CS 563 Advanced Topics in Computer Graphics
Stratified and Low-Discrepancy Sampling

by Stephen Kazmierczak

Overview

- Why?
- Stratified Sampling
- Low-Discrepancy Sampling

Why use a Sampling Pattern?

- A good sample produces a good image with less work
- Help eliminate rendering artifacts, such as aliasing

256 vs. 1

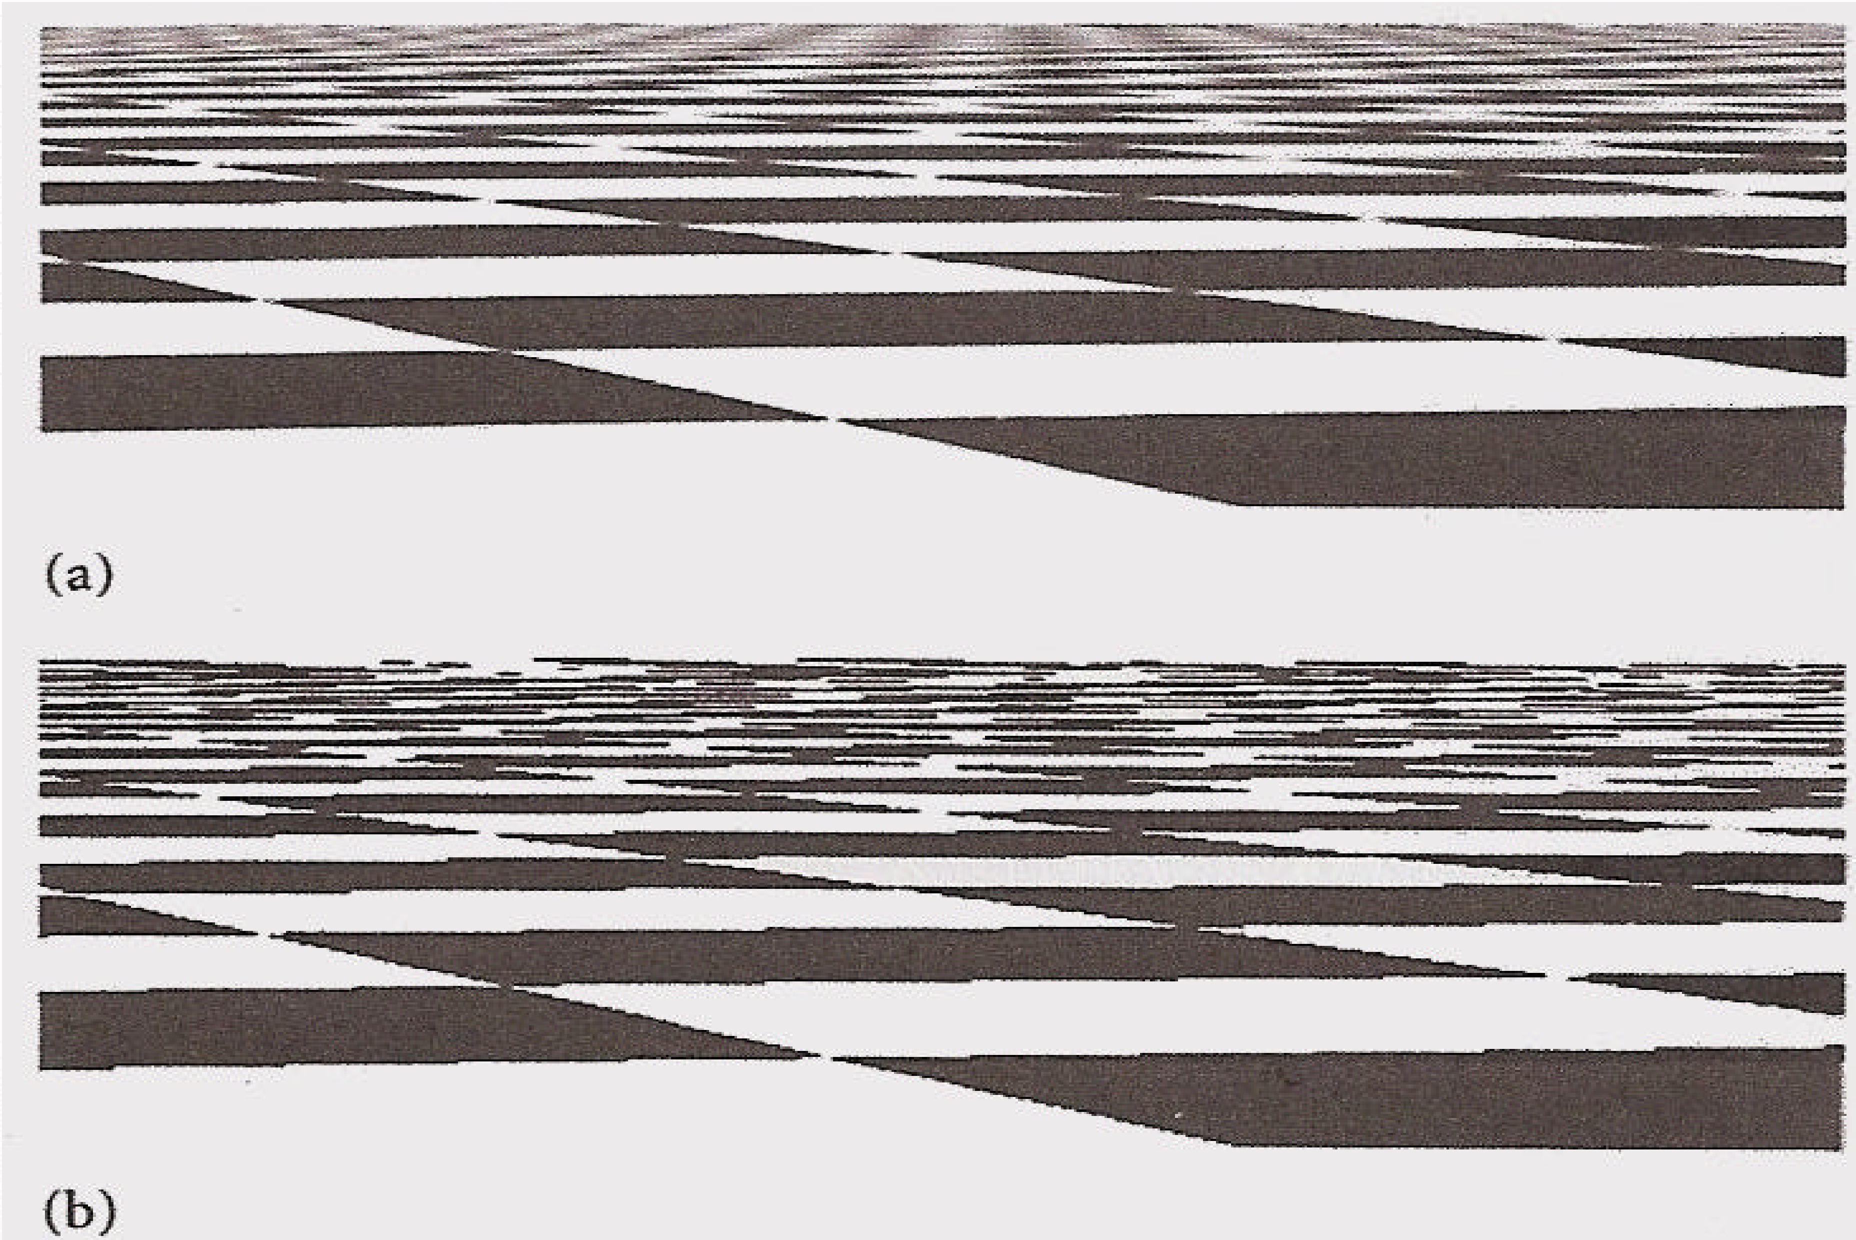

Strata Sampling Overview

- Image plane is evenly divided into non-overlapping rectangular regions called strata
- Within each strata a single sample is taken
- Sample point “jittered”

