

Probability Overview

■ Initialize standard library files:

```
Off[General::spell1];
```

The next package is needed for the add-on multivariate gaussian

Goals

Review the basics of probability distributions and statistics

More on generative modeling: drawing samples

Graphical models for inference

Optimal inference and Task dependence

Probability overview

Random variables, discrete probabilities, probability densities, cumulative distributions

■ Discrete: random variable X can take on a finite set of discrete values

$X = \{x(1), \dots, x(N)\}$

$$\sum_{i=1}^N p_i = \sum_{i=1}^N p(X = x(i)) = 1$$

■ Densities: X takes on continuous values, x, in some range.

Density : $p(x)$

Analogous to material mass,

we can think of the probability over some small domain of the random variable as "probability mass" :

$$\begin{aligned} \text{prob}(x < X < dx + x) &= \int_x^{x+dx} p(x) dx \\ \text{prob}(x < X < dx + x) &\approx p(x) dx \end{aligned}$$

With the mass analogy, however, an object (event space) always "weighs" 1 :

$$\int_{-\infty}^{\infty} p(x) dx = 1$$

Cumulative distribution:

$$\text{prob}(X < x) = \int_{-\infty}^x p(X) dX$$

■ Densities of discrete random variables

The Dirac Delta function, $\delta[\bullet]$, allows us to use the mathematics of continuous distributions for discrete ones, by defining the density as:

$$p[x] = \sum_{i=1}^N p_i \delta[x - x[i]], \text{ where } \delta[x - x[i]] = \begin{cases} \infty & \text{for } x = x[i] \\ 0 & \text{for } x \neq x[i] \end{cases}$$

Think of the delta function, $\delta[\bullet]$, as ϵ wide and $1/\epsilon$ tall, and then let $\epsilon \rightarrow 0$, so that:

$$\int_{-\infty}^{\infty} \delta(y) dy = 1$$

The density, $p[x]$, is a series of spikes. It is infinitely high only at those points for which $x = x[i]$, and zero elsewhere. But "infinity" is scaled so that the local mass or area around each point $x[i]$, is p_i .

■ Joint probabilities

Prob (X AND Y) = $p(X, Y)$

Joint density : $p(x, y)$

Three basic rules of probability

Suppose we know everything there is to know about a set of variables (A,B,C,D,E). What does this mean in terms of probability? It means that we know the joint distribution, $p(A,B,C,D,E)$. In other words, for any particular combination of values (A=a,B=b, C=c, D=d,E=e), we can calculate, look up in a table, or determine some way or another the number $p(A=a,B=b, C=c, D=d,E=e)$.

Deterministic relationships are special cases.

■ Rule 1: Conditional probabilities from joints: The product rule

Probability about an event changes when new information is gained.

Prob(X given Y) = $p(X|Y)$

$$p(X | Y) = \frac{p(X, Y)}{p(Y)}$$

$$p(X, Y) = p(X | Y) p(Y)$$

The form of the product rule is the same for densities as for probabilities.

■ Rule 2: Lower dimensional probabilities from joints: The sum rule (marginalization)

$$p(X) = \sum_{i=1}^N p(X, Y(i))$$

$$p(x) = \int_{-\infty}^{\infty} p(x, y) dx$$

■ Rule 3: Bayes' rule

From the product rule, and since $p[X,Y] = p[Y,X]$, we have:

$$p(Y | X) = \frac{p(X | Y) p(Y)}{p(X)}, \text{ and using the sum rule, } p(Y | X) = \frac{p(X | Y) p(Y)}{\sum_Y p(X, Y)}$$

■ Bayes Terminology in inference

Suppose we have some partial data (see half of someone's face), and we want to recall or complete the whole. Or suppose that we hear a voice, and from that visualize the face. These are both problems of statistical inference. We've already studied how to complete a partial pattern using energy minimization, and how energy minimization can be viewed as probability maximization.

We typically think of the **Y** term as a random variable over the hypothesis space (a face), and **X** as data or a stimulus (partial face, or sound). So for recalling a pattern **Y** from an input stimulus **X**, We'd like to have a function that tells us:

$p(\mathbf{Y} | \mathbf{X})$ which is called the **posterior** probability of the hypothesis (face) given the stimulus (partial face or sound).

-- i.e. what you get when you condition the joint by the stimulus data. The posterior is often what we'd like to base our decisions on, because it can be proved that picking the hypothesis **Y** which maximizes the posterior (i.e. maximum a posteriori or **MAP** estimation) minimizes the average probability of error.

$p(\mathbf{Y})$ is the **prior** probability of the hypothesis. Some hypotheses are more likely than others. Given a context, such as your room, some faces are a priori more likely than others. For me an image patch stimulating my retina in my kitchen is much more likely to be my wife's than my brother's (who lives in another state). This shows that priors are contingent, i.e. conditional on context, $p(\mathbf{Y} | \text{context})$.

$p(\mathbf{X} | \mathbf{Y})$ is the **likelihood** of the hypothesis. Note this is a probability of **X**, but not of **Y**. (The sum over **X** is one, but the sum over **Y** isn't necessarily one.)

■ Bayes Terminology in visual perception

$$p[\mathbf{S} | \mathbf{I}] = \frac{p[\mathbf{I} | \mathbf{S}] p[\mathbf{S}]}{p[\mathbf{I}]}$$

Usually, we will be thinking of the **Y** term as a random variable over the hypothesis space, and **X** as data. So for visual inference, $\mathbf{Y} = \mathbf{S}$ (the scene), and $\mathbf{X} = \mathbf{I}$ (the image data), and $\mathbf{I} = \mathbf{f}(\mathbf{S})$.

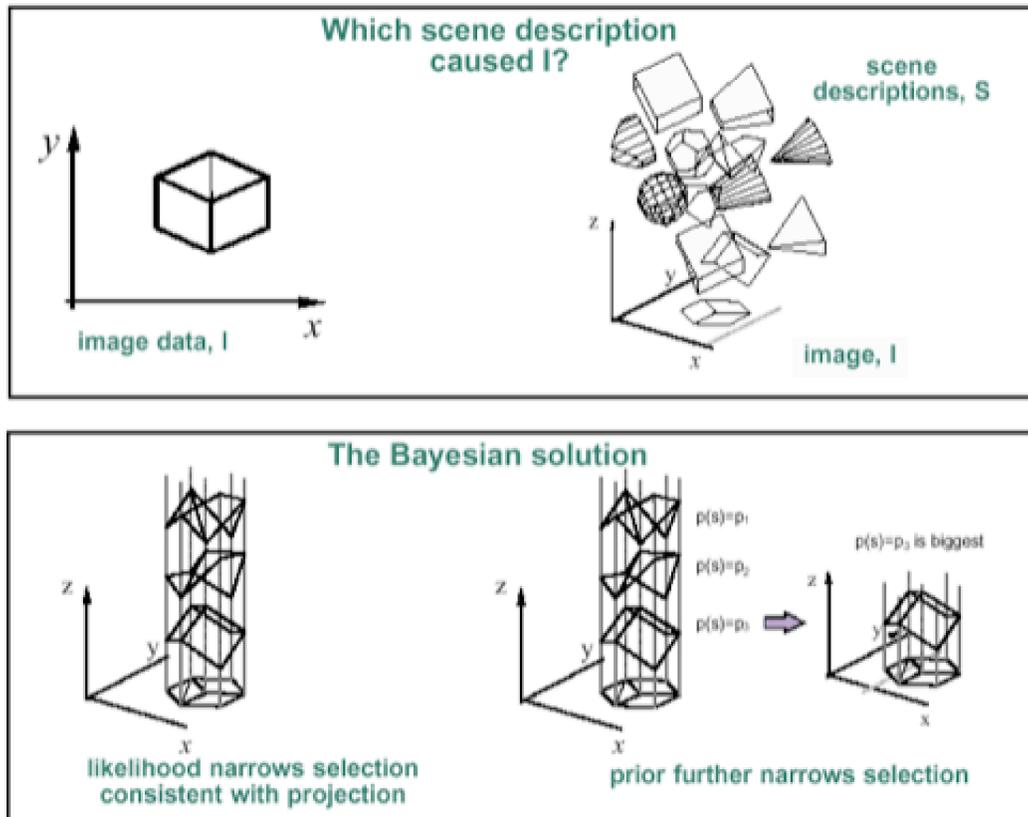
We'd like to have:

$p(\mathbf{S} | \mathbf{I})$ is the **posterior** probability of the scene given the image

-- i.e. what you get when you condition the joint by the image data. The posterior is often what we'd like to base our decisions on, because as we discuss below, picking the hypothesis **S** which maximizes the posterior (i.e. maximum a posteriori or **MAP** estimation) minimizes the average probability of error.

$p(\mathbf{S})$ is the **prior** probability of the scene.

$p(\mathbf{I} | \mathbf{S})$ is the **likelihood** of the scene. Note this is a probability of **I**, but not of **S**.



■ Independence

Knowledge of one event doesn't change the probability of another event.

$$p(X)=p(X|Y)$$

$$p(X,Y)=p(X)p(Y)$$

Density mapping theorem

Suppose we have a change of variables that maps a discrete set of x 's uniquely to y 's: $X \rightarrow Y$.

■ Discrete random variables

No change to probability function. The mapping just corresponds to a change of labels, so the probabilities $p(X)=p(Y)$.

■ Continuous random variables

Form of probability density function does change because we require the probability "mass" to be unchanged: $p(x)dx = p(y)dy$

Suppose, $y=f(x)$

$$p_Y(\mathbf{y}) \delta \mathbf{y} = p_X(\mathbf{x}) \delta \mathbf{x}$$

In higher dimensions, the transformation is done by multiplying the density by the Jacobian, the determinant of the matrix of partial derivatives of the change of coordinates.

$$p_Y(y) = \int \delta(y - f(x)) f^{-1}(x) p_X(x) dx$$

over each monotonic part of f .

