SodaPop

User Manual v1.0

Louis Gauthier
louis.gauthier@umontreal.ca

Adrian W.R. Serohijos
adrian.serohijos@umontreal.ca

Département de Biochimie,
Université de Montréal,
Québec, Canada

This manual accompanies the following mansucript:
L Gauthier, R Di Franco, A Serohijos, **SodaPop: A Forward Simulation Suite for the Evolutionary Dynamics of Asexual Populations on Protein Fitness Landscapes** *Submitted to Bioinformatics*. Also available in *BioRxiv, 189142 (2017)*

Copyright © 2018 Louis Gauthier, Adrian Serohijos

HTTPS://WWW.SEROHIJOSLAB.ORG/SOFTWARE-TOOLS–DATA.HTML

SodaPop is free software. You can redistribute or modify it under the terms of the GNU General Public License published by the Free Software Foundation, either version 3 or later version. SodaPop is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details. You should have received a copy of the GNU General Public License along with SodaPop. If not, see http://www.gnu.org/licenses/.

Last revised 13 April 2018, for SodaPop v1.0

Contents

1	Introduction	5
2	Setup and installation	6
2.1	Downloading SodaPop	6
2.1.1	Cloning via Git	6
2.1.2	Zip archive	6
2.2	Installation	6
3	List of command-line options	8
3.1	Numerical parameters	8
3.2	Input files	8
3.3	Fitness landscape	9
3.4	Analysis and output	10
4	Quick Start	11
4.1	Program workflow	11
4.2	Creating a starting population snapshot	12
4.3	Running a neutral simulation	14
4.4	Folding stability simulation example	15
4.5	Sampled selection coefficients example	15
4.6	Converting binary snapshots to text	16
5	Fitness landscapes	17
5.1	Phenomenological distribution	17
5.2	Computational tools in biophysics	17
5.3	Experimentally derived fitness effects	18
6	Results and analysis	19
6.1	Scripts	19
6.2	Mean fitness plot	19
6.3	Clonal structure	20
6.4	Clonal trajectories	20
6.5	Mutation log	21

6.6 Command log 21

7 Fitness functions 24
7.1 Folding stability 24
7.2 Metabolic flux 24
7.3 Toxicity 24
7.4 Combined metabolic output 25
7.5 Neutral 25
7.6 Multiplicative 25
7.7 No mutation 25

8 Modifying the code 26
8.1 Implementing a fitness function 26
8.2 Adding gene properties 26

9 Implementation and performance 29
9.1 Data structures and complexity 29
9.2 Pseudo-random number generation 29
9.3 Runtime 30

Acknowledgements 31

Annex: Using the command-line 32

Bibliography 33

1. Introduction

Overview

Protein evolution is determined by constraints at several levels of biological organization. Random mutations have an immediate effect on the biophysical properties, structure and function of proteins. These same mutations also affect the fitness of the organism. However, the evolutionary fate of these mutations, whether they succeed to fixation or are purged, also depends on population size and dynamics. There is an emerging interest, both theoretically and experimentally, to integrate these two constraints in protein evolution. Although there are several tools available for simulating protein evolution, most of them focus on either the biophysical or the population-level constraints, but not both.

SodaPop is a new computational suite to simulate protein evolution in the context of the population dynamics of asexual populations. SodaPop accepts as input several fitness landscapes based on protein biochemistry or other user-defined fitness function. The user can also provide as input experimental fitness landscapes derived from deep mutational scanning approaches or theoretical landscapes derived from physical force field estimates. SodaPop is designed such that population geneticists can explore the role of protein biochemistry on genetic variation, and that biochemists and biophysicists can explore the role of population size and demography on protein evolution. This highly flexible tool takes as input a fitness landscape based on protein biochemistry and biophysics, experimental deep mutational scans or a phenomenological distribution of fitness effects. SodaPop was designed such that population geneticists can explore the role of protein biochemistry on genetic variation, and that biochemists and biophysicists can explore the role of population size and demography on protein evolution.

The main software is implemented in C++ and supported on Linux, Mac OS/X and Windows. Source code and binaries are freely available at https://github.com/louisgt/SodaPop under the GNU GPLv3 license.

2. Setup and installation

This section will take you through the installation.

> **Note** This program and the simulation examples were developed for a Unix environment. For use in Windows, we strongly recommend installing a Unix-like environment such as Cygwin or MinGW.

SodaPop is a suite of command-line tools written in the C++ programming language using the C++11 Standard. The source code and compiled binaries can be downloaded from https://github.com/louisgt/SodaPop.

Users unfamiliar with command-line interface are referred to the *Annex: Using the command-line* found at the end of this manual.

2.1 Downloading SodaPop

Directly download SodaPop using the Git command. Alternatively, you can download its zip archive from https://github.com/louisgt/SodaPop.

2.1.1 Cloning via Git

Use the Git command to clone the original SodaPop repository to a local folder with

```
git clone https://github.com/louisgt/SodaPop.git my_folder
```

2.1.2 Zip archive

1. Go to https://github.com/louisgt/SodaPop. Click on the *Clone or download* button and select the *Download ZIP* option.
2. Move the archive to your preferred location and unzip its contents with

```
tar -zxvf [name of the zip file]
```

2.2 Installation

SodaPop comes with compiled binaries (executables). However, you are advised to compile the program on your system.

To compile all components of SodaPop, run

```
make
```

This uses the makefile to build the binaries **sodapop**, **sodasnap** and **sodasumm**, which can be executed from the folder they are located in. If you wish to install the program on your computer, run

```
make install
```

By default, the three binaries above will be appended to `/usr/local/bin`. You can change this option in the makefile by editing the content of the `$INSTALLDIR` variable.

To verify the installation step, run `make` again and the terminal will read

```
make:  Nothing to be done for ‘all’.
```

If the following error arises when compiling

```
error:  unrecognized command line option "-std=c++11"
```

your compiler is likely out-of-date. You can get a new version of gcc/g++ here. SodaPop is compatible with any version starting from gcc 4.7 or higherr.

3. List of command-line options

Command-line flags and parameters are used to toggle among the features, change program behavior and provide input parameters and files for a simulation. To view the available flags and their description, type

```
sodapop -help
```

Flags can also be added or modified in the *evolve.cpp* source file, using the TCLAP library API.

3.1 Numerical parameters

Number of generations `[-m, --maxgen]`
(Integer) The length of the simulation in number of generations.
(Default value) 10.

Population size cutoff `[-n, --size]`
(Integer) The target population size in each generation or the carrying capacity.
(Default value) 1.

Snapshot interval `[-t, --dt]`
(Integer) The time interval at which a population snapshot is saved. The minimum value is 1 and there is no maximum value. However, if *dt* exceeds the specified length of the simulation, only one snapshot is saved.
(Default value) 1.

3.2 Input files

Gene list file `[-g, --gene-list]`
(File) A text file that lists genes and the files that describe them (see **Section 4.2** and **Fig. 4.4**)

Population description file `[-p, --pop-desc]`
(File) The binary snapshot required to initialize the population in the simulation (see **Section 4.2**).

Path to gene library `[-l, --gene-lib]`
(Integer) The path to the folder containing the gene files.
(Default value) `/files/genes/`.

3.3 Fitness landscape

Input fitness landscape `[-i, --input]`

(File) This flag specifies the file(s) containing the effects of nonsynonymous mutations for all genes in the simulation (See section 5). For a given gene and corresponding fitness function, you can specify one or more matrices $M_{a,k}$ describing the change in fitness for a mutation to amino acid a in residue k. The index a is over all 20 amino acid listed in alphabetical order (from 'A' to 'Y'). In the file, this matrix is described by k rows of 20 values. The mutational matrices of all the genes are ordered in the file according to their gene ID.

By default, the values of in the matrix are assumed to be fitness effects (selection coefficients) from deep mutational scan (DMS) experiments. However, these matrices can contain any property of interest. See **Section 7** on defining fitness functions for examples.

Fitness function `[-f, --fitness]`

(Integer) Used to select the appropriate fitness function depending on the input type and the biophysical or biochemical property of interest. For a detailed reference of available functions, refer to **Chapter 7**.

1 Folding stability.
2 Metabolic flux.
3 Misfolding toxicity.
4 Combined metabolic output.
5 Multiplicative.
6 Neutral.
7 No mutations.

Simulation type `[--sim-type]`

(String) This flag is used to specify if mutation effects are interpreted directly as selection coefficients, or if they are effects on biophysical parameters such as stability or activity. By default, values are interpreted as selection coefficients. Alternatively, the type can be set to thermodynamic stability (ΔG), k_{cat}/K_M or another property implemented by the user.

s Selection coefficients.
stability Thermodynamic stability.
efficiency Protein catalytic efficiency.

Use gamma distribution `[--gamma]`

(Flag) This flag will draw the fitness effects from a gamma distribution. The shape and scale parameters are specified by the alpha and beta flags, respectively.

Use normal distribution `[--normal]`

(Flag) This flag will draw fitness effects from a normal distribution. The mean and standard deviation are specified with the alpha and beta flags, respectively.

Alpha parameter `[--alpha]`
(Double) Value of the alpha parameter. If gamma distribution is used, this corresponds to the shape parameter. If normal distribution is used, this corresponds to the mean (μ).

Beta parameter `[--beta]`
(Double) Value of the alpha parameter. If gamma distribution is used, this corresponds to the scale parameter. If normal distribution is used, this corresponds to the standard deviation (σ).

3.4 Analysis and output

Output prefix `[-o, --prefix]`
(String) The name of the output directory to use for the simulation. If the specified directory does not exist, it will be created. If it exists, its current contents are overwritten. This parameter is not required, but strongly recommended since the default setting will create a directory named `sim/` and could overwrite a previous run. You should always specify a new output prefix if working with multiple runs.

Initialize population from single cell `[-c, --create-single]`
(Flag) Using this flag will populate the initial cell vector with copies of a single cell. This can speed up the initialization for large populations (10^6 or higher), at the expense of no clonal diversity. If your starting population is polyclonal, you should not use this flag.

Toggle automatic analysis `[-a, --analyze]`
(Flag) Using this flag will call analysis scripts after the simulation algorithm is complete. Refer to **Chapter 6** for detailed information.

Track mutations `[-e, --track-events]`
(Flag) Using this flag will save a log file with all arising mutations during the course of the simulation (see **Section 6.4**).

Output format `[-s, --seq-output]`
(Flag) Using this flag you can specify the output format of population snapshots. If you do not require sequence information, use the short format as input/output operations will be lighter. If you need sequence data as DNA instead of protein, use the long format with DNA. See **Section 6.5** for details on output formats.
Default Long format.

0 Long format.
1 Short format.
2 Long format with DNA sequence.

4. Quick Start

This chapter introduces users to SodaPop and its applications.

4.1 Program workflow

SodaPop's simulation algorithm is adapted from a Wright-Fisher model with selection (**Fig. 4.1**). Generations are discrete steps in which all N parent cells within the population give birth to a number of daughter cells proportional to their relative fitness. The number of progeny k is drawn from a binomial distribution with N trials and mean w, which is the the fitness of the parent cell normalized by the sum of all cell fitnesses. The offsprings become the parents for the next generation.

