BMC
Systems Biology

Autoimmunity and tumor immunology: two facets of a probabilistic immune system

Jaime Iranzo[1,3] and Pablo Villoslada[2*]

Abstract

Background: The immune system of vertebrates has evolved the ability to mount highly elaborate responses to a broad range of pathogen-driven threats. Accordingly, it is quite a challenge to understand how a primitive adaptive immune system that probably lacked much of its present complexity could provide its bearers with significant evolutionary advantage, and therefore, continue to be selected for.

Results: We have developed a very simple model of the immune system that captures the probabilistic communication between its innate and adaptive components. Probabilistic communication arises specifically from the fact that antigen presenting cells collect and present a range of antigens from which the adaptive immune system must (probabilistically) identify its target. Our results show that although some degree of self-reactivity in the immune repertoire is unavoidable, the system is generally able to correctly target pathogens rather than self-antigens. Particular circumstances that impair correct targeting and that may lead to infection-induced autoimmunity can be predicted within this framework. Notably, the probabilistic immune system exhibits the remarkable ability to detect sudden increases in the abundance of rare self-antigens, which represents a first step towards developing anti-tumoral responses.

Conclusion: A simple probabilistic model of the communication between the innate and adaptive immune system provides a robust immune response, including targeting tumors, but at the price of being at risk of developing autoimmunity.

Keywords: Comparative immunology/evolution, Autoimmunity, Tumor immunity, Antigen presentation/processing, Mathematical modeling

Background

For more than a century, the immune system (IS) of vertebrates has been admired by biologists for its versatility and stunning complexity, which continues to impress as new findings are made in the field. The plethora of cell types, molecular interactions and soluble factors that are involved in the immune response is not only responsible for an organisms defense against infections but also, they play an important role in self-maintenance and antitumor responses [1]. The complexity that is currently observed in the IS of vertebrates is the consequence of an evolutionary process that started more than 500 million years ago [2-4]. Therefore, primitive stages of the IS probably lacked many of the current cell types, lymphoid organs

and regulatory mechanisms [5], although they still had to cope with the risk of potential autoimmunity. Understanding how this simpler IS was sufficiently functional to persist and continue evolving under the pressure of natural selection constitutes a challenge for immunologists and evolutionary biologists alike [6]. In a discipline mainly driven by experimental work, simple theoretical models can be used to provide an intuitive understanding of how the IS works, as well as the conceptual frameworks necessary to organize and integrate experimental findings into a coherent system [7]. Moreover, as new forms of IS are being discovered in organisms other than vertebrates, theoretical studies on the very general properties of the vertebrate IS may help us to formulate hypotheses on what we might expect for other such IS's [8].

One of the first conceptual models of the IS was proposed more than fifty years ago and based the immune response firmly on the foundations of self vs. non-self

* Correspondence: pvilloslada@clinic.ub.es
[2]Institute of Biomedical Research August Pi Sunyer (IDIBAPS), Hospital Clinic of Barcelona, Casanova 145, Cellex Center 3A, 08036 Barcelona, Spain
Full list of author information is available at the end of the article

discrimination [9]. According to a two-signal model developed subsequently [10], discrimination between self and non-self could be achieved if two concurrent triggers were required in order to induce a response. While these models emphasized the importance of correct Antigen (Ag) targeting, they could neither explain tolerance towards certain non-self Ags, nor the responses against self Ags. The danger model introduced in the nineties [11], integrated those phenomena by establishing that the IS must be activated by danger signals in order to initiate a response. Conversely, Ags (self or non-self) that are presented in a non-danger context become tolerated. While focusing on the context required for IS activation, the danger model does not directly address the question of specific Ag targeting (i.e. how the IS decides, in a context of danger, which particular Ag should be targeted). A different perspective was adopted in the tunable activation threshold model [12] in which it was argued that the IS actively tracks Ag abundance and responds against Ags whose abundance changes abruptly [13,14]. Situating the detection of sudden changes in Ag as a central element in the immune response has recently been reformulated as the discontinuity theory of immunity [7]. From a more abstract point of view, the ability of the IS to detect and manage changes in the self was the basis of the cognitive IS theory adopted in the late nineties [15].

Since the discovery of autoimmune diseases and anti-tumor immunity [16,17], self-reactivity has become a central issue that all theoretical models of the IS must deal with. One particularly interesting question is whether such phenomena are exceptional (e.g. due to the malfunction of the IS) or if they are simply natural outcomes displayed by any conceivable IS [18]. We have developed a simple, conceptual model that aims to integrate the basic features of the IS into a unified framework. Inspired by evolutionary considerations, we hypothesize that a simplified model of the IS, even if it neglects some degree of its actual complexity, can shed some light on the intrinsic properties of the immune response. Likewise, such a model might provide hints as to the minimal elements required for an IS to be functional.

In order to construct our model, we focused on the communication between the innate and adaptive components of the IS. On the one hand, the innate immune system (IIS) has evolved to recognize conserved molecular patterns [19] that are usually associated to infections or tissue damage – the so-called danger signals. Once such signals are detected, the IIS is activated and it sets in motion an innate response [20]. Conversely, the adaptive immune system (AIS) can recognize virtually everything, although it must be activated by Ag-presenting cells (APCs) that pertain to the IIS. Communication between the IIS and AIS, materialized in the process of Ag presentation, aims to ensure that the adaptive immune response

