Negative Examples for Sequential Importance Sampling of Binary Contingency Tables

Ivona Bezáková[1], Alistair Sinclair[2], Daniel Štefankovič[3], and Eric Vigoda[4]

[1] Department of Computer Science, University of Chicago, Chicago, IL 60637, ivona@cs.uchicago.edu.
[2] Computer Science Division, University of California, Berkeley, CA 94720, sinclair@cs.berkeley.edu*.
[3] Department of Computer Science, University of Rochester, Rochester, NY 14627, stefanko@cs.rochester.edu.
[4] College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, vigoda@cc.gatech.edu**.

Abstract. The sequential importance sampling (SIS) algorithm has gained considerable popularity for its empirical success. One of its noted applications is to the binary contingency tables problem, an important problem in statistics, where the goal is to estimate the number of 0/1 matrices with prescribed row and column sums. We give a family of examples in which the SIS procedure, if run for any subexponential number of trials, will underestimate the number of tables by an exponential factor. This result holds for any of the usual design choices in the SIS algorithm, namely the ordering of the columns and rows. These are apparently the first theoretical results on the efficiency of the SIS algorithm for binary contingency tables. Finally, we present experimental evidence that the SIS algorithm is efficient for row and column sums that are regular. Our work is a first step in determining rigorously the class of inputs for which SIS is effective.

1 Introduction

Sequential importance sampling is a widely-used approach for randomly sampling from complex distributions. It has been applied in a variety of fields, such as protein folding [8], population genetics [5], and signal processing [7]. Binary contingency tables is an application where the virtues of sequential importance sampling have been especially highlighted; see Chen et al. [4]. This is the subject of this note. Given a set of non-negative row sums $r = (r_1, \dots, r_m)$ and column sums $c = (c_1, \dots, c_n)$, let $\Omega = \Omega_{r,c}$ denote the set of $m \times n$ 0/1 tables with row sums r and column sums c. Our focus is on algorithms for sampling (almost) uniformly at random from Ω, or estimating $|\Omega|$.

* Supported by NSF grant CCR-0121555.
** Supported by NSF grant CCR-0455666.

Sequential importance sampling (SIS) has several purported advantages over the more classical Markov chain Monte Carlo (MCMC) method, such as:

Speed: Chen et al. [4] claim that SIS is faster than MCMC algorithms. However, we present a simple example where SIS requires exponential (in n, m) time. In contrast, a MCMC algorithm was presented in [6, 1] which is guaranteed to require at most time polynomial in n, m for every input.

Convergence Diagnostic: One of the difficulties in MCMC algorithms is determining when the Markov chain of interest has reached the stationary distribution, in the absence of analytical bounds on the mixing time. SIS seemingly avoids such complications since its output is guaranteed to be an unbiased estimator of $|\Omega|$. Unfortunately, it is unclear how many estimates from SIS are needed before we have a guaranteed close approximation of $|\Omega|$. In our example for which SIS requires exponential time, the estimator appears to converge, but it converges to a quantity that is off from $|\Omega|$ by an exponential factor.

Before formally stating our results, we detail the sequential importance sampling approach for contingency tables, following [4]. The general importance sampling paradigm involves sampling from an 'easy' distribution μ over Ω that is, ideally, close to the uniform distribution. At every round, the algorithm outputs a table T along with $\mu(T)$. Since for any μ whose support is Ω we have

$$E[1/\mu(T)] = |\Omega|,$$

we take many trials of the algorithm and output the average of $1/\mu(T)$ as our estimate of $|\Omega|$. More precisely, let $T_1, \ldots, T_t$ denote the outputs from t trials of the SIS algorithm. Our final estimate is

$$X_t = \frac{1}{t} \sum_i \frac{1}{\mu(T_i)}. \tag{1}$$

One typically uses a heuristic to determine how many trials t are needed until the estimator has converged to the desired quantity.

The sequential importance sampling algorithm of Chen et al. [4] constructs the table T in a column-by-column manner. It is not clear how to order the columns optimally, but this will not concern us as our negative results will hold for any ordering of the columns. Suppose the procedure is assigning column j. Let $r'_1, \ldots, r'_m$ denote the residual row sums after taking into account the assignments in the first $j-1$ columns.

The procedure of Chen et al. chooses column j from the correct probability distribution conditional on c_j, $r'_1, \ldots, r'_m$ and the number of columns remaining (but ignoring the column sums $c_{j+1}, \ldots, c_n$). This distribution is easy to describe in closed form. We assign column j the vector $(t_1, \ldots, t_m) \in \{0,1\}^m$, where $\sum_i t_i = c_j$, with probability proportional to

$$\prod_i \left(\frac{r'_i}{n' - r'_i} \right)^{t_i}, \tag{2}$$

where $n' = n - j + 1$. If no valid assignment is possible for the j-th column, then the procedure restarts from the beginning with the first column (and sets $\frac{1}{\mu(T_i)} = 0$ in (1) for this trial).[1] Sampling from the above distribution over assignments for column j can be done efficiently by dynamic programming.