Initializing a simulation with SodaPop requires a combination of input files with a specific hierarchy (**Fig. 4.2**). The next section will take you through the definition and function of these files.

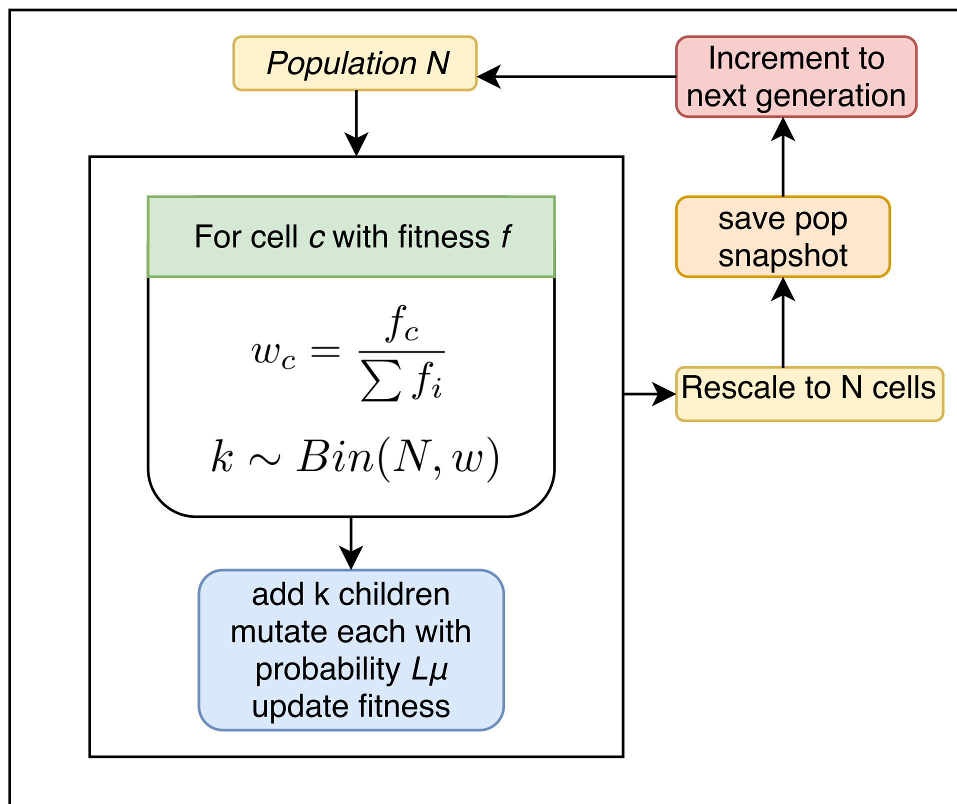

Figure 4.1: Illustration of SodaPop's evolutionary algorithm.

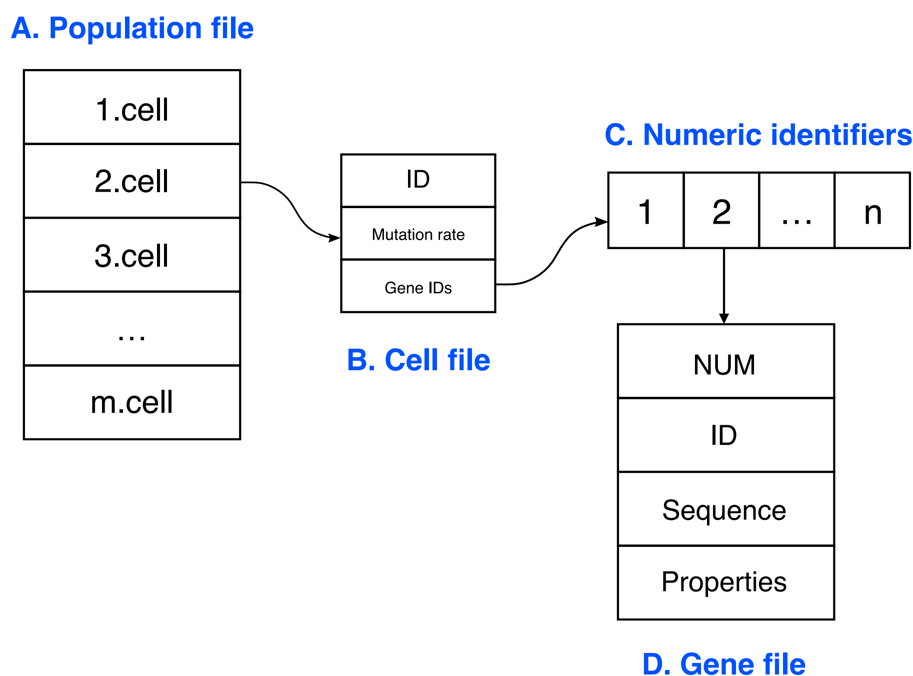

Figure 4.2: Illustration of SodaPop's input file hierarchy. **(A)** The population file indicates the clonal composition of the starting population. This is defined by the number of each different clone to create. These cells are defined individually in cell files **(B)** that contain a numeric identifier, a mutation rate and a list of gene identifiers. These numeric identifiers are indexed in a gene list **(C)** which maps to gene files **(D)** that contain gene sequence and information.

4.2 Creating a starting population snapshot

Prior to running a simulation, the first thing to do is create a population snapshot. In SodaPop, a population is defined by a collection of cells, that are in turned defined by their genes. We will thus start by creating a gene file. An example gene file is shown in **Fig. 4.3**.

It has a numeric identifier as well as an alphanumeric ID. The numeric identifier is important for the proper mapping of primordial sequences. Genes are thus saved as [numericID].*gene*. You can keep an index of the alphanumeric ID corresponding to each gene in the gene list (**Fig. 4.4**). Genes have an explicit nucleotide sequence and the corresponding translated amino acid sequence. The nucleotide sequence should be codon multiples of 3 each corresponding to the protein sequence. There should be no stop codons in the sequence. The following rows define properties of the protein. **E** is the protein's essentiality. DG is the Gibbs folding free energy (ΔG, in $kcal/mol$) or folding stability of the corresponding protein. **CONC** is the intracellular concentration of the protein. **F** is the protein's fitness. **EFF** is the catalytic efficiency of the protein.

The file gene_list.dat keeps an index of all the genes to be used in a simulation. An example gene list is shown in **Fig. 4.4**. You can define as many genes and gene lists as necessary. However, make sure you use the correct list when you run a simulation. If there is a mismatch between the genes of your initial population and the genes in your index list, the program will exit and issue an error.

Now that we have created a gene, we need to define a cell and its properties. Cells possess

```
1 Gene_NUM    0
2 Gene_ID     DHFR
3 A_Seq       MLKPNVAIIVAALKPALGIGYKGKMPWRLRKEIRYFKDVTTRTTK
  PNTRNAVIMGRKTWESIPQKFRPLPDRLNIILSRSYENEIIDDNIIHASSIESSLNL
  VSDVERVFIIGGAEIYNELINNSLVSHLLITEIEHPSPESIEMDTFLKFPLESWTKQ
  PKSELQKFVGDTVLEDDIKEGDFTYNYTLWTRK
4 N_Seq       ATGCTGAAACCAAATGTTGCAATCATTGTCGCGGCATTAAAGCCT
  GCTTTGGGAATTGGATACAAAGGTAAAATGCCTTGGAGACTCCGTAAGGAAATAAGA
  TATTTCAAAGACGTCACCACCAGAACAACGAAACCAAATACTCGTAATGCTGTTATT
  ATGGGAAGAAAGACGTGGGAATCCATACCACAGAAGTTTAGACCTCTTCCAGATAGA
  TTAAACATAATATTATCTAGATCATACGAAAATGAAATTATTGATGACAATATCATT
  CATGCCAGCTCAATTGAGTCGTCACTTAATCTTGTATCTGATGTCGAAAGAGTTTTC
  ATAATTGGAGGGGCAGAAATTTACAATGAATTGATAAATAATTCTTTAGTGAGTCAT
  CTATTAATTACAGAAATCGAACATCCAAGCCCAGAGTCTATTGAAATGGATACCTTT
  TTGAAATTCCCATTGGAAAGTTGGACCAAACAACCAAAATCAGAGTTGCAGAAATTT
  GTTGGAGATACTGTATTAGAAGACGATATCAAGGAAGGTGATTTTACCTATAATTAT
  ACGCTATGGACAAGAAAA
5 E        1
6 EFF      1
7 DG       -3.0
8 CONC             2250
```

Figure 4.3: Example of *.gene* file description.

individual properties such as mutation rate and an organism identifier. The *.cell* file lists all the identifiers of genes in a cell, preceded by a **G**. Again, genes are indexed by their numeric identifier. The same goes for cells. An example cell file is shown in **Fig. 4.5**.

```
 1 Gene_Count  10
 2 G   0.gene  1AI9_A_D1
 3 G   1.gene  1SQL_A
 4 G   2.gene  1K0E_A
 5 G   3.gene  1I2K
 6 G   4.gene  1I2K_D1
 7 G   5.gene  2BMB_A_D2
 8 G   6.gene  2BMB_A_D3
 9 G   7.gene  1O5Z_A_D1
10 G   8.gene  1O5Z_A_D2
11 G   9.gene  1AI9_A_D1
```

Figure 4.4: Example of *gene_list* file description.

Once a cell has been defined, we can create a starting population snapshot. The last layer in the hierarchy is the population description file (**Fig. 4.6**). It lists the composition of the initial population we wish to create. In our case, we will define a polyclonal population with three different clones. The template shows how to define this population (**Fig. 4.6**). The count is the size of the starting subpopulation for each clone. The comment is optional and is ignored by the program.

To recapitulate, we must first create a gene file (**Fig. 4.2D**). We then define a cell file to include our gene(s) (**Fig. 4.2B**). Finally, we create a population summary defining the initial clonal structure

```
1  org_id  1
2  mrate   1.53116e-6
3  G   0
4  G   1
5  G   2
6  G   3
7  G   4
8  G   5
9  G   6
10 G   7
11 G   8
12 G   9
```

Figure 4.5: Example of *.cell* file description. This cell holds 10 different genes.

```
1 //Population summary
2 //  Count   File
3 C   10000   files/start/1.cell
4 C   5000    files/start/2.cell
5 C   10000   files/start/3.cell
```

Figure 4.6: Example of *population* file description. The file defines a polyclonal population with 3 different clones: 10^4 clones of cell 1, $5 \cdot 10^3$ clones of cell 2 and 10^4 clones of cell 3.

we want (**Fig. 4.2A**).