is targeted to the correct Ag (Figure 1). Accordingly, the key property of IIS-AIS communication must be highlighted. In a context of danger, the target Ag is presented by APCs together with a set of non-target Ags from the environment. Because danger signals need not be physically linked to the target Ag (e.g., they may be endogenous factors produced by damaged tissue) there is no way to identify such Ag in a deterministic way. In other words, the IS must choose its target based on incomplete context information, a decision-making process under uncertainty. Control mechanisms, such as central and peripheral tolerance, are necessary to enhance the probability of targeting the correct Ag and it is this probabilistic nature of IIS-AIS communication that we aim to capture with our model.
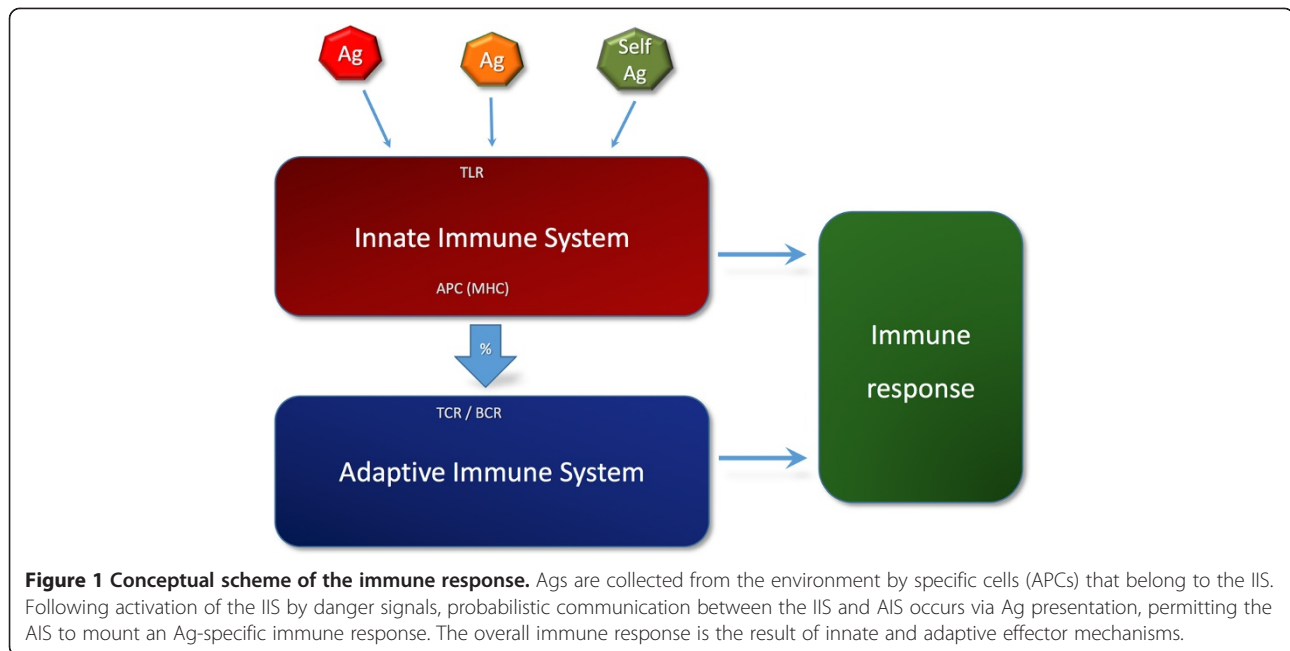
Classical explanations of how infectious agents may provoke autoimmunity [21] fit into the aforementioned conceptual framework at three levels: first, through the incorrect activation of the IIS – bystander activation; second, through a failure in IIS-AIS communication – epitope spreading; and third, through a non-specific adaptive response – molecular mimicry and superantigens. To account for the first two situations, our model allows the role that central and peripheral tolerance play in minimizing the risk of autoimmunity to be evaluated qualitatively. In addition, it shows that some classes of autoimmunity and antitumor responses are related outcomes that arise from the probabilistic nature of IIS-AIS communication.

The rest of the article is structured as follows: we begin Results section by describing the main ingredients of the model, which include an antigenic microenvironment, central tolerance through negative selection, and a simple form of peripheral tolerance. Then we present the results of the mathematical analysis of the model in what concerns self-reactivity, pathogen targeting and tumor detection. When possible, order of magnitude estimations are given to illustrate the performance of the IS in different scenarios. The Discussion highlights the relevance of our results in the context of autoimmunity and tumor immunology, and explores the possibility that the probabilistic model can be extended to non-vertebrate IS's. A comprehensive exposition of the model and its analysis is provided in Methods.

Results
Probabilistic model of the IS
We set out to model IIS-AIS communication as a random process based on the following principles: (i) Antigens, foreign and self, are collected randomly by APCs in a non-selective manner; (ii) APCs sense pathogens or their effects on the environment (e.g. responses provoked by tissue damage) and they are activated; (iii) Ag presentation by active APCs is probabilistic, such that the probability of

Figure 1 Conceptual scheme of the immune response. Ags are collected from the environment by specific cells (APCs) that belong to the IIS. Following activation of the IIS by danger signals, probabilistic communication between the IIS and AIS occurs via Ag presentation, permitting the AIS to mount an Ag-specific immune response. The overall immune response is the result of innate and adaptive effector mechanisms.

presenting an Ag depends on its abundance and the affinity of the MHC: Major Histocompatibility Complex-peptide complex expressed on the surface of the APC. For the sake of simplicity, we refer to the conjunction of both these factors (abundance and affinity) through the term "effective abundance". While inspired by the immune system of vertebrates, our intention was to keep the model sufficiently general so that it could be applied to other hypothetical immune systems with an adaptive component. Therefore, we did not consider further details regarding the specific cell types and molecular interactions involved in the immune response, which would decrease the opportunity to perform specific predictions.

The first step towards building a probabilistic model of the IS lies in defining the antigenic microenvironment in which the IS exists. In a healthy scenario, the antigenic microenvironment can be envisioned as a set of self-Ags, $A = \{A_1, ..., A_i, ..., A_N\}$, each with a particular effective abundance a_i. In analogy to the molar fractions in Chemistry, effective abundances were normalized so that they added up to one. Thus, a_i is the probability of randomly choosing the Ag A_i in the full set A. Pathogens alter the antigenic microenvironment by introducing additional pathogen-associated Ags, $P = \{P_{N+1}, ..., P_M\}$, whose effective abundances equal p. Since the sum of all effective abundances must be equal to one, the presence of a pathogen modifies the effective abundance of the self Ags, which become $a'_i = (1 - p) a_i$. If desired, non-pathogenic, non-self Ags (for instance, those belonging to commensal gut microbiota) could be included in the model as a third set of Ags with their corresponding

abundances, although we did not analyze their instance in this particular model.

The next element in the model was central tolerance. We modeled induction of central tolerance as a stochastic process that mimics thymic negative selection. For each Ag in the set of self-Ags we calculated the probability that a naive, self-reactive cell does not see its cognate Ag during a number t_0 of presentations that are associated to the maturation process (see Methods). This resulted in an expression that links the effective abundance of a self-Ag, a_i, to the probability of it becoming tolerated, f_i, whereby:

$$f_i = 1 - e^{-a_i t_0} \tag{1}$$

In this equation, abundant Ags become tolerated but self-reactivity against rare Ag persists (Figure 2, top). In accordance with intuitive expectations, tolerance increases with t_0, the number of presentations involved in negative selection.