Jittering Effect

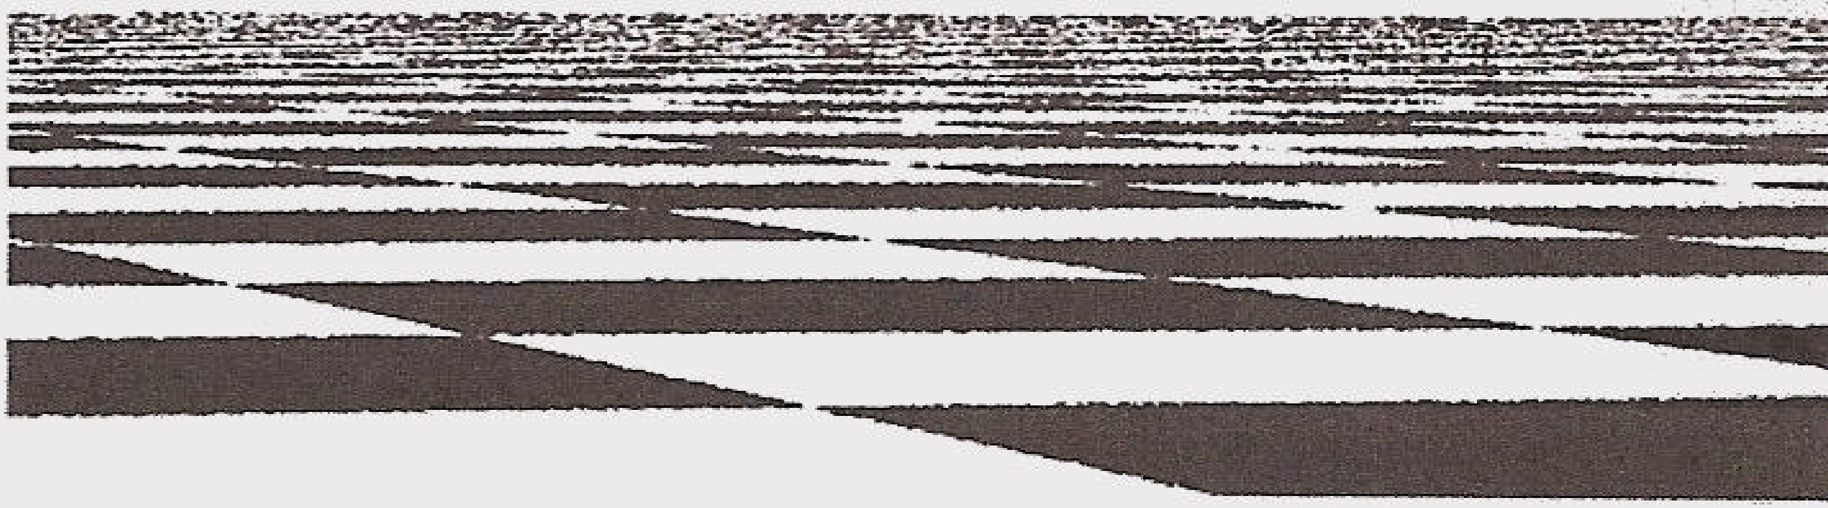

(c)

