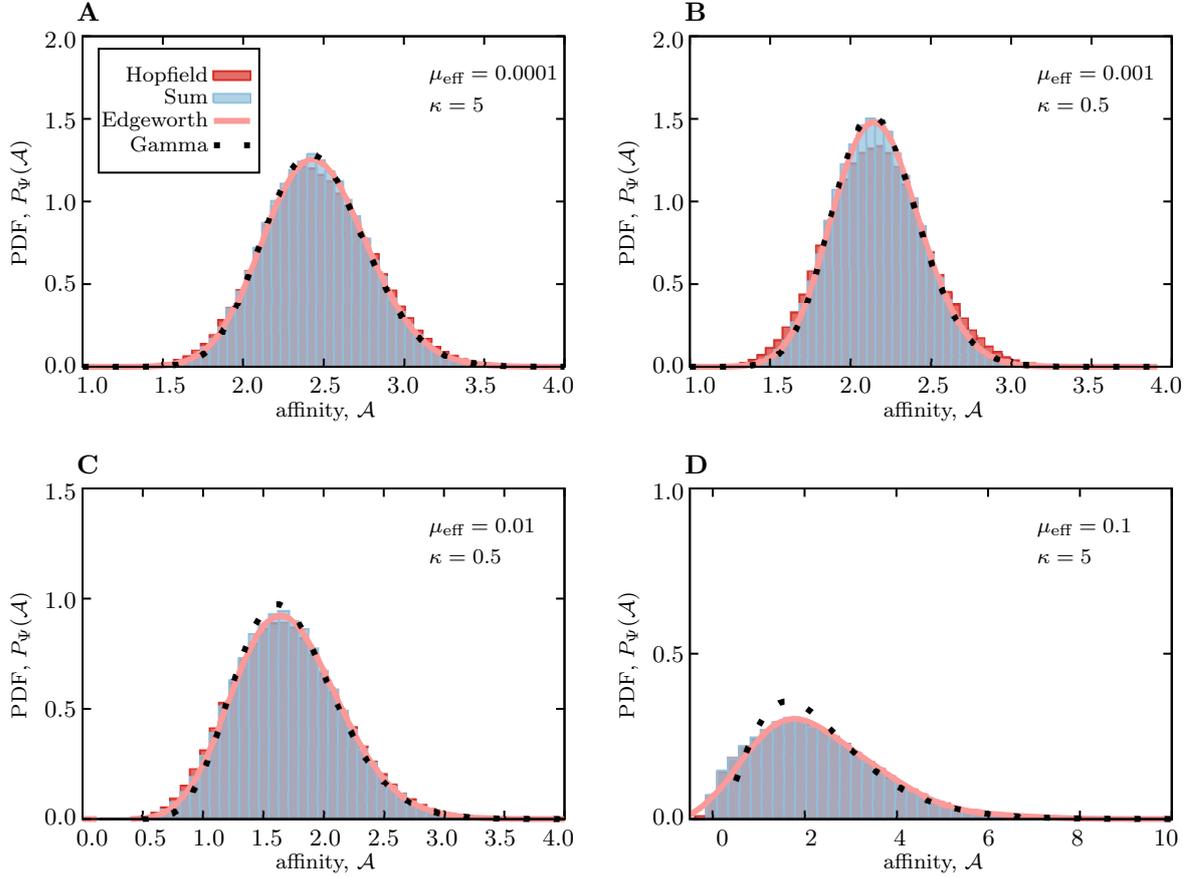


Figure C.3 S3 Distribution of pattern affinities. The distribution of affinities between a memory repertoire and patterns from previously enchanted classes are shown for different effective mutation rates μ_{eff} and risk tolerance κ in different panels. The distributions are characterized based on the simulations using Hopfield network (red; see Appendix C), the process for deriving the cumulant-generating function in Eq S19 (blue), the Edgeworth approximation for the probability density function in Eq. S24 (pink), and the Gamma distribution with the matching mean and variance (dotted lines). For small $(\mu_{\text{eff}} \cdot \kappa)$ in (A-C) all distributions are comparable and as suggested by the Kullback-Leibler divergence in Eq. S28, the Gamma distribution is a good approximation to the underlying distribution of affinities. For larger $(\mu_{\text{eff}} \cdot \kappa)$ in (D), the optimal learning rate becomes large (Eq. 5), and the repertoire only remembers the most recently encountered patterns, resulting in the break down of the analytical approximations. Simulation parameters: $L = 200$, $N = 40$, and $\Theta = 2$.