Convolution theorem for adding rvs

Let x be distributed as $g(x)$, and y as $h(x)$. Then the probability density for $z=x+y$ is, $f(z)$:

$$f(z) = \int g(s) h(z-s) ds \quad (1)$$

Statistics

■ Expectation & variance

Analogous to center of mass:

Definition of expectation or average:

$$\text{Average}[X] = \bar{X} = E[X] = \sum x[i] p[x[i]] \sim \sum_{i=1}^N x_i / N$$

$$\mu = E[X] = \int x p(x) dx$$

Some rules:

$$E[X+Y] = E[X] + E[Y]$$

$$E[aX] = aE[X]$$

$$E[X+a] = a + E[X]$$

Definition of variance:

$$\sigma^2 = \text{Var}[X] = E[(X-\mu)^2] = \sum_{j=1}^N (p(x(j))) (x(j) - \mu)^2 = \sum_{j=1}^N p_j (x_j - \mu)^2$$

$$\text{Var}[X] = \int (x - \mu)^2 p(x) dx \sim \sum_{i=1}^N (x_i - \mu)^2 / N$$

Standard deviation:

$$\sigma = \sqrt{\text{Var}[X]}$$

Some rules:

$$\text{Var}[X] = E[X^2] - E[X]^2$$

$$\text{Var}[aX] = a^2 \text{Var}[X]$$

■ Covariance & Correlation

Covariance:

$$\text{Cov}[X, Y] = E[(X - \mu_X)(Y - \mu_Y)]$$

Correlation coefficient:

$$\rho[X, Y] = \frac{\text{Cov}[X, Y]}{\sigma_X \sigma_Y}$$

■ Cross and Autocovariance matrix

Suppose X and Y are vectors: $\{X_1, X_2, \dots\}$ and $\{Y_1, Y_2, \dots\}$

$$\text{Cov}[X_i, Y_j] = E[(X_i - \mu_{X_i})(Y_j - \mu_{Y_j})] \sim \frac{\sum_{n=1}^N (x_i^n - \mu_{X_i})(y_j^n - \mu_{Y_j})}{N}$$

$$\text{Autocov}[X_i, X_j] = E[(X_i - \mu_{X_i})(X_j - \mu_{X_j})] \sim \frac{\sum_{n=1}^N (x_i^n - \mu_{X_i})(x_j^n - \mu_{X_j})}{N}$$

In other words, the autocovariance matrix can be approximated by the average outer product. It is a Hebbian matrix memory of pair-wise relationships.

■ Independent random variables

If $p(X, Y) = p(X)p(Y)$, then

$$E[XY] = E[X]E[Y] \quad (\text{uncorrelated})$$

$$\text{Cov}[X, Y] = \rho[X, Y] = 0$$

$$\text{Var}[X + Y] = \text{Var}[X] + \text{Var}[Y]$$

If two random variables are uncorrelated, they are not necessarily independent.

Two random variables are said to be orthogonal if their correlation is zero.

Degree of belief vs., relative frequency

What is the probability that the Vikings will win the Superbowl in 2004? Assigning a number between 0 and 1 is assigning a degree of belief. These probabilities are also called subjective probabilities. "Odds" determine subjective probabilities, where the "odds of \mathbf{x} to \mathbf{y} " means probability = $\mathbf{x}/(\mathbf{x}+\mathbf{y})$.

What is the probability that a coin will come up heads? In this case, we can do an experiment. Flip the coin n times, and count the number of heads, say $h[n]$, and then set the probability, $p = h[n]/n$ -- the relative frequency. Of course, if we did it again, we may not get the same estimate of p . One solution often given is:

$$p = \lim_{n \rightarrow \infty} \frac{h(n)}{n}$$

A problem with this, is that there is no guarantee that a well – defined limit exists.

In some domains we can measure statistics, and model probabilities of both inputs and outputs. So the relative frequency interpretation seems reasonable. In practice, the dimensions of many problems in perception, cognition, and memory are so high, that it is impractical to do this. Once we use the statistical framework to model perception, say of a particular cue (say), then probabilities are more like "subjective unconscious beliefs".

Principle of insufficient reason

■ Principle of symmetry

Suppose we have N events, $x[1], x[2], x[3], \dots, x[N]$ that are all physically identical except for the label. Then assume that

$$\text{prob}(x(1)) = \text{prob}(x(2)) = \text{prob}(x(3)) = \text{prob}(x(N)) = \frac{1}{N}$$

In other words, if we have no additional information about the events, we should assume that they are uniformly distributed. I.e., assume a uniform prior.

What about the continuous case where there is no reason to assume any particular value at all between $-\infty$ and $+\infty$?

Improper priors.

■ Information theory and Maximum entropy

Information theory provides a powerful extension to the principle of symmetry. Information of event X is:

$$\text{Information}[X] = -\log_2(p(X))$$

Using the definition of expectation above, we can specify the expectation of information, which is called entropy. Entropy of a random variable X with probability distribution $p[X]$ is:

$$H(X) = \text{Average}(\text{Information}[X]) = -\sum_x p(X) \log_2(p(X))$$

It can be shown that out of all possible probability distributions, $H(X)$ is biggest for the uniform distribution, $p(X)=1/N$. Maximum entropy is looking like symmetry.

It turns out that a more powerful formulation of the principle of symmetry is maximum entropy. For example, out of all possible probability distributions of a random variable with infinity range, but with a specific mean and standard deviation, the Gaussian is unique in having the largest entropy. If the range goes from zero to infinity, and we know the mean, the maximum entropy distribution is an exponential (Cover and Thomas).

An interesting application of the maximum entropy principle is to learning image textures joint probabilities: $p(I[1], \dots, I[N])$, where N is very big, but where one has only a relatively small number of measured statistics relative to the number of possible images (which is really huge). The measurements underdetermine the dimensionality of the probability space-- i.e. there are many different probability distributions which give the same statistics. So the principle of symmetry, or insufficient reason, says to choose the one with the maximum entropy.

More on sampling, generative modeling: Multivariate gaussian, mixtures

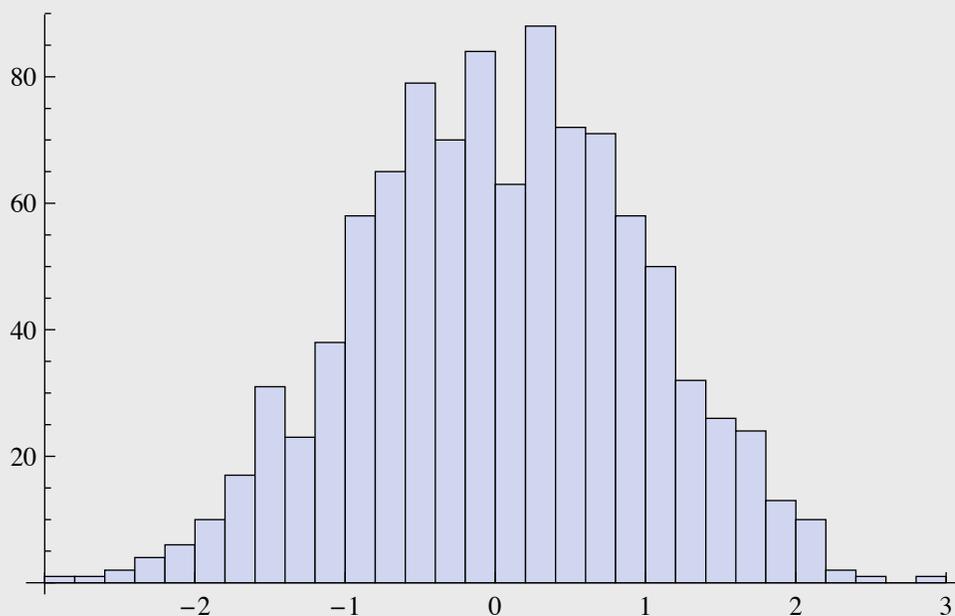
Making a univariate (scalar) gaussian random number generator: Method 1: Just for Gaussian

■ Use Central Limit Theorem

If all we want to do is make a Gaussian random number generator from a uniformly distributed generator, we can use the Central Limit Theorem.

```
Needs["BarCharts`"]  
Needs["Histograms`"]
```

```
nusamples = 10;  
z1 = Table[  
   $\sum_{i=1}^{\text{nusamples}} \text{RandomReal}[] - \frac{\text{nusamples}}{2}$ , {1000}];  
Histogram[z1]
```



Method 2: Use Density Mapping theorem. More general.

We'll use the density mapping theorem to turn uniformly distributed random numbers `Random[]` into gaussian distributed random numbers with mean =0 and standard deviation =1.

$$p_Y (Y) \delta Y = p_X (X) \delta X$$

$$p_Y (Y) \frac{\delta Y}{\delta X} = p_X (X)$$

Suppose $p_Y (Y) = 1$ (over the unit interval, but zero elsewhere). Then

$$Y (X) = \int_{-\infty}^X p_X (X') dX' = P (X) \tag{2}$$

Thus if we sample from the uniform distribution to get y, x should be distributed according to $p_X (X)$. To do this, we need a mapping from y->x. This is given by the inverse cumulative distribution, i.e. $P^{-1}(y)$.

Let's implement this. The quick way is to use *Mathematica's* built-in function to get the inverse cumulative.

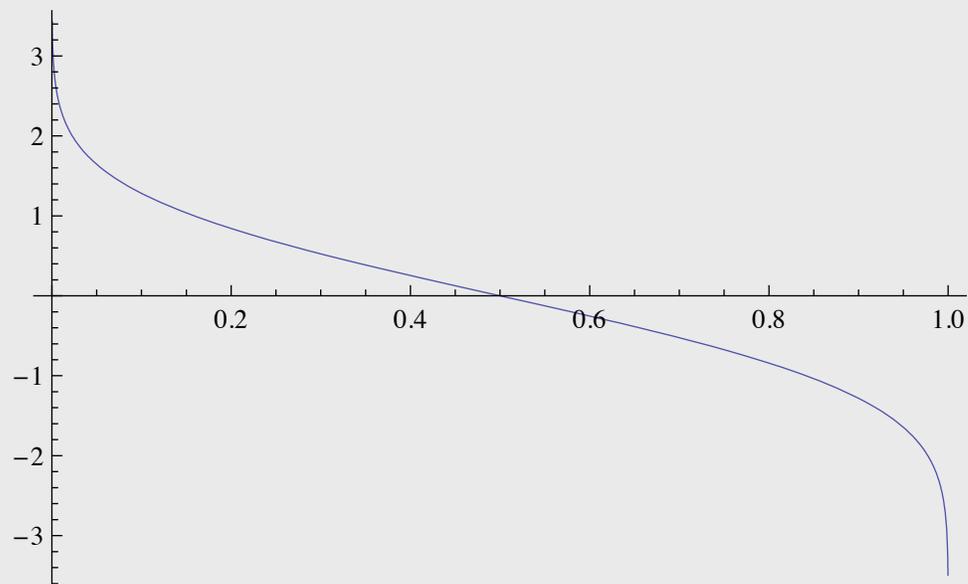
Method 2a: Applied to Gaussian

`InverseErf[]` is the inverse of :

$$\text{erf} (z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$$

We can use this to define a function for the inverse cumulative of a gaussian:

```
zscore[p_] :=  $\sqrt{2}$  InverseErf[1 - 2 p];  
Plot[zscore[y], {y, 0, 1}]
```

