slimr: An R package for integrating data and tailor-made population genomic simulations over space and time

Russell Dinnage[1, 2]
Stephen D. Sarre[1]
Richard P. Duncan[1]
Christopher R. Dickman[3]
Scott V. Edwards[4]
Aaron Greenville[3]
Glenda Wardle[3]
Bernd Gruber[1]

Author Affiliations

1. Centre for Conservation Ecology and Genomics, Institute for Applied Ecology, University of Canberra, Canberra, ACT, Australia
2. Department of Biological Sciences, Florida International University, Miami, FL, USA
3. Desert Ecology Research Group, School of Life and Environmental Sciences, University of Sydney, NSW 2006, Australia
4. Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA

Abstract

- Software for realistically simulating complex population genomic processes is revolutionizing our understanding of evolutionary processes, and providing novel opportunities for integrating empirical data with simulations. However, the integration between simulation software and software designed for working with empirical data is currently not well developed. In particular, SLiM 3.0, which is one of the most powerful population genomic simulation frameworks for linking evolutionary dynamics with ecological patterns and processes is a standalone scripting language with limited data manipulation abilities. Here we present slimr, an R package designed to create a

seamless link between SLiM 3.0 and the R development environment, with its powerful data manipulation and analysis tools.

- We show how slimr, in combination with SliM, facilitates smooth integration between genetic data, ecological data and simulation in a single environment. The package enables pipelines that begin with data reading, cleaning, and manipulation, proceed to constructing empirically-based parameters and initial conditions for simulations, then to running numerical simulations, and finally to retrieving simulation results in a format suitable for comparisons with empirical data – aided by advanced analysis and visualization tools provided by R (such as ABC and deep learning).
- We demonstrate the use of slimr with an example from our own work on the landscape population genomics of desert mammals, highlighting the advantage of having a single integrated tool for both data analysis and simulation.
- slimr makes the powerful simulation ability of SliM 3.0 directly accessible to R users, allowing integrated simulation projects that incorporate empirical data without the need to switch between software environments. This should provide more opportunities for evolutionary biologists and ecologists to use realistic simulations to better understand the interplay between ecological and evolutionary processes. slimr is available at https://rdinnager.github.io/slimr/.

Keywords: population genomics; simulation; landscape genomics; evolution; ecology; evolutionary ecology; application; software

Introduction

Mathematical modelling and simulation are critical cornerstones of population genetic practice. At a fundamental level, empirical datasets demand analytical tool-kits that can accomodate their high complexity, and recent developments in sophisticated simulation software have the potential to provide mechanistic insight into increasingly complex evolutionary scenarios (Strand 2002; Carvajal-Rodríguez

2010; Yuan *et al.* 2012; Messer 2013; Hoban 2014; Kelleher *et al.* 2016; Haller & Messer 2019). However, utilising flexible simulations requires exploration of large parameter space, which often generates large amounts of data that need sophisticated computational tools to unpack, interrogate and synthesize. Likewise, using simulations to model empirical data is an emerging field because it allows researchers to deal with complex situations where it is difficult to obtain a closed likelihood (Beaumont *et al.* 2002; Marjoram *et al.* 2003; Sisson 2018; Torada *et al.* 2019; Brehmer *et al.* 2020; Wang *et al.* 2020; Cranmer *et al.* 2020). To facilitate more rapid and seamless interrogation and synthesis between empirical data and population genetics simulation, we present `slimr` (https://rdinnager.github.io/slimr/). `slimr` is an R package designed to link the very large and widely used ecosystem of analysis and visualization tools in the R statistical language to the SLiM scripting language (Haller & Messer 2019), a popular, powerful and flexible population genetics simulation tool. The package creates a smooth fusion between the computational power and flexible model specification of SLiM with the advanced statistical analysis, visualisation, and metaprogramming tools of R.

Package Description

`slimr` is an R package that interfaces with SLiM 3.0 software for forward population genetics simulations (see Messer 2013; Haller & Messer 2019 for full details on SLiM, as well as the website at https://messerlab.org/slim/).

`slimr` implements a Domain Specific Language (DSL) that mimics the syntax of SLiM, allowing users to write and run SLiM scripts and capture resulting simulation output, all within the R environment. Much of the syntax is identical to SLiM, but `slimr` offers additional R functions that allow users to manipulate SLiM scripts ("`slimr` verbs") by inserting them directly into any SLiM code block. This enables R users to create SLiM scripts that explore large numbers of different parameters and also automatically produce output from SLiM for powerful downstream analysis within R.

