


JENETICS


LIBRARY USER'S MANUAL



Franz Wilhelmstötter
franz.wilhelmstoetter@gmx.at

<http://jenetics.io>

3.1.0—2015/05/04



This work is licensed under a Creative Commons Attribution-ShareAlike 3.0 Austria License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/3.0/at/> or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.

Abstract

Jenetics is an **Genetic Algorithm** and **Evolutionary Algorithm** library, respectively, written in modern day Java. It is designed with a clear separation of the several algorithm concepts, e. g. **Gene**, **Chromosome**, **Genotype**, **Phenotype**, **Population** and fitness **Function**. **Jenetics** allows you to minimize or maximize the given fitness function without tweaking it. In contrast to other GA implementations, the library uses the concept of an evolution *stream* (**EvolutionStream**) for executing the evolution steps. Since the **EvolutionStream** implements the Java **Stream** interface, it works smoothly with the rest of the Java Stream API. This manual describes the concepts implemented in the **Jenetics** project and gives examples and best practice tips.

Contents

1	Introduction	1
2	Architecture	3
3	Base classes	4
3.1	Domain classes	5
3.1.1	Gene	5
3.1.2	Chromosome	6
3.1.3	Genotype	6
3.1.4	Phenotype	7
3.1.5	Population	8
3.2	Operation classes	8
3.2.1	Selector	8
3.2.2	Alterer	12
3.3	Engine classes	15
3.3.1	Fitness function	16
3.3.2	Fitness scaler	17
3.3.3	Engine	17
3.3.4	EvolutionStream	19
3.3.5	EvolutionResult	20
3.3.6	EvolutionStatistics	21
4	Nuts and bolts	22
4.1	Concurrency	22
4.2	Randomness	23
4.3	Serialization	26
4.4	Utility classes	28
5	Extending Jenetics	29
5.1	Genes	29
5.2	Chromosomes	30
5.3	Selectors	30
5.4	Alterers	31
5.5	Statistics	31
5.6	Engine	31
6	Advanced topics	32
6.1	Encoding	32
6.1.1	Real function	33
6.1.2	Scalar function	33
6.1.3	Vector function	34
6.1.4	Affine transformation	35
6.1.5	Graph	36
6.2	Validation	38
6.3	Termination	40
6.3.1	Fixed generation	40
6.3.2	Steady fitness	41
6.3.3	Evolution time	42
6.3.4	Fitness threshold	43

7	Examples	44
7.1	Ones counting	44
7.2	Real function	46
7.3	0/1 Knapsack	48
7.4	Traveling salesman	51
8	Build	53
9	License	55
	References	56

List of Figures

2.1	Evolution engine model	3
2.2	Package structure	4
3.1	Domain model	5
3.2	Genotype structure	7
3.3	Fitness proportional selection	10
3.4	Stochastic-universal selection	12
3.5	Single-point crossover	15
3.6	2-point crossover	15
3.7	3-point crossover	16
3.8	Partially-matched crossover	16
4.1	Block splitting	25
4.2	Leapfrogging	25
4.3	Seq class diagram	28
6.1	Undirected graph and adjacency matrix	37
6.2	Directed graph and adjacency matrix	38
6.3	Weighted graph and adjacency matrix	38
6.4	Fixed generation termination	40
6.5	Steady fitness termination	42
6.6	Execution time termination	43
6.7	Fitness threshold termination	44
7.1	Real function 2D	46

1 Introduction

Jenetics is a library, written in Java¹, which provides an genetic algorithm (GA) implementation. It has no runtime dependencies to other libraries, except the Java 8 runtime. Since the library is available on maven central repository², it can be easily integrated into existing projects. The very clear structuring of the different parts of the GA allows an easy adaption for different problem domains.

This manual is not an introduction or a tutorial for genetic and/or evolutionary algorithms in general. It is assumed that the reader has a knowledge about the structure and the functionality of genetic algorithms. Good introductions to GAs can be found in[8],[5],[7] or[11].

To give a first impression of the library usage, lets start with a simple »Hello World« program. This first example implements the well known bit-counting problem.

```

1 import org.jenetics.BitChromosome;
2 import org.jenetics.BitGene;
3 import org.jenetics.Genotype;
4 import org.jenetics.engine.Engine;
5 import org.jenetics.engine.EvolutionResult;
6 import org.jenetics.util.Factory;
7
8 public class HelloWorld {
9     // 2.) Definition of the fitness function.
10    private static Integer eval(Genotype<BitGene> gt) {
11        return ((BitChromosome)gt.getChromosome()).bitCount();
12    }
13
14    public static void main(String[] args) {
15        // 1.) Define the genotype (factory) suitable
16        //     for the problem.
17        Factory<Genotype<BitGene>> gtf =
18            Genotype.of(BitChromosome.of(10, 0.5));
19
20        // 3.) Create the execution environment.
21        Engine<BitGene, Integer> engine = Engine
22            .builder(HelloWorld::eval, gtf)
23            .build();
24
25        // 4.) Start the execution (evolution) and
26        //     collect the result.
27        Genotype<BitGene> result = engine.stream()
28            .limit(100)
29            .collect(EvolutionResult.toBestGenotype());
30
31        System.out.println("Hello World:\n\t" + result);
32    }
33 }

```

Listing 1: »Hello World« GA

¹The library is build with and depends on Java SE 8: <http://www.oracle.com/technetwork/java/javase/downloads/index.html>

²If you are using Gradle, you can use the following dependency string:
»org.bitbucket.fwilhelm:org.jenetics:3.1.0«.

In contrast to other GA implementations, the library uses the concept of an evolution *stream* (`EvolutionStream`) for executing the evolution steps. Since the `EvolutionStream` implements the Java `Stream` interface, it works smoothly with the rest of the Java Stream API. Now let's have a closer look at listing 1 on the preceding page and discuss this simple program step by step:

1. The probably most challenging part, when setting up a new evolution **Engine**, is to transform the problem domain into a appropriate **Genotype** (factory) representation.³ In our example we want to count the number of *ones* of a **BitChromosome**. Since we are counting only the ones of one chromosome, we are adding only one **BitChromosome** to our **Genotype**. In general, the **Genotype** can be created with 1 to n chromosomes. For detailed description of the genotype's structure have a look at section 3.1.3 on page 6.
2. Once this is done, the fitness function which should be maximized, can be defined. Utilizing the new language features introduced in Java 8, we simply write a private static method, which takes the genotype we defined and calculate it's fitness value. If we want to use the optimized bit-counting method, `bitCount()`, we have to cast the `Chromosome<BitGene>` class to the actual used `BitChromosome` class. Since we know for sure that we created the **Genotype** with a **BitChromosome**, this can be done safely. A reference to the `eval` method is then used as fitness function and passed to the `Engine.build` method.
3. In the third step we are creating the *evolution Engine*, which is responsible for changing, respectively evolving, a given population. The Engine is highly configurable and takes parameters for controlling the evolutionary and the computational environment. For changing the evolutionary behavior, you can set different alterers and selectors (see section 3.2 on page 8). By changing the used **Executor** service, you control the number of threads, the **Engine** is allowed to use. An new **Engine** instance can only be created via its builder, which is created by calling the `Engine.builder` method.
4. In the last step, we can create a new `EvolutionStream` from our **Engine**. The `EvolutionStream` is the model (or view) of the *evolutionary* process. It serves as a »process handle« and also allows you, among other things, to control the termination of the evolution. In our example, we simply truncate the stream after 100 generations. If you don't limit the stream, the `EvolutionStream` will not terminate and run forever. The final result, the best **Genotype** in our example, is then collected with one of the predefined collectors of the `EvolutionResult` class.

As the example shows, **Jenetics** makes heavy use of the `Stream` and `Collector` classes in Java 8. Also the newly introduced lambda expressions and the functional interfaces (SAM types) play an important roll in the library design.

There are many other GA implementations out there and they may slightly differ in the order of the single execution steps. **Jenetics** uses an classical approach. Listing 2 on the following page shows the (imperative) pseudo-code of the **Jenetics** genetic algorithm steps.

³Section 6.1 on page 32 describes some common problem encodings.


```

1 |  $P_0 \leftarrow P_{initial}$ 
2 |  $F(P_0)$ 
3 | while !finished do
4 |    $g \leftarrow g + 1$ 
5 |    $S_g \leftarrow select_S(P_{g-1})$ 
6 |    $O_g \leftarrow select_O(P_{g-1})$ 
7 |    $O_g \leftarrow alter(O_g)$ 
8 |    $P_g \leftarrow filter[g_i \geq g_{max}](S_g) + filter[g_i \geq g_{max}](O_g)$ 
9 |    $F(P_g)$ 

```

Listing 2: Genetic algorithm

Line (1) creates the initial population and line (2) calculates the fitness value of the individuals. The initial population is created implicitly before the first evolution step is performed. Line (4) increases the generation number and line (5) and (6) selects the survivor and the offspring population. The offspring/survivor fraction is determined by the `offspringFraction` property of the `Engine.Builder`. The selected offspring are altered in line (7). The next line combines the survivor population and the altered offspring population—after removing the *died* individuals—to the new population. The steps from line (4) to (9) are repeated until a given termination criterion is fulfilled.

2 Architecture

The basic metaphor of the **Jenetics** library is the *Evolution Stream*, implemented via the Java 8 Stream API. Therefore it is no longer necessary (and advised) to perform the evolution steps in an *imperative* way. An evolution stream is powered by—and bound to—an *Evolution Engine*, which performs the needed *evolution* steps for one generation; the steps are described in the body of the while-loop in listing 2. Once the evolution engine is created, it can be used by multiple evolution streams, which can be safely used in different execution threads. This is possible, because the evolution `Engine` doesn't have any mutable global state. It is practically a stateless function, $f_E : P \rightarrow P$, which maps a start population, P , to an evolved result population. The `Engine` function, f_E , is, of course, *non-deterministic*. Calling it twice with the same start population will lead to different result populations.




Figure 2.1: Evolution engine model

Figure 2.1 illustrates the main *evolution* engine classes, together with its dependencies. Since the `Engine` class itself is immutable, and can't be changed after creation, it is build/configured via a builder. After the `Engine` has been created, it can be used to create an arbitrary number of `EvolutionStreams`. The `EvolutionStream` is used to control the evolutionary process and collect the final result. This is done in the same way as for the normal `java.util.stream.Stream` classes. With the additional `limit(Predicate)` method, it

is possible to truncate the `EvolutionStream` if some termination criteria is fulfilled.




Figure 2.2: Package structure

Diagram 2.2 shows the package structure of the library which consists of the following packages:

org.jenetics This is the base package of the **Jenetics** library and contains all domain classes, like `Gene`, `Chromosome` or `Genotype`. Most of this types are immutable data classes and doesn't implement any behavior. It also contains the `Selector` and `Alterer` interfaces and its implementations. The classes in this package are (almost) sufficient to implement an own GA.

org.jenetics.engine This package contains the actual GA implementation classes, e. g. `Engine`, `EvolutionStream` or `EvolutionResult`. They mainly operate on the domain classes of the `org.jenetics` package.

org.jenetics.stat This package contains additional statistics classes which are not available in the Java core library. Java only includes classes for calculating the sum and the average of a given *numeric* stream (e. g. `DoubleSummaryStatistics`). With the additions in this package it is also possible to calculate the variance, skewness and kurtosis—using the `DoubleMomentStatistics` class. The `EvolutionStatistics` object, which can be calculated for every generation, relies on the classes of this package.

org.jenetics.util This package contains the collection classes (`Seq`, `ISeq` and `MSeq`) which are used in the public interfaces of the `Chromosome` and `Genotype`. It also contains the `RandomRegistry` class, which implements the global PRNG lookup, as well as helper IO classes for serializing `Genotypes` and whole `Populations`.

3 Base classes

This chapter describes the main classes which are needed to setup and run an genetic algorithm with the **Jenetics**⁴ library. They can roughly divided into three types:

⁴The documentation of the whole API is part of the download package or can be viewed online: <http://jenetics.io/javadoc/org.jenetics/index.html>.

Domain classes These classes form the domain model of the evolutionary algorithm and contain the structural classes like **Gene** and **Chromosome**. They are located in the `org.jenetics` package.

Operation classes These classes operate on the domain classes and includes the **Alterer** and **Selector** classes. They are also located in the `org.jenetics` package.

Engine classes These classes implement the actual evolutionary algorithm and reside solely in the `org.jenetics.engine` package.

3.1 Domain classes

Most of the domain classes are pure data classes and can be treated as *value* objects⁵. All **Gene** and **Chromosome** implementations are immutable as well as the **Genotype** and **Phenotype** class. The only exception is the **Population** class, where it is possible to add and/or remove elements after its creation.




Figure 3.1: Domain model

Figure 3.1 shows the class diagram of the domain classes. All domain classes are located in the `org.jenetics` package. The **Gene** is the base of the class structure. **Genes** are aggregated in **Chromosomes**. One to *n* **Chromosomes** are aggregated in **Genotypes**. A **Genotype** and a fitness **Function** form the **Phenotype**, which are collected into a **Population**.

3.1.1 Gene

Genes are the basic building blocks of the **Jenetics** library. They contain the actual information of the encoded solution, the allele. Some of the implementations also contain domain information of the *wrapped* allele. This is the case for all **BoundedGene**, which contain the allowed minimum and maximum values. All **Gene** implementations are final and immutable. In fact, they are all value-based classes and fulfill the properties which are described in the Java 8 API documentation[6].⁶

Beside the container functionality for the allele, every **Gene** is its own factory and is able to create new, random instances of the same type and with the same

⁵https://en.wikipedia.org/wiki/Value_object

⁶It is also worth reading the blog entry from Stephen Colebourne: <http://blog.joda.org/2014/03/valjos-value-java-objects.html>

constraints. The factory methods are used by the **Alterers** for creating new **Genes** from the existing one and play a crucial role by the exploration of the problem space.

```

1 public interface Gene<A, G extends Gene<A, G>>
2     extends Factory<G>, Verifiable
3 {
4     public A getAllele();
5     public G newInstance();
6     public G newInstance(A allele);
7     public boolean isValid();
8 }

```

Listing 3: Gene interface

Listing 3 shows the most important methods of the **Gene** interface. The **isValid** method, introduced by the **Verifiable** interface, allows the gene to mark itself as invalid. All invalid genes are replaced with new ones during the evolution phase.

The available **Gene** implementations in the **Jenetics** library should cover a wide range of problem encodings. Refer to chapter 5.1 on page 29 for how to implement your own **Gene** types.

