Multivariate Normal Distribution

Max Turgeon

STAT 4690–Applied Multivariate Analysis

Building the multivariate density i

- Let $Z \sim N(0, 1)$ be a standard (univariate) normal random variable. Recall that its density is given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}z^2\right).$$

- Now if we take $Z_1, \ldots, Z_p \sim N(0, 1)$ independently distributed, their joint density is

Building the multivariate density ii

$$\begin{aligned}\phi(z_1, \ldots, z_p) &= \prod_{i=1}^{p} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} z_i^2\right) \\ &= \frac{1}{(\sqrt{2\pi})^p} \exp\left(-\frac{1}{2} \sum_{i=1}^{p} z_i^2\right) \\ &= \frac{1}{(\sqrt{2\pi})^p} \exp\left(-\frac{1}{2} \mathbf{z}^T \mathbf{z}\right),\end{aligned}$$

where $\mathbf{z} = (z_1, \ldots, z_p)$.

- More generally, let $\mu \in \mathbb{R}^p$ and let Σ be a $p \times p$ positive definite matrix.

Building the multivariate density iii

- Let $\Sigma = LL^T$ be the Cholesky decomposition for Σ.

- Let $\mathbf{Z} = (Z_1, \ldots, Z_p)$ be a standard (multivariate) normal random vector, and define $\mathbf{Y} = L\mathbf{Z} + \mu$. We know from last lecture that
  - $E(\mathbf{Y}) = LE(\mathbf{Z}) + \mu = \mu$;
  - $\mathrm{Cov}(\mathbf{Y}) = L\mathrm{Cov}(\mathbf{Z})L^T = \Sigma$.
- To get the density, we need to compute the inverse transformation:

$$\mathbf{Z} = L^{-1}(\mathbf{Y} - \mu).$$

Building the multivariate density iv

- The Jacobian matrix J for this transformation is simply L^{-1}, and therefore

$$
\begin{aligned}
|\det(J)| &= |\det(L^{-1})| \\
&= \det(L)^{-1} \qquad (L \text{ is p.d.}) \\
&= \sqrt{\det(\Sigma)}^{-1} \\
&= \det(\Sigma)^{-1/2}.
\end{aligned}
$$

Building the multivariate density v

- Plugging this into the formula for the density of a transformation, we get

$$
\begin{aligned}
f(y_1, \ldots, y_p) &= \frac{1}{\det(\Sigma)^{1/2}} \phi(L^{-1}(\mathbf{y} - \mu)) \\
&= \frac{1}{\det(\Sigma)^{1/2}} \left(\frac{1}{(\sqrt{2\pi})^p} \exp\left(-\frac{1}{2} (L^{-1}(\mathbf{y} - \mu))^T L^{-1}(\mathbf{y} - \mu) \right) \right) \\
&= \frac{1}{\det(\Sigma)^{1/2} (\sqrt{2\pi})^p} \exp\left(-\frac{1}{2} (\mathbf{y} - \mu)^T (LL^T)^{-1} (\mathbf{y} - \mu) \right) \\
&= \frac{1}{\sqrt{(2\pi)^p |\Sigma|}} \exp\left(-\frac{1}{2} (\mathbf{y} - \mu)^T \Sigma^{-1} (\mathbf{y} - \mu) \right).
\end{aligned}
$$

Example i

```r
set.seed(123)

n <- 1000; p <- 2
Z <- matrix(rnorm(n*p), ncol = p)

mu <- c(1, 2)
Sigma <- matrix(c(1, 0.5, 0.5, 1), ncol = 2)
L <- t(chol(Sigma))
```

Example ii

```r
Y <- L %*% t(Z) + mu
Y <- t(Y)

colMeans(Y)
```

```
## [1] 1.016128 2.044840
```

```r
cov(Y)
```

```
##           [,1]      [,2]
## [1,] 0.9834589 0.5667194
## [2,] 0.5667194 1.0854361
```

Example iii

```r
library(tidyverse)
Y %>%
  data.frame() %>%
  ggplot(aes(X1, X2)) +
  geom_density_2d()
```

