

Fast simulation of correlated photon-counting measurements with the Anscombe transformation

Shuting Yu^a, Joakim Andén^a, and Mats U. Persson^{b,c}

^aDepartment of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

^bDepartment of Physics, KTH Royal Institute of Technology, SE-10691, Stockholm, Sweden

^cMedTechLabs, BioClinicum, Karolinska University Hospital, Solna, Sweden

ABSTRACT

With the increasing popularity of photon-counting detectors in X-ray computed tomography, and with virtual clinical trials playing an increasingly important role in the evaluation of imaging systems, it is vital to have accurate simulation models for photon-counting imaging systems. Physical effects such as charge sharing, K-fluorescence and Compton scattering give rise to spatiotemporal correlations between different energy bins in different pixels, and existing methods for simulating such correlations accurately with multi-pixel correlation lengths are too computationally costly to be practical for full-size CT acquisition simulations.

In this work, we propose a fast, accurate method for simulating correlated Poisson noise in photon-counting detectors based on the Anscombe transformation. We use Cholesky factorization to generate correlated Gaussian random numbers and then apply the inverse Anscombe transformation to map these into approximately Poisson-distributed counts. We show that any desired correlation structure of the counts can be obtained by adjusting the mean and covariance matrix used for the Gaussian random number generation.

We evaluate this method in a simulation study with both a one-dimensional “toy” model with a single energy bin and a one-dimensional “realistic” model¹, by using the chi-square statistic to assess the accuracy of the marginal probability distributions and pairwise joint probability distributions. The computational speed is compared to brute-force generation of Poisson random numbers. Our results show that the proposed method achieves reasonable accuracy in approximating the Poisson distribution, with $> 80\%$ lower chi-squared than a Gaussian approximation and that it can decrease the computation time to 1% of the time required for direct Poisson generation.

Keywords: Photon-counting, detector simulation, Poisson statistics, Anscombe transformation

1. INTRODUCTION

Photon-counting detectors play an important role in medical imaging and are in particular becoming more and more commonly used in x-ray computed tomography (CT) scanners.² Accurate simulation models for photon-counting imaging detectors are important tools for development and evaluation of such imaging systems, and their importance is increasing in light of the increasing interest in virtual clinical trials, in which the diagnostic performance of imaging systems are assessed in silico using well-calibrated simulation models.

The number of photons detected at any given pixel is typically Poisson distributed, but cross-talk between detector elements, where one photon triggers counts in multiple pixels, can lead to correlations between pixels and energy bins. These correlations affect the noise characteristics of the reconstructed CT image, including the noise variance, the noise power spectrum, and the covariance between different spectral image channels. Including such correlations in the models is therefore crucial in order to achieve accurate simulations.

Correlated Poisson random variables can be modeled by direct simulation of Poisson random variables, but this requires generating one Poisson random number for each pair of pixels, representing the number of photons counted in both. This is computationally challenging for two-dimensional arrays with large numbers of pixels, and this is further exacerbated by the fact that detectors typically have multiple energy bins and in some cases

Send correspondence to Shuting Yu, yushuting516@163.com or Mats U. Persson, mats.persson@mi.physics.kth.se

multiple layers in the depth direction. Previous methods have been limited to simulating cross-talk between a few neighboring pixels,^{3,4} or based on a Gaussian approximation that does not accurately represent the Poisson nature of counts at low photon counts.⁵

The purpose of this work is to propose a new method for generating correlated Poisson random variables, which combines an accurate approximation of the true probability distribution of photon counts with high computational efficiency. This method is based on first generating multivariate Gaussian random variables, which can be done efficiently, and then applying the inverse Anscombe transformation to make the distribution similar to the Poisson distribution.

2. METHODS

The correlation structure between different pixels (and energy bins) in a detector is uniquely characterized by their covariance matrix. In real PCCT systems, the covariance matrix C of the recorded counts can be obtained either by repeated experimental measurement—scanning a uniform phantom or empty field and computing sample variances and covariances—or by first-principles physical modeling that accounts for photon arrival rates, detector efficiency, pulse pile-up, charge sharing, and related effects. This matrix C encodes both the per-pixel noise strength, $\text{Var}(Y_i) = C_{ii}$, and the pairwise correlations, $\text{Cov}(Y_i, Y_j) = C_{ij}$. To test our samplers under controlled but varied conditions, we construct $n \times n$ covariance matrices for the one-dimensional “toy” model that are guaranteed to be positive definite, nonnegative, and diagonally dominant.