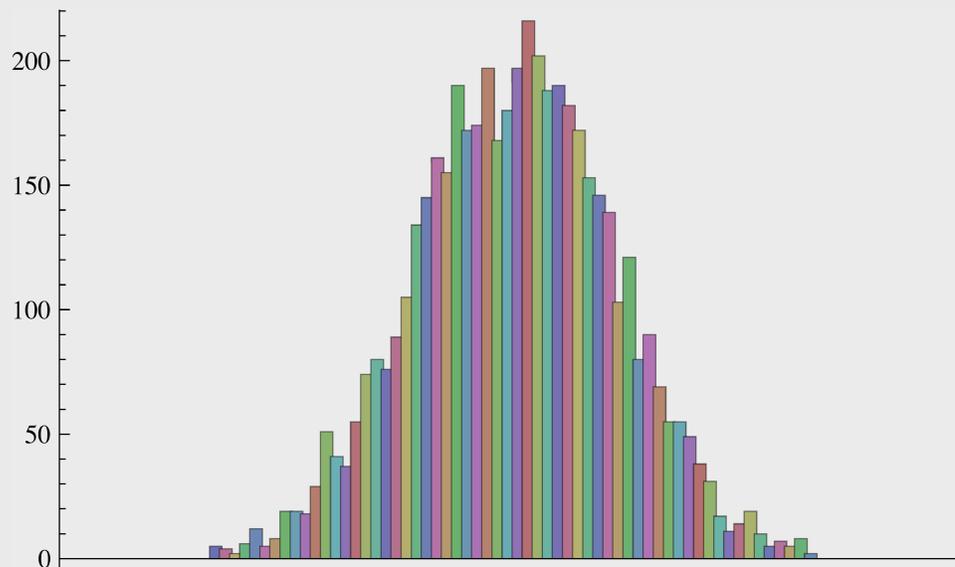


...this time instead of using Histogram[], we could make more work for ourselves, and write our own lines to calculate the histogram--useful to see if you've never calculated one before:

```

binsize = 0.1;
midpoints = Table[i +  $\frac{\text{binsize}}{2}$ , {i, -3, 3 - binsize, binsize}];
z1 = Table[zscore[RandomReal[]], {5000}];
freq = BinCounts[z1, {-3, 3, binsize}];
BarChart[Transpose[{freq}], BarLabels -> None]

```



Method 2b: From scratch: Works for almost any distribution. Do the Gaussian case for a third time

Suppose we have a discrete representation of any cumulative distribution. How can we generate samples? We'll illustrate the method with a discretization of the Gaussian. Our immediate goal is to produce a discrete approximation to the cumulative gaussian. To review where things come from, we'll start with the definition of a Gaussian, and make sure it is normalized.

```

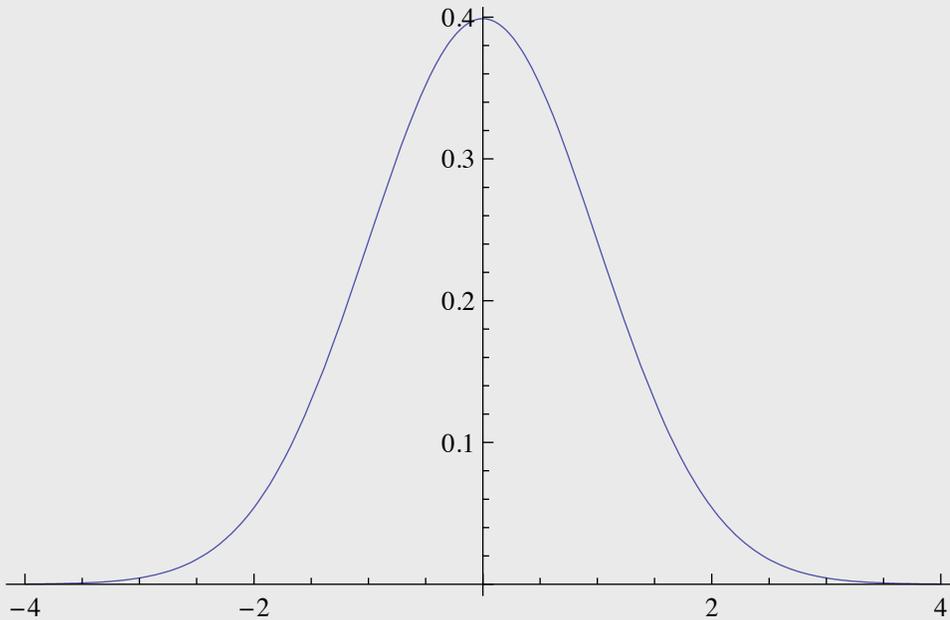
Integrate[Exp[-(x - x0)^2 / (2 * sigma^2)], {x, -Infinity, Infinity}]

```

$$\text{If}[\text{Re}(\sigma^2) > 0, \frac{\sqrt{2\pi}}{\sqrt{\frac{1}{\sigma^2}}}, \text{Integrate}\left[e^{-\frac{(x-x_0)^2}{2\sigma^2}}, \{x, -\infty, \infty\}, \text{Assumptions} \rightarrow \text{Re}(\sigma^2) \leq 0\right]]$$

Let $x_0=0$ and $\sigma=1$:

```
Plot[ $\frac{e^{-\frac{x1^2}{2}}}{\sqrt{2\pi}}$ , {x1, -4, 4}]
```



Plot[PDF[NormalDistribution[0,1],x1],{x1,-4,4}]; gives the same thing using the add-on normal distribution function.

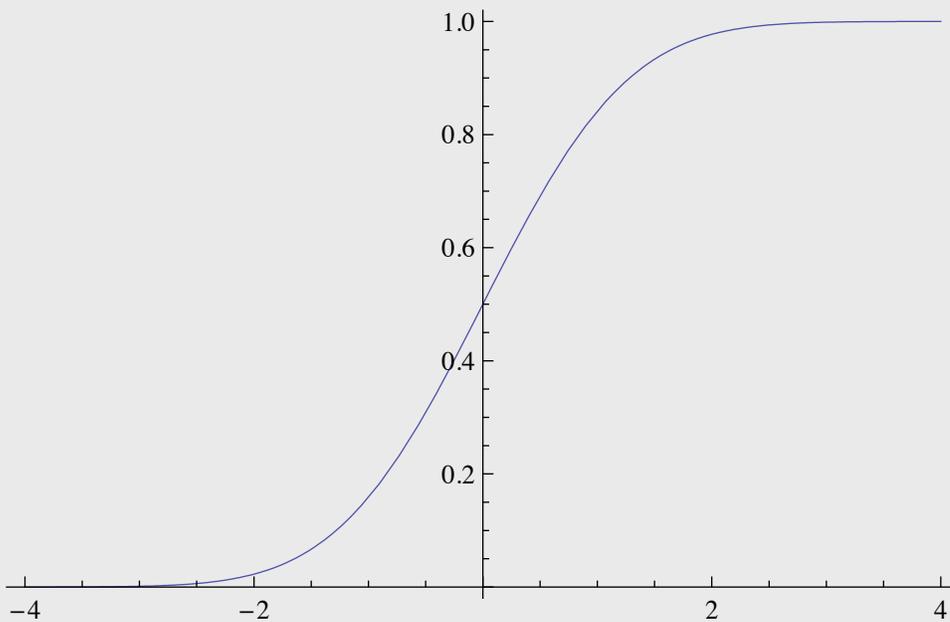
■ Cumulative gaussian

```
Clear[cumulgauss, x, x1];  
cumulgauss[x_] := Integrate[Exp[-(x1^2)/2] / (Sqrt[2 * Pi]),  
  {x1, -Infinity, x}]
```

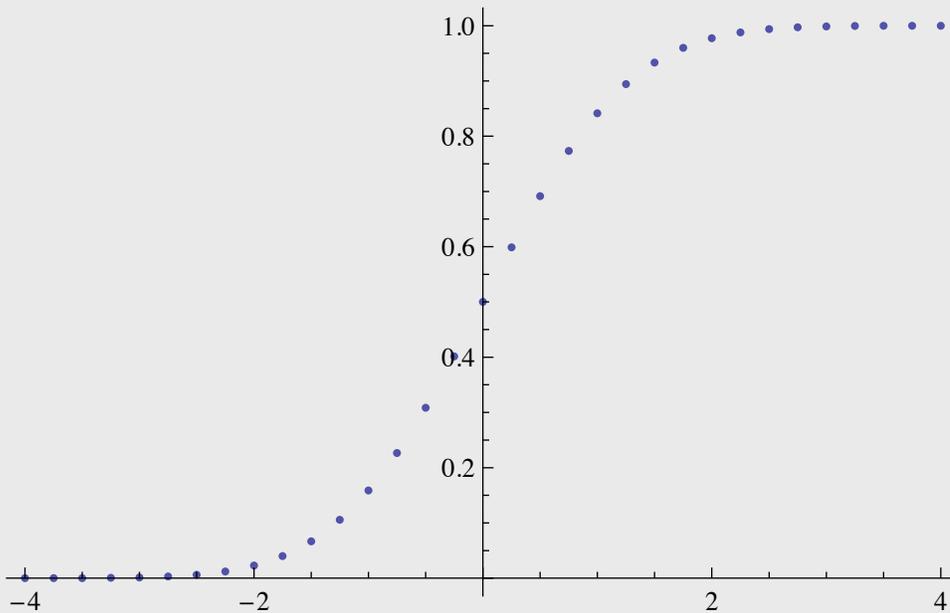
```
cumulgauss[Infinity]
```

```
1
```

```
Plot[cumulgauss[x], {x, -4, 4}]
```



```
lcumulgauss = Table[{x, cumulgauss[x]}, {x, -4., 4., 0.25.}];  
ListPlot[lcumulgauss]
```



■ Make inverse cumulative gaussian table

```
invlcumulgauss = RotateLeft[lcumulgauss, {1, 1}];
```