We now state our negative result. This is a simple family of examples where the SIS algorithm will grossly underestimate $|\Omega|$ unless the number of trials t is exponentially large. Our examples will have the form $(1, 1, \ldots, 1, d_r)$ for row sums and $(1, 1, \ldots, 1, d_c)$ for column sums, where the number of rows is $m + 1$, the number of columns is $n + 1$, and we require that $m + d_r = n + d_c$.

Theorem 1. *Let $\beta > 0, \gamma \in (0, 1)$ be constants satisfying $\beta \neq \gamma$ and consider the input instances $r = (1, 1, \ldots, 1, \lfloor \beta m \rfloor)$, $c = (1, 1, \ldots, 1, \lfloor \gamma m \rfloor)$ with $m + 1$ rows. Fix any order of columns (or rows, if sequential importance sampling constructs tables row-by-row) and let X_t be the random variable representing the estimate of the SIS procedure after t trials of the algorithm. There exist constants $s_1 \in (0, 1)$ and $s_2 > 1$ such that for every sufficiently large m and for any $t \leq s_2^m$,*

$$\mathbf{Pr}\left(X_t \geq \frac{|\Omega_{r,c}|}{s_2^m}\right) \leq 3 s_1^m.$$

In contrast, note that there are MCMC algorithms which provably run in time polynomial in n and m for *any* row/column sums. In particular, the algorithm of Jerrum, Sinclair, and Vigoda [6] for the permanent of a non-negative matrix yields as a corollary a polynomial time algorithm for any row/column sums. The fastest algorithm for the permanent of a $n \times n$ matrix requires $O(n^7 \log^4 n)$ time [3], which implies a running time of $O((nm)^7 \log^4 n)$ time for binary contingency tables. More recently, Bezáková, Bhatnagar and Vigoda [1] have presented a related MCMC algorithm that works directly with binary contingency tables and has an improved polynomial running time. Their algorithm runs in time $O((nm)^2 R^3 s_{\max} \log^5(n + m))$ where $R = \sum_i r_i$ is the sum of the row sums and $s_{\max}$ is the maximum row and column sum. We note that, in addition to being formally asymptotically faster than any exponential time algorithm, a polynomial time algorithm has additional theoretical significance in that it (and its analysis) implies non-trivial insight into the the structure of the problem.

Some caveats are in order here. Firstly, the above results imply only that MCMC outperforms SIS asymptotically *in the worst case*; for many inputs, SIS may well be much more efficient. Secondly, the rigorous worst case upper bounds on the running time of the above MCMC algorithms are still far from practical. Chen et al. [4] showed several examples where SIS outperforms MCMC methods. We present a more systematic experimental study of the performance of SIS, focusing on examples where all the row and column sums are identical as well as on the "bad" examples from Theorem 1. Our experiments suggest that SIS

[1] Chen et al. devised a more subtle procedure which guarantees that there will always be a suitable assignment of every column. We do not describe this interesting modification of the procedure, as the two procedures are equivalent for the input instances which we discuss in this paper.

is extremely fast on the balanced examples, while its performance on the bad examples confirms our theoretical analysis.

We begin in Section 2 by presenting a few basic lemmas that are used in the analysis of our negative example. In Section 3 we present our main example where SIS is off by an exponential factor, thus proving Theorem 1. Finally, in Section 4 we present some experimental results for SIS that support our theoretical analysis.

2 Preliminaries

We will continue to let $\mu(T)$ denote the probability that a table $T \in \Omega_{r,c}$ is generated by sequential importance sampling algorithm. We let $\pi(T)$ denote the uniform distribution over Ω, which is the desired distribution.

Before beginning our main proofs we present two straightforward technical lemmas which are used at the end of the proof of the main theorem. The first lemma claims that if a large set of binary contingency tables gets a very small probability under SIS, then SIS is likely to output an estimate which is not much bigger than the size of the complement of this set, and hence very small. Let $\overline{\mathcal{S}} = \Omega_{r,c} \setminus \mathcal{S}$.