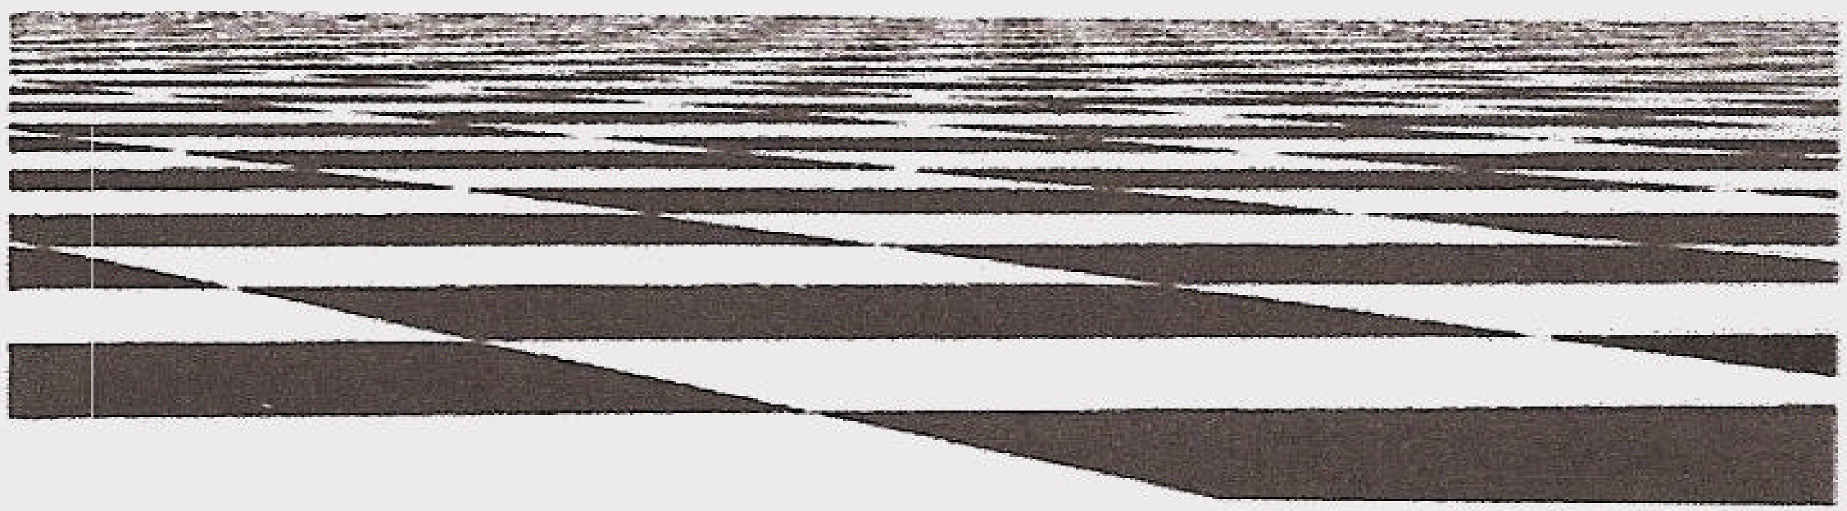

(d)

Advanced Strata Sampling

- Computes image, time, and lens samples for entire pixel all at once
- Beware of Curse of Dimensionality
- Compute patterns for each dimension and then randomly associate samples from each set of dimensions

Latin Hypercube Sampling

- Evenly divides region into grid and generates jittered sample along diagonal
- Samples are randomly shuffled such that only a single sample in each row and column
- Less effective than stratified sampling as grid size increases
- Good for randomizing samples along a single axis (such as shadows)

Low-Discrepancy Theory

Discrepancy: The number of samples within a fraction of a volume compared with the total number of samples in the total volume. The lower the difference, the lower the discrepancy.

Radical Inverse

A positive integer *n* can be expressed in a base *b* with a sequence of digits as follows:

$$N = ?\ d_i b^{i-1}$$

Over the set i = 1 to Infinity

- $F_b(n) = 0.d_1 d_2 d_3 \ldots d_m$
- van der Corput sequence, one-dimensional Radical Inverse in base 2
  - 0.1, 0.01, 0.11, 0.001, 0.101, ...

