Rare event simulation in immunobiology

Florian Lipsmeier * Ellen Baake *

August 27, 2008

1 Introduction

The ability of the adaptive immune system to discriminate safely between foreign and self molecules is a fundamental ingredient to everyday survival of higher vertebrates (such as ourselves). Unlike innate immune responses, which happen in all kinds of organisms, responses of the adaptive immune system are highly specific to the target which induced their activation. Until now it is not quite clear how this process of recognition and activation works. The main problem here is to explain and understand a system which is obviously able to recognise one (or a few) type(s) of foreign molecules against an enormous variety of self molecules, although there is no a priori difference between "self" and "foreign" on the molecular level: There is no such thing as a "self"-marker on self molecules to enable an easy discrimination. (This is obvious, because such a marker could easily be forged by foreign intruders). A classification of molecules into self and foreign is even unique for every individual. This is most obvious if we think of organ transplantation, where, although it is human, the immune system tries to attack the donor organ because it recognizes it as foreign.

A novel approach to this problem of statistical recognition (of one particular foreign signal against a large, fluctuating self background) was established by van den Berg, Rand and Burroughs [26] (henceforth referred to as BRB) and further developed by Zint, Baake and den Hollander [28]. In contrast to many existing deterministic models, they formulate an explicit stochastic model. It describes (random) encounters between two crucial types of white blood cells (see Fig. 2): the antigen-presenting cells (APCs), which display a mixture of self and foreign antigens at their surface (a sample of the molecules around in the body), and the T cells, which "scan" the APCs by means of certain receptors and finally "decide" whether or not to react, i.e. to start an immune response.

Each T cell is characterised by a specific type of T cell receptor (TCR), which is displayed in many identical copies on the surface of the particular T cell. A large number (estimated at 10^7 in [1]) of different receptors, and hence different T cell types, are present in an individual. However, the number of potential antigen types is still vastly larger (roughly 10^{13}; see [16]). Thus, specific recognition is impossible; this is known as Mason's paradox. The task is further complicated by the fact that every APC displays on the order of thousand(s) of different "self" antigen types, in various copy numbers, together with, possibly, one (or a small number of) foreign types.

The probability that a T cell reacts to an encounter with a randomly chosen APC has to be very small in order to avoid autoimmune reactions. Some questions may therefore be answered analytically with the help of large deviation theory; others require simulation, but its use has been limited due to the low probabilities involved, at least with the straightforward simulation methods applied so far [26, 28]. Here we present an efficient method of rare event simulation. The article is organized as follows. Sec. 2 recapitulate the biological model; Sec. 3 the simulation method; and Sec. 4 presents some results obtained by applying this method to the T cell model. More details, as well as further results, may be found in [14].

*Faculty of Technology, Bielefeld University, Germany. E-mail: {flipsmei,ebaake}@techfak.uni-bielefeld.de

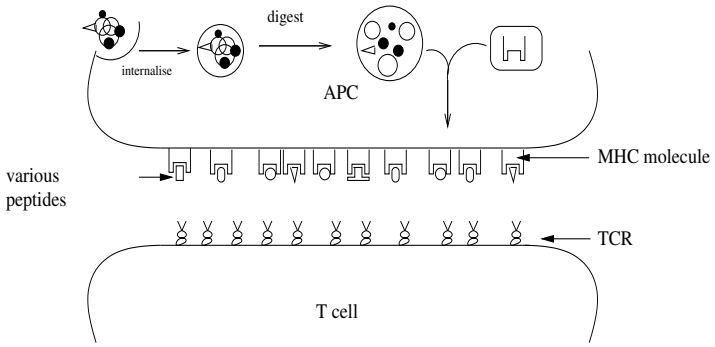

Figure 1: Caricature of a T cell and an APC. The APC takes up proteins from its enviroment, breaks them down and presents a random selection of the resulting fragments as antigens on its surface via so-called MHC molecules, which serve as "anchors". The T cell can bind to these peptide-MHC complexes via its receptors.

2 The T cell model

In this Section, we briefly motivate and introduce the model of T cell recognition as first proposed by BRB in 2001 [26] and further developed by Zint, Baake and den Hollander [28].

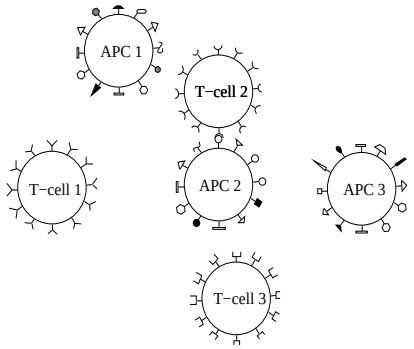

Figure 2: Caricature of T cells and APCs (from [28]). Note that every T cell has many copies of one particular receptor type, but different T cells have different receptor types. In contrast, every APC carries a mixture of antigen types, which may appear in various copy numbers.

When T cells and APCs meet, the T cell receptors bind to the various antigens presented by the APC [5]. For every single receptor-antigen pair, there is an association-dissociation reaction, the rate constant for which depend on the "match" of the molecular structures of receptor and antigen. Assuming that association is much faster than dissociation and that there is an abundance of receptors (so that the antigens are mostly in the bound state), one can describe the reaction in terms of the dissociation rates only.

Every time a receptor unbinds from an antigen, it sends a signal to the T cell, provided the association has lasted for at least one time unit (i.e., we rescale time so that the unit of time is this minimal association time required). If τ is the inverse dissociation rate (i.e. the mean duration of a binding) of a given receptor-antigen pair, the rate of such stimuli induced by the interaction of this individual antigen with the receptors in its vicinity is then given by

$$w(\tau) = \frac{1}{\tau} \exp(-\frac{1}{\tau}) \tag{1}$$

(i.e., the dissociation rate times the probability that the association has lasted long enough). As shown in Fig. 2, $w(\tau)$ first increases and then decreases with τ with a maximum at $\tau = 1$, which

reflects the fact that, for $\tau < 1$, the bindings tend not to last long enough, whereas for $\tau > 1$, they tend to last so long that only few stimuli are expected per time unit.