The features of slimr fall into three categories: 1) SLiM script integrated development, 2) data input/output, and 3) SLiM script metaprogramming. The first set of features is designed to make it easy to develop SLiM scripts in an R development environment such as Rstudio, and mostly recapitulates features that SLiM users already have access to in the form of SLiMgui and QtSLiM (https://messerlab.org/slim/). The second and third features are implemented using what we call "`slimr` verbs", allowing SLiM and R features to be combined in advanced ways. The integration between R and SliM provided by `slimr` compensates knowledgeable users of R for a lack of knowledge of SLiM, helping to lower the barrier to learning and using SLiM.

Each of the 3 categories has subcategories of features as follows:

1) Integrated Development
    a) Autocomplete and Documentation (within R) for SliM code
    b) Code highlighting and pretty printing of SLiM code
    c) Rstudio addins
    d) Run code in SLiM from R

2) Data Input/Output
    a) Automatic output generation and extraction from SLiM to R (`slimr_output()`)
    b) Insert arbitrary R objects into SLiM scripts through inlining (`slimr_inline()`)

3) Metaprogramming
    a) Code templating for SLiM scripts (`slimr_template()`)
    b) Flexible general metaprogramming tools (support for `rlang`'s `!!` and `!!!` forcing operators)

In the next section we describe each of these features in greater detail, showing examples through screenshots and code snippets.

Integrated Development

`slimr` allows the user to write SliM code from within an R integrated development environment (IDE). `slimr` is designed to work well with Rstudio, but can be used in any R IDE. The syntax used to write SliM code is very similar to the native SliM syntax, with a few modifications to make it work with the R interpreter. As an example, here is a minimal SliM program, and its counterpart written in `slimr`.

SliM code:

```
initialize()
{
  initializeMutationRate(1e-7);
  initializeMutationType("m1", 0.5, "f", 0.0);
  initializeGenomicElementType("g1", m1, 1.0);
  initializeGenomicElement(g1, 0, 99999);
  initializeRecombinationRate(1e-8);
}
1
{
  sim.addSubpop("p1", 500);
}
10000
{
  sim.simulationFinished();
}
```

slimr code:

```
slim_script(
  slim_block(initialize(),
             {
               initializeMutationRate(1e-7);
               initializeMutationType("m1", 0.5, "f", 0.0);
               initializeGenomicElementType("g1", m1, 1.0);
               initializeGenomicElement(g1, 0, 99999);
               initializeRecombinationRate(1e-8);
             }),
  slim_block(1,
             {
               sim.addSubpop("p1", 500);
```

```
            }),
  slim_block(10000,
            {
              slimr_output_full();
              sim.simulationFinished();
            })
) -> script_1
```

The above code assigns the script to an R object `script_1`, which can then be further manipulated, printed prettily, and sent to SliM to run. See Fig. 1 to see what the above script looks like in the Rstudio IDE, and examples of things you can do with it. A script is specified using the `slimr_script` function, within which you create `slimr` coding blocks, using the `slimr_block` function. The user can create as many `slimr` code blocks as desired within a `slimr_script`. We've added a `slimr` "verb" (`slimr_output_full`), which tells SliM to output the full state of the simulation and return it to R during the execution of the block. We will discuss `slimr` verbs in more detail in the next section.