3.1.2 Chromosome

A **Chromosome** is a collection of **Genes** which contains at least one **Gene**. This allows to encode problems which requires more than one **Gene**. Like the **Gene** interface, the **Chromosome** is also its own factory and allows to create a new **Chromosome** from a given **Gene** sequence.

```

1 public interface Chromosome<G extends Gene<?, G>>
2     extends Factory<Chromosome<G>>, Iterable<G>, Verifiable
3 {
4     public Chromosome<G> newInstance(ISeq<G> genes);
5     public G getGene(int index);
6     public ISeq<G> toSeq();
7     public int length();
8 }

```

Listing 4: Chromosome interface

Listing 4 shows the main methods of the **Chromosome** interface. This are the methods for accessing single **Genes** by index and as an **ISeq** respectively, and the factory method for creating a new **Chromosome** from a given sequence of **Genes**. The factory method is used by the **Alterer** classes which were able to create altered **Chromosome** from a (changed) **Gene** sequence.

3.1.3 Genotype

The central class, the evolution **Engine** is working with, is the **Genotype**. It is the *structural* and immutable representative of an individual and consists of one to n **Chromosomes**. All **Chromosomes** must be parameterized with the same **Gene** type, but it is allowed to have different lengths and constraints.

Figure 3.2 on the following page shows the **Genotype** structure. A **Genotype** consists of N_G **Chromosomes** and a **Chromosome** consists of $N_{C[i]}$ **Genes** (depending on the **Chromosome**). The overall number of **Genes** of a **Genotype** is




Figure 3.2: Genotype structure

given by the sum of the **Chromosome's Genes**, which can be accessed via the `Genotype.getNumberOfGenes()` method:

$$N_g = \sum_{i=0}^{N_C-1} N_{C[i]} \quad (3.1)$$

As already mentioned, the **Chromosomes** of a **Genotype** doesn't have to have necessarily the same size. It is only required that all genes are from the same type and the **Genes** within a **Chromosome** have the same constraints; e. g. the same min- and max values for numerical **Genes**.

```

1 | Genotype<DoubleGene> genotype = Genotype.of(
2 |     DoubleChromosome.of(0.0, 1.0, 8),
3 |     DoubleChromosome.of(1.0, 2.0, 10),
4 |     DoubleChromosome.of(0.0, 10.0, 9),
5 |     DoubleChromosome.of(0.1, 0.9, 5)
6 | );

```

The code snippet in the listing above creates a **Genotype** with the same structure as shown in figure 3.2. In this example the **DoubleGene** has been chosen as **Gene** type.

3.1.4 Phenotype

The **Phenotype** is the *actual* representative of an individual and consists of the **Genotype** and the **fitness Function**, which is used to (lazily) calculate the **Genotype's** fitness value.⁷ It is *only* a container which forms the *environment* of the **Genotype** and doesn't change the structure. Like the **Genotype**, the **Phenotype** is immutable and can't be changed after creation.

```

1 | public final class Phenotype<
2 |     G extends Gene<?, G>,
3 |     C extends Comparable<? super C>
4 | >
5 |     implements Comparable<Phenotype<G, C>>
6 | {

```

⁷Since the fitness **Function** is shared by all **Phenotypes**, calls to the fitness **Function** must be idempotent. See section 3.3.1 on page 16.

```

7   public C getFitness ();
8   public Genotype<G> getGenotype ();
9   public long getAge (long currentGeneration);
10  public void evaluate ();
11 }

```

Listing 5: Phenotype class

Listing 5 on the previous page shows the main methods of the `Phenotype`. The `fitness` property will return the actual fitness value of the `Genotype`, which can be fetched with the `getGenotype` method. To make the runtime behavior predictable, the fitness value is evaluated lazily. Either by querying the `fitness` property or through the call of the `evaluate` method. The `EvolutionEngine` is calling the `evaluate` method in a separate step and makes the fitness evaluation time available through the `EvolutionDurations` class. Additionally to the fitness value, the `Phenotype` contains the generation when it was created. This allows to calculate the current age and the removal of overaged individuals from the `Population`.

3.1.5 Population

The end of the class hierarchy of the domain model is the `Population`. It is a collection of individuals and forms the start and the end of an evolution step.

```

1  public final class Population<
2     G extends Gene<?, G>,
3     C extends Comparable<? super C>
4 >
5     implements List<Phenotype<G, C>>
6 {
7     public Phenotype<G, C> get (int index);
8     public void add (Phenotype<G, C> phenotype);
9     public void sortWith (Comparator<? super C> comparator);
10 }

```

Listing 6: Population class

Listing 6 gives an overview of the most important methods of the `Population` class. In addition to the `List` methods, it provides a method for sorting the `Phenotypes`. Some `Selector` implementations require a sorted list of individuals according to its fitness value. Calling `population.sortWith(optimization.descending())` will sort the `Population`, so that the first element will be the individual with the best fitness.

3.2 Operation classes

Genetic operators are used for creating *genetic* diversity (`Alterer`) and selecting potentially useful solutions for recombination (`Selector`). This section gives an overview about the genetic operators available in the `Jenetics` library. It also contains some *theoretical* information, which should help you to choose the right combination of operators and parameters, for the problem to be solved.

3.2.1 Selector

Selectors are responsible for selecting a given number of individuals from the population. The selectors are used to divide the population into *survivors* and

offspring. The selectors for offspring and for the survivors can be chosen independently.

The selection process of the **Jenetics** library acts on Phenotypes and indirectly, via the fitness function, on Genotypes. Direct Gene- or Population selection is not supported by the library.

```

1 Engine<DoubleGene, Double> engine = Engine.builder(...)
2   .offspringFraction(0.7)
3   .survivorsSelector(new RouletteWheelSelector<>())
4   .offspringSelector(new TournamentSelector<>())
5   .build();

```

The `offspringFraction`, $f_O \in [0, 1]$, determines the number of selected offspring

$$N_{O_g} = \|O_g\| = \text{rint}(\|P_g\| \cdot f_O) \quad (3.2)$$

and the number of selected survivors

$$N_{S_g} = \|S_g\| = \|P_g\| - \|O_g\|. \quad (3.3)$$

The **Jenetics** library contains the following selector implementations:

- TournamentSelector
- TruncationSelector
- MonteCarloSelector
- ProbabilitySelector
- RouletteWheelSelector
- LinearRankSelector
- ExponentialRankSelector
- BoltzmannSelector
- StochasticUniversalSelector

Beside the well known standard selector implementation the **ProbabilitySelector** is the base of a set of fitness proportional selectors.

Tournament selector In tournament selection the best individual from a random sample of s individuals is chosen from the population P_g . The samples are drawn with replacement. An individual will win a tournament only if the fitness is greater than the fitness of the other $s - 1$ competitors. Note that the worst individual never survives, and the best individual wins in all the tournaments it participates. The selection pressure can be varied by changing the tournament sizes. For large values of s , weak individuals have less chance of being selected.

Truncation selector In truncation selection individuals are sorted according to their fitness. (This is one of the selectors, which relies on the `sortWith` method of the **Population** class.) Only the n best individuals are selected. The truncation selection is a very basic selection algorithm. It has its strength in fast selecting individuals in large populations, but is not very often used in practice.

Monte Carlo selector The Monte Carlo selector selects the individuals from a given population randomly. This selector can be used to measure the performance of a other selectors. In general, the performance of a selector should be better than the selection performance of the Monte Carlo selector.

Probability selectors Probability selectors are a variation of *fitness proportional* selectors and selects individuals from a given population based on it's *selection probability* $P(i)$. Fitness proportional selection works as shown in




Figure 3.3: Fitness proportional selection

figure 3.3. An uniform distributed random number $r \in [0, F)$ specifies which individual is selected, by argument minimization:

$$i \leftarrow \underset{n \in [0, N)}{\operatorname{argmin}} \left\{ r < \sum_{i=0}^n f_i \right\}, \quad (3.4)$$

where N is the number of individuals and f_i the fitness value of the i^{th} individual. The probability selector works the same way, only the fitness value f_i is replaced by the individual's selection probability $P(i)$. It is not necessary to sort the population. The selection probability of an individual i follows a binomial distribution

$$P(i, k) = \binom{n}{k} P(i)^k (1 - P(i))^{n-k} \quad (3.5)$$

where n is the overall number of selected individuals and k the number of individual i in the set of selected individuals. The runtime complexity of the implemented probability selectors is $O(n + \log(n))$ instead of $O(n^2)$ as for the naive approach: *A binary (index) search is performed on the summed probability array.*

Roulette-wheel selector The roulette-wheel selector is also known as fitness proportional selector. In the **Jenetics** library it is implemented as *probability selector*. The fitness value f_i is used to calculate the selection probability of individual i .

$$P(i) = \frac{f_i}{\sum_{j=1}^N f_j} \quad (3.6)$$

Selecting n individuals from a given population is equivalent to play n times on the roulette-wheel. The population don't have to be sorted before selecting the individuals. Roulette-wheel selection is one of the traditional selection strategies.

Linear-rank selector In linear-ranking selection the individuals are sorted according to their fitness values. The rank N is assigned to the best individual and the rank 1 to the worst individual. The selection probability $P(i)$ of individual i is linearly assigned to the individuals according to their rank.

$$P(i) = \frac{1}{N} \left(n^- + (n^+ - n^-) \frac{i-1}{N-1} \right). \quad (3.7)$$

Here $\frac{n^-}{N}$ is the probability of the worst individual to be selected and $\frac{n^+}{N}$ the probability of the best individual to be selected. As the population size is held constant, the condition $n^+ = 2 - n^-$ and $n^- \geq 0$ must be fulfilled. Note that all individuals get a different rank, respectively a different selection probability, even if they have the same fitness value.[4]

Exponential-rank selector An alternative to the *weak* linear-rank selector is to assign survival probabilities to the sorted individuals using an exponential function:

$$P(i) = (c-1) \frac{c^{i-1}}{c^N - 1}, \quad (3.8)$$

where c must within the range $[0 \dots 1)$. A small value of c increases the probability of the best individual to be selected. If c is set to zero, the selection probability of the best individual is set to one. The selection probability of all other individuals is zero. A value near one equalizes the selection probabilities. This selector sorts the population in descending order before calculating the selection probabilities.

Boltzmann selector The selection probability of the Boltzmann selector is defined as

$$P(i) = \frac{e^{b \cdot f_i}}{Z}, \quad (3.9)$$

where b is a parameter which controls the selection intensity and Z is defined as

$$Z = \sum_{i=1}^n e^{f_i}. \quad (3.10)$$

Positive values of b increases the selection probability of individuals with high fitness values and negative values of b decreases it. If b is zero, the selection probability of all individuals is set to $\frac{1}{N}$.

Stochastic-universal selector Stochastic-universal selection [1] (SUS) is a method for selecting individuals according to some given probability in a way that minimizes the chance of fluctuations. It can be viewed as a type of roulette game where we now have p equally spaced points which we spin. SUS uses a single random value for selecting individuals by choosing them at equally spaced




Figure 3.4: Stochastic-universal selection

intervals. The selection method was introduced by James Baker. [2] Figure 3.4 shows the function of the stochastic-universal selection, where n is the number of individuals to select. Stochastic universal sampling ensures a selection of offspring, which is closer to what is deserved than roulette wheel selection.[8]

3.2.2 Alterer

The problem encoding (representation) determines the bounds of the search space, but the **Alterers** determine how the space can be traversed: **Alterers** are responsible for the genetic diversity of the **EvolutionStream**. The two **Alterer** types used in **Jenetics** are

1. mutation and
2. recombination (e. g. crossover).

First we will have a look at the mutation — There are two distinct roles *mutation* plays in the evolution process:

1. **Exploring the search space:** By making small moves, mutation allows a population to explore the search space. This exploration is often slow compared to crossover, but in problems where crossover is disruptive this can be an important way to explore the landscape.
2. **Maintaining diversity:** Mutation prevents a population from correlating. Even if most of the search is being performed by crossover, mutation can be vital to provide the diversity which crossover needs.

The mutation probability, $P(m)$, is the parameter that must be optimized. The optimal value of the mutation rate depends on the role mutation plays. If mutation is the only source of exploration (if there is no crossover), the mutation rate should be set to a value that ensures that a reasonable neighborhood of solutions is explored.

The mutation probability, $P(m)$, is defined as the probability that a specific gene, over the whole population, is mutated. That means, the (average) number of genes mutated by a mutator is

$$\hat{\mu} = N_P \cdot N_g \cdot P(m) \quad (3.11)$$

where N_g is the number of available genes of a genotype and N_P the population size (revere to equation 3.1 on page 7).

Mutator The mutator has to deal with the problem, that the genes are arranged in a 3D structure (see chapter 3.1.3). The mutator selects the gene which will be mutated in three steps:

1. Select a genotype $G[i]$ from the population with probability $P_G(m)$,
2. select a chromosome $C[j]$ from the selected genotype $G[i]$ with probability $P_C(m)$ and
3. select a gene $g[k]$ from the selected chromosome $C[j]$ with probability $P_g(m)$.

The needed *sub*-selection probabilities are set to

$$P_G(m) = P_C(m) = P_g(m) = \sqrt[3]{P(m)}. \quad (3.12)$$

Gaussian mutator The Gaussian mutator performs the mutation of number genes. This mutator picks a new value based on a Gaussian distribution around the current value of the gene. The variance of the new value (before clipping to the allowed gene range) will be

$$\hat{\sigma}^2 = \left(\frac{g_{max} - g_{min}}{4} \right)^2 \quad (3.13)$$

where g_{min} and g_{max} are the valid minimum and maximum values of the number gene. The new value will be cropped to the gene's boundaries.

Swap mutator The swap mutator changes the order of genes in a chromosome, with the hope of bringing related genes closer together, thereby facilitating the production of building blocks. This mutation operator can also be used for combinatorial problems, where no duplicated genes within a chromosome are allowed, e. g. for the TSP.

