22.4UC2

Numerical Computation and Statistics in Engineering

Unit 2: Probability Distributions

Random Variables

A *random variable* is a function taking numerical values which is defined over a sample space. Such a random variable is called *discrete* if it only takes countably many values.

Example. A quality control engineer checks randomly the content of bags, each containing 100 resistors. He selects 2 resistors and measures whether they match the specification (exact value plus or minus 10% tolerance). The number of resistors not matching the specification is a discrete random variable.

Another random variable would be the function taking values 0 and 1, for the outcomes that there are faulty resistors in the bag, or not. □

The *probability distribution* of a random variable is a table, graph, or formula that gives for each possible value of the random variable its probability. The requirements are that

$$0 \leq p(x) \leq 1 \qquad \text{and} \qquad \sum_{\text{all } x} p(x) = 1.$$

Random Variables

The following two diagrams show examples of discrete probability distributions:

Binomial Distribution, n = 10, p = 0.3

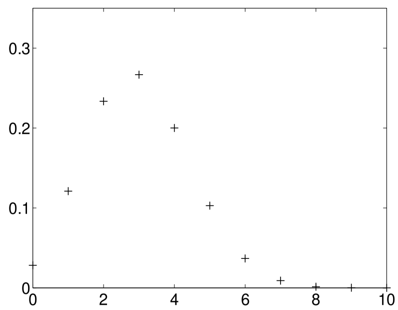

Poisson Distribution, λ = 2

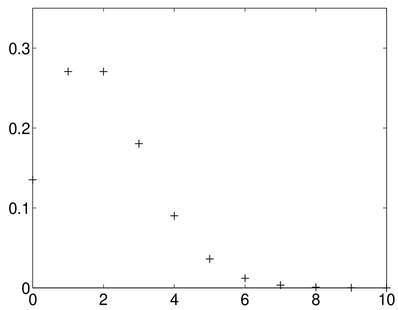

Expected Value

For a discrete random variable x with probability distribution $p(x)$ the *expected value* (or *mean*) is defined as

$$\mu = E(x) = \sum_{\text{all } x} x \cdot p(x)\,.$$

Example. We consider throwing a fair die 6000 times. We expect roughly 1000 outcomes of each possible observations 1, ... , 6. Thus the arithmetic mean of such an experiment will be approximately

$$1\frac{1000}{6000} + 2\frac{1000}{6000} + \cdots + 6\frac{1000}{6000} = 3.5\,.$$

The expected value is $\sum_{i=1}^{6} ip(i) = \frac{1}{6} \cdot 21 = 3.5$, as expected. □

Expected Value

Let x be any discrete random variable with probability distribution $p(x)$, and let g be any function of x. Then the *expected value of* $g(x)$ is defined as

$$E[g(x)] = \sum_{\text{all } x} g(x) \cdot p(x)\,.$$

The *variance* of a discrete random variable x with probability distribution $p(x)$ is defined as

$$\sigma^2 = E[(x-\mu)^2]\,,$$

the *standard deviation* is defined as $\sigma = \sqrt{E[(x-\mu)^2]}$.

Example. We return to the example of throwing a die. For the variance we find

$$\sigma^2 = (1-3.5)^2\frac{1}{6} + (2-3.5)^2\frac{1}{6} + \cdots + (6-3.5)^2\frac{1}{6} \approx 2.917\,.$$

For the standard deviation we find $\sigma \approx 1.708$. □

Properties of the Expected Value

Let x be a discrete random variable with probability distribution $p(x)$.

- $E(c) = c$, for every constant c;
- $E(cx) = cE(x)$, for every constant c;
- $E[g_1(x) + g_2(x)] = E[g_1(x)] + E[g_2(x)]$,
  for any two functions g_1, g_2 on x.

It follows the important formula that

$$\sigma^2 = E[x^2] - \mu^2\,.$$

For the proof of this formula note that

$$\begin{aligned} \sigma^2 &= E[(x-\mu)^2] = E[x^2 - 2\mu x + \mu^2] \\ &= E[x^2] - 2\mu E[x] + \mu^2 E[1] \\ &= E[x^2] - 2\mu\mu + \mu^2\,. \end{aligned}$$

The Binomial Probability Distribution

Example.

- Tossing a coin 10 times.
- Questioning 100 people on Princess Street in Edinburgh whether they know that Madonna's wedding takes place in a Scottish castle.
- Checking whether lots of transistors contain faulty transistors or not.

□

These experiments or observations are all examples of what is called a *binomial experiment* (the corresponding discrete random variable is called a *binomial random variable*).

The Binomial Probability Distribution

The examples have the following common characteristics:

- The experiment consists of n identical trials.
- In each trial there are exactly two possible outcomes (yes/no, pass/failure, or success/failure), denoted here 0 and 1 (for success).
- The probabilities for the outcomes 0 and 1 are the same in each trial (the trials are independent). They are usually denoted $p = P(‘1’)$ and $q = 1 - p = P(‘0’)$.
- The discrete (binomial) random variable is the number of successes (i.e., of 1's) in the n trials.

