CellPress

Trends in Immunology

Special issue: Systems immunology II

Opinion

Towards a quantitative theory of tolerance

Thierry Mora[1,*] and Aleksandra M. Walczak[1,*]

A cornerstone of the classical view of tolerance is the elimination of self-reactive T cells via negative selection in the thymus. However, high-throughput T cell receptor (TCR) sequencing data have so far failed to detect substantial signatures of negative selection in the observed repertoires. In addition, quantitative estimates as well as recent experiments suggest that the elimination of self-reactive T cells is at best incomplete. We discuss several recent theoretical ideas that might explain tolerance while being consistent with these observations, including collective decision-making through quorum sensing, and sensitivity to change through dynamic tuning and adaptation. We propose that a unified quantitative theory of tolerance should combine these elements to help to explain the plasticity of the immune system and its robustness to autoimmunity.

Negative selection: all or nothing?

Thymic selection in mammals is an important step in the generation of mature T cells that can protect us against foreign pathogens while avoiding autoimmunity. **Negative selection** (see Glossary) is said to eliminate T cells that bind too strongly to self peptides presented on the major histocompatibility complexes (MHCs) of antigen-presenting cells (APCs) in the thymus through apoptosis [1]. This ensures that T cells cannot trigger autoimmune reactions. The original evidence for this elimination was based on a transgenic mouse model where males and females were compared for the survival of T cells reactive to a peptide encoded by the male chromosome [2,3].

Recent studies have questioned the role of deletion as the main mechanism to avert autoimmunity. Abundant self-reactive T cells can be found in the blood [4] and tissues [5] of healthy humans and proliferate following the ablation of **regulatory T cells (Tregs)** [6]. In fact, turning the concept of negative selection into a quantitative theory of the entire repertoire is a challenging task [7], notably because of the physical constraints imposed by the large number of precise decision-making processes that it implies. With the advent of high-throughput, quantitative experimental techniques, there is currently a need for theories to quantitatively test the plausibility of these decision-making processes. Nevertheless, most of the time, tolerance is guaranteed and self-immunity is avoided. By checking whether the numbers 'add up', such theories can help us to reveal previously unobserved mechanisms and better understand tolerance at the systems level. We first outline two quantitative puzzles that have recently been identified, and then discuss possible partial solutions.

The failure to distinguish negatively selected TCRs

Ex vivo experiments suggest that selection in the thymus is TCR-specific and depends sharply on the affinity between the TCR and self-peptides [8]. Therefore, we should expect to observe clear statistical signatures of this selection in the DNA sequence identities of TCRs from different T cell populations. Such analyses are now possible with the recent rise of **high-throughput repertoire sequencing** [9]. Specifically, we would not only expect differences between sequences that pass or fail TCR selection but also between peripheral $CD4^+$ and $CD8^+$ T cells – given that these would be subjected to distinct selection forces

Highlights

Not all self-reactive T cells are deleted during negative selection. Accordingly, the α and β chains of the selected T cell repertoire are indistinguishable from those of the deleted repertoire.

There are no general rules that distinguish self from foreign peptides.

T cells make collective decisions to determine their fate, thereby providing a mechanism to decide more accurately how to respond to the presence of self or foreign peptides.

T cells adapt their response through feedback mechanisms, such as those conferred by regulatory T cells, allowing them to be robust against sustained antigenic stimulation from the organism's own proteins.

Significance

The classical theory of T lymphocyte negative selection, which explains why cells of the adaptive immune system do not attack our own tissues, has recently been complemented by the ideas of quorum sensing and adaptation to help to explain tolerance to self-antigens. This offers directions to explain current holes in our understanding of how autoimmunity is avoided.

[1]Laboratoire de Physique de l'Ecole Normale Supérieure, Centre National de la Recherche Scientifique (CNRS), Université Paris Sciences et Lettres (PSL University), Sorbonne Université, and Université Paris-Cité, 75005 Paris, France

*Correspondence:
thierry.mora@phys.ens.fr
(T. Mora) and
aleksandra.walczak@phys.ens.fr
(A.M. Walczak).

 https://doi.org/10.1016/j.it.2023.04.008

[10]. If the decision to eliminate a cell is a deterministic function of its TCR, then negatively selected cells should in theory have TCR sequences that are distinguishable from those of cells that fail to undergo negative selection.