Lemma 1. *Let $p \leq 1/2$ and let $\mathcal{S} \subseteq \Omega_{r,c}$ be such that $\mu(\mathcal{S}) \leq p$. Then for any $a > 1$, and any t, we have*

$$\mathbf{Pr}\left(X_t \leq a\pi(\overline{\mathcal{S}})|\Omega|\right) \geq 1 - 2pt - 1/a.$$

Proof. The probability that all t SIS trials are not in $\mathcal{S}$ is at least

$$(1-p)^t > e^{-2pt} \geq 1 - 2pt,$$

where the first inequality follows from $\ln(1-x) > -2x$, valid for $0 < x \leq 1/2$, and the second inequality is the standard $e^{-x} \geq 1 - x$ for $x \geq 0$.

Let $T_1, \ldots, T_t$ be the t tables constructed by SIS. Then, with probability $> 1 - 2pt$, we have $T_i \in \overline{\mathcal{S}}$ for all i. Notice that for a table T constructed by SIS from $\overline{\mathcal{S}}$, we have

$$\mathbf{E}\left(\frac{1}{\mu(T)} \mid T \in \overline{\mathcal{S}}\right) = |\overline{\mathcal{S}}|.$$

Let $\mathcal{F}$ denote the event that $T_i \in \overline{\mathcal{S}}$ for all i, $1 \leq i \leq t$; hence,

$$\mathbf{E}\left(X_t \mid \mathcal{F}\right) = |\overline{\mathcal{S}}|.$$

We can use Markov's inequality to estimate the probability that SIS returns an answer which is more than a factor of a worse than the expected value, conditioned on the fact that no SIS trial is from $\mathcal{S}$:

$$\mathbf{Pr}\left(X > a|\overline{\mathcal{S}}| \mid \mathcal{F}\right) \leq \frac{1}{a}.$$

Finally, removing the conditioning we get:

$$\begin{aligned}\Pr(X \leq a|\overline{\mathcal{S}}|) &\geq \mathbf{Pr}\left(X \leq a|\overline{\mathcal{S}}| \mid \mathcal{F}\right)\mathbf{Pr}\left(\mathcal{F}\right)\\ &\geq \left(1-\frac{1}{a}\right)(1-2pt)\\ &\geq 1-2pt-\frac{1}{a}.\end{aligned}$$

The second technical lemma shows that if in a row with large sum (linear in m) there exists a large number of columns (again linear in m) for which the SIS probability of placing a 1 at the corresponding position differs significantly from the correct probability, then in any subexponential number of trials the SIS estimator will very likely exponentially underestimate the correct answer.

Lemma 2. *Let $\alpha < \beta$ be positive constants. Consider a class of instances of the binary contingency tables problem, parameterized by m, with $m+1$ row sums, the last of which is $\lfloor \beta m \rfloor$. Let $\mathcal{A}_i$ denote the set of all valid assignments of $0/1$ to columns $1, \ldots, i$. Suppose that there exist constants $f < g$ and a set I of cardinality $\lfloor \alpha m \rfloor$ such that one of the following statements is true:*

(i) for every $i \in I$ and any $A \in \mathcal{A}_{i-1}$,

$$\pi(A_{m+1,i} = 1 \mid A) \leq f < g \leq \mu(A_{m+1,i} = 1 \mid A),$$

(ii) for every $i \in I$ and any $A \in \mathcal{A}_{i-1}$,

$$\mu(A_{m+1,i} = 1 \mid A) \leq f < g \leq \pi(A_{m+1,i} = 1 \mid A).$$

Then there exists a constant $b_1 \in (0,1)$ such that for any constant $1 < b_2 < 1/b_1$ and any sufficiently large m, for any $t \leq b_2^m$,

$$\mathbf{Pr}\left(X_t \geq \frac{|\Omega_{r,c}|}{b_2^m}\right) \leq 3(b_1 b_2)^m.$$

In words, in b_2^m trials of sequential importance sampling, with probability at least $1-3(b_1b_2)^m$ the output is a number which is at most a b_2^{-m} fraction of the total number of corresponding binary contingency tables.

Proof. We will analyze case (i); the other case follows from analogous arguments. Consider indicator random variables U_i representing the event that the uniform distribution places 1 in the last row of the i-th column. Similarly, let V_i be the corresponding indicator variable for the SIS. The random variable U_i is dependent on U_j for $j < i$ and V_i is dependent on V_j for $j < i$. However, each U_i is stochastically dominated by U_i' which has value 1 with probability f, and each V_i stochastically dominates the random variable V_i' which takes value 1 with probability g. Moreover, the U_i' and V_i' are respectively i.i.d.

