ENMTools User Manual v1.0

Dan Warren, Rich Glor, and Michael Turelli

danwarren@ucdavis.edu

Thanks for using ENMTools! At present, this software should be considered an alpha release – some bits aren't meant to be used yet, and it hasn't been extensively user-proofed yet. It's still entirely possible to tell the software to do things that it shouldn't do, which will produce nonsense results.

I. Installation

I-a. Installing Perl

ENMTools is a Perl script with a graphical user interface that is implemented via the Tk+ package. This means that you need to have Perl and the Tk+ package installed before you can use it. It is unlikely that the versions of Perl and Tk+ needed to run ENMTools are already installed on your computer. First and foremost, you need Perl. If you don't have it, go to:

http://www.activestate.com

Choose Language Distributions -> ActivePerl. The Standard distribution has everything you need, and is free. Download it and install it in just like a regular windows or Mac OSX program. Where you install it doesn't matter to ENMTools, but you'll want to make sure you can easily find this location later.

I-b. Installing Tk+

Once you've got the newest Activestate Perl distribution installed, ENMTools should run right away. If it doesn't, you may need the newest Tk+ distribution. For this, you need to launch the Perl Package Manager. You can find this in the Start menu folder that was created during the Perl installation. You can then either press [ctrl-1] or click the button highlighted in green in the image to the right to show all packages.

Scroll down until you find the package "Tk+" and then right click on it to install.
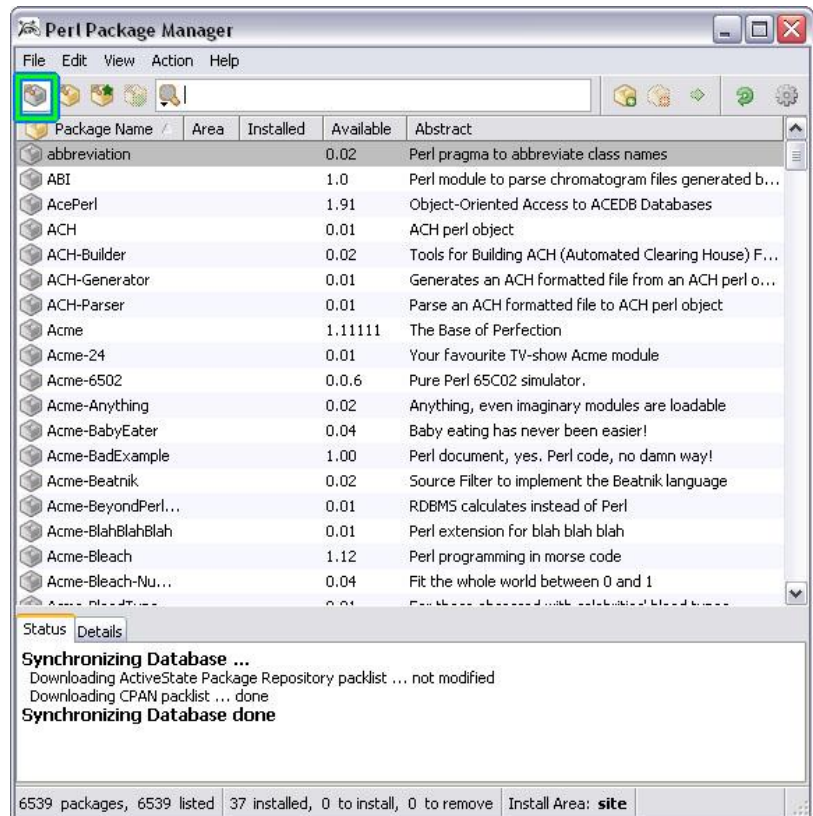
I-c. Launching ENMTools

ENMTools can be launched in several ways. If you are using one of the binary versions, it can be launched simply by double-clicking on the icon for the executable file. The Perl version can be launched this way as well if your system is configured so that Perl is the default application for files with the .pl extension. If this is the case, you can also launch ENMTools from the command prompt by typing the appropriate path (e.g., ,.ENMTools.pl if you are currently in the directory with the perl script).