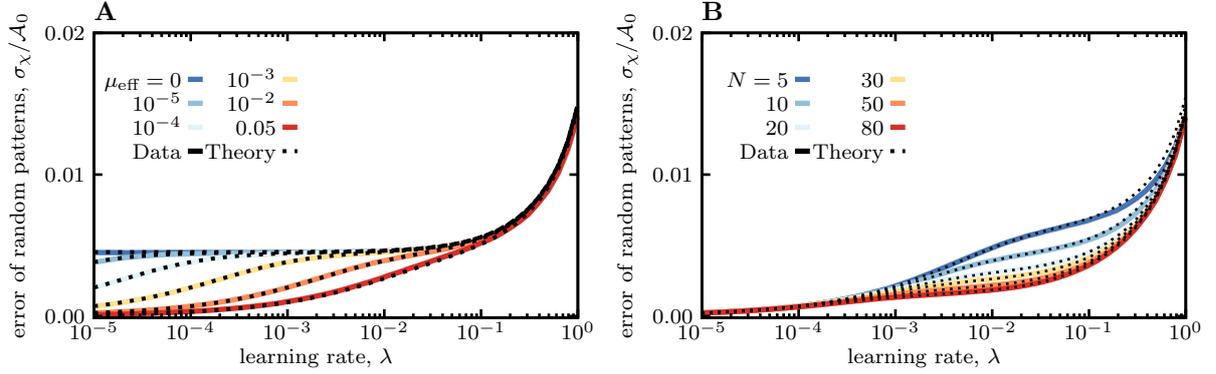


Figure C.4 S4 Standard deviation of affinities for random patterns. Solid lines show simulation results for $\Theta = 2$ and dashed lines give the analytic result from Eq. S42. **(A)** shows results for constant number of pattern classes $N = 30$ and for different effective mutation rates μ_{eff} (colors). **(B)** shows results for constant $\mu_{\text{eff}} = 0.01$ and for different numbers of pattern classes N (colors). In both cases we observe, that when the learning rate λ is large and only one pattern class is stored in the system the fluctuations of random affinities are the largest as there is no self averaging over all stored patterns. When the system can learn close to the optimum (maximal affinity for stored patterns with small fluctuations) the prefactor in Eq. S42 reaches the naïve expectation $1/N$ and the fluctuations of random overlaps are averaged over all stored patterns. When the learning rate become small and the system can no longer follow the evolution of the patterns, the system effectively stores many additional states and the random fluctuations are averaged over all these states. As a consequence, the fluctuations σ_χ go to zero. Simulation parameters: $L = 100$, $\Theta = 2$.

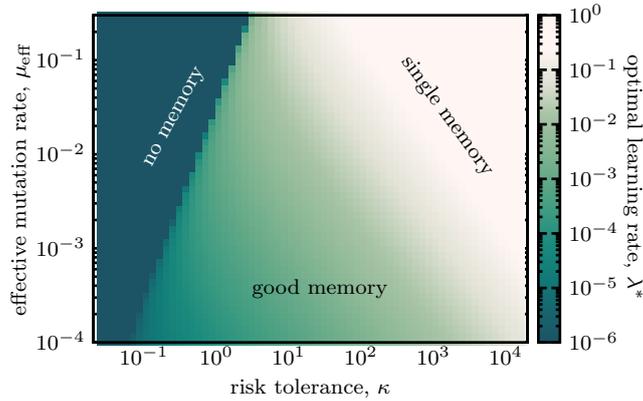


Figure C.5 S5 Optimal learning rate for three phases of memory. The optimal learning rate λ^* corresponding to the different strategies in Fig. 3 is shown for different values of risk tolerance and mutation rate. Fast learning with $\lambda \sim 1$ corresponds to the phase of single memory storage (light), where only the memory of most recent encounter is retained. On the other hand, slow learning $\lambda \ll 1$ corresponds to the phase where effectively no memory is stored (dark). The triangle of good memory is associated with intermediate rates of learning. Simulation parameters: $L = 200$, $N = 40$, and $\Theta = 2$.