The Binomial Probability Distribution

The *binomial probability distribution* is given by the formula

$$p(x) = \binom{n}{x} p^x q^{n-x}, \qquad x \in \{0, \dots, n\},$$

where

- p is the probability of a success in a single trial, and $q = 1 - p$;
- n is the number of trials; and
- x is the number of successes.

The expected value (mean) and standard deviation are given by

$$\mu = np \qquad \text{and} \qquad \sigma = \sqrt{npq}\,.$$

The Binomial Probability Distribution

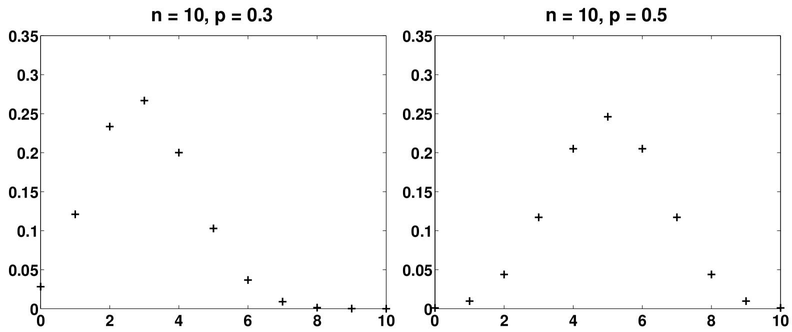

The Binomial Probability Distribution

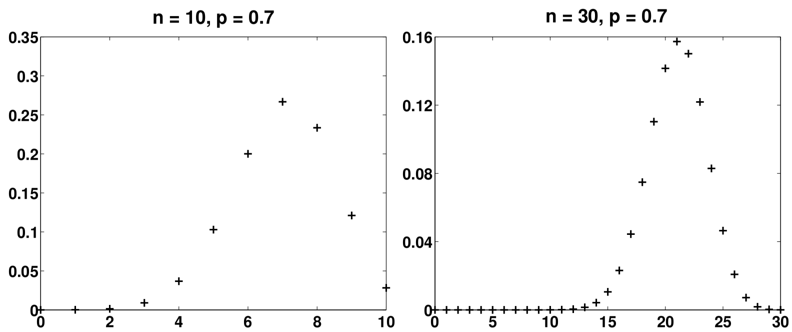

The Binomial Probability Distribution

Example. Tests show that about 20% of all private wells in some specific region are contaminated. What are the probabilities that in a random sample of 4 wells exactly 2, fewer than 2, or at least 2 wells are contaminated?

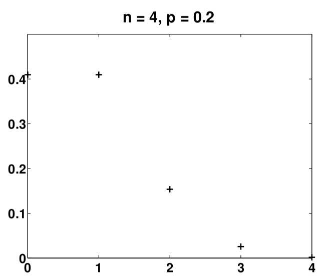

Here $n = 4$, $p = 0.2$ (success for being contaminated). We find

$$\begin{aligned} P(\text{‘}x=2\text{’}) &= \binom{4}{2}0.2^2 0.8^{4-2} = 0.1536\,, \\ P(\text{‘}x<2\text{’}) &= P(\text{‘}x=0\text{’}) + P(\text{‘}x=1\text{’}) = \binom{4}{0}0.2^0 0.8^4 + \binom{4}{1}0.2^1 0.8^3 = 0.8192\,, \\ P(\text{‘}x\geq 2\text{’}) &= P(\text{‘}x=2\text{’}) + P(\text{‘}x=3\text{’}) + P(\text{‘}x=4\text{’}) \\ &= 0.1536 + \binom{4}{3}0.2^3 0.8^1 + \binom{4}{0}0.2^4 0.8^0 = 0.1808\,. \end{aligned}$$

The Geometric Probability Distribution

Example. Customers wait in line to be served at a wicket. Per time interval the probability that a customer is served is 10%. What is the probability that a customer has to wait 15 time intervals before being served? □

Such and similar events are modeled by the *geometric probability distribution*. Each time interval we have an 'independent experiment' which can succeed or fail with success probability p (as for the binomial probability distribution). To be successful in the xth try we need $x-1$ failures (with probability $q = 1-p$) and one success (with probability p).

The Geometric Probability Distribution

The data for the geometric probability distribution are

- $p(x) = pq^{x-1}, \qquad x = 1, 2, \ldots,$
  where x is the number of trials until the first success; and
- $\mu = \frac{1}{p}$, and $\sigma = \sqrt{\frac{q}{p^2}}$.

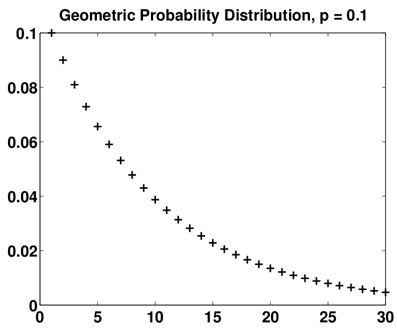

The Geometric Probability Distribution

Example. The average life expectancy of a fuse is 15 months. What is the probability that the fuse will last exactly 20 months?