We will use the program called **sodasumm** to create a starting population snapshot. If you run the program without any flag or argument, it will display the options:

```
sodasumm <population summary> [ 0-full | 1-single cell | 2-randomize
barcodes | 3-introduce variation]
```

As input, we need the population summary created above (**Fig. 4.6**), and we must specify whether we want to build the full population with clones sharing the same barcode **(option 0)**, a randomized population from the first cell in the file **(option 1)**, the full population with each cell having a randomized unique barcode **(option 2)**, or the full population with clones sharing the same barcode, with fitness effects randomly introduced in clusters of cells **(option 3)**. After running **sodasumm**, a *.snap* file will be created in the same directory and with the same prefix as the population description file. This is your starting population snapshot. You can rename it, but keep the *.snap* extension to distinguish between different input files.

4.3 Running a neutral simulation

Using the created population snapshot, we can run a very simple neutral simulation, that is, without selection. We can do this using the sample files packaged with the program. Try running the following command:

```
sodapop --sim-type stability -f 6 -p files/start/pop1K.snap
  -g files/genes/gene_list.dat -o neutral -n 1000 -m 1000 -t 10
```

The command-line display will read:

```
 Begin ...
Initializing matrix ...
Loading primordial genes file ...
Opening starting population snapshot ...
Creating directory out/neutral/snapshots ...  OK
Constructing population from source files/start/pop1K.snap ...
Saving initial population snapshot ...
Starting evolution ...
Done.
Total number of mutation events:  10767
```

Breaking down this command step by step:

- `sodapop` is the main simulation utility (or program).
- `--sim-type stability` defines the type of input to be used for the simulation. In our case, we chose stability.
- `-f 6` is the choice of fitness function. For more information on fitness functions, see **Chapter 7**. For now, just know that 6 is the mapping for the neutral fitness function.
- `-p files/start/pop1K.snap` is the the starting population snapshot file.
- `-g files/genes/gene_list.dat` is the gene list file.
- `-o neutral` is the prefix for output.
- `-n 1000` is the population size.
- `-m 1000` is the number of generations we want to simulate.
- `-t 10` is the interval at which the program saves a population snapshot.

4.4 Folding stability simulation example

In the `files/` folder of the SodaPop archive, you will find examples and data to help you get started. Here, we will use some of these files to run a simple simulation based on protein folding stability of the dihydrofolate reductase enzyme (DHFR) which is encoded by the *folA* gene in *E. coli*.

The file *DHFR_DDG.matrix* is the fitness landscape of the protein folding stability of *E. coli* DHFR. We will use it as input, along with a starting population snapshot of $N_e = 10^4$ initially monoclonal cells. We will use the first fitness function, based on metabolic flux (see **Chapter 7**). The following command will start a simulation and direct output to the directory `out/foldsim/`:

```
sodapop --sim-type stability -i files/start/DHFR_DDG.matrix
  -f 1 -p files/start/DHFR_start.snap -g files/genes/gene_list.dat
  -o foldsim -t 100 -n 10000 -m 5000 -a
```

Refer to **Chapter 6** for detailed information on simulation output.

4.5 Sampled selection coefficients example

Here, we will use the multiplicative fitness function (**-f 5**) and sample selection coefficients from a gaussian distribution with $\mu = -0.02$ and $\sigma = 0.01$:

```
sodapop --sim-type s --normal --alpha -0.02 --beta 0.01
  -p files/start/population.snap -g files/genes/gene_list.dat
  -o codesim -t 100 -n 10000 -m 1500 -a
```

Note that the multiplicative fitness function is implicit here because we are working with selection coefficients. It is thus not required to specify it.

4.6 Converting binary snapshots to text

sodasumm and **sodapop** both create snapshots in binary format (.snap). To manually convert binary files to text, use **sodasnap**:

```
sodasnap <snap-binary> <out-ascii>
```

Where `<snap-binary>` is the binary snapshot to convert and `<out-ascii>` is the name of the text file to create. Alternatively, adding the `-a` flag to your **sodapop** command will call analysis scripts once the simulation is done. This includes automatically converting your snapshots to text, extracting barcodes and generating basic plots. Refer to **Chapter 6** for more details.

5. Fitness landscapes

When a random mutation occurs in SodaPop, its effect on organismal fitness depends on the fitness landscape chosen by the user. The following sections describe the three models of fitness landscape that can be used in the program.

5.1 Phenomenological distribution

The effects of random mutations on organismal fitness can come from user-defined distributions. In the current version of SodaPop, the DFE can be modelled by a two-parameter distribution of either normal form $N(\mu,\sigma)$ or gamma form $\Gamma(\alpha,\beta)$ using the `--normal` and `--gamma` flags. The `--alpha` and `--beta` flags are used to specify μ and σ or α and β, respectively. These flags are used in tandem with `--sim-type` to determine the nature of the drawn values. For instance, the command

```
sodapop --sim-type s --normal --alpha -0.02 --beta 0.01
  -p files/start/population.snap -g files/genes/gene_list.dat
  -o codesim -t 100 -n 10000 -m 1500 -a
```

will draw selection coefficients from a distribution $N(-0.02,0.01)$.

However, the command

```
sodapop --sim-type stability --normal --alpha 1 --beta 1 -f 1
  -p files/start/population.snap -g files/genes/gene_list.dat
  -o codesim -t 100 -n 10000 -m 1500 -a
```

will draw $\Delta\Delta G$ values from a distribution $N(1,1)$ and calculate fitness effects based on the resulting folding stability, with all other things kept equal. The same distinction is applied to the `--gamma` option.