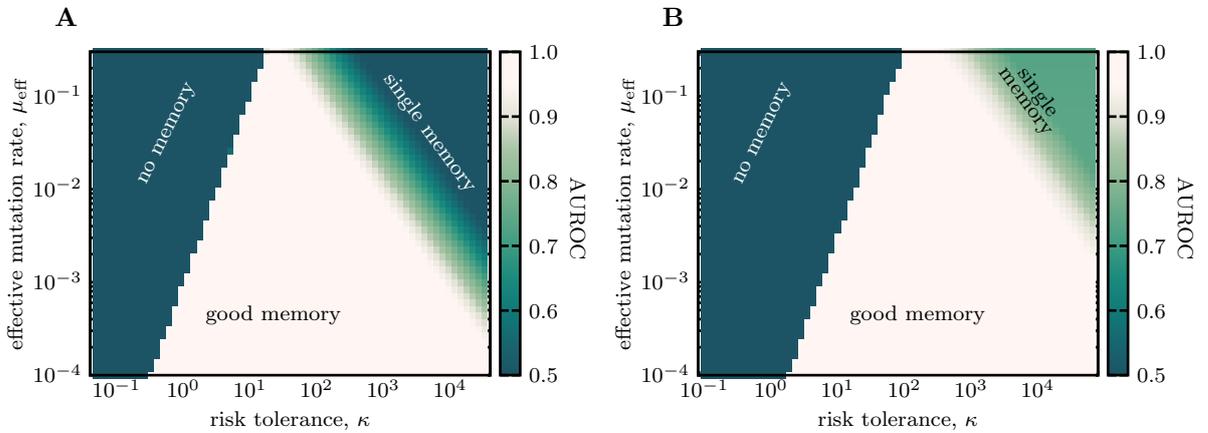


Figure C.6 S6 Three phases of memory discrimination for different shape parameters of the affinity function. Similar to Fig. 3 the phase diagram shows the discrimination ability of repertoires (AUROC) between familiar patterns with prior encounter history and random patterns, for different shape parameters of the affinity function (A) $\Theta = 4$, and (B) $\Theta = 8$. Other parameters: Parameters: $L = 200$, $N = 40$.

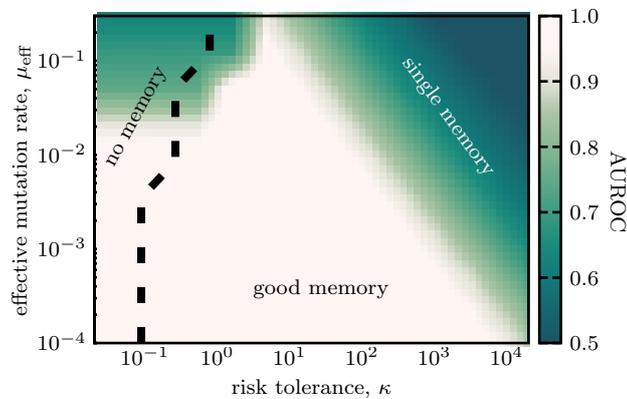


Figure C.7 S7 Three phases of memory discrimination uncovered with simulations of Hopfield networks. Similar to Fig. 3 the phase diagram shows discrimination ability of repertoires (AUROC) between familiar patterns with prior encounter history and random patterns. The phase diagram is acquired by direct simulation of memory, using the correspondence between repertoires with shape parameter $\Theta = 2$ and Hopfield network; see Appendix C for numeral technique. The low risk region on the left side of the dashed line is not accessible by simulations, which explains the differences between the simulations and results of the analytic approximation shown in Fig. 3. Parameters: $L = 200$, $N = 40$, and $\Theta = 2$.

Acknowledgements

This thesis would not have been possible without the help and support of many others.

First of all, I want to thank my supervisor Armita Nourmohammad for giving me the opportunity and funding to pursue my doctoral research in her group. Thank you for suggesting and letting me pursue many different projects while always redirecting me when needed. I am very grateful for all the advice you gave me and especially for supporting all my applications to summer schools, conferences, and grants. It has been a great pleasure to work with you over the last four years, and I hope the feeling is mutual.

Beyond the advice of Armita, all other group members helped me with feedback and suggestions over the past years. In particular, I want to thank Colin, Giulio, Mike, Jakub, and Zach for all the advice I received and the working environment we have created over the past years. Indeed, I will always cherish my time in this group, and I will miss our discussions during lunch and intense foosball matches that kept us from returning to work. I also want to thank Colin for our virtual coffee dates and the Sunday morning workout during the pandemic—both have definitely helped me to stay focused and sane during the past months. A special thanks also needs to go to Mike and Zach for welcoming me during my Fulbright in Seattle. I hope that our friendships will keep us in touch far beyond our academic careers.

In addition to all group members, I also want to thank my collaborator Luca Peliti for all his advice. It was an honor to work with you and I feel very fortunate to have pursued this project with you.

I also want to thank Filippo and Florian for hosting many coffee breaks in their office at the MPIDS. These breaks were always a welcome opportunity to clear the mind from research and instead argue about politics. In addition to this daily distraction, I am very grateful to Florian and Marcel for our Bundesliga/ Champions League ritual, which we could even keep alive over 8.000 km and a nine-hour time difference.

Beyond my scientific family, I want to thank my parents, step-parents, and all my siblings for their support not only during my Ph.D., but for helping me get there in the first place.

Finally, I want to thank my partner Melanie for supporting me during all my ups and downs over the past years and especially for bearing me and keeping me sane during the last months. I love you, and I cannot wait for all the adventures that await us.