We have that $\mu = 15$ (months), or $p = \frac{1}{15}$, which is the probability that a fuse will break. For $x = 20$ we obtain

$$P(\text{'}x = 20\text{'}) = \frac{1}{15}(1 - \frac{1}{15})^{20-1},$$

which is approximately 0.018. For σ we find $\sqrt{210} = 14.49$. □

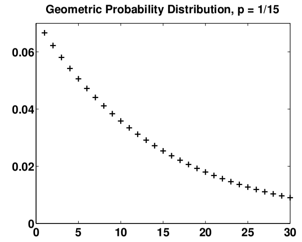

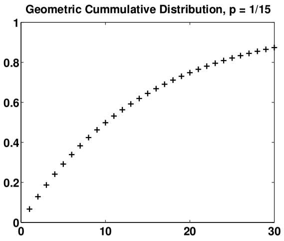

The Hypergeometric Distribution

The binomial and the geometric probability distribution are to be applied if, after observing a result, the sample is put back into the population. However, in practice, we often sample without replacement:

Example.

- If we test a bag of 1000 resistors whether they meet the specification we usually won't put back the tested items.
- Suppose people are randomly selected at Princess Street in Edinburgh to fill in a questionnaire about a new product. When people are approached they are usually first asked whether they have already taken part in this marketing research.
- A big manufacturing company maintains their machines on a regular basis. Suppose that on average 15% of the machines need repair. What is the probability that among the five machines inspected this week, one of them needs repair?

The Hypergeometric Distribution

- A box of 1000 fuses is tested one by one until the first defective fuse is found. Supposing that about 5% of the fuses are defective, what is the probability that a defective fuse is among the first 5 fuses tested?

$\square$

Such and similar random variables have a *hypergeometric* probability distribution:

- The populations consists if N objects.
- The possible outcomes of the experiment are success or failure.
- Each sample of size n is equally likely to be drawn.

The Hypergeometric Distribution

The data for the hypergeometric probability distribution are

$$p(x) = \frac{\binom{n}{x}\binom{N-r}{n-x}}{\binom{N}{n}}, \qquad 0, n-N+r \le x \le n, r,$$

where

- N is the number of elements in the population;
- r is the number in the population for success;
- n is the number of elements drawn; and
- x is the number of successes in the n randomly drawn elements.

The mean and standard deviation are given by

$$\mu = n\frac{r}{N}, \qquad \text{and} \qquad \sigma = \sqrt{\frac{r(N-r)n(N-n)}{N^2(N-1)}}.$$

The Hypergeometric Distribution

If we write $p = \frac{r}{N}$ then $\mu = np$ and $\sigma = \sqrt{\frac{N-n}{N-1} np(1-p)}$. This shows that the binomial and the hypergeometric distributions have the same expected value, but different standard deviations. The correction factor $\frac{N-n}{N-1}$ is less than 1, but close to 1 if n is small relative to N.

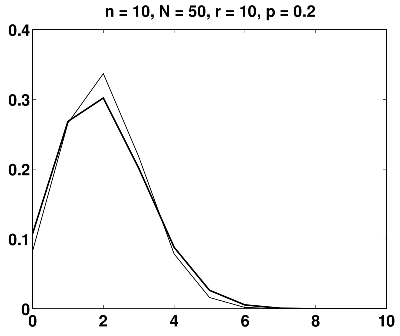

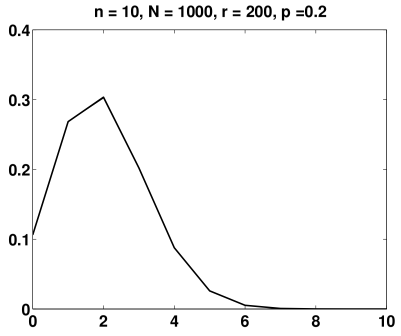

The Hypergeometric Distribution

Example. A retailer sells computers. He buys lots of 10 motherboards from a manufacturer who sells them cheaply, but also offers low quality. Suppose the current lot contains one defective item. If the retailer usually tests 4 items per lot, what is the probability that the lot is accepted?

Here $N = 10$, $r = 1$, and $n = 4$, and we are looking for $P(\text{‘}x = 0\text{’})$, which is

$$P(\text{‘}x = 0\text{’}) = \frac{\binom{1}{0}\binom{9}{4}}{\binom{10}{4}} = \frac{1 \cdot 9 \cdot 8 \cdot 7 \cdot 6}{1 \cdot 2 \cdot 3 \cdot 4} \frac{1 \cdot 2 \cdot 3 \cdot 4}{10 \cdot 9 \cdot 8 \cdot 7} = \frac{6}{10}.$$

We would use the same calculation if we would only know that *on average* 10% of the motherboards are faulty. □

The Hypergeometric Distribution

Example. We test lots of 100 fuses. On average 5% of the fuses are defective If we test 4 fuses, what is the probability that we accept the current lot?