5.2 Computational tools in biophysics

SodaPop can take as input one or more substitution matrices that describe the molecular effects of mutations on a particular protein property for all 1-away mutants (**Fig. 5.1**). The entries in these tables can be derived from computational protein engineering tools such as Rosetta [1], Eris [5] or FoldX [3]. Each matrix must be tab-separated and preceded by a header line containing the prefix "Gene_NUM", the gene index matching the corresponding gene file and the protein length in amino acids. Then, each row in the matrix must start by indicating the property (e.g. "DDG"), the residue position (starting at 1), followed by a list of values for each amino acid substitution, including the identity substitution. If you are simulating several proteins at once, each with a different matrix, you must list the matrices one after the other in a single file. To input this file in a simulation, you must use the `-i` flag **Section 3.3**.

```
Gene_NUM 0 192
DDG 1	-0.01400	8.33133	-0.67133	-0.40333	0.34067	1.98400	0.74733	-0.63400	-1.84533	-1.62733	0.00000	0.00733	10.50000	-1.56800	0.72133	-0.48933	-1.76800	-0.54933	2.06733	0.82933
DDG 2	-0.32400	8.24733	-0.92600	-0.17933	1.74867	0.75400	2.02400	0.58200	-0.03667	0.00000	1.28333	0.63067	0.00867	0.47267	0.99067	0.14667	-0.32200	-0.55067	3.50600	1.32000
DDG 3	0.72600	8.69400	0.12467	-0.35133	2.53400	0.92333	0.96400	-0.13800	0.00000	0.62400	1.05600	-0.07867	-0.32000	0.23667	0.70200	0.82933	0.37867	-0.29267	4.00933	1.53867
DDG 4	1.37000	10.50000	1.37600	0.57400	6.71467	1.55533	4.42133	1.81133	4.77200	9.31467	2.78933	1.69133	0.00000	2.23000	4.05467	0.71667	0.43667	1.28933	0.02867	10.50000
DDG 5	-0.41933	7.53800	-0.51600	-1.81600	1.05200	0.11667	2.08133	0.13067	-0.51600	0.23733	0.25667	0.00000	-0.03933	0.20600	0.70533	0.08400	-0.03000	0.75533	3.05933	2.14267
DDG 6	2.29800	10.50000	3.57200	4.38933	9.44333	4.49067	6.02133	-0.26533	6.96667	0.64333	3.50333	2.82533	0.36667	3.85067	10.50000	3.12467	2.47000	0.00000	10.50000	10.50000
DDG 7	0.00000	8.49267	4.00200	0.57733	5.86333	2.42533	3.71200	-1.21533	0.78467	4.90667	1.44733	1.98267	-0.00000	0.52067	5.43467	0.78667	-0.72067	-1.71600	5.86067	6.51133
DDG 8	3.70800	10.50000	5.06800	3.08533	8.14800	5.16600	7.03533	0.00000	4.86067	1.31267	2.41400	3.43133	-0.18333	2.50000	8.10067	4.15733	2.72467	1.11867	10.50000	10.50000
DDG 9	3.24600	10.50000	3.70067	3.54000	3.07533	4.72667	2.35333	0.00000	3.95333	0.71133	1.69267	2.23533	-0.23667	2.49333	5.77600	3.31267	2.58600	0.66867	6.63867	3.19800
DDG 10	0.69000	10.50000	2.29333	3.79600	7.80133	3.15733	5.47533	2.41800	7.53467	0.69400	3.90733	1.16600	0.16267	3.14200	10.50000	2.25200	0.30333	0.00000	10.50000	10.50000
DDG 11	0.00000	9.28600	5.70600	5.64600	10.50000	1.79667	10.50000	3.55333	7.49533	5.30867	9.59933	3.35267	0.06133	3.79533	10.50000	1.76133	1.78867	2.10733	10.50000	10.50000
DDG 12	0.00000	10.50000	3.92333	6.59533	10.50000	1.33000	7.94000	1.93267	9.43533	1.35800	4.11800	3.03800	0.02067	6.30133	10.50000	1.23067	2.30800	0.95600	10.50000	10.50000
DDG 13	2.15867	10.50000	2.28733	1.68067	1.61867	4.11800	0.89667	0.09800	2.39733	0.00000	0.67067	2.05267	7.75800	-0.35600	3.90400	2.34867	1.22067	0.55267	2.98867	6.78333
DDG 14	2.52333	10.50000	3.99733	1.86200	3.21267	4.14467	2.94800	2.15133	0.00000	1.30667	2.65733	3.12000	9.84133	2.10467	2.18067	1.36800	1.38000	2.16933	4.11667	4.82600
DDG 15	-0.25133	7.93800	-0.92467	-1.07200	1.90867	-0.27733	1.08667	0.81867	-0.82600	-0.69533	0.78533	-0.21400	0.00000	-0.08133	0.93533	-0.12933	-0.27600	-0.01733	2.97867	1.48333
DDG 16	0.00000	7.90000	-0.59733	-0.92733	0.84533	0.47400	1.03467	0.31067	-0.38933	-0.42733	0.08800	-0.49333	2.81867	-0.54467	-0.63200	-0.01067	-0.46800	-0.23200	2.01067	1.08600
DDG 17	1.88800	9.67733	1.37800	0.87733	1.00533	1.24800	-0.62667	0.06733	1.20533	0.00000	-0.32333	0.59067	0.24800	0.51933	0.67667	2.49133	1.49000	-0.22000	-0.82467	0.39133
DDG 18	-0.36667	10.50000	6.48267	6.58533	5.65267	0.00000	5.98733	10.50000	4.50000	5.09467	4.37533	4.96867	5.17000	6.13600	5.94533	1.51600	5.29533	6.22733	6.41400	5.31400
DDG 19	1.82333	10.50000	3.60333	2.33333	10.50000	2.70067	1.72333	0.00000	5.55200	2.82133	2.76000	3.38800	-0.95933	5.26333	7.88333	2.36133	0.89400	0.00600	10.50000	1.33867
DDG 20	0.24267	9.58800	2.21133	2.19867	10.50000	0.00000	10.50000	-0.06200	2.51733	3.56933	1.91600	6.49667	-0.44733	0.18133	2.77267	0.96000	3.59400	-0.52800	-0.49733	10.50000
DDG 21	2.65600	10.50000	2.06467	1.66533	2.44333	2.93000	1.65533	1.90933	1.75933	1.89867	3.20000	0.92667	-0.13600	1.98067	1.29333	2.64000	2.47933	1.23133	3.67667	0.00000
DDG 22	0.55933	8.61133	-0.63000	-0.80267	2.39733	0.00267	1.67267	0.86600	0.00000	0.35200	1.27867	0.80533	4.78933	0.79867	0.69000	0.34533	1.33000	0.88867	-0.08600	2.82533
DDG 23	0.12000	8.04333	-0.01667	-0.69133	1.54067	0.00000	0.86400	0.87667	0.28733	1.24067	1.54267	-1.64533	4.00067	0.14733	-0.15533	0.41933	0.75267	0.33000	1.97467	1.10600
DDG 24	-0.22800	7.12733	-1.02800	-0.57333	1.62533	-0.31200	0.66600	-0.22867	0.00000	0.03667	0.39800	-0.12533	-0.05733	-0.14200	0.42733	-0.68733	-1.03800	-0.18133	2.48800	0.70467
DDG 25	-0.37600	7.20933	0.18067	-1.23533	0.41400	0.65600	-0.09600	5.38333	-0.06200	-0.56867	0.00000	-0.57067	-0.40133	-0.54400	-0.45867	-0.31400	0.11600	4.47800	0.81067	0.16467
DDG 26	-0.77800	8.60667	0.98333	-0.48200	6.26667	0.37267	2.84800	1.64400	-0.38733	1.41667	-0.98600	0.30733	0.00000	-1.14800	-0.06267	0.92333	0.93333	0.81200	-0.06133	7.70067
DDG 27	2.42933	10.50000	3.76733	3.66533	0.34733	4.48667	1.55400	0.38600	4.85867	2.38667	1.89200	1.99000	0.61867	3.34067	4.35333	2.61933	2.12267	3.14067	0.00000	3.85467
DDG 28	-1.42467	6.31867	-1.49000	-1.47800	-0.29200	-1.77133	-1.03067	-1.17400	-1.18133	-1.79933	-0.49800	-1.33933	-0.13333	-1.03800	0.00000	-1.24533	-0.88600	-1.15800	1.39400	0.05733
DDG 29	3.24600	10.50000	3.21667	3.21267	3.27067	4.96200	1.56400	6.30067	5.81200	0.00000	3.28467	1.99333	-0.71867	3.14133	10.50000	4.17933	4.88400	7.42467	10.50000	10.50000
DDG 30	-1.69933	7.36133	-1.24600	-2.85133	0.89200	-0.61733	-0.42533	-0.39000	-0.63200	-0.59200	0.00200	-1.07067	7.53200	-0.99067	0.00000	-0.35600	-0.48333	-1.11133	1.97733	0.70800
DDG 31	1.89867	10.50000	1.56267	2.69733	3.18400	1.22400	2.76000	1.84267	0.00000	2.25667	1.30333	2.24200	10.50000	1.22133	2.11600	3.85467	2.17133	1.22267	3.80600	2.54733
DDG 32	0.37667	10.50000	0.67067	0.00000	-0.68600	2.46000	-0.66467	4.47267	0.99067	-1.50467	-0.08933	-0.33867	10.50000	-0.02667	-0.30467	1.57133	1.77267	1.12867	-1.13800	-0.73733
DDG 33	0.72667	9.58467	0.99000	0.75800	1.44267	1.22733	0.69400	0.00000	0.41800	0.35133	1.41400	0.06400	10.50000	0.76667	1.23400	0.59467	-0.26733	1.09133	2.45467	1.11400
DDG 34	-0.86467	7.72067	-0.48067	-1.51800	1.30267	0.10600	0.18800	-0.53467	-0.65333	0.54667	0.05133	-0.97000	6.27867	-0.91733	0.00000	-0.95000	-1.13067	-0.92867	1.29267	0.82667
DDG 35	4.56067	10.50000	5.85400	5.06600	1.33533	6.37933	3.11867	6.18133	7.44667	9.53533	4.23200	5.22867	10.50000	5.68333	10.50000	4.05733	2.64600	5.80667	1.27933	0.00000
DDG 36	0.26467	8.92533	0.06200	-0.05467	0.00000	0.72333	0.35267	0.68667	0.15533	-0.76867	0.49400	0.00800	10.50000	-0.12733	0.15333	-0.81867	-1.43267	-0.10200	-0.09000	0.20200
DDG 37	-0.41600	8.51867	0.20267	-0.33733	0.90600	0.40400	0.76667	0.21133	0.00000	0.22667	-0.02200	-0.55267	10.50000	-0.67667	0.00000	-1.14600	-1.62533	-0.11067	1.16267	0.52400
DDG 38	0.46200	9.98400	0.00000	-1.16000	4.52667	2.27600	1.47667	2.21733	1.35600	0.76267	1.28067	0.49267	10.50000	0.08667	1.33533	1.37667	1.01733	2.47400	4.31200	8.44333
DDG 39	1.77267	10.50000	4.43533	2.51667	10.50000	3.55267	7.20267	1.81267	5.50667	2.36800	3.38533	3.81067	10.50000	3.41800	10.50000	2.46400	0.46067	0.00000	10.50000	10.50000
DDG 40	3.13467	10.50000	3.76333	7.82267	10.50000	4.78000	4.43200	6.54333	9.69267	2.42133	6.28800	3.03733	10.50000	9.07733	10.50000	3.06733	0.00000	4.34533	10.50000	10.50000
DDG 41	0.16467	8.41933	1.36333	0.40800	1.87133	0.43400	1.36733	0.48067	-0.31667	1.49600	0.29000	0.20200	10.50000	0.71800	0.78733	-0.31133	0.00000	0.38133	4.41200	1.96667
DDG 42	-0.93267	7.03133	-0.58533	-0.87000	0.39067	-0.22733	-0.35067	-0.80200	-1.71800	-0.19800	0.55800	-0.08667	9.98133	-1.37667	0.00000	-0.98467	-1.64533	-0.38733	1.06867	0.11133
DDG 43	1.29600	9.80400	2.19667	1.93000	2.63533	1.08133	1.50267	2.82933	10.50000	5.35133	5.68667	2.23733	0.71000	2.91800	6.78200	0.64733	0.00000	-0.14067	4.20533	2.84667
DDG 44	0.80333	9.76467	0.35267	0.28133	3.56067	1.02733	1.32267	0.84867	1.05867	1.73333	2.69933	0.72800	9.44533	1.73600	1.68467	-1.00867	0.00000	0.43133	2.82867	3.53267
DDG 45	0.40667	9.09867	0.87600	-0.15733	2.82933	-0.61467	0.53667	4.62800	0.00000	0.16333	0.20133	0.41467	8.77333	-0.68667	0.40467	1.04667	4.15333	3.45933	-0.25400	2.00733
DDG 46	-1.25800	6.66000	-1.29667	-1.95067	1.92467	-0.94533	0.78067	-0.50733	-1.36333	0.31333	0.34867	-1.07800	0.00000	-0.24867	-0.69933	-1.69933	-0.60467	0.25133	3.38467	1.81600
DDG 47	0.48533	8.27133	-0.21467	0.01867	2.65333	0.05600	2.08400	1.27333	0.60867	0.78667	1.59400	0.00000	0.43800	1.12600	0.72133	0.20667	-0.11800	1.89067	3.97933	1.60867
DDG 48	0.99933	9.01400	1.32733	0.62200	0.12333	1.95067	0.32200	6.28667	0.38667	1.39400	1.27800	1.05800	-0.23333	0.14533	0.10800	0.55533	0.00000	0.00133	2.44200	-0.28133
DDG 49	1.98600	9.33667	2.50133	1.56067	3.05467	1.91600	3.06267	0.72667	1.74067	1.60800	2.08000	1.54467	0.42267	1.86800	0.00000	2.39800	1.10533	0.54867	4.17733	2.97400
```

Figure 5.1: Example of substitution matrix for folding stability.