In recent years several new mechanisms associated with the maintenance of peripheral tolerance have been described, involving regulatory T and B cells, tolerogenic dendritic cell presentation, inhibitory and pro-apoptotic co-stimulatory signals, etc. [22]. However, as our aim was to build a model that was as simple as possible, and from an evolutionary perspective, we hypothesized that a much simpler implementation of peripheral tolerance might be sufficient to yield a functional IS. Several mechanisms have been described to carry out this function: on the one hand, minimal stimulation through the TCR: T cell receptor-Ag interaction must be received to activate the lymphocyte [23-25]; on the other, anergic and regulatory
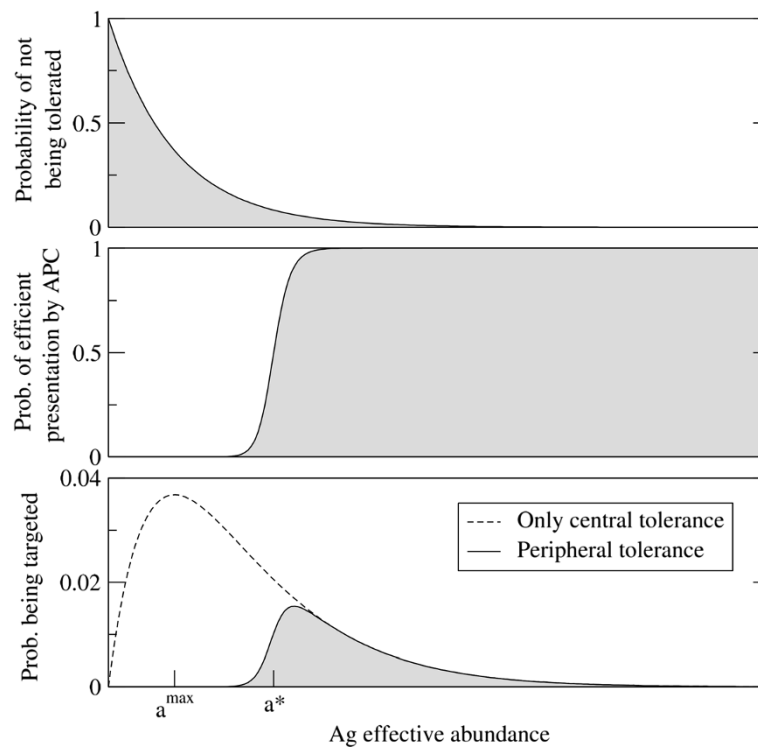
Figure 2 Effect of central and peripheral tolerance on immune self-reactivity. The x-axis represents the effective abundance of a self-Ag, in arbitrary units. Top: probability that self-reactive lymphocytes towards a given self-Ag remain after the induction of central tolerance. The curve corresponds to $1 - f_i$ from eqn. 1 and it can be also interpreted as the fraction of self-Ags with a given effective abundance that remain untolerated. Middle: probability that an Ag (self or non-self) is efficiently presented by an APC. Bottom: combined effect of central and peripheral tolerance on the probability of targeting a self-Ag. In our simple model, peripheral tolerance is determined by the threshold a^*, the minimum effective abundance required for efficient Ag presentation. In the absence of peripheral tolerance, the probability of targeting a self-Ag becomes greatest for those self-antigens with effective abundances close to a^{max}.

T cells [26,27] exert an inhibitory role, preventing the triggering of responses against low concentration Ags. These findings suggest that a new source of regulation should be introduced into the model, namely the need for a minimum number n^* of MHC-peptide complexes at the surface of the APC in order to induce lymphocyte activation. Since the frequency at which a given Ag is presented depends on its effective abundance, this minimum number of MHC-peptide complexes can be translated into an effective abundance threshold a^* (see Methods). Accordingly, only those Ags with effective abundances greater than a^* are presented efficiently and induce an adaptive response, whereas Ags with sub-threshold abundance can still be presented but they do not produce a response (Figure 2, middle). This approach to modeling efficient Ag presentation constitutes a simple yet functional form of (unspecific) peripheral tolerance, because it controls lymphocyte activation avoiding autoimmune responses against rare, non-centrally-tolerated Ags. It should be kept in mind, though, that peripheral tolerance in vertebrates is an active and complex process that fine-tunes the immune response through multiple regulatory interactions.

Remaining self-reactivity in the immune repertoire

The probabilistic model allows the effect that tolerance mechanisms exert on the immune repertoire to be evaluated. While negative selection removes lymphocytes that are highly reactive towards a set of self-Ags, rare self-Ags may not be centrally tolerated if they are not confronted by lymphocytes that are reactive to them. In conditions where the IIS is active (e.g. as a result of an ongoing infection or tissue damage), those self-Ags with an effective abundance above the efficient presentation threshold are susceptible to induce immune responses. The probability that this occurs depends on the abundance of each self-Ag and specifically, it is the product of the effective Ag abundance and the probability of not being tolerated. Like very rare self-Ags, common ones are considered to be at low risk of inducing immune responses, essentially because the former are hardly ever presented and the latter are generally tolerated. Rather, the model predicts a maximum reaction risk for self-Ags whose effective abundances lie around a particular value a^{max} (Figure 2, bottom). Moreover, if central tolerance is complemented with peripheral tolerance (as described in

Methods) the risk can be greatly minimized. The only condition is that the abundance threshold for efficient presentation lies above the abundance of those Ags, i.e. a^* must be greater than a^{max}. In such a case, when self-Ags and self-reactive lymphocytes are present, the APC-lymphocyte interaction is not strong enough to trigger an adaptive immune response. Note that when peripheral tolerance comes into play, the riskiest self-Ags shift from those with effective abundances around a^{max} to those around the presentation threshold a^*, for which the probability of being erroneously targeted is lower.

Correct targeting of pathogens and the risk of bystander autoimmunity

In the context of an infection, a primary goal for the AIS is correctly targeting those Ags that correspond to the pathogen, with the complication of this having to be accomplished in the presence of non-target self-Ag towards which a fraction of the immune repertoire is potentially reactive. Wrong targeting may lead to activation of self-reactive lymphocytes and eventually trigger an autoimmune response, or at least immune responses that consume resources and decrease fitness. In order to evaluate how feasible it is for the IS to achieve this goal, we introduce the correct targeting ratio, R, as the quotient between the probability of targeting an Ag from the pathogen and the probability of selecting one of those self-Ag for which the self-targeting risk is greatest (i.e., those self-Ag with effective abundance equal to a^*). Provided that the IIS has been activated by the pathogen and the effective abundance of the target Ag is greater than the efficient presentation threshold, the correct targeting ratio R becomes (see Methods):

$$R = \frac{pe^{a^*t_0}}{(1-p)a^*} \qquad (2)$$

whereby the ratio R quantifies how much more probable it is for the AIS to target a pathogen Ag rather than a self-Ag. By examining its expression, it becomes evident that the greater the pathogen abundance p, the more easily the AIS can target it. Conversely, low concentrations of pathogens (in the context of an activated IS) are more prone to induce autoimmunity. In the least favorable scenario in terms of Ag targeting, expression (2) can be approximated by $R = e^{Ksn^*}$, where n^* is the number of MHC-peptide complexes required to activate the lymphocyte, K is the number of APCs that a lymphocyte scans during maturation, and s is the fraction of the APC surface that is scanned by the lymphocyte (see Methods). Available experimental data assign n^* a typical value between 5 and 10 [28], while estimates of K range from a few hundred to several thousand [29]. Under the conservative

assumption that lymphocytes only scan as little as a 5% of the APC surface, those parameter values render a correct targeting ratio $R > 10^{20}$, which implies that Ag targeting is extremely reliable. Note that estimates of R become even greater if larger surface scanning fractions are considered.

In the light of our previous calculations, the wrong targeting of self-Ags that would lead to bystander activation is very unlikely, although there are two instances where that might occur with higher probability. First, wrong targeting can occur if the IS faces pathogens capable of activating the IIS while remaining at abundances too low to be detected by the AIS (i.e., below the efficient presentation threshold a^*). Such a scenario trivially corresponds to a null value of the correct targeting ratio, $R = 0$, which means that the Ag targeted by the AIS, if any, will be a self-Ag. The second possibility is that the infection produces changes in the effective abundance of self-Ags, such that Ags that are usually kept out of sight of the AIS suddenly become exposed. The magnitude of this effect depends on the relative changes experienced by self-Ags. For example, if the basal abundance of a self-Ag lies around a^{max} (undetectable) and it raises to a^* (detectable), the probability that it be accidentally targeted by the AIS can be as high as one third of the probability of actually targeting the pathogen. This is equivalent to a correct targeting ratio equal to three, to be compared to 10^{20} if the basal abundance of the self-Ag is already a^* (detectable) but it remains unchanged during the infection. This observation suggests that autoimmunity would be more prone to occur against self-Ags the presence of which increases during certain pathologies. Indeed, this process is very similar to the mechanism through which the IS can recognize tumors, as will be explained below.