Nevertheless, recent analysis based on TCR repertoire sequencing of subsets of thymocytes failed to reveal TCR features that might accurately predict whether a given cell will pass thymic selection or not [11]. Briefly, the α and β chains of the TCR were separately sequenced from mouse thymocytes that were sorted into double-positive ($CD4^+CD8^+$), activated double-positive (as marked by a Nur77 reporter in male inbred Nur77–GFP/Foxp3–mCherry mice on a C57BL/6 background), dying double-positive, and apoptotic $CD4^+$ and $CD8^+$ single-positive (SP) cells as well as splenic SP cells [11]. **Classifiers** were then **trained** on each TCR chain from these subsets using either **neural networks** or distributions of ***k*-mers** (where $k \leq 4$) within the third complementarity-determining region (CDR3) of the TCR. However, these classifiers did not distinguish apoptotic thymic SP cells from mature splenic SP cells at the single sequence level [11], suggesting that negative selection is imprecise or incomplete.

A limitation of that study was that apoptotic cells were sorted using an annexin V marker (a proxy for cell death), and some of these cells could be dying for other reasons, such as cell manipulation, which could confound the classification task. More importantly, each of the TCR α and β chains was sequenced in bulk (all cells), and thus analyzed separately. However, TCR affinity for its cognate peptides is determined by the combination of both chains. Although the analysis of antigen-specific TCR subsets suggests that the sequence of each single chain carries information about antigen specificity [12], part of the signal may be lost when considering the chains separately. Considering both chains together is thus an important missing aspect that could resolve this discriminability puzzle. Another issue is that achieving discrimination may be computationally difficult. For instance, one might think of the effect of negative selection by each self-peptide as 'digging a hole' in the potential repertoire. Given the significantly large number and diversity of self-peptides, detecting all these 'holes' may be too difficult for our computational techniques. Structural data [13] suggest that the TCR makes contact with the peptide at only a few amino acid sites, where hydrophobic residues are suppressed by negative selection, consistent with theoretical predictions [14]; however, it is not clear how one might identify those sites. Such obstacles in peptide recognition and T cell selection may reduce the sequence-encoded signal of negative selection, and thus the poor discriminability between 'passing' versus 'failing' TCR sequences during selection remains puzzling.

Detecting different selection pressures placed on $CD4^+$ and $CD8^+$ T cell TCRs should also be possible, given that these receptors interact with MHC-II and MHC-I, respectively, and hence with separate sets of self-peptides during negative selection. Indeed, these interaction differences are statistically significant at the repertoire level [10], but are too small to allow accurate discrimination of $CD4^+$ or $CD8^+$ T cells based solely on their TCR [11,15]. Unlike the discrimination between dying and surviving cells, $CD4^+$ and $CD8^+$ TCR repertoires do not need to be exclusive: in principle, the same TCR might be found in both repertoires, perhaps partially explaining why such a discrimination task may not be perfect.

By the same logic, we should also be able to distinguish the TCR of Treg subsets. Tregs are $CD4^+$ T cells, and presumably might recognize the same antigens as conventional T cells (Tconv), but with a bias for self-antigens, thus suppressing the inflammatory response caused by Tconv activation. For this reason, we posit that Tconv and Treg cells might share the same TCRs. Accordingly, statistical scores of TCRs have been learned from repertoire data that can identify modest but measurable differences in the TCR sequences of Treg and Tconv populations