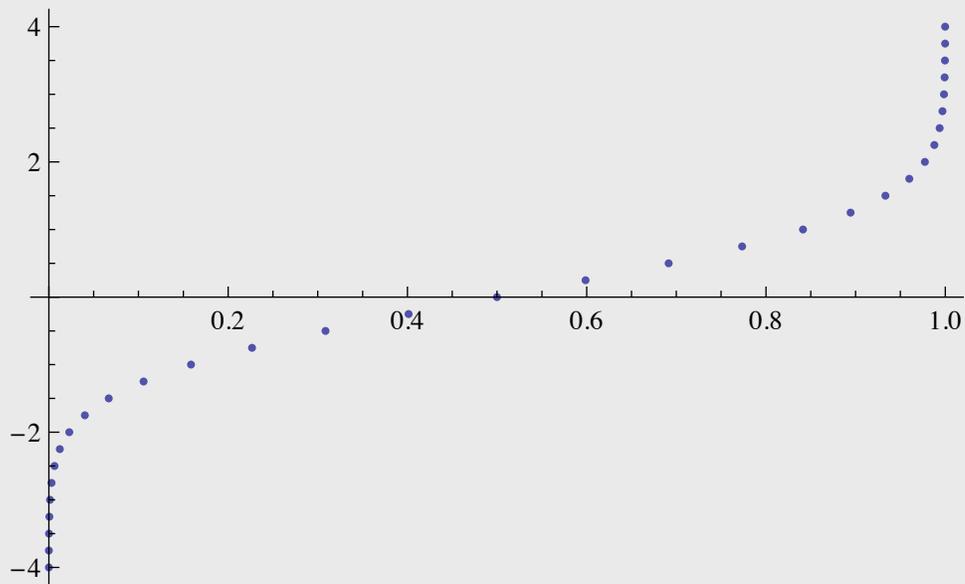
To see what this does, evaluate:

```
{{x1, y1}, {x2, y2}, {x3, y3}}
RotateLeft[{{x1, y1}, {x2, y2}, {x3, y3}}, {1, 1}]
```

$$\begin{pmatrix} x1 & y1 \\ x2 & y2 \\ x3 & y3 \end{pmatrix}$$

$$\begin{pmatrix} y2 & x2 \\ y3 & x3 \\ y1 & x1 \end{pmatrix}$$

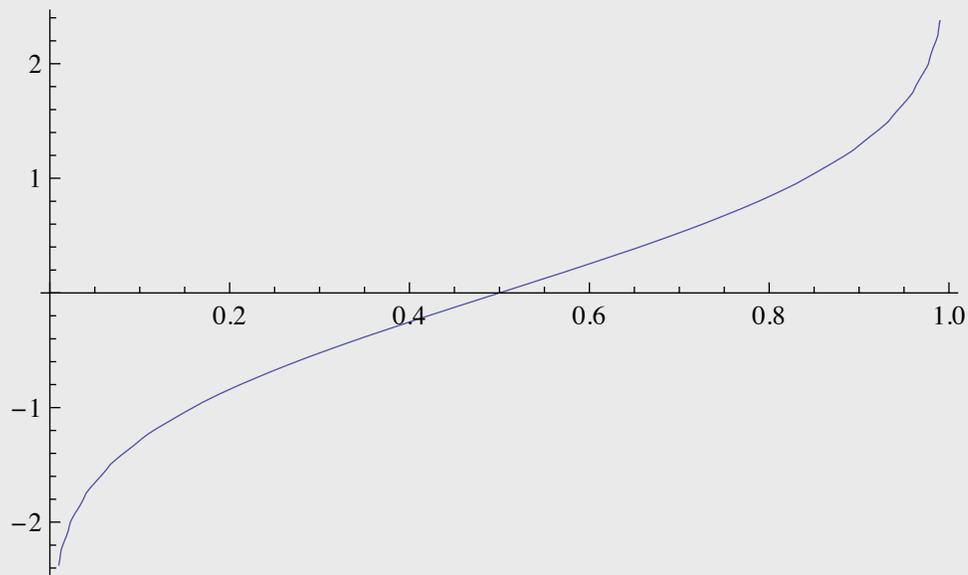
```
ListPlot[invlcumulgauss]
```



■ Make interpolated function of the inverse cumulative

```
interinvlcumulgauss = Interpolation[invlcumulgauss];
```

```
Plot[interinvcumulgauss[x], {x, 0.01, 0.99}]
```



- Draw samples with a standard deviation of $\text{Sqrt}[10]$

```
Round[10 interinvcumulgauss[RandomReal[]]]
```

```
-11
```

- Draw a bunch of samples, and plot up histogram

```
z = Table[Round[10 interinvcumulgauss[RandomReal[]]], {10 000}];
domain = Range[-20, 20];
Freq = (Count[z, #1] &) /@ domain;
```

```
InterpolatingFunction::dmval:
```

```
Input value {0.000010826} lies outside the range of data
in the interpolating function. Extrapolation will be used. >>
```

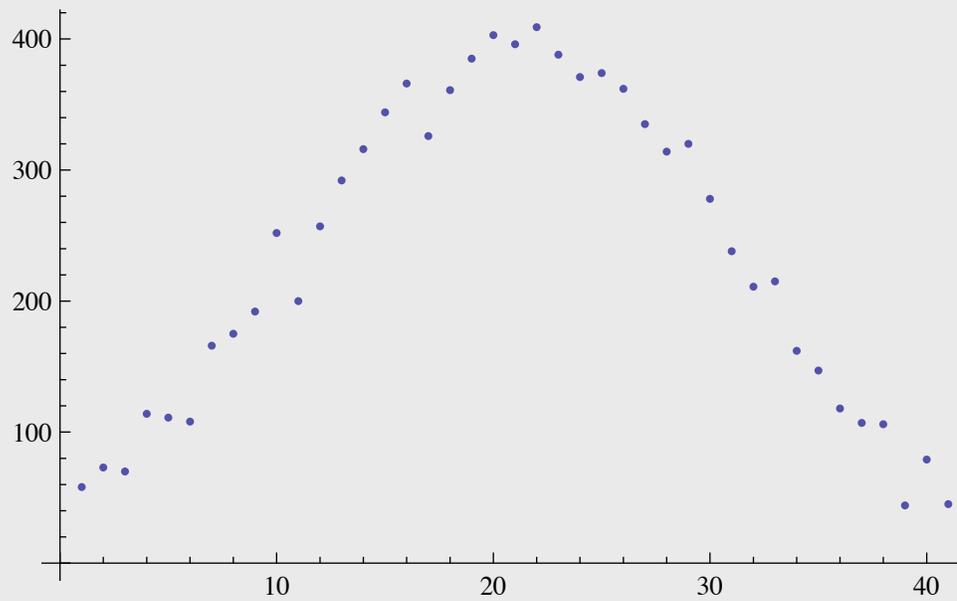
```
a
```

```
a
```

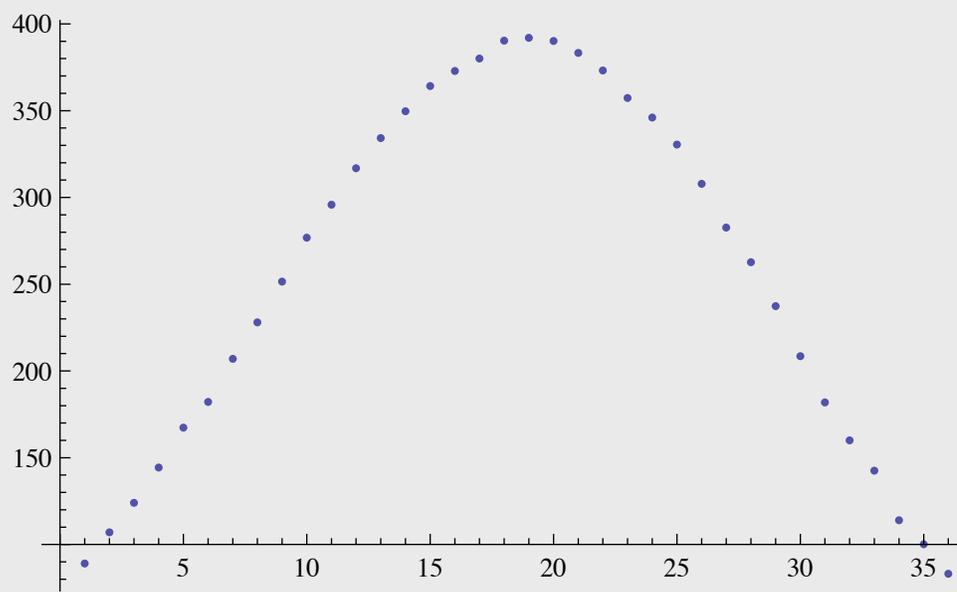
```
Sqrt[10.]
```

```
3.16228
```

```
ListPlot[Freq]
```

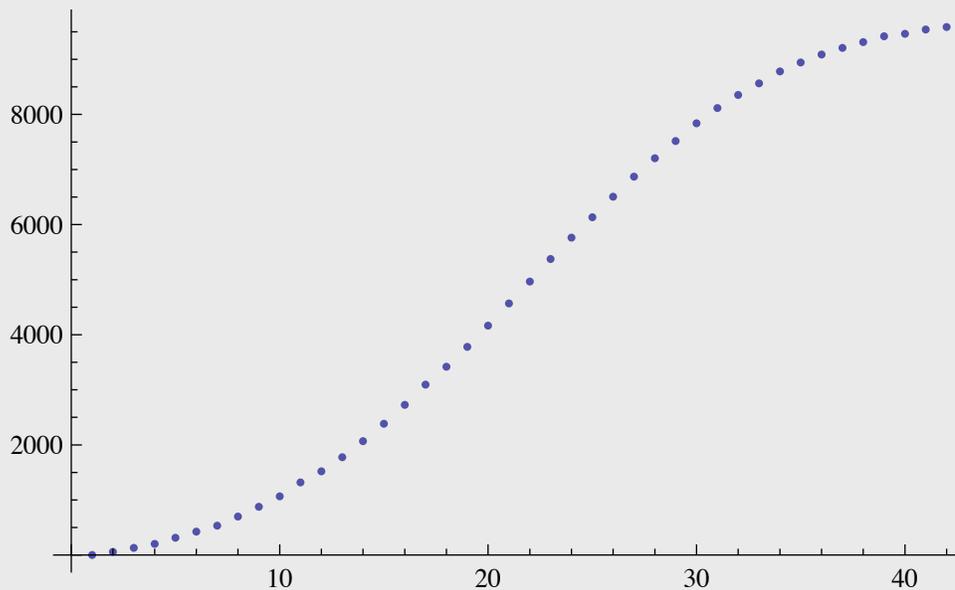


```
ListPlot[MovingAverage[Freq, 6]]
```



Plot up cumulative histogram

```
CumFreq = FoldList[Plus, 0, Freq];
ListPlot[CumFreq]
```



Multivariate (vector) gaussian distributions

■ Define multivariate gaussian probability density

An n -variate multivariate gaussian (multinormal) distribution with mean vector μ and covariance matrix Σ is denoted $N_n(\mu, \Sigma)$. The density is:

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{n/2} \text{Det}[\Sigma]^{1/2}} \text{Exp}\left[-\frac{1}{2} (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu)\right] \quad (3)$$