Tumor detection

The probabilistic induction of central tolerance implies that some reactivity remains against rare self-Ags. Detecting self-Ags implies an intrinsic risk of autoimmunity but at the same time, it provides the immune system with a valuable asset that if well managed, enables it to monitor its own microenvironment and opens the door to antitumor responses. The fact that the IS can recognize tumor cells by sensing specific tumor Ags is well documented and underlies the concepts of cancer immunosurveillance and immunoediting [1,30]. We hypothesize that immune responses against tumors are based on the targeting of rare, non-tolerated self-Ags whose abundance increases during tumor progression [31,32].

Within the framework of the probabilistic IS model, there are three requisites that must be fulfilled for the IS to detect changes in the abundance of self-Ags: (1) The IS must be able to organize a response against self-Ags; (2) conserved self-reactivity must depend on the abundance of each Ag; and (3) Ag presentation must resemble

Ag abundance. The first condition is an obvious prerequisite, while the second and third conditions are necessary for meaningful Ag detection. This is a tantamount to distinguishing between noisy serendipitous activation (the risk of bystander autoimmunity) and meaningful changes in concentration. As an additional requirement, the IIS must be activated before a response can be mounted.

The fact that there is a threshold for efficient presentation implies that self-Ags with an effective abundance below the threshold must increase in abundance in order to induce immune responses. For example, the probability of detecting self-Ags when some of them have increased in abundance several fold can be seen in Figure 3. In basal conditions only those Ags with an abundance greater than a^* (red line) can induce an auto-reactive response (natural autoimmunity), of which there are only a few and their probability of being detected is very low. In pathological conditions, changes in abundance make presentation more likely for some Ags the basal abundance of which is otherwise below a^* (red points). In this way the IS can detect changes in self-Ags.

The ability of the IS to detect self-Ags that increase their abundance can be characterized through two parameters: the minimum detectable increment (MDI) and the discrimination ability (Δ). We introduced the MDI to address the question of how sensitive the IS is towards changes in self-Ags. Specifically, the MDI quantifies the magnitude of
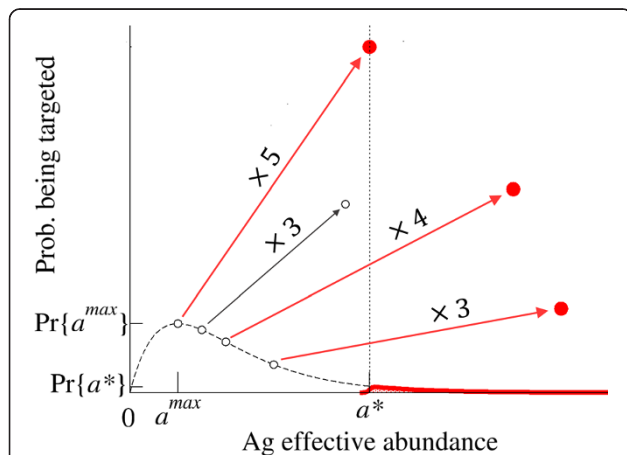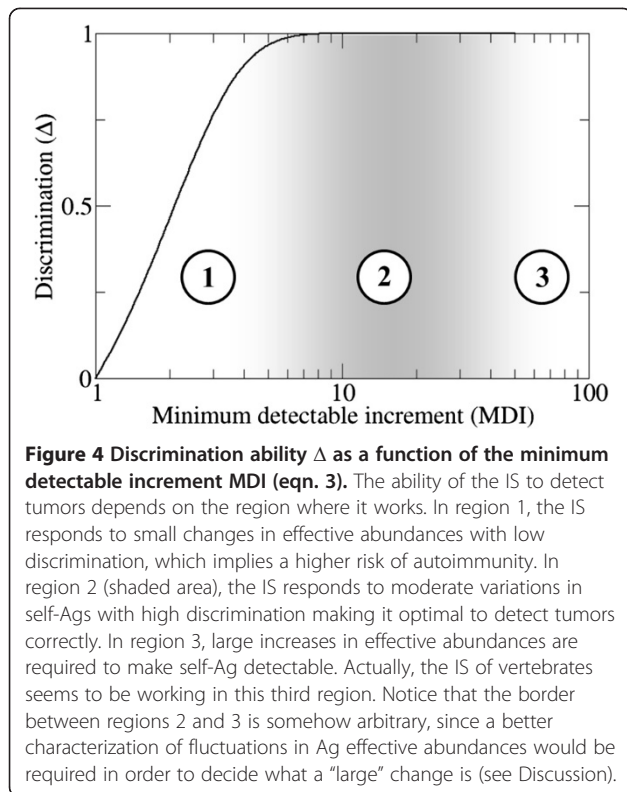

Figure 3 Probability of targeting self-Ag in basal conditions (curve) and after pathological over-expression (isolated points). The arrows indicate multiplicative changes in the effective abundances of several self-Ag (increase factors are written on the arrows). Only those Ags with effective abundances greater than the threshold a^* can induce a response (red symbols). Note that, due to peripheral tolerance, the targeting probability for Ag with abundances below a^* is effectively zero; the values represented in black (dashed line and open symbols) correspond to the targeting probabilities in the absence of peripheral tolerance as given by eqn. (8b). The MDI in this example has been set to 5, which means that an Ag with basal abundance equal to a^{max} has to be over-expressed by a factor 5 in order to reach the threshold abundance a^*.

the changes that the IS can detect. According to our model, the IS is not equally sensitive to changes in all self-Ags but rather, to those Ags whose basal effective abundances lie around a^{max}. Thus, the MDI is defined as the factor by which the abundance of such Ags must increase in order to reach the detection threshold a^*, whereby $MDI = a^*/a^{max}$ (see Figure 3). By using the previous result $a^{max} = 1/t_0$ it is possible to express the MDI in the alternative way $MDI = t_0 a^*$.