Again, the random variable is hypergeometric, and since $N = 100$ is large we can assume that there are 5 defective fuses in this lot. We find

$$P(\text{‘}x = 0\text{’}) = \frac{\binom{5}{0}\binom{95}{5}}{\binom{100}{5}} = \frac{\frac{5!}{0!5!}\frac{95!}{5!90!}}{\frac{100!}{5!95!}} = \frac{95 \cdot 94 \cdot 93 \cdot 92 \cdot 91}{100 \cdot 99 \cdot 98 \cdot 97 \cdot 96} \approx 0.7696.$$

Later we will see how reliable this value is, as we don't know the exact number of faulty fuses in this lot. □

The Poisson Distribution

The *Poisson probability distribution* provides a model for the frequency of events, like the number people arriving at a counter, the number of plane crashes per month, or the number of micro-cracks in steel. (Micro-cracks in steel wheels of the German high-speed train ICE led to a disastrous rail accident in 1998.) The characteristics of a Poisson random variable are as follows:

- The experiment consists of counting events in a particular unit (time, area, volume, etc.).
- The probability that an event occurs in a given unit is the same for every unit.
- The number of events that occur in one unit is independent of the number of events that occur in other units.

The Poisson Distribution

The Poisson probability distribution with mean λ is given by the formula

$$p(x) = \frac{\lambda^x e^{-\lambda}}{x!}, \qquad (x = 0, 1, 2, \dots),$$

where e is the constant 2.71828. The expected value and standard deviation are

$$\mu = \lambda, \qquad \text{and} \qquad \sigma = \sqrt{\lambda}.$$

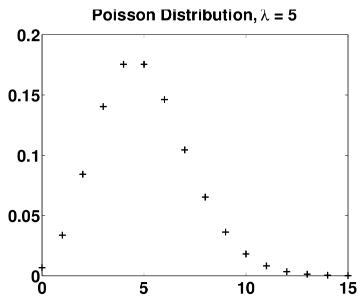

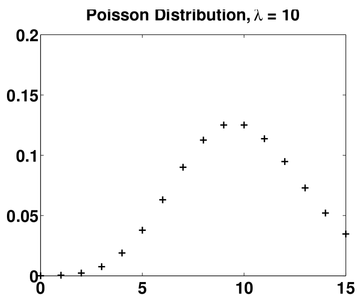

The Poisson Distribution

Example. Suppose customers arrive at a counter at an average rate of 6 per minute, and suppose that the random variable 'customer arrival' has a Poisson distribution. What is the probability that in a half-minute interval at most one new customer arrives? Here $\lambda = \frac{6}{2} = 3$ customers per half-minute. So

$$\begin{aligned} P(`x \le 1') &= P(`x = 0') + P(`x = 1') \\ &= \tfrac{e^{-3}3^0}{0!} + \tfrac{e^{-3}3^1}{1!} \\ &= \tfrac{4}{e^3}, \end{aligned}$$

which equals approximately 0.199. □

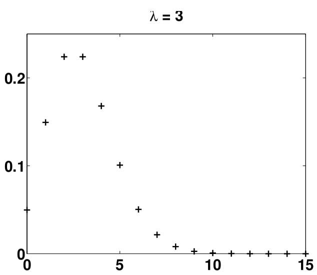

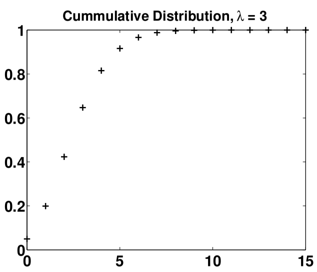

The Poisson Distribution

As an example we will verify that the Poisson probability distribution $p(x)$ really is a distribution, and that the mean is λ.

We note first that $0 \le p(x)$ for all values of x. Also, since

$$e^{\lambda} = \sum_{x=0}^{\infty} \frac{\lambda^x}{x!}\,,$$

$1 = e^{-\lambda}e^{\lambda} = e^{-\lambda}\sum_{x=0}^{\infty}\frac{\lambda^x}{x!} = \sum_{x=0}^{\infty}\frac{e^{-\lambda}\lambda^x}{x!} = \sum_{x=0}^{\infty}p(x)$, which shows that $p(x) \le 1$ and $\sum_{\text{all } x} p(x) = 1$.

For the mean we calculate

$$\begin{aligned} E(x) &= \sum_{x=0}^{\infty} x\frac{e^{-\lambda}\lambda^x}{x!} = 0 + \sum_{x=1}^{\infty} x\frac{e^{-\lambda}\lambda^x}{x!} \\ &= \sum_{x=1}^{\infty}\frac{\lambda e^{-\lambda}\lambda^{x-1}}{(x-1)!} = \lambda\sum_{x=0}^{\infty}\frac{e^{-\lambda}\lambda^x}{x!} = \lambda\,. \end{aligned}$$

Continuous Random Variables

Many random variables arising in practice are not discrete. Examples are the strength of a beam, the height of a person, or the capacity of a conductor. Such random variables are called *continuous*.

A practical problem arises, as it is *impossible* to assign finite amounts of probabilities to uncountably many values of the real line (or some interval) so that the values add up to 1. Thus, continuous probability distributions are usually based on *cumulative distribution functions*.