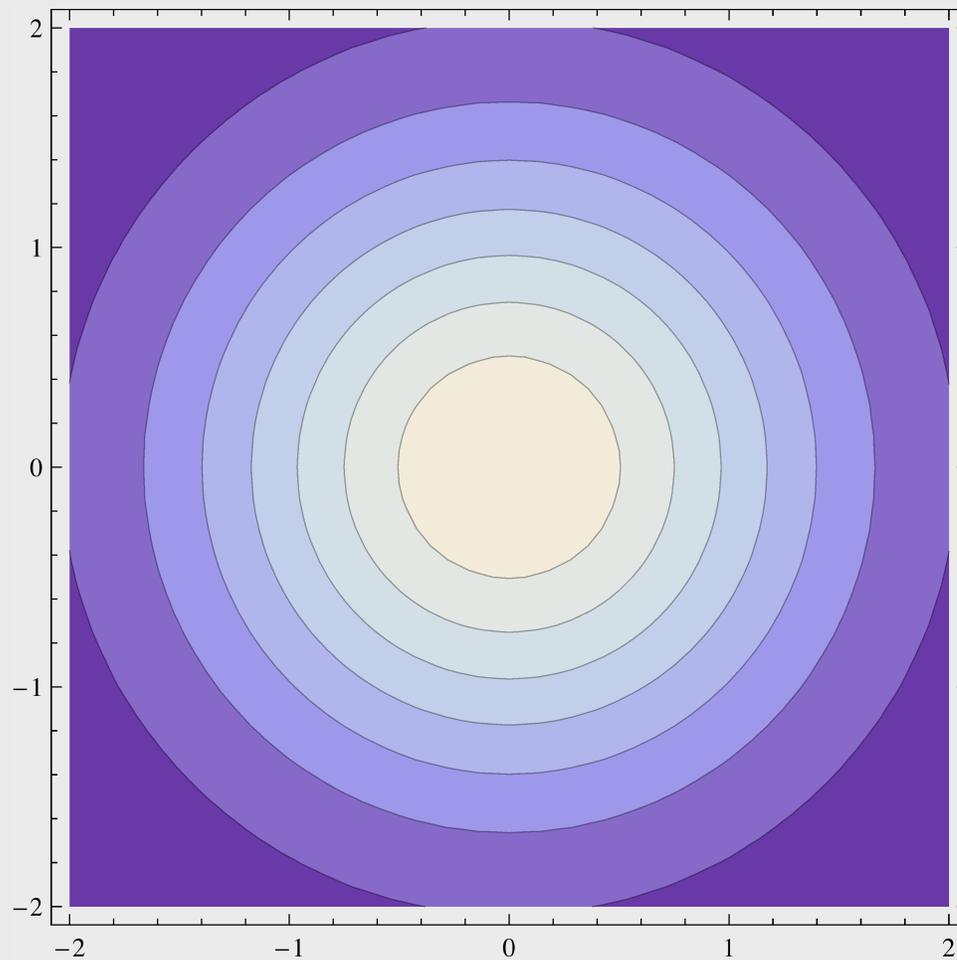
We define a 2-variate density:

```
multigaus[x_, m_, cov_] :=
  Module[{IC, detCov, norm, p},
    IC = Inverse[cov];
    detCov = Abs[Det[cov]];
    norm = N[Sqrt[(2Pi)^2 detCov]];
    p = Exp[-0.5 (x-m).IC.(x-m)]/norm;
    Return[p];
  ];
```

Two variable examples

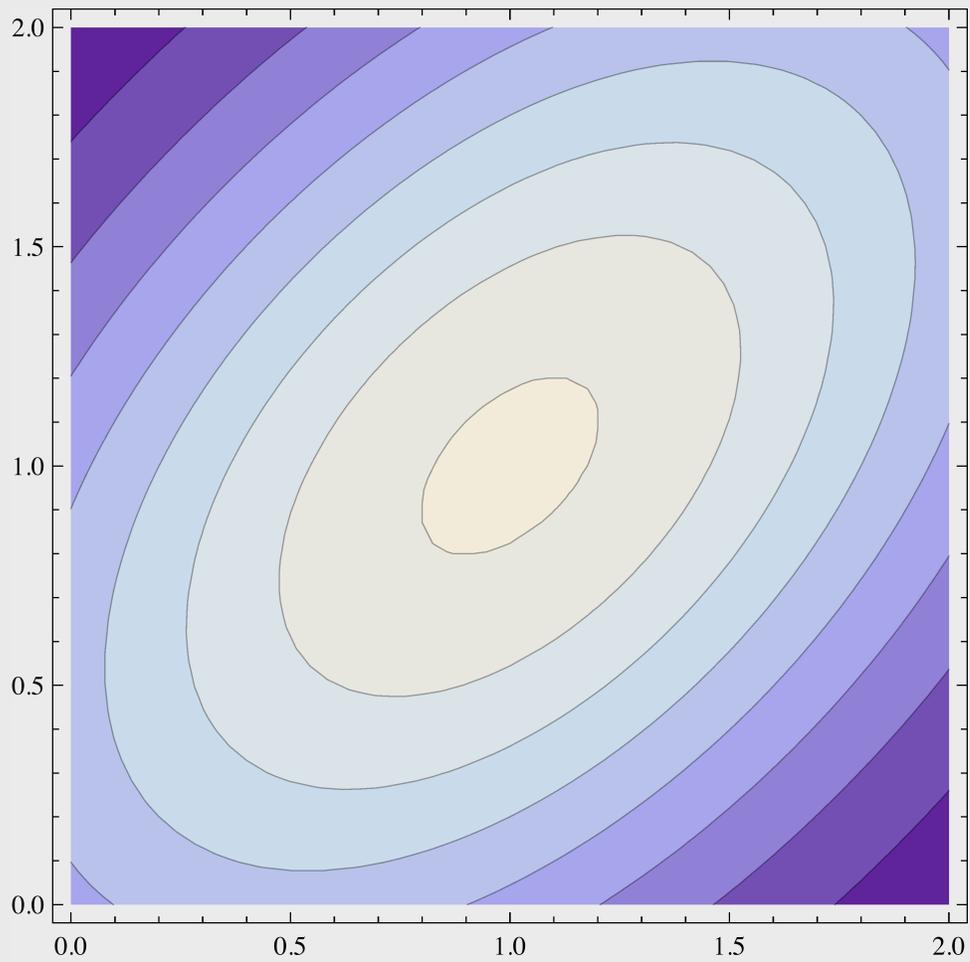
■ Zero mean, zero correlation

```
m1 = {0, 0}; Cov = {{1, 0}, {0, 1}};  
ContourPlot[multigauss[{x1, x2}, m1, Cov], {x1, -2, 2}, {x2, -2, 2}]
```



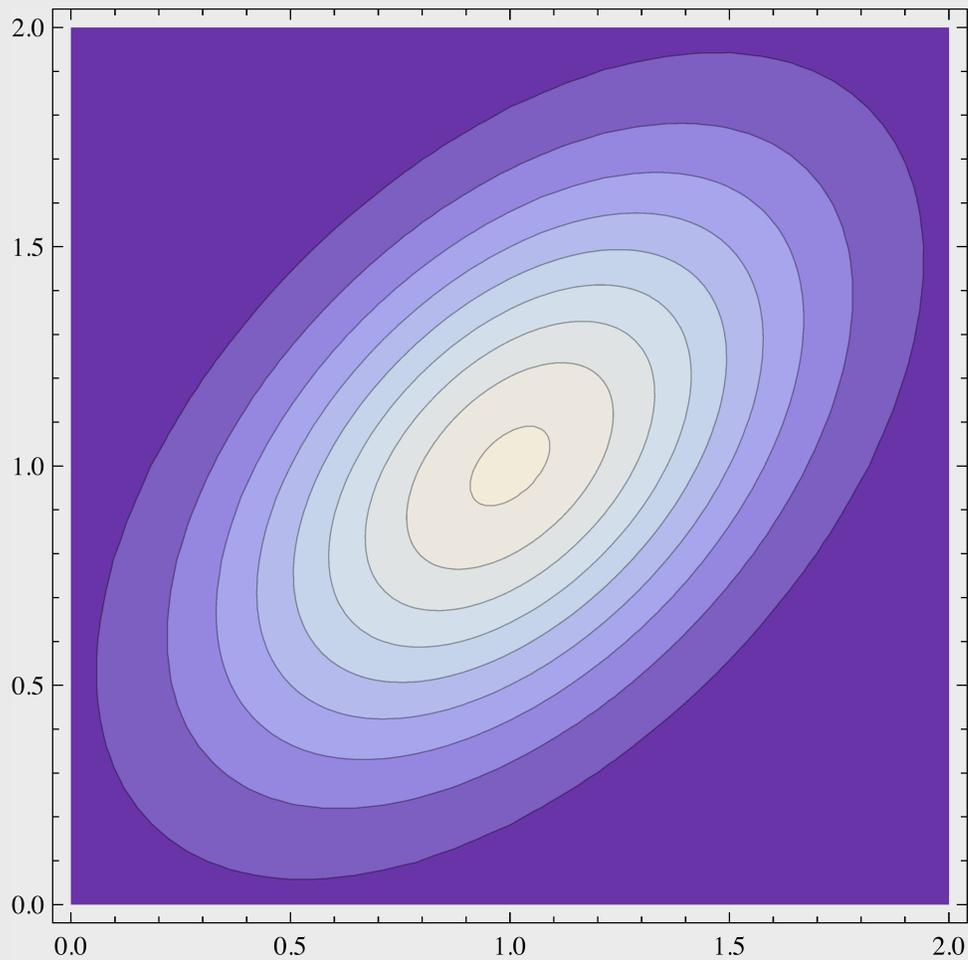
■ Mean = {1,1}, positive correlation

```
m1 = {1, 1}; Cov = {{1, 0.5`}, {0.5`, 1}};  
ContourPlot[multigauss[{x1, x2}, m1, Cov], {x1, 0, 2}, {x2, 0, 2}]
```



■ Mean = {1,1}, positive correlation, small variance

```
m1 = {1, 1}; Cov = 0.2` {{1, 0.5`}, {0.5`, 1}};
ContourPlot[multigauss[{x1, x2}, m1, Cov], {x1, 0, 2}, {x2, 0, 2}]
```