The discrimination ability assesses the extent to which self-recognition can be correctly interpreted as a result of a change in the self. Provided that the IIS has been previously activated by other means, detection of self-Ags by random presentation is indistinguishable from that derived from genuine changes in effective abundance. In relation to self-recognition, it is possible to solve this conundrum if it can be assumed to be more probable that a true change in concentration, rather than accidental noise-driven detection, has occurred. We define the discrimination ability (Δ) as the certainty with which detection of a self-Ag can be attributed to a change in its effective abundance. The discrimination ability adopts a value between zero (when there is no certainty that detection is due to a change in abundance) and one (when the change in abundance is certain). In order to calculate it, we use the same Ag for which the MDI was defined, such that:

$$\Delta = \frac{\Pr\{a^{max}(\times MDI)\} - \Pr\{a^*\}}{\Pr\{a^{max}(\times MDI)\} + \Pr\{a^*\}} \qquad (3)$$

As shown in Methods, it is possible to derive a relation that links both the MDI and the discrimination ability:

$$\Delta = \frac{e^{MDI-1} - 1}{e^{MDI-1} + 1} \qquad (4)$$

The dependence between Δ and the MDI means that the IS discriminates strongly if it only responds to big changes in the self. When trying to detect small changes, the IS has to deal with a weaker capacity to distinguish between changes and random presentations. This trade-off between the MDI and discrimination can be represented (Figure 4) and the most remarkable feature is that the IS can achieve almost full discrimination by establishing its MDI just over a factor 10 (notice that this holds regardless of all other parameters of the model). An order of magnitude estimation of the actual MDI in mammals can be carried out by recalling that $a^* t_0 = Ksn^*$ (see Methods) and using for K, s and n^* the same conservative estimates as in the previous section. This renders an MDI of the order of 100 (in a range from 50 to 500), which corresponds to a Δ virtually equal to one (the difference is of the order of 10^{-20} for MDI =50 and even smaller for larger MDI). Thus, it

Figure 4 Discrimination ability Δ as a function of the minimum detectable increment MDI (eqn. 3). The ability of the IS to detect tumors depends on the region where it works. In region 1, the IS responds to small changes in effective abundances with low discrimination, which implies a higher risk of autoimmunity. In region 2 (shaded area), the IS responds to moderate variations in self-Ags with high discrimination making it optimal to detect tumors correctly. In region 3, large increases in effective abundances are required to make self-Ag detectable. Actually, the IS of vertebrates seems to be working in this third region. Notice that the border between regions 2 and 3 is somehow arbitrary, since a better characterization of fluctuations in Ag effective abundances would be required in order to decide what a "large" change is (see Discussion).

seems that the IS of (at least) mammals works in a safety zone, where discrimination is almost perfect and detection of self-Ags requires that their effective abundances increase in a large factor.

Discussion

The strength of an adaptive immune system lies in its ability to recognize a virtually infinite number of antigens in a highly specific manner. However, such specific recognition becomes useless, or even dangerous, if the targeted Ag is not related to the cause of IS activation. Consequently, targeting the correct Ag in a universe of self and non-self Ags, damaging and non-damaging Ags, is a fundamental task for the IS. In order to gain an intuitive understanding as to how an active IS accomplishes this task, we have explored a simple model that focuses on how the innate and adaptive components of the IS communicate with each other. The main feature of this model is that communication occurs in a probabilistic way, based on the fact that the IIS does not exactly tell the AIS which the target Ag is but rather, it presents it along with other, non-target Ags. In order to provide the AIS with the ability to discriminate Ags, our model only considers two control mechanisms: first, central tolerance, where induction is modeled as probabilistic self-Ag presentation leading to inactivation; and second, the need of a minimum number of MHC-peptide-TCR complexes to be formed in order to induce lymphocyte

activation (the effective presentation threshold). Provided that the IIS only becomes activated when signals of danger are present, these simple control mechanisms are sufficient to ensure reliable targeting of the pathogen Ag by the AIS. Such a finding is striking and involves evolutionary implications that are not trivial, particularly as it shows that a very simplified IS can still be functional. It is conceivable that the addition of further complexity during evolution (e.g. regulatory feedback loops, memory, specific peripheral tolerance) will have optimized the dynamics of immune responses and enlarged their functional spectrum.

Conceptual models of the IS usually assume that every Ag present will be detected by the AIS. In the probabilistic model, we explicitly consider the possibility that rare Ags remain undetected and that if any of such Ag experiences a sudden increase in abundance it will then become visible to the IS. In the presence of activation −danger− signals, such events may lead to autoimmune or antitumor responses, the latter if a tumor Ag is involved [31,33]. Experimental evidence supporting this prediction includes detection of autoantibodies and self-reactive T-cells specific to certain types of cancer, which has been related to the overexpression of their cognate tumor-associated Ags [30,32,34,35]. It should be noted that in this model, detecting changes in Ag abundance is not based on any kind of active tracking but rather on probabilistic issues. Thus, autoimmunity and responses to tumors are two sides of the same coin: revolving around the probabilistic nature of the IIS-AIS communication. This intrinsic relation that links autoimmunity and antitumor immunity manifests itself, for instance, in similarities between the Ag involved in both kinds of responses [36,37]. In the context of immunosenescence, the model predicts that depletion of the self-reactive immune repertoire with aging would impair the IS from detecting tumors.

An order of magnitude estimation of the sensitivity of the IS to changes in self-Ag suggests that a tumor Ag's effective abundance should increase approximately in a factor 100 in order to become detectable. While such a large change may not always be experienced by serum markers typically used for monitoring cancer [38,39], instances of cell markers whose expression in tumors raises several orders of magnitude have been reported [40]. From a systems optimization perspective, the 100-fold increase in effective abundance required for detection seems surprisingly conservative, since almost certain discrimination is compatible with detecting changes of as a small magnitude as 10-fold. A possible explanation is that ordinary fluctuations in Ag effective abundances often raise over a factor of 10, e.g. due to local changes in concentration or episodes during which they become more accessible to APCs. Therefore, the high MDI observed would contribute to filtering significant changes.

Testing this idea would require a better characterization of the MHC-bound peptide profiles and their fluctuations, an information that is just starting to become available [41,42].

According to our results, the very existence of self-reactive lymphocytes can be at least partially explained by the probabilistic nature of Ag presentation during lymphocyte maturation. Complete tolerance cannot be achieved since the period of maturation is limited. This constraint on central tolerance would still hold even if all self-Ags were expressed in the thymus and therefore, incomplete tolerance and residual self-reactivity are not faults but intrinsic properties of the vertebrate IS. In our model the mechanism of peripheral tolerance is very basic, static and unspecific, and could be considered a sort of 'peripheral ignorance'. However, it is easy to envisage what would happen with a more sophisticated mechanism able to induce tolerance to those Ags that are efficiently presented in the absence of activation signals. In such circumstances, increases in the abundance of rare Ags would switch them from untolerated to tolerated, and such specific peripheral tolerance would have dramatic consequences in terms of tumor detection. Indeed, the lack of activating signals may allow tumors to escape the IS and induce tolerance to the over-expressed tumor Ags [43]. This represents an instance of how a mechanism evolved to reduce the risk of autoimmunity entails a cost at the level of tumor detection.

A connection between genes involved in Ag presentation and susceptibility to autoimmune diseases has already been suggested on the basis of the results of genome wide association studies [44,45]. The probabilistic model of the IS in turn captures some of the mechanisms through which infections can trigger autoimmune responses [21], identifying two risk factors of interest. First, the risk of an autoimmune response increases in the presence of pathogens that activate the IIS while remaining at concentrations too low to be effectively presented to the AIS. The second risk factor is the massive release of otherwise rare self-Ags as a result of infection. In both cases, autoimmunity would follow from incorrect Ag targeting and its onset might possibly require additional failings of other regulatory mechanisms not included in our model [46]. An experimental setting that simulates the aforementioned scenario is the induction of experimental autoimmune encephalomyelitis, an animal model of multiple sclerosis, after priming with central nervous system-restricted Ags plus an adjuvant [47]. Alternatively, autoimmunity can follow massive Ag release after traumatic lesions provided that they lead to immune activation. This may be the case of sympathetic ophthalmia, a rare autoimmune pathology that affects both eyes after accidental or surgical trauma in one of the eyes [48].