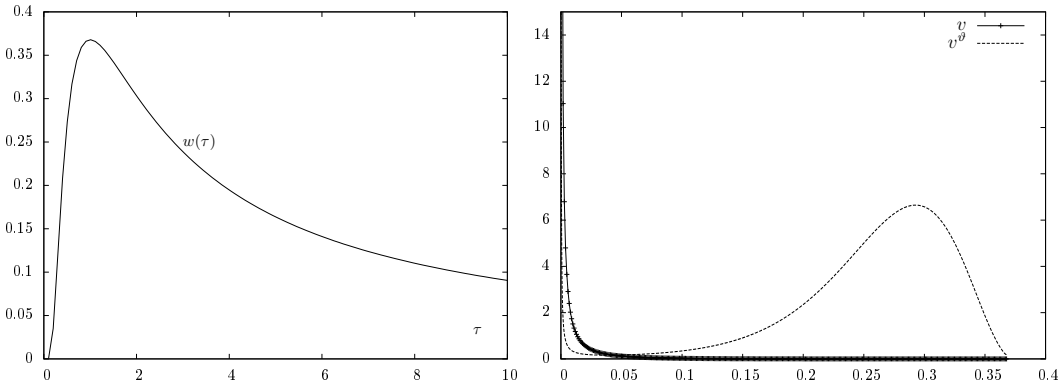

Figure 3: The function w (left) and the densities v and v^ϑ of $W = w(\mathcal{T})$ and W^ϑ, with tilting parameter $\vartheta = 46$ (right). The densities have poles at 0 and $w(1) = 0.3679$, but the latter is invisible because, in fact, it supports practically no probability mass.

The T cell sums up the signals induced by the various antigens on the APC, and if the total stimulation rate reaches a certain threshold value, the cell initiates an immune response. This model relies on several hypothesis, which are known as kinetic proofreading [17, 18, 15, 11], serial triggering [25, 24, 22, 3, 23, 9], counting of stimulated TCRs [27, 20], and the optimal dwell-time hypthoses [12, 10].

Due to the huge amount of different receptor and antigen types, it is impossible (and unnecessary) to prescribe the binding durations for all pairs of receptor and antigen types individually. Therefore, BRB chose a probabilistic approach to describe the meeting of APCs and T cells. A randomly chosen T cell (that is, a randomly chosen type of receptor) encounters a randomly chosen APC (that is, a random mixture of antigens). The mean binding time that governs the binding of this random receptor to the jth type of antigen is taken to be a random variable denoted by $\mathcal{T}_j$. The $\mathcal{T}_j$ are independent and identically distributed (i.i.d.) and are assumed to follow the $\mathrm{Exp}(1/\bar{\tau})$ distribution, i.e., the exponential distribution with mean $\bar{\tau}$, where $\bar{\tau}$ is a free parameter. Note that there are two exponential distributions (and two levels of averaging) involved here: The duration of an individual binding between a type-j antigen and a random receptor is $\mathrm{Exp}(1/\mathcal{T}_j)$ distributed. But $\mathcal{T}_j$, the mean duration of such a binding (where the receptor is chosen once and the times are averaged over repeated bindings with a j antigen) is itself an exponential random variable, with realisation τ_j and mean $\bar{\tau}$ when random receptors are considered (that is, $\bar{\tau}$ is the mean binding time of a j-antigens (and, due to the i.i.d. assumption, of any antigen) averaged over all encounters with the various receptor types). The exponential distribution of the individual binding time is an immediate consequence of the (first-order) unbinding kinetics (compare the discussion of Eq. (1)). In contrast, the corresponding assumption for the $\mathcal{T}_j$ is made for simplicity; the concept is compatible with various other distributions as well, see [26] and [28]. The i.i.d. assumption, however, is crucial, since it implies, in particular, that there is no difference between self and foreign antigens here; i.e., no a priori distinction is built into the model.

The total stimulation a T cell receives is the sum over all stimulus rates $W_j = w(\mathcal{T}_j)$ that emerge from antigens of the j'th type. It is further assumed that there is at most one type of foreign antigen in $z^{(f)}$ copies on an APC, whose signal must be discriminated against the signals of a huge amount of self antigens. The self antigens are here divided into two distinct classes, c and v, that are present in different copy numbers $z^{(c)}$ and $z^{(v)}$. An APC displays m_c and m_v different types of class c and v. The indices c and v stand for "constitutive" and for "variable", respectively; but for the purpose of this article, only the abundancies are relevant, in particular, $z^{(c)} > z^{(v)}$ and $m_c < m_v$. Over the whole APC the total number of antigens is then $m_c z^{(c)} + m_v z^{(v)} =: M$ if no

foreign antigen is present. If $z^{(f)}$ foreign molecules are also present, the self molecules are assumed to be proportionally displaced (via the factor $q := (M - z^{(f)})/M$), so that the total number of antigens remains unchanged at

$$z^{(f)} + m_c q z^{(c)} + m_v q z^{(v)} = M. \tag{2}$$

The total stimulation rate in a random encounter of T cell and APC can then be described as a function of $z^{(f)}$:

$$G(z^{(f)}) := \left(\sum_{j=1}^{m_c} q z^{(c)} W_j \right) + \left(\sum_{j=m_c+1}^{m_c+m_v} q z^{(v)} W_j \right) + z^{(f)} W_{m_c+m_v+1}, \tag{3}$$

i.e., a weighted sum of i.i.d. random variables.