Glossary

Classifier: in machine learning, an algorithm that orders data into two or more categories based on a set of input features.
Condorcet jury theorem: a mathematical result stating that a large collective making a decision using the majority rule can be highly reliable, even if each person is individually fallible.
High-throughput repertoire sequencing: massively parallel DNA sequencing of a large number of sequences of T and B cell receptors from a biological sample, obtained through the targeted amplification of mRNA or genomic DNA of the receptor genes by PCR.
***k*-mer:** a sequence of *k* consecutive amino acids in a protein sequence.
Learnability: given a dataset, the ability to learn a rule from a portion of the dataset that can be used to make predictions about unseen data.
Negative selection: the removal of self-reactive T cells during thymic development.
Neural network: in machine learning, a multilayered algorithmic structure mimicking biological neural networks that processes input data into a predictive output.
Overfitting: in machine learning, an undesirable property of a model where the algorithm learns the idiosyncratic features of the data used to train the model, but fails to learn general rules predictive of unseen data.
Regulatory T cells (Tregs): a subpopulation of $CD4^+$ T cells that suppress the immune response and are generally more self-reactive than conventional T (Tconv) cells.
Trained: in machine learning, training means fitting the parameters of the algorithm to optimize its predictive performance.

[15,16]. Again, these scores are not sufficiently powerful to classify individual TCRs. A similar discriminability was obtained when trying to predict self-reactivity (as measured by CD5 expression) from the TCR sequence, showing modest but statistically significant power [17].

In summary, although our concept of how thymic selection works suggests that there should be strong signatures of fate at the sequence level, in practice it seems difficult to identify TCR sequence signatures that can unambiguously determine T cell survival during thymic selection, or predict which subpopulation a T cell will join.

Can T cells screen all self-peptides?

The second puzzle, first noted by Butler *et al.* [18], concerns the numbers and timescales involved in the negative selection process. In principle, to avoid autoimmunity, each T cell should be screened against every presentable self-peptide–MHC complex. Do the cells have sufficient time to do that? Each T cell spends ~4–5 days in the mouse thymus, which gives them time to interact with ~500 APCs [19]. During each of these encounters, multiple copies of the TCR on the cell surface may bind to distinct peptide–MHC complexes presented by APCs. In principle, any of those engagements could result in T cell activation and subsequent apoptosis during negative selection. This means that the number of self-peptides that may be screened during each encounter with an APC during negative selection could be large.

Unfortunately, it is difficult to estimate that number. Previous literature calculated that the total number of screened peptides (across APC encounters) is in the thousands [20]. Butler *et al.* [18] used a self-consistency argument to estimate this number by quantifying its impact on the specificity of the selected TCR: their model predicted that the fraction of peptides recognized by a peripheral TCR is inversely related to the number of screened self-peptides during negative selection. Using peptide-specific precursor frequencies measured in mice [21], they estimated that the number of peptides screened is <7000, or less than 2–10% of the self-peptidome (assuming 20 000 mammalian coding genes, each comprising 300 antigenic peptides, and that 1–5% of peptides are presented by MHC molecules). These numbers suggest that T cells may not have sufficient time to ensure that they are not self-reactive, calling for additional mechanisms of tolerance.

The unlearnability of the self-peptidome

The issue of insufficient screening time might be solved if the ensemble of self-peptides could be 'learned' – in other words, if there are general rules or properties that distinguish self- from pathogen-derived and foreign peptides. TCRs might then generalize their 'knowledge' of self-peptides from their interactions with a few self-peptides and might not need to scan every single peptide. This idea of **learnability** has been previously proposed [22,23]. Both studies concluded that there are minute statistical differences in amino acid composition between self- and pathogen-derived peptides. However, they are too small to be used for efficient self versus non-self discrimination. Redundancy in the self-peptidome could in principle be exploited to reduce the number of scanned peptides by around half [22], but this number is still too large for the allotted time of negative selection. In other words, the difference between self- and non-self-peptides cannot be learned through rules, and TCRs actually do need to scan peptides comprehensively (i.e., memorize them) – akin to **overfitting** in machine-learning language.