The second alterer type is the recombination — An enhanced genetic algorithm (EGA) combine elements of existing solutions in order to create a new solution, with some of the properties of each parents. Recombination creates a new chromosome by combining parts of two (or more) parent chromosomes. This combination of chromosomes can be made by selecting one or more crossover points, splitting these chromosomes on the selected points, and merge those portions of different chromosomes to form new ones.

```

1 void recombine(final Population<G, C> pop) {
2     // Select the Genotypes for crossover.
3     final Random random = RandomRegistry.getRandom();
4     final int i1 = random.nextInt(pop.length());
5     final int i2 = random.nextInt(pop.length());
6     final Phenotype<G, C> pt1 = pop.get(i1);
7     final Phenotype<G, C> pt2 = pop.get(i2);
8     final Genotype<G> gt1 = pt1.getGenotype();
9     final Genotype<G> gt2 = pt2.getGenotype();
10
11     //Choosing the Chromosome for crossover.
12     final int chIndex = random.nextInt(gt1.length());

```

```

13  final MSeq<Chromosome<G>> c1 = gt1.toSeq().copy();
14  final MSeq<Chromosome<G>> c2 = gt2.toSeq().copy();
15  final MSeq<G> genes1 = c1.get(chIndex).toSeq().copy();
16  final MSeq<G> genes2 = c2.get(chIndex).toSeq().copy();
17
18  // Perform the crossover.
19  crossover(genes1, genes2);
20  c1.set(chIndex, c1.get(chIndex).newInstance(genes1.toISeq()));
21  c2.set(chIndex, c2.get(chIndex).newInstance(genes2.toISeq()));
22
23  //Creating two new Phenotypes and replace the old one.
24  pop.set(i1, pt1.newInstance(gt1.newInstance(c1.toISeq())));
25  pop.set(i2, pt2.newInstance(gt1.newInstance(c2.toISeq())));
26 }

```

Listing 7: Chromosome selection for recombination

Listing 7 on the preceding page shows how two chromosomes are selected for *recombination*. It is done this way for preserving the given *constraints* and to avoid the creation of invalid individuals.

Because of the possible different Chromosome length and/or Chromosome constraints within a Genotype, only Chromosomes with the same Genotype position are recombined (see listing 7 on the previous page).

The recombination probability, $P(r)$, determines the probability that a given individual (genotype) of a population is selected for recombination. The (mean) number of changed individuals depend on the concrete implementation and can be vary from $P(r) \cdot N_G$ to $P(r) \cdot N_G \cdot O_R$, where O_R is the order of the recombination, which is the number of individuals involved in the `combine` method.

Single-point crossover The single-point crossover changes two children chromosomes by taking two chromosomes and cutting them at some, randomly chosen, site. If we create a child and its complement we preserve the total number of genes in the population, preventing any genetic drift. Single-point crossover is the classic form of crossover. However, it produces very slow mixing compared with multi-point crossover or uniform crossover. For problems where the site position has some intrinsic meaning to the problem single-point crossover can lead to smaller disruption than multiple-point or uniform crossover.

Figure 3.5 shows how the `SinglePointCrossover` class is performing the crossover for different crossover points—in the given example for the chromosome indexes 0, 1, 3, 6 and 7.

Multi-point crossover If the `MultiPointCrossover` class is created with one crossover point, it behaves exactly like the single-point crossover. The following picture shows how the Multi-point crossover works with two crossover points, defined at index 1 and 4.

Figure 3.7 you can see how the crossover works for an odd number of crossover points.

Partially-matched crossover The partially-matched crossover guarantees that all genes are found exactly once in each chromosome. No gene is dupli-




Figure 3.5: Single-point crossover




Figure 3.6: 2-point crossover

cated by this crossover strategy. The partially-matched crossover (PMX) can be applied usefully in the TSP or other permutation problem encodings. Permutation encoding is useful for all problems where the fitness only depends on the ordering of the genes within the chromosome. This is the case in many combinatorial optimization problems. Other crossover operators for combinatorial optimization are:

- order crossover
- edge recombination crossover
- cycle crossover
- edge assembly crossover

The PMX is similar to the two-point crossover. A crossing region is chosen by selecting two crossing points (see figure 3.8 *a*). After performing the crossover we normally got two invalid chromosomes (figure 3.8 *b*). Chromosome 1 contains the value 6 twice and misses the value 3. On the other side chromosome 2 contains the value 3 twice and misses the value 6. We can observe that this crossover is equivalent to the exchange of the values 3→6, 4→5 and 5→4. To repair the two chromosomes we have to apply this exchange outside the crossing region (figure 3.8 *b*). At the end figure 3.8 *c*) shows the repaired chromosome.

3.3 Engine classes

The *executing* classes, which perform the actual evolution, are located in the `org.jenetics.engine` package. The *evolution stream* (`EvolutionStream`) is the base metaphor for performing an GA. On the `EvolutionStream` you can define the termination predicate and than collect the final `EvolutionResult`.




Figure 3.7: 3-point crossover




Figure 3.8: Partially-matched crossover

This decouples the static data structure from the executing evolution part. The `EvolutionStream` is also very flexible, when it comes to collecting the final result. The `EvolutionResult` class has several predefined collectors, but you are free to create your own one, which can be seamlessly *plugged* into the existing stream.

3.3.1 Fitness function

The fitness `Function` is also an important part when modeling an genetic algorithm. It takes a `Genotype` as argument and returns, at least, a `Comparable` object as result—the fitness value. This allows the evolution `Engine`, respectively the selection operators, to select the offspring- and survivor `Population`. Some selectors have stronger requirements to the fitness value than a `Comparable`, but this constraints is checked by the Java type system at compile time.

Since the fitness `Function` is shared by all `Phenotypes`, calls to the fitness `Function` has to be idempotent. A fitness `Function` is idempotent if, whenever it is applied twice to any `Genotype`, it returns the same fitness value as if it were applied once. In the simplest case, this is achieved by `Functions` which doesn't contain any global *mutable* state.

The following example shows the simplest possible fitness `Function`. This `Function` simply returns the allele of a 1x1 *float* `Genotype`.

```

1 public class Main {
2     static Double identity( final Genotype<DoubleGene> gt ) {

```

```

3     return gt.getGene().getAllele();
4   }
5
6   public static void main(final String[] args) {
7       // Create fitness function from method reference.
8       Function<Genotype<DoubleGene>, Double>> ff1 =
9           Main::identity;
10
11      // Create fitness function from lambda expression.
12      Function<Genotype<DoubleGene>, Double>> ff2 = gt ->
13          gt.getGene().getAllele();
14  }
15 }

```

The first type parameter of the `Function` defines the kind of `Genotype` from which the fitness value is calculated and the second type parameter determines the return type, which must be, at least, a `Comparable` type.

3.3.2 Fitness scaler

The fitness value, calculated by the fitness `Function`, is treated as the *raw*-fitness of an individual. The `Jenetics` library allows you to apply an additional scaling function on the raw-fitness to form the fitness value which is used by the selectors. This can be useful when using probability selectors (see chapter 3.2.1 on page 10), where the actual amount of the fitness value influences the selection probability. In such cases, the fitness scaler gives you additional flexibility when selecting offspring and survivors. In the default configuration the raw-fitness is equal to the actual fitness value, that means, the used fitness scaler is the identity function.

```

1 class Main {
2     public static void main(final String[] args) {
3         Engine<DoubleGene, Double> engine = Engine.builder(...)
4             .fitnessScaler(Math::sqrt)
5             .build();
6     }
7 }

```

The given listing shows a fitness scaler which reduces the the raw-fitness to its square root. This gives weaker individuals a greater changes being selected and weakens the influence of *super*-individuals.

When using a fitness scaler you have to take care that your scaler doesn't *destroy* your fitness value. This can be the case when your fitness value is negative and your fitness scaler squares the value. Trying to find the minimum will not work in this configuration.

3.3.3 Engine

The *evolution* `Engine` controls how the evolution steps are executed. Once the `Engine` is created, via a `Builder` class, it can't be changed. It doesn't contain any mutable global state and can therefore safely used/called from different threads. This allows to create more than one `EvolutionStreams` from the `Engine` and execute them in parallel.

```

1 public final class Engine<
2     G extends Gene<?, G>,
3     C extends Comparable<? super C>
4 >
5 {
6     // The evolution function, performs one evolution step.
7     EvolutionResult<G,C> evolve(
8         Population<G,C> population,
9         long generation
10    );
11
12    // Evolution stream for "normal" evolution execution.
13    EvolutionStream<G,C> stream();
14
15    // Evolution iterator for external evolution iteration.
16    Iterator<EvolutionResult<G,C>> iterator();
17 }

```

Listing 8: Engine class

Listing 8 shows the main methods of the `Engine` class. It is used for performing the actual evolution of a give population. One evolution step is executed by calling the `Engine.evolve` method, which returns an `EvolutionResult` object. This object contains the evolved `Population` plus additional information like execution duration of the several evolution sub-steps and information about the killed and as invalid marked individuals. With the `stream` method you create a new `EvolutionStream`, which is used for controlling the evolution process—see section 3.3.4 on the following page. Alternatively it is possible to iterate through the evolution process in an imperative way (for whatever reasons this should be necessary). Just create an `Iterator` of `EvolutionResult` object by calling the `iterator` method.

As already shown in previous examples, the `Engine` can only be created via its `Builder` class. Only the fitness `Function` and the `Chromosomes`, which represents the problem encoding, must be specified for creating an `Engine` instance. For the rest of the parameters default values are specified. This are the `Engine` parameters which can configured:

alterers A list of `Alterers` which are applied to the offspring `Population`, in the defined order. The default value of this property is set to `SinglePointCrossover<>(0.2)` followed by `Mutator<>(0.15)`.

clock The `java.time.Clock` used for calculating the execution durations. A `Clock` with nanosecond precision (`System.nanoTime()`) is used as default.

executor With this property it is possible to change the `java.util.concurrent.Executor` engine used for evaluating the evolution steps. This property can be used to define an application wide `Executor` or for controlling the number of execution threads. The default value is set to `ForkJoinPool.commonPool()`.

fitnessFunction This property defines the fitness `Function` used by the evolution `Engine`. (See section 3.3.1 on page 16.)

fitnessScaler This property defines the fitness scaler used by the evolution `Engine`. The default value is set to the identity function. (See section 3.3.2 on the preceding page.)

- genotypeFactory** Defines the **Genotype Factory** used for creating new individuals. Since the **Genotype** is its own **Factory**, it is sufficient to create a **Genotype**, which then serves as template.
- genotypeValidator** This property lets you *override* the default implementation of the **Genotype.isValid** method, which is useful if the **Genotype** validity not only depends on valid property of the elements it consists of.
- maximalPhenotypeAge** Set the maximal allowed age of an individual (**Phenotype**). This prevents *super* individuals to live *forever*. The default value is set to 70.
- offspringFraction** Through this property it is possible to define the fraction of offspring (and survivors) for evaluating the next generation. The fraction value must within the interval [0, 1]. The default value is set to 0.6.
- offspringSelector** This property defines the **Selector** used for selecting the offspring **Population**. The default values is set to **TournamentSelector**<>(3).
- optimize** With this property it is possible to define whether the fitness **Function** should be maximized of minimized. By default, the fitness **Function** is maximized.
- phenotypeValidator** This property lets you *override* the default implementation of the **Phenotype.isValid** method, which is useful if the **Phenotype** validity not only depends on valid property of the elements it consists of.
- populationSize** Defines the number of individuals of a **Population**. The evolution **Engine** keeps the number of individuals constant. That means, the **Population** of the **EvolutionResult** always contains the number of entries defined by this property. The default value is set to 50.
- survivorsSelector** This property defines the **Selector** used for selecting the survivors **Population**. The default values is set to **TournamentSelector**<>(3).
- individualCreationRetries** The evolution **Engine** tries to create only valid individuals. If a newly created **Genotype** is not valid, the **Engine** creates another one, till the created **Genotype** is valid. This parameter sets the maximal number of retries before the **Engine** gives up and accept invalid individuals. The default value is set to 10.

3.3.4 EvolutionStream

The **EvolutionStream** controls the execution of the evolution process and can be seen as a kind of execution *handle*. This handle can be used to define the termination criteria and to *collect* the final evolution result. Since the **EvolutionStream** extends the Java **Stream** interface, it integrates smoothly with the rest of the Java Stream API.⁸

⁸It is recommended to make yourself familiar with the Java Stream API. A good introduction can be found here: <http://winterbe.com/posts/2014/07/31/java8-stream-tutorial-examples/>

```

1 public interface EvolutionStream<
2   G extends Gene<?, G>,
3   C extends Comparable<? super C>
4 >
5   extends Stream<EvolutionResult<G, C>>
6 {
7   public EvolutionStream<G, C>
8     limit(Predicate<? super EvolutionResult<G, C>> proceed);
9 }

```

Listing 9: EvolutionStream class

Listing 9 shows the whole `EvolutionStream` interface. As it can be seen, it only adds one additional method. But this additional `limit` method allows to truncate the `EvolutionStream` based on a `Predicate` which takes an `EvolutionResult`. Once the `Predicate` returns `false`, the evolution process is stopped. Since the `limit` method returns an `EvolutionStream`, it is possible to define more than one `Predicate`, which both must be fulfilled to continue the evolution process.

```

1 Engine<DoubleGene, Double> engine = ...
2 EvolutionStream<DoubleGene, Double> stream = engine.stream()
3   .limit(predicate1)
4   .limit(predicate2)
5   .limit(100);

```

The `EvolutionStream`, created in the example above, will be truncated if one of the two predicates is `false` or if the maximal allowed generations, of 100, is reached. An `EvolutionStream` is usually created via the `Engine.stream()` method. The *immutable* and *stateless* nature of the evolution `Engine` allows to create more than one `EvolutionStream` with the same `Engine` instance.

In cases where you appreciate the usage of the `EvolutionStream` but need a different `Engine` implementation, you can use the `EvolutionStream.of` factory method for creating a new `EvolutionStream`.

```

1 static <G extends Gene<?, G>, C extends Comparable<? super C>>
2 EvolutionStream<G, C> of(
3   Supplier<EvolutionStart<G, C>> start,
4   Function<? super EvolutionStart<G, C>, EvolutionResult<G, C>> f
5 );

```

This factory method takes a start value, of type `EvolutionStart`, and an evolution `Function`. The evolution `Function` takes the start value and returns an `EvolutionResult` object. To make the runtime behavior more predictable, the start value is fetched/created lazily at the evolution start time.

```

1 final Supplier<EvolutionStart<DoubleGene, Double>> start = ...
2 final EvolutionStream<DoubleGene, Double> stream =
3   EvolutionStream.of(start, new MySpecialEngine());

```

3.3.5 EvolutionResult

The `EvolutionResult` collects the result data of an evolution step into an immutable *value* class. This class is the type of the stream elements of the `EvolutionStream`, as described in section 3.3.4 on the preceding page.

```

1 public final class EvolutionResult<
2   G extends Gene<?, G>,
3   C extends Comparable<? super C>

```

```

4 >
5   implements Comparable<EvolutionResult<G, C>>
6 {
7     Population<G,C> getPopulation();
8     long getGeneration();
9 }

```

Listing 10: EvolutionResult class

Listing 3.3.5 on the previous page shows the two most important properties, the **population** and the **generation** the result belongs to. These are also the two properties needed for the next evolution step. The **generation** is, of course, incremented by one. To make collecting the **EvolutionResult** object easier, it also implements the **Comparable** interface. Two **EvolutionResults** are compared by its best **Phenotype**.