In line with [26, 28], we numerically specify the model parameters as follows: $\bar{\tau} = 0.04$; $m_c = 50$, $m_v = 1500$, $z^{(c)} = 500$, $z^{(v)} = 50$, $M = 10^5$.

The relevant quantity for us is now the probability

$$\mathbb{P}\big(G(z^{(f)}) \geq g_{\text{act}}\big) \tag{4}$$

that the activation rate reaches or surpasses a threshold g_{act}. To achieve a good foreign-self discrimination, there must be a large difference in probability between the activation rate in the case with self antigens only ($z^{(f)} = 0$), and the activation rate with the foreign antigen present, i.e.,

$$\mathbb{P}\big(G(z^{(f)}) \geq g_{\text{act}}\big) \gg \mathbb{P}\big(G(0) \geq g_{\text{act}}\big) \tag{5}$$

for realistic values of $z^{(f)}$. Note that both events must be rare events—otherwise, the immune system would "fire" all the time. Thus g_{act} must be much larger than $\mathbb{E}(G(z^{(f)}))$ (which, due to (2) and the identical distribution of the W_j, is independent of $z^{(f)}$). Evaluating these small probabilities is a challenge. So far, two routes have been used: analytic (asymptotic) theory based on large deviations (LD) and straightforward simulation (so-called simple sampling). Both have their shortcomings: The LD approach is only exact in the limit of infinitely many antigen types; the simulation strategy, on the other hand, is so time-consuming that it becomes simply impossible to obtain sample sizes large enough for a detailed analysis, in particular for large values of g_{act}. Therefore, an importance sampling approach is required.

3 Rare event simulation

Consider the problem of estimating the probability $\mathbb{P}(A)$ of a (rare) event A under a probability measure P. As is well-known, the essential idea behind importance sampling is to find a sampling distribution Q instead of the original P which is taylored to the problem so that the variance of the importance sampling estimate

$$\widehat{(P_Q(A))}_N := \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\{T^{(i)} \in A\} \frac{dP}{dQ}(T^{(i)}) \tag{6}$$

(where the $\{T^{(i)}\}_{1 \leq i \leq N}$ are i.i.d. random variables with distribution P, $\mathbb{1}\{.\}$ denotes the indicator function and N is the sample size) is reduced compared to the variance of the simple sampling estimate (obtained from (6) by setting $Q = P$). Finding a good sampling distribution is highly dependent on the specific structure of the problem at hand. Apart from some general purpose methods, many ad hoc strategies (see [4]) and solutions which are tailored to specific problems, there exists the more systematic technique of large deviation simulation as introduce by Sadowski and Bucklew [21]. It is suitable for problems that can be embedded into a sequence of problems characterised by a sequence of random variables $\{S_n\}$ with probability measures $\{P_n\}$ for which a so-called large deviation principle is valid [7, 6]. Van den Berg et al. and Zint et al. showed that

this applies to the BRB model [26, 28]. We therefore use this theory following ideas mainly from Bucklew [4] and Dieker and Mandjes [8]. For the theoretical background (and the proofs) we refer to [14].

In large deviation simulation the usual way to construct a sampling distribution P_n^ϑ is via so-called tilting of the original distribution P_n, such that the event we want to estimate becomes a typical event. That is, if P_n is a probability measure on $\mathbb{R}$, P_n^ϑ is obtained from P_n via exponential reweighting,

$$\frac{dP_n^\vartheta}{dP_n}(x) = \frac{e^{n\vartheta x}}{\mathbb{E}_{P_n}(e^{n\vartheta S_n})}, \tag{7}$$

where $\frac{dP_n^\vartheta}{dP_n}$ denotes the Radon-Nikodym derivative and ϑ is chosen so that

$$\mathbb{E}_{P_n^\vartheta}(S_n) = a \tag{8}$$

if we want to estimate the probability of the rare event $\{S_n \in A\}$, $A := [a, \infty)$, $a > \mathbb{E}(S_n)$.

To apply this to the BRB model, we consider $n = m := m_c + m_v + m_f$, where $m_f = 0$ or $m_f = 1$ depending on whether foreign antigen is absent or present, and identify S_m with $G(z^{(f)})/m$ and a with g_{act}/m. Tilting S_m with $m\vartheta$ then corresponds to tilting $G(z^{(f)})$ with ϑ. This, in turn, is equivalent to tilting every summand in (3) with ϑ (since these summands are independent). The only difficulty lies in sampling from the tilted W-distribution: No direct method (via transformation) is available, and the numerical calculation of the distribution poses difficulties due to the singularities at the boundaries of the support (see Fig. 2 (right)). We therefore propose to do the tilting at the level of the $\mathcal{T}_j$, rather than the $W_j = w(\mathcal{T}_j)$, via the following simple result.

Fact 1 *Let X be a real-valued random variable with probability measure F, and let $Y := h(X)$, where $h : \mathbb{R} \to \mathbb{R}$ is an F-measurable function. Y then has measure $V := F \circ h^{-1}$, where $h^{-1}(y)$ denotes the preimage of y. Assume now that $\mathbb{E}_F(e^{\vartheta h(X)})$ exists and let $\tilde{X}^\vartheta$ be a random variable with probability measure $\tilde{F}^\vartheta$ related to F via*

$$\frac{d\tilde{F}^\vartheta}{dF}(x) = \frac{e^{\vartheta h(x)}}{\mathbb{E}_F(e^{\vartheta h(X)})}, \tag{9}$$

and let $\tilde{Y}^\vartheta := h(\tilde{X}^\vartheta)$. Then the measures $\tilde{V}^\vartheta$ (of $\tilde{Y}^\vartheta$) and V^ϑ (for the tilted version of V, the measure of Y^ϑ) are equal, with Radon-Nikodym derivative

$$\frac{d\tilde{V}^\vartheta}{dV}(y) = \frac{dV^\vartheta}{dV}(y) = \frac{e^{\vartheta h(y)}}{\mathbb{E}_V(e^{\vartheta h(Y)})}. \tag{10}$$

