STAT 2593

Lecture 018 - The Normal Distribution

Dylan Spicker

The Normal Distribution

Learning Objectives

1. Understand the normal distribution, its use cases, and its structure.
2. Understand the standard normal distribution, and how to transform to/from it.
3. Understand and apply the empirical rule.
4. Understand the normal approximations to other probability distributions.


Why cant you just be normal?
Histogram of Percent Fat
Normal
Frequency
Percent Fat
Screams

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical "bell curve" shape.

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical “bell curve” shape.
  - It characterizes *many* natural phenomena (e.g., heights, weights, measurement errors, reaction times, etc.)

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical "bell curve" shape.
  - It characterizes *many* natural phenomena (e.g., heights, weights, measurement errors, reaction times, etc.)
- It is a symmetric, bell-shaped curve, characterized by two parameters

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical “bell curve” shape.
  - It characterizes *many* natural phenomena (e.g., heights, weights, measurement errors, reaction times, etc.)
- It is a symmetric, bell-shaped curve, characterized by two parameters
  - The mean μ.

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical "bell curve" shape.
  - It characterizes *many* natural phenomena (e.g., heights, weights, measurement errors, reaction times, etc.)
- It is a symmetric, bell-shaped curve, characterized by two parameters
  - The mean μ.
  - The variance σ^2

The Normal Distribution

- The **normal distribution** is the most important distribution in all of statistics and probability.
- It is characterized the the typical "bell curve" shape.
  - It characterizes *many* natural phenomena (e.g., heights, weights, measurement errors, reaction times, etc.)
- It is a symmetric, bell-shaped curve, characterized by two parameters
  - The mean μ.
  - The variance σ^2
  - Sometimes, researchers will use the standard deviation, σ, in place of the variance. This is equivalent.

The Normal Distribution, Mathematically

- The normal density function is given by

$$f(x;\mu,\sigma^2) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

The Normal Distribution, Mathematically

- The normal density function is given by

$$f(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

- This is defined on $-\infty < x < \infty$.

The Normal Distribution, Mathematically

- The normal density function is given by

$$f(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

- This is defined on $-\infty < x < \infty$.
  - It can be shown that $E[X] = \mu$ and that $\text{var}(X) = \sigma^2$.

The Normal Distribution, Mathematically

- The normal density function is given by

$$f(x;\mu,\sigma^2) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

- This is defined on $-\infty < x < \infty$.
  - It can be shown that $E[X] = \mu$ and that $\text{var}(X) = \sigma^2$.
  - Moreover, this is a valid density.

The Normal Distribution, Mathematically

- The normal density function is given by

$$f(x;\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

- This is defined on $-\infty < x < \infty$.
  - It can be shown that $E[X] = \mu$ and that $\text{var}(X) = \sigma^2$.
  - Moreover, this is a valid density.
- There is no closed form expression for the CDF.

The Normal Distribution, Graphically

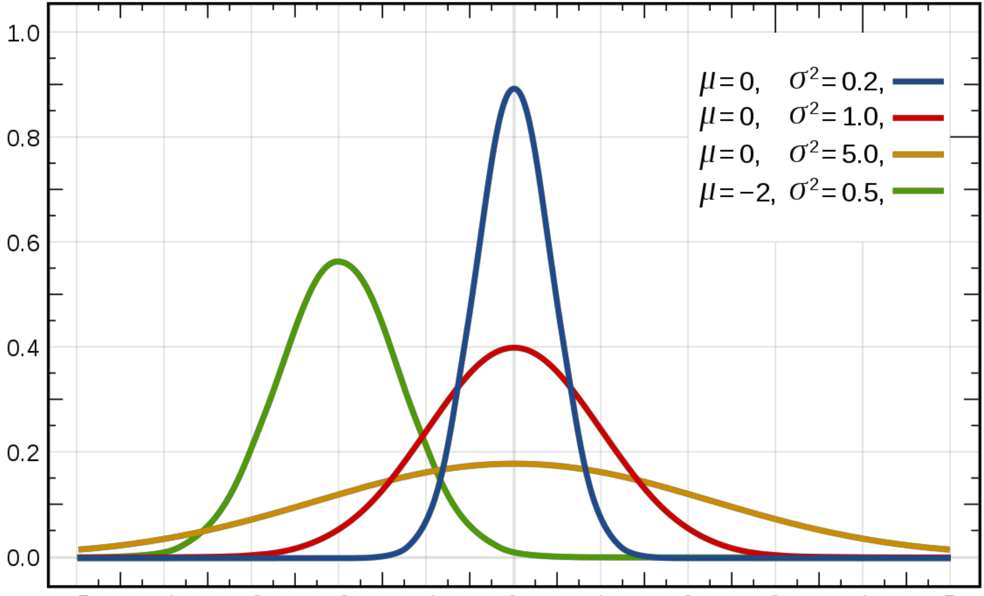

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.
- A common strategy for dealing with the normal distribution is to transform results to follow the **standard normal distribution**.