For the toy model construction, we begin by sampling a random $n \times n$ matrix A with independent $\text{Uniform}(0, 1)$ entries and symmetrize it via $A \leftarrow \frac{1}{2}(A + A^T)$. We then add a multiple of the identity, $C = A + \text{scale} \cdot I_n$, where $\text{scale} \geq n$ ensures strict diagonal dominance and hence positive definiteness by Gershgorin’s theorem. The resulting C is entrywise nonnegative, symmetric, and positive definite.

For the realistic model¹, the silicon and tungsten sheets are 60 mm, and the autocovariance is a 41×1 matrix. We convert the autocovariance function directly into a covariance matrix with a size of either 4×4 or 21×21 , where the reduced size is obtained by neglecting longer-range correlations. This is normalized to counts per incident photons so we can multiply with some suitable number to study the effects of varying the incident fluence rate. This matrix is symmetric, positive-definite, and has non-negative entries, thereby serving as a realistic covariance structure for the subsequent evaluation.

Next, we tried four methods for simulating PCCT Poisson random variables. The first method realizes exact statistical modeling by decomposing the target covariance into independent and shared Poisson rate components, guaranteeing precise reproduction of the desired covariance as well as the probabilities dictated by the Poisson distribution. Specifically, the covariance matrix C is decomposed into a symmetric rate matrix R with off-diagonal entries $R_{ij} = C_{ij}$ ($i \neq j$) and diagonal entries $R_{ii} = C_{ii} - \sum_{k \neq i} C_{ik}$, ensuring $R_{ii} \geq 0$. Correlated Poisson counts Y_i are then generated as $Y_i = \sum_{j=1}^n Z_{ij}$, where $Z_{ij} \sim \text{Poisson}(R_{ij})$ and $Z_{ij} = Z_{ji}$ for all $i \neq j$, thereby preserving the covariance structure $\text{Cov}(Y_i, Y_j) = C_{ij}$. After repeating the process m times, the $\mathbf{Y} = [Y_1, \dots, Y_m]$ is the correlated Poisson random variables we want to get and the individual Y_i represents the total photon count detected at the i -th pixel (or the i -th energy channel).

The second method, Gaussian truncation method, samples from a multivariate normal distribution with mean equal to the diagonal of the covariance and variance equal to the covariance matrix, then truncates and rounds values to nonnegative integers. The third and fourth methods employ variance-stabilizing transforms: the Anscombe transform⁶ and the Freeman–Tukey transform,⁷ respectively. Both map Poisson variables into a Gaussian domain for multivariate normal sampling, then apply inverse transforms to recover integer counts. Moreover, we set the Anscombe transformation as $g_A(y) = 2\sqrt{y + \frac{3}{8}}$, $\mu_{A,i} = 2\sqrt{\lambda_i + \frac{3}{8}}$, $\alpha_i = (\lambda_i + 3/8)^{-1/2}$, $(C_A)_{ij} = \alpha_i \alpha_j C_{ij}$. Given the target mean vector $\boldsymbol{\lambda}$ and covariance matrix C (with $\text{diag}(C) = \boldsymbol{\lambda}$), we approximate $g_A(Y)$ for $Y \sim \text{Poisson}(\boldsymbol{\lambda})$ as Gaussian by matching mean $\mu_{A,i}$ and covariance $(C_A)_{ij}$, the α_i is the scaling factor. We then draw a sample $Z \sim \mathcal{N}(\boldsymbol{\mu}_A, C_A)$ and apply the (approximate) inverse Anscombe transform to recover integer counts. A simple inverse is $\hat{Y}_i = \max\left(0, \left\lfloor \left(\frac{Z_i}{2}\right)^2 - \frac{3}{8} + 0.5 \right\rfloor\right)$, i.e., compute $\left(\frac{Z_i}{2}\right)^2 - \frac{3}{8}$, round to the nearest integer, and truncate negatives to zero. For Freeman–Tukey transform, the transformation is defined as $g_{\text{FT}}(y) = \sqrt{y} + \sqrt{y+1}$. Given the target mean vector $\boldsymbol{\lambda}$ and covariance C (with $\text{diag}(C) = \boldsymbol{\lambda}$), we compute the