Now we may use the Chernoff bound. Let $k = \lfloor \alpha m \rfloor$. Then

$$\mathbf{Pr}\left(\sum_{i \in I} U_i' - kf \geq \frac{g-f}{2}k\right) < e^{-(g-f)^2 k/8}$$

and

$$\mathbf{Pr}\left(kg - \sum_{i\in I} V'_i \geq \frac{g-f}{2}k\right) < e^{-(g-f)^2 k/8}.$$

Let S be the set of all tables which have less than $kf+(g-f)k/2 = kg-(g-f)k/2$ ones in the last row of the columns in I. Let $b_1 := e^{-(g-f)^2\alpha/16} \in (0,1)$. Then $e^{-(g-f)^2k/8} \leq b_1^m$ for $m \geq 1/\alpha$. Thus, by the first inequality, under uniform distribution over all binary contingency tables the probability of the set S is at least $1-b_1^m$. However, by the second inequality, SIS constructs a table from the set S with probability at most b_1^m.

We are ready to use Lemma 1 with $\mathcal{S} = S$ and $p = b_1^m$. Since under uniform distribution the probability of S is at least $1-b_1^m$, we have that $|\mathcal{S}| \geq (1-b_1^m)|\Omega_{r,c}|$. Let $b_2 \in (1, 1/b_1)$ be any constant and consider $t \leq b_2^m$ SIS trials. Let $a = (b_1b_2)^{-m}$. Then, by Lemma 1, with probability at least $1-2pt-1/a \geq 1-3(b_1b_2)^m$ the SIS procedure outputs a value which is at most an $ab_1^m = b_2^{-m}$ fraction of $|\Omega_{r,c}|$.

3 Proof of Main Theorem

In this section we prove Theorem 1. Before we analyze the input instances from Theorem 1, we first consider the following simpler class of inputs.

3.1 Row sums $(1,1,\dots,1,d)$ and column sums $(1,1,\dots,1)$

The row sums are $(1,\dots,1,d)$ and the number of rows is $m+1$. The column sums are $(1,\dots,1)$ and the number of columns is $n = m+d$. We assume that sequential importance sampling constructs the tables column-by-column. Note that if SIS constructed the tables row-by-row, starting with the row with sum d, then it would in fact output the correct number of tables exactly. However, in the next subsection we will use this simplified case as a tool in our analysis of the input instances $(1,\dots,1,d_r)$, $(1,\dots,1,d_c)$, for which SIS must necessarily fail regardless of whether it works row-by-row or column-by-column, and regardless of the order it chooses.

Lemma 3. *Let $\beta > 0$, and consider an input of the form $(1,\dots,1,\lfloor\beta m\rfloor)$, $(1,\dots,1)$ with $m+1$ rows. Then there exist constants $s_1 \in (0,1)$ and $s_2 > 1$, such that for any sufficiently large m, with probability at least $1-3s_1^m$, column-wise sequential importance sampling with s_2^m trials outputs an estimate which is at most a s_2^{-m} fraction of the total number of corresponding binary contingency tables. Formally, for any $t \leq s_2^m$,*

$$\mathbf{Pr}\left(X_t \geq \frac{|\Omega_{r,c}|}{s_2^m}\right) \leq 3s_1^m.$$

The idea for the proof of the lemma is straightforward. By the symmetry of the column sums, for large m and d and $\alpha \in (0,1)$ a uniform random table will have about αd ones in the first αn cells of the last row, with high probability. We will show that for some $\alpha \in (0,1)$ and $d = \beta m$, sequential importance sampling is very unlikely to put this many ones in the first αn columns of the last row. Therefore, since with high probability sequential importance sampling will not construct any table from a set that is a large fraction of all legal tables, it will likely drastically underestimate the number of tables.

Before we prove the lemma, let us first compare the column distributions arising from the uniform distribution over all binary contingency tables with the SIS distributions. We refer to the column distributions induced by the uniform distribution over all tables as the *true* distributions. The true probability of 1 in the first column and last row can be computed as the number of tables with 1 at this position divided by the total number of tables. For this particular sequence, the total number of tables is $Z(m,d) = \binom{n}{d}m! = \binom{m+d}{d}m!$, since a table is uniquely specified by the positions of ones in the last row and the permutation matrix in the remaining rows and corresponding columns. Therefore,

$$\pi(A_{m+1,1} = 1) = \frac{Z(m,d-1)}{Z(m,d)} = \frac{\binom{m+d-1}{d-1}m!}{\binom{m+d}{d}m!} = \frac{d}{m+d}.$$

On the other hand, by the definition of sequential importance sampling, $\Pr(A_{i,1} = 1) \propto r_i/(n - r_i)$, where r_i is the row sum in the i-th row. Therefore,

$$\mu(A_{m+1,1} = 1) = \frac{\frac{d}{n-d}}{\frac{d}{n-d} + m\frac{1}{n-1}} = \frac{d(m+d-1)}{d(m+d-1)+m^2}.$$

Observe that if $d \approx \beta m$ for some constant $\beta > 0$, then for sufficiently large m we have

$$\mu(A_{m+1,1} = 1) > \pi(A_{m+1,1} = 1).$$

As we will see, this will be true for a linear number of columns, which turns out to be enough to prove that in polynomial time sequential importance sampling exponentially underestimates the total number of binary contingency tables with high probability.

