Advanced Topics in Software Engineering: Performance Modelling and Evaluation

Prof. Michele Loreti

Advanced Topics in Software Engineering
Corso di Laurea in Informatica (L31)
Scuola di Scienze e Tecnologie

An example. . .

The municipality of *Neverland* want to activate a bike-sharing system.

An example...

The municipality of *Neverland* want to activate a bike-sharing system.

A bicycle-sharing system is a service in which:

- bicycles are made available for shared use to individuals on a very short term basis for a price;

An example...

The municipality of *Neverland* want to activate a bike-sharing system.

A bicycle-sharing system is a service in which:

- bicycles are made available for shared use to individuals on a very short term basis for a price;
- people can borrow a bike from point A and return it at point B;

An example...

The municipality of *Neverland* want to activate a bike-sharing system.

A bicycle-sharing system is a service in which:

- bicycles are made available for shared use to individuals on a very short term basis for a price;
- people can borrow a bike from point A and return it at point B;
- smartphone mapping apps show nearby stations with available bikes and open docks.

An example...

The municipality of *Neverland* want to activate a bike-sharing system.

A bicycle-sharing system is a service in which:

- bicycles are made available for shared use to individuals on a very short term basis for a price;
- people can borrow a bike from point A and return it at point B;
- smartphone mapping apps show nearby stations with available bikes and open docks.

An example...

The municipality of *Neverland* want to activate a bike-sharing system.

A bicycle-sharing system is a service in which:

- bicycles are made available for shared use to individuals on a very short term basis for a price;
- people can borrow a bike from point A and return it at point B;
- smartphone mapping apps show nearby stations with available bikes and open docks.

The civil council ask you to develop a software architecture that...

- supports users in station selections while improving their satisfaction;
- guarantees a balanced use of resources;
- identifies anomalous situations.

An example...

Question...

How can we design this kind of software architectures?

An example...

Question...

How can we design this kind of software architectures?

Answer...

We have to...

... build a model of our system...

An example...

Question...

How can we design this kind of software architectures?

Answer...

We have to...

... build a model of our system...

... define possible scenario of use...

An example...

Question...

How can we design this kind of software architectures?

Answer...

We have to...

... build a model of our system...

... define possible scenario of use...

... study system behaviour and the impact of implementation choices on these scenarios.

Key Notions

A model can be constructed to represent some aspect of the dynamic behaviour of a system.

Key Notions

A model can be constructed to represent some aspect of the dynamic behaviour of a system.

Once constructed, such a model becomes a tool with which we can investigate the behaviour of the system.

The discrete event view

In this course we will consider discrete event systems.

The discrete event view

In this course we will consider discrete event systems.

The state of the system is characterised by variables which take distinct values and which change by discrete events, i.e. at a distinct time something happens within the system which results in a change in one or more of the state variables.

Example: Bike Sharing System

Discrete event system

We might be interested in the number of available bikes and available slots at each station.

Example: Bike Sharing System

Discrete event system

We might be interested in the number of available bikes and available slots at each station.

Let $\{\ell_1, \dots \ell_n\}$ be the bike stations in our system, we can count the number of bikes (resp. slots) B_{ℓ_i} (resp. S_{ℓ_i}) available at each station ℓ_i.

- When a bike is retrieved form ℓ_i, B_{ℓ_i} is decreased by 1 and S_{ℓ_i} increased;

Example: Bike Sharing System

Discrete event system

We might be interested in the number of available bikes and available slots at each station.

Let $\{\ell_1, \dots \ell_n\}$ be the bike stations in our system, we can count the number of bikes (resp. slots) B_{ℓ_i} (resp. S_{ℓ_i}) available at each station ℓ_i.

- When a bike is retrieved form ℓ_i, B_{ℓ_i} is decreased by 1 and S_{ℓ_i} increased;
- When a bike is returned at ℓ_i, B_{ℓ_i} is incremented and S_{ℓ_i} decremented.

Discrete time vs Continuous time

Within discrete event systems there is a distinction between a discrete time representation and a continuous time representation:

Discrete time vs Continuous time

Within discrete event systems there is a distinction between a discrete time representation and a continuous time representation:

Discrete time: such models only consider the system at predetermined moments in time, which are typically evenly spaced, eg. at each clock "tick".