transformed mean $\mu_{\text{FT},i} = g_{\text{FT}}(\lambda_i)$ and derivative $\beta_i = g'_{\text{FT}}(\lambda_i) = \frac{1}{2\sqrt{\lambda_i}} + \frac{1}{2\sqrt{\lambda_i+1}}$. The transformed covariance is obtained via linear propagation: $(C_{\text{FT}})_{ij} = \beta_i\beta_j C_{ij}$. We then sample $Z \sim \mathcal{N}(\boldsymbol{\mu}_{\text{FT}}, C_{\text{FT}})$ and apply the approximate inverse $\hat{Y}_i = \max\left(0, \left\lfloor \left(\frac{Z_i^2-1}{2Z_i}\right)^2 + 0.5 \right\rfloor\right)$, which rounds $((Z_i^2-1)/(2Z_i))^2$ to the nearest integer and truncates negatives to zero. All three methods must round the continuous sampled values to the nearest integer and truncate to non-negative values, because photon counts are inherently discrete and non-negative physical observables. The truncated Gaussian method rounds the Gaussian samples directly, while the Anscombe and Freeman–Tukey methods round after applying their respective inverse transforms, thereby ensuring that the generated counts are both physically realistic and statistically consistent.


Finally, we compare the four methods by using 1D histogram visualization, marginal χ^2 statistics, and pairwise χ^2 statistics. We start by overlaying each method’s simulated count histogram on the theoretical Poisson curve to check how closely these line up, compute the χ^2 statistic for each pixel to quantify its deviation from the Poisson distribution, and compute another χ^2 statistic for pixel pairs to see how well each method reproduces the pairwise marginal distributions, verifying that the desired correlations are obtained.

3. RESULTS


To evaluate each method, for the toy model we use it to generate a large number m of samples, each of which consists of an 1D array of size n that measures the number of photons detected in each of the n pixels. For the realistic model, we convert the autocovariance function (ACF) to covariance matrix. For the covariance matrix, since this is normalized to counts per incident photons so we can multiply with some suitable number to study the effects of varying the incident fluence rate. Here, we use $5 \times$ covariance matrix with selected n and m . For each of those pixels, we then compute a histogram, measuring the number of times a specific number of photons is registered in that pixel. Figure 1 shows the comparison results for toy model when $n = 4$. Figure 2 shows the comparison results for realistic model when $n = 4$. From those plots, all 2 models and four methods visually track the target Poisson marginals closely, with only minor deviations. To quantify these small differences, we calculate the χ^2 statistics for each pixel, given by $\sum_{k=0}^{+\infty} (h_k - mp_k)^2 / mp_k$, where m is the number of samples, h_k is the histogram value for count k , and p_k is the probability of observing that count according to the marginal Poisson distribution of that pixel. The same procedure is performed using 2D histograms, where pairs of counts are recorded and compared with the bivariate Poisson marginal distribution. In both cases, a lower χ^2 statistic indicates closer agreement with the true Poisson model.