From the point of view of comparative immunology, the results we have obtained with this simple model provide hints as to what might be expected when studying different immune systems. Incomplete tolerance, some risk of autoimmunity and the ability to detect tumors are intrinsic properties of immune systems whose innate and adaptive components communicate probabilistically. Nevertheless, it is worth noting that there are alternatives to probabilistic communication. For instance, the intracellular RNA interference system of plants detects and processes viral double-stranded RNA to produce target-specific effector molecules (siRNA) [49]. From a conceptual perspective, the RNA interference system involves innate –dsRNA recognition– and adaptive –effector siRNA– components that communicate deterministically through RNA processing. Alternatively, immune systems with no clear distinction between innate and adaptive components may also behave as a probabilistic IS if: (1) their highly diversified receptors are potentially capable of binding self-molecules; and (2) they respond to activation (danger) signals that are not physically linked to their "target Ags". As research on the immunology of jawless vertebrates, protochordates and invertebrates progresses [2,8,50,51], it will be of great interest to see whether their highly diversified immune systems are based on, and share, the characteristics of those with probabilistic communication strategies [18]. If that is the case, we expect that the general conclusions obtained in this work will also hold for them.

Conclusion

A simple probabilistic model of the communication between the innate and adaptive immune system provides a robust immune response, including targeting tumors, but at the price of being at risk of developing autoimmunity.

Methods

Effective abundance

As explained above, our model deals with an IS immersed in a complex antigenic universe. Each Ag A_i is characterized by its effective abundance a_i, which is a normalized quantity accounting for the probability that A_i is presented by APCs. Ideally, the effective abundance would be the product of the actual abundance, the degree of access that APCs have to the Ag, and the affinity of the MHC-peptide complex, normalized in such a way that the sum of effective abundances for all Ags equals one. According to that, the effective abundance increases not only if the actual abundance does, but also if the Ag becomes more accessible to the APCs (e.g. after infiltration of APCs in a tissue), or if a particular MHC allele has greater affinity for that Ag.

Self-reactivity of the immune repertoire after thymic negative selection

Central tolerance was modeled as a stochastic process that mimics thymic negative selection on the immune repertoire. To that end, we first calculated the probability that a self-reactive lymphocyte finds its cognate Ag during maturation and experiences negative selection. The maturation process itself is modeled as a number t_0 of Ag presentations, where the probability of antigen A_i being presented is proportional to its effective abundance a_i. For this process, the probability that A_i is presented k times to the self-reactive lymphocyte follows a binomial distribution:

$$\Pr\{A_i \text{ presented } k \text{ times}\} = \binom{t_o}{k} a_i^k (1-a_i)^{t_o-k} \qquad (5)$$

Provided that the number of Ag presentations during lymphocyte maturation is large, the binomial distribution can be approximated by a Poisson distribution. Therefore, the previous probability becomes:

$$\Pr\{A_i \text{ presented } k \text{ times}\} \approx \frac{(a_i t_o)^k}{k!} e^{-a_i t_o} \qquad (6)$$

The probability that a self-reactive lymphocyte escapes negative selection can be obtained by making $k = 0$. If there are m naive lymphocytes reactive to the same self-Ag, all of them must experience negative selection so that the self-Ag becomes tolerated. That would happen with probability f_i, whereby:

$$f_i = \Pr\{A_i \text{ tolerated}\} = (1 - e^{-a_i t_o})^m \qquad (7)$$

which, for $m = 1$, is the expression reported in the Results section as eqn. (1).

In the context of an ongoing infection, there is a possibility that self-antigens which are not tolerated be erroneously targeted by the IS. For a given self-Ag, the probability of being targeted is proportional to its effective abundance and conditional to not being tolerated. In mathematical terms, that can be expressed as the product of the effective abundance and the probability of not being tolerated. If there are m self-reactive clones per each self-Ag the targeting probability becomes:

$$\Pr\{A_i \text{ targeted}\} = 1 - (1 - a_i (1 - f_i))^m \qquad (8a)$$

In the simplest case with $m = 1$, equation (8a) takes a more compact form:

$$\Pr\{A_i \text{ targeted}\} = a_i e^{-a_i t_0} \qquad (8b)$$

The probability of being targeted is maximum for self-Ag with effective abundance equal to a^{max}. The exact expression for a^{max} can be obtained by taking the derivative of equation (8) with respect to a_i and making it equal to zero. This results in $a^{max} = 1/t_0 m$.

Effective abundance threshold for efficient Ag presentation

In order to incorporate a simple form of peripheral tolerance in our model, we introduce the following premise: lymphocyte activation by APCs requires the formation of a minimum number n^* of MHC-peptide-TCR complexes in the surface of the APC. In other words, only those antigens that appear in (at least) n^* MHC-peptide complexes on the same APC will be able to induce a response. For an antigen A_i with effective abundance a_i, the probability that this happens is equal to:

$$\Pr\{A_i \text{ represented at least } n^* \text{ times}\} = 1 - e^{-a_i N} \sum_{j=0}^{n^*-1} \frac{(a_i N^j)}{j!}$$

$$(9)$$

where N is the total number of MHC-peptide complexes on the surface of the APC. Provided that N is moderately large, this expression grows rapidly from zero to one when the effective abundance reaches a value $a^* = n^*/N$. An example of this behavior is shown in Figure 2 (middle panel). The effective abundance a^* can be interpreted as a threshold for efficient Ag presentation.

Once peripheral tolerance is taken into account, the probability of targeting a given self-Ag can be calculated by multiplying the expressions in equations (8a) and (9). This is how the solid curve in the lower panel of Figure 2 was obtained. However, for practical purposes equation (9) can be approximated by a Heaviside step function centered in a^*. Under that approximation, the probability of targeting a self-Ag becomes:

$$\Pr\{A_i \text{ targeted}\} = \begin{cases} 0 & \text{if } a_i < a^* \\ 1 - (1 - a_i (1 - f_i))^m & \text{if } a_i \geq a^* \end{cases}$$

$$(10)$$

Notice that, provided that $a^* > a^{max}$, the probability of being targeted is maximum for those self-Ag with effective abundance equal to a^*.

Pathogen vs. self-Ag targeting

We introduce the correct targeting ratio, R, as a tool to evaluate the performance of the IS when it comes to targeting pathogen Ag. R is defined as the quotient between the probabilities of targeting an antigen from the pathogen and a self-Ag with effective abundance a^* (notice that those are the self-Ags for which the self-targeting probability is greatest). In order to compute R it must be recalled that the presence of a pathogen modifies the effective abundances of self-Ag, multiplying them by a factor $(1 - p)$, were p is the effective abundance of pathogen Ag (this comes from the fact that effective

abundances are normalized). As we work with normalized effective abundances, the probability of targeting an antigen from the pathogen is simply equal to the effective abundance p. Moreover, we assume that $p \geq a^*$, so that the pathogen is efficiently presented (otherwise, the AIS would be unable to "see" the pathogen). Taking all that into account, the correct targeting ratio becomes:

$$R = \frac{\Pr\{\text{Pathogen targeted}\}}{\Pr\{A^* \text{ targeted}\}}$$
$$= \frac{p}{(1-p)\left[1-(1-a^*(1-f^*))^m\right]} \approx \frac{p\, e^{a^* t_0}}{(1-p)\, a^*\, m^2} \quad (11)$$

The last approximation has been obtained by discarding terms of lower order in $\left(e^{a^* t_0}\right)$. Biological data allow to estimate the order of magnitude of this factor around 10^{20} (see Results) which supports the validity of the truncation. In the particular case $m = 1$ equation (2) reported in the Results section is exactly recovered.