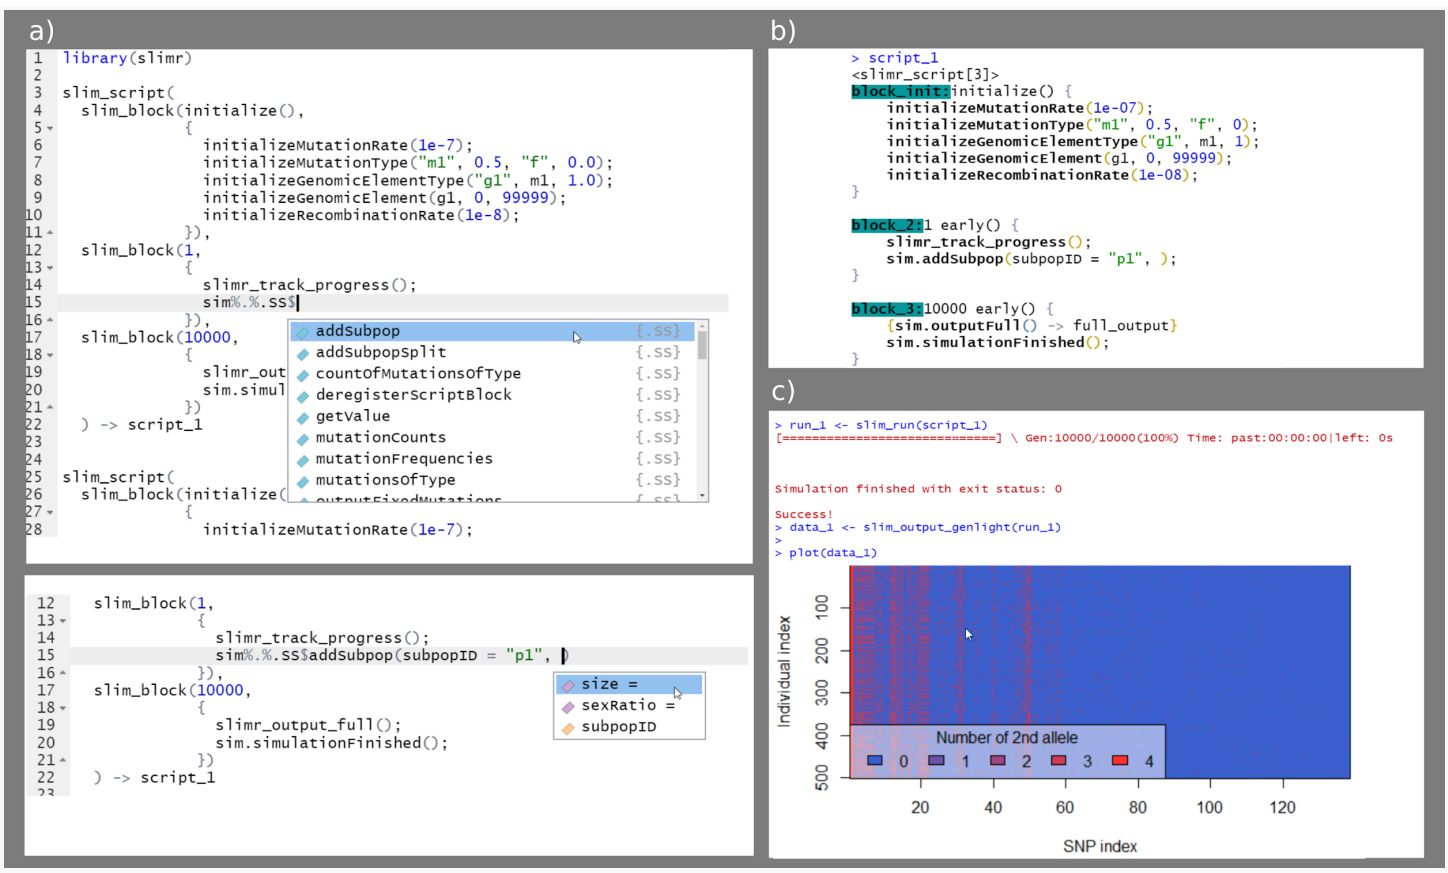

Figure 1. Screenshots of working with `slimr` in Rstudio. An example of autocomplete for SLiM code (a), an example of pretty printing of a `slimr_script` object (b), and an example of running a script, converting its output to a standard R format for genetic data (adegenet package's genlight), and plotting it.

`slimr` makes it easy to write SliM code in R after the user learns a few differences between SliM and slimr. This means users can learn how to write complex SliM simulations by reading SliM 3.0 documentation and the examples found within it (https://messerlab.org/slim/). To make this process easier for `slimr` users, the entire reference documentation for functions in SliM 3.0 and Eidos scripting language (on which SliM is based) is included in `slimr` (with the original author's permission). Hence, not only can R users look up relevant SliM functions in their R session, but the R IDE can perform autocompletion.

`slimr` also provides several Rstudio addins to make common tasks simple. These include an addin that converts SliM code to `slimr` code automatically by pasting from the clipboard, and an addin to easily send `slimr_script` calls to be run in SliM. Converting code from `slimr` to SliM is also supported, including the ability to open the converted code in SLiMGUI or QtSLiM if installed.

`slimr_script` objects (and `slimr_script_coll`, which contain lists of `slimr_script` objects) can be run in SliM, and their results collected and returned using the `slim_run` function.

Data Input/Output

The input/output and metaprogramming features of `slimr` are achieved using special `slimr` "verbs" that can be inserted directly into `slimr` coding blocks (Fig. 2). These verbs are pure R functions that modify how the SliM script will be generated and run in SliM. They are not passed directly to SliM, but make it easy for R to interact with SliM. In this way, `slimr` code appears to be a hybrid between SliM and R code. `slimr` verbs allow all setup and logic required to use SliM with R to occur inside the coding blocks comprising the slimr_script object, thus requiring fewer arguments to be set in preparation for downstream analysis (for example `slim_run` does not require many complex arguments because most

of what it needs to know is embedded in the `slimr_script` object). In our experience, this leads to a very smooth experience using `slimr` by reducing the frequency of switches between different mental modes. By convention, all `slimr` verbs have the prefix `slimr_`, and are meant to be used only inside `slim_script` calls. All other `slimr` functions are prefixed with `slim_`, which means they are to be used on `slimr_script` objects, and not inside a `slim_script` call. The following are the main `slimr` verbs supported:

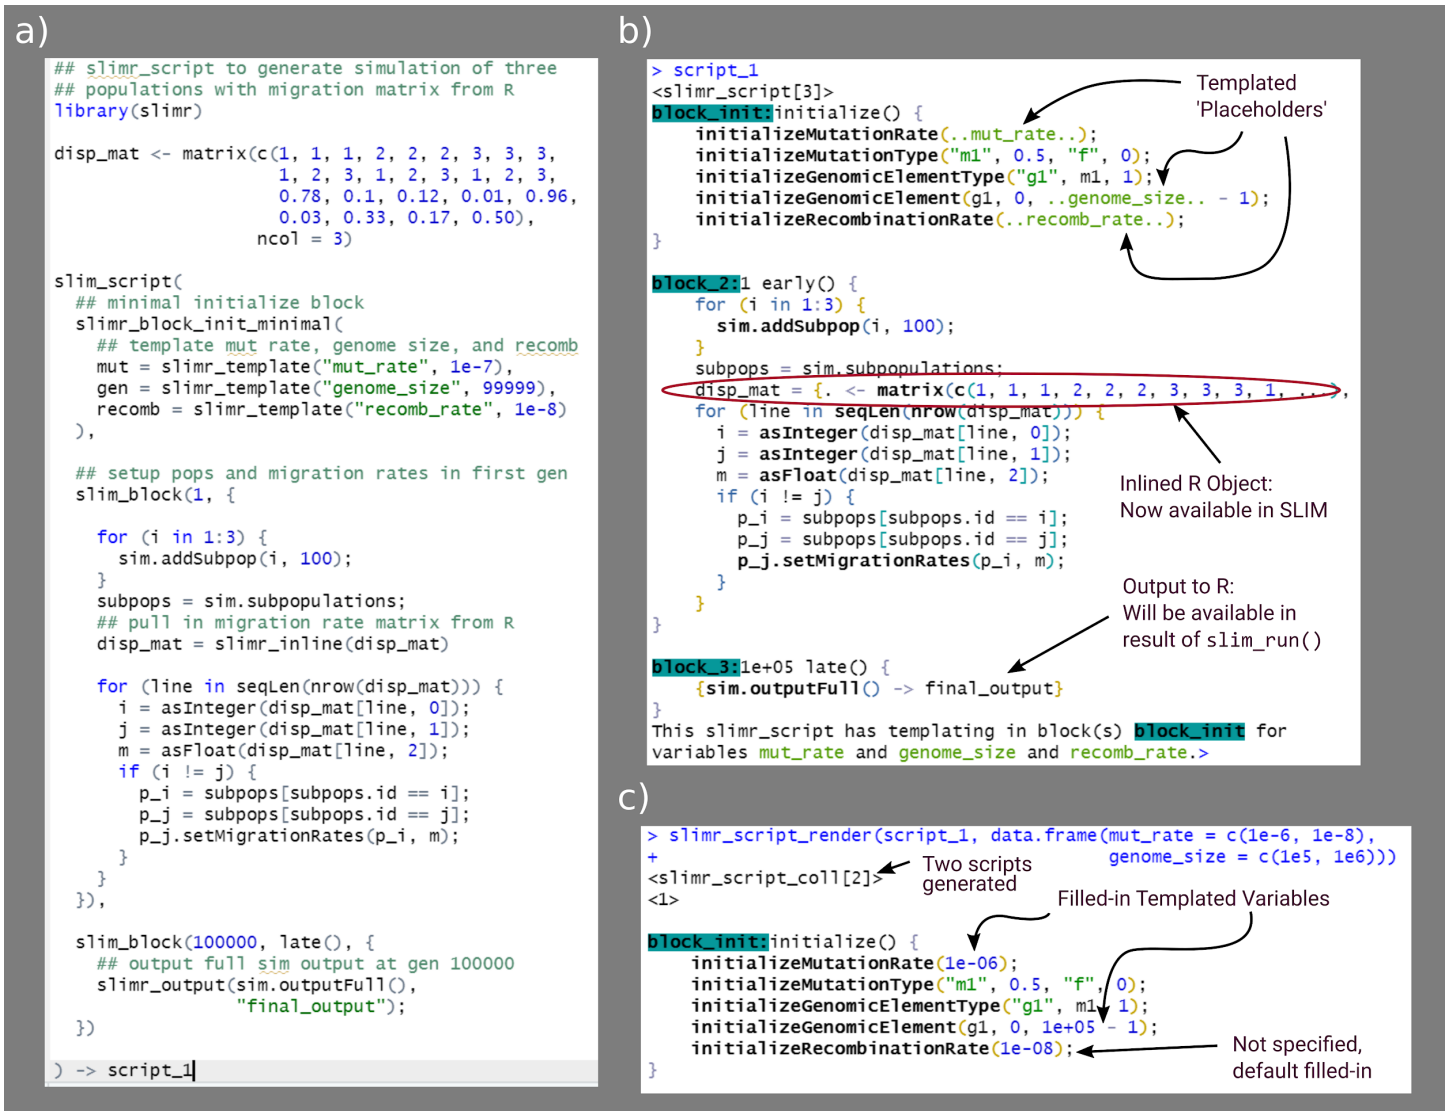

Figure 2. Example of a single script using the main `slimr` verbs (`slimr_template`, `slimr_output`, and `slimr_inline`). A) Code to specify the `slimr_script`. B) Pretty printing of the script, showing special `slimr` syntax. C) Example of running `slim_script_render` on the `slimr_script` object, demonstrating how placeholder variables specified in slimr_template are replaced with provided values. All code from the above example can be accessed as a package vignette (https://rdinnager.github.io/slimr/articles/simple_example.html).

`slimr_inline(object, delay = FALSE)`

`slimr_inline` allows slimr users to embed (or “inline”) an R object directly into a SLiM script so that it can be accessed from within a SliM simulation. This is a powerful way to use empirical data that has been generated, loaded, and / or cleaned from within R within a simulation. `slimr_inline` automatically detects the type of R object and attempts to coerce it into a format compatible with SliM. Currently supported types are all atomic vectors, matrices, arrays, and Raster* objects from the raster package, which will allow users to insert maps for use in spatial simulations.

slimr_output(slimr_expr, name, do_every = 1, callbacks = NULL)

`slimr_output` makes it simple to output data from the simulation by wrapping a SLIM expression. Where it is called in the `slimr_script`, it will produce SLiM code to take the output of the expression and send it to R. The output will be available after running `slim_run` in the returned object as a data.frame. Output can even be accessed live during the simulation run via the use of callback functions. The do_every argument tells `slimr_output` not to output every time it is called, but rather only after every `do_every` generations.

`slimr` includes several functions to create different commonly desired outputs and visualizations, which use the `slimr_output_` prefix (e.g. `slimr_output_nucleotides()`, which outputs DNA sequences data for nucleotide-based simulations).

Metaprogramming

Metaprogramming is programming that generates or manipulates programming code itself. `slimr` has facilities for manipulating SLiM programming code and generating scripts. The main `slimr` verb for doing this is `slimr_template`. `slimr` also supports the metaprogramming operators for forcing (!!) and splicing