For convenience, Table 1 shows the average χ^2 statistics for the marginal and pairwise marginal distributions for toy model. Table 3 and Table 4 show the average χ^2 statistics for the marginal and pairwise marginal distributions for realistic model with different Poisson rate(λ). From the toy-model results, we observe that Method 1 (exact statistical decomposition) consistently achieves the smallest χ^2 values across all pixels and choices of n . This confirms that it reproduces the Poisson distribution most faithfully, as expected for an exact algorithm. In contrast, Method 2 (the truncated Gaussian method) yields by far the largest χ^2 values in every case, indicating the poorest fit to the Poisson statistics. Methods 3 and 4, based on the Anscombe and Freeman–Tukey variance-stabilizing transforms respectively, exhibit intermediate performance. Although both methods are designed to stabilize the variance of Poisson-distributed data, our results demonstrate that the Anscombe-based sampler (Method 3) consistently outperforms the Freeman–Tukey variant (Method 4) in the toy-model setting. The results obtained from the realistic model model largely agree with those from the toy model, with one notable exception. According to figure 3, At low Poisson rates, Method 2 produces more accurate estimates than Method 4, whereas the opposite behavior is observed at higher Poisson rates. Specifically, when the Poisson rate is below approximately 2.5, Method 2 outperforms Method 4. Around a Poisson rate of 2.5, the two methods exhibit comparable accuracy, while for higher Poisson rates, Method 4 becomes more accurate than Method 2.


Table 1, Table 3 and Table 4 also show the joint χ^2 statistics for every pixel pair under different n . As with the marginal evaluation, for larger poisson rate, method 1 (exact statistical decomposition) consistently attains the smallest χ^2 values—close to zero for all pairs—confirming that it reproduces both marginal and cross-pixel correlations with near-perfect fidelity. Methods 3 (Anscombe) and 4 (Freeman–Tukey) deliver moderate χ^2 statistics (in the low- 10^2 to mid- 10^2 range), indicating small but measurable deviations in the joint law; notably,




Method 1: Exact statistical decomposition



Method 2: Truncated Gaussian




Method 4: Freeman-Tukey transformation




Method 3: Anscombe transformation


Figure 1: Toy model: Histograms for pixel 1 in the $n = 4$ case ($m = 10^6$).




Method 1: Exact statistical decomposition



Method 2: Truncated Gaussian



Method 4: Freeman-Tukey transformation



Method 3: Anscombe transformation

Figure 2: Realistic model: Histograms for pixel 1 in the $n = 4$ case ($m = 10^6$).




Figure 3: Realistic model: Comparison of Method 2 Truncated Gaussian and Method 4 Freeman-Tukey

Anscombe is uniformly slightly closer to the ideal than Freeman–Tukey. Method 2 (truncated Gaussian) again shows very large χ^2 statistics (on the order of 10^4), demonstrating its inability to capture the correct Poisson structure. These results mirror the findings for the marginal χ^2 statistics and further reinforce the conclusion that exact decomposition is the most accurate, followed by Anscombe, Freeman–Tukey, and finally the Gaussian approximation.

Table 2 and Table 5 reports total-sample times (s) for the four methods at different covariance dimensions with $N_{\text{SAMPLES}} = 10^6$. Method 1 is much slower and grows steeply with dimension; Methods 2-4 remain far faster and are only slightly different from each other.

Across all our evaluations in both toy model and realistic model model, the exact statistical decomposition (method 1) unsurprisingly delivers the closest match to the true Poisson model, with χ^2 values near zero for both marginal and joint tests. For larger poisson rate λ ($\lambda > 2.5$), the truncated Gaussian (method 2) consistently fares worst, exhibiting very large χ^2 errors. The transform-based approaches (methods 3 and 4) occupy the middle ground: both greatly improve over the Gaussian approximation, with the Anscombe-based sampler (method 3) slightly outperforming the Freeman–Tukey variant (Method 4). Overall, these results highlight a clear accuracy-vs.-speed trade-off among the four techniques.

Table 1: Toy model: Average marginal and pairwise joint χ^2 statistics across all pixels / pixel pairs ($m = 10^6$).

n	Exact statistical		Truncated Gaussian		Anscombe		Freeman-Tukey	
	marginal	joint	marginal	joint	marginal	joint	marginal	joint
3	19.77	173.87	38 675.43	78 652.32	6 860.20	14 931.03	22 568.22	39 348.79
4	22.14	156.35	38 452.37	78 287.85	6 838.26	14 626.60	22 596.78	39 189.41
5	29.22	181.72	38 741.56	78 918.63	6 821.33	14 444.56	22 768.12	39 861.64

Table 2: Toy model: Total-sample times (s) for $m = 10^6$ at covariance dimensions.