A lower bound for R can be calculated if the concentration of pathogen is as low as the minimum detectable, $p = a^*$. In this least favorable scenario, and provided that most Ag effective abundances (a^* among them) are much smaller than one, the correct targeting ratio can be approximated by $R \approx e^{a_i t_0}$. Let us denote by K the number of APCs that a lymphocyte scans during maturation and s the fraction of the APC surface scanned by the lymphocyte. Then it is possible to express the number of presentations during maturation as $t_0 = KNs$, which combined with $a^* = n^*/N$ renders $R \approx e^{Ksn^*}$.

Detection of tumors by the IS

Because the probability of targeting an Ag depends on its effective abundance, changes in Ag abundances modify targeting probabilities. Specifically, let us suppose that the effective abundance of antigen A_i experiences a k-fold increase. Then the targeting probability for that Ag increases from the value in equation (10) to the new value:

$$\Pr\{A_i\,(\times k)\text{ targeted}\} = \begin{cases} 0 & \text{if } ka_i < a^* \\ 1-(1-ka_i\,(1-f_i))^m & \text{if } ka_i \geq a^* \end{cases} \quad (12)$$

Two observations can be made from this expression. First, the IS is not equally sensitive to changes in all self-Ag: by taking the derivative with respect to k and finding its maximum it comes out that the IS is most sensitive to changes in self-Ag whose basal effective abundances are equal to $a^{max} = 1/t_0$ (notice that, for $m = 1$, that coincides with the Ag for which the targeting probability in the absence of peripheral tolerance is greatest). The second observation is that such Ag will be only detected if its effective abundance becomes greater than a^*, i.e. if

$k \geq a^*/a^{max}$. The value of k such that $ka^{max} = a^*$ defines a characteristic factor that we termed Minimum Detectable Increment (MDI). By using previous results, it is possible to express it as MDI = $a^* t_0$.

The discrimination ability (Δ) is introduced with the aim of evaluating the performance of the IS when trying to detect tumors. It quantifies the certainty with which detection of a self-Ag can be attributed to a real increase in its effective abundance and not just to chance. In order to calculate Δ we compare the probability of targeting an Ag with basal abundance a^{max} that has experienced an increase of the magnitude of the MDI with the probability of targeting an Ag with basal abundance a^* (the most likely to be targeted in the absence of changes). Let us denote those probabilities $\Pr\{a^{max}(\times\text{MDI})\}$ and $\Pr\{a^*\}$, respectively.

$$\Pr\{a^{max}(\times\text{MDI})\} = 1-\left[1-a^*\left(1-(1-e^{-1})^m\right)\right]^m$$
$$\Pr\{a^*\} = 1-\left[1-a^*\left(1-\left(1-e^{-a^* t_0}\right)^m\right)\right]^m$$

The discrimination ability is defined as:

$$\Delta = \frac{\Pr\{a^{max}(\times\text{MDI})\} - \Pr\{a^*\}}{\Pr\{a^{max}(\times\text{MDI})\} + \Pr\{a^*\}} \quad (3\text{bis})$$

and after some manipulations it becomes

$$\Delta = \frac{\delta-1}{\delta+1}. \quad (13)$$

Where

$$\delta \approx \frac{1-(1-e^{-1})^m}{1-(1-e^{-\text{MDI}})^m}$$

The approximation has been done by discarding terms of order $(a^*)^2$ and higher. In the particular case $m = 1$ this expression is exact and can be expressed as

$$\Delta = \frac{e^{\text{MDI}-1} - 1}{e^{\text{MDI}-1} + 1}. \quad (4\text{bis})$$

Notice that, because of its definition, Δ takes values between zero (detecting "noise" is as probable as detecting changes) and one (detection is always due to meaningful changes). The relation between Δ and the MDI does not depend on any of the parameters of the model (a_*, t_0).

Author details
[1]Centro de Astrobiología, INTA – CSIC, Madrid, Spain. [2]Institute of Biomedical Research August Pi Sunyer (IDIBAPS), Hospital Clinic of Barcelona, Casanova 145, Cellex Center 3A, 08036 Barcelona, Spain. [3]Current address: National Center for Biotechnology Information, National Library of Medicine, National Institutes of Health, Bethesda, MD, USA.