The cumulative distribution $F(x)$ of a random variable x is the function

$$F(x_0) = P(\text{'}x \le x_0\text{'})\,.$$

Continuous Random Variables

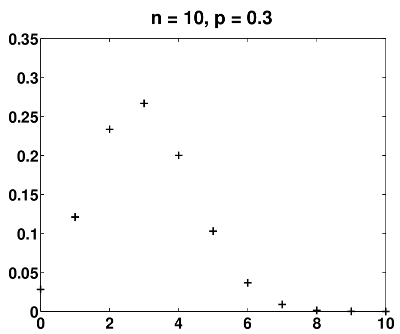

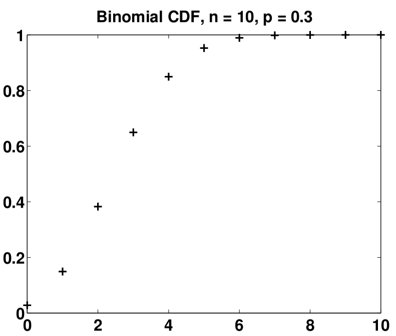

Continuous Random Variables

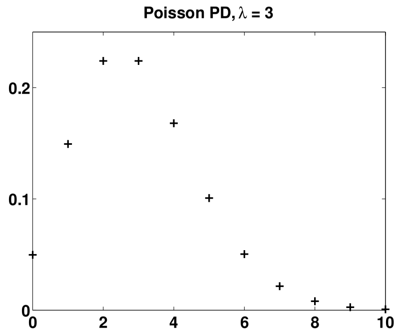

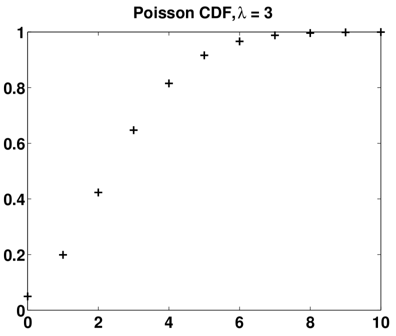

Density Functions

If F is the cumulative distribution of a continuous random variable x then the *density* function $\rho(x)$ for x is given by

$$\rho(x) = \frac{dF}{dx}$$

(provided that F is differentiable). It follows that

$$F(x) = \int_{-\infty}^{x} \rho(t)\, dt\,.$$

Moreover, the density function always satisfies the following two properties:

- $\rho(x) \geq 0$; and
- $\int\limits_{-\infty}^{\infty} \rho(t)\, dt = 1.$

In particular, $P(a < x < b) = P(a \leq x \leq b) = \int\limits_{a}^{b} \rho(t)\, dt.$

Density Functions

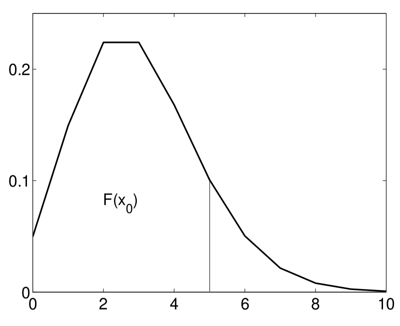

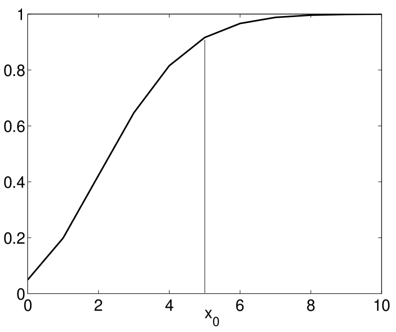

Expected Values

Let us recall from calculus that an integral is a limit process of a summation. Finding

$$F(x_0) = \int_{-\infty}^{x_0} \rho(x)\,dx$$

for a continuous random variable is analogous to finding

$$F(x_0) = \sum_{x \le x_0} p(x)$$

for a discrete random variable. Thus, we define the *expected value* analogous to the discrete case.

The *expected value* of a continuous random variable x with density function $\rho(x)$ is given by

$$\mu = E(x) = \int_{-\infty}^{\infty} t\rho(t)\,dt\,.$$

Expected Values

If g is any function we define the *expected value of* $g(x)$ as

$$E[g(x)] = \int_{-\infty}^{\infty} g(t)\rho(t)\,dt\,,$$

provided that these integrals exist. The *standard deviation* is $\sigma = \sqrt{E[(x-\mu)^2]}$. Note that

- $E(c) = c$, for every constant c;
- $E(cx) = cE(x)$, for every constant c;
- $E[g_1(x) + g_2(x)] = E[g_1(x)] + E[g_2(x)]$, for any two functions g_1, g_2 on x.
- $\sigma^2 = E[x^2] - \mu^2$.