The **EvolutionResult** class has three predefined factory methods, which will return **Collectors** usable for the **EvolutionStream**:

toBestEvolutionResult() Collects the best **EvolutionResult** of an **EvolutionStream** according to the defined optimization strategy.

toBestPhenotype() This collector can be used if you are only interested in the best **Phenotype**.

toBestGenotype() Use this collector if you only need the best **Genotype** of the **EvolutionStream**.

The following code snippets show how to use the different **EvolutionStream** collectors.

```

1 // Collecting the best EvolutionResult of the EvolutionStream.
2 EvolutionResult<DoubleGene, Double> result = stream
3     .collect(EvolutionResult.toBestEvolutionResult());
4
5 // Collecting the best Phenotype of the EvolutionStream.
6 Phenotype<DoubleGene, Double> result = stream
7     .collect(EvolutionResult.toBestPhenotype());
8
9 // Collecting the best Genotype of the EvolutionStream.
10 Genotype<DoubleGene> result = stream
11     .collect(EvolutionResult.toBestGenotype());

```

3.3.6 EvolutionStatistics

The **EvolutionStatistics** class allows you to gather additional statistical information from the **EvolutionStream**. This is especially useful during the development phase of the application, when you have to find the right parametrization of the evolution **Engine**. Besides other information, the **EvolutionStatistics** contains (statistical) information about the fitness, invalid and killed **Phenotypes** and runtime information of the different evolution steps. Since the **EvolutionStatistics** class implements the **Consumer<EvolutionResult<?, C>>** interface, it can be easily plugged into the **EvolutionStream**, adding it with the **peek** method of the stream.

```

1 Engine<DoubleGene, Double> engine = ...
2 EvolutionStatistics<?, Double> statistics =
3     EvolutionStatistics.ofNumber();

```

```

4 engine.stream()
5   .limit(100)
6   .peek(statistics)
7   .collect(toBestGenotype());

```

Listing 11: EvolutionStatistics usage

Listing 11 on the preceding page shows how to add the the `EvolutionStatistics` to the `EvolutionStream`. Once the algorithm tuning is finished, it can be removed in the production environment.

There are two different specializations of the `EvolutionStatistics` object available. The first is the general one, which will be working for every kind of `Genes` and fitness types. It can be created via the `EvolutionStatistics.ofComparable()` method. The second one collects additional statistical data for numeric fitness values. This can be created with the `EvolutionStatistics.ofNumber()` method.

```

1  +-----+
2  | Time statistics |
3  +-----+
4  |           Selection: sum=0.046538278000 s; mean=0.003878189833 s |
5  |           Altering: sum=0.086155457000 s; mean=0.007179621417 s |
6  | Fitness calculation: sum=0.022901606000 s; mean=0.001908467167 s |
7  | Overall execution: sum=0.147298067000 s; mean=0.012274838917 s |
8  +-----+
9  | Evolution statistics |
10 +-----+
11 |           Generations: 12 |
12 |           Altered: sum=7,331; mean=610.916666667 |
13 |           Killed: sum=0; mean=0.000000000 |
14 |           Invalids: sum=0; mean=0.000000000 |
15 +-----+
16 | Population statistics |
17 +-----+
18 |           Age: max=11; mean=1.951000; var=5.545190 |
19 |           Fitness: |
20 |               min = 0.000000000000 |
21 |               max = 481.748227114537 |
22 |               mean = 384.430345078660 |
23 |               var = 13006.132537301528 |
24 +-----+

```

A typical output of an number `EvolutionStatistics` object will look like the example above.

4 Nuts and bolts

4.1 Concurrency

The **Genetics** library parallelizes independent task whenever possible. Especially the evaluation of the fitness function is done concurrently. That means that the fitness function must be thread safe, because it is shared by all phenotypes of a population. The easiest way for achieving thread safety is to make the fitness function immutable and re-entrant. The used `Executor` can be defined when building the evolution `Engine` object.

```

1 import java.util.concurrent.Executor;
2 import java.util.concurrent.Executors;
3
4 public class Main {
5     private static Double eval(final Genotype<DoubleGene> gt) {
6         // calculate and return fitness

```

```

7     }
8
9     public static void main(final String[] args) {
10        // Creating an fixed size ExecutorService
11        final ExecutorService executor = Executors
12            .newFixedThreadPool(10)
13        final Factory<Genotype<DoubleGene>> gtf = ...
14        final Engine<DoubleGene, Double> engine = Engine
15            .builder(Main::eval, gtf)
16            // Using 10 threads for evolving.
17            .executor(executor)
18            .build()
19        ...
20    }
21 }

```

If no `Executor` is given, **Jenetics** uses a common `ForkJoinPool`⁹ for concurrency.

4.2 Randomness

In general, GAs heavily depends on *pseudo* random number generators (PRNG) for creating new individuals and for the selection- and mutation-algorithms. **Jenetics** uses the Java `Random` object, respectively sub-types from it, for generating random numbers. To make the random engine pluggable, the `Random` object is always fetched from the `RandomRegistry`. This makes it possible to change the implementation of the random engine without changing the client code. The central `RandomRegistry` also allows to easily change `Random` engine even for specific parts of the code.

The following example shows how to change and restore the `Random` object. When opening the `with` scope, changes to the `RandomRegistry` are only visible within this scope. Once the `with` scope is left, the original `Random` object is restored.

```

1 List<Genotype<DoubleGene>> genotypes =
2     RandomRegistry.with(new Random(123), r -> {
3         Genotype.of(DoubleChromosome.of(0.0, 100.0, 10))
4             .instances()
5             .collect(Collectors.toList())
6     });

```

With the previous listing, a random, but reproducible, list of genotypes is created. This might be useful while testing your application or when you want to evaluate the `EvolutionStream` several times with the same initial population.

```

1 Engine<DoubleGene, Double> engine = ...;
2 // Create a new evolution stream with the given
3 // initial genotypes.
4 Phenotype<DoubleGene, Double> best = engine.stream(genotypes)
5     .limit(10)
6     .collect(EvolutionResult.toBestPhenotype());

```

The example above uses the generated genotypes for creating the `EvolutionStream`. Each created stream uses the same starting population, but will, most likely, create a different result. This is because the stream evaluation is still non-deterministic.

⁹<https://docs.oracle.com/javase/8/docs/api/java/util/concurrent/ForkJoinPool.html>

Setting the PRNG to a `Random` object with a defined seed has the effect, that every evolution *stream* will produce the same result—in an single threaded environment.

The parallel nature of the GA implementation requires the creation of streams $t_{i,j}$ of random numbers which are statistically independent, where the streams are numbered with $j = 1, 2, 3, \dots, p$, p denotes the number of processes. We expect statistical independence between the streams as well. The used PRNG should enable the GA to *play fair*, which means that the outcome of the GA is strictly independent from the underlying hardware and the number of parallel processes or threads. This is essential for reproducing results in parallel environments where the number of parallel tasks may vary from run to run.

The *Fair Play* property of a PRNG guarantees that the quality of the genetic algorithm (evolution stream) does not depend on the degree of parallelization.

When the `Random` engine is used in an multi-threaded environment, there must be a way to parallelize the sequential PRNG. Usually this is done by taking the elements of the sequence of pseudo-random numbers and distribute them among the threads. There are essentially four different parallelizations techniques used in practice: *Random seeding*, *Parameterization*, *Block splitting* and *Leapfrogging*.

Random seeding Every thread uses the same kind of PRNG but with a different seed. This is the default strategy used by the **Jenetics** library. The `RandomRegistry` is initialized with the `ThreadLocalRandom` class from the `java.util.concurrent` package. Random seeding works well for the most problems but without theoretical foundation.¹⁰ If you assume that this strategy is responsible for some *non-reproducible* results, consider using the `LCG64ShiftRandom` PRNG instead, which uses *block splitting* as parallelization strategy.

Parameterization All threads uses the same kind of PRNG but with different parameters. This requires the PRNG to be parameterizable, which is not the case for the `Random` object of the JDK. You can use the `LCG64ShiftRandom` class if you want to use this strategy. The theoretical foundation for these method is weak. In a massive parallel environment you will need a reliable set of parameters for every random stream, which are not trivial to find.

Block splitting With this method each thread will be assigned a non-overlapping contiguous block of random numbers, which should be enough for the whole runtime of the process. If the number of threads is not known in advance, the length of each block should be chosen much larger then the maximal

¹⁰This is also expressed by Donald Knuth's advice: »Random number generators should not be chosen at random.«

expected number of threads. This strategy is used when using the `LCG64-ShiftRandom.ThreadLocal` class. This class assigns every thread a block of $2^{56} \approx 7,2 \cdot 10^{16}$ random numbers. After 128 threads, the blocks are recycled, but with changed seed.




Figure 4.1: Block splitting

Leapfrog With the leapfrog method each thread $t \in [0, P)$ only consumes the P^{th} random number and jump ahead in the random sequence by the number of threads, P . This method requires the ability to jump very quickly ahead in the sequence of random numbers by a given amount. Figure 4.2 graphically shows the concept of the *leapfrog* method.




Figure 4.2: Leapfrogging

LCG64ShiftRandom The `LCG64ShiftRandom` class is a part of the `trng::-lcg64_shift` PRNG class of the `TRNG`¹¹ library, implemented in C++.[3] It implements additional methods, which allows to implement the *block splitting*—and also the *leapfrog*—method.

```

1 public class LCG64ShiftRandom extends Random {
2     public void split(final int p, final int s);
3     public void jump(final long step);
4     public void jump2(final int s);
5     ...
6 }

```

Listing 12: LCG64ShiftRandom class

¹¹<http://numbercrunch.de/trng/>

Listing 12 shows the interface used for implementing the block splitting and leapfrog parallelizations technique. This methods have the following meaning:

split Changes the internal state of the PRNG in a way that future calls to `nextLong()` will generated the s^{th} sub-stream of p^{th} sub-streams. s must be within the range of $[0, p - 1)$. This method is used for parallelization via *leapfrogging*.

jump Changes the internal state of the PRNG in such a way that the engine jumpss steps ahead. This method is used for parallelization via *block splitting*.

jump2 Changes the internal state of the PRNG in such a way that the engine jumps 2^s steps ahead. This method is used for parallelization via *block splitting*.

Runtime performance Table 4.1 shows the random number generation speed for the different PRNG implementations.¹²

	int/s	long/s	float/s	double/s
Random	$67 \cdot 10^6$	$33 \cdot 10^6$	$67 \cdot 10^6$	$33 \cdot 10^6$
ThreadLocalRandom	$333 \cdot 10^6$	$344 \cdot 10^6$	$237 \cdot 10^6$	$220 \cdot 10^6$
LCG64ShiftRandom	$243 \cdot 10^6$	$251 \cdot 10^6$	$196 \cdot 10^6$	$199 \cdot 10^6$

Table 4.1: Performance of various PRNG implementations.

The default PRNG used by the **Jenetics** has the best runtime performance behavior (for generating `int` values).

4.3 Serialization

Jenetics supports serialization for a number of classes, most of them are located in the `org.jenetics` package. Only the concrete implementations of the **Gene** and the **Chromosome** interfaces implements the **Serializable** interface. This gives a greater flexibility when implementing own **Genes** and **Chromosomes**.

- `BitGene`
- `BitChromosome`
- `CharacterGene`
- `CharacterChromosome`
- `IntegerGene`
- `IntegerChromosome`
- `LongGene`
- `LongChromosome`
- `DoubleGene`
- `DoubleChromosome`
- `EnumGene`
- `PermutationChromosome`
- `Genotype`
- `Phenotype`
- `Population`

¹²Measured on a Intel(R) Core(TM) i5-3427U CPU @ 1.80GHz with Java(TM) SE Runtime Environment (build 1.8.0_25-b17)—Java HotSpot(TM) 64-Bit Server VM (build 25.25-b02, mixed mode)—, using the JHM micro-benchmark library.

With the serialization mechanism you can write a population to disk and load it into an new `EvolutionStream` at a later time. It can also be used to transfer populations to evolution engines, running on different hosts, over a network link. The `IO` class, located in the `org.jenetics.util` package, supports native Java serialization and JAXB XML serialization.

```

1 // Creating result population.
2 EvolutionResult<DoubleGene, Double> result = stream
3   .collect(toBestEvolutionResult());
4
5 // Writing the population to disk.
6 final File file = new File("population.xml");
7 IO.jaxb.write(result.getPopulation(), file);
8
9 // Reading the population from disk.
10 Population<DoubleGene, Double> population =
11   (Population<DoubleGene, Double>)IO.jaxb.read(file);
12 EvolutionStream<DoubleGene, Double> stream = Engine
13   .build(ff, gtf)
14   .stream(population, 1);

```

The following listing shows the XML serialization of a `Population` which consists of `Genotypes` as shown in figure 3.2 on page 6; only the first `Phenotype` is shown.

```

1 <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
2 <org.jenetics.Population size="5">
3   <phenotype
4     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
5     xsi:type="org.jenetics.Phenotype" generation="294"
6   >
7     <genotype length="5" ngenes="100">
8       <chromosome
9         xsi:type="org.jenetics.DoubleChromosome"
10        length="20" min="0.0" max="1.0"
11      >
12        <allele>0.27251556008507416</allele>
13        <allele>0.003140816229067145</allele>
14        <allele>0.43947528327497376</allele>
15        <allele>0.10654807463069327</allele>
16        <allele>0.19696530915810317</allele>
17        <allele>0.7450003838065538</allele>
18        <allele>0.5594416969271359</allele>
19        <allele>0.02823782430152355</allele>
20        <allele>0.5741102315010789</allele>
21        <allele>0.4533651041367144</allele>
22        <allele>0.81148141800367</allele>
23        <allele>0.5710456351848858</allele>
24        <allele>0.30166768355230955</allele>
25        <allele>0.5455492865240272</allele>
26        <allele>0.21068427527733102</allele>
27        <allele>0.5265067943902246</allele>
28        <allele>0.273549098065591</allele>
29        <allele>0.2648197379297126</allele>
30        <allele>0.8732775776362911</allele>
31        <allele>0.9498003919007005</allele>
32      </chromosome>
33      ...
34    </genotype>
35    <fitness
36      xmlns:xs="http://www.w3.org/2001/XMLSchema"
37      xsi:type="xs:double"
38    >234.23443</fitness>
39    <raw-fitness
40      xmlns:xs="http://www.w3.org/2001/XMLSchema"
41      xsi:type="xs:double"
42    >34.2498</raw-fitness>
43  </phenotype>
44  ...
45 </org.jenetics.Population>

```

When serializing a whole population the fitness function and fitness scaler are not part of the serialized XML file. If an `EvolutionStream` is initialized with a previously serialized `Population`, the `Engine`'s current fitness function and fitness scaler are used for *re-calculating* the fitness values.