Halton Sequence

- Uses radial inverse base b, with a different base for each dimension of the pattern
- Bases are increasing prime numbers
- Very good even if total number of samples isn’t known in advance
  - Photon Integrator

$$x_i = (F_2(i), F_3(i), F_5(i), \ldots, F_{p(n)}(i))$$

Hammersley Sequence

- Requires a fixed number of samples, but produces slightly lower discrepancy than Halton

$$x_i = (i/N, F_2(i), F_3(i), F_5(i), \ldots, F_{p(n)}(i))$$

Folded Radical Inverse

$$? _b(n) = ? \ ((d_i + i - 1) \bmod b) \ / \ (b^i))$$

- Halton -> Halton-Zaremba
- Hammersley -> Hammersley-Zaremba

(0,2)-Sequences

- Uses van der Corput in one dimension and a radical inverse sequence in the other dimension
- Each sample is well-stratified with other samples in the same pixel, as well as sample positions in other image samples around the pixel
- Pbrt uses a scrambled (0, 2)-sequence in its LD Sampler

The LD Sampler

- Could use a Hammersley sequence in all dimensions, but Hammersley is prone to aliasing in images
- Instead PBRT uses a (0, 2)-sequence

Strat vs. Hammersley

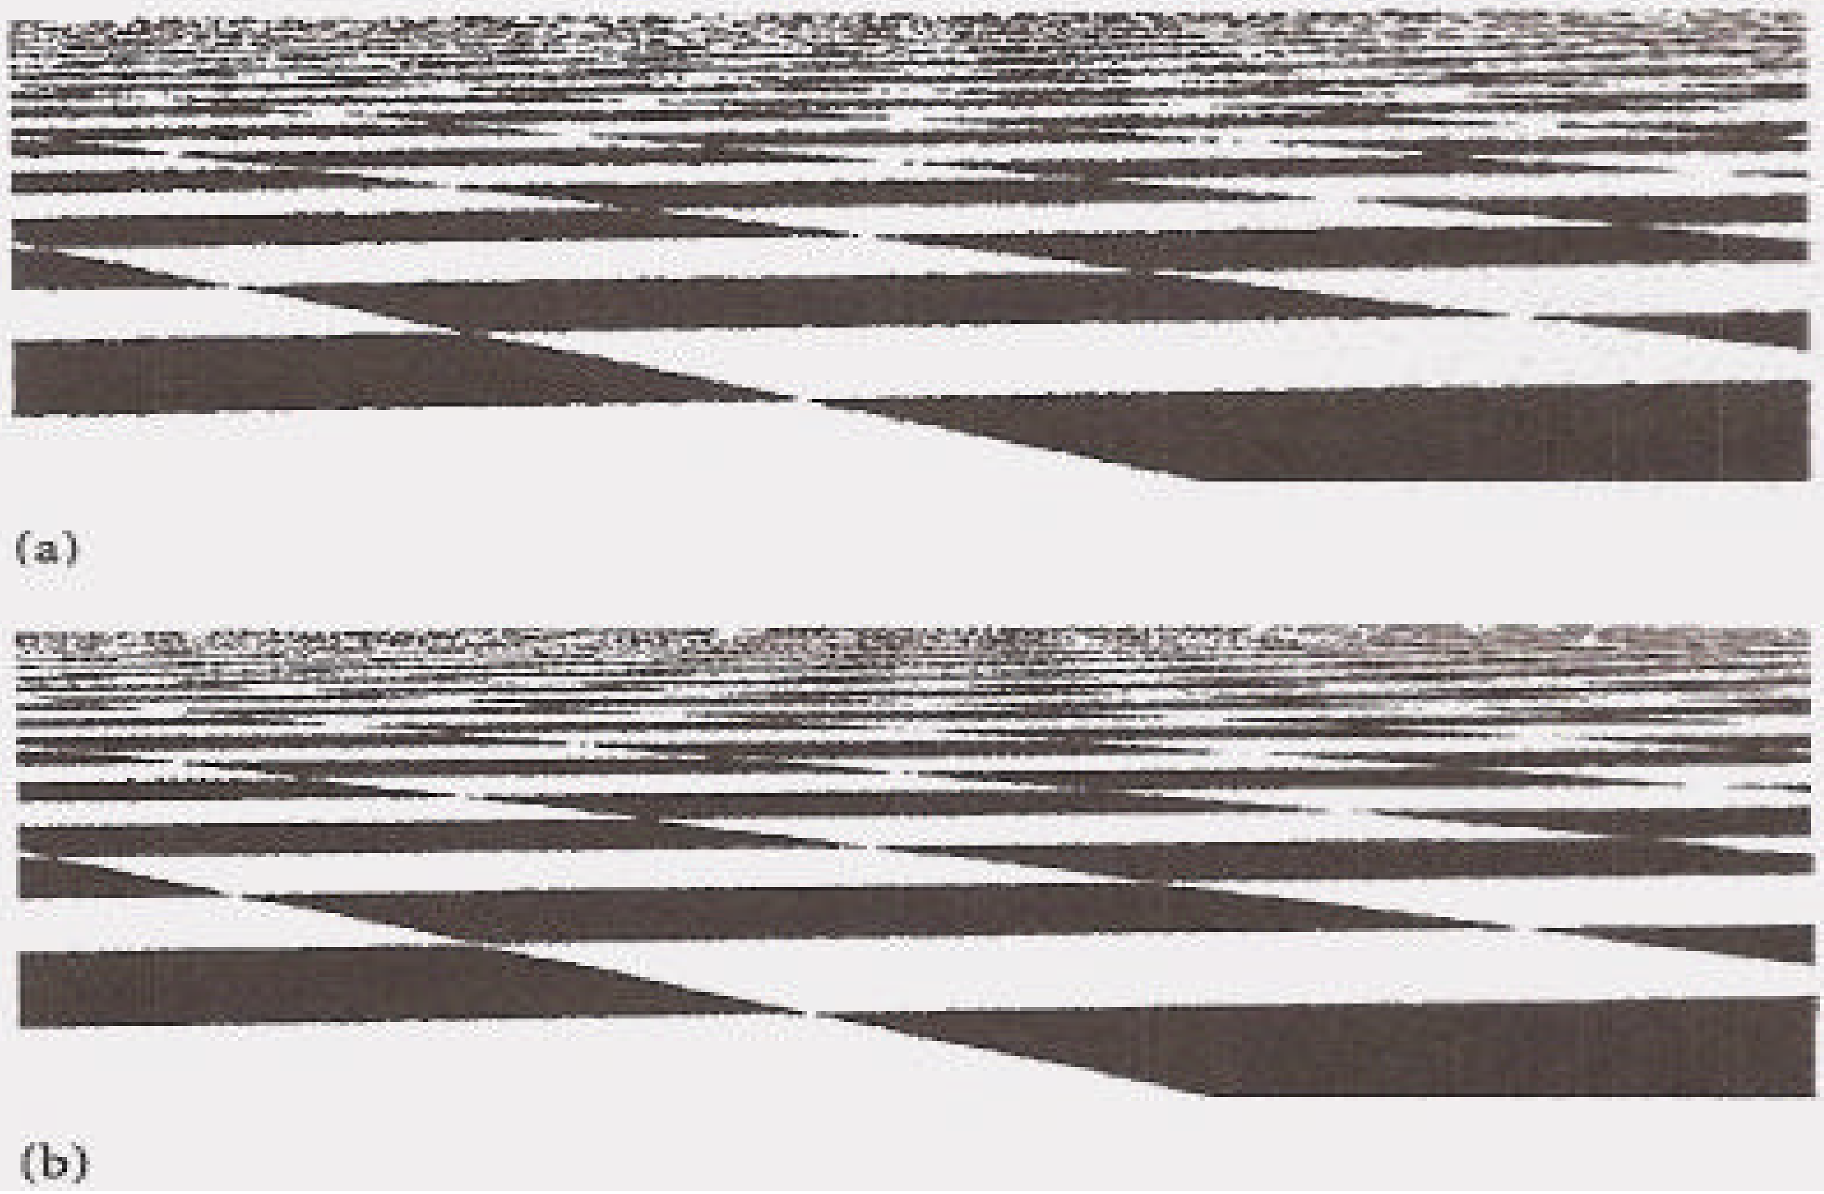

LD Sampler Internals

- (0, 2)-sequence to generate samples for 2-dimensional image and lens samples
- Van der Corput sequence for 1-dimensional time and integrator samples
- Number of samples must be power of 2
- Generates an entire pixel's worth of samples at once