Example iv

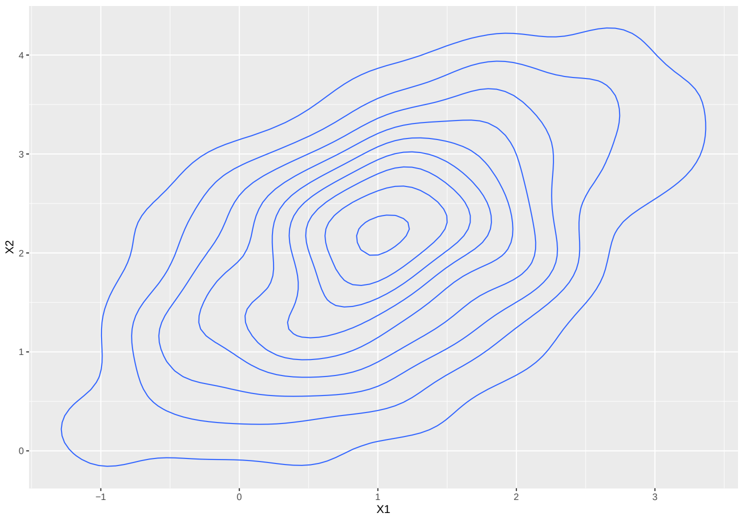

Example v

```r
library(mvtnorm)

Y <- rmvnorm(n, mean = mu, sigma = Sigma)

colMeans(Y)
```

```
## [1] 0.9812102 1.9829380
```

```r
cov(Y)
```

Example vi

```
##               [,1]      [,2]
## [1,] 0.9982835 0.4906990
## [2,] 0.4906990 0.9489171
```

```
Y %>%
  data.frame() %>%
  ggplot(aes(X1, X2)) +
  geom_density_2d()
```

Example vii

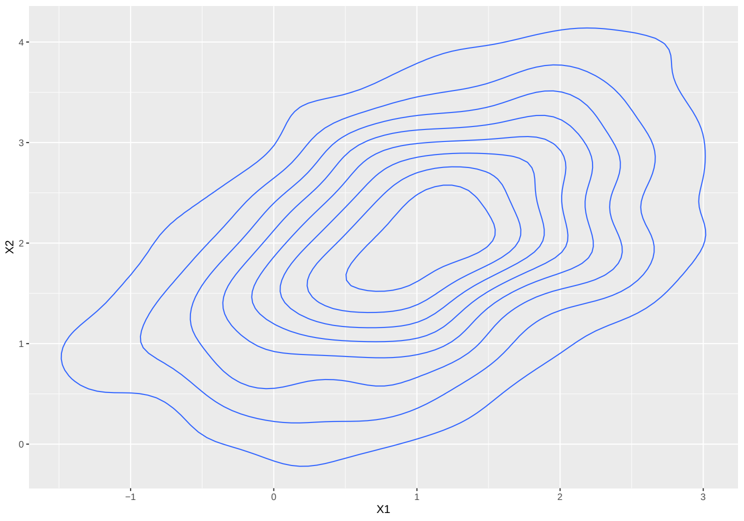

Other characterizations

There are at least two other ways to define the multivariate random distribution:

1. A p-dimensional random vector $\mathbf{Y}$ is said to have a multivariate normal distribution if and only if every linear combination of $\mathbf{Y}$ has a *univariate* normal distribution.
2. A p-dimensional random vector $\mathbf{Y}$ is said to have a multivariate normal distribution if and only if its distribution maximises entropy over the class of random vectors with fixed mean μ and fixed covariance matrix Σ and support over $\mathbb{R}^p$.

Useful properties i

- If $\mathbf{Y} \sim N_p(\mu, \Sigma)$, A is a $q \times p$ matrix, and $b \in \mathbb{R}^q$, then

$$A\mathbf{Y} + b \sim N_q(A\mu + b, A\Sigma A^T).$$

- If $\mathbf{Y} \sim N_p(\mu, \Sigma)$ then all subsets of $\mathbf{Y}$ are normally distributed; that is, write
  - $\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{pmatrix}$, $\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$;
  - $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$.
  - Then $\mathbf{Y}_1 \sim N_r(\mu_1, \Sigma_{11})$ and $\mathbf{Y}_2 \sim N_{p-r}(\mu_2, \Sigma_{22})$.

