Lecture 1: How random number generators (don't) work

ACM 206, April 4, 2023

Samples inside a circle

Let's generate samples inside a circle.

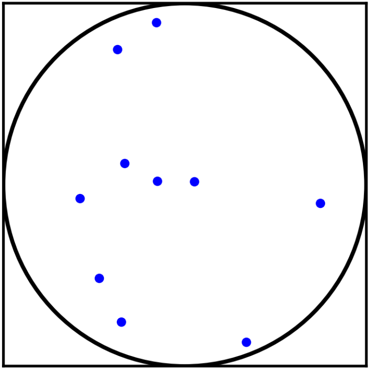

Figure: A Monte Carlo approach (10 samples)

Samples inside a circle

Let's generate samples inside a circle.

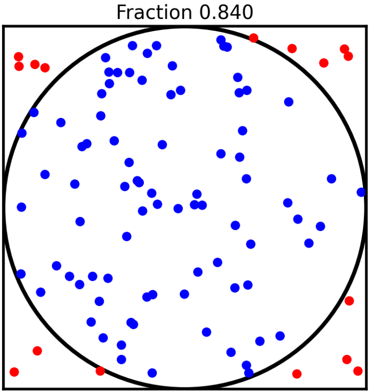

Figure: A Monte Carlo approach (100 samples)

Samples inside a circle

Let's generate samples inside a circle.

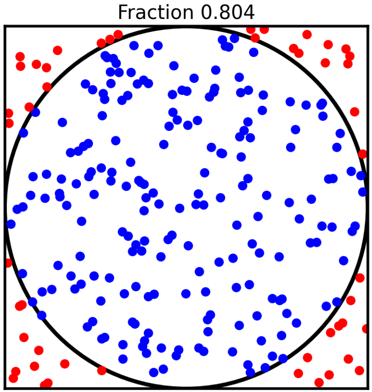

Figure: A Monte Carlo approach (250 samples)

Samples inside a circle

Let's generate samples inside a circle.

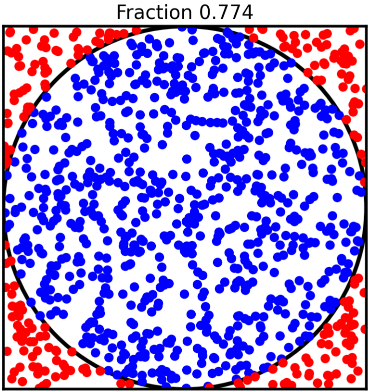

Figure: A Monte Carlo approach (1000 samples)

A simple rejection sampling approach

A simple rejection sampling approach

1. For $t = 1, 2, \ldots, T$, draw a random vector $X^t = (X_1^t, X_2^t)$ where
$$X_1^t, X_2^t \sim \text{Unif}\,(-1, 1)\,.$$
2. Accept if
$$(X_1^t)^2 + (X_2^t)^2 \leq 1.$$

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

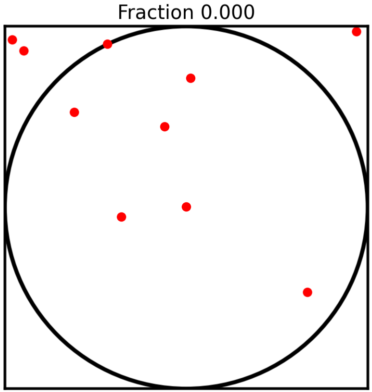

Figure: A Monte Carlo approach (10 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

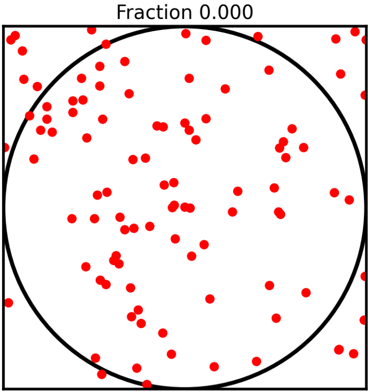

Figure: A Monte Carlo approach (100 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