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.
- A common strategy for dealing with the normal distribution is to transform results to follow the **standard normal distribution**.
  - The standard normal has $\mu = 0$ and $\sigma^2 = 1$.

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.
- A common strategy for dealing with the normal distribution is to transform results to follow the **standard normal distribution**.
  - The standard normal has $\mu = 0$ and $\sigma^2 = 1$.
  - If X is distributed as $N(\mu, \sigma^2)$ then

$$Z = \frac{X - \mu}{\sigma} \sim N(0, 1).$$

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.
- A common strategy for dealing with the normal distribution is to transform results to follow the **standard normal distribution**.
  - The standard normal has $\mu = 0$ and $\sigma^2 = 1$.
  - If X is distributed as $N(\mu, \sigma^2)$ then

$$Z = \frac{X - \mu}{\sigma} \sim N(0, 1).$$

  - We denote the PDF of the standard normal as $\varphi(z)$ and the CDF of the standard normal as $\Phi(z) = P(Z \leq z)$.

Computing with the Normal Distribution

- The form of the PDF makes probability calculations challenging as they generally cannot be solved analytically.
  - Instead, we typically use software or tables.
- A common strategy for dealing with the normal distribution is to transform results to follow the **standard normal distribution**.
  - The standard normal has $\mu = 0$ and $\sigma^2 = 1$.
  - If X is distributed as $N(\mu, \sigma^2)$ then

$$Z = \frac{X - \mu}{\sigma} \sim N(0, 1).$$

  - We denote the PDF of the standard normal as $\varphi(z)$ and the CDF of the standard normal as $\Phi(z) = P(Z \leq z)$.
- This transformation can be undone to revert back to $N(\mu, \sigma^2)$.

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).
- Standard normal tables (values for $\Phi(z)$) are commonly published in statistics textbooks.

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).
- Standard normal tables (values for $\Phi(z)$) are commonly published in statistics textbooks.
  - While you will see these tables, you will never be expected to use them in this course.

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).
- Standard normal tables (values for $\Phi(z)$) are commonly published in statistics textbooks.
  - While you will see these tables, you will never be expected to use them in this course.
  - In the modern days, using software is more accurate and easier to do.

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).
- Standard normal tables (values for $\Phi(z)$) are commonly published in statistics textbooks.
    - While you will see these tables, you will never be expected to use them in this course.
    - In the modern days, using software is more accurate and easier to do.
- Percentiles of the normal distribution can be found as with any continuous distribution.

Properties of the Standard Normal

- Nearly all of the density falls in the interval ± 3 (> 0.99).
- Standard normal tables (values for $\Phi(z)$) are commonly published in statistics textbooks.
  - While you will see these tables, you will never be expected to use them in this course.
  - In the modern days, using software is more accurate and easier to do.
- Percentiles of the normal distribution can be found as with any continuous distribution.
  - If Z_p is the p-th percentile of Z, then $X_p = \mu + Z_p \sigma$.

Critical Values for the Standard Normal

- We most often care about *tail values* of the normal distribution.

Critical Values for the Standard Normal

- We most often care about *tail values* of the normal distribution.
- If we take z_α to be the value such that $P(Z \leq z_\alpha) = \alpha$, this is typically called the α level **critical value**.

Critical Values for the Standard Normal

- We most often care about *tail values* of the normal distribution.
- If we take z_α to be the value such that $P(Z \leq z_\alpha) = \alpha$, this is typically called the α level **critical value**.
  - We will often be concerned with $z_{0.05}$ / $z_{0.95}$, $Z_{0.025}$ and $Z_{0.975}$.

Critical Values for the Standard Normal

- We most often care about *tail values* of the normal distribution.
- If we take z_α to be the value such that $P(Z \leq z_\alpha) = \alpha$, this is typically called the α level **critical value**.
  - We will often be concerned with $z_{0.05}$ / $z_{0.95}$, $Z_{0.025}$ and $Z_{0.975}$.
  - Note that $Z_\alpha = -Z_{1-\alpha}$ due to symmetry of the standard normal.