Stratified vs. LD

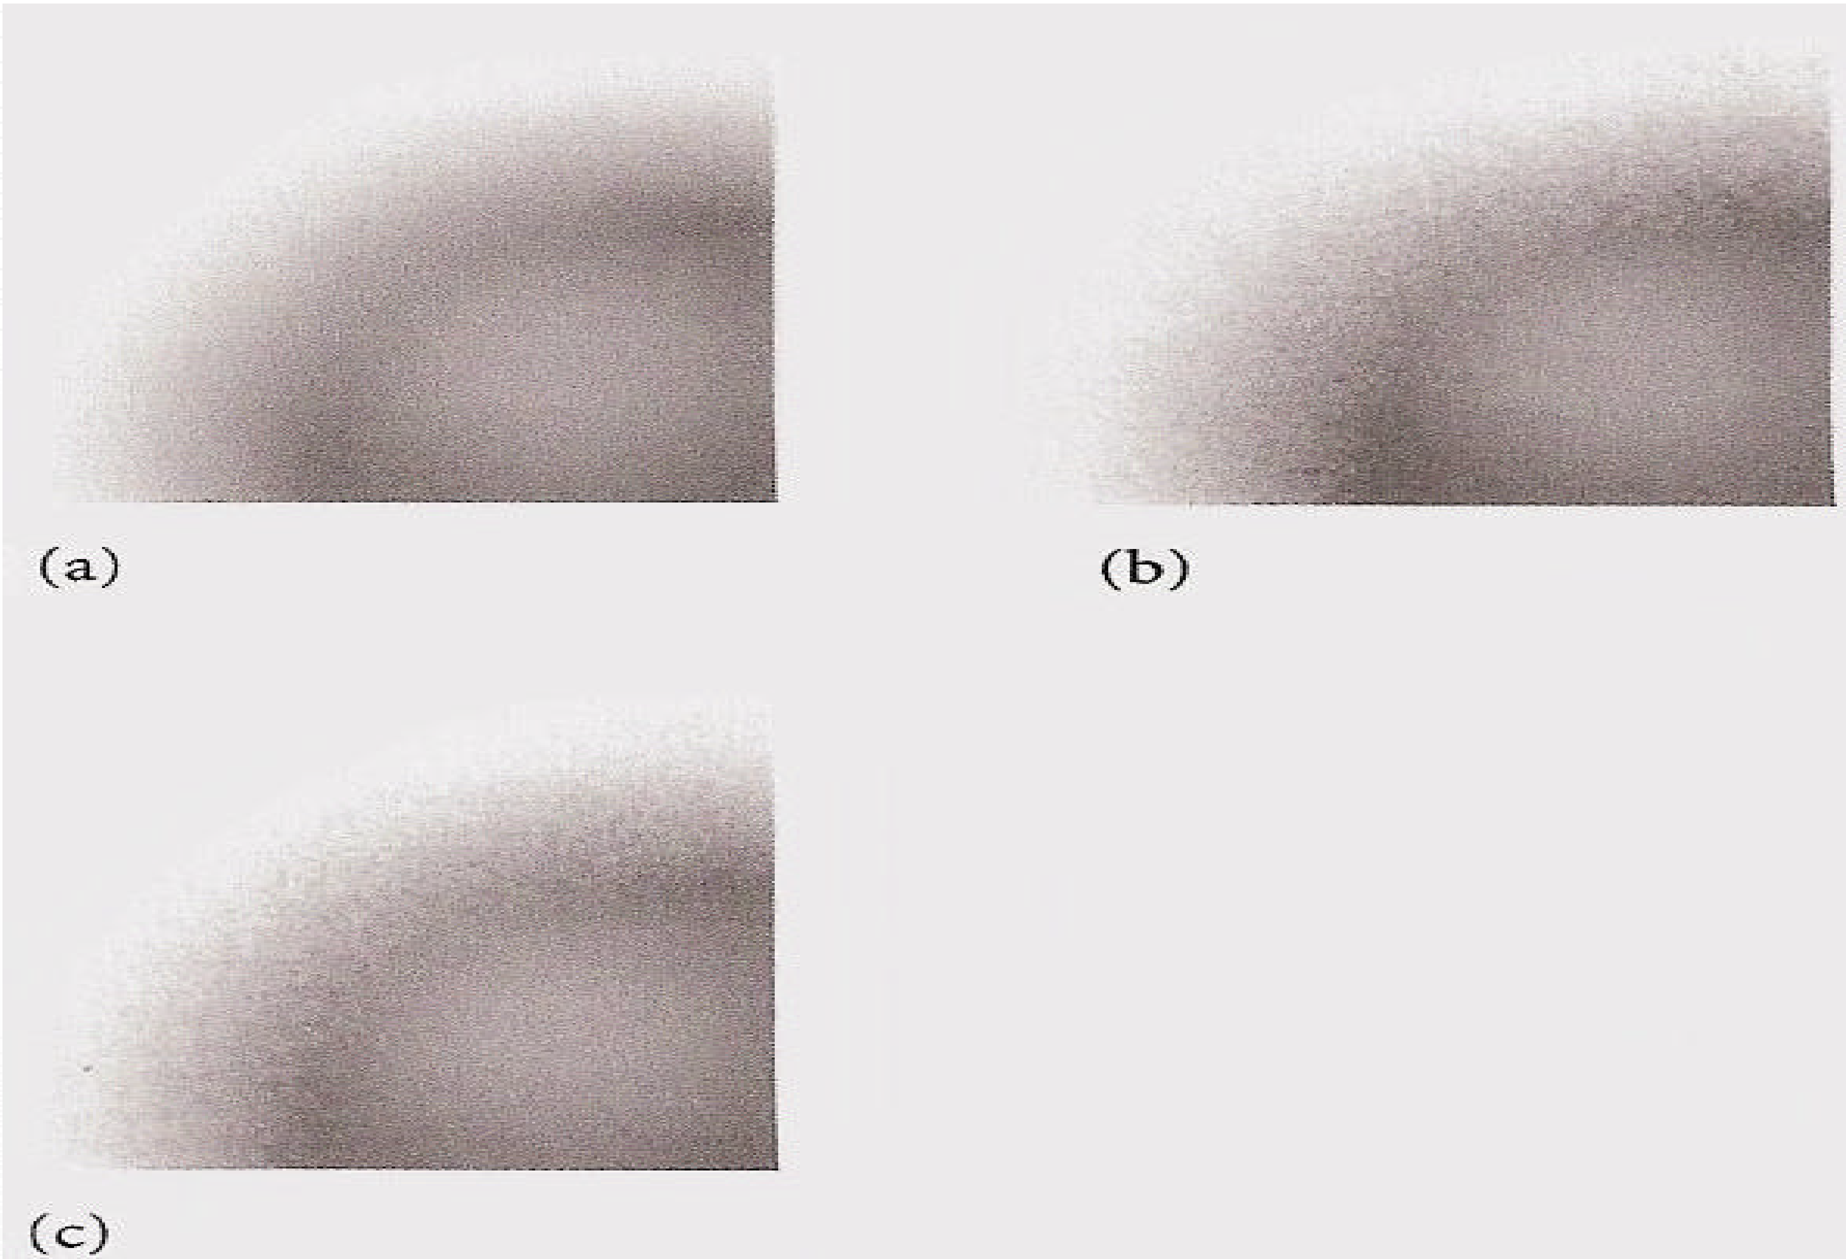

Random Shuffling

- In practice, sometimes correlation can still remain between sample elements even after random scrambling
- So, PBRT independently shuffles the sequences after creation

What Next?

- Both the Stratified Sampler and the LD Sampler generate good samples for a single pixel, but neither takes into account surrounding pixels
- Best-Candidate Sampling patterns overcome this shortcoming, but this is a topic for another week...