■ Mixture of gaussians

α is a mixing parameter

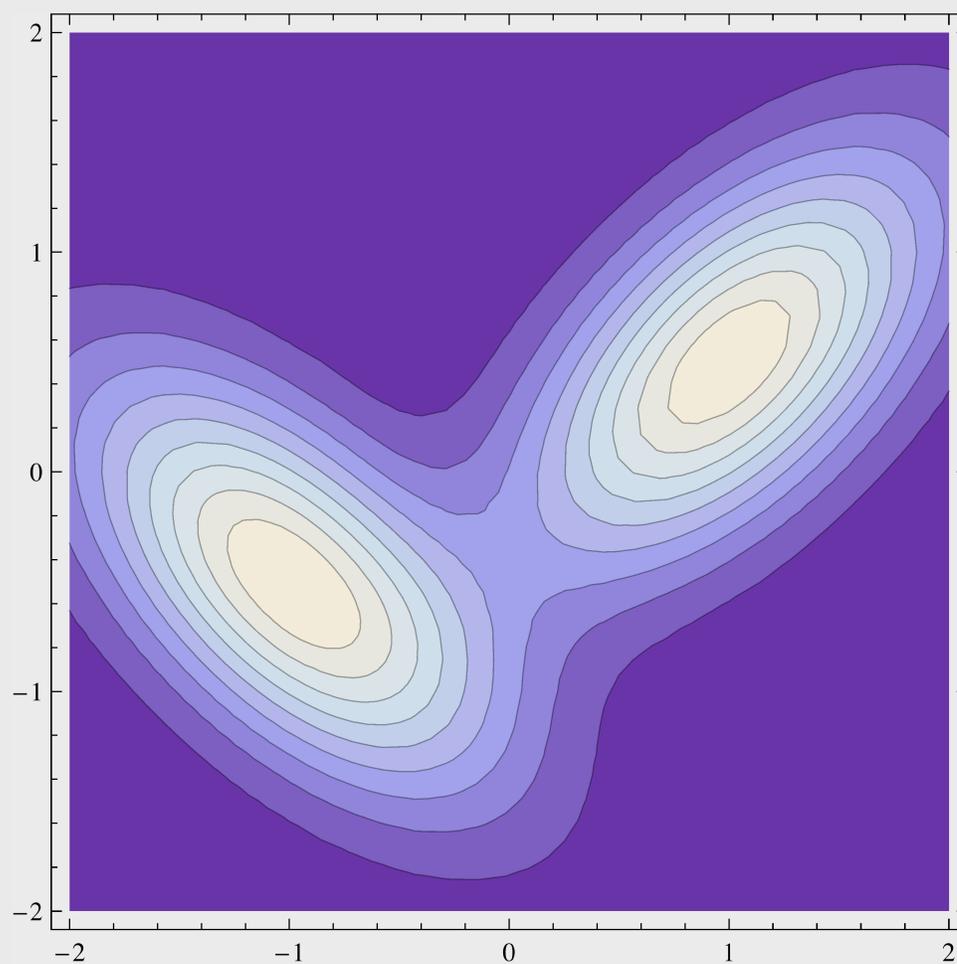
$$p(\mathbf{x}) = \alpha p_1(\mathbf{x}) + (1 - \alpha) p_2(\mathbf{x}) \quad \text{where } 0 \leq \alpha \leq 1 \quad (4)$$

α can be interpreted in terms of a prior probability of choosing which of two distributions a sample will be drawn from.

Starting with $p(x,a)$ use the sum and product rules to derive the above mixture density where $p(a=a_1) = \alpha$, and $p(a = a_2)=(1-\alpha)$.

```
m1 = {1,.5}; m2 = {-1,-.5};  
Cov1 = 0.4*{{1,.6},{.6,1}};  
Cov2 = 0.4*{{1,-.6},{-.6,1}};  
mix[x_] := 0.5 (multigauss[x,m1,Cov1] + multigauss[x,m2,Cov2]);
```

```
ContourPlot[mix[{x1, x2}], {x1, -2, 2}, {x2, -2, 2}]
```



- Drawing samples from the density--draw from a hat method
- We'll simulate the process of filling a hat with slips of paper, where the number of slips is proportional to the probability the number being in some range (dx1,dx2)

```

m1 = {0,0};
Cov = {{1,.8},{.8,1}};
dx1 = 0.1;
dx2 = dx1;

Nslips=100;
hat = {};
For[x1=-2,x1<=2,x1=x1+dx1,
  For[x2=-2,x2<=2,x2=x2+dx2,
    np = Nslips*multigauss[{x1,x2},m1,Cov];
    For[i=1,i<np,i=i+1,
      hat = Append[hat,{x1,x2}];
    ];
  ];
];
];

```

hat is a list of pairs of numbers for which the frequency of occurrence of pairs is determined by multigauss.

```
Dimensions[hat]
```

```
{8760, 2}
```

- Now let's do a check, where we compile a histogram representing the frequencies of each slip

First, define the "bins" in domain, that we'll use to check for matches:

```

domain = {};
For[x1=-2,x1<=2,x1=x1+dx1,
  For[x2=-2,x2<=2,x2=x2+dx2,
    domain = Append[domain,{x1,x2}];
  ];
];
];

```

An alternate way of specifying the domain using Outer[], and Range[]:

```
domain2 = Flatten[Outer[List,Range[-2,2,dx1],Range[-2,2,dx1]],1];
```

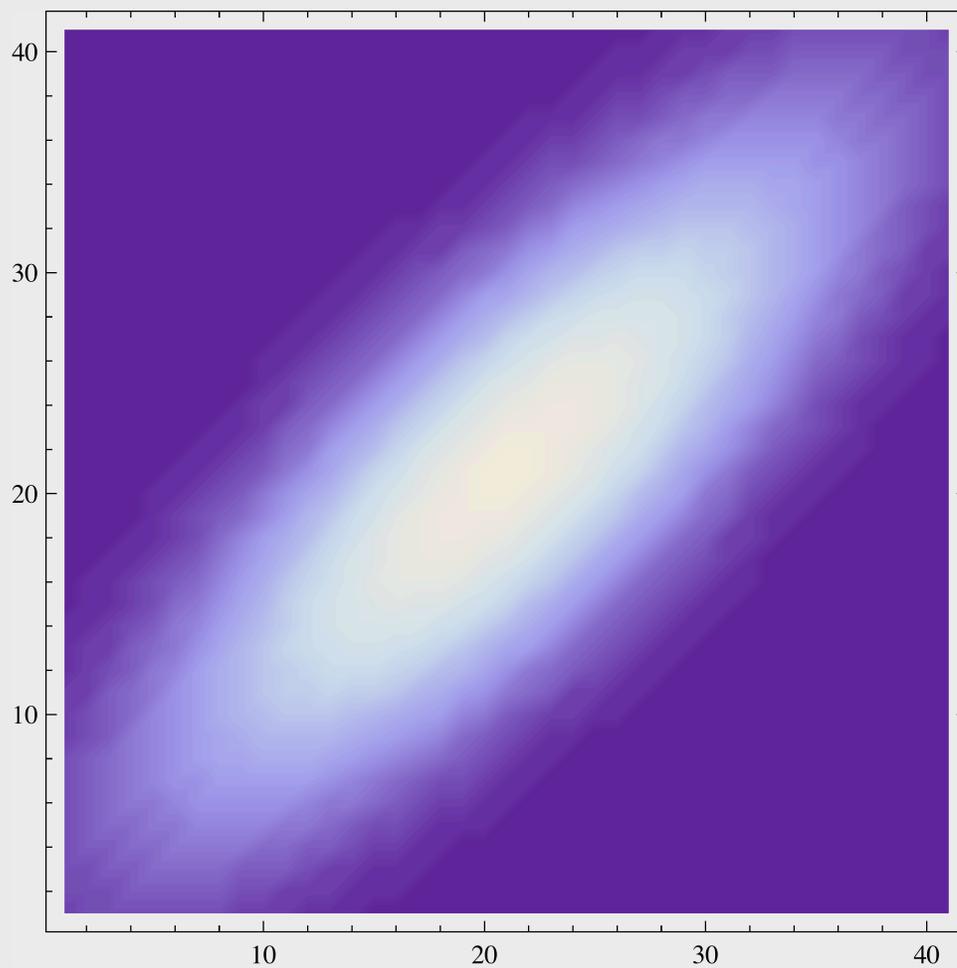
Now we'll count how many times we find that an element of hat matches a domain element:

```
Freq = Map[Count[hat,#]&,domain];
```

```
width = Sqrt[Dimensions[Freq]]
```

```
{41}
```

```
ListDensityPlot[Partition[Freq, width]]
```

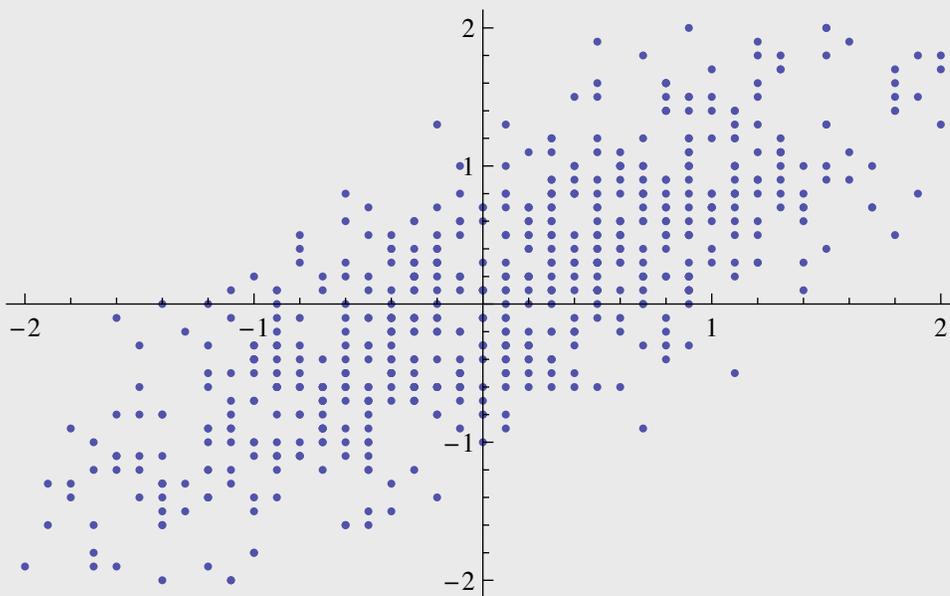


- Now draw a sample--simulate pulling a slip from hat

```
rv := hat[[RandomInteger[{1, Length[hat]}]]];
```

```
test = Table[hat[[RandomInteger[{1, Length[hat]}]]], {600}];
```