This conclusion is consistent with observations that a single mutation in an epitope can turn a non-immunogenic self-peptide into an immunogenic neoantigen [24], and, more generally, that self- and pathogen-derived antigens from databases such as the Immune Epitope Database (IEDB, https://www.iedb.org/) [25] are promiscuous in sequence space [23]. Self-peptide

learnability also seems implausible from an evolutionary perspective. To avoid immunity, viruses should be under strong selective pressure to make themselves statistically indistinguishable from the host (self). If that is the case, then there are no rules stating that the immune system can learn to decide what is self and what is viral. Note that even if rules of the self-peptidome could be learned, we should still expect the sequences of self-reactive and non-self-reactive TCR to be different, meaning that our first puzzle hypothesis about sequence discriminability would hold.

Quorum sensing

A solution that might explain both puzzles was proposed by Butler *et al.* [18]. It relies on the idea (similar to the **Condorcet jury theorem**) that if N cells must make a decision, and the probability that each cell makes the right decision is better than chance, then the probability that a collective majority vote makes the right decision – or, equivalently, that a 'quorum' of activated cells is reached [26] – quickly approaches 1 as N increases. What distinguishes a self from a foreign peptide is that at least some of the self-reactive TCRs are removed, meaning that the number of precursor T cells that are specific for self-peptides falls below the quorum that is necessary to induce an immune reaction, while the number of precursor T cells that are specific for foreign peptides rises above such quorum. The optimal quorum has been estimated to be ~40 T cells, separating the case of self-peptides that are recognized by 10–30 T cells from that of foreign peptides which would recruit 50–100 T cells [18]. Recent *in vitro* mouse cell culture experiments provide evidence for this concept. By measuring the rate of differentiation and activation signals (phosphorylated STAT5) of progenitor central memory [27] and $CD4^+CD25^+$ T cells [28], cells were shown to sense the number of activated cells through the secretion and sensing of cytokines such as IL-2. By examining the commitment of activated T cells to memory cells [28], as well as the commitment of $CD8^+$ T cells to regulating the balance between proliferation and apoptosis [29], further evidence has been provided to suggest that T cell fate is indeed collective and requires a quorum of activated cells.

That argument assumes that the preselection frequency of peptide-specific precursors is relatively constant across peptides. However, this frequency is known to vary by one or two orders of magnitude, even across foreign epitopes [21,30,31]. The frequency may also be predictive of immune response magnitude and immunodominance, as suggested by the correlation between naive precursor frequencies and response measured using peptide–MHC tetramers in mice [21,30]. Imagine a self-peptide with a preselection frequency of 200 precursor T cells, half of which are removed by negative selection. From the quorum perspective, this is indistinguishable from a foreign peptide with a preselection frequency of 100 T cells, all of which survive selection because 100 T cells would be activated in both cases. Thus, to avoid autoimmunity, it is important that the absolute number of self-peptide-specific T cells is controlled, not merely their probability of removal. Therefore, this implies the existence of other mechanisms of regulation.

Regulation, dynamic tuning, and adaptation

We propose a third solution to our puzzles which relies on the adjusted responses of the immune system via dynamic tuning [32,33]. In its simplest form, immune perturbations activate both excitatory and inhibitory pathways. These two pathways may either compensate for each other, resulting in adaptation and tolerance, or trigger an immune response if the excitatory signal is dominant. Dynamic tuning is reminiscent of mechanisms of adaptation discussed in the context of simple signaling networks, such as the chemotactic network of the flagellar bacterium *Escherichia coli* [34]. In these systems activation of the pathway on short timescales is accompanied by repressive action on a longer timescale, resulting in transient activation followed by return to a low level of activity, even if the perturbation is sustained. A similar idea has been theorized in

the 'discontinuity theory' of immunity [35] that has mostly been discussed in the context of the innate immune system. According to this theory, the state of the immune system does not depend on immune stimuli *per se*, provided that these stimuli do not change over time. Instead, it responds to the kinetics of the antigenic or immune environment [36], thus allowing the system to detect perturbations and to adapt to sustained changes.

For T cells, obvious candidates for repressive regulatory signals are Tregs and anergy [37,38]. Anergy is a T cell state in which the ability to proliferate is impaired. This can occur when cells are stimulated weakly or in the absence of costimulatory signals, as when engaging self-peptides in the absence of inflammation. Tregs are selected in the thymus to be more self-reactive than Tconv cells, and they are essential for inducing tolerance and preventing autoimmune diseases. They proliferate in response to immune challenge and tend to suppress the immune response through cytokine signaling [27,39] by inducing anergy [38] or by pruning self-activated T cells [40].