This fact allows us to "pull back" the tilting from the level of W to the level of $\mathcal{T}$, which is computationally unproblematic. If f is the density of $\mathcal{T}$ (i.e., $f(\tau) = \frac{1}{\bar{\tau}} e^{-\tau/\bar{\tau}}$) this yields three different densities $\tilde{f}_\alpha^\vartheta$ depending on the weighting factors $\alpha \in \{qz^{(c)}, qz^{(v)}, z^{(f)}\}$, namely,

$$\frac{e^{\alpha\vartheta w(\tau)}}{\mathbb{E}_f(e^{\alpha\vartheta w(\tau)})} f(\tau). \tag{11}$$

Note that these densities do not coincide with the "usual" tilted versions of f (tilting is with $w(\tau)$ rather than τ). They are thus not of the form of any known standard distribution (in particular, they are not exponential). Simulating from the tilted distribution requires numerical integration (which is well-behaved since f is numerically well-behaved). The resulting distribution functions $\tilde{F}_\alpha^\vartheta$ are discretized and tabulated, followed by "looking up" the solution $\tilde{\mathcal{T}}^\vartheta$ of $\tilde{F}_\alpha^\vartheta(\tilde{\mathcal{T}}^\vartheta) = U$ for $U \sim \text{Uni}_{[0,1]}$ (the uniform distribution on the unit interval), to finally yield $\alpha W^\vartheta = \alpha w(\tilde{\mathcal{T}}^\vartheta)$. To circumvent the speed limiting step of searching through the table we apply the so-called alias method for discrete random number generation (see [14, 19, 13]).

Now with everything at hand we formulate the algorithm to simulate realisations of $G(z^{(f)})$. (For notational convenience, we will not distinguish between random variables and their realisations here).

Algorithm 1

compute ϑ numerically so that (8) is satisfied; see [14] for the explicit procedure
calculate the tilted densities $\tilde{f}^{\vartheta}_{\alpha}$, $\alpha \in \{qz^{(c)}, qz^{(v)}, z^{(f)}\}$, via (11)
for** i=1 till sample size N **do
for every summand j of (3) generate a sample $(\tilde{\mathcal{T}}^{\vartheta}_j)^{(i)}$ according to its density $\tilde{f}^{\vartheta}_{\alpha(j)}$ with the help of the alias method (here, the upper index (i) is added to reflect sample i, and $\alpha(j)$ is the weighting factor of the sum to which j belongs)
calculate

$$\big(G(z^{(f)})\big)^{(i)} = \left(\sum_{j=1}^{m_c} qz^{(c)} w\big((\tilde{\mathcal{T}}^{\vartheta}_j)^{(i)}\big)\right) + \left(\sum_{j=m_c+1}^{m_c+m_v} qz^{(v)} w\big((\tilde{\mathcal{T}}^{\vartheta}_j)^{(i)}\big)\right) + z^{(f)} w\big((\tilde{\mathcal{T}}^{\vartheta}_{m_c+m_v+1})^{(i)}\big)$$

calculate the indicator function times the reweigthing factor (i.e., the i-th summand in Eq. (6))
if** $(G(z^{(f)}))^{(i)} \geq g_{\text{act}}$ **then

$$R^{(i)} = \prod_{j=1}^{m} \frac{f_{\alpha(j)}((\tilde{\mathcal{T}}^{\vartheta}_j)^{(i)})}{\tilde{f}^{\vartheta}_{\alpha(j)}((\tilde{\mathcal{T}}^{\vartheta}_j)^{(i)})}$$

else
$R^{(i)} = 0$
end if
end for
calculate $\big(\widehat{P^{\vartheta}_{P_m}(A)}\big)_N = \dfrac{\sum_{i=1}^{N} R^{(i)}}{N}$, *as estimate of* $\mathbb{P}(G(z^{(f)}) > g_{\text{act}})$.

4 Results

In this Section we briefly examine the performance of our sampling method and then extract some insights about the T cell model obtained with the help of simulation.

4.1 Performance analysis

We compare our method to the simple sampling method and the exact asymptotics based on large deviation theory, as used in [28]. The goal is to estimate the probability $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$ as a function of g_{act} for various values of $z^{(f)}$. In immunobiology the corresponding graph is known as an activation curve. The results of the three methods are summarized in Figure 4.

The main difference between our method and simple sampling is that, for every value of g_{act}, we calculate a new tilting factor and generate a fresh sample for the estimation, whereas this is not required for simple sampling. Here we just generate a huge amount of samples first and then calculate estimates. At first sight this seems to be a big disadvantage for our method, but reality proves different. Since $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$ decreases exponentially with g_{act}, the number of samples required to obtain sufficientily precise estimates with simple sampling increases exponentially. With LD sampling only a linear increase in sample size is necessary. To generate the graphs in Figure 4 we used 10000 samples per threshold value for LD sampling, that is $1.9 * 10^5$ samples altogether. For simple sampling we generated $1.3 * 10^8$ samples. Within this sample, there were no realizations of $G(z^{(f)})$ beyond 400 (for $z^{(f)} = 1000$) and 800 (for $z^{(f)} = 2000$), respectively. With LD sampling, however, we get realizations (and thus estimates) for the whole range of g_{act} values. The huge difference in sample size is of course reflected in the runtime. Estimation via LD sampling took only a few minutes, whereas simple sampling required 48 hours on a standard PC.

Last not least, Fig. 4 shows that the results of simple sampling and importance sampling agree wherever they can be compared - as was to be expected. Also, the agreement with exact asymptotics is excellent for large threshold values, whereas some (small) finite size effects are revealed at smaller g_{act}.