(!!!), as used in the `{rlang}` R package. Here we briefly describe `slimr_template`, designed to help users easily generate many versions of a slimr_script with different parameters.

slimr_template(var_name, default = NULL)

`slimr_template` allows the user to insert "templated" variables into a `slimr_script;` the call to `slimr_template` will be replaced in the SliM script with a placeholder with the name `var_name`. This placeholder can be replaced with values of the user's choice by calling `slim_script_render(slm_scrpt, template = tmplt)`, and providing a template – a list or data.frame containing values with names matching `var_name`. This action can be performed on multiple `slimr_template` variables simultaneously, as well as producing multiple replicate scripts with different combinations of replacements. This feature can create a swathe of parameter values to be run (automatically) in parallel to explore parameter space, conduct sensitivity analyses, or fit data to simulation output using methods requiring many simulation runs, such as Approximate Bayesian Computation (ABC). Users can provide a default value for each templated variable, which will be used if the user does not specify a replacement for that variable.

These features together make `slimr` far more than a simple wrapper for SliM – its goal is to enhance and complement SliM by creating a hybrid domain specific language for R. We plan to continue to increase integration of our package with SliM, and to continuously update it as new SliM versions are released in the future.

`slim_run(slimr_ob, slim_path, parallel = FALSE)`

Once a `slimr_script` or `slimr_script_coll` object has been created, with all SLiM simulation logic and `slimr` verbs for interacting with R, it can be sent to the SLiM software to be run using the `slim_run` function. To access this functionality, users must install SLiM on their computers. This is facilitated by the

`slim_setup()` function, which will attempt to automatically install a platform appropriate version of SLiM on the user's system and set it up to work with `slimr`.