This double mechanism of activation by effector T cells and repression by Treg cells is reminiscent of an incoherent feedforward loop. This regulatory architecture causes the system to respond to fold-changes rather than to the absolute value of the perturbation [41] and allows a rapid but controlled response [42]. It is consistent with the ideas of dynamic tuning and discontinuity theory. The same tug-of-war phenomenon between effector and Treg cells is central to the mechanism of quorum sensing discussed earlier [26,27], which is based on the local balance between cytokine secretion and consumption by the two types of cells in the vicinity.

Marshland *et al.* [43] proposed a theory of the balance between effector and Tregs in the stationary state (constant antigenic environment) based on an ecological description of interactions between T cells and antigens. In their model, self-peptides stimulate Tconv and Treg cells in the same way, such that the numbers of Tconv and Treg cells always counterbalance each other. Their calculations suggest that a minimum number of distinct Treg specificities is necessary to satisfy this balance and thus avoid autoimmunity [43].

For these theories, the fact that Tregs are more self-reactive than Tconv cells provides a natural mechanism for selectively suppressing autoimmunity, regardless of negative selection. Self-reactive TCRs are acceptable provided that they are not massively stimulated [44]. Accordingly, self-tolerance may be disrupted when a self-antigen or a crossreactive foreign antigen is overexpressed.

Nevertheless, it is not clear how one might reconcile these theories of adaptation with the immunology of asymptomatic chronic infections such as with cytomegalovirus, where antigen stimulation is stable but a large growing fraction of the TCR repertoire is mobilized [45]. Presumably, such antigen stimulation might be explained, at least in part, by a high adaptation plateau where the immune response is repressed; however, this hypothesis remains to be quantitatively checked. Another caveat of these theories is that Treg cells constitute a small fraction of the repertoire (about 5–10% of $CD4^+$ T cells); whether these small numbers are sufficient to regulate and repress autoimmunity is a quantitative question that warrants further investigation.

Concluding remarks

We have presented quantitative observations that seem to challenge the old view of negative thymic selection – which holds that self-reactive T cells should be largely eliminated from the periphery. Because these observations are based on indirect estimates and incomplete data, it is possible that these contradictions will evaporate upon more rigorous inspection (see Outstanding questions). In the modern view of tolerance, autoimmunity is averted by a combination

Outstanding questions

What are the molecular and structural determinants of self-reactivity, and how do they impact on the sequence statistics of selected repertoires? This knowledge could help us to better understand and diagnose autoimmune diseases using repertoire sequencing.

Does quorum sensing play an active role in self versus non-self discrimination by T cells?

Is dynamic tuning or adaptation an important mechanism of tolerization? Can the immune system learn to not respond to arbitrary peptides if delivered with the appropriate kinetics? Harnessing the timing of antigen exposure could inform future vaccination strategies.

Can a single theoretical modeling that combines the mechanisms of (partial) negative selection, quorum sensing, and dynamic tuning explain immune tolerance at a quantitative level and recapitulate existing data?

of mechanisms, among which is that negative selection is an important but not sufficient element. We have argued for the importance of evaluating possible scenarios of tolerance in a quantitative way to assess, by the numbers, whether these alone can explain how the immune system solves the task of self versus non-self discrimination, which is viewed as an information problem. A future theory will need to combine a few of the solutions outlined in the preceding text, particularly the ideas of quorum sensing and adaptation, with the hope that these ideas can mutually correct each other's inconsistencies. This theory should aim to explain existing data and make new testable predictions. This would help to provide missing evidence to complete our understanding of tolerance and to better predict how its balance is broken in autoimmune diseases.

Acknowledgments

The authors thank Grégoire Altan-Bonnet, Benny Chain, Arup Chakraborty, and Olivier Lantz for useful discussions. This research was supported by the European Research Council COG 724208 and ANR-19-CE45-0018 'RESP-REP' from the Agence Nationale de la Recherche.