The precision of the estimates is best analysed in terms of the relative error in the sense of Dieker and Mandjes [8], i.e. the standard deviation of the estimate divided by the estimate. As

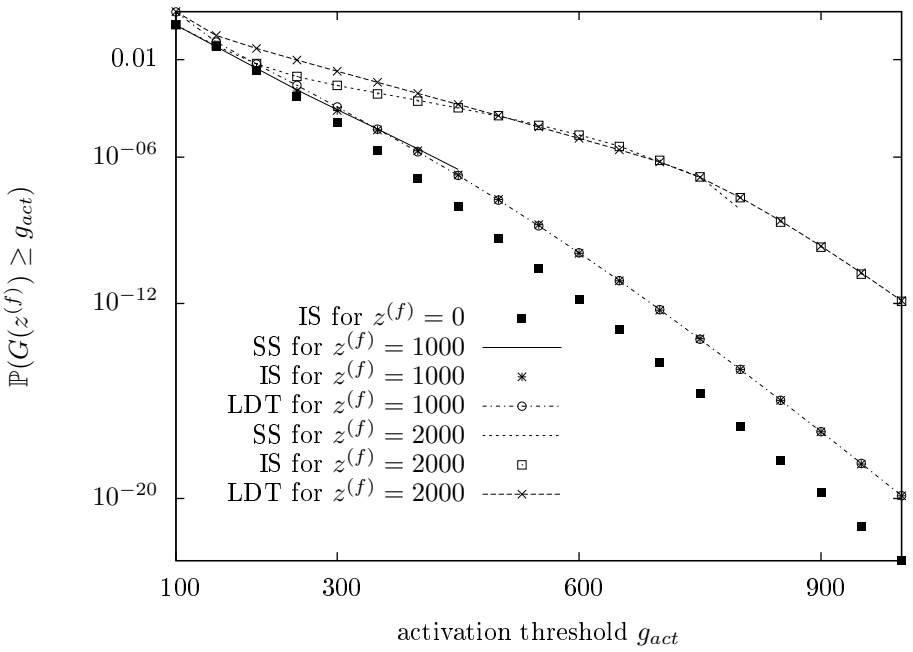

Figure 4: Estimates of the activation curve, $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$, in the model (3) for $z^{(f)} = 1000$ and $z^{(f)} = 2000$, as well as for the self background ($z^{(f)} = 0$), on logarithmic scale. The probablities were estimated independently with three different methods at 19 values of g_{act} (from 100 to 1000 in steps of 50). SS and IS denote simple and importance sampling, respectively, and LDT stands for the exact asymptotics based on large deviation theory as used in [28]. For the IS simulation, $N = 1.9 * 10^5$ samples were generated (10000 per threshold), whereas for the SS simulation, $N = 1.3 * 10^8$ samples were used over the entire range. The SS curves end at $g_{\text{act}} = 400$ and $g_{\text{act}} = 800$, respectively, because larger values were not hit in the given samples. The IS and SS graphs agree perfectly until the SS simulation lacks precision. For larger threshold values, we see a perfect agreement of the IS and LDT graphs. Note the general feature that, for threshold values that are not too small, the activation probability in the presence of foreign antigens is several orders of magnitude larger than the self background, i.e. Eq. (5) is satisfied.

is to be expected, this increases exponentially with g_{act} for simple sampling, whereas it reamins more or less constant for importance sampling (if comparable sample sizes are used).

4.2 Analysis of the T cell model

In this Section, we will use our simulation method to obtain more detailed insight into the phenomenon of statistical recognition in the T cell model. As discussed before, the task is to discriminate one foreign antigen type against a "noisy" background of a large number of self antigens. We already know from Fig. 4 that, for threshold values that are not too small, the activation probability in the presence of foreign antigens is several orders of magnitude larger than the self-background, i.e. Eq. (5) is satisfied. As discussed in [28], this distinction relies on $z^{(f)} > z^{(c)}, z^{(v)}$—basically, what happens is that larger copy numbers of the foreign antigen "thicken" the tail of the distribution of $G(z^{(f)})$ (without changing its mean), so that the threshold is more easily surpassed. The self-nonself distinction may, according to this model, be roughly described as follows. For a given antigen (foreign or self), finding a "highly-stimulating" T cell receptor is a rare event; but if it occurs to a foreign antigen, it occurs manifold since there are numerous copies, which all contribute the same large signal, since all receptors of the T cell involved are identical; the resulting activation rate is thus high. In contrast, if it is a self antigen that finds a highly-stimulating

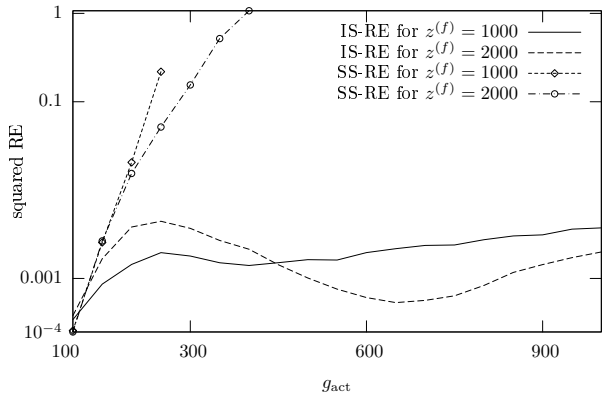

Figure 5: Estimated squared relative error (RE) for simple sampling (SS; $N = 10000$ times the number of steps contained in the considered interval), and importance sampling (IS; $N = 10000$ per threshold value) simulations of $\mathbb{P}(G(z^{(f)}) \geq g_{\mathrm{act}})$. Note that the vertical axis is on logarithmic scale.

receptor, the effect is less pronounced due to the smaller copy numbers. Put this way, the toy model "explains" the distinction solely on the basis of copy numbers; in particular, the distinction requires $z^{(f)} > z^{(c)}, z^{(v)}$ [26, 28].