5.3 Experimentally derived fitness effects

SodaPop can take as input a substitution matrix that describes the fitness effects of all 1-away mutants derived from deep mutagenesis scanning (DMS) experiments (**Fig. 5.2**). Here, we describe the format of these matrices. Each matrix must be tab-separated and preceded by a header line containing the prefix "Gene_NUM", the gene index matching the corresponding gene file and the protein length in amino acids. Then, each row in the matrix must start by indicating "DMS", the residue position (starting at 1), followed by a list of values for each amino acid substitution, including the identity substitution. If you are simulating several proteins at once, each with a different matrix, you must list the matrices one after the other in a single file. To input this file in a simulation, you must use the `-i` flag **Section 3.3**.

```
Gene_NUM	0	211
DMS 1	-0.061012	-0.051554	-0.050201	-0.036683	-0.021648	-0.062298	-0.034156	-0.024325	-0.037651	-0.024325	0	-0.046309	-0.038143	-0.038268	-0.033732	-0.074003	-0.033432	-0.03362	-0.048549	-0.02905
DMS 2	0	-0.041417	-0.045628	-0.045631	-0.049666	-0.032799	-0.046542	-0.049647	-0.051251	-0.051241	-0.044474	-0.045628	-0.022912	-0.045205	-0.051251	-0.035963	-0.031204	-0.033573	-0.057602	-0.049666
DMS 3	-0.034349	-0.054748	-0.046424	-0.049064	-0.057602	0	-0.050634	-0.062654	-0.052264	-0.060104	-0.051251	-0.044941	-0.044677	-0.046542	-0.052264	-0.033109	-0.048768	-0.052454	-0.058593	-0.058576
DMS 4	-0.034349	-0.054748	-0.046424	-0.049064	-0.057602	0	-0.050634	-0.062654	-0.052264	-0.060104	-0.051251	-0.044941	-0.044677	-0.046542	-0.052264	-0.033109	-0.048768	-0.052454	-0.058593	-0.058576
DMS 5	-0.061398	-0.063158	-0.04784	-0.032799	-0.03658	-0.071168	-0.02942	-0.038254	0	-0.038865	-0.036066	-0.042722	-0.045941	-0.032563	-0.020875	-0.060215	-0.038006	-0.038421	-0.04907	-0.030976
DMS 6	-0.07479	-0.067773	-0.064844	-0.054251	-0.030227	-0.098733	-0.04825	0	-0.049808	-0.033612	-0.03131	-0.063304	-0.056569	-0.049588	-0.059062	-0.083139	-0.044043	-0.031618	-0.05717	-0.048666
DMS 7	-0.036224	-0.06025	-0.050769	-0.047561	-0.06463	-0.051865	-0.05868	-0.053679	-0.054389	-0.055264	-0.046051	-0.046336	0	-0.045267	-0.064037	-0.045044	-0.041705	-0.045756	-0.074396	-0.064023
DMS 8	-0.104845	-0.101134	-0.098205	-0.086198	-0.056199	-0.134232	-0.074181	0	-0.075631	-0.058112	-0.056421	-0.096528	-0.083352	-0.075941	-0.07573	-0.117383	-0.071785	-0.05726	-0.07286	-0.067105
DMS 9	-0.06194	-0.07928	-0.078548	-0.063105	-0.061786	-0.098206	-0.058588	-0.045359	-0.063081	-0.054531	-0.044028	-0.074545	-0.061005	-0.062761	-0.064813	-0.065317	-0.060031	0	-0.062802	-0.061478
DMS 10	-0.062646	-0.083236	-0.066348	-0.067104	-0.082456	0	-0.068675	-0.08715	-0.070439	-0.084958	-0.070069	-0.064062	-0.064602	-0.065008	-0.069631	-0.061535	-0.066323	-0.070464	-0.084058	-0.08343
DMS 11	-0.106426	-0.102422	-0.099801	-0.085794	-0.054972	-0.135812	-0.073777	0	-0.075228	-0.057577	-0.053919	-0.096125	-0.082949	-0.075538	-0.07555	-0.118963	-0.071382	-0.055448	-0.071853	-0.064096
DMS 12	-0.073968	-0.090697	-0.090577	-0.075133	-0.073815	-0.106048	-0.070617	-0.059563	-0.075109	-0.068561	-0.058232	-0.086574	-0.073034	-0.074789	-0.076842	-0.077345	-0.07206	0	-0.075675	-0.074351
DMS 13	0	-0.053163	-0.049168	-0.048363	-0.050863	-0.036267	-0.047606	-0.050233	-0.051837	-0.051826	-0.046148	-0.04836	-0.027288	-0.046924	-0.051837	-0.039334	-0.034142	-0.036067	-0.065919	-0.050863
DMS 14	-0.063614	0	-0.060439	-0.060212	-0.073783	-0.06601	-0.060625	-0.065003	-0.06458	-0.065003	-0.063664	-0.061322	-0.061014	-0.062196	-0.074772	-0.052666	-0.060997	-0.06222	-0.075177	-0.074165
DMS 15	-0.106332	-0.101657	-0.099036	-0.088143	-0.053475	-0.131286	-0.07271	-0.055839	-0.072757	0	-0.053154	-0.09536	-0.08121	-0.074019	-0.074397	-0.118198	-0.070617	-0.064265	-0.071088	-0.063331
DMS 16	-0.085406	-0.091766	-0.070322	-0.055354	-0.076792	-0.107249	-0.057281	-0.075753	-0.070529	-0.076727	-0.074957	-0.067295	-0.070878	0	-0.066167	-0.081365	-0.067919	-0.075951	-0.078006	-0.07411
DMS 17	-0.056011	-0.087287	-0.072757	-0.070388	-0.076662	-0.078068	-0.0728	-0.073973	-0.07417	-0.074947	-0.068944	-0.069101	0	-0.066722	-0.074727	-0.067513	-0.063251	-0.067694	-0.089492	-0.075324
DMS 18	-0.082703	-0.092952	-0.068269	0	-0.090185	-0.108881	-0.066322	-0.087326	-0.071972	-0.091172	-0.076738	-0.06811	-0.076789	-0.058177	-0.071223	-0.075811	-0.072079	-0.077549	-0.09197	-0.076267
DMS 19	-0.107604	-0.096758	-0.096053	-0.07493	-0.054839	-0.10981	-0.070178	-0.056864	-0.073248	-0.056864	0	-0.089728	-0.075025	-0.073834	-0.069883	-0.126492	-0.069397	-0.066336	-0.072314	-0.061534
DMS 20	-0.049566	-0.073304	-0.06139	-0.063375	-0.075502	0	-0.064945	-0.080555	-0.067383	-0.078005	-0.06637	-0.06043	-0.058915	-0.061376	-0.066575	-0.048326	-0.062691	-0.066766	-0.076493	-0.076476
DMS 21	-0.058784	-0.052276	-0.049347	-0.039559	-0.020156	-0.08087	-0.036608	0	-0.038165	-0.023541	-0.021239	-0.048613	-0.044926	-0.037946	-0.046596	-0.064084	-0.032721	-0.021547	-0.046591	-0.037772
DMS 22	-0.025062	-0.046158	-0.037136	-0.040473	-0.049011	0	-0.042044	-0.054064	-0.043673	-0.051514	-0.042661	-0.035654	-0.03539	-0.037255	-0.043673	-0.023822	-0.039481	-0.043864	-0.050002	-0.049985
DMS 23	-0.112695	-0.126115	-0.122678	-0.097159	0	-0.135062	-0.077499	-0.077183	-0.079085	-0.077183	-0.058596	-0.127165	-0.114708	-0.079766	-0.076918	-0.126493	-0.114118	-0.077544	-0.067252	-0.058752
DMS 24	-0.107673	-0.125496	-0.108817	-0.070044	-0.073309	-0.116982	-0.055482	-0.078079	-0.056758	-0.079639	-0.070291	-0.096097	-0.10649	-0.066126	0	-0.107673	-0.084324	-0.079199	-0.073499	-0.068197
DMS 25	-0.049566	-0.073304	-0.06139	-0.063375	-0.075502	0	-0.064945	-0.080555	-0.067383	-0.078005	-0.06637	-0.06043	-0.058915	-0.061376	-0.066575	-0.048326	-0.062691	-0.066766	-0.076493	-0.076476
DMS 26	-0.069534	-0.091173	-0.074285	-0.075041	-0.090393	0	-0.076612	-0.095088	-0.078376	-0.092895	-0.078006	-0.071173	-0.07102	-0.072945	-0.077568	-0.068422	-0.073434	-0.078401	-0.091995	-0.091367
DMS 27	-0.074799	-0.067773	-0.064844	-0.057909	-0.030227	-0.095346	-0.047285	-0.032027	-0.048834	0	-0.03131	-0.063304	-0.054984	-0.048223	-0.058094	-0.083139	-0.044043	-0.039696	-0.056559	-0.048666
DMS 28	-0.023607	-0.043012	-0.03658	-0.033373	-0.050273	-0.039248	-0.043989	-0.039491	-0.040856	-0.041076	-0.032182	-0.032468	0	-0.03231	-0.04968	-0.032087	-0.029088	-0.032799	-0.056991	-0.049666
DMS 29	-0.135392	-0.134044	-0.134044	-0.123095	-0.090519	-0.134044	-0.119988	-0.115539	-0.116242	-0.115539	-0.116378	-0.134044	-0.135392	-0.117665	-0.101833	-0.134317	-0.126906	-0.112625	0	-0.08603
DMS 30	-0.10737	-0.125193	-0.108513	-0.06974	-0.073006	-0.116679	-0.055179	-0.077775	-0.056455	-0.079335	-0.069988	-0.095794	-0.106187	-0.065822	0	-0.10737	-0.084021	-0.078896	-0.073196	-0.067894
DMS 31	-0.073786	-0.070319	-0.06739	-0.060456	-0.032773	-0.095527	-0.049831	-0.034573	-0.05138	0	-0.033856	-0.06585	-0.05753	-0.050769	-0.06064	-0.082127	-0.04659	-0.042242	-0.059106	-0.051212
DMS 32	-0.047818	-0.078266	-0.063736	-0.061368	-0.067419	-0.069874	-0.064384	-0.064953	-0.065149	-0.065927	-0.059924	-0.060081	0	-0.058529	-0.065484	-0.059319	-0.055058	-0.059501	-0.080249	-0.066081
DMS 33	-0.073895	-0.072608	-0.066797	-0.070488	-0.090121	-0.06963	-0.073099	-0.089778	-0.077352	-0.089778	-0.092097	-0.062874	-0.067807	-0.071922	-0.076148	0	-0.06719	-0.073917	-0.090726	-0.089299
DMS 34	-0.059145	-0.064673	-0.044016	0	-0.062717	-0.079139	-0.04329	-0.060376	-0.048548	-0.064034	-0.053265	-0.043856	-0.053342	-0.036158	-0.047671	-0.051528	-0.048125	-0.054611	-0.074969	-0.054389
DMS 35	-0.08038	-0.074938	-0.073585	-0.053096	-0.038697	-0.081951	-0.050284	-0.041374	-0.053257	-0.041374	0	-0.069693	-0.053953	-0.053873	-0.04986	-0.100268	-0.049561	-0.050669	-0.062472	-0.044119
DMS 36	-0.09939	-0.103794	-0.078831	-0.060763	-0.064469	-0.108775	-0.054644	-0.065368	0	-0.064636	-0.063985	-0.071519	-0.07717	-0.060527	-0.044724	-0.098206	-0.068703	-0.066096	-0.066216	-0.059449
DMS 37	-0.103782	-0.115774	-0.115184	-0.080441	-0.047437	-0.125617	-0.064345	-0.064732	-0.06806	-0.064732	-0.058502	-0.114881	-0.101618	-0.072023	-0.062881	-0.116371	-0.100827	-0.074464	-0.058356	0
DMS 38	-0.054761	-0.060083	-0.05729	-0.042582	0	-0.070834	-0.029604	-0.033973	-0.032526	-0.033973	-0.015542	-0.062225	-0.056774	-0.032676	-0.029554	-0.061073	-0.056184	-0.031139	-0.027642	-0.015643
DMS 39	-0.109467	-0.12729	-0.11061	-0.071837	-0.075103	-0.118776	-0.057276	-0.079872	-0.058552	-0.081432	-0.072085	-0.097891	-0.108284	-0.067919	0	-0.109467	-0.086118	-0.080993	-0.075293	-0.069991
DMS 40	-0.084426	-0.090786	-0.069342	-0.054374	-0.075813	-0.106269	-0.056301	-0.074774	-0.06955	-0.075748	-0.073977	-0.066316	-0.069898	0	-0.065187	-0.080385	-0.06694	-0.074971	-0.077249	-0.073131
DMS 41	-0.060099	-0.077439	-0.076708	-0.061264	-0.059945	-0.096057	-0.056747	-0.043519	-0.06124	-0.052691	-0.042187	-0.072705	-0.059165	-0.06092	-0.062973	-0.063476	-0.058191	0	-0.060962	-0.059637
DMS 42	-0.027377	-0.034898	-0.024625	-0.024332	-0.041682	-0.034938	-0.024083	-0.027772	-0.026855	-0.027772	-0.025474	-0.025465	-0.023288	-0.022567	-0.039796	-0.02582	0	-0.02582	-0.043673	-0.03819
DMS 43	-0.063183	-0.061896	-0.056912	-0.060381	-0.079187	-0.059745	-0.062165	-0.079066	-0.066417	-0.079066	-0.081163	-0.052989	-0.057922	-0.061814	-0.065214	0	-0.057305	-0.063205	-0.079792	-0.078364
DMS 44	-0.106332	-0.101657	-0.099036	-0.088143	-0.053475	-0.131286	-0.07271	-0.055839	-0.072757	0	-0.053154	-0.09536	-0.08121	-0.074019	-0.074397	-0.118198	-0.070617	-0.064265	-0.071088	-0.063331
DMS 45	-0.059987	-0.075797	-0.057988	-0.057644	-0.065578	-0.073686	-0.054614	-0.060524	-0.059292	-0.060524	-0.058354	-0.057305	-0.055087	-0.053803	-0.060682	-0.058191	0	-0.05766	-0.066392	-0.059988
DMS 46	-0.108102	-0.112506	-0.087544	-0.068957	-0.073489	-0.117487	-0.062838	-0.074388	0	-0.073656	-0.073006	-0.080539	-0.08619	-0.068721	-0.054193	-0.106918	-0.076896	-0.075117	-0.075459	-0.068469
DMS 47	-0.071653	-0.078851	0	-0.055223	-0.079065	-0.069355	-0.061778	-0.081043	-0.066723	-0.081043	-0.080344	-0.046611	-0.066394	-0.061327	-0.067075	-0.061259	-0.061158	-0.079729	-0.081852	-0.082685
DMS 48	-0.047818	-0.078266	-0.063736	-0.061368	-0.067419	-0.069874	-0.064384	-0.064953	-0.065149	-0.065927	-0.059924	-0.060081	0	-0.058529	-0.065484	-0.059319	-0.055058	-0.059501	-0.080249	-0.066081
DMS 49	-0.095887	-0.090383	-0.055617	-0.068062	-0.091419	-0.084812	-0.073201	-0.092248	-0.076626	-0.092248	-0.091664	0	-0.073564	-0.069418	-0.075628	-0.069866	-0.071878	-0.087291	-0.093474	-0.093785
```