Proof (Proof of Lemma 3). We will find a constant α such that for every column $i < \alpha m$ we will be able to derive an upper bound on the true probability and a lower bound on the SIS probability of 1 appearing at the $(m+1,i)$ position.

For a partially filled table with columns $1,\dots,i-1$ assigned, let d_i be the remaining sum in the last row and let m_i be the number of other rows with remaining row sum 1. Then the true probability of 1 in the i-th column and last row can be bounded as

$$\pi(A_{m+1,i} = 1 \mid A_{(m+1)\times(i-1)}) = \frac{d_i}{m_i + d_i} \le \frac{d}{m+d-i} =: f(d,m,i),$$

while the probability under SIS can be bounded as

$$\mu(A_{m+1,i} = 1 \mid A_{(m+1)\times(i-1)}) = \frac{d_i(m_i + d_i - 1)}{d_i(m_i + d_i - 1) + m_i^2}$$
$$\geq \frac{(d-i)(m+d-i-1)}{d(m+d-1)+m^2} =: g(d,m,i).$$

Observe that for fixed m, d, the function f is increasing and the function g is decreasing in i, for $i < d$.

Recall that we are considering a family of input instances parameterized by m with $d = \lceil \beta m \rceil$, for a fixed $\beta > 0$. We will consider $i < \alpha m$ for some $\alpha \in (0, \beta)$. Let

$$f^\infty(\alpha, \beta) := \lim_{m\to\infty} f(d, m, \alpha m) = \frac{\beta}{1+\beta-\alpha}; \tag{3}$$

$$g^\infty(\alpha, \beta) := \lim_{m\to\infty} g(d, m, \alpha m) = \frac{(\beta-\alpha)(1+\beta-\alpha)}{\beta(1+\beta)+1}; \tag{4}$$

$$\triangle_\beta := g^\infty(0,\beta) - f^\infty(0,\beta) = \frac{\beta^2}{(1+\beta)(\beta(1+\beta)+1)} > 0, \tag{5}$$

and recall that for fixed β, f^∞ is increasing in α and g^∞ is decreasing in α, for $\alpha < \beta$. Let $\alpha < \beta$ be such that $g^\infty(\alpha,\beta) - f^\infty(\alpha,\beta) = \triangle_\beta/2$. Such an α exists by continuity and the fact that $g^\infty(\beta,\beta) < f^\infty(\beta,\beta)$.

By the above, for any $\epsilon > 0$ and sufficiently large m, and for any $i < \alpha m$, the true probability is upper-bounded by $f^\infty(\alpha,\beta) + \epsilon$ and the SIS probability is lower-bounded by $g^\infty(\alpha,\beta) - \epsilon$. For our purposes it is enough to fix $\epsilon = \triangle_\beta/8$. Now we can use Lemma 2 with α and β defined as above, $f = f^\infty(\alpha,\beta) + \epsilon$ and $g = g^\infty(\alpha,\beta) - \epsilon$ (notice that all these constants depend only on β), and $I = \{1, \ldots, \lfloor \alpha m \rfloor\}$. This finishes the proof of the lemma with $s_1 = b_1 b_2$ and $s_2 = b_2$.

Note 1. Notice that every contingency table with row sums $(1, 1, \ldots, 1, d)$ and column sums $(1, 1, \ldots, 1)$ is binary. Thus, this instance proves that the column-based SIS procedure for general (non-binary) contingency tables has the same flaw as the binary SIS procedure. We expect that the negative example used for Theorem 1 also extends to general (i. e., non-binary) contingency tables, but the analysis becomes more cumbersome.

3.2 Proof of Theorem 1

Recall that we are working with row sums $(1, 1, \ldots, 1, d_r)$, where the number of rows is $m + 1$, and column sums $(1, 1, \ldots, 1, d_c)$, where the number of columns is $n + 1 = m + 1 + d_r - d_c$. We will eventually fix $d_r = \lfloor \beta m \rfloor$ and $d_c = \lfloor \gamma m \rfloor$, but to simplify our expressions we work with d_r and d_c for now.