An Example

Example. We consider the density function

$$\rho = \begin{cases} \frac{1}{2}e^{-\frac{x}{2}} & \text{if } 0 \le x < \infty, \\ 0 & \text{else.} \end{cases}$$

This density function is everywhere positive, and for $x \le 0$, $F(x) = \int_{-\infty}^{x} \rho(t)\,dt = 0$, whereas for $x \ge 0$

$$\begin{aligned} F(x) &= \int_{-\infty}^{x} \rho(t)\,dt = \int_0^x \frac{1}{2}e^{-\frac{t}{2}}\,dt \\ &= \left[-e^{-\frac{t}{2}}\right]_0^x = e^0 - e^{-\frac{x}{2}} \end{aligned}$$

In particular,

$$\begin{aligned} \int_{-\infty}^{\infty} \rho(t)\,dt &= \lim_{x\to\infty} F(x) \\ &= e^0 - \lim_{x\to\infty} e^{-\frac{x}{2}} = e^0 - 1 = 1\,. \end{aligned}$$

An Example

For the expected value we find

$$\begin{aligned}\mu &= \int_{-\infty}^{\infty} t\rho(t)\,dt \;=\; \int_{-\infty}^{\infty} t\rho(t)\,dt \\ &= \lim_{x\to\infty}\int_0^x t\frac{1}{2}e^{-\frac{t}{2}}\,dt \qquad\qquad \Big|\int te^{at} = \frac{e^{at}}{a^2}(at-1) \\ &= \frac{1}{2}\lim_{x\to\infty}\left[4e^{-\frac{t}{2}}(2\frac{t}{2}-1)\right]_0^x \\ &= \frac{1}{2}[0-4e^0(-0-1)] \;=\; 2\,.\end{aligned}$$

A similar calculation shows that $E(x^2)=8$, so that $\sigma=\sqrt{E(x^2)-\mu^2}=\sqrt{8-4}=2$. Finally, to do another calculation,

$$\begin{aligned}P(\mu-\sigma<x<\mu+\sigma) &= \int_0^4 \frac{1}{2}e^{-\frac{t}{2}}\,dt \;=\; -e^{-\frac{t}{2}}\Big|_0^4 \\ &= -e^{-2}+e^0 \;=\; 1-\frac{1}{e^2} \;=\; 0.8647\,.\end{aligned}$$

An Example

The empirical rule of Unit 1 suggested 68%. □

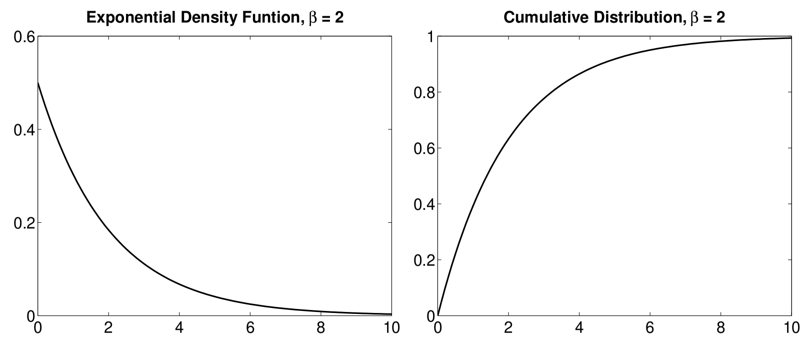

The Uniform Probability Distribution

If we select randomly a number in the interval $[a,b]$ then the corresponding random variable x is called a *uniform random variable*. Its density function is

$$\rho = \begin{cases} \frac{1}{b-a} & \text{if } a\le x\le b, \\ 0 & \text{else.}\end{cases}$$

For the mean and standard deviation one finds

$$\mu=\frac{a+b}{2} \qquad \text{and} \qquad \sigma=\frac{b-a}{2\sqrt{3}}=\frac{\sqrt{3}}{6}(b-a)\,.$$

The Uniform Probability Distribution

Example. A manufacturer of wires believes that one of her machines makes wires with diameter uniformly distributed between 0.98 and 1.03 millimeters.

The mean of the thickness is $\frac{1.03+0.98}{2} = 1.005$ millimeters, and the standard deviation is $\sigma = \frac{\sqrt{3}}{6}(1.03 - 0.98) \approx 0.014$ millimeters. The density function for this uniform random variable is $\rho = \frac{1}{.05} = 20$.

For example,

$$\begin{aligned} P(\text{`}x \leq 1.00\text{'}) &= \int_{0.98}^{1.00} 20\, dt \\ &= 20[1.00 - 0.98] \\ &= 0.4\,. \end{aligned}$$

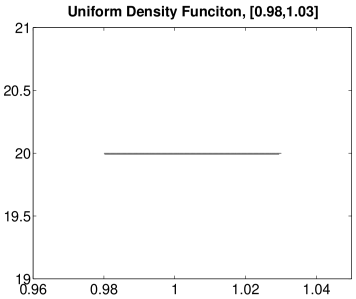

The Normal Probability Distribution

The *normal* probability distribution was suggested by C. F. Gauss as a model of the relative frequency distribution of errors (for example in measurements). The density function of this probability distribution is

$$\rho(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \qquad -\infty < x < \infty\,,$$

where μ and σ denote the mean and standard deviation, respectively (so these two values are parameters of the normal probability distribution).