Figure 5.2: Example of substitution matrix derived from deep mutational scanning data.

6. Results and analysis

The package includes a small bioinformatics pipeline to extract, analyze and plot your data. The scripts are written in Bash and R* and can be used separately from the main program. Those familiar with scripting languages can also use the scripts as templates for other analyses.

By default, SodaPop will always output population snapshots following the format chosen by the user. To toggle on automatic analysis, use the `-a [-analyze]` boolean flag with your command (see List of command-line options).

> **Note** As a requirement for plotting, the R programming language must be installed on your machine. You can download the latest version of R here. The script will automatically install the required R packages for you at runtime. ■

6.1 Scripts

barcodes.sh : this is a Bash script that uses Unix utilities such as `awk`, `grep`, `uniq`, `sort` and `join`. It executes the following operations:

1. Detect snapshot file format
2. Unzip and convert binary snapshots to text
3. Extract and sort the barcodes for each time point
4. Compute the average population fitness for each time point
5. Count the number of occurrences of a given barcode at each time point
6. Identify the time point corresponding to the fixation of a single barcode
7. Combine all time points in a dataframe
8. Call *polyclonal_structure.R*

This script can executed independently from the main program using the command-line. It requires the following parameters: [name of the simulation directory] [number of generations] [population size] [step size (dt)] [encoding format] [gene count per cell].

As an example, the following would be the command to analyze a long format simulation with n = 10000, m = 10000 and a time step of t = 25 generations:

```
./barcodes.sh test_sim 10000 10000 25 0
```

polyclonal_structure.R : this is a R script that imports the time series data and the fitness table to generate four distinct plots. Examples of these plots are found throughout this chapter.

6.2 Mean fitness plot

The first plot is named *fitness* and shows the evolution of the average population fitness during the course of the simulation.

In the example below (**Fig. 6.1**), we observe a steep drop in fitness in the first few generations. This is a typical case of Muller's ratchet: under a high mutation rate, the population sees its mean

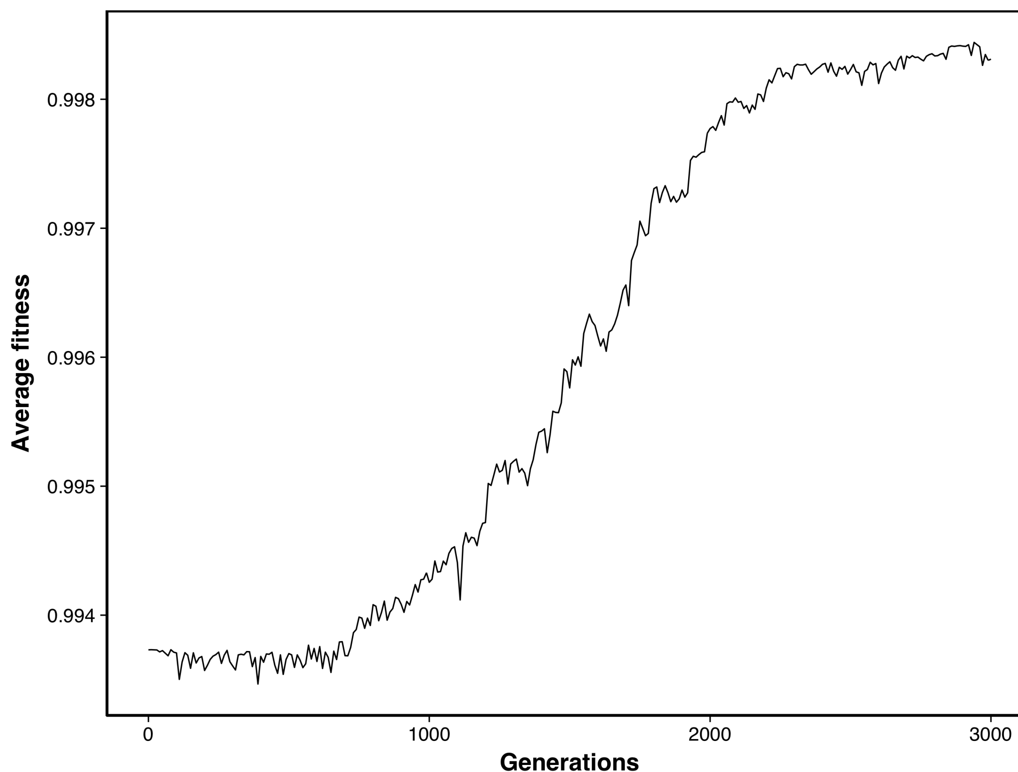

Figure 6.1: Example trajectory showing the evolution of the average population fitness over 10^3 generations for $N_e = 10^4$.

fitness decrease as the fittest individuals are hit by marginally deleterious mutations. Over time, marginally beneficial mutations arise, conferring a selective advantage that is driving the mean fitness higher.

6.3 Clonal structure

The second plot is named *clonal_structure*. It represents the share of each unique barcode over the course of the simulation, up to the point of fixation of a single lineage. In **Fig. 6.2**, the vertical axis represents the density of each lineage. The sum of all lineages is always equal to the size of the population (in this case, 10^4). The slight slope in the top left area of the plot is a result of discarding all lineages that are lost immediately after the first time point. The reason is that as the ratio of unique barcodes to the population size approaches 1, a majority of barcodes will be naturally lost to stochastic drift. Removing these lineages speeds up plotting and makes for a sharper image.

6.4 Clonal trajectories

The two remaining plots are named *clonal_trajectories* and *log_clonal_trajectories*. They are temporal representations of the structure of the population in terms of each segregating lineage – that is, cells sharing a common barcode.

In **Fig. 6.3**, the vertical axis represents the total count of each lineage as it segregates. Again, this plot will show lineages up to the point of fixation, if such a point exists. This plot is useful to quickly identify selective sweeps. **Fig. 6.4** shows the same plot with the vertical axis in log10-scale. This is useful to visualize the segregation of low-frequency lineages and estimate the selection

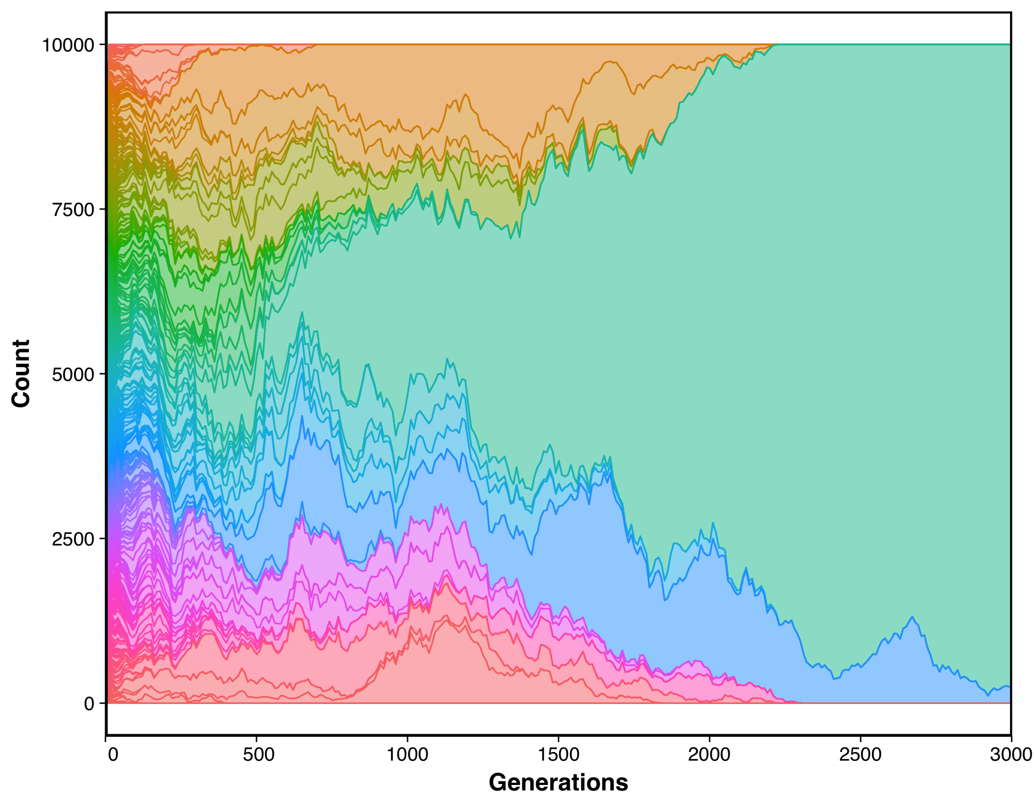

Figure 6.2: Example run of $N_e = 10^4$ cells showing the clonal structure spanning 10^3 generations.

coefficient and establishment time of beneficial mutation by extrapolating the exponential growth regime of a lineage. These estimations can then be correlated with the exact numbers found in the mutation log (see section below).

6.5 Mutation log

The file *MUTATION_LOG* contains the information on arising mutations in a simulation run. It can be toggled on with the `-e [-track-events]` flag.