The theorem claims that the SIS procedure fails for an arbitrary order of columns with high probability. We first analyze the case when the SIS procedure starts with columns of sum 1; we shall address the issue of arbitrary column order later. As before, under the assumption that the first column has sum 1, we compute the probabilities of 1 being in the last row for uniform random tables and for SIS respectively. For the true probability, the total number of tables can be computed as $\binom{m}{d_c}\binom{n}{d_r}(m-d_c)! + \binom{m}{d_c-1}\binom{n}{d_r-1}(m-d_c+1)!$, since a table is uniquely determined by the positions of ones in the d_c column and d_r row and a permutation matrix on the remaining rows and columns. Thus we have

$$\begin{aligned}\pi(A_{(m+1),1}) &= \frac{\binom{m}{d_c}\binom{n-1}{d_r-1}(m-d_c)! + \binom{m}{d_c-1}\binom{n-1}{d_r-2}(m-d_c+1)!}{\binom{m}{d_c}\binom{n}{d_r}(m-d_c)! + \binom{m}{d_c-1}\binom{n}{d_r-1}(m-d_c+1)!} \\ &= \frac{d_r(n-d_r+1)+d_c d_r(d_r-1)}{n(n-d_r+1)+n d_c d_r} =: f_2(m,d_r,d_c);\end{aligned}$$

$$\mu(A_{(m+1),1}) = \frac{\frac{d_r}{n-d_r}}{\frac{d_r}{n-d_r}+m\frac{1}{n-1}} = \frac{d_r(n-1)}{d_r(n-1)+m(n-d_r)} =: g_2(m,d_r,d_c).$$

Let $d_r = \lfloor \beta m \rfloor$ and $d_c = \lfloor \gamma m \rfloor$ for some constants $\beta > 0, \gamma \in (0,1)$ (notice that this choice guarantees that $n \geq d_r$ and $m \geq d_c$, as required). Then, as m tends to infinity, f_2 approaches

$$f_2^{\infty}(\beta,\gamma) := \frac{\beta}{1+\beta-\gamma},$$

and g_2 approaches

$$g_2^{\infty}(\beta,\gamma) := \frac{\beta(1+\beta-\gamma)}{\beta(1+\beta-\gamma)+1-\gamma}.$$

Notice that $f_2^{\infty}(\beta,\gamma) = g_2^{\infty}(\beta,\gamma)$ if and only if $\beta = \gamma$. Suppose that $f_2^{\infty}(\beta,\gamma) < g_2^{\infty}(\beta,\gamma)$ (the opposite case follows analogous arguments and uses the second part of Lemma 2). As in the proof of Lemma 3, we can define α such that if the importance sampling does not choose the column with sum d_c in its first αm choices, then in any subexponential number of trials it will exponentially underestimate the total number of tables with high probability. Formally, we derive an upper bound on the true probability of 1 being in the last row of the i-th column, and a lower bound on the SIS probability of the same event (both conditioned on the fact that the d_c column is not among the first i columns assigned). Let $d_r^{(i)}$ be the current residual sum in the last row, m_i be the number of rows with sum 1, and n_i the remaining number of columns with sum 1. Notice that $n_i = n-i+1$, $m \geq m_i \geq m-i+1$, and $d_r \geq d_r^{(i)} \geq d_r - i + 1$. Then

$$\begin{aligned}\pi(A_{(m+1),i} \mid A_{(m+1)\times(i-1)}) &= \frac{d_r^{(i)}(n_i - d_r^{(i)}+1) + d_c d_r^{(i)}(d_r^{(i)}-1)}{n_i(n_i-d_r^{(i)}+1)+n_i d_c d_r^{(i)}} \\ &\leq \frac{d_r(n-d_r+1)+d_c d_r^2}{(n-i)(n-i-d_r)+(n-i)d_c(d_r-i)} \\ &=: f_3(m,d_r,d_c,i);\end{aligned}$$

$$\mu(A_{(m+1),i} \mid A_{(m+1)\times(i-1)}) = \frac{d_r^{(i)}(n_i-1)}{d_r^{(i)}(n_i-1)+m_i(n_i-d_r^{(i)})}$$
$$\geq \frac{(d_r-i)(n-i)}{d_r n + m(n-d_r)} =: g_3(m,d_r,d_c,i).$$

As before, notice that if we fix $m, d_r, d_c > 0$ satisfying $d_c < m$ and $d_r < n$, then f_3 is an increasing function and g_3 is a decreasing function in i, for $i < \min\{n - d_r, d_r\}$. Recall that $n - d_r = m - d_c$. Suppose that $i \leq \alpha m < \min\{m - d_c, d_r\}$ for some α which we specify shortly. Thus, the upper bound on f_3 in this range of i is $f_3(m, d_r, d_c, \alpha m)$ and the lower bound on g_3 is $g_3(m, d_r, d_c, \alpha m)$. If $d_r = \lfloor \beta m \rfloor$ and $d_c = \lfloor \gamma m \rfloor$, then the upper bound on f_3 converges to

$$f_3^\infty(\alpha,\beta,\gamma) := \lim_{m\to\infty} f_3(m,d_r,d_c,\alpha m) = \frac{\beta^2}{(1+\beta-\gamma-\alpha)(\beta-\alpha)}$$

and the lower bound on g_3 converges to

$$g_3^\infty(\alpha,\beta,\gamma) := \lim_{m\to\infty} g_3(m,d_r,d_c,\alpha m) = \frac{(\beta-\alpha)(1+\beta-\gamma-\alpha)}{\beta(1+\beta-\gamma)+1-\gamma}$$