Declaration of interests

The authors declare no conflicts of interest.

References

1. Palmer, E. (2003) Negative selection – clearing out the bad apples from the T-cell repertoire. *Nat. Rev. Immunol.* 3, 383–391
2. von Boehmer, H. (2008) Positive and negative selection in Basel. *Nat. Immunol.* 9, 571–573
3. Kisielow, P. *et al.* (1988) Tolerance in T-cell-receptor transgenic mice involves deletion of nonmature $CD4^+8^+$ thymocytes. *Nature* 333, 742–746
4. Yu, W. *et al.* (2015) Clonal deletion prunes but does not eliminate self-specific Aβ $CD8^+$ T lymphocytes. *Immunity* 42, 929–941
5. Legoux, F.P. *et al.* (2015) $CD4^+$ T cell tolerance to tissue-restricted self antigens is mediated by antigen-specific regulatory T cells rather than deletion. *Immunity* 43, 896–908
6. Lee, V. *et al.* (2023) The endogenous repertoire harbors self-reactive $CD4^+$ T cell clones that adopt a follicular helper T cell-like phenotype at steady state. *Nat. Immunol.* 24, 487–500
7. Yates, A.J. (2014) Theories and quantification of thymic selection. *Front. Immunol.* 5, 13
8. Daniels, M.A. *et al.* (2006) Thymic selection threshold defined by compartmentalization of Ras/MAPK signalling. *Nature* 444, 724–729
9. Robins, H. (2013) Immunosequencing: applications of immune repertoire deep sequencing. *Curr. Opin. Immunol.* 25, 646–652
10. Emerson, R. *et al.* (2013) Estimating the ratio of $CD4^+$ to $CD8^+$ T cells using high-throughput sequence data. *J. Immunol. Methods* 391, 14–21
11. Camaglia, F. *et al.* (2023) Quantifying changes in the T cell receptor repertoire during thymic development. *eLife* 12, e81622
12. Dash, P. *et al.* (2017) Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature* 547, 89–93
13. Stadinski, B.D. *et al.* (2016) Hydrophobic CDR3 residues promote the development of self-reactive T cells. *Nat. Immunol.* 17, 946–955
14. Kosmrlj, A. *et al.* (2008) How the thymus designs antigen-specific and self-tolerant T cell receptor sequences. *Proc. Natl. Acad. Sci. U. S. A.* 105, 16671–16676
15. Isacchini, G. *et al.* (2021) Deep generative selection models of T and B cell receptor repertoires with soNNia. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2023141118
16. Lagattuta, K.A. *et al.* (2022) Repertoire analyses reveal T cell antigen receptor sequence features that influence T cell fate. *Nat. Immunol.* 23, 446–457
17. Textor, J. *et al.* (2022) Machine learning analysis of the T cell receptor repertoire identifies sequence features that predict self-reactivity. *bioRxiv.* Published online March 22, 2023. https://doi.org/10.1101/2022.11.23.517563
18. Butler, T.C. *et al.* (2013) Quorum sensing allows T cells to discriminate between self and nonself. *Proc. Natl. Acad. Sci. U. S. A.* 110, 11833–11838
19. Le Borgne, M. *et al.* (2009) The impact of negative selection on thymocyte migration in the medulla. *Nat. Immunol.* 10, 823–830
20. Faro, J. *et al.* (2004) The impact of thymic antigen diversity on the size of the selected T cell repertoire. *J. Immunol. (Baltimore, Md.: 1950)* 172, 2247–2255
21. Jenkins, M.K. and Moon, J.J. (2012) The role of naive T cell precursor frequency and recruitment in dictating immune response magnitude. *J. Immunol.* 188, 4135–4140
22. Wortel, I.M.N. *et al.* (2020) Is T cell negative selection a learning algorithm? *Cells* 9, 690
23. Mayer, A. *et al.* (2022) How different are self and nonself? *arXiv.* Published online December 22, 2022. http://dx.doi.org/10.48550/arXiv.2212.12049
24. Matsushita, H. *et al.* (2012) Cancer exome analysis reveals a T-cell-dependent mechanism of cancer immunoediting. *Nature* 482, 400–404
25. Vita, R. *et al.* (2015) The immune epitope database (IEDB) 3.0. *Nucleic Acids Res.* 43, D405–D412
26. Burroughs, N. *et al.* (2006) Regulatory T cell adjustment of quorum growth thresholds and the control of local immune responses. *J. Theor. Biol.* 241, 134–141
27. Feinerman, O. *et al.* (2010) Single-cell quantification of IL-2 response by effector and regulatory T cells reveals critical plasticity in immune response. *Mol. Syst. Biol.* 6, 437
28. Polonsky, M. *et al.* (2018) Induction of CD4 T cell memory by local cellular collectivity. *Science* 360, eaaj1853
29. Zenke, S. *et al.* (2020) Quorum regulation via nested antagonistic feedback circuits mediated by the receptors CD28 and CTLA-4 confers robustness to T cell population dynamics. *Immunity* 52, 313–327
30. Moon, J.J. *et al.* (2007) Naive $CD4^+$ T cell frequency varies for different epitopes and predicts repertoire diversity and response magnitude. *Immunity* 27, 203–213
31. Legoux, F. *et al.* (2010) Impact of TCR reactivity and HLA phenotype on naive CD8 T cell frequency in humans. *J. Immunol.* 184, 6731–6738
32. Grossman, Z. and Paul, W.E. (2001) Autoreactivity, dynamic tuning and selectivity. *Curr. Opin. Immunol.* 13, 687–698
33. Bhandoola, A. *et al.* (2002) Peripheral expression of self-MHC-II influences the reactivity and self-tolerance of mature $CD4^+$ T cells. *Immunity* 17, 425–436
34. Tu, Y. *et al.* (2008) Modeling the chemotactic response of *Escherichia coli* to time-varying stimuli. *Proc. Natl. Acad. Sci. U. S. A.* 105, 14855–14860
35. Pradeu, T. *et al.* (2013) The speed of change: towards a discontinuity theory of immunity? *Nat. Rev. Immunol.* 13, 764–769

