

Sliced Optimal Transport Sampling: *Supplemental Material*

LOIS PAULIN, Univ Lyon, CNRS
NICOLAS BONNEEL, Univ Lyon, CNRS
DAVID COEURJOLLY, Univ Lyon, CNRS
JEAN-CLAUDE IEHL, Univ Lyon, CNRS
ANTOINE WEBANCK, Univ Lyon, CNRS
MATHIEU DESBRUN, ShanghaiTech/Caltech
VICTOR OSTROMOUKHOV, Univ Lyon, CNRS

This supplemental material provides additional results to complement our paper from ACM SIGGRAPH 2020.

1 SPECTRAL ANALYSIS IN DIMENSION 4

Figure 1 illustrates spectral properties of 2D projections of various samplers in dimension 4. As illustrated in dimension 6 in the paper, our Projective SOT exhibits better Fourier spectrum than the other alternatives, even those designed to have projective distributions such as the Projective Blue-Noise sampler of [RRSG16].

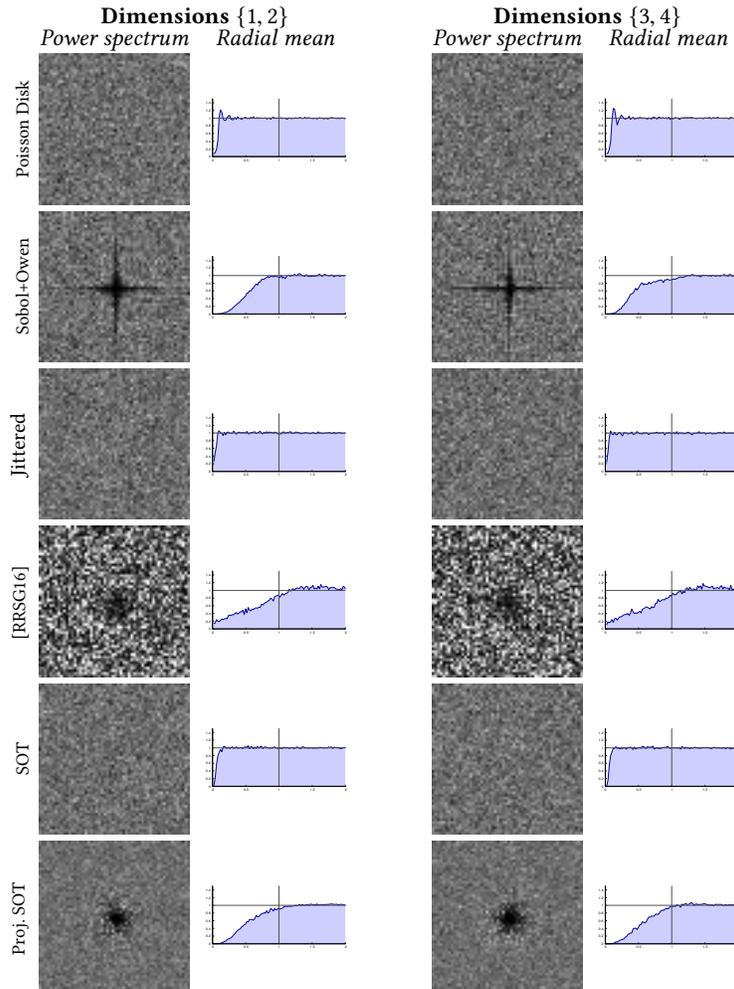


Fig. 1. **Fourier spectra on projections.** In dimension 4, we project various pointsets of 8k samples generated with our projective SOT sampler onto the first two and the last two dimensions, in order to evaluate their projective equi-distribution via 2D power spectrum and radial power spectrum.

2 INTEGRATION ACCURACY IN DIMENSION 6: EXTRA SAMPLERS

We also complement the integration error graphs of the main paper with convergence results of Monte Carlo integration with extra samplers for completeness (same experimental settings as our plots in dimension 6): Fig. 2 shows Projective Blue-Noise [RRSG16], Spoke-Darts [MEA⁺18], and BNLDs [PCX⁺18]. Projective Blue-Noise shows similar integration results as dart throwing (at higher computational cost due to sample rejection in projective subspaces). The curve for Spoke-Darts exhibits convergence artifacts on the Gaussian test which can be explained by the low quality near the boundary of the unit domain of the point distributions obtained from the authors source code¹ (see Fig. 16 for a clear visualization of these artifacts). Finally, BNLDs shows interesting convergence results (as the pointset has some low discrepancy properties) but samples are only available for sample counts of 64, 4096 and 262144 in dimension 6.

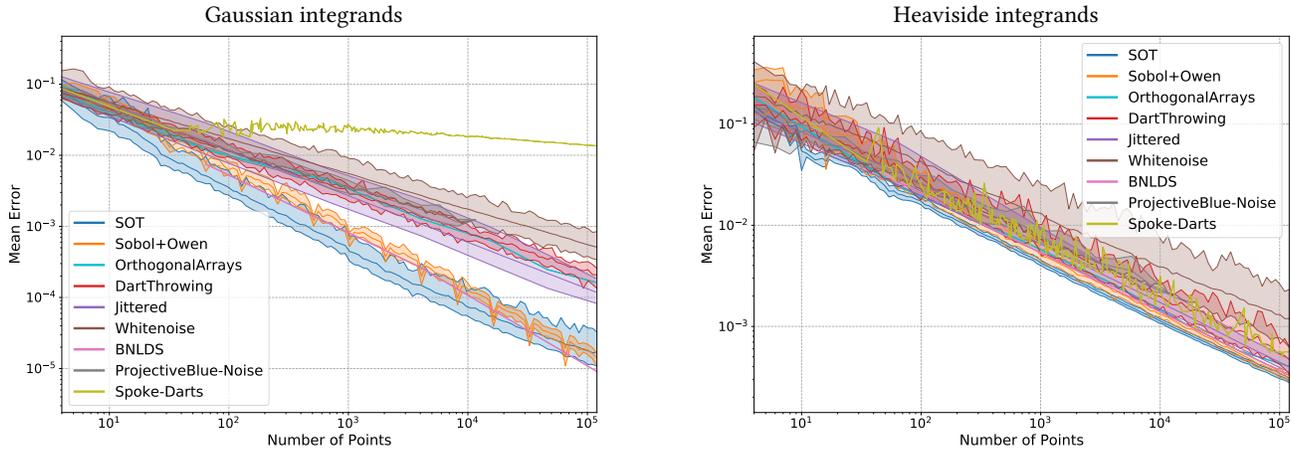


Fig. 2. **Monte-Carlo integration in dimension 6.** To complement the convergence graphs in 6D of the paper, we add few more samplers: Projective Blue-Noise [RRSG16], Spoke-Darts [MEA⁺18], BNLDSeq [PCX⁺18].

3 SOT ENERGY EVOLUTION

Figure 3 presents the sliced energy behavior of SOT as a function of the number of steps of the optimization scheme (in dimensions 2, 4, 6, 8, and 20). As the dimension increases, more steps are required to reach a plateau in the energy function.

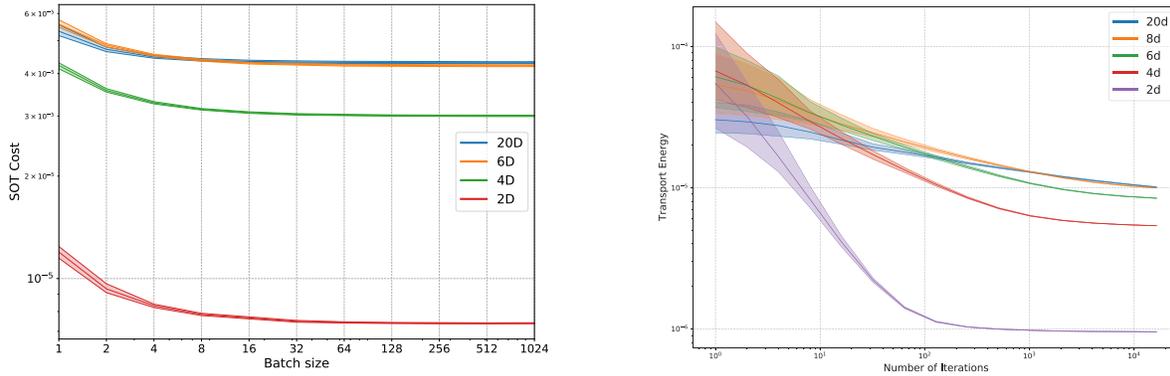


Fig. 3. **SOT energy during optimization.** Evolution of sliced optimal transport energy averaged over 64 realisations (mean, min and max are shown) as a function of (left) the number of slices per batch for 4096 batches on 1024 samples, and (right) the number of batches for 32 slices for 4096 samples.

4 SOT ON VARYING DENSITIES

Arbitrary sampling densities are rather simple to handle in low dimensions as the Radon transform can be directly evaluated via the sliced Fourier theorem. We illustrate a preliminary result of this feature using a 2D distribution defined by an input image in Fig. 4.

¹We used the 1line-spoke version without toricity.

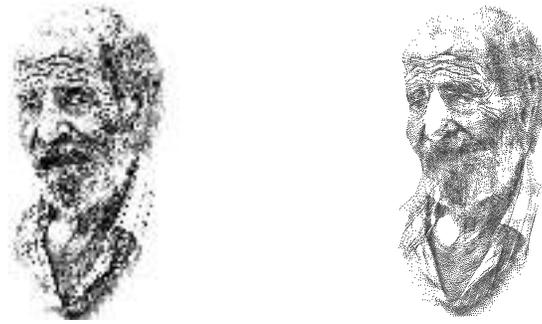


Fig. 4. **Varying densities.** Our SOT sampling supports arbitrary densities in low dimensions, as illustrated on this 2D example. This can be done via a Fourier-based evaluation (or through other numerical integration techniques) of the Radon transform.

5 EXTRA RENDERING RESULTS

We also add more rendering examples along with their comparisons: one with indirect lighting and 3 bounces for the Cornell scene, one with 1 bounce for the San Miguel scene, and one with 2 bounces for the San Miguel scene again — each time preceded by their reference image for comparison purposes. Finally, we also provide more volumetric rendering results.

5.1 Cornell scene, with 3 bounces

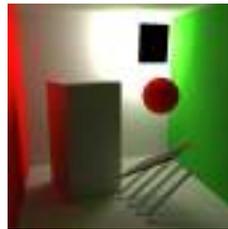


Fig. 5. **Reference image for 10D rendering.** Cornell scene with 3 bounces of indirect lighting (10D samples).