Discrete time vs Continuous time

Within discrete event systems there is a distinction between a discrete time representation and a continuous time representation:

Discrete time: such models only consider the system at predetermined moments in time, which are typically evenly spaced, eg. at each clock "tick".

Continuous time: such models consider the system at the time of each event so the time parameter in such models is conceptually continuous.

Discrete time vs Continuous time

Within discrete event systems there is a distinction between a discrete time representation and a continuous time representation:

Discrete time: such models only consider the system at predetermined moments in time, which are typically evenly spaced, eg. at each clock "tick".

Continuous time: such models consider the system at the time of each event so the time parameter in such models is conceptually continuous.

Discrete time vs Continuous time

Within discrete event systems there is a distinction between a discrete time representation and a continuous time representation:

Discrete time: such models only consider the system at predetermined moments in time, which are typically evenly spaced, eg. at each clock "tick".

Continuous time: such models consider the system at the time of each event so the time parameter in such models is conceptually continuous.

The use of discrete time or continuous time mainly depend on the levels of abstraction considered in the model.

Performance Modelling

Performance modelling is concerned with the dynamic behaviour of systems and quantified assessment of that behaviour.

Performance Modelling

Performance modelling is concerned with the dynamic behaviour of systems and quantified assessment of that behaviour.

There are often conflicting interests at play:

- Users typically want to optimise external measurements of the dynamics such as response time (as small as possible), throughput (as high as possible) or blocking probability (preferably zero);

Performance Modelling

Performance modelling is concerned with the dynamic behaviour of systems and quantified assessment of that behaviour.

There are often conflicting interests at play:

- Users typically want to optimise external measurements of the dynamics such as response time (as small as possible), throughput (as high as possible) or blocking probability (preferably zero);
- In contrast, system managers may seek to optimize internal measurements of the dynamics such as utilisation (reasonably high, but not too high), idle time (as small as possible) or failure rates (as low as possible).

Performance Modelling: Motivation

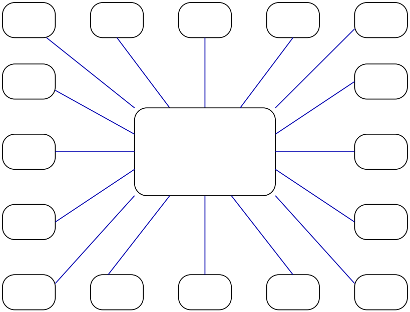

Capacity planning

- How many clients can the existing server support and maintain reasonable response times?

Performance Modelling: Motivation

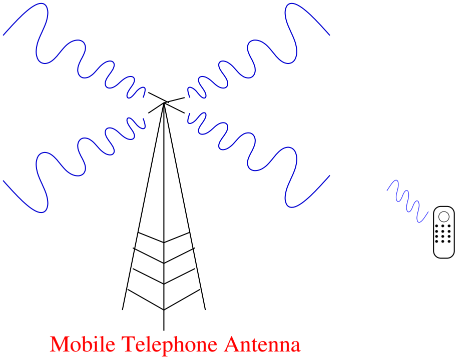

System Configuration

- How many frequencies do you need to keep blocking probabilities low?

Performance Modelling: Motivation

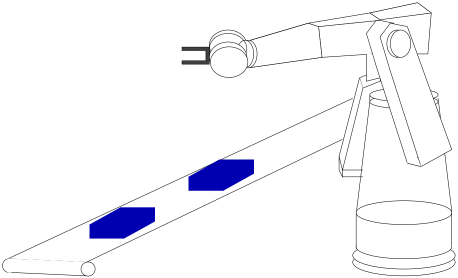

System Tuning

- What speed of conveyor belt will minimize robot idle time and maximize throughput whilst avoiding lost widgets?

Performance Modelling: Response time analysis

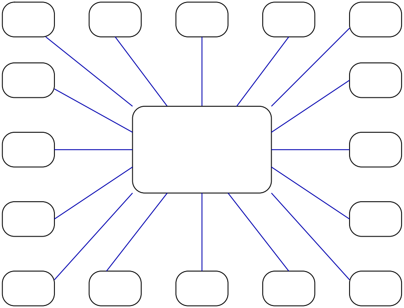

Quality of Service issues

- Can the server maintain reasonable response times?