Following these intuitive arguments, we now aim at a more detailed picture of how the self background looks, and how the foreign type stands out against it. To investigate this, it is useful to consider the histograms of the total constitutive, variable, and foreign activation rates, i.e., the contributions of the "constitutive sum", the "variable sum", and the individual "foreign term" in the sum (3), conditional on $\{G(z_f) \geq g_{\mathrm{act}}\}$ for various g_{act}. Since this requires a higher resolution (and thus larger sample size) than the calculation of the activation probabilities alone, such analysis would be practically impossible with simple sampling. With importance sampling, we again generated 10000 samples per g_{act} value, from which between 30 and 70 percent turned out to reach the threshold.

Figure 6 shows the resulting histograms for four different g_{act} values. The first observation from these histograms is that the variable activation rate is closely peaked around a small mean, unchanged under conditioning. In terms of signal discrimination this is of course favourable, as it is easy to predict in advance. The picture is different for the constitutive and foreign activation rate. For $g_{\mathrm{act}} = 100$ the situation is still similar to the unconditional case. The foreign activation rate is closely peaked at 0 and more or less follows the W-density from Figure 2. For $g_{\mathrm{act}} = 250$ where, according to Fig. 4, foreign-self distinction sets in, the foreign activation rate becomes prominent: The right branch of the W-distribution now becomes populated, and the activation rates are large due to the large copy number $z^{(f)}$. Still for $g_{\mathrm{act}} = 250$, the foreign activation rate is close to 0 in a sizeable fraction of the cases in which an immune reaction occurs - here, the reaction is brought about by the constitutive background, which moves to the right. Fig. 7 shows that the constitutive and foreign activation rates are, indeed, negatively correlated: as is to be expected, low foreign rates are compensated by high constitutive rates and vice versa. If g_{act} is increased further, every T cell beyond the threshold displays high stimuli for the foreign antigen, their distribution shifting even further to the right and concentrating near the maximal stimulation rate given by the maximum of the function w of Eq. (1), more precisely, by $z^{(f)}w(1)$. This maximum can, of course, not change due to conditioning; thus any further increase of g_{act} must then be matched by a shift of the constitutive background.

As pointed out before, this model is a very simplistic one, but it serves as a basis for the development of more realistic ones. To enhance foreign-self discrimination, the mechanism must work for far lower numbers $z^{(f)}$ of foreign antigen and for lower activation thresholds. As our observations show, this requires a reduction of the constitutive background. This is achieved in

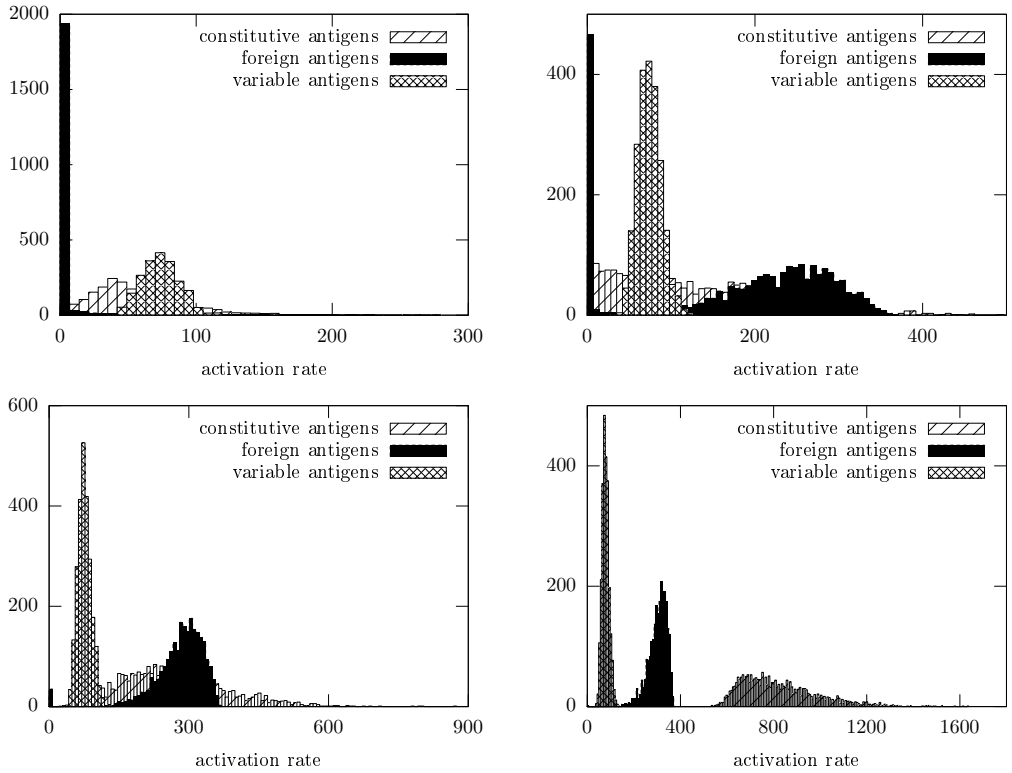

Figure 6: Histograms of the total constitutive, variable and foreign activation rates for for $z^{(f)} = 1000$. Sample size is 10000, and the vertical axis holds the number of samples that reach the threshold g_{act} and whose total constitutive (variable, foreign) activation rate falls into a given interval for $g_{\text{act}} = 100$ (upper left), $g_{\text{act}} = 250$ (upper right), $g_{\text{act}} = 500$ (lower left), $g_{\text{act}} = 1000$ (lower right). The maximal activation rate for the foreign antigens is $z^{(f)}w(1) = 367.9$. Note that the scaling of both axes varies across diagrams.

an advanced version of the BRB model [26] that includes negative selection, a mechanism that prevents the release of T cells which react too strongly to self activation. This model is even harder to analyze analytically, but it should be possible to further develop our simulation method to cope with the problem.