4.4 Utility classes

The `org.jenetics.util` and the `org.jenetics.stat` package of the library contains utility and helper classes which are essential for the implementation of the GA.

org.jenetics.util.Seq Most notable are the `Seq` interfaces and its implementation. They are used, among others, in the `Chromosome` and `Genotype` classes and holds the `Genes` and `Chromosomes`, respectively. The `Seq` interface itself represents a fixed-sized, ordered sequence of elements. It is an abstraction over the Java build-in `array`-type, but much safer to use for *generic* elements, because there are no casts needed when using *nested* generic types.




Figure 4.3: Seq class diagram

Figure 4.3 shows the `Seq` class diagram with their most important methods. The interfaces `MSeq` and `ISeq` are mutable, respectively immutable specializations of the basis interface. Creating instances of the `Seq` interfaces is possible via the static factory methods of the interfaces.

```

1 // Create "different" sequences.
2 final Seq<Integer> a1 = Seq.of(1, 2, 3);
3 final MSeq<Integer> a2 = MSeq.of(1, 2, 3);
4 final ISeq<Integer> a3 = MSeq.of(1, 2, 3).toISeq();
5 final MSeq<Integer> a4 = a3.copy();
6
7 // The 'equals' method performs element-wise comparison.
8 assert(a1.equals(a2) && a1 != a2);
9 assert(a2.equals(a3) && a2 != a3);
10 assert(a3.equals(a4) && a3 != a4);
  
```

How to create instances of the three `Seq` types is shown in the listing above. The `Seq` classes also allows a more *functional* programming style. For a full method description refer to the Javadoc.

org.jenetics.stat This package contains classes for calculating statistical moments. They are designed to work smoothly with the Java Stream API and are divided into mutable (number) consumers and immutable value classes, which holds the statistical moments. The additional classes calculate the

- *minimum*,
- *maximum*,
- *sum*,
- *mean*,
- *variance*,
- *skewness* and
- *kurtosis* value.

Numeric type	Consumer class	Value class
<code>int</code>	<code>IntMomentStatistics</code>	<code>IntMoments</code>
<code>long</code>	<code>LongMomentStatistics</code>	<code>LongMoments</code>
<code>double</code>	<code>DoubleMomentStatistics</code>	<code>DoubleMoments</code>

Table 4.2: Statistics classes

Table 4.2 contains the available statistical moments for the different numeric types. The following code snippet shows an example on how to collect double statistics from an given `DoubleGene` stream.

```

1 // Collecting into an statistics object.
2 DoubleChromosome chromosome = ...
3 DoubleMomentStatistics statistics = chromosome.toSeq().stream()
4   .collect(DoubleMomentStatistics
5     .toDoubleMomentStatistics(v -> v.doubleValue()));
6
7 // Collecting into an moments object.
8 DoubleMoments moments = chromosome.toSeq().stream()
9   .collect(DoubleMoments.toDoubleMoments(v -> v.doubleValue()));

```

5 Extending Jenetics

The **Jenetics** library was designed to give you a great flexibility in transforming your problem into a structure that can be solved by an GA. It also comes with different implementations for the base data-types (genes and chromosomes) and operators (alterers and selectors). If it is still some functionality missing, this section describes how you can extend the existing classes. Most of the *extensible* classes are defined by an interface and have an abstract implementation which makes it easier to extend it.

5.1 Genes

Genes are the starting point in the class hierarchy. They hold the actual information, the alleles, of the problem domain. Beside the *classical* bit-gene, **Jenetics** comes with gene implementations for numbers (`double`-, `int`- and `long` values), characters and enumeration types.

For implementing your own gene type you have to implement the `Gene` interface with three methods: (1) the `getAllele()` method which will return the wrapped data, (2) the `newInstance` method for creating new, random instances of the gene—must be of the same type and have the same constraint—and (3) the `isValid()` method which checks if the gene fulfill the expected constraints. The gene constraint might be violated after mutation and/or recombination. If you want to implement a new number-gene, e. g. a gene which holds complex

values, you may want extend it from the abstract `NumericGene` class. Every `Gene` extends the `Serializable` interface. For *normal* genes there is no more work to do for using the Java serialization mechanism.

The custom `Genes` and `Chromosomes` implementations must use the `Random` engine available via the `RandomRegistry.getRandom` method when implementing their factory methods. Otherwise it is not possible to seamlessly change the `Random` engine by using the `RandomRegistry.setRandom` method.

5.2 Chromosomes

A new gene type normally needs a corresponding chromosome implementation. The most important part of a chromosome is the factory method `newInstance`, which lets the evolution `Engine` create a new `Chromosome` instance from a sequence of `Genes`. This method is used by the `Alterers` when creating new, combined `Chromosomes`. The other methods should be self-explanatory. The chromosome has the same serialization mechanism as the gene. For the minimal case it can extend the `Serializable` interface.

5.3 Selectors

If you want to implement your own selection strategy you only have to implement the `Selector` interface with the `select` method.

```

1 @FunctionalInterface
2 public interface Selector<
3     G extends Gene<?, G>,
4     C extends Comparable<? super C>
5 >
6 {
7     public Population<G, C> select(
8         Population<G, C> population,
9         int count,
10        Optimize opt
11    );
12 }

```

Listing 13: Selector interface

The first parameter is the original `population` from which the *sub*-population is selected. The second parameter, `count`, is the number of individuals of the returned sub-population. Depending on the selection algorithm, it is possible that the sub-population contains more elements than the original one. The last parameter, `opt`, determines the optimization strategy which must be used by the selector. This is exactly the point where it is decided whether the GA minimizes or maximizes the fitness function.

Before implementing a selector from scratch, consider to extend your selector from the `ProbabilitySelector` (or any other available `Selector` implementation). It is worth the effort to try to express your selection strategy in terms of selection property $P(i)$.

5.4 Alterers

For implementing a new alterer class it is necessary to implement the `Alterer` interface. You might do this if your new `Gene` type needs a special kind of alterer not available in the `Jenetics` project.

```

1 @FunctionalInterface
2 public interface Alterer<
3     G extends Gene<?, G>,
4     C extends Comparable<? super C>
5 >
6 {
7     public int alter(
8         Population<G, C> population,
9         long generation
10    );
11 }

```

Listing 14: `Alterer` interface

The first parameter of the `alter` method is the `Population` which has to be altered. Since the `Population` class is mutable, the altering is performed in place. The second parameter is the `generation` of the newly created individuals and the return value is the number of genes that has been altered.

5.5 Statistics

During the developing phase of an application which uses the `Jenetics` library, additional statistical data about the evolution process is crucial. Such data can help to optimize the parametrization of the evolution Engine. A good starting point is to use the `EvolutionStatistics` class in the `org.jenetics.-engine` package. If the data in the `EvolutionStatistics` class doesn't fit your needs, you simply have to write your own statistics class. It is not possible to derive from the existing `EvolutionStatistics` class. This is not a real restriction, since you still can use the class by delegation. Just implement the Java `Consumer<EvolutionResult<G, C>>` interface.

5.6 Engine

The evolution Engine itself can't be extended, but it is still possible to create an `EvolutionStream` without using the `Engine` class.¹³ Because the `EvolutionStream` has no direct dependency to the `Engine`, it is possible to use an different, special evolution `Function`.

```

1 public final class SpecialEngine {
2     // The Genotype factory.
3     private static final Factory<Genotype<DoubleGene>> GTF =
4         Genotype.of(DoubleChromosome.of(0, 1));
5
6     // The fitness function.
7     private static Double fitness(final Genotype<DoubleGene> gt) {
8         return gt.getGene().getAllele();
9     }
10
11     // Create new evolution start object.

```

¹³Also refer to section 3.3.4 on page 19 on how to create an `EvolutionStream` from an evolution `Function`.

```

12 private static EvolutionStart<DoubleGene, Double>
13 start(final int populationSize, final long generation) {
14     final Population<DoubleGene, Double> population = GTF
15         .instances()
16         .map(gt -> Phenotype
17             .of(gt, generation, SpecialEngine::fitness))
18         .limit(populationSize)
19         .collect(Population.toPopulation());
20
21     return EvolutionStart.of(population, generation);
22 }
23
24 // The special evolution function.
25 private static EvolutionResult<DoubleGene, Double>
26 evolve(final EvolutionStart<DoubleGene, Double> start) {
27     return ...; // Add implementation!
28 }
29
30 public static void main(final String[] args) {
31     final Genotype<DoubleGene> best = EvolutionStream
32         .of(() -> start(50, 0), SpecialEngine::evolve)
33         .limit(limit.bySteadyFitness(10))
34         .limit(100)
35         .collect(EvolutionResult.toBestGenotype());
36
37     System.out.println("Best Genotype: " + best);
38 }
39 }

```

Listing 15: Special evolution engine

Listing 15 on the preceding page shows a *complete* implementation stub for using an own special evolution `Function`.

6 Advanced topics

This section describes some advanced topics for setting up an evolution `Engine` or `EvolutionStream`. It contains some problem encoding examples and how to override the default validation strategy of the given `Genotypes`. The last section contains a detailed description of the implemented termination strategies.

6.1 Encoding

This section presents some encoding examples for common problems. The encoding should be a complete and minimal expression of a solution to the problem. An encoding is complete if it contains enough information to represent every solution to the problem. An minimal encoding contains only the information needed to represent a solution to the problem. If an encoding contains more information than is needed to uniquely identify solutions to the problem, the search space will be larger than necessary.

Whenever possible, the encoding should not be able to represent infeasible solutions. If a genotype can represent an infeasible solution, care must be taken in the fitness function to give partial credit to the genotype for its »good« genetic material while sufficiently penalizing it for being infeasible. Implementing a specialized `Chromosome`, which won't create invalid encodings can be a solution to this problem. In general, it is much more desirable to design a representation

that can only represent valid solutions so that the fitness function measures only fitness, not validity. An encoding that includes invalid individuals enlarges the search space and makes the search more costly.

6.1.1 Real function

Jenetics contains three different numeric gene and chromosome implementations, which can be used to encode a real function, $f : \mathbb{R} \rightarrow \mathbb{R}$:

- IntegerGene/Chromosome,
- LongGene/Chromosome and
- DoubleGene/Chromosome.

It is quite easy to encode a real function. Only the minimum and maximum value of the function domain must be defined. The `DoubleChromosome` of length 1 is then wrapped into a `Genotype`.

```
1 Genotype.of(
2     DoubleChromosome.of(min, max, 1)
3 );
```

Decoding the double value from the `Genotype` is also straight forward. Just get the first gene from the first chromosome, with the `getGene()` method, and convert it to a `double`.

```
1 static double toDouble(final Genotype<DoubleGene> gt) {
2     return gt.getGene().doubleValue();
3 }
```

When the `Genotype` only contains *scalar* chromosomes¹⁴, it should be clear, that it can't be altered by every `Alterer`. That means, that none of the `Crossover` alterers will be able to create modified `Genotypes`. For *scalars* the appropriate alterers would be the `MeanAlterer`, `GaussianAlterer` and `Mutator`.

Scalar Chromosomes and/or Genotypes can only be altered by `MeanAlterer`, `GaussianAlterer` and `Mutator` classes. Other alterers are allowed, but will have no effect on the Chromosomes.

6.1.2 Scalar function

Optimizing a function $f(x_1, \dots, x_n)$ of one or more variable whose range is one-dimensional, we have two possibilities for the `Genotype` encoding.[9] For the *first* encoding we expect that all variables, x_i , have the same minimum and maximum value. In this case we can simply create a `Genotype` with a `NumericChromosome` of the desired length n .

```
1 Genotype.of(
2     DoubleChromosome.of(min, max, n)
3 );
```

¹⁴Scalar chromosomes contain only one gene.

The decoding of the **Genotype** requires a cast of the first **Chromosome** to a **DoubleChromosome**. With a call to the **DoubleChromosome.toArray()** method we return the variables (x_1, \dots, x_n) as **double[]** array.

```
1 static double[] toScalars(final Genotype<DoubleGene> gt) {
2     return ((DoubleChromosome)gt.getChromosome()).toArray();
3 }
```

With the *first* encoding you have the possibility to use all available alterers, including all **Crossover** alterer classes.

The *second* encoding *must* be used if the minimum and maximum value of the variables x_i can't be the same for all i . For the different domains, each variable x_i is represented by a *NumericChromosome* with length one. The final **Genotype** will consist of n **Chromosomes** with length one.

```
1 Genotype.of(
2     DoubleChromosome.of(min1, max1, 1),
3     DoubleChromosome.of(min2, max2, 1),
4     ...
5     DoubleChromosome.of(minn, maxx, 1)
6 );
```

With the help of the new Java Stream API, the decoding of the **Genotype** can be done in a view lines. The **DoubleChromosome** stream, which is created from the chromosome **Seq**, is first mapped to **double** values and then collected into an array.

```
1 static double[] toScalars(final Genotype<DoubleGene> gt) {
2     return gt.toSeq().stream()
3         .mapToDouble(c -> c.getGene().doubleValue())
4         .toArray();
5 }
```

As already mentioned, with the use of scalar chromosomes we can only use the **MeanAlterer**, **GaussianAlterer** or **Mutator** alterer class.

If there are performance issues in converting the **Genotype** into a **double[]** array, or any other numeric array, you can access the **Genes** directly via the **Genotype.get(i, j)** method and than convert it to the desired numeric value, by calling **intValue()**, **longValue()** or **doubleValue()**.