Performance Modelling: Capacity planning

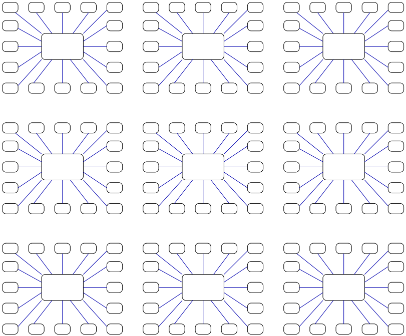

Scalability and capacity planning issues

- How many times do we have to replicate this service to support all of the subscribers?

Performance Modelling: Scalability analysis

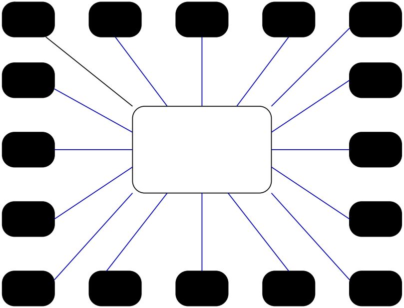

Robustness and scalability issues

- Will the server withstand a distributed denial of service attack?

Performance Modelling: Service Level Agreements

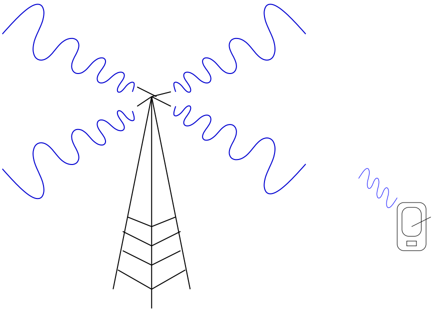

Service-level agreements

- What percentage of downloads do complete within the time we advertised?

Quantitative modelling

When systems are modelled to verify their functional behaviour (correctness), all definite values are abstracted away — qualitative modelling.

Quantitative modelling

When systems are modelled to verify their functional behaviour (correctness), all definite values are abstracted away — qualitative modelling.

In contrast, performance modelling is quantitative modelling as we must take into account explicit values for time (latency, service time etc.) and probability (choices, alternative outcomes, mixed workload).

Quantitative modelling

When systems are modelled to verify their functional behaviour (correctness), all definite values are abstracted away — qualitative modelling.

In contrast, performance modelling is quantitative modelling as we must take into account explicit values for time (latency, service time etc.) and probability (choices, alternative outcomes, mixed workload).

Probability will be used to represent randomness (e.g. from human users) but also as an abstraction over unknown values (e.g. service times).

Probability Theory

Sample space

A sample space is an arbitrary non empty set Ω, containing of all possible outcomes or results of an experiment.

A sample space is an arbitrary non empty set Ω, containing of all possible outcomes or results of an experiment.

Examples:

Toss of a coin: $\Omega_C = \{\mathsf{H}, \mathsf{T}\}$;

Toss of two coins: $\Omega_{2C} = \Omega_C \times \Omega_C = \{(\mathsf{H}, \mathsf{H}), (\mathsf{H}, \mathsf{T}), (\mathsf{T}, \mathsf{H}), (\mathsf{T}, \mathsf{T})\}$;

Roll of a dice: $\Omega_D = \{1, 2, 3, 4, 5, 6\}$.

Probability Theory

Events

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Examples: Toss of a coin

Let Σ_C be the σ-algebra on $\Omega_C = \{H, T\}$ containing:

- Neither *head* nor *tail*: $\{\}$

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Examples: Toss of a coin

Let Σ_C be the σ-algebra on $\Omega_C = \{\mathrm{H}, \mathrm{T}\}$ containing:

- Neither *head* nor *tail*: $\{\}$
- *Head*: $\{\mathrm{H}\}$

Probability Theory

Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Examples: Toss of a coin

Let Σ_C be the σ-algebra on $\Omega_C = \{\mathrm{H}, \mathrm{T}\}$ containing:

- Neither *head* nor *tail*: $\{\}$
- *Head*: $\{\mathrm{H}\}$
- *Tail*: $\{\mathrm{T}\}$

Probability Theory
Events

A σ-algebra $\Sigma \subseteq 2^{\Omega}$ (the powerset of Ω), called events, such that:

1. Σ contains the sample space, $\Omega \in \Sigma$;
2. Σ is closed under complement: if $A \in \Sigma$ then $(\Omega - A) \in \Sigma$;
3. Σ is close under countable unions.