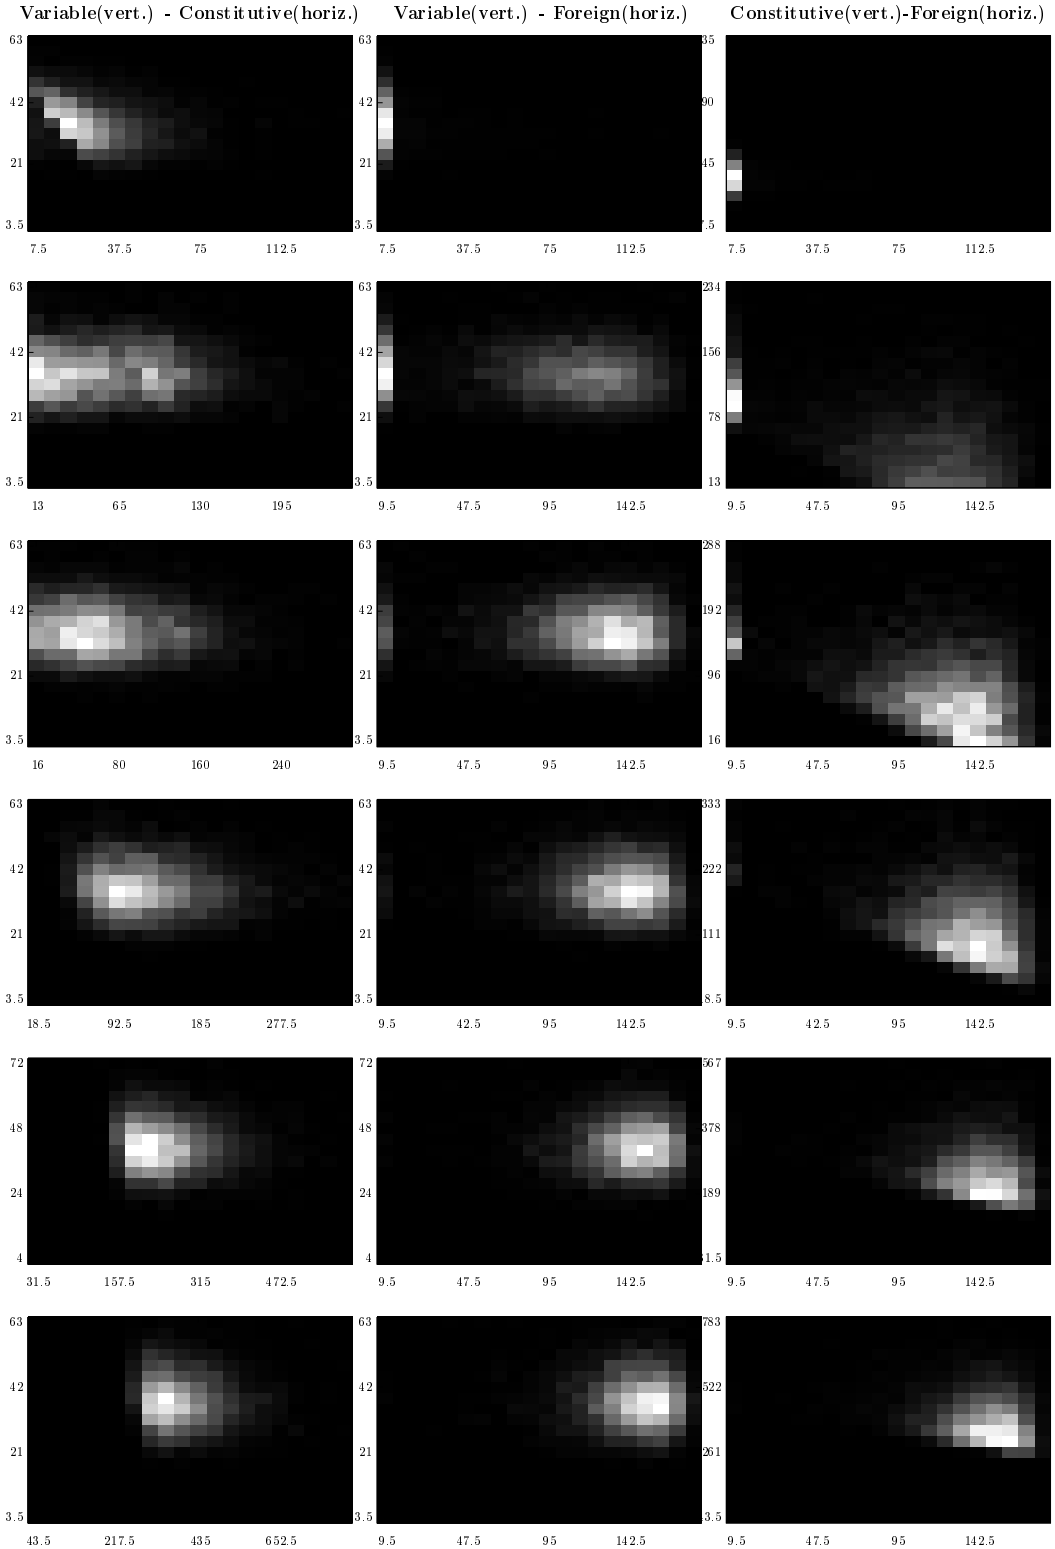

Figure 7: Joint empirical distribution of the total constitutive, variable, and foreign activation rates. Greyscales correspond to normalised number of samples falling into 2D-intervals defined by total activation rates of pairs of antigen types. Rows (from top to bottom): $g_{\text{act}} = 100, 250, 350, 500, 750, 1000$; columns (from left to right): variable (vertical)–constitutive (horizontal); variable (vertical)–foreign (horizontal); constitutive (vertical)–foreign (horizontal). Lighter shading corresponds to higher frequencies.