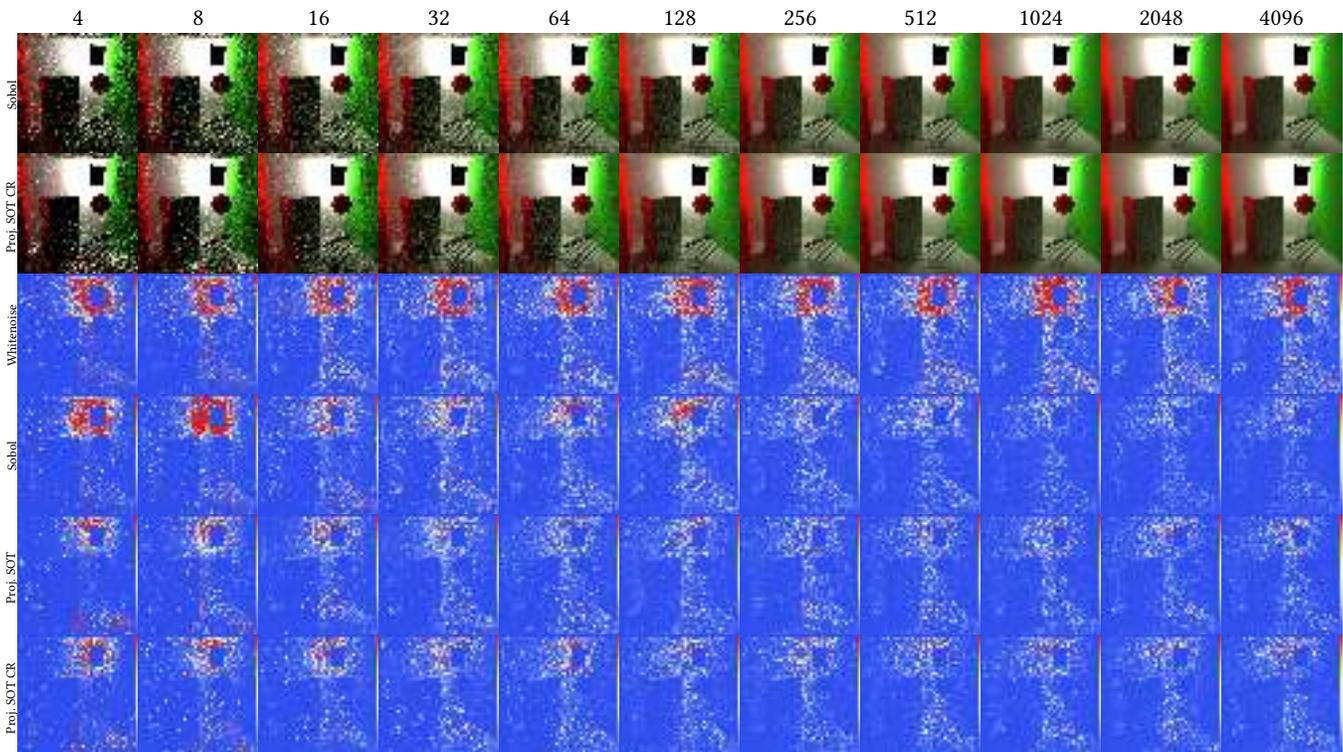


Fig. 6. **Rendering with 10D sampling.** Cornell scene with 3 bounces of indirect lighting; absolute error measured with respect to reference image. Colormaps: from blue = 0, to red = $0.002/spp$.

5.2 San Miguel scene, with only 1 bounce of indirect lighting



Fig. 7. **Reference image for 6D rendering.** San Miguel scene with 1 bounce and indirect lighting.

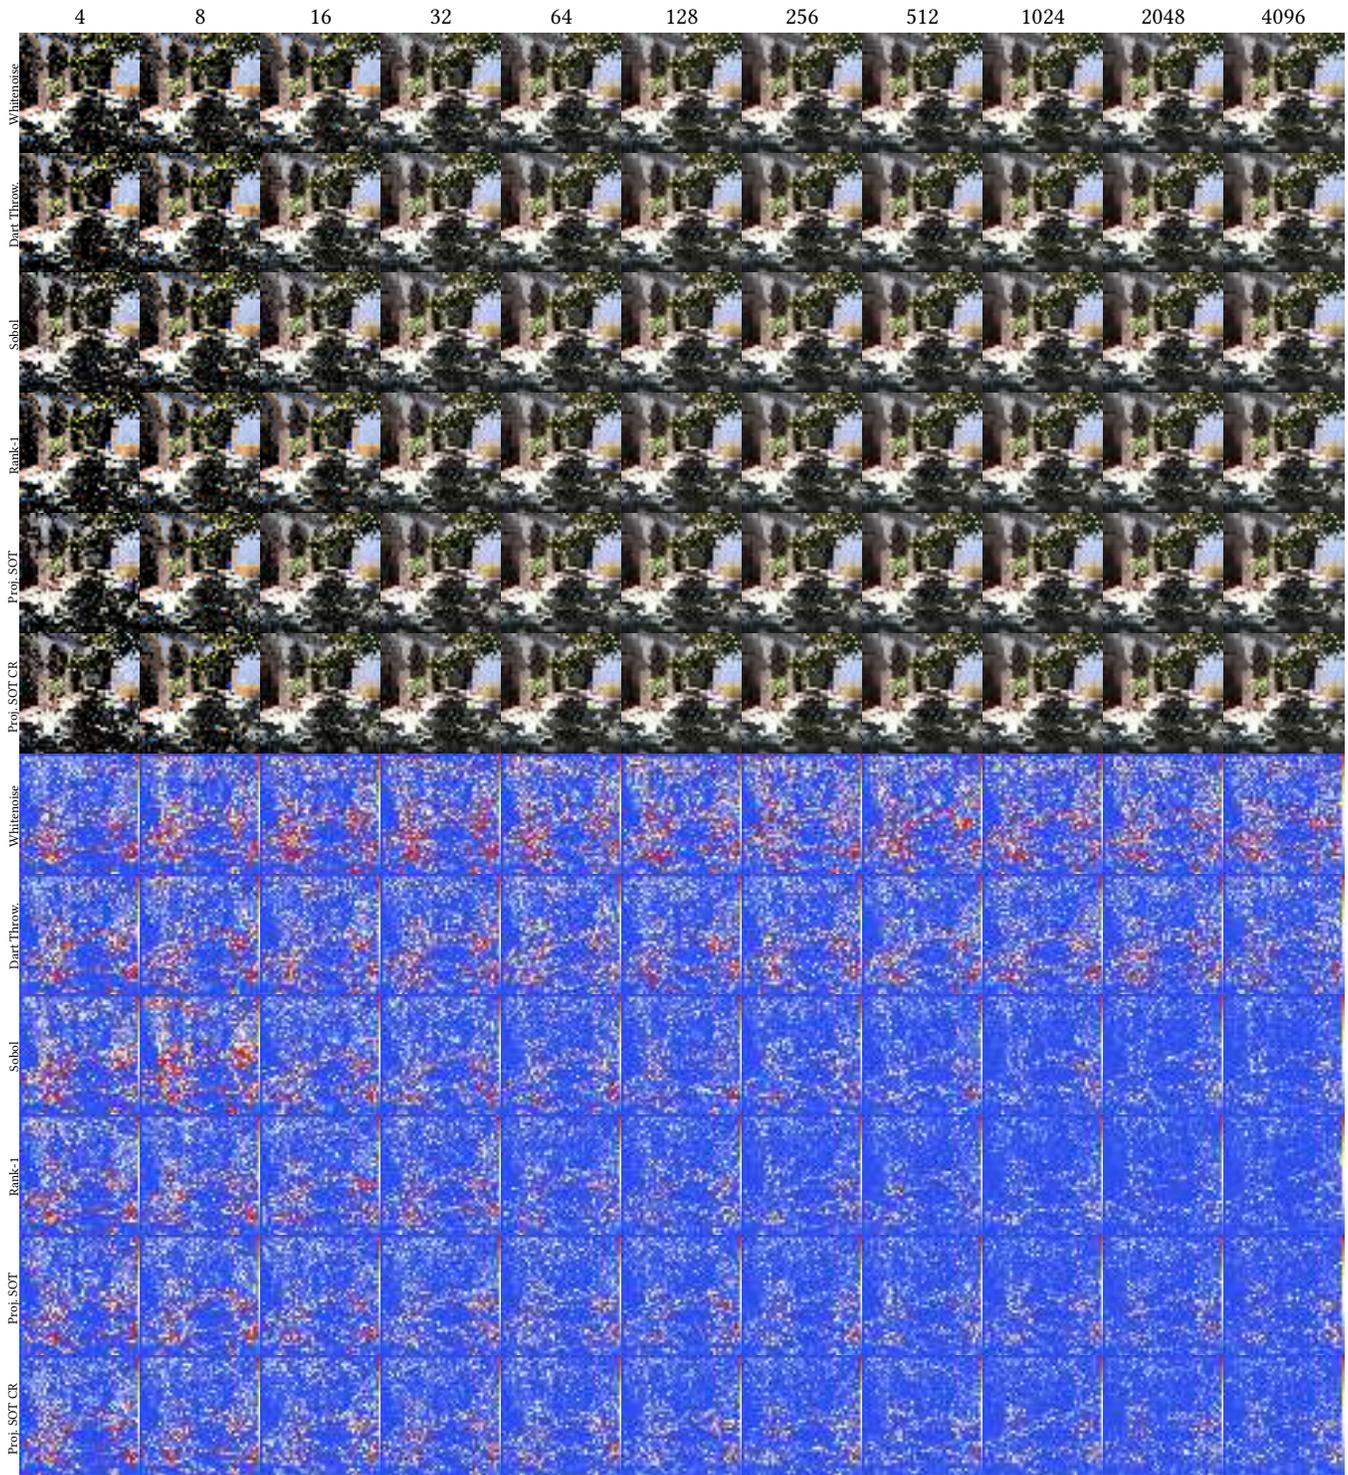


Fig. 8. **Rendering with 6D sampling.** San Miguel scene with 1 bounce and indirect lighting; absolute error measured with respect to the reference image. Colormaps: from blue = 0, to red = $0.0005/spp$.

5.3 San Miguel with 2 bounces of indirect lighting



Fig. 9. **Reference image for 8D rendering.** San Miguel scene with 2 bounces and indirect lighting.