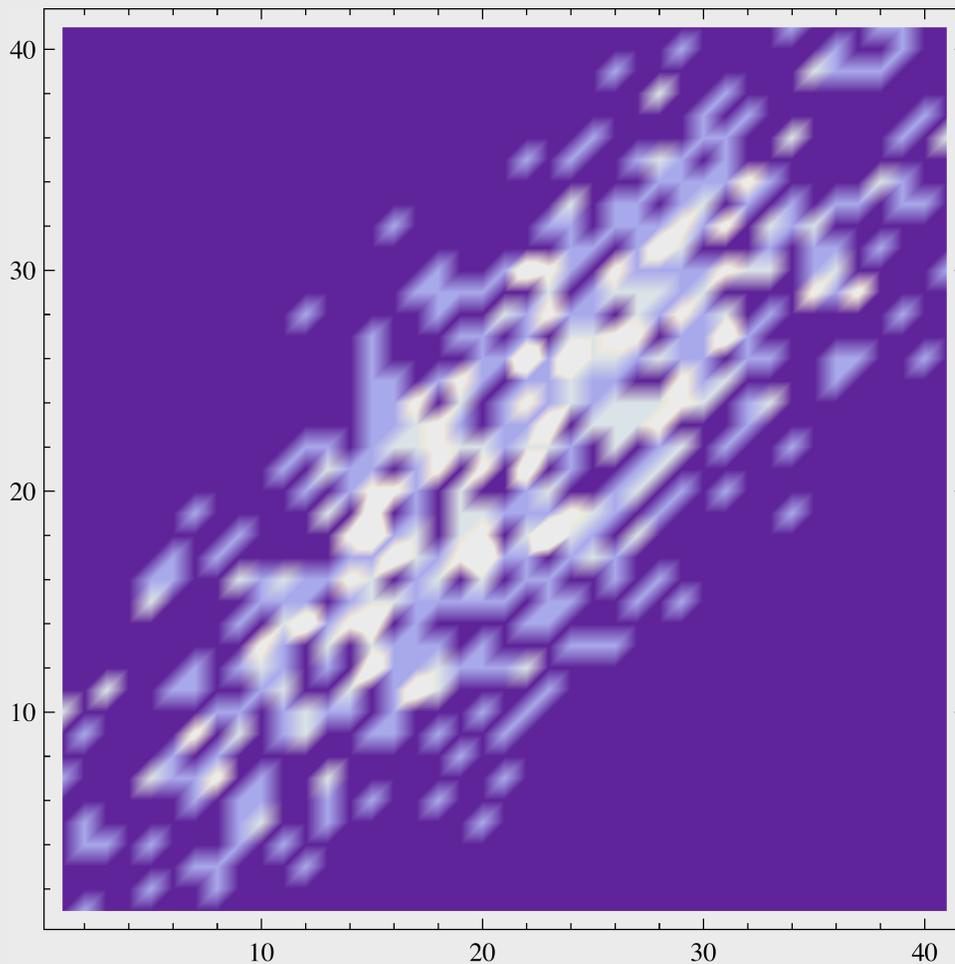
```
g1 = ListPlot[test]
```



Of course, we can't see the frequency of draws in this plot, so let's count up the number of occurrences per bin, and plot up the results as we did above.

```
Freq2 = Map[Count[test, #]&, domain];  
width = Sqrt[Dimensions[Freq2]];
```

```
ListDensityPlot[Partition[Freq2, width]]
```



Exercise: Drawing multivariate samples from the density -- use the inverse cumulative distribution

Add-on *Mathematica* functions for gaussian multivariates & exploring marginals

■ Define PDF, CDF

The add – on function,

`MultinormalDistribution[μ , σ]` specifies a multinormal (multivariate Gaussian) distribution with mean

```
Needs["MultivariateStatistics`"]
```

```
m1 = {1, .5};  
r=0.4*{{1, .6},{.6,4}};  
ndist = MultinormalDistribution[m1, r];
```

```
pdf = PDF[ndist, {x1, x2}]
```

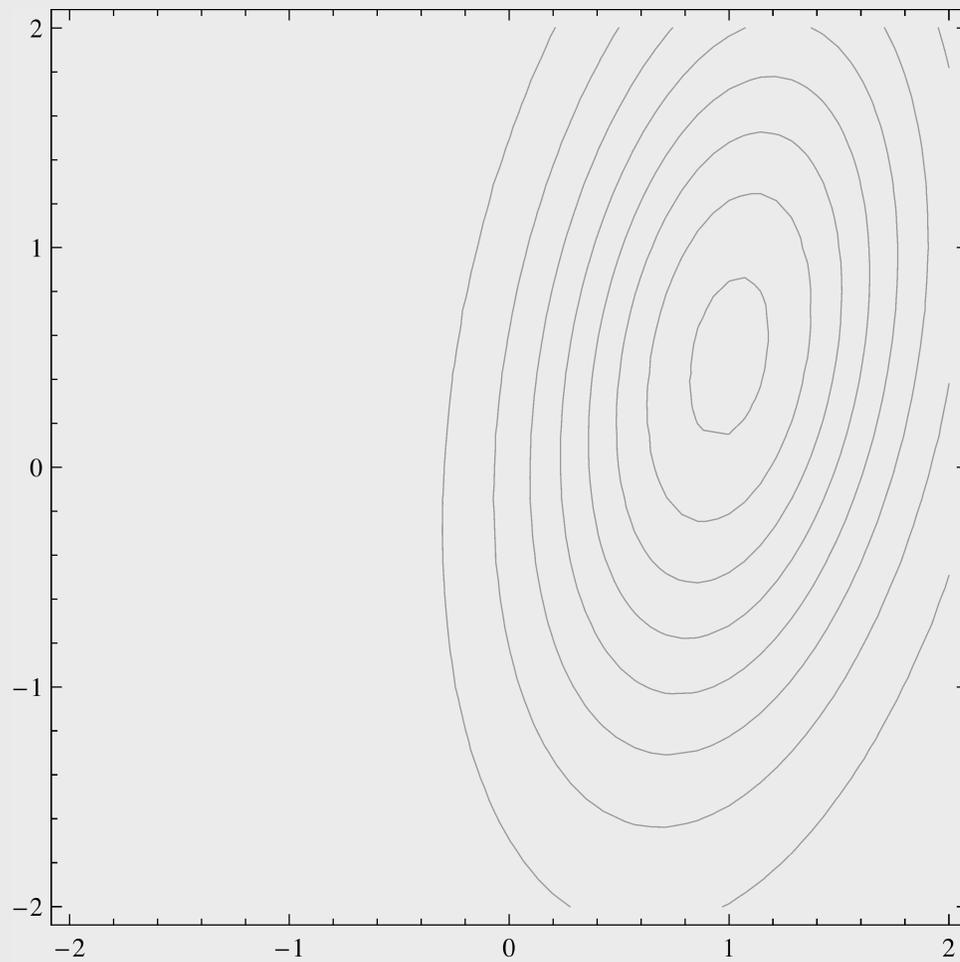
```
0.0339281
```

What is the probability of the distribution in the region $x_1 < -2 \cap x_2 < 1$.

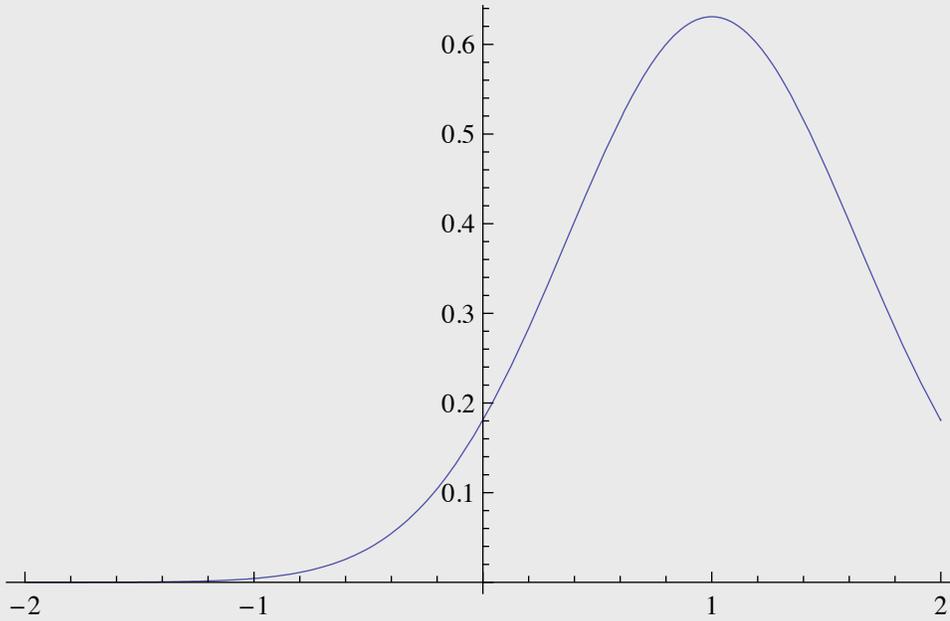
```
CDF[ndist, {-2, 1}]
```

```
 $1.02471 \times 10^{-6}$ 
```

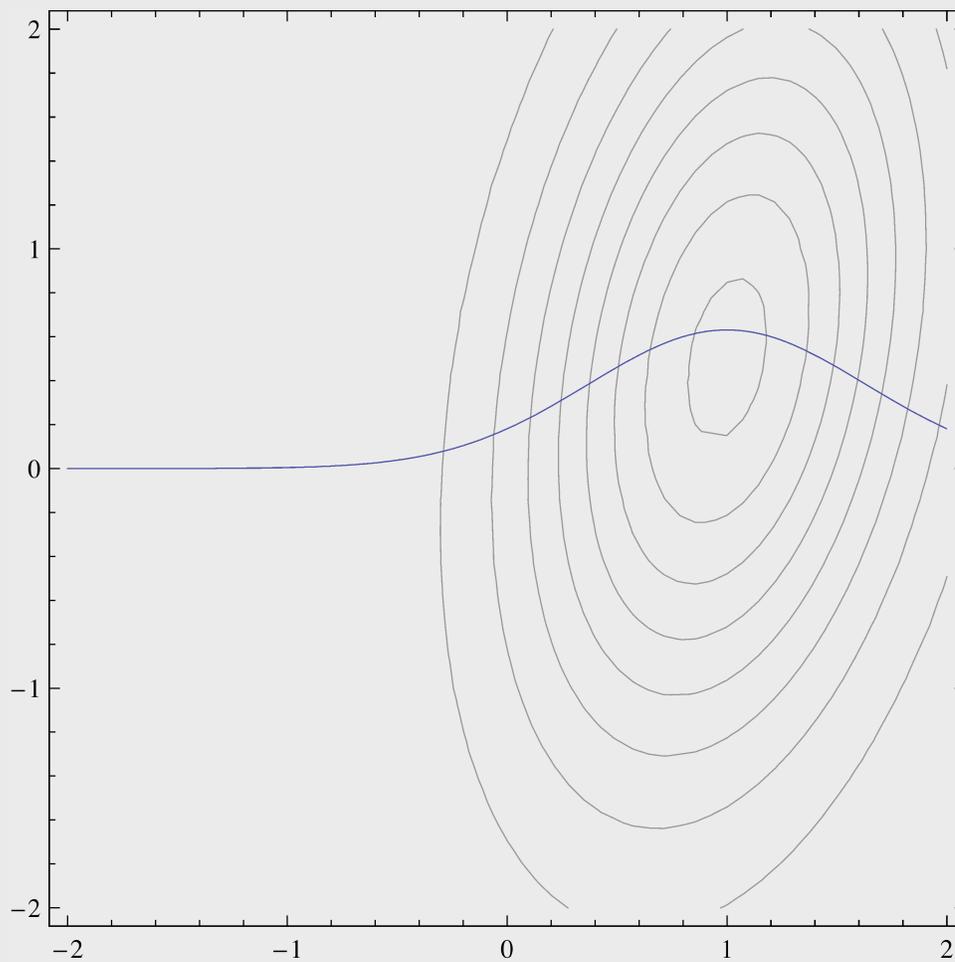
```
g1 = ContourPlot[PDF[ndist, {x1, x2}], {x1, -2, 2}, {x2, -2, 2},  
ContourShading -> False]
```



```
marginal[x11_] := Evaluate[ $\int_{-\infty}^{\infty}$  PDF[ndist, {x11, x22}] dx22];  
g2 = Plot[marginal[x11], {x11, -2, 2}]
```



```
Show[{g1, g2}]
```