Examples: Toss of a coin

Let Σ_C be the σ-algebra on $\Omega_C = \{\mathrm{H}, \mathrm{T}\}$ containing:

- Neither *head* nor *tail*: $\{\}$
- *Head*: $\{\mathrm{H}\}$
- *Tail*: $\{\mathrm{T}\}$
- Either *head* or *tail*: $\{\mathrm{H}, \mathrm{T}\}$

Probability Theory
Probability measure

Let $\Sigma \subseteq 2^{\Omega}$, be a σ-algebra. A probability measure is a function $\Pr : \Sigma \rightarrow [0,1]$ associating elements in Σ with a real value in $[0,1]$ such that:

- Pr is countably additive, if $\{A_i\}_{i \in I}$ is a countable collection of pairwise disjoint set, then:

$$\Pr\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \Pr(A_i)$$

Let $\Sigma \subseteq 2^{\Omega}$, be a σ-algebra. A probability measure is a function $Pr : \Sigma \rightarrow [0,1]$ associating elements in Σ with a real value in $[0,1]$ such that:

- Pr is countably additive, if $\{A_i\}_{i \in I}$ is a countable collection of pairwise disjoint set, then:

$$Pr\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} Pr(A_i)$$

- the measure of the entire sample space is 1:

$$Pr(\Omega) = 1$$

A probability space is a triple $(\Omega, \Sigma, \Pr)$ such that:

- Ω is a sample space;
- Σ is a σ-algebra on Ω;
- $\Pr$ is a probability measure for Σ.

A probability space is a triple $(\Omega, \Sigma, \Pr)$ such that:

- Ω is a sample space;
- Σ is a σ-algebra on Ω;
- $\Pr$ is a probability measure for Σ.

Examples: Toss of a coin

We can consider the probability space $(\Omega_C, \Sigma_C, \Pr_C)$ such that:

- $\Pr_C(\{\}) = 0$;
- $\Pr_C(\{H\}) = \Pr_C(\{T\}) = \frac{1}{2}$;
- $\Pr_C(\{H, T\}) = 1$.

Probability Space: Some Properties

Let $(\Omega, \Sigma, \Pr)$ be a probability space. The following properties hold:

- For any $A \in \Sigma$:

$$\Pr(\Omega - A) = 1 - \Pr(A)$$

Probability Space: Some Properties

Let $(\Omega, \Sigma, \Pr)$ be a probability space. The following properties hold:

- For any $A \in \Sigma$:
$$\Pr(\Omega - A) = 1 - \Pr(A)$$
- For any $A, B \in \Sigma$:
$$\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$$

Conditional Probability

Let $(\Omega, \Sigma, \Pr)$ be a probability space

Let $A, B \in \Sigma$, the Conditional Probability of A occurring, give that B has occurred, is:

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

Conditional Probability

Let $(\Omega, \Sigma, \Pr)$ be a probability space

Let $A, B \in \Sigma$, the Conditional Probability of A occurring, give that B has occurred, is:

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

- If A and B are mutually exclusive $\Pr(A \mid B) = 0$.

Conditional Probability

Let $(\Omega, \Sigma, \Pr)$ be a probability space

Let $A, B \in \Sigma$, the Conditional Probability of A occurring, give that B has occurred, is:

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

- If A and B are mutually exclusive $\Pr(A \mid B) = 0$.
- If B is a precondition for A, then $\Pr(A \cap B) = \Pr(A)$.

Conditional Probability

Let $(\Omega, \Sigma, \Pr)$ be a probability space

Let $A, B \in \Sigma$, the Conditional Probability of A occurring, give that B has occurred, is:

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}$$

- If A and B are mutually exclusive $\Pr(A \mid B) = 0$.
- If B is a precondition for A, then $\Pr(A \cap B) = \Pr(A)$.
- Two events are independent if knowledge of the occurrence of one of them tells us nothing about the probability of the other, i.e. $\Pr(A \mid B) = \Pr(A)$, or

$$\Pr(A \cap B) = \Pr(A) \times \Pr(B)$$

Random variables

We are interested in the dynamics of a system as events happen over time.

Random variables

We are interested in the dynamics of a system as events happen over time.

A function which associates a (real-valued) number with the outcome of an experiment is known as a random variable.