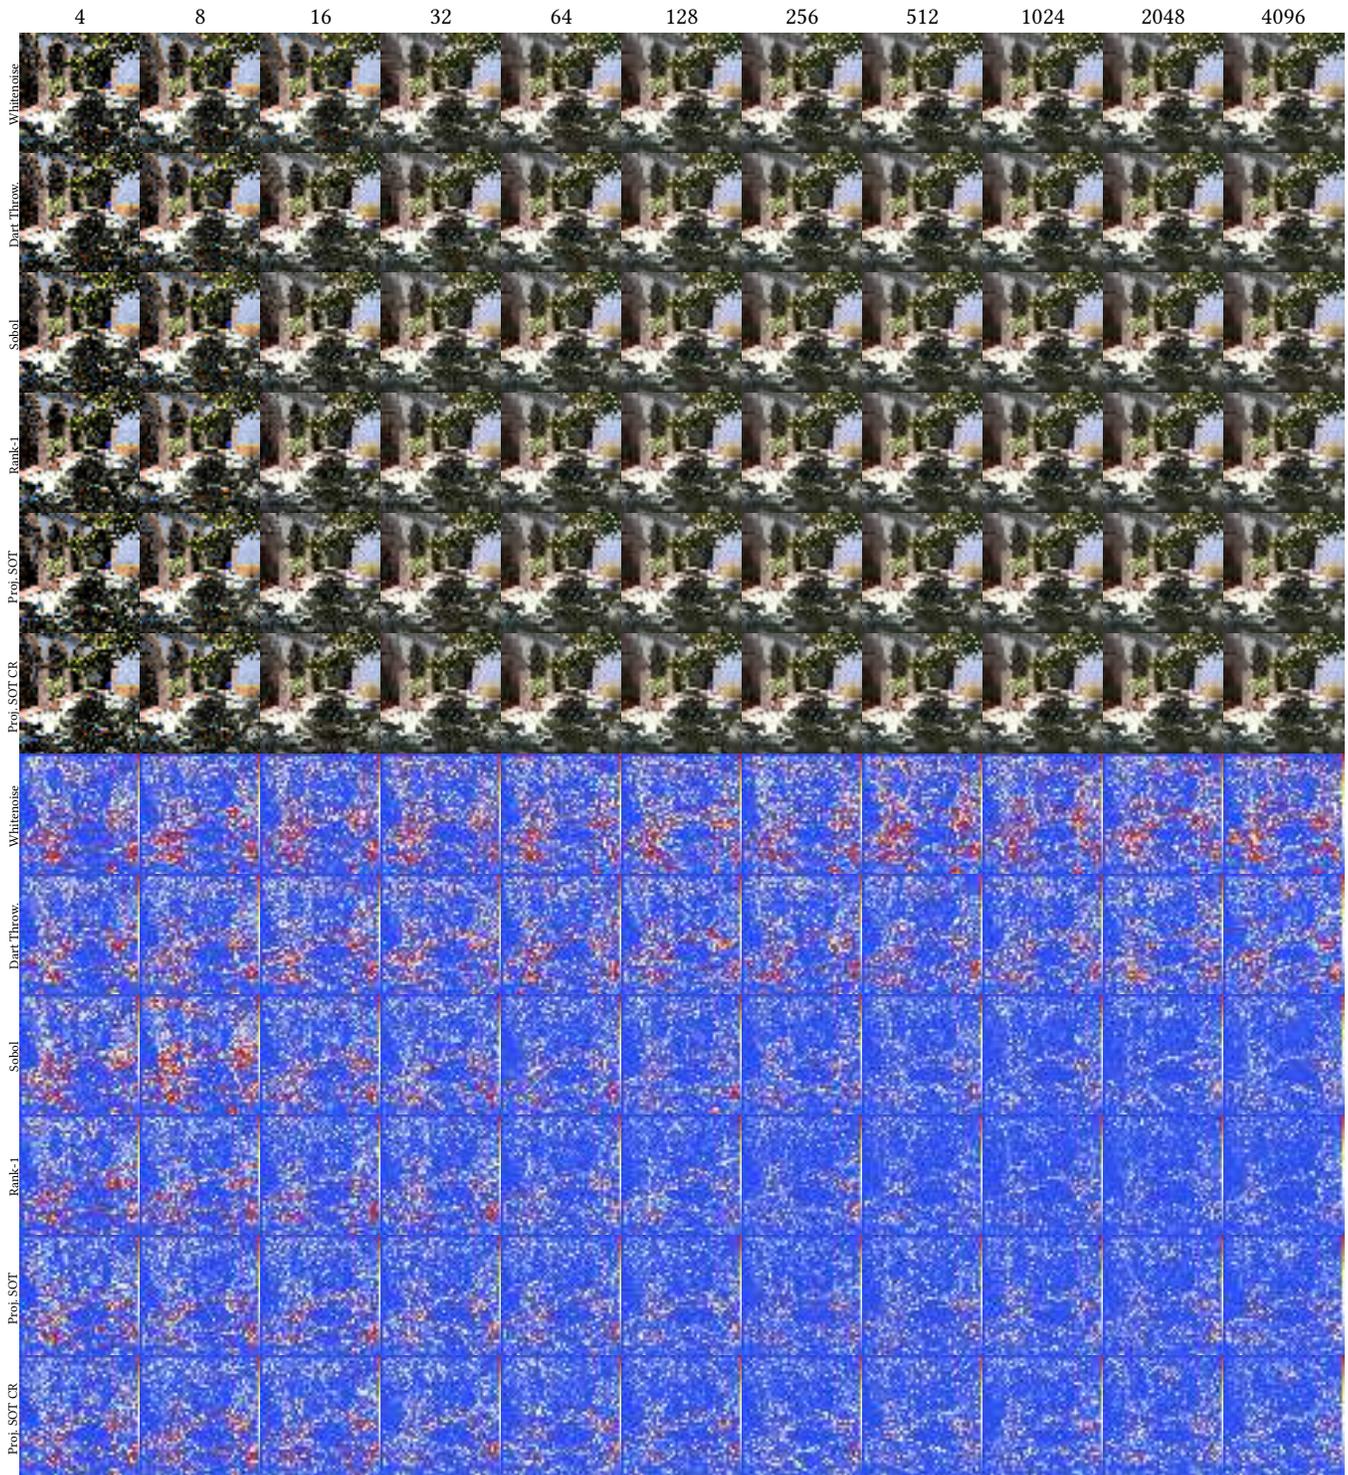


Fig. 10. **Rendering with 8D sampling.** San Miguel scene with 2 bounces and indirect lighting; absolute error measured with respect to the reference image. Colormaps: from blue = 0, to red = $0.0005/spp$.

5.4 Volume rendering with 2 bounces

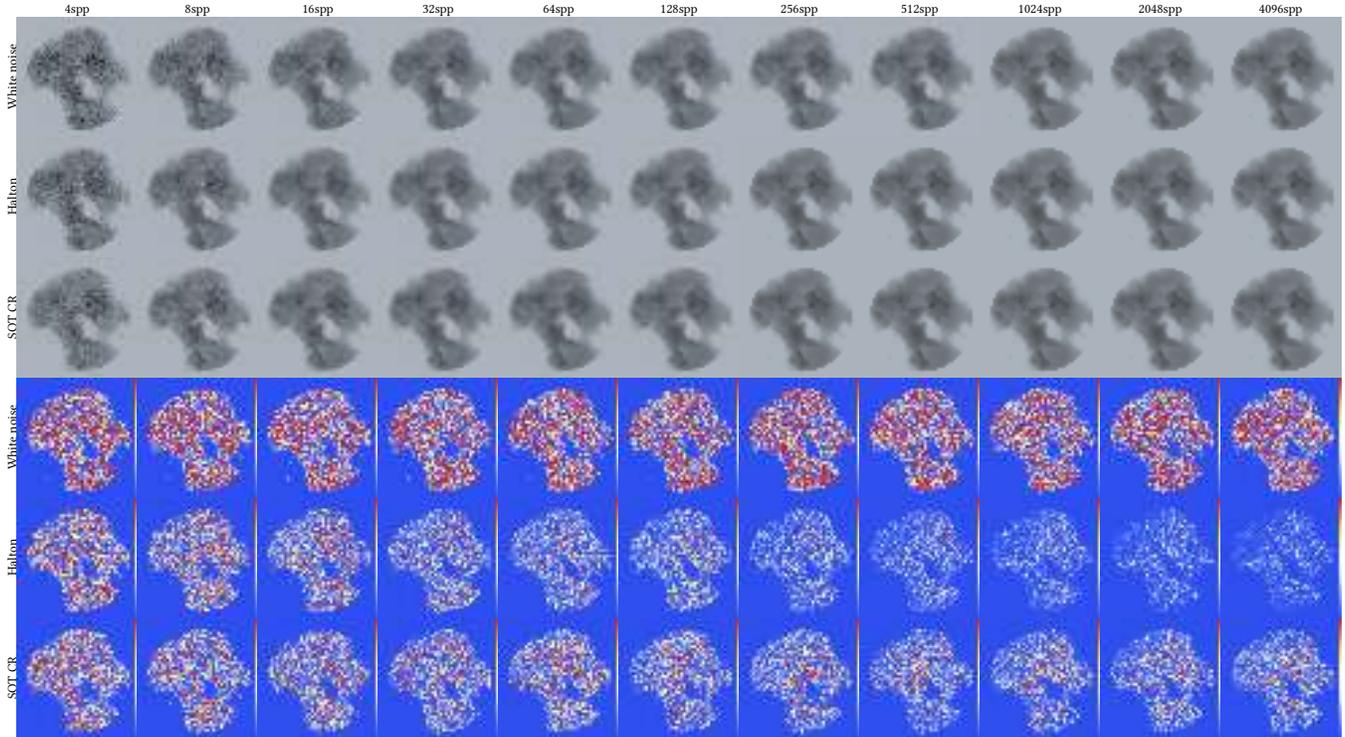


Fig. 11. Extra volumetric rendering results in 8 dimensions (two bounces). Complement of Fig. 16 in the main paper.

6 VISUALIZATION OF POINTSET PROJECTIONS

Finally, we present in this section a visualization of all the 2D projections of 6D 4k sample distributions generated with various samplers.

In Figure 12, the Sobol samples scrambled by the Owen method [Owe98] seem well distributed in each projection except in the projection (4, 6) where patterns, clusters, and voids are observed. The same kind of problems can be seen in different projections of the BNLDSeq samples [PCX⁺18] in Figure 13, namely the projections (2, 3), and (3, 6). Note that the latter optimizes the samples along the (1, 2), (3, 4) and (5, 6) 2D spaces which is clearly visible in the sample distribution. No particular patterns are observed in the projections of the Projective Blue Noise samples [RRSG16] in Figure 18) or in the projections of the Orthogonal Arrays samples [JEK⁺19] in Figure 17. The projections of these two latter pointsets look like 2D white noise samples. In Figure 19, the samples of the rank-1 lattice [Kel04] avoid any clustering by design, but strong alignments are present in some projections even in subsequent pairs of dimensions, such as (5, 6), or, to a lesser extent, in projection (3, 4). In Figure 20, alignments are noticeable near the edges of each projection of our SOT samples. Finally, in Figure 21, the projections of our Projective SOT samples are distributed in a more blue-noise way in subsequent pairs of dimensions (i.e., the ones optimized by design), namely the projections (1, 2), (3, 4), and (5, 6).

Notice that the distributions of samples of the rank-1 lattice in projections (1, 2), (3, 4), and (5, 6) does vary considerably, according to the number of samples being projected; please refer to Figures 22, 23 and 19, where projections of 1K, 2K and 4K samples respectively are shown. For example, the distribution of samples in the projection (3, 4) is close to an hexagonal grid for 1K and 2K samples, whereas it forms a very uneven distribution for 4K samples. This uneven behavior of the distribution of samples in projections, proper to the rank-1 lattice method, may partly explain the behavior of our error graphs shown in Figs. 12 and 15 in the main paper.

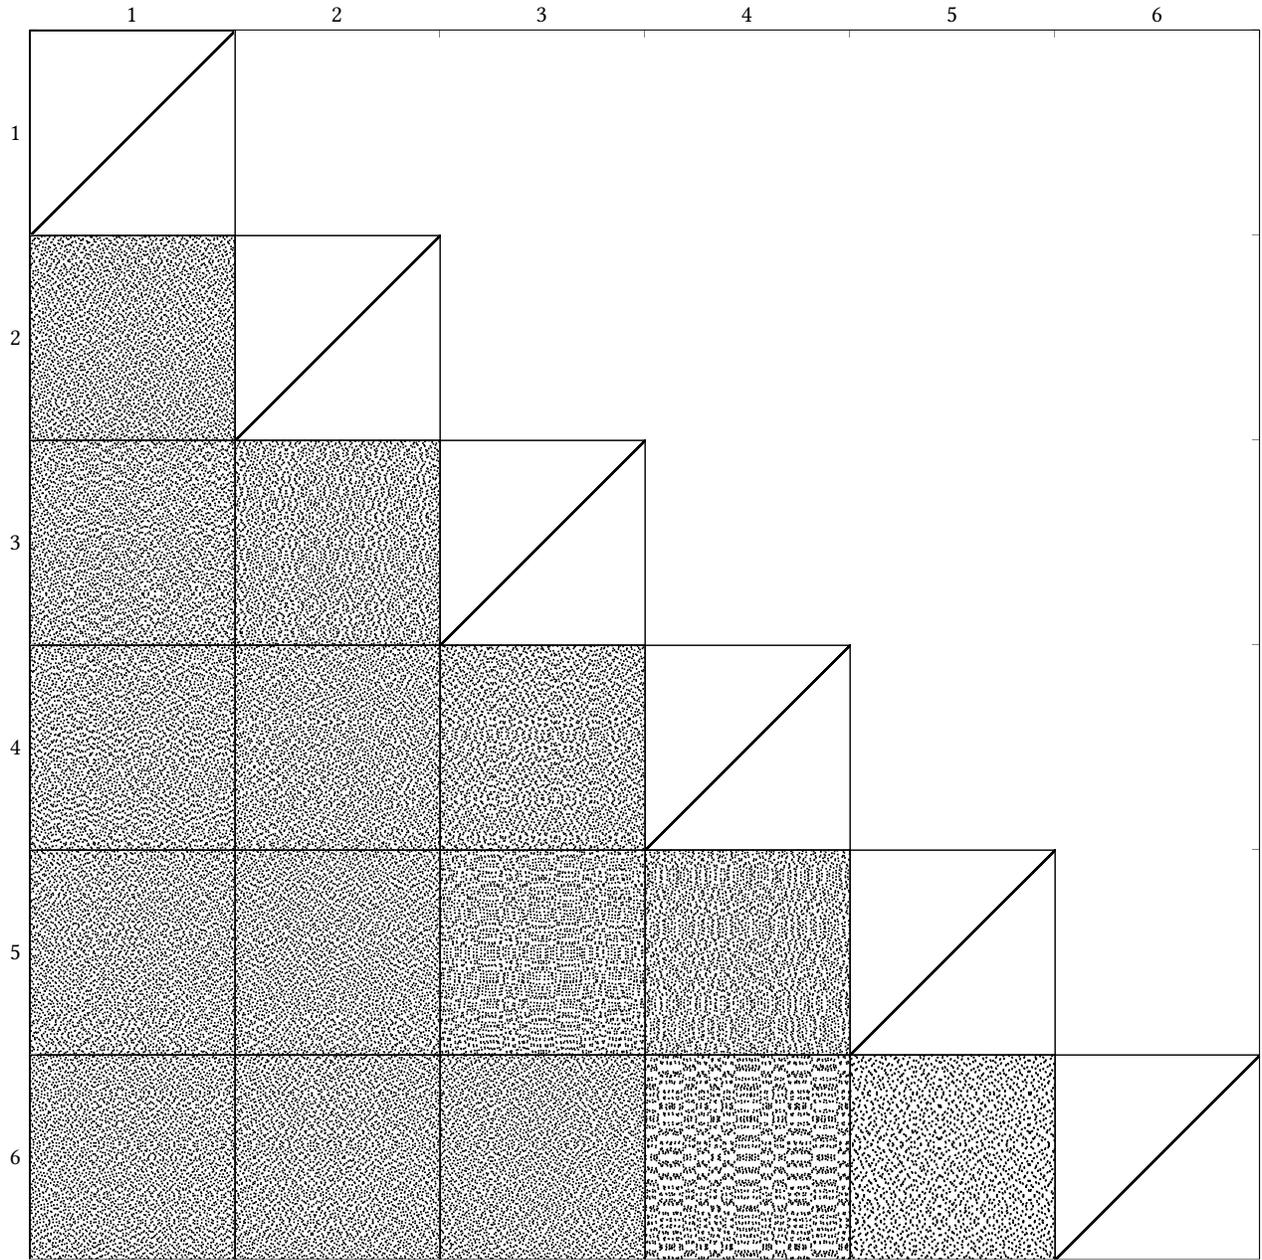


Fig. 12. [Owe98]. A grid of 2D projections for 4096 samples of Sobol+Owen in 6D.