The tab-separated file lists mutations line-by-line:

```
AGACTCAAGTGTGAC 6 K 272 E 0.001078 1
TGCAAGCAAACGGGC 5 D 160 H -0.363828 9
CGAGACTGTGGGAGT 4 T 48 R -0.198414 21
...
```

The columns respectively correspond to the barcode, the gene ID, the previous residue identity, its position in the sequence, the new residue identity (the mutation), its selection coefficient and the generation at which it occurred.

6.6 Command log

The *command.log* file is created by default for each simulation run and contains the full command and its corresponding standard output.

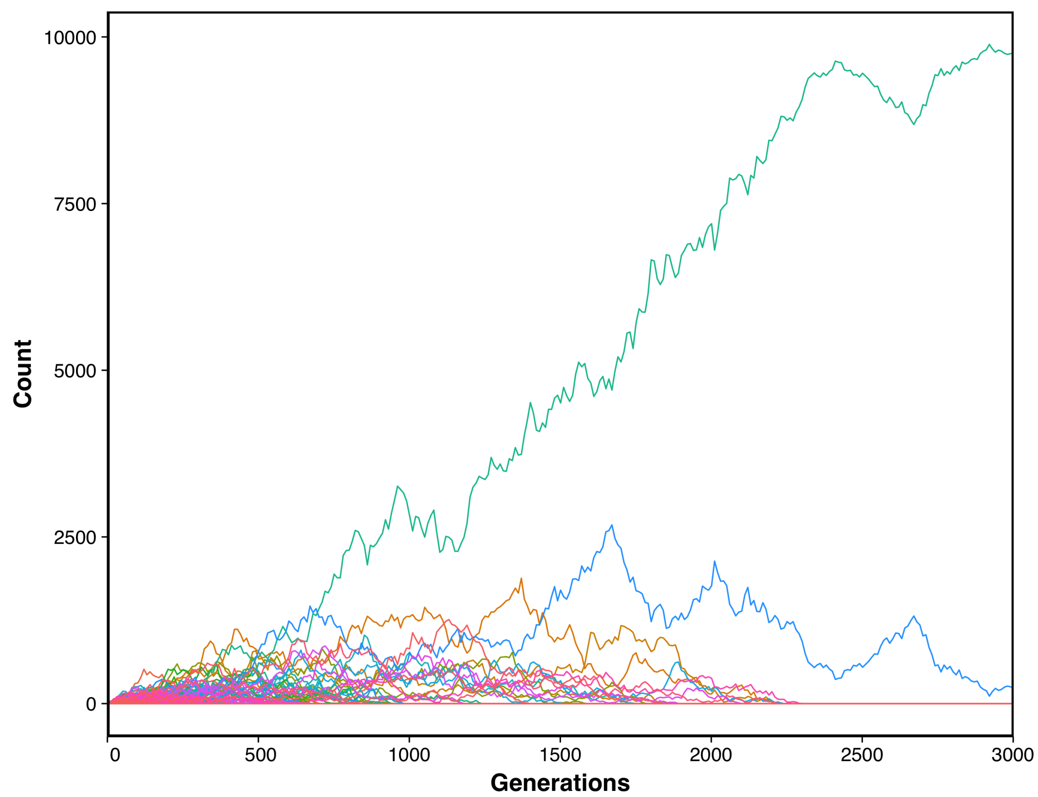

Figure 6.3: Example run of $N_e = 10^5$ cells showing the clonal trajectories spanning 10^4 generations.

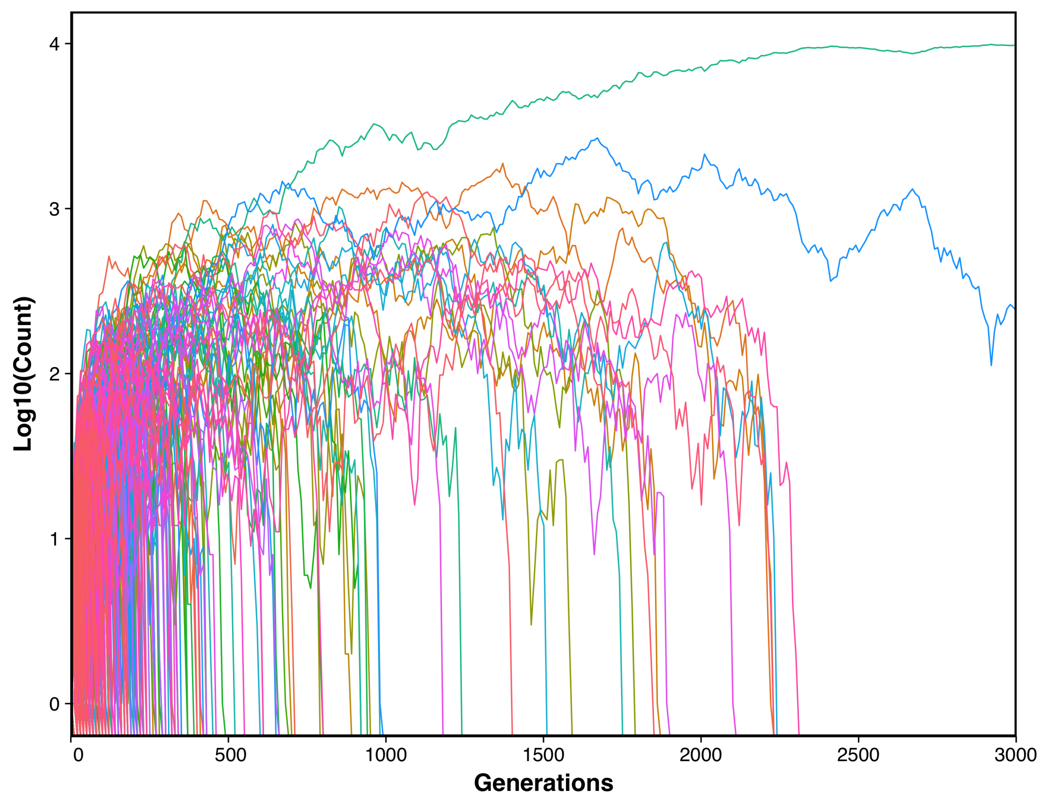

Figure 6.4: Example run of $N_e = 10^5$ cells showing the clonal trajectories spanning 10^4 generations. While the majority of trajectories fall to extinction, the topmost trajectory (in green) sweeps the population.

7. Fitness functions

SodaPop allows you to specify how organismal fitness is calculated by choosing a fitness function.

There are currently seven native functions in this version. However, it is easy to implement a custom fitness function (for details, see **Section 8.1**). Here we introduce the native fitness functions of the program.

7.1 Folding stability

This function assumes that organismal fitness is given by the number of correctly folded proteins in the cell. Assuming a two-state model of protein folding [4], the fraction of proteins in the native state P_{nat} is given by

$$P_{(nat,i)} = \frac{1}{1+e^{\beta \Delta G_i}} \tag{7.1}$$

where $\beta = \frac{1}{k_B T}$ ($k_B T = 0.593 kcal/mol$) and ΔG_i is the folding free energy of the protein. Summing over all genes and accounting for gene expression level, we define fitness as

$$fitness = \sum_i (A_i \cdot P_{(nat,i)}) \tag{7.2}$$

where i is the index for each gene and A_i is the corresponding cellular abundance of the protein.

7.2 Metabolic flux

This function expresses organismal fitness as proportional to metabolic flux of the cell. Assuming that proteins are enzymes in a linear metabolic pathway, and accounting for gene expression level and enzymatic efficiency, we obtain

$$fitness = \frac{a_0}{\sum_i (\varepsilon_i \cdot A_i \cdot P_{(nat,i)})^{-1}} \tag{7.3}$$

where a_0 is a normalizing factor, i is the index for each gene, ε_i is the catalytic efficiency of the protein and A_i is its intracellular abundance. By requiring that fitness be optimally 1 at very stable regimes ($\Delta G \rightarrow -\infty$), we get

$$a_0 = \frac{1}{\sum_{i=1}^{N} (A_i)^{-1}} \tag{7.4}$$

7.3 Toxicity

This function expresses fitness the fitness cost due to misfolded proteins in the cell, given by

$$fitness = c \sum_i A_i (1 - P_{(nat,i)}) \tag{7.5}$$

where c is a negative constant describing the fitness cost per misfolded protein [2].

7.4 Combined metabolic output

This function describes fitness as a combination of the two previous functions, namely metabolic flux and toxicity due to misfolding. It is expressed as the difference

$$fitness = f_{flux} - f_{toxicity} \tag{7.6}$$

If $f_{toxicity} > f_{flux}$, the cell is not viable and the fitness is 0.

7.5 Neutral

In this scenario, the fitness effect of all arising mutations is null. That is, the selection coefficient of a mutation is 0. Thus, organismal fitness is given by

$$fitness = f_{neutral} = 1 \tag{7.7}$$

7.6 Multiplicative

Consider a gene with wild-type fitness w_f that acquires n mutations with selection coefficient s_i. The multiplicative fitness of this gene is given by

$$fitness = F_n = w_f \prod_{i}^{n} (1 + s_i) \tag{7.8}$$

By design, only a single mutation can arise in a gene at once, so this computation is straightforward and only requires the previous fitness state of the sequence. In this case, if multiple genes are present, the organismal fitness is the arithmetic mean of all sequences' fitnesses.

7.7 No mutation

Mutations can be turned off completely by using this dummy fitness function. This is useful to simulate the effects of standing genetic variation on population dynamics.

8. Modifying the code

SodaPop is written in an object-oriented programming (OOP) framework, making it easily hackable by anyone with elementary knowledge of OOP and/or C++.

8.1 Implementing a fitness function

```
// METABOLIC FLUX FITNESS FUNCTION
double PolyCell::flux()
{
    double sum_func = 0;
    double a = 0;
    for(auto gene_it = Gene_arr_.begin(); gene_it != Gene_arr_.end(); ++gene_it){
        // efficiency * Pnat * abundance
        sum_func += 1/(gene_it->eff()*gene_it->functional());
        a += gene_it->A_factor();
    }
    return static_cast<double>(a)/sum_func;
}
```

Figure 8.1: Example of fitness function. The function returns a double.

Fitness functions are located in the *PolyCell.h* source file. All functions have the same return type (double) and are indexed by increasing integers starting from 1. A static function pointer selects the correct fitness function to use during simulations (**Fig. 8.2**). Users can implement a new fitness function by following this signature and adding their new function to the selectFitness() switch case statement.