Random variables

We are interested in the dynamics of a system as events happen over time.

A function which associates a (real-valued) number with the outcome of an experiment is known as a random variable.

Let $(\Omega, \Sigma, \Pr)$ be a probability space, a random variable $X : \Omega \to \mathbb{R}$ is a measurable function from Ω to $\mathbb{R}$.

Random variables

We are interested in the dynamics of a system as events happen over time.

A function which associates a (real-valued) number with the outcome of an experiment is known as a random variable.

Let $(\Omega, \Sigma, \Pr)$ be a probability space, a random variable $X : \Omega \to \mathbb{R}$ is a measurable function from Ω to $\mathbb{R}$.

The probability that X takes value in a measurable set $S \subseteq \mathbb{R}$ is written as:

$$\Pr(X \in S) = \Pr(\{\omega \in \Omega | X(\omega) \in S\})$$

Distribution function

Let X a random variable on the probability space $(\Omega, \Sigma, \Pr)$, we define the distribution function F_X for each real $x \in \mathbb{R}$ by

$$F_X(x) = \Pr[X \leq x] = \Pr(\{\omega | X(\omega) \leq x\})$$

Distribution function

Let X a random variable on the probability space $(\Omega, \Sigma, \Pr)$, we define the distribution function F_X for each real $x \in \mathbb{R}$ by

$$F_X(x) = \Pr[X \leq x] = \Pr(\{\omega | X(\omega) \leq x\})$$

We associate another function $p_X(\cdot)$, called the probability mass function, with X (pmf), for each $x \in \mathbb{R}$:

$$p(x) = \Pr[X = x] = \Pr(\{\omega | X(\omega) \leq x\})$$

Distribution function

Let X a random variable on the probability space $(\Omega, \Sigma, \Pr)$, we define the distribution function F_X for each real $x \in \mathbb{R}$ by

$$F_X(x) = \Pr[X \leq x] = \Pr(\{\omega | X(\omega) \leq x\})$$

We associate another function $p_X(\cdot)$, called the probability mass function, with X (pmf), for each $x \in \mathbb{R}$:

$$p(x) = \Pr[X = x] = \Pr(\{\omega | X(\omega) \leq x\})$$

A random variable X is continuous if $p(x) = 0$ for all real x.

Distribution function

Let X a random variable on the probability space $(\Omega, \Sigma, \Pr)$, we define the distribution function F_X for each real $x \in \mathbb{R}$ by

$$F_X(x) = \Pr[X \leq x] = \Pr(\{\omega | X(\omega) \leq x\})$$

We associate another function $p_X(\cdot)$, called the probability mass function, with X (pmf), for each $x \in \mathbb{R}$:

$$p(x) = \Pr[X = x] = \Pr(\{\omega | X(\omega) \leq x\})$$

A random variable X is continuous if $p(x) = 0$ for all real x.

NB: If X is a continuous random variable, then X can assume infinitely many values, and so it is reasonable that the probability of its assuming any specific value we choose beforehand is zero.

Example: Dice Roll

A random variable can be used to describe the process of rolling two (fair) dice and the possible outcomes.

Example: Dice Roll

A random variable can be used to describe the process of rolling two (fair) dice and the possible outcomes.

We can consider the probability space $(\Omega_{2D}, \Sigma_{2D}, \mathrm{Pr}_{2D})$ such that:

$$\Omega_{2D} = \{(n_1, n_2) | 1 \leq n_1, n_2 \leq 6\} \qquad \Sigma_{2D} = 2^{\Omega_{2D}} \qquad Pr(A) = \frac{|A|}{36}$$

Example: Dice Roll

A random variable can be used to describe the process of rolling two (fair) dice and the possible outcomes.

We can consider the probability space $(\Omega_{2D}, \Sigma_{2D}, \mathrm{Pr}_{2D})$ such that:

$$\Omega_{2D} = \{(n_1, n_2) | 1 \leq n_1, n_2 \leq 6\} \qquad \Sigma_{2D} = 2^{\Omega_{2D}} \qquad Pr(A) = \frac{|A|}{36}$$

The total number rolled is then a random variable X given by the function that maps the pair to the sum: $X((n_1, n_2)) = n_1 + n_2$

Example: Dice Roll

A random variable can be used to describe the process of rolling two (fair) dice and the possible outcomes.