Let

$$\triangle_{\beta,\gamma} := g_3^\infty(0,\beta,\gamma) - f_3^\infty(0,\beta,\gamma) = g_2^\infty(\beta,\gamma) - f_2^\infty(\beta,\gamma) > 0.$$

We set α to satisfy $g_3^\infty(\alpha,\beta,\gamma) - f_3^\infty(\alpha,\beta,\gamma) \geq \triangle_{\beta,\gamma}/2$ and $\alpha < \min\{1-\gamma, \beta\}$. Now we can conclude this part of the proof identically to the last paragraph of the proof of Lemma 3.

It remains to deal with the case when sequential importance sampling picks the d_c column within the first $\lfloor \alpha m \rfloor$ columns. Suppose d_c appears as the k-th column. In this case we focus on the subtable consisting of the last $n+1-k$ columns with sum 1, m' rows with sum 1, and one row with sum d', an instance of the form $(1,1,\dots,1,d'), (1,\dots,1)$. We will use arguments similar to the proof of Lemma 3.

First we express d' as a function of m'. We have the bounds $(1-\alpha)m \leq m' \leq m$ and $d - \alpha m \leq d' \leq d$ where $d = \lfloor \beta m \rfloor \geq \beta m - 1$. Let $d' = \beta' m'$. Thus, $\beta - \alpha - 1/m \leq \beta' \leq \beta/(1-\alpha)$.

Now we find α' such that for any $i \leq \alpha' m'$ we will be able to derive an upper bound on the true probability and a lower bound on the SIS probability of 1 appearing at position $(m'+1, i)$ of the $(n+1-k) \times m'$ subtable, no matter how the first k columns were assigned. In order to do this, we might need to decrease α – recall that we are free to do so as long as α is a constant independent of m. By the derivation in the proof of Lemma 3 (see expressions (3) and (4)), as m' (and thus also m) tends to infinity, the upper bound on the true probability approaches

$$f^\infty(\alpha',\beta') = \lim_{m\to\infty} \frac{\beta'}{1+\beta'-\alpha'}$$
$$\leq \lim_{m\to\infty} \frac{\frac{\beta}{1-\alpha}}{1+\beta-\alpha-\frac{1}{m}-\alpha'} = \frac{\frac{\beta}{1-\alpha}}{1+\beta-\alpha-\alpha'} =: f_4^\infty(\alpha,\beta,\alpha')$$

and the lower bound on the SIS probability approaches

$$
\begin{aligned}
g^{\infty}(\alpha',\beta') &= \lim_{m\to\infty} \frac{(\beta'-\alpha')(1+\beta'-\alpha')}{\beta'(1+\beta')+1} \\
&\geq \lim_{m\to\infty} \frac{(\beta-\alpha-\frac{1}{m}-\alpha')(1+\beta-\alpha-\frac{1}{m}-\alpha')}{\frac{\beta}{1-\alpha}(1+\frac{\beta}{1-\alpha})+1} \\
&= \frac{(\beta-\alpha-\alpha')(1+\beta-\alpha-\alpha')}{\frac{\beta}{1-\alpha}(1+\frac{\beta}{1-\alpha})+1} \\
&\geq \frac{(\beta-\alpha-\alpha')(1+\beta-\alpha-\alpha')}{\frac{\beta}{1-\alpha}(\frac{1}{1-\alpha}+\frac{\beta}{1-\alpha})+\frac{1}{(1-\alpha)^2}} =: g_4^{\infty}(\alpha,\beta,\alpha'),
\end{aligned}
$$

where the last inequality holds as long as $\alpha < 1$. Notice that for fixed α, β satisfying $\alpha < \min\{1, \beta\}$, the function f_4^{∞} is increasing and g_4^{∞} is decreasing in α', for $\alpha' < \beta - \alpha$. Similarly, for fixed α', β satisfying $\alpha' < \beta$, the function f_4^{∞} is increasing and g_4^{∞} is decreasing in α, for $\alpha < \min\{1, \beta - \alpha'\}$. Therefore, if we take $\alpha = \alpha' < \min\{1, \beta/2\}$, we will have the bounds

$$
f_4^{\infty}(x,\beta,y) \leq f_4^{\infty}(\alpha,\beta,\alpha) \qquad \text{and} \qquad g_4^{\infty}(x,\beta,y) \geq g_4^{\infty}(\alpha,\beta,\alpha)
$$