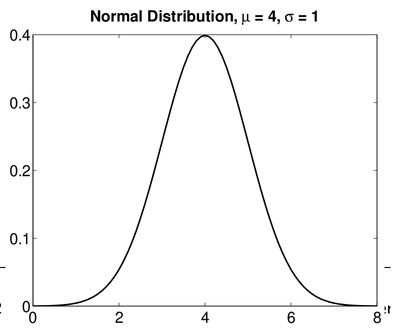

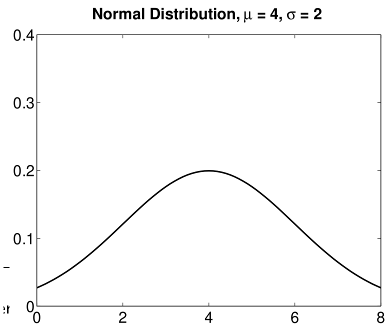

The Normal Probability Distribution

The *standard normal random variable* has mean 0 and variance 1:

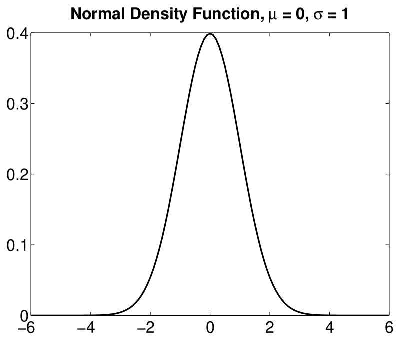

In practice it is enough to have tables for the standard normal probability distribution: Given a random variable x the variable $z = \frac{x-\mu}{\sigma}$ has mean 0 and standard deviation 1.

The Normal Probability Distribution

Example. Suppose a normally distributed random variable x has mean 10 and standard deviation 3. Find $P(`x \le 11')$ using tables.

We set $z = \frac{x-10}{3}$, which has standard normal distribution. The x-value 11 corresponds to the z-value $\frac{11-10}{3} = \frac{1}{3}$. Then the table says $P(`x \le 11') = P(`z \le \frac{1}{3}') = P(`z \le 0') + P(`0 \le z \le \frac{1}{3}') \approx .5 + 0.1293 = 0.6293.$ □

Why is this justified? In the integral calculating $P(`x \le 11')$ we substitute $z = \frac{x-\mu}{\sigma}$. Then $\frac{dz}{dx} = \frac{1}{\sigma}$, and

$$\begin{aligned} P(`x \le 10') &= \int_{-\infty}^{11} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}\, dx \\ &= \int_{-\infty}^{\frac{1}{3}} \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}\, dx \\ &= P(`z \le \tfrac{1}{3}'). \end{aligned}$$

The Normal Probability Distribution

Example. An amplifier is built using two integrated circuits. Both have a life-length that is normally distributed, the first with mean 36000 hours and standard deviation 8000 hours, the second with mean 38000 hours and standard deviation 85000 hours. Which of the two integrated circuits is more likely to last at least 40.000 hours?

In both cases we ask for $P(`x \ge 40000') = 1 - P(`x \le 40000')$. The corresponding values for the standardised normal random variables z_1 and z_2 are $z_1 = \frac{1}{2}$, and $z_2 = \frac{1}{5}$. Thus

$$P(`x_1 \ge 40000') = 1 - P(`z_1 < \tfrac{1}{2}') \approx 1 - (0.5 + 0.1915) = 0.3085\,,$$

and similarly, $P(`x_2 \ge 40000') = 1 - P(`z_2 < \frac{1}{5}') \approx 1 - (0.5 + 0.0793) = 0.4207$. Thus, the second integrated circuit is more likely to last more than 40000 hours. □

The Gamma Distribution

Many continuous random variables can only take positive values, like height, thickness, life expectations of transistors, etc. Such random variables are often modelled by *gamma type random variables*. The corresponding density functions contain two parameters α, β. The first is known as the *shape* parameter, the second as the *scale* parameter.

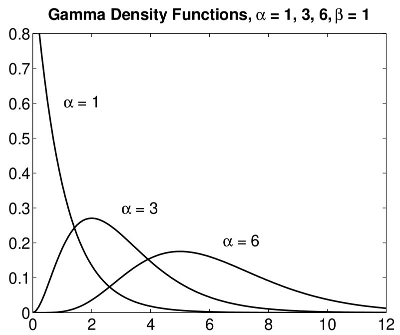

The Gamma Distribution

The density function is given by

$$\rho(x) = \begin{cases} \frac{x^{\alpha-1}e^{-\frac{x}{\beta}}}{\beta^\alpha\Gamma(\alpha)} & \text{if } 0 \le x < \infty,\ \alpha,\beta > 0, \\ 0 & \text{else,} \end{cases}$$

where $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1}e^{-t}\,dt$. The mean and standard deviation are

$$\mu = \alpha\beta \qquad \text{and} \qquad \sigma = \sqrt{\alpha\beta^2}\,.$$

The gamma function plays an important role in mathematics. It holds that $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$, and $\Gamma(1) = 1$, so that for integer values of α, $\Gamma(\alpha) = \alpha!$. In general there is no closed form for the gamma function, and its values are approximated and taken from tables.

The Gamma Distribution

Example. A manufacturer of CPU's knows that the relative frequency complaints from customers (in weeks) about total failures is modeled by a gamma distribution with $\alpha = 2$ and $\beta = 4$. Exactly 12 weeks after the quality control department was restructured the next (first) major complaint arrives. Does this suggest that the restructuring resulted in an improvement of quality control?