We can consider the probability space $(\Omega_{2D}, \Sigma_{2D}, \mathrm{Pr}_{2D})$ such that:

$$\Omega_{2D} = \{(n_1, n_2) | 1 \leq n_1, n_2 \leq 6\} \qquad \Sigma_{2D} = 2^{\Omega_{2D}} \qquad Pr(A) = \frac{|A|}{36}$$

The total number rolled is then a random variable X given by the function that maps the pair to the sum: $X((n_1, n_2)) = n_1 + n_2$

The pms function p_X and the df F_X function can be defined as:

$$p_X(x) = \begin{cases} \frac{\min(x-1, 13-x)}{36} & 2 \leq x \leq 12 \\ 0 & \text{otherwise} \end{cases} \qquad F_X(x) = \sum_{y \leq x} p_X(y)$$

Mean, or expected value

If X is a discrete random variable with probability mass function $p(\cdot)$, we define the mean or expected value of $X \in S$, $\mu = E[X]$ by

$$E(X) = \sum_{x \in S} x \cdot p(x)$$

Mean, or expected value

If X is a discrete random variable with probability mass function $p(\cdot)$, we define the mean or expected value of $X \in S$, $\mu = E[X]$ by

$$E(X) = \sum_{x \in S} x \cdot p(x)$$

If X is a continuous random variable with density function $f(\cdot) = \frac{dF(\cdot)}{dx}$, we define the mean or expected value of X, $\mu = E[X]$ by

$$\mu = E[X] = \int_{-\infty}^{\infty} x \cdot f(x) dx$$

Variance

The expectation only gives us an idea of the average value assumed by a random variable, not how much individual values may differ from this average.

Variance

The expectation only gives us an idea of the average value assumed by a random variable, not how much individual values may differ from this average.

The variance, $Var(X)$, gives us an indication of the "spread" of values:

$$Var(X) = E\left[(X - E[X])^2\right] = E\left[X^2\right] - E\left[X\right]^2$$

Random variables

A random variable is a function that assigns a numerical quantity to an event in a give probability space.

Random variables

A random variable is a function that assigns a numerical quantity to an event in a give probability space.

Random variables enables the use of standard operations on functions to model randomness of a syste.

Random variables

A random variable is a function that assigns a numerical quantity to an event in a give probability space.

Random variables enables the use of standard operations on functions to model randomness of a syste.

We can focus on the distribution function of a random variable without consider a specific probability space.

Exponential random variables, distribution function

The random variable X is said to be an exponential random variable with parameter λ ($\lambda > 0$) or to have an exponential distribution with parameter λ if it has the distribution function

$$F(x) = \begin{cases} 1 - e^{-\lambda x} & \text{for } x > 0 \\ 0 & \text{for } x \leq 0 \end{cases}$$

Exponential random variables, distribution function

The random variable X is said to be an exponential random variable with parameter λ ($\lambda > 0$) or to have an exponential distribution with parameter λ if it has the distribution function

$$F(x) = \begin{cases} 1 - e^{-\lambda x} & \text{for } x > 0 \\ 0 & \text{for } x \leq 0 \end{cases}$$

Some authors call this distribution the negative exponential distribution.

Exponential random variables, density function

The density function $f = dF/dx$ is given by

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases}$$

Mean, or expected value, of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\mu = E[X] = \int_{-\infty}^{\infty} x\lambda e^{-\lambda x} dx$$

Mean, or expected value, of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\mu = E[X] = \int_{-\infty}^{\infty} x\lambda e^{-\lambda x} dx = \frac{1}{\lambda}$$

Variance of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$Var(X) = E\left[X^2\right] - E[X]^2$$

Variance of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\begin{aligned} Var(X) &= E\left[X^2\right] - E[X]^2 \\ &= \int_0^\infty x^2 \lambda e^{-\lambda x} dx - \left(\frac{1}{\lambda}\right)^2 \end{aligned}$$

Variance of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\begin{aligned} Var(X) &= E\left[X^2\right] - E[X]^2 \\ &= \int_0^\infty x^2 \lambda e^{-\lambda x} dx - \left(\frac{1}{\lambda}\right)^2 \\ &= \frac{2}{\lambda^2} - \frac{1}{\lambda^2} \end{aligned}$$

Variance of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\begin{aligned} Var(X) &= E\left[X^2\right] - E[X]^2 \\ &= \int_0^\infty x^2 \lambda e^{-\lambda x} dx - \left(\frac{1}{\lambda}\right)^2 \\ &= \frac{2}{\lambda^2} - \frac{1}{\lambda^2} \\ &= \frac{1}{\lambda^2} \end{aligned}$$

Exponential inter-event time distribution

The time interval between successive events can also be deduced.