Calling `slim_run` will run the simulation. While the simulation is running, slimr_run produces progress updates if requested, as well as any output generated by calls to `slimr_output` with custom callbacks. If called on a `slimr_script_coll` containing multiple `slimr_script` objects, each `slimr_script` object will be run, optionally in parallel, and the result returned in a list.

Once finished, `slim_run` wil return a `slimr_results` object, which contains information about the simulation run, such as whether it succeeded or failed, any error messages produced, all output generated from slimr_output calls, and any file names where additional data from the run are stored. This can then be used for any downstream analysis the user desires.

Example: Investigating population genomics of small mammals in a periodic environment

In this section we provide a brief description of a full example analysis using simulation (Fig. 3, Fig. 4). Full code for the example can be found in the supplementary material.

The context for this example is a long-term ecological study in the Simpson Desert in central Australia. Several authors of this paper have studied the population dynamics of small mammals and reptiles in this desert for more than 30 years (Dickman *et al.* 2014; Greenville *et al.* 2016, 2017). Recently, we have begun sequencing tissue samples taken from animals captured during the past 15 years, and obtained single nucleotide polymorphism (SNP) data using DArT technology. Here, we use SNP data from 167 individuals of a common native rodent species, the sandy inland mouse *Pseudomys hermannsburgensis*, sampled at 7 sites over three years (2006-2008), and subsequently aggregated to 3 subpopulations for

analysis. The three sample years span periods before and after a major rainfall event at the end of 2006; big rains occur infrequently in the study region (every 8-12 years) (Greenville *et al.* 2012) but drive major population eruptions.

We used the SNP data to calculate pairwise Fst values among the three subpopulations in each year, revealing that pairwise Fst values dropped rapidly to nearly zero immediately after the rainfall event from a high recorded just prior to the event when the populations were more genetically differentiated. We interpreted this result to mean that the rainfall event, which caused the sandy inland mouse population to rapidly increase, also allowed animals to move out of spatially scattered refuge patches to which they had been confined during the preceding dry period (Dickman *et al.* 2011). This movement allowed the subpopulations to mix, leading to a decrease in population genetic structure as measured by Fst.

In the example, we use simulations to evaluate our interpretation regarding the processes driving changes in Fst values. We found that our initial hypothesis, that the rainfall event led to the mixing of previously unconnected populations in refuge patches, provided a good match to the data when we simulated the population and genomic processes. However, we also identified several other processes that could generate similar outcomes, which raises the question as to what data or analyses would be required to distinguish among these competing processes.

The results from these preliminary simulations will thus be invaluable in guiding which individuals and time periods we should focus sequencing on to maximize the chances of distinguishing among competing hypotheses that might explain the combined population and genetic patterns in the data. Ultimately, we aim to use this approach to understand how future climate change could alter the population and genetic structure of desert animals, highlighting the value of `slimr` in a scientific workflow.

a)

```
pop_sim <- slim_script(

  slim_block(initialize(), {

    initializeMutationRate(slimr_template("mut_rate", 1e-6));
    initializeMutationType("m1", 0.5, "n", 0, slimr_template("selection_strength", 0.1));
    initializeGenomicElementType("g1", m1, 1.0);
    initializeGenomicElement(g1, 0, slimr_template("genome_size", 50000) - 1);
    initializeRecombinationRate(slimr_template("recomb_rate", 1e-8));
    initializeSex("A");
    defineConstant("abund", slimr_inline(pop_abunds, delay = TRUE));
    defineConstant("sample_these", slimr_inline(sample_these, delay = TRUE));

  }),
  slim_block(1, {

    init_pop = slimr_inline(init_popsize, delay = TRUE)

    ## set populations to initial size
    sim.addSubpop("p1", asInteger(init_pop[0]));
    sim.addSubpop("p2", asInteger(init_pop[1]));
    sim.addSubpop("p3", asInteger(init_pop[2]));

  }),

  slim_block(1, late(), {
    ## get starting population from a file which we will fill-in later
    sim.readFromPopulationFile(slimr_inline(starting_pop, delay = TRUE));
    ## migration on or off flags for pops 1-3 (using tag)
    p1.tag = 0;
    p2.tag = 0;
    p3.tag = 0;
  }),

  slim_block(1, 1000, late(), {

    ## update generation number
    gen = sim.generation %% 50
    if(gen == 0) {[redacted]}

    ## set population size to observed levels
    p1.setSubpopulationSize(asInteger(ceil(abund[0, gen - 1] * slimr_template("popsize_scaling", 100))));
    p2.setSubpopulationSize(asInteger(ceil(abund[1, gen - 1] * ..popsize_scaling..)));
    p3.setSubpopulationSize(asInteger(ceil(abund[2, gen - 1] * ..popsize_scaling..)));

    ## increase migration when above abundance threshold
    if(p1.tag == 0 & abund[0, gen - 1] > slimr_template("abund_threshold", 5)) {
      p2.setMigrationRates(p1, slimr_template("migration_rate", 0))
      p3.setMigrationRates(p1, ..migration_rate..)
      p1.tag = 1;
    }
    if(p1.tag == 1 & abund[0, gen - 1] <= ..abund_threshold..) {
      p2.setMigrationRates(p1, 0)
      p3.setMigrationRates(p1, 0)
      p1.tag = 0;
    }

    if(p2.tag == 0 & abund[1, gen - 1] > ..abund_threshold..) {[redacted]}
    if(p2.tag == 1 & abund[1, gen - 1] <= ..abund_threshold..) {[redacted]}

    if(p3.tag == 0 & abund[2, gen - 1] > ..abund_threshold..) {[redacted]}
    if(p3.tag == 1 & abund[2, gen - 1] <= ..abund_threshold..) {[redacted]}

    if(any(match(sample_these, sim.generation) >= 0)) {
      ind_sample = sample(sim.subpopulations.individuals, 50)
      slimr_output(ind_sample.genomes.output(), "pop_sample", do_every = 1);
    }

  }),

  slim_block(1000, late(), {[redacted]})

)
```