■ Drawing samples

Drawing samples is done by:

```
RandomReal[ndist]
```

```
{1.16838, -0.538701}
```

■ Mixtures of gaussians

```
Clear[mix];
```

```

r1=0.4*{{1,.6},{.6,1}};
r2=0.4*{{1,-.6},{-.6,1}};
m1 = {1,.5}; m2 = {-1,-.5};
ndist1 = MultinormalDistribution[m1, r1];
ndist2 = MultinormalDistribution[m2, r2];

```

```

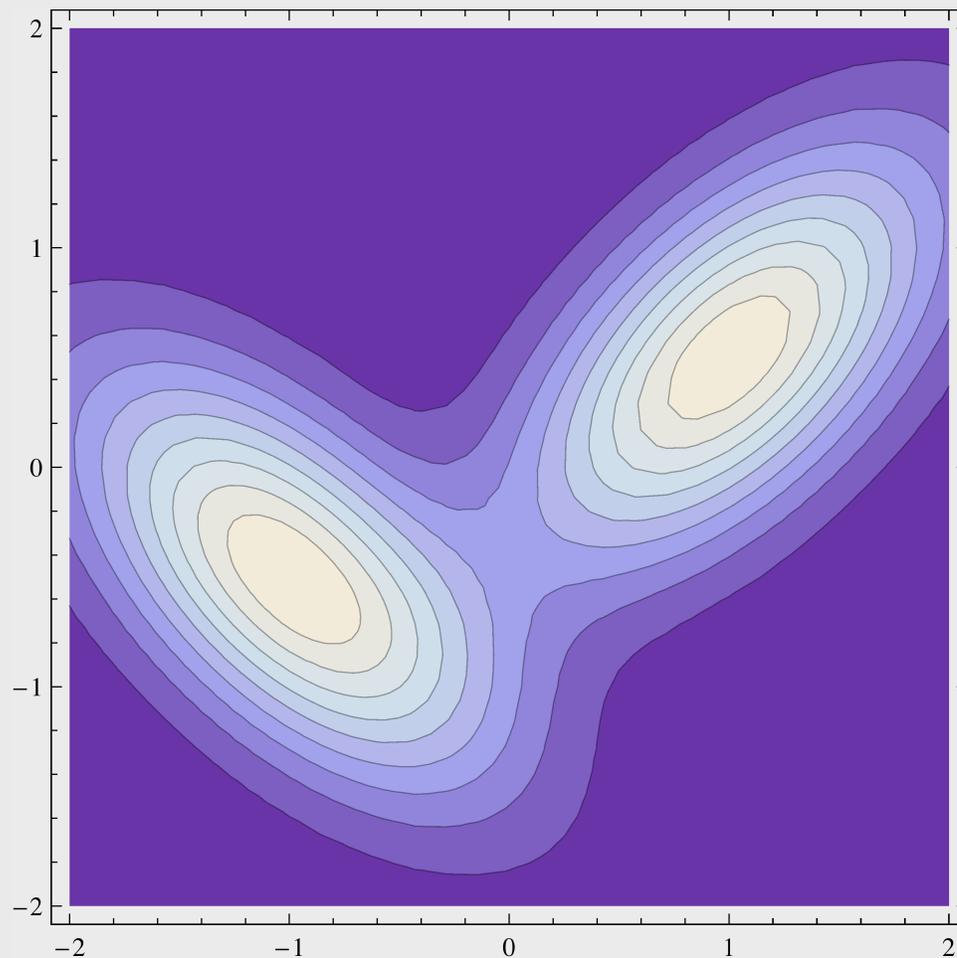
mix[x_] := 0.5 (PDF[ndist1, x] + PDF[ndist2, x]);

```

```

ContourPlot[mix[{x1, x2}], {x1, -2, 2}, {x2, -2, 2}]

```



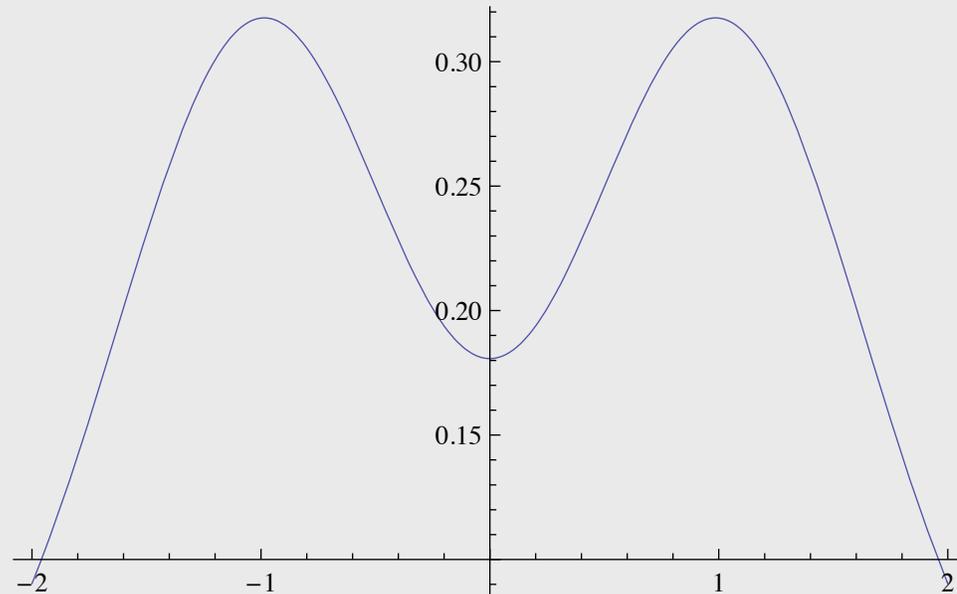
■ Marginals for mixture

Project onto the x_1 axis:

$$\text{marginal}[x1_] := \text{Integrate}[\text{mix}[\{x1, x2\}], \{x2, -\text{Infinity}, \text{Infinity}\}] \quad (5)$$

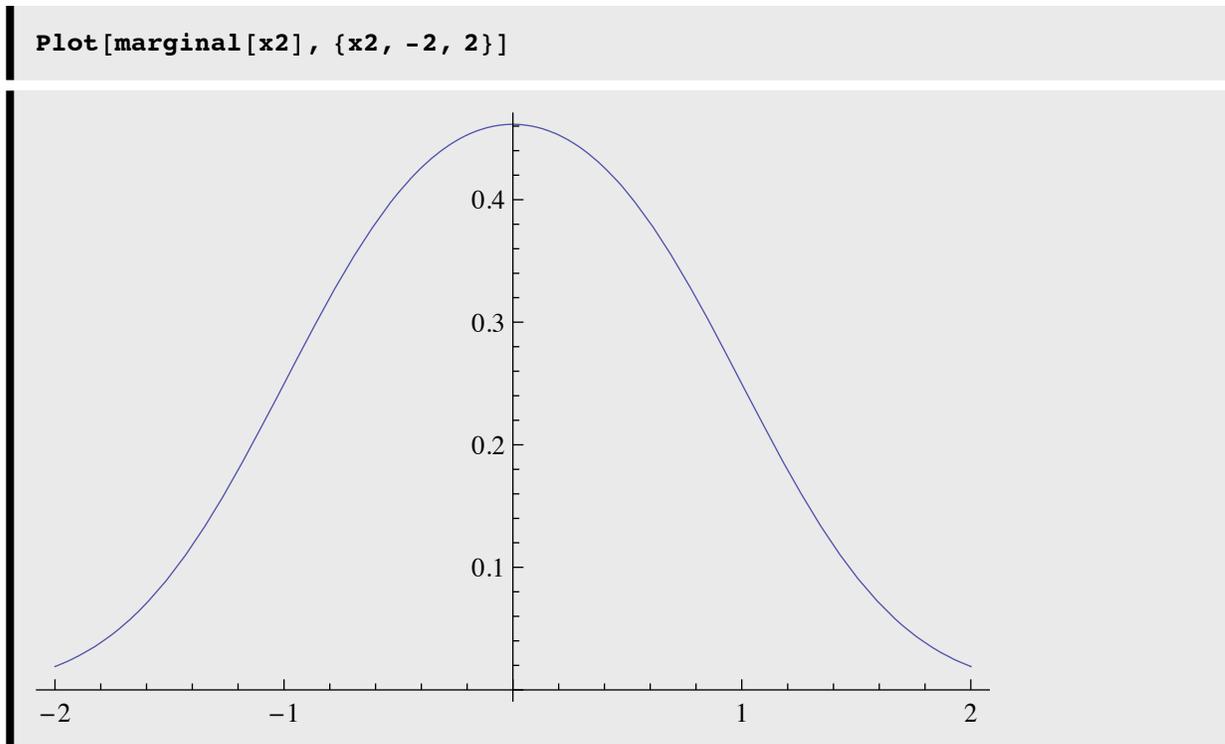
```
Clear[marginal, x1, x2];  
marginal[x1_] :=  
  Evaluate[  
    0.5 * (Integrate[PDF[ndist1, {x1, x2}], {x2, -Infinity, Infinity}] +  
    Integrate[PDF[ndist2, {x1, x2}], {x2, -Infinity, Infinity}]);
```

```
Plot[marginal[x1], {x1, -2, 2}]
```



Project onto the x_2 axis:

```
Clear[marginal, x1, x2];  
marginal[x2_] :=  
  Evaluate[  
    0.5 * (Integrate[PDF[ndist1, {x1, x2}], {x1, -Infinity, Infinity}] +  
    Integrate[PDF[ndist2, {x1, x2}], {x1, -Infinity, Infinity}]);
```



Which projection (marginal) is more "interesting"--the one onto x_1 or onto x_2 ?

Exercises

Exercise: Use density mapping theorem to make random number generator for density $p(y)$

References

- Applebaum, D. (1996). Probability and Information . Cambridge, UK: Cambridge University Press.
- Cover, T. M., & Joy, A. T. (1991). *Elements of Information Theory*. New York: John Wiley & Sons, Inc.
- Duda, R. O., & Hart, P. E. (1973). Pattern classification and scene analysis . New York.: John Wiley & Sons.
- Kersten, D. and P.W. Schrater (2000), *Pattern Inference Theory: A Probabilistic Approach to Vision*, in *Perception and the Physical World*, R. Mausfeld and D. Heyer, Editors. , John Wiley & Sons, Ltd.: Chichester. (pdf)
- Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2)
<http://gandalf.psych.umn.edu/~kersten/kersten-lab/papers/KerstenYuilleApr2003.pdf>
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge, UK: Cambridge University Press.
- Van Trees, H. L. (1968). Detection, Estimation and Modulation Theory . New York: John Wiley and Sons.
- Yuille, A., Coughlan J., Kersten D.(1998) (pdf)
- © 2000, 2001, 2003, 2004, 2008 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.