Let $F(t)$ be the distribution function of T, the time between events.
Consider $\Pr(T > t) = 1 - F(t)$:

$$\Pr(T > t) \quad = \quad \Pr(\text{No events in an interval of length } t)$$

Exponential inter-event time distribution

The time interval between successive events can also be deduced.

Let $F(t)$ be the distribution function of T, the time between events. Consider $\Pr(T > t) = 1 - F(t)$:

$$\begin{aligned}\Pr(T > t) &= \Pr(\text{No events in an interval of length } t)\\ &= 1 - F(t)\end{aligned}$$

Exponential inter-event time distribution

The time interval between successive events can also be deduced.

Let $F(t)$ be the distribution function of T, the time between events.
Consider $\Pr(T > t) = 1 - F(t)$:

$$\begin{aligned}
\Pr(T > t) &= \Pr(\text{No events in an interval of length } t) \\
&= 1 - F(t) \\
&= 1 - (1 - e^{-\lambda t})
\end{aligned}$$

Exponential inter-event time distribution

The time interval between successive events can also be deduced.

Let $F(t)$ be the distribution function of T, the time between events. Consider $\Pr(T > t) = 1 - F(t)$:

$$\begin{aligned}
\Pr(T > t) &= \Pr(\text{No events in an interval of length } t) \\
&= 1 - F(t) \\
&= 1 - (1 - e^{-\lambda t}) \\
&= e^{-\lambda t}
\end{aligned}$$

Memoryless property exponential distribution

The exponential distribution is said to have the memoryless property because the time to the next event is independent of when the last event occurred.

Memoryless property exponential distribution

Suppose the last event occurred at time 0.

Memoryless property exponential distribution

Suppose the last event occurred at time 0.

What is the probability that the next event will be after $t + s$, given that time t has elapsed since the last event, and no events have occurred?

Memoryless property exponential distribution

Suppose the last event occurred at time 0.

What is the probability that the next event will be after $t+s$, given that time t has elapsed since the last event, and no events have occurred?

$$\Pr(T > t + s \mid T > t) \quad = \quad \frac{\Pr(T > t + s \text{ and } T > t)}{\Pr(T > t)}$$

Memoryless property exponential distribution

Suppose the last event occurred at time 0.
What is the probability that the next event will be after $t+s$, given that time t has elapsed since the last event, and no events have occurred?

$$\begin{aligned}\Pr(T > t+s \mid T > t) &= \frac{\Pr(T > t+s \text{ and } T > t)}{\Pr(T > t)} \\ &= \frac{\Pr(T > t+s)}{\Pr(T > t)}\end{aligned}$$

Memoryless property exponential distribution

Suppose the last event occurred at time 0.
What is the probability that the next event will be after $t + s$, given that time t has elapsed since the last event, and no events have occurred?

$$
\begin{aligned}
\Pr(T > t + s \mid T > t) &= \frac{\Pr(T > t + s \text{ and } T > t)}{\Pr(T > t)} \\
&= \frac{\Pr(T > t + s)}{\Pr(T > t)} \\
&= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}}
\end{aligned}
$$

Memoryless property exponential distribution

Suppose the last event occurred at time 0.

What is the probability that the next event will be after $t+s$, given that time t has elapsed since the last event, and no events have occurred?

$$
\begin{aligned}
\Pr(T > t+s \mid T > t) &= \frac{\Pr(T > t+s \text{ and } T > t)}{\Pr(T > t)} \\
&= \frac{\Pr(T > t+s)}{\Pr(T > t)} \\
&= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}} \\
&= e^{-\lambda s}
\end{aligned}
$$

Memoryless property exponential distribution

Suppose the last event occurred at time 0.
What is the probability that the next event will be after $t + s$, given that time t has elapsed since the last event, and no events have occurred?