6.1.3 Vector function

A function $f(X_1, \dots, X_n)$, of one to n variables whose range is m -dimensional, is encoded by m **DoubleChromosomes** of length n .^[10] The domain–minimum and maximum values–of one variable X_i are the same in this encoding.

```
1 Genotype.of(
2     DoubleChromosome.of(min1, max1, m),
3     DoubleChromosome.of(min2, max2, m),
4     ...
5     DoubleChromosome.of(minn, maxx, m)
6 );
```

The decoding of the vectors is quite easy with the help of the Java Stream API. In the first **map** we have to cast the **Chromosome<DoubleGene>** object to the actual **DoubleChromosome**. The second **map** then converts each **DoubleChromosome** to an **double[]** array, which is collected to an 2-dimensional **double[n][m]** array afterwards.


```

1 static double [][] toVectors(final Genotype<DoubleGene> gt) {
2     return gt.toSeq().stream()
3         .map(DoubleChromosome.class::cast)
4         .map(DoubleChromosome::toArray)
5         .toArray(double [][]::new);
6 }

```

For the special case of $n = 1$, the decoding of the **Genotype** can be simplified to the decoding we introduced for scalar functions in section 6.1.2.

```

1 static double [] toVector(final Genotype<DoubleGene> gt) {
2     return ((DoubleChromosome)gt.getChromosome()).toArray();
3 }

```

6.1.4 Affine transformation

An affine transformation^{15, 16} is usually performed by a matrix multiplication with a transformation matrix—in a homogeneous coordinates system¹⁷. For a transformation in \mathbb{R}^2 , we can define the matrix A ¹⁸:

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix}. \quad (6.1)$$

A simple representation can be done by creating a **Genotype** which contains two **DoubleChromosomes** with a length of 3.

```

1 Genotype.of(
2     DoubleChromosome.of(min, max, 3),
3     DoubleChromosome.of(min, max, 3)
4 );

```

The drawback with this kind of encoding is, that we will create a lot of *invalid* (non-affine transformation matrices) during the evolution process, which must be detected and discarded. It is also difficult to find the right parameters for the *min* and *max* values of the **DoubleChromosomes**.

A better approach will be to encode the transformation parameters instead of the transformation matrix. The affine transformation can be expressed by the following parameters:

- s_x – the scale factor in x direction
- s_y – the scale factor in y direction
- t_x – the offset in x direction
- t_y – the offset in y direction
- θ – the rotation angle clockwise around origin
- k_x – shearing parallel to x axis
- k_y – shearing parallel to y axis

¹⁵https://en.wikipedia.org/wiki/Affine_transformation

¹⁶<http://mathworld.wolfram.com/AffineTransformation.html>

¹⁷https://en.wikipedia.org/wiki/Homogeneous_coordinates

¹⁸https://en.wikipedia.org/wiki/Transformation_matrix

This parameters can then be represented by the following **Genotype**.

```

1 Genotype.of(
2     // Scale
3     DoubleChromosome.of(0.1, 5.0, 2),
4     // Translation
5     DoubleChromosome.of(-300.0, 300.0, 2),
6     // Rotation
7     DoubleChromosome.of(-Math.PI, Math.PI),
8     // Shear
9     DoubleChromosome.of(-0.5, 0.5, 2)
10 );

```

This encoding ensures that no invalid **Genotype** will be created during the evolution process, since the crossover will be only performed on the same kind of chromosome (same chromosome index). To convert the **Genotype** back to the transformation matrix A , we can use the following equations:

$$\begin{aligned}
 a_{11} &= s_x ((1 + k_x k_y) \cos \theta + k_y \sin \theta) \\
 a_{12} &= s_x (k_x \cos \theta + \sin \theta) \\
 a_{13} &= t_x \\
 a_{21} &= s_y (-(1 + k_x k_y) \sin \theta + k_y \cos \theta) \\
 a_{22} &= s_y (-k_x \sin \theta + \cos \theta) \\
 a_{23} &= t_y
 \end{aligned} \tag{6.2}$$

In Java code, the conversion from the **Genotype** to the transformation matrix will look like this:

```

1 static double [][] toMatrix(final Genotype<DoubleGene> gt) {
2     final double sx = gt.get(0, 0).doubleValue();
3     final double sy = gt.get(0, 1).doubleValue();
4     final double tx = gt.get(1, 0).doubleValue();
5     final double ty = gt.get(1, 1).doubleValue();
6     final double th = gt.get(2, 0).doubleValue();
7     final double kx = gt.get(3, 0).doubleValue();
8     final double ky = gt.get(3, 1).doubleValue();
9
10    final double a11 = sx*((1 + kx*ky)*cos(th) + ky*sin(th));
11    final double a12 = sx*(kx*cos(th) + sin(th));
12    final double a21 = sy*(-(1 + kx*ky)*sin(th) + ky*cos(th));
13    final double a22 = sy*(-kx*sin(th) + cos(th));
14
15    return new double [][] {
16        {a11, a12, tx},
17        {a21, a22, ty},
18        {0.0, 0.0, 1.0}
19    };
20 }

```

For the introduced encoding all kind of alterers can be used. Since we have one scalar **DoubleChromosome**, the rotation angle θ , it is recommended also to add an **MeanAlterer** or **GaussianAlterer** to the list of alterers.

6.1.5 Graph

A graph can be represented in many different ways. The most known graph representation is the adjacency matrix. The following encoding examples uses adjacency matrices with different characteristics.

Undirected graph In an undirected graph the edges between the vertices have no direction. If there is a path between nodes i and j , it is assumed that there is also path from j to i .




Figure 6.1: Undirected graph and adjacency matrix

Figure 6.1 shows an undirected graph and its corresponding matrix representation. Since the edges between the nodes have no direction, the values of the lower diagonal matrix are not taken into account. An application which optimizes an undirected graph has to ignore this part of the matrix.¹⁹

```

1 final int n = 6;
2 final Genotype<BitGene> gt = Genotype
3   .of(BitChromosome.of(n), n);

```

The code snippet above shows how to create an adjacency matrix for a graph with $n = 6$ nodes. It creates a genotype which consists of n `BitChromosomes` of length n each. Whether the node i is connected to node j can be easily checked by calling `gt.get(i-1, j-1).booleanValue()`. For extracting the whole matrix as `int[]` array, the following code can be used.

```

1 final int [][] array = gt.toSeq().stream()
2   .map(c -> c.toSeq().stream()
3     .mapToInt(BitGene::ordinal)
4     .toArray())
5   .toArray(int [][]::new);

```

Directed graph A directed graph (digraph) is a graph where the path between the nodes have a direction associated with them. The encoding of a directed graph looks exactly like the encoding of an undirected graph. This time the whole matrix is used and the second diagonal matrix is no longer ignored.

Figure 6.2 on the next page shows the adjacency matrix of a digraph. This time the whole matrix is used for representing the graph.

Weighted directed graph A weighted graph associates a weight (label) with every path in the graph. Weights are usually real numbers. They may be restricted to rational numbers or integers.

¹⁹This property violates the *minimal* encoding requirement we mentioned at the beginning of section 6.1 on page 32. For simplicity reason this will be ignored for the undirected graph encoding.




Figure 6.2: Directed graph and adjacency matrix




Figure 6.3: Weighted graph and adjacency matrix

The following code snippet shows how the **Genotype** of the matrix is created.

```

1 final int n = 6;
2 final double min = -1;
3 final double max = 20;
4 final Genotype<DoubleGene> gt = Genotype
5   .of(DoubleChromosome.of(min, max, n), n);

```

For accessing the single matrix elements, you can simply call **Genotype.get(i, j).doubleValue()**. If the interaction with another library requires an **double[][]** array, the following code can be used.

```

1 final double [][] array = gt.toSeq().stream()
2   .map(DoubleChromosome.class::cast)
3   .map(DoubleChromosome::toArray)
4   .toArray(double [][]::new);

```

6.2 Validation

A given problem should usually be encoded in a way, that it is not possible for the evolution **Engine** to create *invalid* individuals (**Genotypes**). Some possible encodings for common data-structures are described in section 6.1 on page 32. The **Engine** creates new individuals in the *altering* step, by rearranging (or creating new) **Genes** within a **Chromosome**. Since a **Genotype** is treated as *valid*

if every single **Gene** in every **Chromosome** is *valid*, the validity property of the **Genes** determines the validity of the whole **Genotype**.

The **Engine** tries to create only valid individuals when creating the initial **Population** and when it replaces **Genotypes** which has been *destroyed* by the altering step. Individuals which has exceeded its lifetime are also replaced by new valid ones. To guarantee the termination of the **Genotype** creation, the **Engine** is parameterized with the maximal number of retries (**individualCreationRetries**)²⁰.

If the described validation mechanism doesn't fulfill your needs, you can *override* the validation mechanism by creating the **Engine** with an external **Genotype validator**.

```

1 final Predicate<? super Genotype<DoubleGene>> validator = gt -> {
2     // Implement advanced Genotype check.
3     boolean valid = ...;
4     return valid;
5 };
6 final Engine<DoubleGene, Double> engine = Engine.builder(gtf, ff)
7     .limit(100)
8     .genotypeValidator(validator)
9     .individualCreationRetries(15)
10    .build();

```

Having the possibility to replace the default validation check is a nice thing, but it is better to not create invalid individuals in the first place. For achieving this goal, you have two possibilities:

1. Creating an explicit **Genotype** factory and
2. implementing new **Gene/Chromosome/Alterer** classes.

Genotype factory The usual mechanism for defining an encoding is to create a **Genotype prototype**²¹. Since the **Genotype** implements the **Factory** interface, an prototype instance can easily passed to the **Engine.builder** method. For a more advanced **Genotype** creation, you *only* have to create an explicit **Genotype** factory.

```

1 final Factory<Genotype<DoubleGene>> gtf = () -> {
2     // Implement your advanced Genotype factory.
3     Genotype<DoubleGene> genotype = ...;
4     return genotype;
5 };
6 final Engine<DoubleGene, Double> engine = Engine.builder(gtf, ff)
7     .limit(100)
8     .individualCreationRetries(15)
9     .build();

```

With this method you can avoid that the **Engine** creates invalid individuals in the first place, but it is still possible that the alterer step will destroy your **Genotypes**.

Gene/Chromosome/Alterer Creating your own **Gene**, **Chromosome** and **Alterer** classes is the most heavy-wighted possibility for solving the *validity* problem. Refer to section 5 on page 29 for a more detailed description on how to implement this classes.

²⁰See section 3.3.3 on page 17.

²¹https://en.wikipedia.org/wiki/Prototype_pattern

6.3 Termination

Termination is the criterion by which the evolution stream decides whether to continue or truncate the stream. This section gives a deeper insight into the different ways of terminating or truncating the evolution stream, respectively.

The termination strategies in the following sub-sections were tested by solving the Knapsack problem (see section 7.3 on page 48) with 250 items. This makes it a real problem with a search-space size of $2^{250} \approx 10^{75}$ elements. To make the tests comparable, all runs uses the very same set of knapsack items.

6.3.1 Fixed generation

The simplest way for terminating the evolution process, is to define a maximal number of generations on the `EvolutionStream`. It just uses the existing `limit` method of the Java `Stream` interface.

```

1 final long MAX_GENERATIONS = 100;
2 EvolutionStream<DoubleGene, Double> stream = engine.stream()
3   .limit(MAX_GENERATIONS);

```

This kind of termination method should always be applied—usually additional with other evolution terminators—, to guarantee the truncation of the evolution stream and to define an upper limit of the executed generations.




Figure 6.4: Fixed generation termination

Diagram 6.4 shows the best fitness values of the used Knapsack problem after a given number of generations, whereas the candle-stick points represents the *min*, *25th percentile*, *median*, *75th percentile* and *max* fitness after 250 repetitions per generation. The solid line shows for the *mean* of the best fitness values. For a small increase of the fitness value, the needed generations grows exponentially. This is especially the case when the fitness is approaching to its *maximal* value.

6.3.2 Steady fitness

The *steady fitness* strategy truncates the evolution stream if its best fitness hasn't changed after a given number of generations. The predicate maintains an internal state, the number of generations with non increasing fitness, and must be newly created for every evolution stream.

```

1 final class SteadyFitnessLimit<C extends Comparable<? super C>>
2     implements Predicate<EvolutionResult<?, C>>
3 {
4     private final int _generations;
5     private boolean _proceed = true;
6     private int _stable = 0;
7     private C _fitness;
8
9     public SteadyFitnessLimit(final int generations) {
10         _generations = generations;
11     }
12
13     @Override
14     public boolean test(final EvolutionResult<?, C> er) {
15         if (!_proceed) return false;
16         if (_fitness == null) {
17             _fitness = er.getBestFitness();
18             _stable = 1;
19         } else {
20             final Optimize opt = result.getOptimize();
21             if (opt.compare(_fitness, er.getBestFitness()) >= 0) {
22                 _proceed = ++_stable <= _generations;
23             } else {
24                 _fitness = er.getBestFitness();
25                 _stable = 1;
26             }
27         }
28         return _proceed;
29     }
30 }

```

Listing 16: Steady fitness

Listing 16 shows the implementation of the `limit.bySteadyFitness(int)` in the `org.jenetics.engine` package. It should give you an impression of how to implement own termination strategies, which possible holds and internal state.

```

1 Engine<DoubleGene, Double> engine = ...
2 EvolutionStream<DoubleGene, Double> stream = engine.stream()
3     .limit(limit.bySteadyFitness(15));

```

The steady fitness terminator can be created by the `bySteadyFitness` factory method of the `org.jenetics.engine.limit` class. In the example above, the evolution stream is terminated after 15 stable generations.

Diagram 6.5 on the following page shows the actual total executed generation depending on the desired number of steady fitness generations. The variation of the total generation is quite big, as shown by the candle-sticks. Though the variation can be quite big—the termination test has been repeated 250 times for each data point—the tests showed that the *steady fitness* termination strategy always terminated, at least for the given test setup. The lower diagram give an overview of the fitness progression. Only the mean values of the maximal fitness is shown.




Figure 6.5: Steady fitness termination

6.3.3 Evolution time

This termination strategy stops the evolution when the elapsed evolution time exceeds an user-specified maximal value. The evolution stream is only truncated at the end of an generation and will not interrupt the current evolution step. An maximal evolution time of zero ms will at least evaluate one generation. In an time-critical environment, where a solution must be found within a maximal time period, this terminator let you define the desired guarantees.

```

1 Engine<DoubleGene, Double> engine = ...
2 EvolutionStream<DoubleGene, Double> stream = engine.stream()
3   .limit(limit.byExecutionTime(Duration.ofMillis(500)));

```

In the code example above, the `byExecutionTime(Duration)` method is used for creating the termination object. Another method, `byExecutionTime(Duration, Clock)`, lets you define the `java.time.Clock`, which is used for measure the execution time. **Jenetics** uses the nano precision clock `org.jenetics.util.NanoClock` for measuring the time. To have the possibility to define a different `Clock` implementation is especially useful for testing purposes.