for any $x, y \leq \alpha$. Recall that $\triangle_\beta = g^{\infty}(0,\beta) - f^{\infty}(0,\beta) = g_4^{\infty}(0,\beta,0) - f_4^{\infty}(0,\beta,0) > 0$. If we choose α so that $g_4^{\infty}(\alpha,\beta,\alpha) - f_4^{\infty}(\alpha,\beta,\alpha) \geq \triangle_\beta/2$, then in similar fashion to the last paragraph of the proof of Lemma 3, we may conclude that the SIS procedure likely fails. More precisely, let $\epsilon := \triangle_\beta/8$ and let $f := f_4^{\infty}(\alpha,\beta,\alpha)+\epsilon$ and $g := g_4^{\infty}(\alpha,\beta,\alpha)-\epsilon$ be the upper bound (for sufficiently large m) on the true probability and the lower bound on the SIS probability of 1 appearing at the position $(m+1, i)$ for $i \in I := \{k+1, \ldots, k+\lfloor\alpha' m'\rfloor\}$. Therefore Lemma 2 with parameters $\alpha(1-\alpha)$, β, I of size $|I| = \lfloor\alpha' m'\rfloor \geq \lfloor\alpha(1-\alpha)m\rfloor$, f, and g implies the statement of the theorem.

Finally, if the SIS procedure constructs the tables row-by-row instead of column-by-column, symmetrical arguments hold. This completes the proof of Theorem 1.

□

4 Experiments

We performed several experimental tests which show sequential importance sampling to be a promising approach for certain classes of input instances. We discuss the experiments in more detail and present supporting figures in the full version of this paper [2].

Our first set of experiments tested the SIS technique on regular input sequences, i.e., $r_i = c_j$ for all i, j. It appears the approach is very efficient for these input sequences. We considered input sequences which were 5, 10, $\lfloor 5\log n\rfloor$ and $\lfloor n/2\rfloor$-regular, and $n \times n$ matrices with $n = 10, 15, 20, \ldots, 100$. The required number of SIS trials until the algorithm converged resembled a linear function of n.

In contrast, we examined the evolution of SIS on the negative example from Theorem 1. In our simulations we used the more delicate sampling mentioned in footnote 1, which guarantees that the assignment in every column is valid, i. e., such an assignment can always be extended to a valid table (or, equivalently, the random variable X_t is always strictly positive). We ran the SIS algorithm under three different settings: first, we constructed the tables column-by-column where the columns were ordered in decreasing order of their sums, as suggested in the paper of Chen et al. [4]; second, we ordered the columns in increasing order of their sums; and third, we constructed the tables row-by-row where the rows were ordered in decreasing order of their sums.

The experiments confirmed the poor performance described in Theorem 1. For $m = 300$, $\beta = .6$ and $\gamma = .7$, even the best of the three estimators differed from the true value by about a factor of 40, while some estimates were off by more than a factor of 1000.

References

1. I. Bezáková, N. Bhatnagar, and E. Vigoda, Sampling Binary Contingency Tables with a Greedy Start. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2006.
2. I. Bezáková, A. Sinclair, D. Štefankovič, and E. Vigoda. Negative Examples for Sequential Importance Sampling of Binary Contingency Tables. Submitted. Available from Mathematics arXiv math.ST/0606650.
3. I. Bezáková, D. Štefankovič, V. Vazirani, and E. Vigoda. Accelerating Simulated Annealing for the Permanent and Combinatorial Counting Problems. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2006.
4. Y. Chen, P. Diaconis, S. Holmes, and J.S. Liu, Sequential Monte Carlo Methods for Statistical Analysis of Tables. *Journal of the American Statistical Association*, 100:109-120, 2005.
5. M. De Iorio, R. C. Griffiths, R. Lebois, and F. Rousset, Stepwise Mutation Likelihood Computation by Sequential Importance Sampling in Subdivided Population Models. *Theor. Popul. Biol.*, 68:41-53, 2005.
6. M.R. Jerrum, A. Sinclair and E. Vigoda, A Polynomial-time Approximation Algorithm for the Permanent of a Matrix with Non-negative Entries. *Journal of the Association for Computing Machinery*, 51(4):671-697, 2004.
7. J. Miguez and P. M. Djuric, Blind Equalization by Sequential Importance Sampling. *Proceedings of the IEEE International Symposium on Circuits and Systems*, 845-848, 2002.
8. J. L. Zhang and J. S. Liu, A New Sequential Importance Sampling Method and its Application to the Two-dimensional Hydrophobic-Hydrophilic Model. *J. Chem. Phys.*, 117:3492-8, 2002.