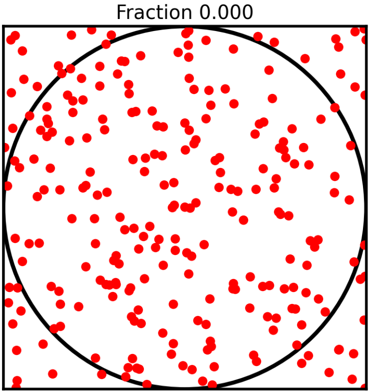

Figure: A Monte Carlo approach (250 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

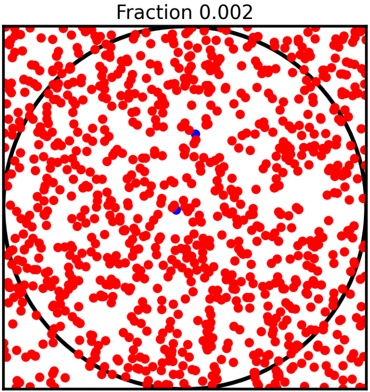

Figure: A Monte Carlo approach (1000 samples)

Higher dimensions

Here's a better algorithm for higher dimensions.

Higher dimensions

Here's a better algorithm for higher dimensions.

Don't throw it like this.

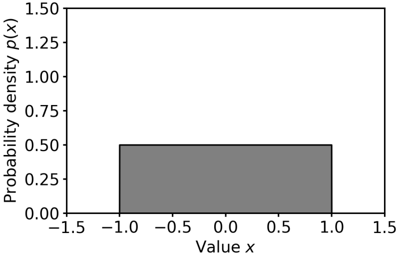

Figure: Unif$(-1, 1)$ distribution

Higher dimensions

Here's a better algorithm for higher dimensions.

Don't throw it like this.

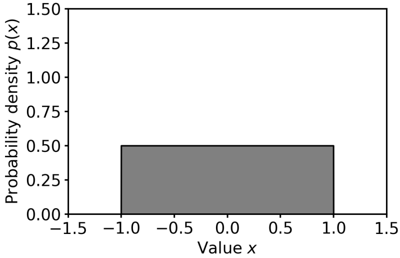

Figure: Unif$(-1, 1)$ distribution

Throw it like this instead.

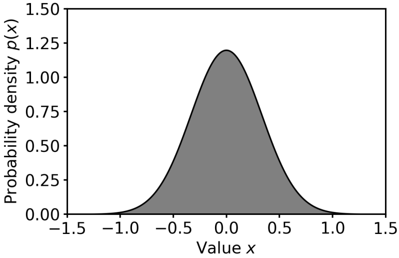

Figure: $\mathrm{N}(0, 1/(d-1))$ distribution

Higher dimensions

Here's a better algorithm for higher dimensions.

Higher dimensions

Here's a better algorithm for higher dimensions.

A faster Monte Carlo approach

1. For $t = 1, 2, \ldots, T$, draw a random vector $X^t = (X_1^t, \ldots, X_d^t)$, where

$$X_1^t, \ldots, X_d^t \sim \mathcal{N}(0, \tfrac{1}{d-1}).$$

2. Accept with probability

$$\exp\left(\frac{d-1}{2}((X_1^t)^2 + \cdots + (X_d^t)^2 - 1)\right) \mathbb{1}\left\{(X_1^t)^2 + \cdots + (X_d^t)^2 \leq 1\right\}$$

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

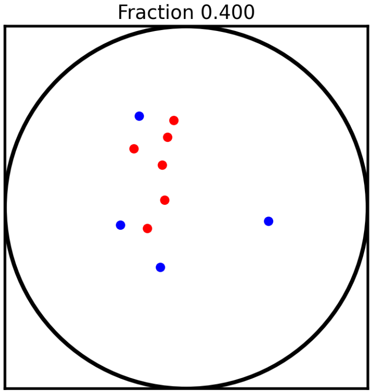

Figure: A faster Monte Carlo approach (10 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