n	Exact statistical	Truncated Gaussian	Anscombe	Freeman-Tukey
3	4.119	0.060	0.061	0.077
4	5.924	0.078	0.082	0.090
5	8.249	0.096	0.101	0.107

Table 3: Realistic model ($\lambda = 0.82$): Average marginal and pairwise joint χ^2 statistics across all pixels / pixel pairs ($m = 10^6$).

n	λ	Exact statistical		Truncated	Gaussian	Anscombe		Freeman-Tukey	
		marginal	joint	marginal	joint	marginal	joint	marginal	joint
4	0.82	9.09	33.16	4533.73	9176.07	1194.01	2329.11	33533.61	65653.47
21	0.82	7.39	38.80	4511.53	9224.08	1214.33	2426.87	35893.84	67413.13

Table 4: Realistic model 2 ($\lambda = 4.10$): Average marginal and pairwise joint χ^2 statistics across all pixels / pixel pairs ($m = 10^6$).

n	λ	Exact statistical		Truncated	Gaussian	Anscombe		Freeman-Tukey	
		marginal	joint	marginal	joint	marginal	joint	marginal	joint
4	4.10	13.34	184.01	44126.48	90475.69	6607.92	13104.47	19242.57	36544.53
21	4.10	19.99	206.48	44379.02	90654.64	6522.09	12960.23	19182.68	37372.16

Table 5: Realistic model: Total-sample times (s) for $m = 10^6$ at covariance dimensions.

n	Exact statistical	Truncated Gaussian	Anscombe	Freeman-Tukey
21	104.770	0.613	0.541	0.576
4	5.811	0.079	0.085	0.103

4. CONCLUSIONS

In this work, we have for the first time presented a new, computationally efficient method for generate correlated Poisson random variables able to model correlations in photon-counting CT detectors. Our results show that the proposed method using the inverse Anscombe transform can generate correlated random numbers with a distribution closer to the Poisson distribution compared to both the truncated Gaussian and Freeman–Tukey methods while being significantly faster to compute than generating raw Poisson random numbers. Future work will include extending this work to realistic detector models with long-range correlations, and investigating the accuracy of the model for image-domain noise metrics. This work can have important implications for improving simulation models for photon-counting x-ray imaging systems.

ACKNOWLEDGMENTS

This work received financial support from MedTechLabs, The Swedish Research Council (grants 2021-05103 and 2023-04143) and GE HealthCare. Mats U. Persson discloses research collaboration and licence fees, GE HealthCare.

REFERENCES

- [1] Persson, M., Wang, A., and Pelc, N. J., “Detective quantum efficiency of photon-counting cdte and si detectors for computed tomography: a simulation study,” *Journal of Medical Imaging* **7**, 043501 (July 2020). Open Access.
- [2] Danielsson, M., Persson, M., and Sjölin, M., “Photon-counting x-ray detectors for CT,” *Phys. Med. Biol.* **66**, 03TR01 (1 2021).
- [3] Taguchi, K., Polster, C., Lee, O., Stierstorfer, K., and Kappler, S., “Spatio-energetic cross talk in photon counting detectors: Detector model and correlated poisson data generator,” *Medical Physics* **43**(12), 6386–6404 (2016).
- [4] Faby, S., Maier, J., Sawall, S., Simons, D., Schlemmer, H.-P., Lell, M., and Kachelrieß, M., “An efficient computational approach to model statistical correlations in photon counting x-ray detectors,” *Medical Physics* **43**(7), 3945–3960 (2016).

- [5] Grönberg, F., Yin, Z., Maltz, J. S., Pelc, N. J., and Persson, M., “The effects of intra-detector compton scatter on low-frequency dqe for photon-counting ct using edge-on-irradiated silicon detectors,” *Med. Phys.* **51**(7), 4948–4969 (2024).
- [6] Anscombe, F. J., “The transformation of Poisson, binomial and negative-binomial data,” *Biometrika* **35**(3/4), 246–254 (1948).
- [7] Freeman, M. F. and Tukey, J. W., “Transformations related to the angular and the square root,” *The Annals of Mathematical Statistics* **21**(4), 607–611 (1950).