References
1. Dunn GP, Old LJ, Schreiber RD: **The immunobiology of cancer immunosurveillance and immunoediting.** *Immunity* 2004, **21**(2):137–148.
2. Litman GW, Cannon JP, Dishaw LJ: **Reconstructing immune phylogeny: new perspectives.** *Nat Rev Immunol* 2005, **5**(11):866–879.
3. Cooper MD, Alder MN: **The evolution of adaptive immune systems.** *Cell* 2006, **124**(4):815–822.
4. Boehm T: **Evolution of vertebrate immunity.** *Curr Biol* 2012, **22**(17):R722–732.
5. Boehm T, Bleul CC: **The evolutionary history of lymphoid organs.** *Nat Immunol* 2007, **8**(2):131–135.
6. Hedrick SM: **The acquired immune system: a vantage from beneath.** *Immunity* 2004, **21**(5):607–615.
7. Pradeu T, Jaeger S, Vivier E: **The speed of change: towards a discontinuity theory of immunity?** *Nat Rev Immunol* 2013, **13**(10):764–769.
8. Rimer J, Cohen IR, Friedman N: **Do all creatures possess an acquired immune system of some sort?** *Bioessays* 2014, **36**(3):273–281.
9. Burnett FM: *The clonal selection theory of acquired immunity.* Cambridge, England: Cambridge University Press; 1959.
10. Bretscher P, Cohn M: **A theory of self-nonself discrimination.** *Science* 1970, **169**(3950):1042–1049.
11. Matzinger P: **Tolerance, danger, and the extended family.** *Annu Rev Immunol* 1994, **12**:991–1045.
12. Grossman Z, Paul WE: **Adaptive cellular interactions in the immune system: the tunable activation threshold and the significance of subthreshold responses.** *Proc Natl Acad Sci U S A* 1992, **89**(21):10365–10369.
13. van den Berg HA, Rand DA: **Quantitative theories of T-cell responsiveness.** *Immunol Rev* 2007, **216**:81–92.
14. van den Berg HA, Rand DA: **Dynamics of T cell activation threshold tuning.** *J Theor Biol* 2004, **228**(3):397–416.
15. Atlan H, Cohen IR: **Immune information, self-organization and meaning.** *Int Immunol* 1998, **10**(6):711–717.
16. Smith HR, Steinberg AD: **Autoimmunity–a perspective.** *Annu Rev Immunol* 1983, **1**:175–210.
17. Dunn GP, Bruce AT, Ikeda H, Old LJ, Schreiber RD: **Cancer immunoediting: from immunosurveillance to tumor escape.** *Nat Immunol* 2002, **3**(11):991–998.
18. Bailey M, Christoforidou Z, Lewis M: **Evolution of immune systems: specificity and autoreactivity.** *Autoimmun Rev* 2013, **12**(6):643–647.
19. Janeway CA Jr: **Approaching the asymptote? Evolution and revolution in immunology.** *Cold Spring Harb Symp Quant Biol* 1989, **54**(Pt 1):1–13.
20. Medzhitov R, Janeway CA Jr: **Innate immunity: the virtues of a nonclonal system of recognition.** *Cell* 1997, **91**(3):295–298.
21. Wucherpfennig KW: **Mechanisms for the induction of autoimmunity by infectious agents.** *J Clin Invest* 2001, **108**(8):1097–1104.
22. Mueller DL: **Mechanisms maintaining peripheral tolerance.** *Nat Immunol* 2010, **11**(1):21–27.
23. Smith KA: **The quantal theory of immunity.** *Cell Res* 2006, **16**(1):11–19.
24. Carreno LJ, Riquelme EM, Gonzalez PA, Espagnolle N, Riedel CA, Valitutti S, Kalergis AM: **T-cell antagonism by short half-life pMHC ligands can be mediated by an efficient trapping of T-cell polarization toward the APC.** *Proc Natl Acad Sci U S A* 2010, **107**(1):210–215.
25. Huppa JB, Davis MM: **T-cell-antigen recognition and the immunological synapse.** *Nat Rev Immunol* 2003, **3**(12):973–983.
26. Vendetti S, Chai JG, Dyson J, Simpson E, Lombardi G, Lechler R: **Anergic T cells inhibit the antigen-presenting function of dendritic cells.** *J Immunol* 2000, **165**(3):1175–1181.
27. Carneiro J, Leon K, Caramalho I, van den Dool C, Gardner R, Oliveira V, Bergman ML, Sepúlveda N, Paixão T, Faro J, Demengeot J: **When three is not a crowd: a crossregulation model of the dynamics and repertoire selection of regulatory CD4+ T cells.** *Immunol Rev* 2007, **216**:48–68.
28. Irvine DJ, Purbhoo MA, Krogsgaard M, Davis MM: **Direct observation of ligand recognition by T cells.** *Nature* 2002, **419**:845–849.
29. Klein L, Kyewski B, Allen PM, Hogquist KA: **Positive and negative selection of the T cell repertoire: waht thymocytes see (and don't see).** *Nat Rev Immunol* 2014, **14**:377–391.
30. Ramirez-Montagut T, Turk MJ, Wolchok JD, Guevara-Patino JA, Houghton AN: **Immunity to melanoma: unraveling the relation of tumor immunity and autoimmunity.** *Oncogene* 2003, **22**(20):3180–3187.
31. Tan EM, Zhang J: **Autoantibodies to tumor-associated antigens: reporters from the immune system.** *Immunol Rev* 2008, **222**:328–340.
32. Coulie PG, Van den Eynde BJ, van der Bruggen P, Boon T: **Tumour antigens recognized by T lymphocytes: at the core of cancer immunotherapy.** *Nat Rev Cancer* 2014, **14**:135–146.
33. Lleo A, Invernizzi P, Gao B, Podda M, Gershwin ME: **Definition of human autoimmunity–autoantibodies versus autoimmune disease.** *Autoimmun Rev* 2010, **9**(5):A259–266.
34. Tan EM: **Autoantibodies as reporters identifying aberrant cellular mechanisms in tumorigenesis.** *J Clin Invest* 2001, **108**:1411–1415.
35. Lacombe J, Mangé A, Solassol J: **Use of autoantibodies to detect the onset of breast cancer.** *J Immunol Res* 2014. in press.
36. Tan EM, Shi FD: **Relative paradigms between autoantibodies in lupus and autoantibodies in cancer.** *Clin Exp Immunol* 2003, **134**(2):169–177.
37. Bei R, Masuelli L, Palumbo C, Modesti M, Modesti A: **A common repertoire of autoantibodies is shared by cancer and autoimmune disease patients: inflammation in their induction and impact on tumor growth.** *Cancer Lett* 2009, **281**(1):8–23.
38. Tonack S, Jenkinson C, Cox T, Elliott V, Jenkins RE, Kitteringham NR, Greenhalf W, Shaw V, Michalski CW, Friess H, Neoptolemos JP, Costello E: **iTRAQ reveals candidate pancreatic cancer serum biomarkers: influence of obstructive jaundice on their performance.** *Br J Cancer* 2013, **108**(9):1846–1853.
39. Barlow LJ, Badalato GM, McKiernan JM: **Serum tumor markers in the evaluation of male germ cell tumors.** *Nat Rev Urol* 2010, **7**(11):610–617.
40. Norgard O, Singh G, Solberg S, Jorgensen L, Halvorsen AR, Smaaland R, Brustugun OT, Helland A: **Novel molecular tumor cell markers in regional lymph nodes and blood samples from patients undergoing surgery for non-small cell lung cancer.** *Plos One* 2013, **8**:e62153.
41. Croft NP, Smith SA, Wong YC, Tan CT, Dudek NL, Flesch IE, Lin LC, Tscharke DC, Purcell AW: **Kinetics of antigen expression and epitope presentation during virus infection.** *PLoS Pathog* 2013, **9**(1):e1003129.
42. Hassan C, Kester MG, de Ru AH, Hombrink P, Drijfhout JW, Nijveen H, Leunissen JA, Heemskerk MH, Falkenburg JH, van Veelen PA: **The human leukocyte antigen-presented ligandome of B lymphocytes.** *Mol Cell Proteomics* 2013, **12**(7):1829–1843.
43. Vesely MD, Kershaw MH, Schreiber RD, Smyth MJ: **Natural innate and adaptive immunity to cancer.** *Annu Rev Immunol* 2011, **29**:235–271.
44. Lettre G, Rioux JD: **Autoimmune diseases: insights from genome-wide association studies.** *Hum Mol Genet* 2008, **17**(R2):R116–121.
45. Zenewicz LA, Abraham C, Flavell RA, Cho JH: **Unraveling the genetics of autoimmunity.** *Cell* 2010, **140**(6):791–797.
46. Goodnow CC: **Multistep pathogenesis of autoimmune disease.** *Cell* 2007, **130**(1):25–35.
47. Baxter AG: **The origin and application of experimental autoimmune encephalomyelitis.** *Nat Rev Immunol* 2007, **7**(11):904–912.

48. Damico FM, Kiss S, Young LH: **Sympathetic ophthalmia.** *Semin Ophthalmol* 2005, **20**(3):191–197.
49. Bergstrom CT, Antia R: **How do adaptive immune systems control pathogens while avoiding autoimmunity?** *Trends Ecol Evol* 2006, **21**(1):22–28.
50. Dishaw LJ, Litman GW: **Changing views of the evolution of immunity.** *Front Immunol* 2013, **4**:122.
51. Kasamatsu J: **Evolution of innate and adaptive immune systems in jawless vertebrates.** *Microbiol Immunol* 2013, **57**(1):1–12.