8.2 Adding gene properties

The present version implements folding stability (ΔG) and enzymatic efficiency (K_{cat}/K_m) as the fitness defining properties for each gene product. Other biochemical properties can be implemented here, such as or binding affinity (K_d). Additionally, other properties like protein concentration and essentiality can be used to model specific genotype-to-phenotype relationships (fitness functions).

```
void PolyCell::selectFitness()
{
    switch(PolyCell::ff_){
        case 1: fit = &PolyCell::fold;
            break;
        case 2: fit = &PolyCell::flux;
            break;
        case 3: fit = &PolyCell::toxicity;
            break;
        case 4: fit = &PolyCell::metabolicOutput;
            break;
        case 5: fit = &PolyCell::multiplicative;
            break;
        case 6: fit = &PolyCell::neutral;
            break;
        case 7: fit = &PolyCell::noMut;
            break;
        default:;
    }
}
```

Figure 8.2: The control structure used to select the proper fitness function. New functions can be added to the switch statement with a function pointer.

```
void ch_dg(const double a){dg_ = a;}
void ch_f(const double a){f_ = a;}
void ch_eff(const double a){eff_ = a;}
void ch_conc(const double c){conc_ = c;}
void ch_Na(const int a){Na_ = a;}
void ch_Ns(const int a){Ns_ = a;}
void ch_e(const double e){e_ = e;}

private:
    int g_num_;     //numeric ID pointing to primordial gene
    int ln_;        //length nuc seq
    int la_;        //length aa seq

    int Na_;        //number of non-synonymous substitutions
    int Ns_;        //number of sysnonymous substitutions

    std::string nucseq_;    //nucleotide sequence

    double dg_;     //stability
    double f_;      //gene "fitness"
    double eff_;    //enzymatic efficiency

    double conc_;    //concentration
    double e_;       //essentiality: between 0 and 1, can be used as a coefficient

    static std::gamma_distribution<> gamma_;
    static std::normal_distribution<> normal_;
```

Figure 8.3: Part of the *gene.h* class definition. New private properties can be added here, with the corresponding accessory functions.

9. Implementation and performance

SodaPop is built on streamlined data structures and a fast algorithm to achieve high computational efficiency and to minimize the general trade-off between flexibility and runtime. Here we briefly describe this software's design decisions.

9.1 Data structures and complexity

The population is implemented as a vector (*std::vector* container) of cells. This is because vectors store elements contiguously in physical memory, making elements accessible sequentially through iterators and pointer arithmetic, and providing spatial locality of reference. The most common operation is appending a cell at the end of the vector, which is done in amortized constant time [$O(1)$]. This operation is linear $O(k)$ when appending k cells at once. Random iterator access is also done in constant time, making the random rescaling of the population an efficient operation that is linear in the number of cells to be added/removed. Vectors also confer efficient memory handling since memory can be reserved ahead of time, provided a good estimate of capacity is known. Because we use a constant population size between generations, and because generation overhead never exceeds $2N_e$, the upper bound on population size is a known variable, preventing costly $O(n)$ vector doubling reallocations when the container reaches full capacity.

Whenever it is possible, we use safe C++ standard library operators such as *std::shuffle* and *std::swap*. The first operator is used when the new generation exceeds the intended size (as a result of binomial drawings). It performs a fast random shuffle of the population vector indexes. This is followed by a standard resizing of the vector whose complexity is linear on the number of elements erased. Essentially, instead of picking k cells to erase at random and downshifting the vector every time, we shuffle the vector and then trim off the last k cells. The second operator is used to assign the new vector of cells to the population vector, overwriting the previous generation. Because cells are large objects containing a lot of information distributed over different containers, overwriting the population one by one would not only be inefficient, but potentially pose the risk of losing information by shallow copies of certain object members. The *swap* operator uses move semantics to swap the two cell vectors efficiently. We can then reset the old vector for the new generation.

The program reads and writes binary files because they are smaller in size and can be handled with buffers reading discrete chunks of data. Since the memory overhead for I/O is significantly higher than that of type conversion, the computational cost incurred by converting types to binary is negligible. Moreover, while using binary over raw text might not be significantly advantageous for small-scale simulations, it can trim down the size of snapshots in larger runs where each population can reach several hundred megabytes in size.

9.2 Pseudo-random number generation

We are working with the assumption that arising mutations are uniformly distributed along the genome, and a single simulation run can easily produce millions of mutations. Hence, we need a reliable pseudo-random number generator (PRNG) that lacks bias. Additionally, we want our

PRNG to be relatively fast as it is being used heavily in the algorithm. Finally, the PRNG should be lightweight and compatible with C/C++. Following these criteria, we opted to use the PCG family of random generators by Melissa E. O'Neill, as a substitute to the widely used Mersenne Twister with a native implementation in C++11. The latter has a large state space and is known to be uneven in its output, while PCG is compact, relatively fast, and uniform in its output, which makes it a good candidate for simulation.

9.3 Runtime

SodaPop is the first publicly available tool that can effectively simulate multi-scale molecular evolution and polyclonal population dynamics. We benchmarked SodaPop for multiple population sizes and number of generations. All simulations were run on a standard iMac desktop with a 3.2GHz Intel Core i5 processor and 16GB memory. **Figure 8.1** shows that runtime is quasi-monomial with respect to population size. Simulating up to a million cells for long time periods is entirely tractable using standard desktop computers. We limited our desktop benchmarking to $N_e = 10^6$ cells, as higher orders of magnitude induce a shift in performance due to a RAM bottleneck. The simulation of populations with higher orders of magnitude requires a larger amount of memory than the current standard in commercial desktop computers.However, larger population sizes can be simulated on high-performance computing clusters where memory allocation is not as limiting.

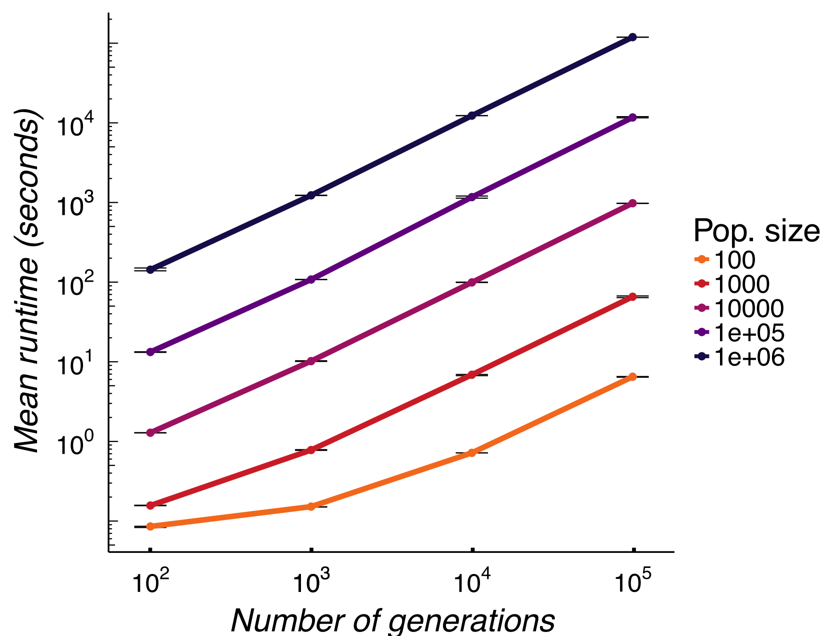

Figure 9.1: Average runtimes per population size for different time scales. The initial population was characterized by a monoclonal cell with a single 194 bp gene. The time step for snapshots was set at 0.01∗(number of generations) for each scenario. Each point represents the average runtime over 100 simulations for a particular condition. Error bars represent standard error of the mean.

Acknowledgements

We would like to thank members of the Serohijos lab for testing the program and contributing valuable questions and ideas. Part of this research was funded by NSERC & start-up funds from UdeM.

Annex: Using the command-line

The command-line interface is a powerful way to interact with the computer. In its simplest form, the command-line is a space where you type commands for the computer to execute. On Mac OS/X, the command-line is an application called Terminal. It is located in the `/Applications/Utilities/` folder.

The overwhelming majority of command-line programs follow the same simple syntax. A command can be broken down into three basic components:

- The *utility*, also referred to as the program. Utilities execute a sequence of actions depending on their input. In some cases, you may use a utility without any flag or argument.
- The *flags*: flags are like options. They allow you to modify the behavior of the utility.
- The *arguments*: some utilities take arguments. These are commonly files, but they can be numbers, words or special characters.

Here is a simple command. We will break it down by component.

```
ls -l Documents/
```

`ls` is the utility. It lists the content of directories in the command-line interface.

`-l` is a flag. It indicates that we want more information than is provided by default. Flags are most often preceded by a hyphen ('-') and consist of single characters. Some flags can also be specified using a word preceded by a double hyphen ('--'). Using one or the other will have the interpretation by the command-line. For instance, typing

```
man -h
```

or

```
man -help
```

will output the same response from the terminal. The `man` utility displays the manual pages for a specific utility. The command is used in conjuction with another utility as its *argument*. For example,

```
man ls
```

will display the manual page for `ls`.

The last component of our initial command is the argument 'Documents/'. It tells `ls` that we want to show the content of that directory. And that?s it! Besides `ls`, you will need to know how to use a few other basic utilities in order to get started.

```
cd
```

is the command to **c**hange **d**irectory. By default, without any flag or argument, `cd` will move up one folder. You can navigate down a folder by giving the name of the folder you wish to move to as an argument. If you are unsure of the folder you are currently in, type

```
pwd
```

The command stands for **p**rint **w**orking **d**irectory, and it does exactly that.

If you want to view the contents of a text file, say 'some_file.txt', try

```
less some_file.txt
```

This will display the content of the file in the terminal. You can scroll through the file using the up and down arrows. Pressing Q will quit the utility and bring you back to the command-line prompt.

If you want to copy a file to a specific location, say 'Data/', type

```
cp some_file.txt Data/
```

Likewise, you can instead *move* the file using a similar syntax:

```
mv some_file.txt Data/
```

Finally, you can remove a file using

```
rm some_file.txt
```

Note A handy feature of the terminal is TAB autocompletion. Whenever you are typing an argument, say a path or a filename, you can type the first few letters and press TAB. This will list all the files and folders accessible from your current directory corresponding to that prefix. If there is only one match, it will autocomplete the argument for you. Getting familiar with this option will help you to use the command-line swiftly. ■

Bibliography

[1] Rhiju Das and David Baker. Macromolecular modeling with rosetta. *Annual Review of Biochemistry*, 77(1):363–382, 2008. PMID: 18410248.

[2] K. A. Geiler-Samerotte, M. F. Dion, B. A. Budnik, S. M. Wang, D. L. Hartl, and D. A. Drummond. Misfolded proteins impose a dosage-dependent fitness cost and trigger a cytosolic unfolded protein response in yeast. *Proc Natl Acad Sci U S A*, 108(2):680–5, 2011.

[3] Raphael Guerois, Jens Erik Nielsen, and Luis Serrano. Predicting changes in the stability of proteins and protein complexes: A study of more than 1000 mutations. *Journal of Molecular Biology*, 320(2):369 – 387, 2002.

[4] A. W. Serohijos and E. I. Shakhnovich. Contribution of selection for protein folding stability in shaping the patterns of polymorphisms in coding regions. *Mol Biol Evol*, 31(1):165–76, 2014.

[5] S. Yin, F. Ding, and N. V. Dokholyan. Eris: an automated estimator of protein stability. *Nat Methods*, 4(6):466–7, 2007.