Sampling from nonuniform distributions

CE 391F

April 9, 2013

REVIEW

What kinds of distributions are we considering in this class?

How do we generate a $U(0,1)$ number?

What makes for a good random number generator?

OUTLINE

Generating other kinds of random number

1. Uniform between A and B
2. Uniform integers
3. Generic distributions
4. Exponential distribution
5. Standard normal distribution
6. Nonstandard normal distribution

From here on out, we'll assume that we have a good generator for $U(0, 1)$ variates, which we will refer to as U (possibly with subscripts).

OTHER TYPES OF UNIFORM DISTRIBUTION

What if we wanted a random number uniformly drawn between any real numbers A and B, where $A \leq B$?

Not difficult, generate a $U(0,1)$ and scale it:

$$A + (B - A)U$$

What if we wanted a random *integer* uniform between two *integers* C and D, inclusive ($C \leq D$)?

Requires a little bit more care. Define the "floor" function $\lfloor x \rfloor$ to be the greatest integer less than or equal to x.

Then we can do

$$\lfloor (D - C + 1)U + C \rfloor$$

What about sampling from non-uniform distributions?

Let's try to leverage the generator U. What else ranges from 0 to 1 in the context of probability distributions?

One idea is the cdf. If X has cdf $F_X(x)$, it ranges between 0 and 1.

Furthermore, wherever the pdf is positive, F_X is strictly increasing and thus invertible.

So, one idea is to generate U, then see what value of x satisfies $F_X(x) = U$.

This method actually works, and is called "inverse transform sampling."

Proof: Not that difficult.

- $P(U \leq p) = p$ for any $p \in [0, 1]$.
- We want $P(X \leq x) = F_X(x)$ for any x.
- Fix x. Since $F_X(x) \in [0, 1]$, $P(U \leq F_X(x)) = F_X(x)$
- Therefore $P(F_X^{-1}(U) \leq x) = F_X(x)$
- So, if $X = F_X^{-1}(U)$, then X has the right distribution.

Note that this requires F_X to be invertible. This isn't a major problem with continuous random variables, since F_X is only noninvertible in regions with zero probability. This derivation relies on the fact that F is nondecreasing. Why?

We can use this to develop a formula for sampling the exponential distribution.

Recall that the *exponential* distribution with mean λ has the pdf $f_X(x) = \frac{1}{\lambda}\exp(-x/\lambda)$ if $x \geq 0$ and 0 otherwise, and cdf $F_X(x) = 1 - \exp(-x/\lambda)$ if $x \geq 0$ and 0 otherwise.

What is the inverse of the CDF?

Directly applying inverse transform sampling would give $X = -\lambda \log(1 - U)$.

We can simplify further... if U is $U(0, 1$, so is $1 - U$, so the formula can be reduced to $X = -\lambda \log U$.

Can we visualize what inverse transform sampling is doing?

Inverse transform sampling works for any distribution where we can compute the inverse cdf. Usually, that's not a problem, but in one common case it is:

The *normal* distribution with mean mean μ and variance σ^2 has the pdf $f_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-(x-\mu)^2/2\sigma^2)$.

Try any integration tricks you like, there is no way to obtain the cdf in closed form (let alone its inverse).

The Box-Muller method is an ingenious way of generating normal samples from $U(0, 1)$ samples.

Oddly enough, even though $\int \exp(-x^2)\ dx$ can't be written in closed form, we *can* show that

$$\int_{-\infty}^{\infty} \exp(-x^2)\ dx = \sqrt{\pi}$$

The trick is to introduce a second copy of the pdf, convert to polar coordinates, integrate, and then take the square root to get the original pdf back.

A similar method can be applied to generate independent standard normal samples, when $\mu = 0$ and $\sigma^2 = 1$.

We will use two $U(0,1)$ samples U_1 and U_2 to generate two standard normal samples X and Y

At a high level, we will use U_1 and U_2 to generate R^2 and Θ using inverse transform sampling, then convert to rectangular coordinates to obtain Z_1 and Z_2.

If we multiply the pdfs for X and Y together (valid since they are independent), we obtain

$$f_{X,Y}(x,y) = \frac{1}{2\pi}\exp(-(x^2+y^2))$$

We then apply the transformation $r^2 = x^2 + y^2$, $\theta = \tan^{-1}(y/x)$

Recall that when transforming PDFs, we need to multiply by the determinant of the transformation's Jacobian:

$$|J| = \begin{vmatrix} \frac{\partial r^2}{\partial x} & \frac{\partial r^2}{\partial y} \\ \frac{\partial \theta}{\partial x} & \frac{\partial \theta}{\partial x} \end{vmatrix}$$

(This is the same thing that happens when we apply transformations to multiple integrals.)

It turns out that $|J| = 2$, so the appropriate joint pdf for R^2 and Θ is

$$f_{R,\Theta}(r,\theta) = \frac{2}{2\pi}\exp(-r^2/2)$$

We can factor this into marginal pdfs for R^2 and Θ:'

$$f_{\Theta}(\theta) = \frac{1}{2\pi}$$

$$f_{R^2}(r^2) = 2\exp(-r^2/2)$$

Notice that Θ has a uniform distribution between 0 and 2π, and R^2 has an exponential distribution with mean 2.

We know how to generate both of those, and from here the rest of the procedure is simple:

- Generate two uniform(0,1) variates U_1 and U_2.
- Generate $\Theta = 2\pi U_1$
- Generate $R^2 = -2 \log U_2$.
- Generate $X = R \cos \Theta$ and $Y = R \sin \Theta$

Presto! You know have two random numbers X and Y with independent standard normal distributions.

What about a nonstandard normal distribution with arbitrary μ and σ^2?

Other distributions

In traffic simulation, we are often concerned with other types of distributions in addition to the uniform. A typical example:

- Vehicles enter the facility according to a Poisson process (i.e., time between new vehicles is exponentially distributed)
- v_{max} for different vehicles follows a uniform distribution between 4 and 6.
- l for different vehicles follows a normal distribution with mean 10 and standard deviation 2.
- and so on...

If we were doing a car-following simulation, we may have distributions for T and λ, and each time a vehicle is generated, we need to pick an appropriate value.