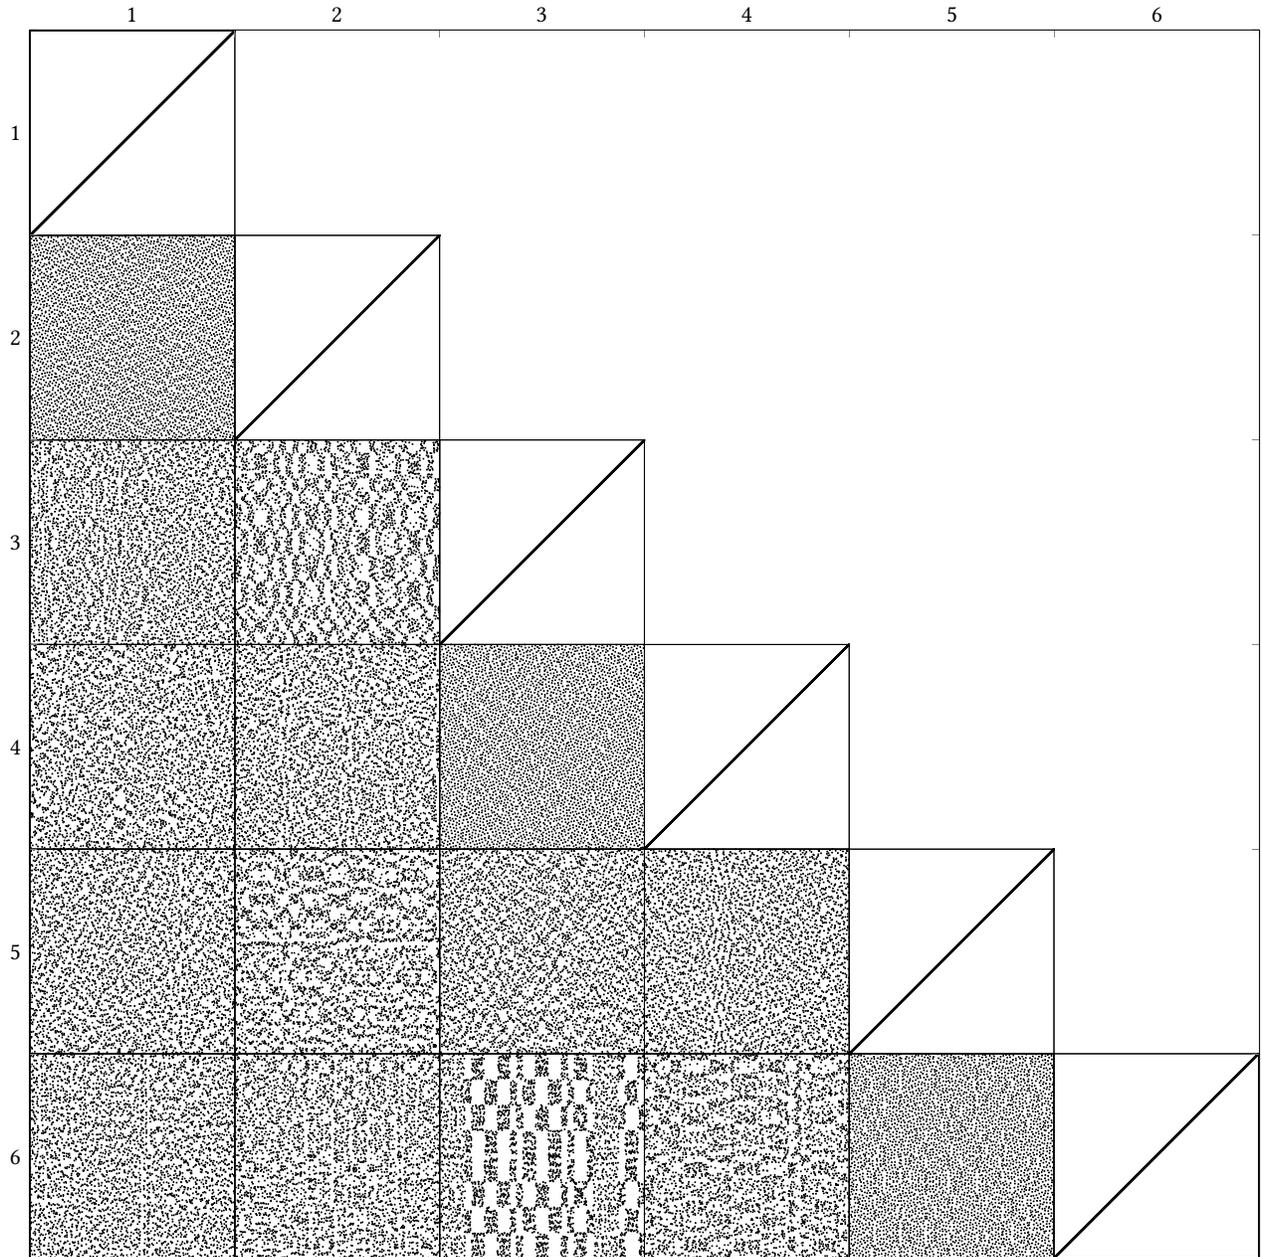


Fig. 13. [PCX⁺18]. A grid of 2D projections for 4096 samples of BNLDSeq in 6D.

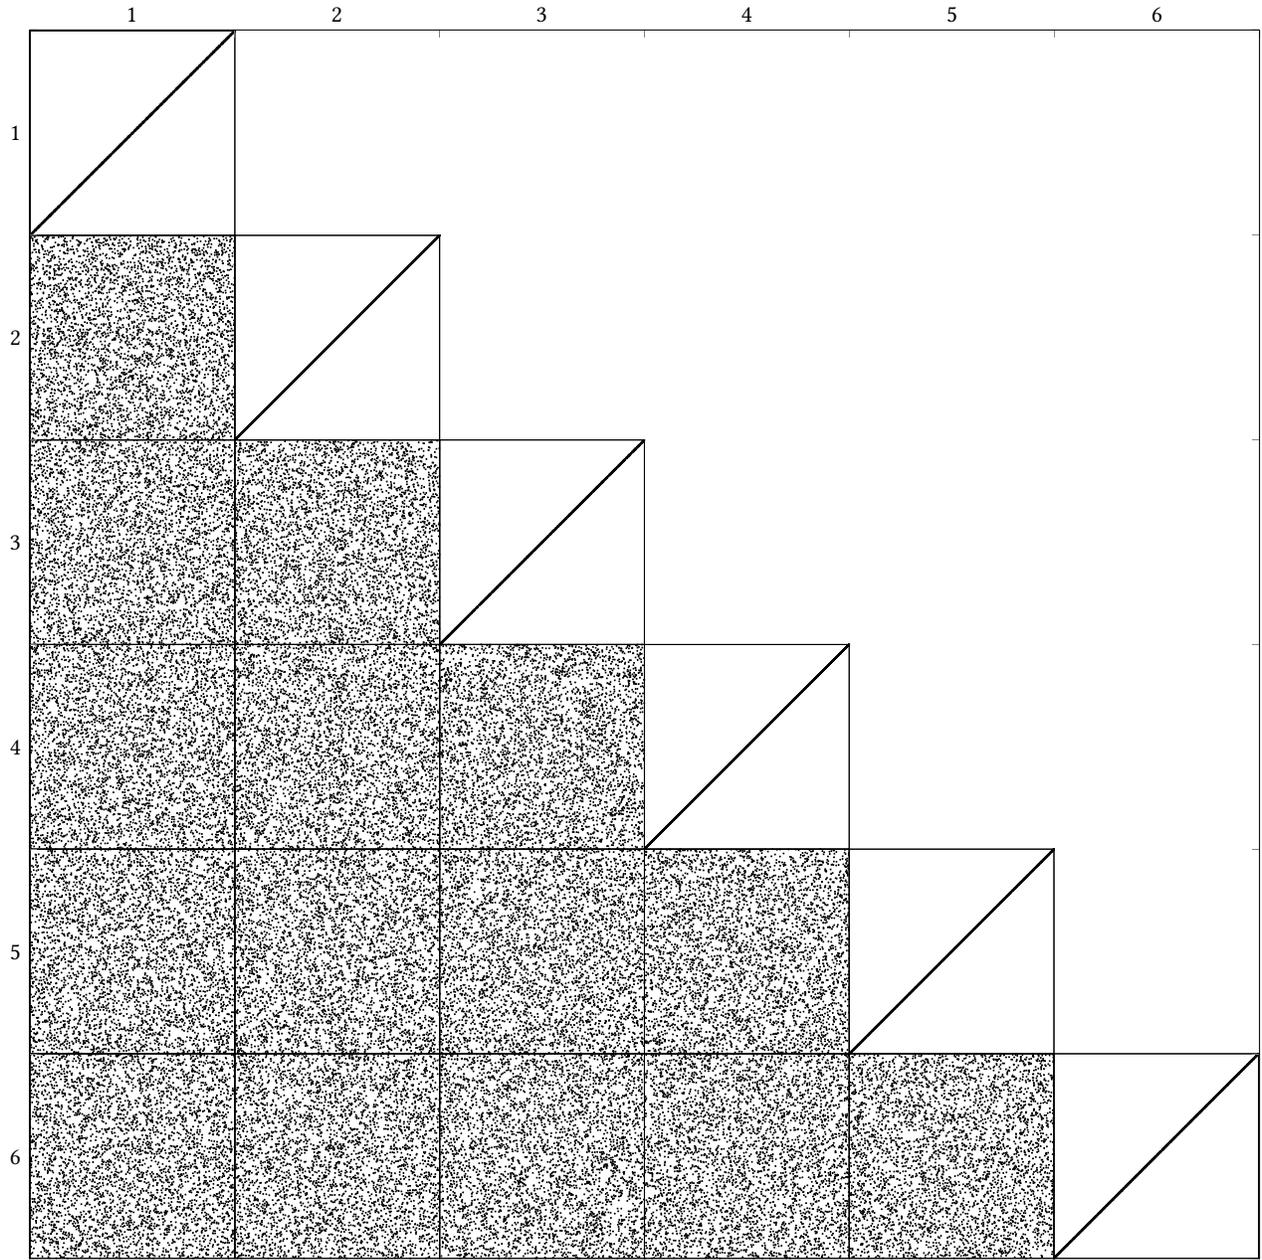


Fig. 14. **Stratified sampling.** A grid of 2D projections for 4096 samples in 6D