The Empirical Rule

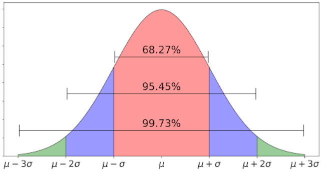

- The empirical rule gives us a quick way to determine what to expect to see

The Empirical Rule

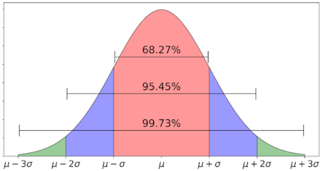

- The empirical rule gives us a quick way to determine what to expect to see
  - 68% of observations fall within $\mu \pm \sigma$.

The Empirical Rule

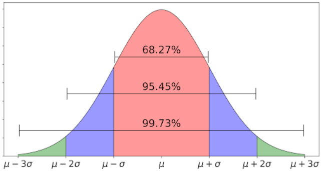

- The empirical rule gives us a quick way to determine what to expect to see
  - 68% of observations fall within $\mu \pm \sigma$.
  - 95% of observations fall within $\mu \pm 2\sigma$.

The Empirical Rule

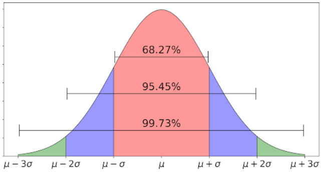

- The empirical rule gives us a quick way to determine what to expect to see
  - 68% of observations fall within $\mu \pm \sigma$.
  - 95% of observations fall within $\mu \pm 2\sigma$.
  - 99.7% of observations fall within $\mu \pm 3\sigma$.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.
  - This works well if $np \geq 10$ and $n(1-p) \geq 10$.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.
  - This works well if $np \geq 10$ and $n(1-p) \geq 10$.
- Note, because the distribution is discrete we need to use a continuity correction.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.
  - This works well if $np \geq 10$ and $n(1-p) \geq 10$.
- Note, because the distribution is discrete we need to use a continuity correction.
  - If we wish to consider $X \leq x$, then we should really consider $X \leq x + 0.5$

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.
  - This works well if $np \geq 10$ and $n(1-p) \geq 10$.
- Note, because the distribution is discrete we need to use a continuity correction.
  - If we wish to consider $X \leq x$, then we should really consider $X \leq x + 0.5$
  - If we wish to consider $X \geq x$, then we should really consider $X \geq x - 0.5$.

Normal Approximation to the Binomial Distribution

- If n is large, and p is closed to 0.5, then the binomial distribution is well approximated by a normal distribution.
- Specifically, we would have $X \dot{\sim} N(np, np(1-p))$.
  - This works well if $np \geq 10$ and $n(1-p) \geq 10$.
- Note, because the distribution is discrete we need to use a continuity correction.
  - If we wish to consider $X \leq x$, then we should really consider $X \leq x + 0.5$
  - If we wish to consider $X \geq x$, then we should really consider $X \geq x - 0.5$.
  - If we want $X > x$ this is $X \geq \lfloor x - \epsilon \rfloor$ (and vice versa for $X < x$).

Why Use the Normal Approximation?

- It is generally easier to reason about normally distributed variables.

Why Use the Normal Approximation?

- It is generally easier to reason about normally distributed variables.
  - As you progress through statistics and probability, your intuition for normal distributions will improve.

Why Use the Normal Approximation?

- It is generally easier to reason about normally distributed variables.
  - As you progress through statistics and probability, your intuition for normal distributions will improve.
- Because of normal tables, it is generally easier computationally to use the approximation.

Why Use the Normal Approximation?

- It is generally easier to reason about normally distributed variables.
  - As you progress through statistics and probability, your intuition for normal distributions will improve.
- Because of normal tables, it is generally easier computationally to use the approximation.
  - With modern software, this is becoming less of an issue.

Why Use the Normal Approximation?

- It is generally easier to reason about normally distributed variables.
  - As you progress through statistics and probability, your intuition for normal distributions will improve.
- Because of normal tables, it is generally easier computationally to use the approximation.
  - With modern software, this is becoming less of an issue.
- It gives us access to the empirical rule.

Summary

- The normal distribution is an incredibly important distribution for characterizing natural processes.
- It is characterized by its mean and variance, and has a closed form PDF (but not CDF).
- Translations to the standard normal permit easier calculations with standard critical values and the empirical rule.
- The normal distribution can approximate certain binomial distributions, so long as continuity corrections are applied.