We calculate $\mu = \alpha\beta = 8$ and $\sigma = 4\sqrt{2} \approx 5.657$. The value $x = 12$ lies well within one standard deviation from the (old) mean, so we would not consider it an exceptional value. Thus there is insufficient evidence to indicate an improvement in quality control given just this data. □

The Chi-Square Distribution

The χ^2 *(chi-square) probability distribution* plays an important role in statistics. The distribution is a special case of the gamma distribution for $\alpha = \frac{\nu}{2}$ and $\beta = 2$ (ν is called the *number of degrees of freedom*):

$$\rho(\chi^2) = c(\chi^2)^{\frac{\nu}{c}-1}e^{-\frac{\chi^2}{2}}\,,$$

where $c(\chi^2) = \frac{1}{2^{\frac{\nu}{2}}\Gamma(\frac{\nu}{2})}$. For mean and standard deviation one finds

$$\mu = \nu \qquad \text{and} \qquad \sigma = \sqrt{2\nu}\,.$$

The Exponential Density Function

The *exponential density function* is a gamma density function with $\alpha = 1$,

$$\rho(x) = \frac{e^{-\frac{x}{\beta}}}{\beta}, \qquad x \geq 0\,,$$

with mean $\mu = \beta$ and standard deviation $\sigma = \beta$. The corresponding random variable models for example the length of time *between* events (arrivals at a counter, requests to a CPU, etc) when the probability of an arrival in an interval is independent from arrivals in other intervals. This distribution also models the life expectancy of equipment or products, provided that the probability that the equipment will last t more time intervals is the same as for a new product (this holds for well-maintained equipment).

If the arrival of events follows a Poisson distribution with mean $\frac{1}{\beta}$ (arrivals per unit interval), then the time interval between two successive arrivals is modeled by the exponential distribution with mean β.

The Weibull Density Function

As the gamma probability distribution the *Weibull probability distribution* is often used to model length of life of products, equipment, or components. The density function is

$$\rho(x) = \begin{cases} \frac{\alpha}{\beta} x^{\alpha-1} e^{-\frac{x^\alpha}{\beta}} & \text{if } x \geq 0, \\ 0 & \text{else,} \end{cases}$$

with *shape parameter* α and *scale parameter* β. Moreover,

- $\mu = \beta^{\frac{1}{\alpha}} \Gamma(\frac{\alpha+1}{\alpha})$,
- $\sigma = \sqrt{\beta^{\frac{2}{\alpha}} [\Gamma(\frac{\alpha+2}{\alpha}) - \Gamma(\frac{\alpha+1}{\alpha})]}$.

For $\alpha = 1$ we get the exponential density function.

The Weibull Density Function

The Weibull cumulative distribution has a closed form; after substituting $y = x^\alpha$ and $dy = \alpha x^{\alpha-1}\, dx$ we find

$$\begin{aligned} F(x \leq x_0) &= \int_0^{x_0} \frac{\alpha}{\beta} x^{\alpha-1} e^{-\frac{x^\alpha}{\beta}}\, dx \\ &= \int_0^{x_0^\alpha} \frac{1}{\beta} e^{-\frac{y}{\beta}}\, dy \\ &= -e^{-\frac{y}{\beta}} \Big|_0^{x_0^\alpha} \\ &= 1 - e^{-\frac{x_0^\alpha}{\beta}} \end{aligned}$$

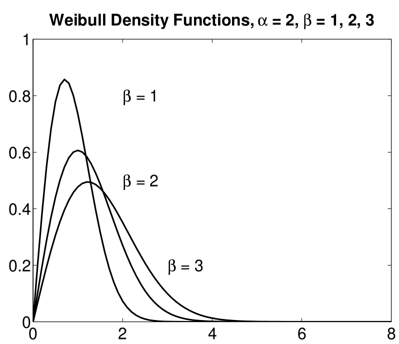

The Weibull Density Function

Example. The length of life in years of a component in a camera is known to have a Weibull distribution with $\alpha = 2$ and $\beta = 100$. What is the probability that the component will last at least 6 years?

We are looking for $P(`x \geq 6')$ which is

$$\begin{aligned} P(`x \geq 6') &= 1 - P(`x \leq 6') \\ &= 1 - (1 - e^{-\frac{6^2}{100}}) \\ &= \frac{1}{e^{\frac{36}{100}}} \\ &\approx 0.698\,. \end{aligned}$$

□

Summary

- Random variables are functions assigning numerical values to each simple event of a sample space. We distinguish discrete and continuous random variables.
- The probability distribution of a discrete random variable is a function that gives for each event the probability that the event occurs.
- The expected value $E(x)$ is the mean, the standard deviation the square root of $E[(x - E(x))^2]$.
- Examples of discrete probability distribution are the binomial, geometric, hypergeometric and the Poisson distribution.
- For continuous random variables we have to give the cumulative probability distribution.

Summary

- The relative frequency distribution for a population with continuous random variable can be modeled using a density function $\rho(x)$ (usually a smooth curve) such that

$$\rho(x) \geq 0 \qquad \text{and} \qquad \int_{\infty}^{-\infty} \rho(x)\, dx = 1\,.$$

- Examples are the uniform distribution, normal distribution, gamma distribution, the exponential distribution and the Weibull distribution.