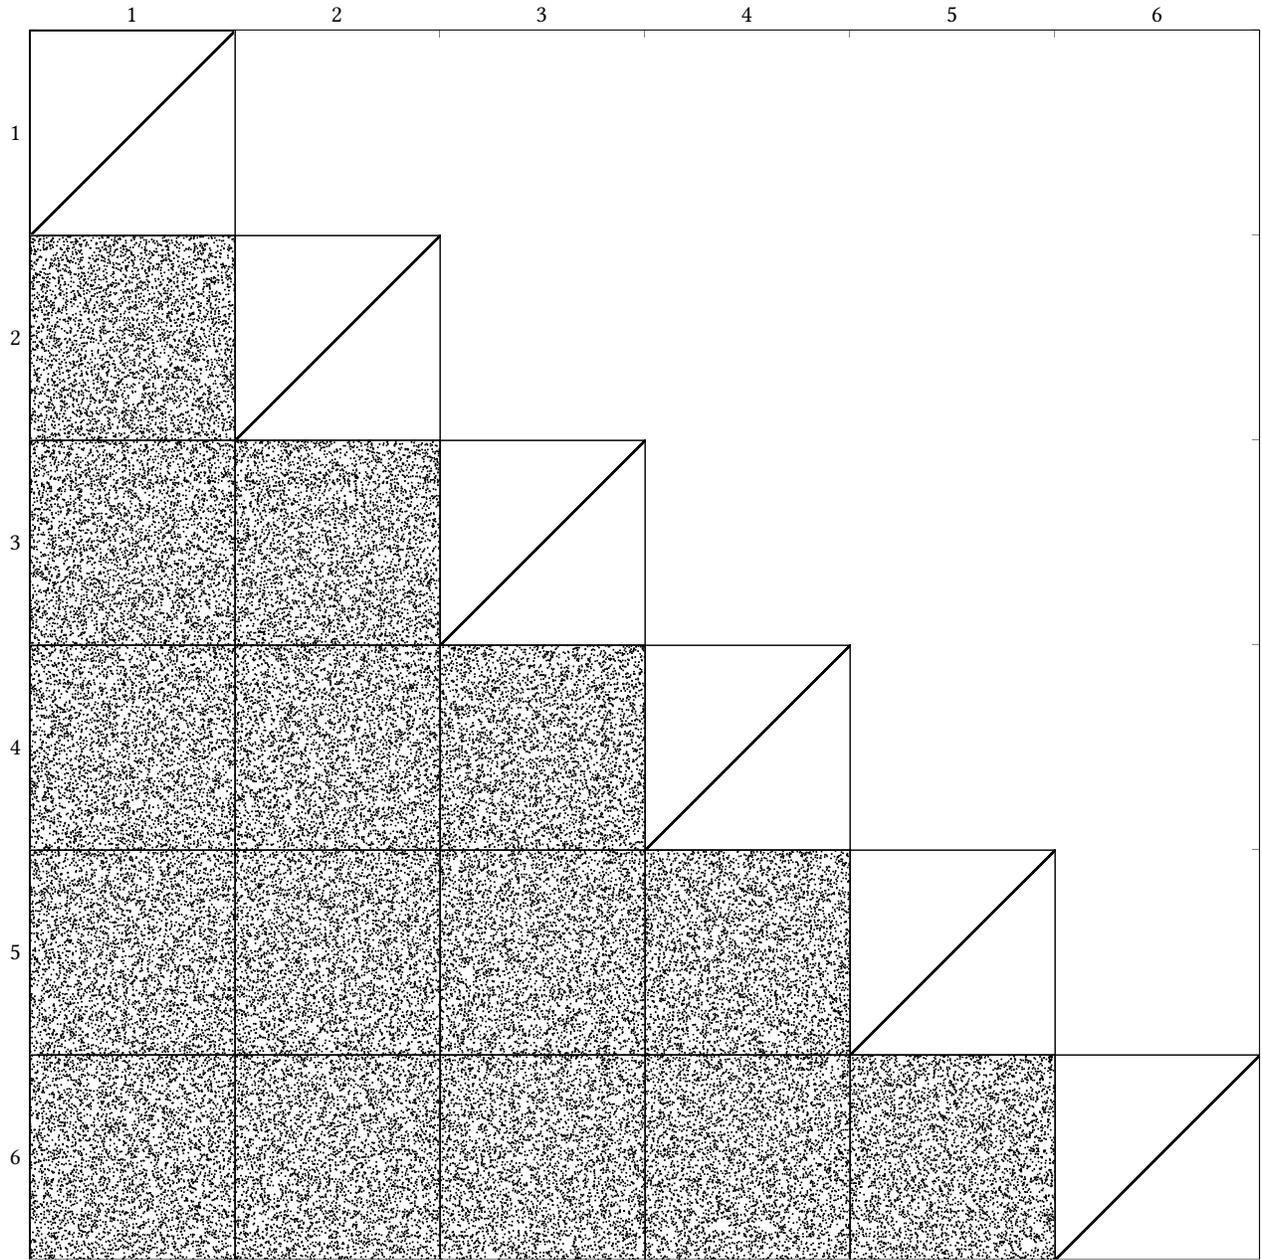


Fig. 15. **Dart throwing.** A grid of 2D projections for 4096 samples in 6D.

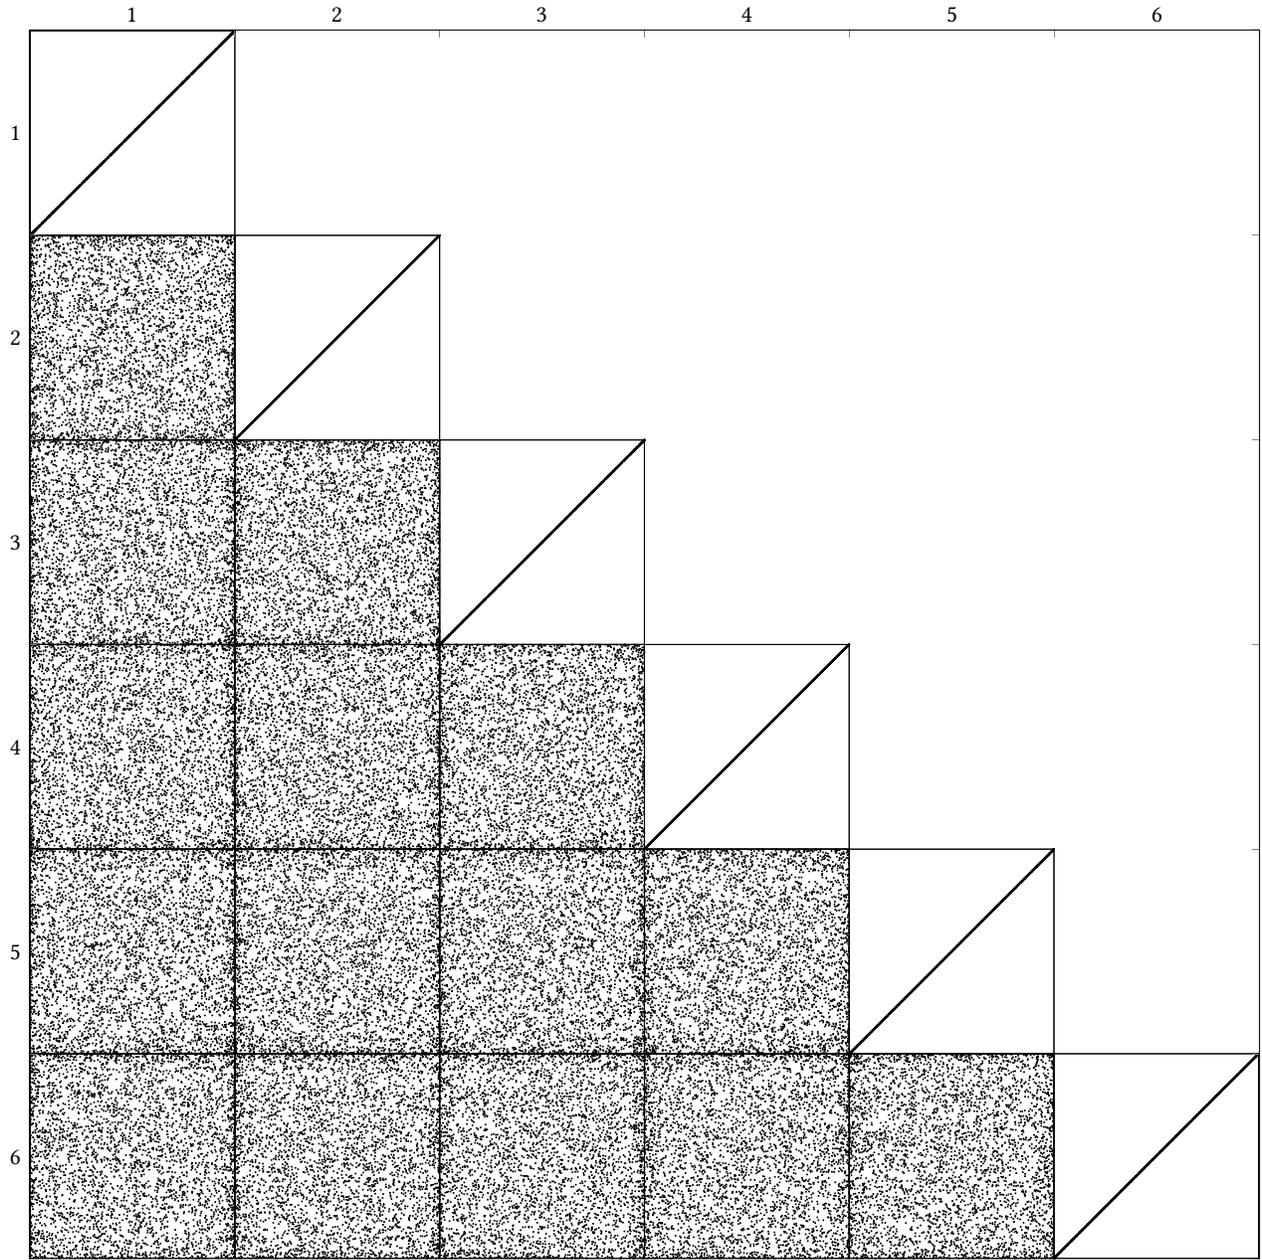


Fig. 16. [MEA⁺18]. A grid of 2D projections for 4056 samples in 6D.