Diagram 6.6 on the next page shows the evaluated generations depending on the execution time. Except for very small execution times, the evaluated generations per time unit stays quite stable.²² That means that a doubling of the execution time will double the number of evolved generations.

²²While running the tests, all other CPU intensive process has been stopped. The measuring started after a warm-up phase.




Figure 6.6: Execution time termination

6.3.4 Fitness threshold

A termination method that stops the evolution when the best fitness in the current population becomes less than the user-specified fitness threshold and the objective is set to minimize the fitness. This termination method also stops the evolution when the best fitness in the current population becomes greater than the user-specified fitness threshold when the objective is to maximize the fitness.

Diagram 6.7 on the following page shows executed generations depending on the minimal fitness value. The total generations grows exponentially with the desired fitness value. This means, that this termination strategy will (practically) not terminate, if the value for the fitness threshold is chosen to high. And it will definitely not terminate if the fitness threshold is higher than the *global* maximum of the fitness function. It will be a *perfect* strategy if you can define some *good enough* fitness value, which can be *easily* achieved.




Figure 6.7: Fitness threshold termination

Appendix

7 Examples

This section contains some coding examples which should give you a feeling of how to use the **Genetics** library. The given examples are complete, in the sense that they will compile and run and produce the given example output.

Running the examples delivered with the **Genetics** library can be started with the `run-examples.sh` script.

```
$ ./run-examples.sh
```

Since the script uses JARs located in the build directory you have to build it with the `jar Gradle` target first; see section 8 on page 53.

7.1 Ones counting

Ones counting is one of the simplest model-problem. It uses a binary chromosome and forms a classic genetic algorithm²³. The fitness of a Genotype is proportional to the number of ones.

```
1 import static org.jenetics.engine.EvolutionResult.toBestPhenotype;
2 import static org.jenetics.engine.limit.bySteadyFitness;
3
```

²³In the classic genetic algorithm the problem is a maximization problem and the fitness function is positive. The domain of the fitness function is a bit-chromosome.

```

4 import org.jenetics.BitChromosome;
5 import org.jenetics.BitGene;
6 import org.jenetics.Genotype;
7 import org.jenetics.Mutator;
8 import org.jenetics.Phentype;
9 import org.jenetics.RouletteWheelSelector;
10 import org.jenetics.SinglePointCrossover;
11 import org.jenetics.engine.Engine;
12 import org.jenetics.engine.EvolutionStatistics;
13
14 public class OnesCounting {
15
16     // This method calculates the fitness for a given genotype.
17     private static Integer count(final Genotype<BitGene> gt) {
18         return ((BitChromosome)gt.getChromosome()).bitCount();
19     }
20
21     public static void main(String[] args) {
22         // Configure and build the evolution engine.
23         final Engine<BitGene, Integer> engine = Engine
24             .builder(
25                 OnesCounting::count,
26                 BitChromosome.of(20, 0.15))
27             .populationSize(500)
28             .selector(new RouletteWheelSelector<>())
29             .alterers(
30                 new Mutator<>(0.55),
31                 new SinglePointCrossover<>(0.06))
32             .build();
33
34         // Create evolution statistics consumer.
35         final EvolutionStatistics<Integer, ?>
36             statistics = EvolutionStatistics.ofNumber();
37
38         final Phenotype<BitGene, Integer> best = engine.stream()
39             // Truncate the evolution stream after 7 "steady"
40             // generations.
41             .limit(bySteadyFitness(7))
42             // The evolution will stop after maximal 100
43             // generations.
44             .limit(100)
45             // Update the evaluation statistics after
46             // each generation
47             .peek(statistics)
48             // Collect (reduce) the evolution stream to
49             // its best phenotype.
50             .collect(toBestPhenotype());
51
52         System.out.println(statistics);
53         System.out.println(best);
54     }
55 }

```

The genotype in this example consists of one `BitChromosome` with a ones probability of 0.15. The altering of the offspring population is performed by mutation, with mutation probability of 0.55, and then by a single-point crossover, with crossover probability of 0.06. After creating the initial population, with the `ga.setup()` call, 100 generations are evolved. The tournament selector is used for both, the offspring- and the survivor selection—this is the default selector.²⁴

²⁴For the other default values (population size, maximal age, ...) have a look at the Javadoc:<http://jenetics.sourceforge.net/javadoc/index.html>

```

1 | +-----+
2 | | Time statistics |
3 | +-----+
4 | |           Selection: sum=0.017594839000 s; mean=0.001172989267 s |
5 | |           Altering: sum=0.100199190000 s; mean=0.006679946000 s |
6 | |           Fitness calculation: sum=0.007447917000 s; mean=0.000496527800 s |
7 | |           Overall execution: sum=0.131435736000 s; mean=0.008762382400 s |
8 | +-----+
9 | | Evolution statistics |
10 | +-----+
11 | |           Generations: 15 |
12 | |           Altered: sum=50,469; mean=3364.600000000 |
13 | |           Killed: sum=0; mean=0.000000000 |
14 | |           Invalids: sum=0; mean=0.000000000 |
15 | +-----+
16 | | Population statistics |
17 | +-----+
18 | |           Age: max=10; mean=0.855733; var=1.534325 |
19 | |           Fitness: |
20 | |               min = 1.000000000000 |
21 | |               max = 18.000000000000 |
22 | |               mean = 10.201466666667 |
23 | |               var = 7.237042787927 |
24 | +-----+
25 | [00001101|10111111|11111111] --> 18

```

The given example will print the overall timing statistics onto the console. In the *Evolution statistics* section you can see that it actually takes 15 generations to fulfill the termination criteria—finding no better result after 7 consecutive generations.

7.2 Real function

In this example we try to find the minimum value of the function

$$f(x) = \cos\left(\frac{1}{2} + \sin(x)\right) \cdot \cos(x). \quad (7.1)$$




Figure 7.1: Real function 2D

The graph of function 7.1, in the range of $[0, 2\pi]$, is shown in figure 7.1 and the listing beneath shows the GA implementation which will minimize the function.

```

1 import static java.lang.Math.PI;
2 import static java.lang.Math.cos;
3 import static java.lang.Math.sin;
4 import static org.jenetics.engine.EvolutionResult.toBestPhenotype;
5 import static org.jenetics.engine.limit.bySteadyFitness;
6
7 import org.jenetics.DoubleChromosome;
8 import org.jenetics.DoubleGene;
9 import org.jenetics.Genotype;
10 import org.jenetics.MeanAlterer;
11 import org.jenetics.Mutator;
12 import org.jenetics.Optimize;
13 import org.jenetics.Phenotype;
14 import org.jenetics.engine.Engine;
15 import org.jenetics.engine.EvolutionStatistics;
16
17 public class RealFunction {
18
19     // This method calculates the fitness for a given genotype.
20     private static Double eval(final Genotype<DoubleGene> gt) {
21         final double x = gt.getGene().doubleValue();
22         return cos(0.5 + sin(x))*cos(x);
23     }
24
25     public static void main(String[] args) {
26         final Engine<DoubleGene, Double> engine = Engine
27             // Create a new builder with the given fitness
28             // function and chromosome.
29             .builder(
30                 RealFunction::eval,
31                 DoubleChromosome.of(0.0, 2.0*PI))
32             .populationSize(500)
33             .optimize(Optimize.MINIMUM)
34             .alterers(
35                 new Mutator<>(0.03),
36                 new MeanAlterer<>(0.6))
37             // Build an evolution engine with the
38             // defined parameters.
39             .build();
40
41         // Create evolution statistics consumer.
42         final EvolutionStatistics<Double, ?>
43             statistics = EvolutionStatistics.ofNumber();
44
45         final Phenotype<DoubleGene, Double> best = engine.stream()
46             // Truncate the evolution stream after 7 "steady"
47             // generations.
48             .limit(bySteadyFitness(7))
49             // The evolution will stop after maximal 100
50             // generations.
51             .limit(100)
52             // Update the evaluation statistics after
53             // each generation
54             .peek(statistics)
55             // Collect (reduce) the evolution stream to
56             // its best phenotype.
57             .collect(toBestPhenotype());
58
59         System.out.println(statistics);
60         System.out.println(best);
61     }
62 }

```

The GA works with 1×1 DoubleChromosomes whose values are restricted to the range $[0, 2\pi]$.

```

1 | +-----+
2 | | Time statistics |
3 | +-----+
4 | |           Selection: sum=0.018522842000 s; mean=0.000974886421 s |
5 | |           Altering: sum=0.088635838000 s; mean=0.004665044105 s |
6 | | Fitness calculation: sum=0.009183011000 s; mean=0.000483316368 s |
7 | | Overall execution: sum=0.125399375000 s; mean=0.006599967105 s |
8 | +-----+
9 | | Evolution statistics |
10 | +-----+
11 | |           Generations: 19 |
12 | |           Altered: sum=3,557; mean=187.210526316 |
13 | |           Killed: sum=0; mean=0.000000000 |
14 | |           Invalids: sum=0; mean=0.000000000 |
15 | +-----+
16 | | Population statistics |
17 | +-----+
18 | |           Age: max=8; mean=1.102105; var=1.919461 |
19 | |           Fitness: |
20 | |               min = -0.938171897696 |
21 | |               max = 0.916616962601 |
22 | |               mean = -0.906090313372 |
23 | |               var = 0.021851810400 |
24 | +-----+
25 | [[[3.3891257853138868]]] --> -0.9381718976956661

```

The GA will generate a console output like above. The *exact* result of the function—for the given range—will be 3.389, 125, 782, 8907, 939... You can also see, that we reached the final result after 19 generations.

7.3 0/1 Knapsack

In the knapsack problem²⁵ a set of items, together with its size and value, is given. The task is to select a disjoint subset so that the total size does not exceed the knapsack size. For solving the 0/1 knapsack problem we define a BitChromosome, one bit for each item. If the i^{th} bit is set to one the i^{th} item is selected.

```

1 | import static org.jenetics.engine.EvolutionResult.toBestPhenotype;
2 | import static org.jenetics.engine.limit.bySteadyFitness;
3 |
4 | import java.util.Random;
5 | import java.util.function.Function;
6 | import java.util.stream.Collectors;
7 | import java.util.stream.Stream;
8 |
9 | import org.jenetics.BitChromosome;
10 | import org.jenetics.BitGene;
11 | import org.jenetics.Genotype;
12 | import org.jenetics.Mutator;
13 | import org.jenetics.Phenotype;
14 | import org.jenetics.RouletteWheelSelector;
15 | import org.jenetics.SinglePointCrossover;
16 | import org.jenetics.TournamentSelector;
17 | import org.jenetics.engine.Engine;
18 | import org.jenetics.engine.EvolutionStatistics;
19 | import org.jenetics.util.RandomRegistry;
20 |
21 | // This class represents a knapsack item, with a specific
22 | // "size" and "value".

```

²⁵https://en.wikipedia.org/wiki/Knapsack_problem

```

23 final class Item {
24     public final double size;
25     public final double value;
26
27     Item(final double size, final double value) {
28         this.size = size;
29         this.value = value;
30     }
31
32     // Create a new random knapsack item.
33     static Item random() {
34         final Random r = RandomRegistry.getRandom();
35         return new Item(r.nextDouble()*100, r.nextDouble()*100);
36     }
37
38     // Create a new collector for summing up the knapsack items.
39     static Collector<Item, ?, Item> toSum() {
40         return Collector.of(
41             () -> new double[2],
42             (a, b) -> {a[0] += b.size; a[1] += b.value;},
43             (a, b) -> {a[0] += b[0]; a[1] += b[1]; return a;},
44             r -> new Item(r[0], r[1])
45         );
46     }
47 }
48
49 // The knapsack fitness function class, which is parametrized with
50 // the available items and the size of the knapsack.
51 final class FF
52     implements Function<Genotype<BitGene>, Double>
53 {
54     private final Item[] items;
55     private final double size;
56
57     public FF(final Item[] items, final double size) {
58         this.items = items;
59         this.size = size;
60     }
61
62     @Override
63     public Double apply(final Genotype<BitGene> gt) {
64         final Item sum = ((BitChromosome)gt.getChromosome()).ones()
65             .mapToObj(i -> items[i])
66             .collect(Item.toSum());
67
68         return sum.size <= this.size ? sum.value : 0;
69     }
70 }
71
72 // The main class.
73 public class Knapsack {
74
75     public static void main(final String[] args) {
76         final int nitems = 15;
77         final double kssize = nitems*100.0/3.0;
78
79         final FF ff = new FF(
80             Stream.generate(Item::random)
81                 .limit(nitems)
82                 .toArray(Item[]::new),
83             kssize
84         );

```

```

85
86 // Configure and build the evolution engine.
87 final Engine<BitGene, Double> engine = Engine
88     .builder(ff, BitChromosome.of(nitems, 0.5))
89     .populationSize(500)
90     .survivorsSelector(new TournamentSelector<>(5))
91     .offspringSelector(new RouletteWheelSelector<>())
92     .alterers(
93         new Mutator<>(0.115),
94         new SinglePointCrossover<>(0.16))
95     .build();
96
97 // Create evolution statistics consumer.
98 final EvolutionStatistics<Double, ?>
99     statistics = EvolutionStatistics.ofNumber();
100
101 final Phenotype<BitGene, Double> best = engine.stream()
102     // Truncate the evolution stream after 7 "steady"
103     // generations.
104     .limit(bySteadyFitness(7))
105     // The evolution will stop after maximal 100
106     // generations.
107     .limit(100)
108     // Update the evaluation statistics after
109     // each generation
110     .peek(statistics)
111     // Collect (reduce) the evolution stream to
112     // its best phenotype.
113     .collect(toBestPhenotype());
114
115 System.out.println(statistics);
116 System.out.println(best);
117 }
118 }

```

The console out put for the Knapsack GA will look like the listing beneath.

```

1 | +-----+
2 | | Time statistics |
3 | +-----+
4 | |           Selection: sum=0.046538278000 s; mean=0.003878189833 s |
5 | |           Altering: sum=0.086155457000 s; mean=0.007179621417 s |
6 | | Fitness calculation: sum=0.022901606000 s; mean=0.001908467167 s |
7 | | Overall execution: sum=0.147298067000 s; mean=0.012274838917 s |
8 | +-----+
9 | | Evolution statistics |
10 | +-----+
11 | |           Generations: 12 |
12 | |           Altered: sum=7,331; mean=610.916666667 |
13 | |           Killed: sum=0; mean=0.000000000 |
14 | |           Invalids: sum=0; mean=0.000000000 |
15 | +-----+
16 | | Population statistics |
17 | +-----+
18 | |           Age: max=11; mean=1.951000; var=5.545190 |
19 | |           Fitness: |
20 | |               min = 0.000000000000 |
21 | |               max = 481.748227114537 |
22 | |               mean = 384.430345078660 |
23 | |               var = 13006.132537301528 |
24 | +-----+
25 | [01110110|10101111] --> 481.74822711453686

```


7.4 Traveling salesman

The Traveling Salesman problem²⁶ is one of the classical problems in computational mathematics and it is the most notorious NP-complete problem. The goal is to find the shortest distance, or the path, with the least costs, between N different cities. Testing all possible path for N cities would lead to $N!$ checks to find the shortest one.