$$
\begin{aligned}
\Pr(T > t + s \mid T > t) &= \frac{\Pr(T > t + s \text{ and } T > t)}{\Pr(T > t)} \\
&= \frac{\Pr(T > t + s)}{\Pr(T > t)} \\
&= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}} \\
&= e^{-\lambda s} = \Pr(T > s)
\end{aligned}
$$

Memoryless property exponential distribution

Suppose the last event occurred at time 0.
What is the probability that the next event will be after $t + s$, given that time t has elapsed since the last event, and no events have occurred?

$$
\begin{aligned}
\Pr(T > t + s \mid T > t) &= \frac{\Pr(T > t + s \text{ and } T > t)}{\Pr(T > t)} \\
&= \frac{\Pr(T > t + s)}{\Pr(T > t)} \\
&= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}} \\
&= e^{-\lambda s} = \Pr(T > s)
\end{aligned}
$$

This value is independent of t (and so the time already spent has not been remembered).

The Poisson distribution

The exponential distribution function is closely related to a discrete random variable, the Poisson distribution.

The Poisson distribution

The exponential distribution function is closely related to a discrete random variable, the Poisson distribution.

This random variable takes values in the set $\{0, 1, 2, \ldots\}$ and has the mass function

$$p_i = e^{-\mu} \frac{\mu^i}{i!} \qquad i \geq 0.$$

The Poisson distribution

The exponential distribution function is closely related to a discrete random variable, the Poisson distribution.

This random variable takes values in the set $\{0, 1, 2, \dots\}$ and has the mass function

$$p_i = e^{-\mu} \frac{\mu^i}{i!} \qquad i \geq 0.$$

The expectation of a Poisson random variable with parameter μ is also μ.

The Poisson random variable

The Poisson random variable is typically used as a counting variable, recording the number of events that occur in a given period of time.

The Poisson random variable

The Poisson random variable is typically used as a counting variable, recording the number of events that occur in a given period of time.

If we observe a Poisson process with parameter μ for some short time period of length h then:

- the probability that one event occurs is $\mu h + o(h)$.
- the probability that two or more events occur is $o(h)$.
- the probability that no events occur is $1 - \mu h + o(h)$.

Poisson vs exponential distributions

If we observe a Poisson process for a infinitesimal time period dt the probability that an event occurs is μdt.

Poisson vs exponential distributions

If we observe a Poisson process for a infinitesimal time period dt the probability that an event occurs is μdt.

If the occurrence of events is governed by a Poisson distribution then the inter-event times are governed by an exponential distribution with the same parameter, and vice versa.

Poisson vs exponential distributions

If we observe a Poisson process for a infinitesimal time period dt the probability that an event occurs is μdt.

If the occurrence of events is governed by a Poisson distribution then the inter-event times are governed by an exponential distribution with the same parameter, and vice versa.

Therefore, if we know that the delay until an event is exponentially distributed then the probability that it will occur in the infinitesimal time interval of length dt, is μdt.

Exponential distributions: properties

Let X and Y two exponentially distributed random variables, with parameters λ_X and λ_Y respectively.

Exponential distributions: properties

Let X and Y two exponentially distributed random variables, with parameters λ_X and λ_Y respectively.

The random variable $Z = min(X, Y)$ is also an exponentially distributed random variable, with parameter $\lambda_X + \lambda_Y$.

Exponential distributions: properties

Let X and Y two exponentially distributed random variables, with parameters λ_X and λ_Y respectively.

The random variable $Z = min(X, Y)$ is also an exponentially distributed random variable, with parameter $\lambda_X + \lambda_Y$.

Consider a stream of events which has events of two types — type A and type B — and assume that the probability that an event has type A is p_A and the probability it has type B is p_B ($p_A + p_B = 1$).

Exponential distributions: properties

Let X and Y two exponentially distributed random variables, with parameters λ_X and λ_Y respectively.

The random variable $Z = min(X, Y)$ is also an exponentially distributed random variable, with parameter $\lambda_X + \lambda_Y$.

Consider a stream of events which has events of two types — type A and type B — and assume that the probability that an event has type A is p_A and the probability it has type B is p_B ($p_A + p_B = 1$).

Then if the inter-event time for any events is exponentially distributed with parameter λ, then the inter-event time for type A events is $p_A \times \lambda$ and similarly for type B events it is $p_B \times \lambda$.

To be continued...