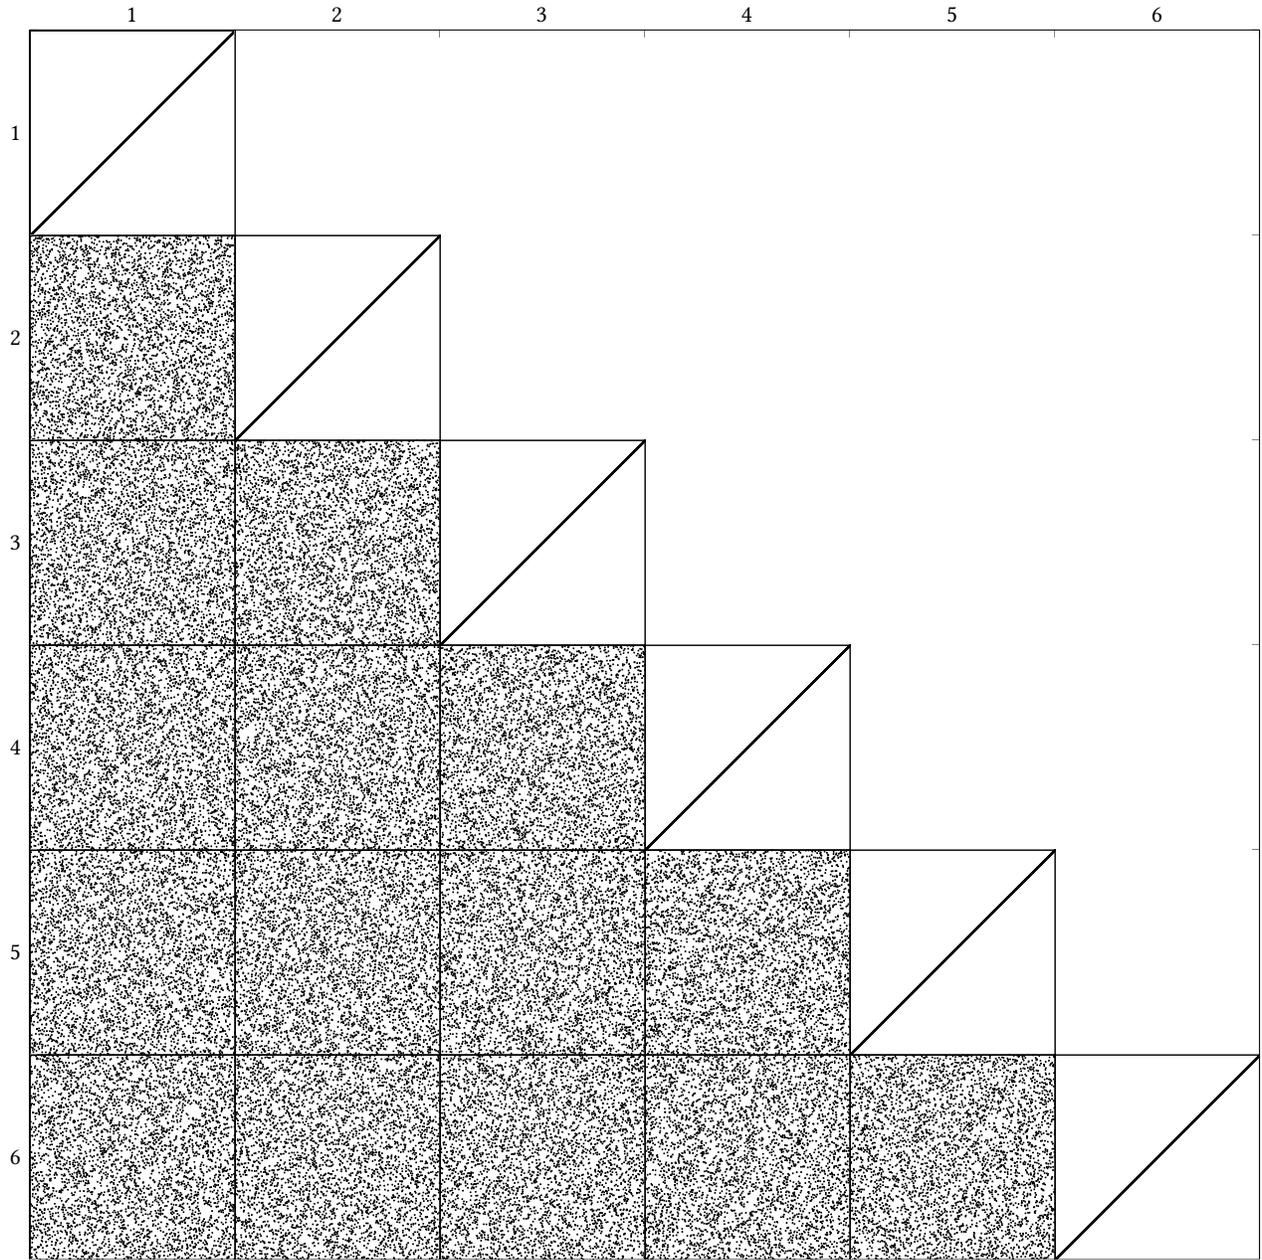


Fig. 17. [JEK⁺19]. A grid of 2D projections for 4096 samples of Orthogonal Arrays in 6D.

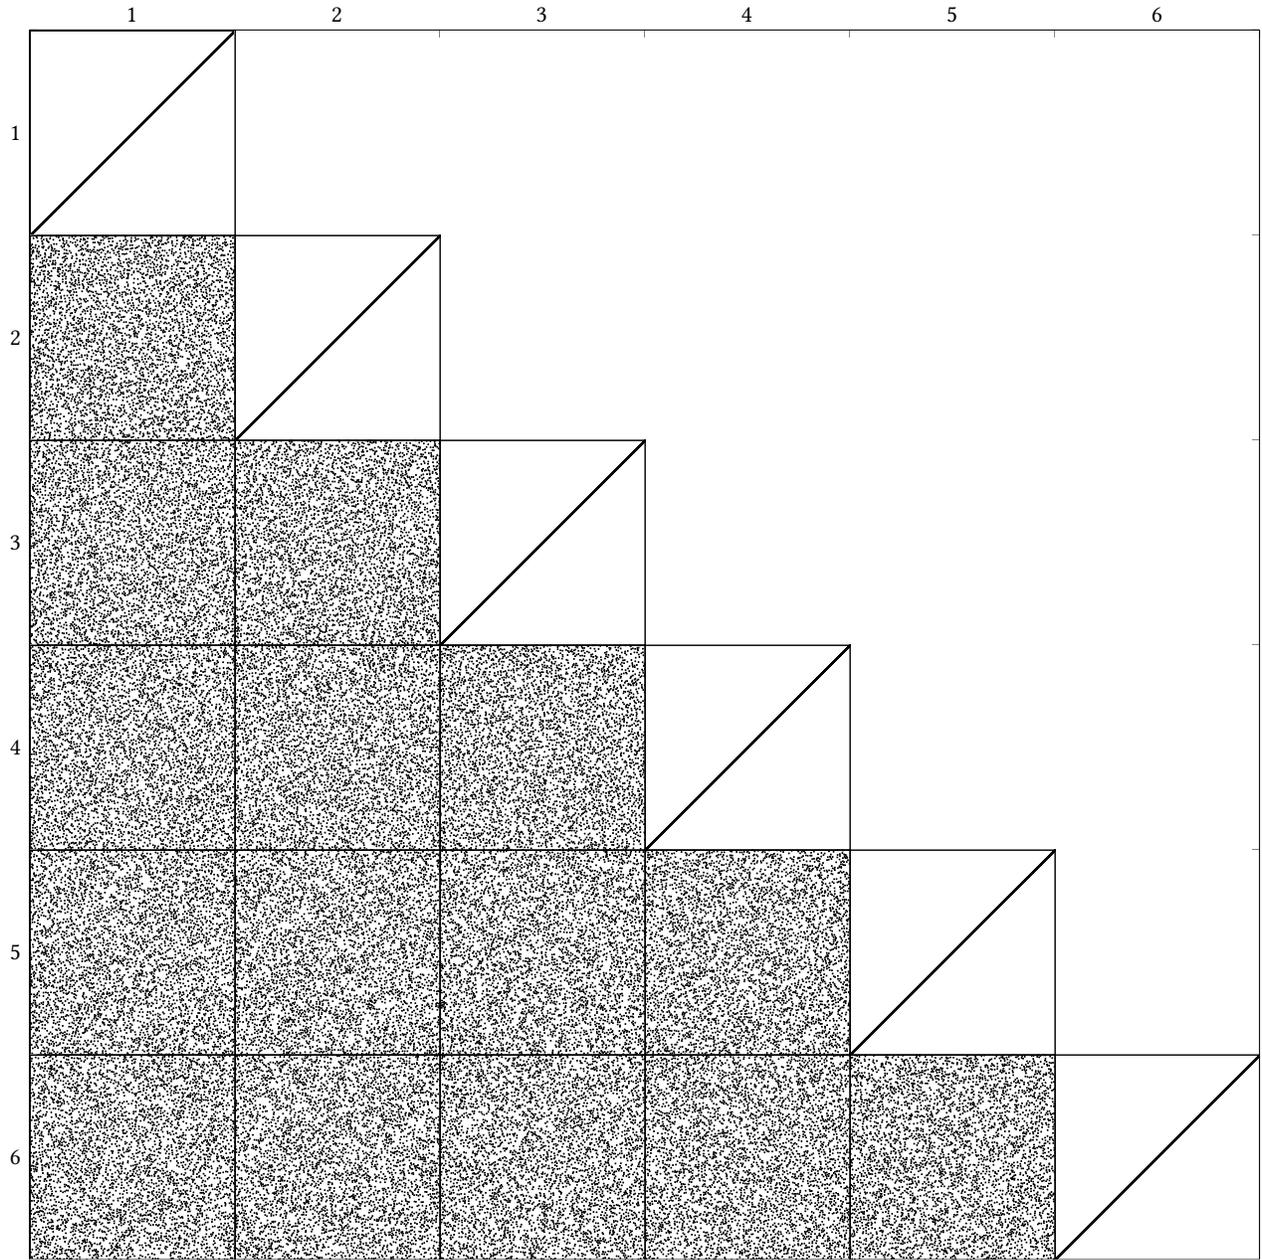


Fig. 18. [RRSG16]. A grid of 2D projections for 4096 samples of Projective Blue Noise in 6D.

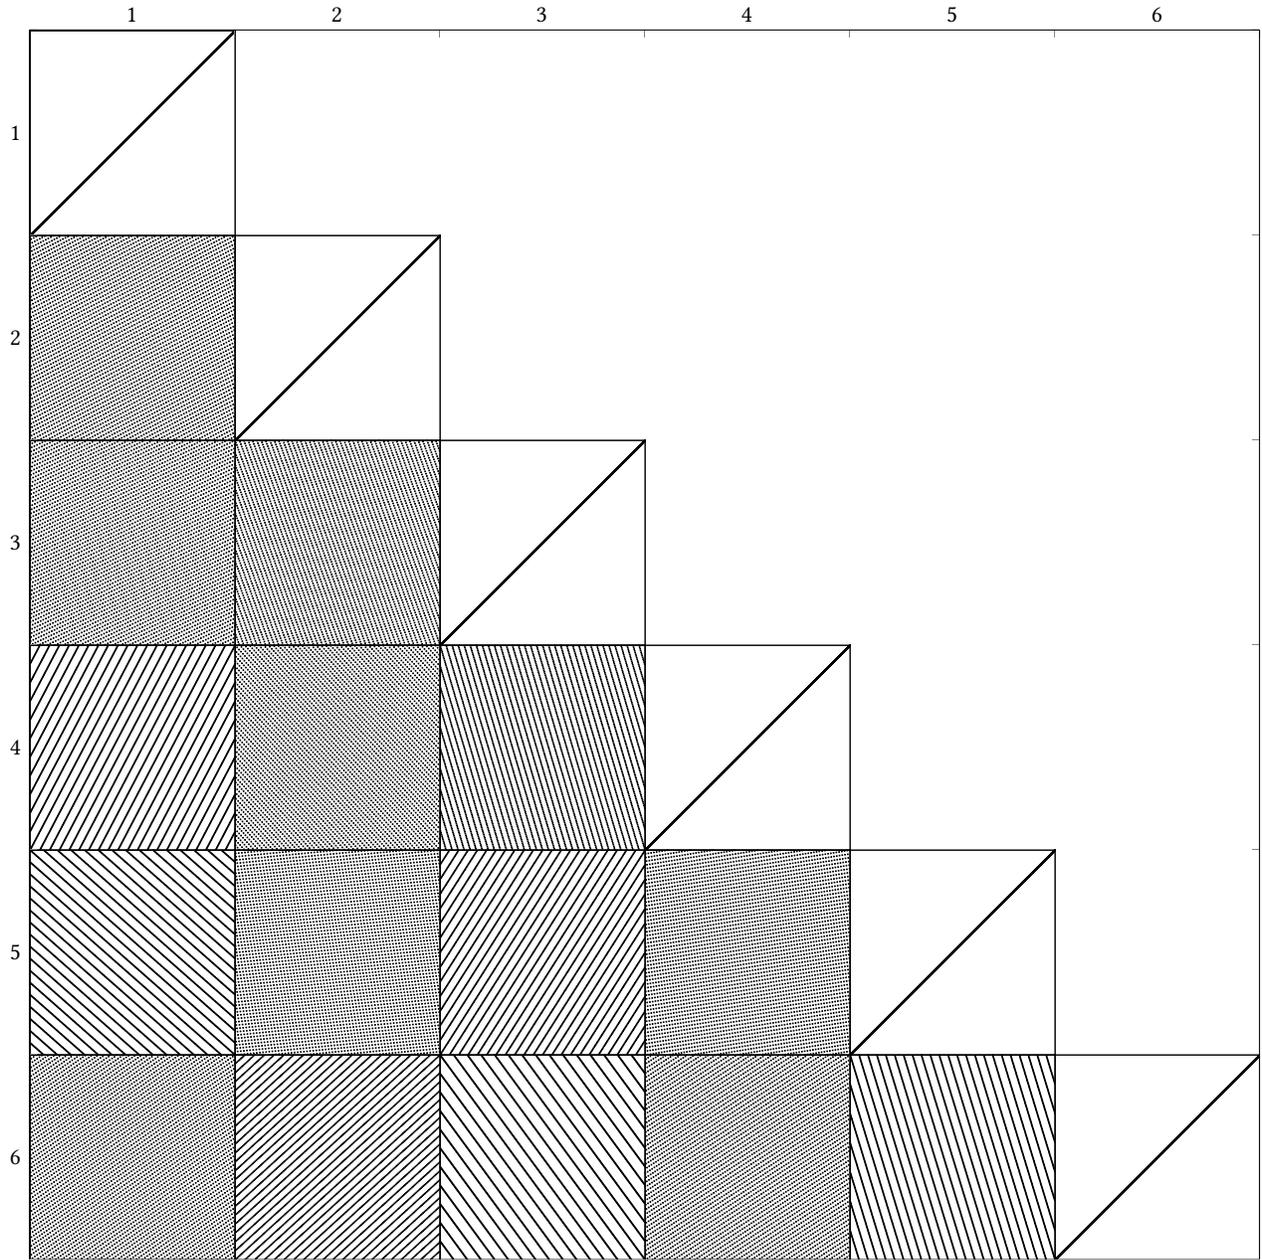


Fig. 19. [Kel04]. A grid of 2D projections for 4096 samples of rank-1 in 6D using the implementation of [LM16].