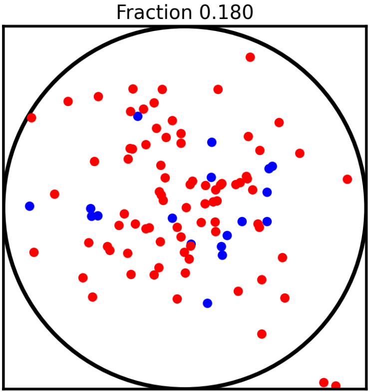

Figure: A faster Monte Carlo approach (100 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

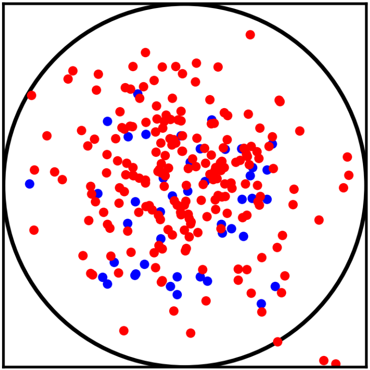

Figure: A faster Monte Carlo approach (250 samples)

10-dimensional hypersphere

Let's generate samples inside a 10-dimensional hypersphere.

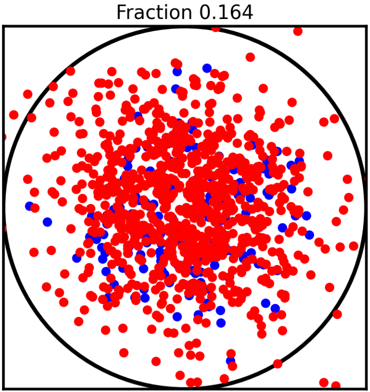

Figure: A faster Monte Carlo approach (1000 samples)

Conclusion

Why is the Gaussian sampler so effective?

Conclusion

Why is the Gaussian sampler so effective?

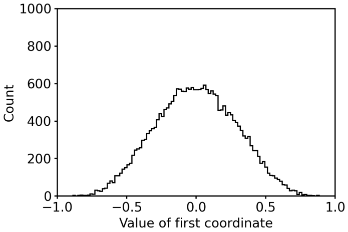

Figure: Each coordinate is quite close to a $N\left(0, \frac{1}{d-1}\right)$ distribution.

Conclusion

Why is the Gaussian sampler so effective?

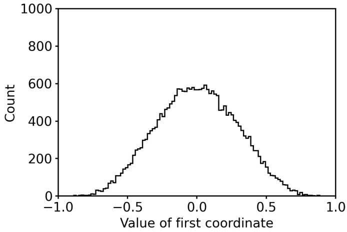

Figure: Each coordinate is quite close to a $N\left(0, \frac{1}{d-1}\right)$ distribution.

Takeaway message: to speed up Monte Carlo, use everything you know about the events you're sampling!

Conclusion

Why is the Gaussian sampler so effective?

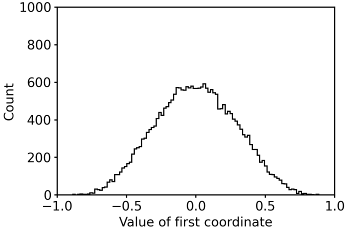

Figure: Each coordinate is quite close to a $N\left(0, \frac{1}{d-1}\right)$ distribution.

Takeaway message: to speed up Monte Carlo, use everything you know about the events you're sampling!

Faster option: There is also a perfect sampler (Barthe et al., 2005):

$$(Z_1, \dots, Z_n)/\left(\sum Z_i^2 + S\right)^{1/2}$$

where Z_i are independent Gaussians, $S \sim \exp(1/2)$.