b)

```
Console   Terminal ×   Jobs ×
D:/Projects/slimr_manuscript/
> result <- slim_run(slimr_script_render(pop_sim))
[==========================================================] | Gen:1000/1000(100%) Time: past:00:00:00|left: 0s


Simulation finished with exit status: 0

Success!
```

Figure 3. Screenshots of the main `slimr` example from this manuscript. **a)** Rstudio screenshot showing completed slimr script for specifying the model. **Note:** some code sections have been collapsed for brevity. **b)** Screen shot of the Rstudio console after running the example using default values for templated variables. This shows how fast SLiM can run – this example took less than 1 second!

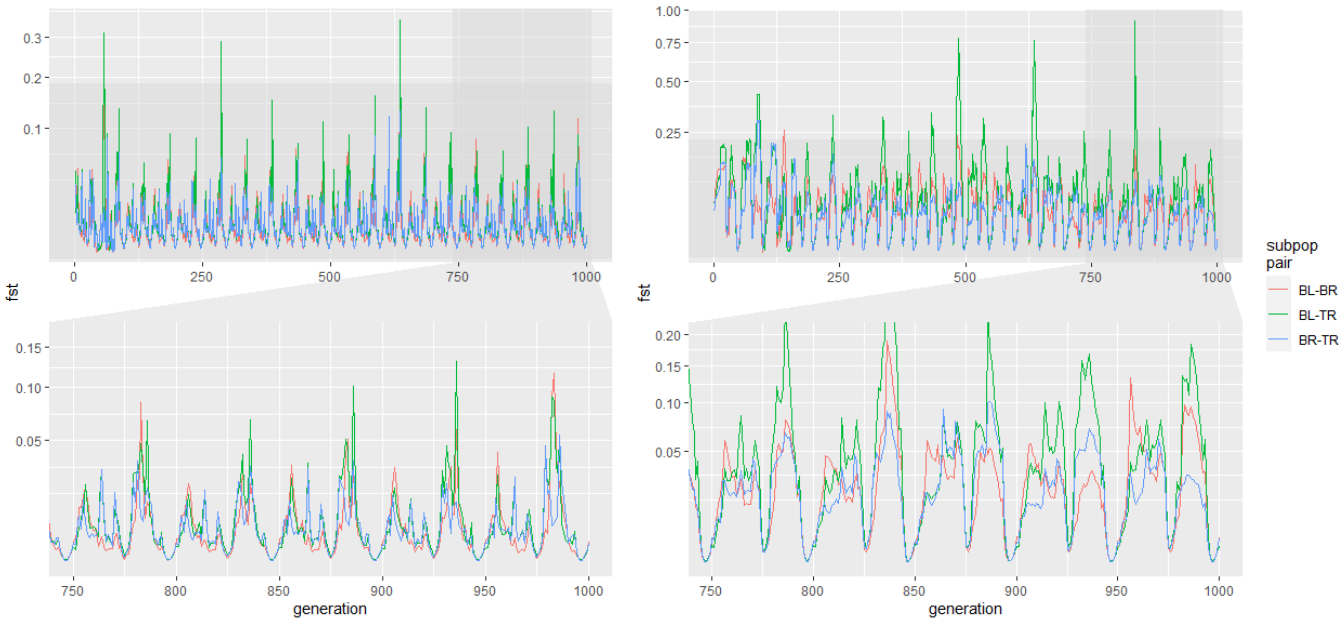

Figure 4. Two very different evolutionary scenarios that produce similar patterns of fluctuations in pairwise Fst between three subpopulations when simulated using `slimr`. **Left Panel:** Maximum migration rate between subpopulations is at its maximum value (panmictic), selection is relatively high, and migration is "on" all the time. **Right Panel:** Maximum migration rate is 0.2 (20% population exchange), there is no selection, and migration is only "on" when total population size is high (e.g. during and just after rainfall events), otherwise there is no migration.