Useful properties ii

- Assume the same partition as above. Then the following are equivalent:
    - $\mathbf{Y}_1$ and $\mathbf{Y}_2$ are independent;
    - $\Sigma_{12} = 0$;
    - $\mathrm{Cov}(\mathbf{Y}_1, \mathbf{Y}_2) = 0$.

Exercise (J&W 4.3)

Let $(Y_1, Y_2, Y_3) \sim N_3(\mu, \Sigma)$ with $\mu = (3, 1, 4)$ and

$$\Sigma = \begin{pmatrix} 1 & -2 & 0 \\ -2 & 5 & 0 \\ 0 & 0 & 2 \end{pmatrix}.$$

Which of the following random variables are independent? Explain.

1. Y_1 and Y_2.
2. Y_2 and Y_3.
3. (Y_1, Y_2) and Y_3.
4. $0.5(Y_1 + Y_2)$ and Y_3.
5. Y_2 and $Y_2 - \frac{5}{2}Y_1 - Y_3$.

Conditional Normal Distributions i

- **Theorem**: Let $\mathbf{Y} \sim N_p(\mu, \Sigma)$, where
  - $\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{pmatrix}$, $\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$;
  - $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$.
- Then the *conditional distribution* of $\mathbf{Y}_1$ given $\mathbf{Y}_2 = y_2$ is multivariate normal $N_r(\mu_{1|2}, \Sigma_{1|2})$, where
  - $\mu_{1|2} = \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(y_2 - \mu_2)$
  - $\Sigma_{1|2} = \Sigma_{11} + \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$.

Conditional Normal Distributions ii

- **Corrolary**: Let $\mathbf{Y}_2 \sim N_{p-r}(\mu_2, \Sigma_{22})$ and assume that $\mathbf{Y}_1$ given $\mathbf{Y}_2 = y_2$ is multivariate normal $N_r(Ay_2 + b, \Omega)$, where Ω does not depend on y_2. Then $\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{pmatrix} \sim N_p(\mu, \Sigma)$, where
    - $\mu = \begin{pmatrix} A\mu_2 + b \\ \mu_2 \end{pmatrix}$;
    - $\Sigma = \begin{pmatrix} \Omega + A\Sigma_{22}A^T & A\Sigma_{22} \\ \Sigma_{22}A^T & \Sigma_{22} \end{pmatrix}$.

Exercise

- Let $\mathbf{Y}_2 \sim N_1(0, 1)$ and assume

$$\mathbf{Y}_1 \mid \mathbf{Y}_2 = y_2 \sim N_2\left(\begin{pmatrix} y_2 + 1 \\ 2y_2 \end{pmatrix}, I_2\right).$$

Find the joint distribution of $(\mathbf{Y}_1, \mathbf{Y}_2)$.

Another important result i

- Let $\mathbf{Y} \sim N_p(\mu, \Sigma)$, and let $\Sigma = LL^T$ be the Cholesky decomposition of Σ.
- We know that $\mathbf{Z} = L^{-1}(\mathbf{Y} - \mu)$ is normally distributed, with mean 0 and covariance matrix

$$\mathrm{Cov}(\mathbf{Z}) = L^{-1}\Sigma(L^{-1})^T = I_p.$$

- Therefore $(\mathbf{Y} - \mu)^T\Sigma^{-1}(\mathbf{Y} - \mu)$ is the sum of *squared* standard normal random variables.
  - In other words, $(\mathbf{Y} - \mu)^T\Sigma^{-1}(\mathbf{Y} - \mu) \sim \chi^2(p)$.
  - This can be seen as a generalization of the univariate result $\left(\frac{X-\mu}{\sigma}\right)^2 \sim \chi^2(1)$.

Another important result ii

- From this, we get a result about the probability that a multivariate normal falls within an *ellipse*:

$$P\left((\mathbf{Y}-\mu)^T\Sigma^{-1}(\mathbf{Y}-\mu) \leq \chi^2(\alpha;p)\right) = 1-\alpha.$$

    - We can use this to construct a confidence region around the sample mean.