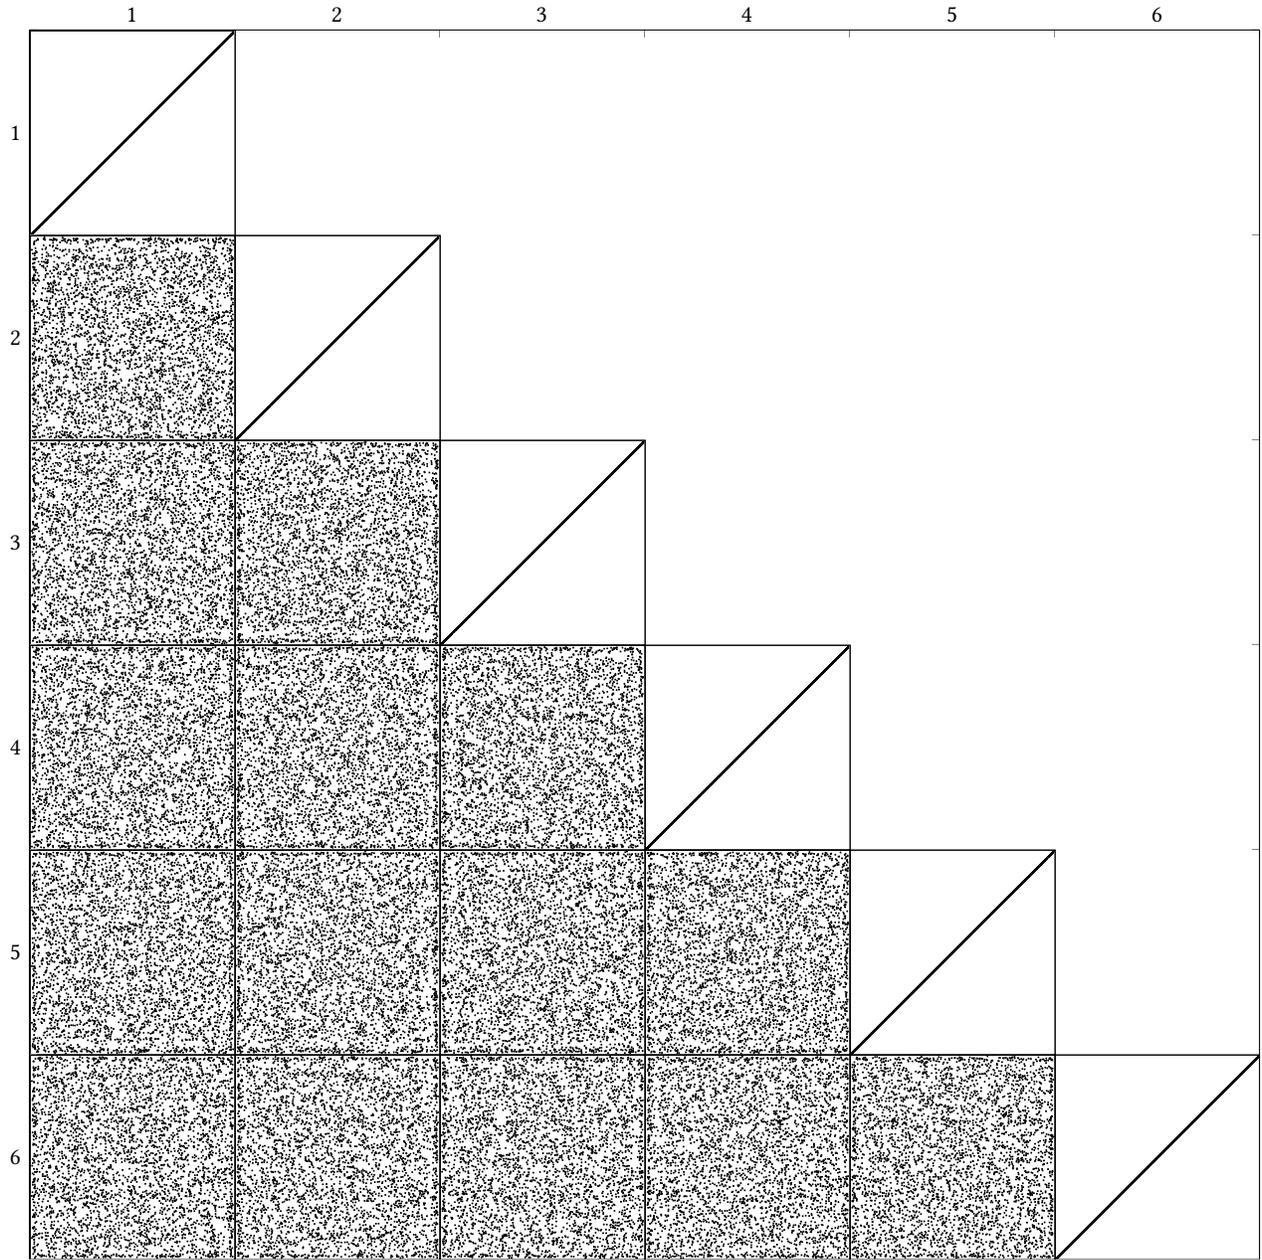


Fig. 20. **SOT sampling.** A grid of 2D projections for 4096 samples of SOT in 6D.

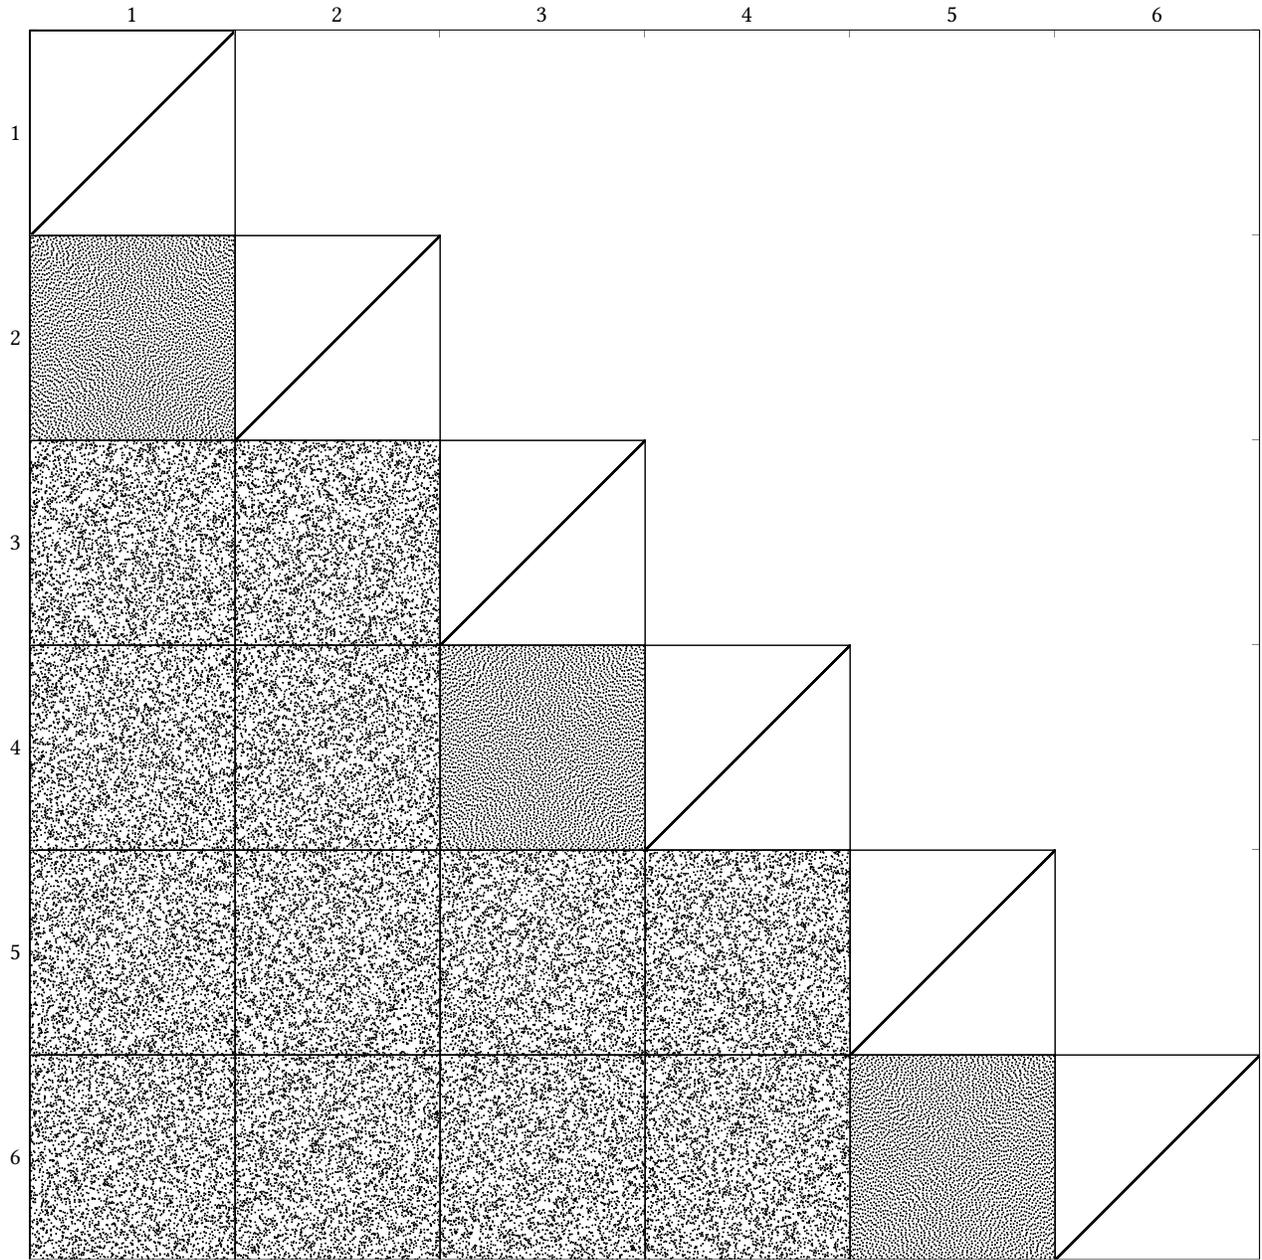


Fig. 21. **Projective SOT.** A grid of 2D projections for 4096 samples of Projective SOT in 6D.