Acknowledgments

We thank Benjamin C. Haller and Phillip W. Messer for permission to reproduce the documentation and examples of SLiM in `slimr`, and for valuable feedback on the package (from B.C.H.). Thanks to Emily Stringer for providing additional test data to help develop the methods used in this manuscript.

Author Contributions

RD, BG, SS, and RPD developed the concept for the package. SE, CD, GW, and AG provided feedback on the package design. RD coded the package and wrote the manuscript draft. CD, GW, and AG contributed data for testing of the package, and BG helped test the package as a user. All authors contributed critically to manuscript drafts and gave final approval for publication.

References

Beaumont, M.A., Zhang, W. & Balding, D.J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, **162**, 2025–2035.

Brehmer, J., Louppe, G., Pavez, J. & Cranmer, K. (2020). Mining gold from implicit models to improve likelihood-free inference. *Proceedings of the National Academy of Sciences of the United States of America*, **117**, 5242–5249.

Carvajal-Rodríguez, A. (2010). Simulation of genes and genomes forward in time. *Current Genomics*, **11**, 58–61.

Cranmer, K., Brehmer, J. & Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences of the United States of America*, **117**, 30055–30062.

Dickman, C.R., Greenville, A.C., Tamayo, B. & Wardle, G.M. (2011). Spatial dynamics of small mammals in central Australian desert habitats: the role of drought refugia. *Journal of mammalogy*, **92**, 1193–1209.

Dickman, C., Wardle, G., Foulkes, J. & de Preu, N. (2014). Desert complex environments. *Biodiversity and Environmental Change: Monitoring, Challenges and Direction*, 379–438.

Greenville, A.C., Dickman, C.R. & Wardle, G.M. (2017). 75 years of dryland science: Trends and gaps in arid ecology literature. *Plos One*, **12**, e0175014.

Greenville, A.C., Wardle, G.M. & Dickman, C.R. (2012). Extreme climatic events drive mammal irruptions: regression analysis of 100-year trends in desert rainfall and temperature. *Ecology and Evolution*, **2**, 2645–2658.

Greenville, A.C., Wardle, G.M., Nguyen, V. & Dickman, C.R. (2016). Population dynamics of desert mammals: similarities and contrasts within a multispecies assemblage. *Ecosphere*, **7**, e01343.

Haller, B.C. & Messer, P.W. (2019). SLiM 3: Forward Genetic Simulations Beyond the Wright-Fisher Model. *Molecular Biology and Evolution*, **36**, 632–637.

Hoban, S. (2014). An overview of the utility of population simulation software in molecular ecology. *Molecular Ecology*, **23**, 2383–2401.

Kelleher, J., Etheridge, A.M. & McVean, G. (2016). Efficient coalescent simulation and genealogical analysis for large sample sizes. *PLoS Computational Biology*, **12**, e1004842.

Marjoram, P., Molitor, J., Plagnol, V. & Tavare, S. (2003). Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the United States of America*, **100**, 15324–15328.

Messer, P.W. (2013). SLiM: simulating evolution with selection and linkage. *Genetics*, **194**, 1037–1039.

Sisson, S.A. (2018). *Handbook of approximate bayesian computation*. Chapman and Hall/CRC, Boca Raton, Florida : CRC Press, [2019].

Strand, A.E. (2002). metasim 1.0: an individual-based environment for simulating population genetics of

complex population dynamics. *Molecular ecology notes*, **2**, 373–376.

Torada, L., Lorenzon, L., Beddis, A., Isildak, U., Pattini, L., Mathieson, S. & Fumagalli, M. (2019). ImaGene: a convolutional neural network to quantify natural selection from genomic data. *BMC Bioinformatics*, **20**, 337.

Wang, Z., Wang, J., Kourakos, M., Hoang, N., Lee, H.H., Mathieson, I. & Mathieson, S. (2020). Automatic inference of demographic parameters using generative adversarial networks. *BioRxiv*.

Yuan, X., Miller, D.J., Zhang, J., Herrington, D. & Wang, Y. (2012). An overview of population genetic data simulation. *Journal of Computational Biology*, **19**, 42–54.