The following example uses a path where the cities are lying on a circle. That means, the optimal path will be a polygon. This makes it easier to check the quality of the found solution.

```

1 import static java.lang.Math.PI;
2 import static java.lang.Math.abs;
3 import static java.lang.Math.sin;
4 import static org.jenetics.engine.EvolutionResult.toBestPhenotype;
5 import static org.jenetics.engine.limit.bySteadyFitness;
6
7 import java.util.stream.IntStream;
8
9 import org.jenetics.EnumGene;
10 import org.jenetics.Genotype;
11 import org.jenetics.Optimize;
12 import org.jenetics.PartiallyMatchedCrossover;
13 import org.jenetics.PermutationChromosome;
14 import org.jenetics.Phenotype;
15 import org.jenetics.SwapMutator;
16 import org.jenetics.engine.Engine;
17 import org.jenetics.engine.EvolutionStatistics;
18
19 public class TravelingSalesman {
20
21     // Problem initialization:
22     // Calculating the adjacency matrix of the "city" distances.
23
24     private static final int STOPS = 20;
25     private static final double[][] ADJACENCE = matrix(STOPS);
26
27     private static double[][] matrix(int stops) {
28         final double radius = 10.0;
29         double[][] matrix = new double[stops][stops];
30
31         for (int i = 0; i < stops; ++i) {
32             for (int j = 0; j < stops; ++j) {
33                 matrix[i][j] = chord(stops, abs(i - j), radius);
34             }
35         }
36         return matrix;
37     }
38
39     private static double chord(int stops, int i, double r) {
40         return 2.0*r*abs(sin((PI*i)/stops));
41     }
42
43     // Calculate the path length of the current genotype.
44     private static
45     Double dist(final Genotype<EnumGene<Integer>> gt) {
46         // Convert the genotype to the traveling path.
47         final int[] path = gt.getChromosome().toSeq().stream()
48             .mapToInt(EnumGene<Integer>::getAllele)

```

²⁶https://en.wikipedia.org/wiki/Travelling_salesman_problem

```

49     .toArray();
50
51     // Calculate the path distance.
52     return IntStream.range(0, STOPS)
53         .mapToDouble(i ->
54             ADJACENCE[path[i]][path[(i + 1)%STOPS]])
55         .sum();
56 }
57
58 public static void main(String[] args) {
59     final Engine<EnumGene<Integer>, Double> engine = Engine
60         .builder(
61             TravelingSalesman::dist,
62             PermutationChromosome.ofInteger(STOPS))
63         .optimize(Optimize.MINIMUM)
64         .maximalPhenotypeAge(11)
65         .populationSize(500)
66         .alterers(
67             new SwapMutator<>(0.2),
68             new PartiallyMatchedCrossover<>(0.35))
69         .build();
70
71     // Create evolution statistics consumer.
72     final EvolutionStatistics<Double, ?>
73         statistics = EvolutionStatistics.ofNumber();
74
75     final Phenotype<EnumGene<Integer>, Double> best =
76         engine.stream()
77             // Truncate the evolution stream after 15 "steady"
78             // generations.
79             .limit(bySteadyFitness(15))
80             // The evolution will stop after maximal 250
81             // generations.
82             .limit(250)
83             // Update the evaluation statistics after
84             // each generation
85             .peek(statistics)
86             // Collect (reduce) the evolution stream to
87             // its best phenotype.
88             .collect(toBestPhenotype());
89
90     System.out.println(statistics);
91     System.out.println(best);
92 }
93
94 }

```

The Traveling Salesman problem is a very good example which shows you how to solve combinatorial problems with an GA. **Jenetics** contains several classes which will work very well with this kind of problems. Wrapping the base *type* into an **EnumGene** is the first thing to do. In our example, every city has a unique number, that means we are wrapping an **Integer** into an **EnumGene**. Creating a genotype for integer values is very easy with the factory method of the **PermutationChromosome**. For other data types you have to use one of the constructors of the permutation chromosome. As alterers, we are using a swap-mutator and a partially-matched crossover. These alterers guarantees that no invalid solutions are created—every city exists exactly once in the altered chromosomes.

```

1 | +-----+
2 | Time statistics |

```

```

3 |-----+-----|
4 |           Selection: sum=0.072050903000 s; mean=0.000878669549 s |
5 |           Altering: sum=0.253133165000 s; mean=0.003086989817 s |
6 |           Fitness calculation: sum=0.082208402000 s; mean=0.001002541488 s |
7 |           Overall execution: sum=0.412332982000 s; mean=0.005028451000 s |
8 |-----+-----|
9 |           Evolution statistics |
10 |-----+-----|
11 |           Generations: 82 |
12 |           Altered: sum=114,639; mean=1398.036585366 |
13 |           Killed: sum=41; mean=0.500000000 |
14 |           Invalids: sum=0; mean=0.000000000 |
15 |-----+-----|
16 |           Population statistics |
17 |-----+-----|
18 |           Age: max=11; mean=1.593951; var=4.688995 |
19 |           Fitness: |
20 |               min = 68.677087189481 |
21 |               max = 346.356353231264 |
22 |               mean = 147.897084643156 |
23 |               var = 4605.952361305273 |
24 |-----+-----|
25 | [1|2|3|4|5|6|7|8|9|10|11|12|13|14|15|17|16|18|19|0] --> 68.67708718948101

```

The listing above shows the output generated by our example. The last line represents the phenotype of the best solution found by the GA, which represents the traveling path. As you can see, the GA has found the shortest path, in reverse order.

8 Build

For building the **Jenetics** library from source, download the most recent, stable package version from <https://sourceforge.net/projects/jenetics/files/latest/download> or <https://github.com/jenetics/jenetics/releases> and extract it to some build directory.

```
$ unzip jenetics-<version>.zip -d <builddir>
```

<version> denotes the actual **Jenetics** version and <builddir> the actual build directory. Alternatively you can check out the latest version from the Git master branch.

```
$ git clone https://github.com/jenetics/jenetics.git \
  <builddir>
```

Jenetics uses Gradle²⁷ as build system and organizes the source into *sub*-projects (*modules*).²⁸ Each *sub*-project is located in its own *sub*-directory:

- **org.jenetics**: This project contains the source code and tests for the **Jenetics** *core*-module.
- **org.jenetics.example**: This project contains example code for the *ore*-module.
- **org.jenetics.doc**: Contains the *code* of the web-site and *this* manual.

²⁷<http://gradle.org/downloads>

²⁸If you are calling the `gradlew` script (instead of `gradle`), which are part of the downloaded package, the proper Gradle version is automatically downloaded and you don't have to install Gradle explicitly.

For building the library change into the `<builddir>` directory (or one of the *module* directory) and call one of the available *tasks*:

- **compileJava**: Compiles the **Jenetics** sources and copies the class files to the `<builddir>/<module-dir>/build/classes/main` directory.
- **jar**: Compiles the sources and creates the JAR files. The artifacts are copied to the `<builddir>/<module-dir>/build/libs` directory.
- **test**: Compiles and executes the unit tests. The test results are printed onto the console and a test-report, created by *TestNG*, is written to `<builddir>/<module-dir>` directory.
- **javadoc**: Generates the API documentation. The Javadoc is stored in the `<builddir>/<module-dir>/build/docs` directory
- **clean**: Deletes the `<builddir>/build/*` directories and removes all generated artifacts.

For building the library from the source, call

```
$ cd <build-dir>
$ gradle jar
```

or

```
$ ./gradlew jar
```

if you don't have the the Gradle build system installed—calling the the Gradle wrapper script will download all needed files and trigger the build task afterwards.

IDE integration Gradle has tasks which creates the project file for Eclipse²⁹ and IntelliJ IDEA³⁰. Call

```
$ ./gradlew <eclipse|idea>
```

for creating the project files for Eclipse or IntelliJ, respectively.

External library dependencies The following external projects are used for running and/or building the **Jenetics** library.

- **TestNG**
 - **Version:** *6.8.21*
 - **Homepage:** *<http://testng.org/doc/index.html>*
 - **License:** *Apache License, Version 2.0*
 - **Scope:** *test*
- **Apache Commons Math**

²⁹<http://www.eclipse.org/>

³⁰<http://www.jetbrains.com/idea/>

- **Version:** *3.4.1*
- **Homepage:** *<http://commons.apache.org/proper/commons-math/>*
- **Download:** *<http://tweedo.com/mirror/apache/commons/math/binaries/commons-math3-3.4.1-bin.zip>*
- **License:** *Apache License, Version 2.0*
- **Scope:** *test*

- **Java2Html**

- **Version:** *5.0*
- **Homepage:** *<http://www.java2html.de/>*
- **Download:** *http://www.java2html.de/java2html_50.zip*
- **License:** *GPL or CPL1.0*
- **Scope:** *javadoc*

- **Gradle**

- **Version:** *2.3 (or later)*
- **Homepage:** *<http://gradle.org/>*
- **Download:** *<http://services.gradle.org/distributions/gradle-2.2-bin.zip>*
- **License:** *Apache License, Version 2.0*
- **Scope:** *build.*

Maven Central The whole Jenetics package can also be downloaded from the *Maven Central* repository:

```
<dependency>
  <groupId>org.bitbucket.fwilhelm</groupId>
  <artifactId>org.jenetics</artifactId>
  <version>3.1.0</version>
</dependency>
```

9 License

The library itself is licensed under the Apache License, Version 2.0.

Copyright 2007-2015 Franz Wilhelmstötter

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

References

- [1] Thomas Back. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, 1996.
- [2] James E. Baker. Reducing bias and inefficiency in the selection algorithm. *Proceedings of the Second International Conference on Genetic Algorithms and their Application*, pages 14–21, 1987.
- [3] Heiko Bauke. Tina’s random number generator library. <http://numbercrunch.de/trng/trng.pdf>, 2011.
- [4] Tobias Blickle and Lothar Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4:361–394, 1997.
- [5] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution*. Springer, 1996.
- [6] Oracle. Value-based classes. <https://docs.oracle.com/javase/8/docs/api/-java/lang/doc-files/ValueBased.html>.
- [7] Daniel Shiffman. *The Nature of Code*. The Nature of Code, 1 edition, 12 2012.
- [8] S. N. Sivanandam and S. N. Deepa. *Introduction to Genetic Algorithms*. Springer, 2010.
- [9] Eric W. Weisstein. Scalar function. <http://mathworld.wolfram.com/ScalarFunction.html>.
- [10] Eric W. Weisstein. Vector function. <http://mathworld.wolfram.com/VectorFunction.html>.
- [11] Darrell Whitley. A genetic algorithm tutorial. *Statistics and Computing*, 4:65–85, 1994.

Index

- 0/1 Knapsack, 48
- 2-point crossover, 15
- 3-point crossover, 16

- Allele, 5, 29
- Alterer, 12, 31
- Apache Commons Math, 54

- Base classes, 4
- Block splitting, 24
- Boltzmann selector, 11
- Build, 53
 - Gradle, 53
 - gradlew, 53

- Chromosome, 6, 30
 - recombination, 14
 - scalar, 33
- Compile, 54
- Concurrency, 22
- Crossover
 - 2-point crossover, 15
 - 3-point crossover, 16
 - Multiple-point crossover, 14
 - Partially-matched crossover, 14, 16
 - Single-point crossover, 14, 15

- Directed graph, 37
- Domain classes, 5
- Domain model, 5
- Download, 53

- Encoding, 32
 - Affine transformation, 35
 - Directed graph, 37
 - Graph, 36
 - Real function, 33
 - Scalar function, 33
 - Undirected graph, 37
 - Vector function, 34
 - Weighted graph, 37
- Engine, 17, 31
- Engine classes, 15
- Evolution
 - Engine, 17
 - Stream, 15, 19
- Evolution time, 42
- EvolutionResult, 20

- EvolutionStatistics, 21
- EvolutionStream, 19
- Examples, 44
 - 0/1 Knapsack, 48
 - Ones counting, 44
 - Real function, 46
 - Traveling salesman, 51
- Exponential-rank selector, 11

- Fitness function, 16
- Fitness scaler, 17
- Fitness threshold, 43
- Fixed generation, 40

- Gaussian mutator, 13
- Gene, 5, 29
 - validation, 6
- Genetic algorithm, 3
- Genotype, 6, 7
 - scalar, 33
 - Validation, 19
- Git repository, 53
- Gradle, 53, 55
- gradlew, 53
- Graph, 36

- Hello World, 1

- Installation, 53

- Java2Html, 55

- LCG64ShiftRandom, 25
- Leapfrog, 25
- License, i, 55
- Linear-rank selector, 11

- Monte Carlo selector, 10
- Multiple-point crossover, 14
- Mutation, 12
- Mutator, 13

- Ones counting, 44
- Operation classes, 8

- Package structure, 4
- Partially-matched crossover, 14, 16
- Phenotype, 7

- Validation, 19
- Population, 8
- PRNG, 23
 - Block splitting, 24
 - LCG64ShiftRandom, 25
 - Leapfrog, 25
 - Parameterization, 24
 - Performance, 26
 - Random seeding, 24
- Probability selector, 10

- Random, 23
 - Engine, 23
 - LCG64ShiftRandom, 25
 - Registry, 23
- Random seeding, 24
- Randomness, 23
- Real function, 46
- Recombination, 13
- Roulette-wheel selector, 10

- Scalar chromosome, 33
- Scalar genotype, 33
- Selector, 8, 30
- Seq, 28
- Serialization, 26
- Single-point crossover, 14, 15
- Source code, 53
- Statistics, 28, 31
- Steady fitness, 41
- Stochastic-universal selector, 11
- Swap mutator, 13

- Termination, 40
 - Evolution time, 42
 - Fitness threshold, 43
 - Fixed generation, 40
 - Steady fitness, 41
- TestNG, 54
- Tournament selector, 9
- Traveling salesman, 51
- Truncation selector, 9

- Undirected graph, 37

- Validation, 6, 19, 38

- Weighted graph, 37