36. Johansen, P. *et al.* (2008) Antigen kinetics determines immune reactivity. *Proc. Natl. Acad. Sci. U. S. A.* 105, 5189–5194
37. Kalekar, L.A. and Mueller, D.L. (2017) Relationship between CD4 regulatory T cells and anergy in vivo. *J. Immunol.* 198, 2527–2533
38. Alonso, R. *et al.* (2018) Induction of anergic or regulatory tumor-specific $CD4^+$ T cells in the tumor-draining lymph node. *Nat. Commun.* 9, 2113
39. Busse, D. *et al.* (2010) Competing feedback loops shape IL-2 signaling between helper and regulatory T lymphocytes in cellular microenvironments. *Proc. Natl. Acad. Sci. U. S. A.* 107, 3058–3063
40. Wong, H.S. *et al.* (2021) A local regulatory T cell feedback circuit maintains immune homeostasis by pruning self-activated T cells. *Cell* 184, 3981–3997
41. Goentoro, L. *et al.* (2009) The incoherent feedforward loop can provide fold-change detection in gene regulation. *Mol. Cell* 36, 894–899
42. Rosenfeld, N. *et al.* (2002) Negative autoregulation speeds the response times of transcription networks. *J. Mol. Biol.* 323, 785–793
43. Marsland, R. *et al.* (2021) Tregs self-organize into a computing ecosystem and implement a sophisticated optimization algorithm for mediating immune response. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2011709118
44. Zehn, D. and Bevan, M.J. (2006) T cells with low avidity for a tissue-restricted antigen routinely evade central and peripheral tolerance and cause autoimmunity. *Immunity* 25, 261–270
45. Schober, K. *et al.* (2018) TCR repertoire evolution during maintenance of CMV-specific T-cell populations. *Immunol. Rev.* 283, 113–128