References

[1] Arstila, T., Casrouge, A., Baron, V., Even, J., Kannelopoulos, J., Kourilsky, P.: A direct estimate of the hman $\alpha\beta$ T cell receptor diversity. Science **286**, 958–961 (1999)

[2] van den Berg, H.A., Molina-París, C.: Thymic presentation of autoantigens and the efficiency of negative selection. J. Theor. Med. **5**, 1–22 (2003)

[3] Borovsky, Z., Mishan-Eisenberg, G., Yaniv, E., Rachmilewitz, J.: Serial triggering of T cell receptors results in incremental accumulation of signaling intermediates. J. Biol. Chem. **277**, 21529–21536 (2002).

[4] Bucklew, J.A.: Introduction to Rare Event Simulation. Springer, New York (2004)

[5] Davis, S.J., Ikemizu, S., Evans, E.J., Fugger, L., Bakker, T.R., van der Merwe, P.A.: The nature of molecular recognition by T cells. Nat. Immunol. **4**, 217–224 (2003).

[6] Dembo, A., Zeitouni, O.: Large Deviations Techniques and Applications. Springer, New York (1998)

[7] den Hollander, F.: Large Deviations. AMS, Providence, RI (2000)

[8] Dieker, A., Mandjes, M.: On asymptotically efficient simulation of large deviation probabilities. Adv. Appl. Prob. **37**, 539–552 (2005)

[9] Dushek, O., Coombs, D.: Analysis of serial engagement and peptide-MHC transport in T cell receptor microclusters. Biophys. J. **94**, 3447–3460 (2008).

[10] Gonzalez, P.A., Carreno, L.J., Coombs, D., Mora, J.E., Palmieri, E., Goldstein, B., Nathenson, S.G., Kalergis, A.M.: T cell receptor binding kinetics required for T cell activation depend on the density of cognate ligand on the antigen-presenting cell. Proc. Natl. Acad. Sci. U.S.A **102**, 4824–4829 (2005).

[11] Hlavacek, W.S., Redondo, A., Wofsy, C., Goldstein, B.: Kinetic proofreading in receptor-mediated transduction of cellular signals: receptor aggregation, partially activated receptors, and cytosolic messengers. Bull. Math. Biol. **64**, 887–911 (2002).

[12] Kalergis, A.M., Boucheron, N., Doucey, M.A., Palmieri, E., Goyarts, E.C., Vegh, Z., Luescher, I.F., Nathenson, S.G.: Efficient T cell activation requires an optimal dwell-time of interaction between the TCR and the pMHC complex. Nat. Immunol. **2**, 229–234 (2001).

[13] Kronmal, R.A., Peterson, A.J.: On the alias method for generating random variables from a discrete distribution. Amer. Stat. **33**, 214–218 (1979)

[14] Lipsmeier, F., Baake, E.: Rare event simulation for T cells, in preparation

[15] Lord, G.M., Lechler, R.I., George, A.J.: A kinetic differentiation model for the action of altered TCR ligands. Immunol. Today **20**, 33–39 (1999)

[16] Mason, D.: A very high level of crossreactivity is an essential feature of the T-cell receptor. Immunol. Today **19**, 395–404 (1998)

[17] McKeithan, T.W.: Kinetic proofreading in T-cell receptor signal transduction. Proc. Natl. Acad. Sci. U.S.A. **92**, 5042–5046 (1995)

[18] Rabinowitz, J.D., Beeson, C., Wulfing, C., Tate, K., Allen, P.M., Davis, M.M., McConnell, H.M.: Altered T cell receptor ligands trigger a subset of early T cell signals. Immunity **5**, 125–135 (1996)

[19] Ross, S.M.: Simulation. Academic Press (2002)

[20] Rothenberg, E.V.: How T cells count. Science **273**, 78–79 (1996).

[21] Sadowsky, J.S., Bucklew, J.A.: On large deviations theory and asymptotically efficient Monte Carlo estimation. IEEE TIT **36**, 579–588 (1990)

[22] Sousa, J., Carneiro, J.: A mathematical analysis of TCR serial triggering and down-regulation. Eur. J. Immunol. **30**, 3219–3227 (2000)

[23] Utzny, C., Coombs, D., Muller, S., Valitutti, S.: Analysis of peptide/MHC-induced TCR downregulation: deciphering the triggering kinetics. Cell Biochem. Biophys. **46**, 101–111 (2006).

[24] Valitutti, S., Lanzavecchia, A.: Serial triggering of TCRs: a basis for the sensitivity and specificity of antigen recognition. Immunol. Today **18**, 299–304 (1997)

[25] Valitutti, S., Muller, S., Cella, M., Padovan, E., Lanzavecchia, A.: Serial triggering of many T-cell receptors by a few peptide-MHC complexes. Nature **375**, 148–151 (1995).

[26] van den Berg, H.A., Rand, D.A., Burroughs, N.J.: A reliable and safe T cell repertoire based on low-affinity T cell receptors. J. Theor. Biol. **209**, 465–486 (2001).

[27] Viola, A., Lanzavecchia, A.: T cell activation determined by T cell receptor number and tunable thresholds. Science **273**, 104–106 (1996).

[28] Zint, N., Baake, E., den Hollander, F.: How T-cells use large deviations to recognize foreign antigens. J. Math. Biol., in press (2008)