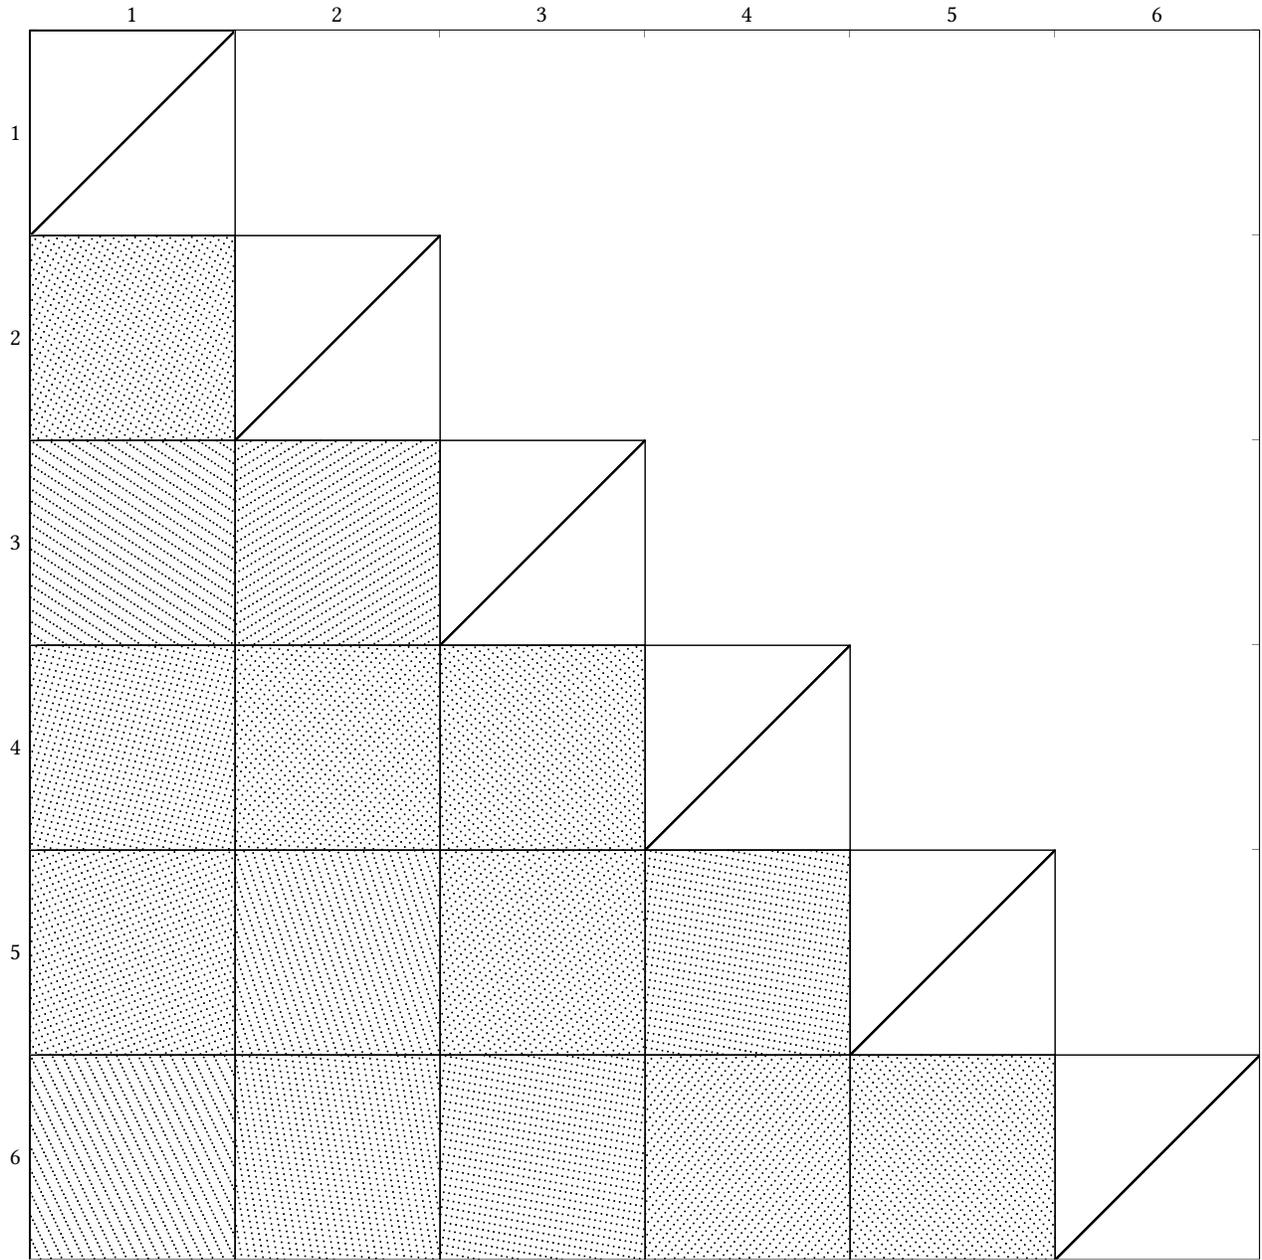


Fig. 22. [Kel04]. A grid of 2D projections for 1024 samples of rank-1 in 6D using the implementation of [LM16].

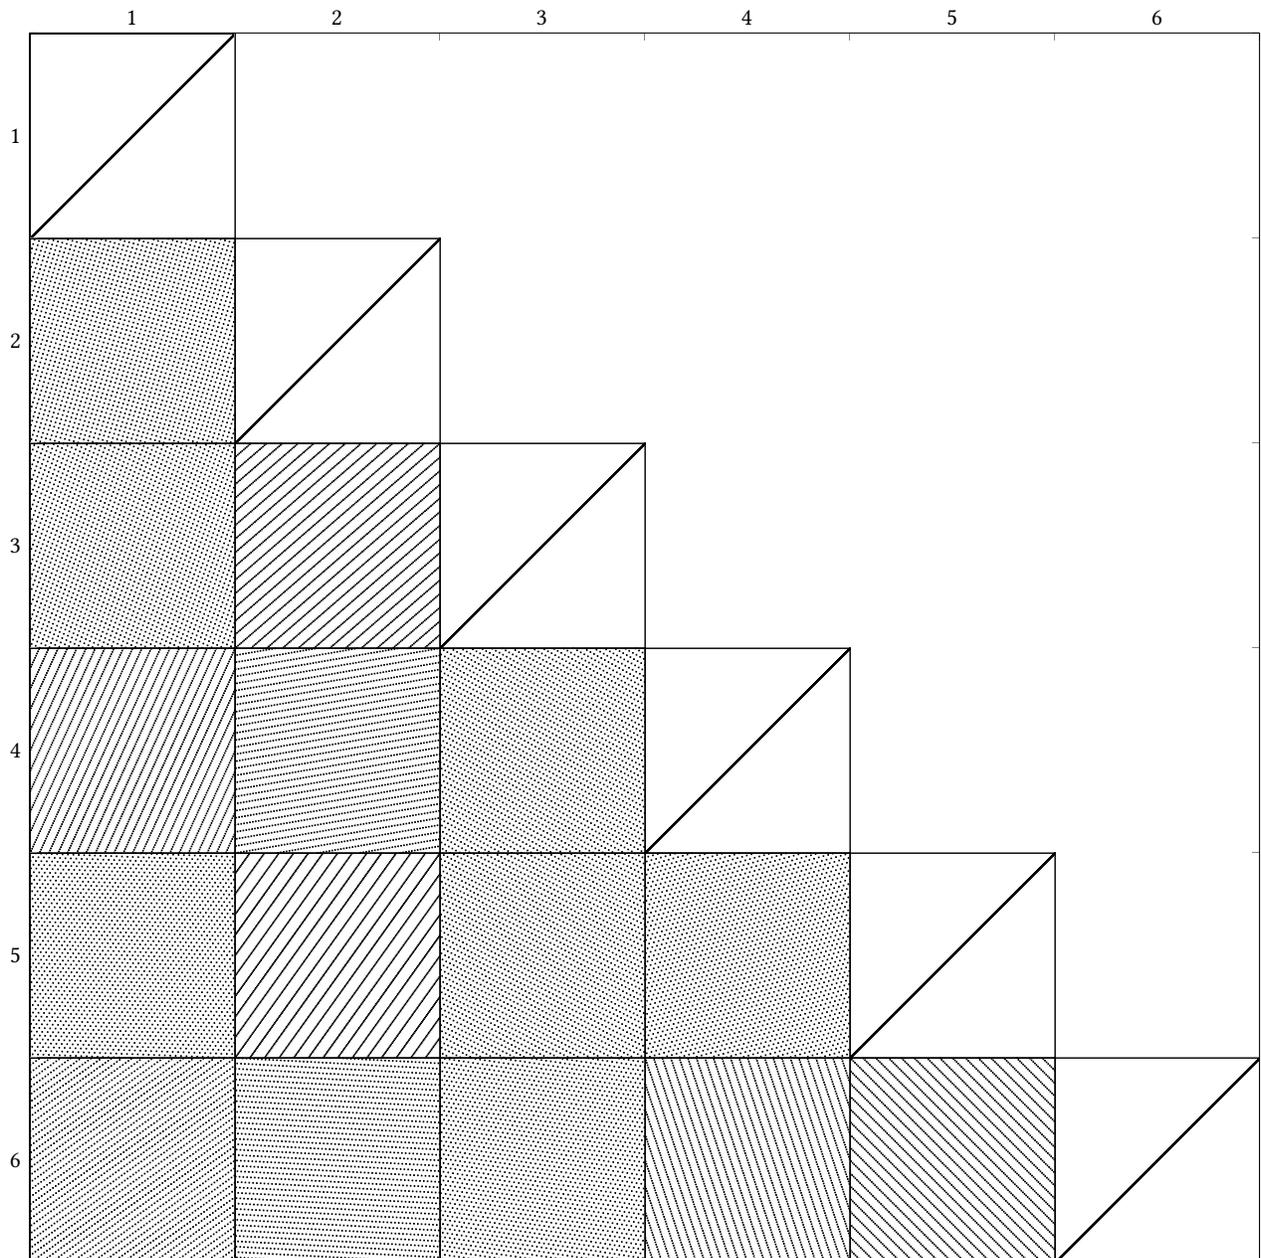


Fig. 23. [Kel04]. A grid of 2D projections for 2048 samples of rank-1 in 6D using the implementation of [LM16].

REFERENCES

- [JEK⁺19] Wojciech Jarosz, Afnan Enayet, Andrew Kensler, Charlie Kilpatrick, and Per Christensen. Orthogonal array sampling for monte carlo rendering. *Computer Graphics Forum*, 38(4):135–147, 2019.
- [Kel04] Alexander Keller. Stratification by rank-1 lattices. In Harald Niederreiter, editor, *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pages 299–313. Springer, 2004.
- [LM16] Pierre L’Ecuyer and David Mungler. Lattice builder: A general software tool for constructing rank-1 lattice rules. <https://github.com/umontreal-simul/latnetbuilder>, 2016.
- [MEA⁺18] Scott A Mitchell, Mohamed S Ebeida, Muhammad A Awad, Chonhyon Park, Anjul Patney, Ahmad A Rushdi, Laura P Swiler, Dinesh Manocha, and Li-Yi Wei. Spoke-darts for high-dimensional blue-noise sampling. *ACM Trans. Graph.*, 37(2):22, 2018.
- [Owe98] Art B. Owen. Scrambling sobol’ s and niederreiter-xing points. *Journal of Complexity*, 14(4):466–489, 1998.
- [PCX⁺18] Hélène Perrier, David Coeurjolly, Feng Xie, Matt Pharr, Pat Hanrahan, and Victor Ostromoukhov. Sequences with low-discrepancy blue-noise 2-d projections. In *Computer Graphics Forum*, volume 37, pages 339–353. Wiley Online Library, 2018.
- [RRSG16] Bernhard Reinert, Tobias Ritschel, Hans-Peter Seidel, and Iliyan Georgiev. Projective blue-noise sampling. In *Computer Graphics Forum*, volume 35, pages 285–295. Wiley Online Library, 2016.