If your system isn't configured to open .pl files with Perl, you can launch ENMTools by opening a command prompt and typing *<perlpath> <enmtoolspath>*, replacing each of those with the appropriate directory path (e.g., /usr/bin/perl ./ENMTools.pl).

If ENMTools is failing to launch or exiting unexpectedly, try launching it from a command line. This won't fix the problem, but will keep the console from automatically closing so that any error messages can be seen.

II. Running ENMTools

ENMTools menubar has four basic options, each of which provides you with a pull-down menu and suite of associated options: (1) ENM Measurements, (2) Hypothesis Testing, (3) Resampling, and (4) Options.

II.a. The Options Menu

You will need to begin by familiarizing yourself with the Options menu because some basic features of ENMTools and an associated program – Maxent – need to be configured properly before you go any further.

II.a.i. Configuring ENMTools
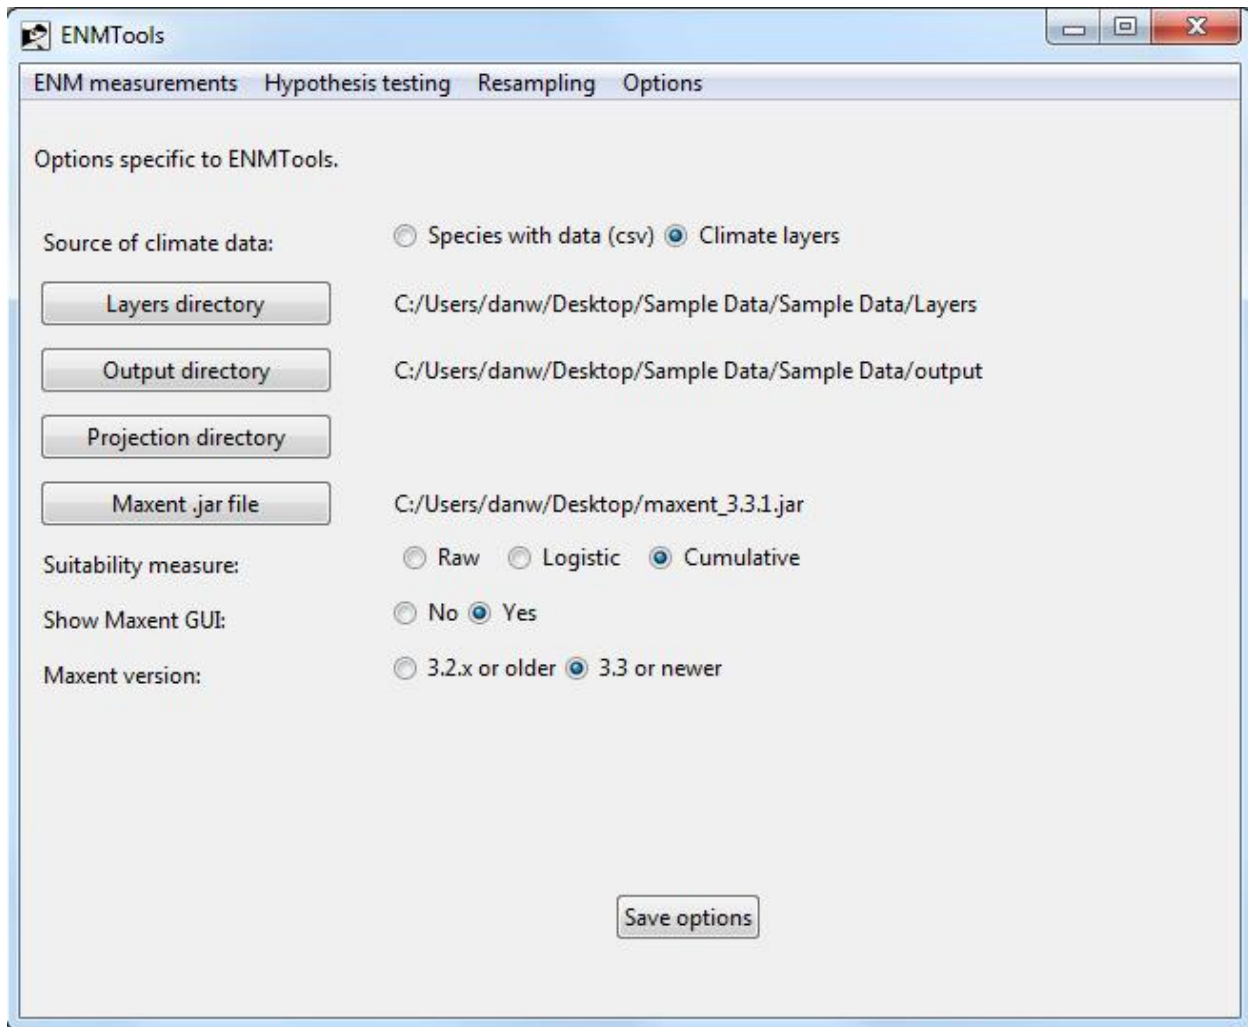
Once ENMTools has launched for the first time, it is essential that you complete a few basic steps to configure the program before proceeding any further. This is necessary because ENMTools needs to know a few things about where your data and other files are before it can do anything (trust us, this required configuration stage will save you time down the road). If it is not properly configured, ENMTools will confess its ignorance at start-up by complaining about a missing config file. To configure ENMTools, you need to go to select the options tab from the ENMTools menu:

1. Tell ENMTools where your environmental data can be found. There are two possible locations for your environmental data: (1) embedded directly in your list of occurrence points ("species with data"), or (2) as a separate file or set of files ("climate layers"). If you select the latter (more frequently used) option, you will then need to tell ENMTools where your climate layers are located by clicking on the "Layers directory" button and navigating to the appropriate folder (once this step is completed, the directory listed to the right of the "Layers directory" button should correspond with the location of your layers).

2. Specify an output directory for your results. ENMTools also needs an output directory where it can create, manipulate, and analyze files. It's a VERY good idea to generate and specify a separate directory

for each analysis you conduct (some ENMTools analyses generate lots of output files and keeping track of these files can quickly become very confusing if you don't keep them organized). You can create a directory in the same fashion that you would create any new directory with your operating system. Select this directory in ENMTools using the "Output directory" button.

3 (necessary only if you have selected the "Species with data" option). Select a set of layers for projection. If you are using the "Species with data" (SWD) option, you will need to have a set of layers to project onto. If you are using SWD, be sure to set your "Projection directory" button.

4. Specify location of Maxent. Many of the analyses conducted in ENMTools are done in association with the program Maxent. For this collaboration to be successful, ENMTools needs to know where Maxent is. You can set this location using the "Maxent .jar file" button. **Make sure you select Maxent's .jar file,** *not* **the .bat file.** Due to changes in the command line arguments for later versions of Maxent, ENMTools also needs to know which version of Maxent you're using. Select the appropriate radio button at the bottom for your Maxent version.

5. Select the suitability measure that you want to use for all Maxent analyses. Maxent is capable of providing output in the form of several different types of suitability measures. To specify the option used for analyses conducted in Maxent, select your preferred option using the "suitability measure" radio button.

6. Set Maxent visibility option. You have the option of seeing, or not seeing, Maxent's interface when ENMTools sends it data. Viewing the interface can be useful when diagnosing problems, but it may also slow down analyses somewhat. You can choose whether to show the interface or not using the "Show Maxent GUI" radio button.

7. Save your configuration. Once you've made all of your choices, save your configuration using the "save options" button. This creates a text file named "ENMTools.config", which will automatically be loaded whenever you start ENMTools.

II.a.ii Maxent Options
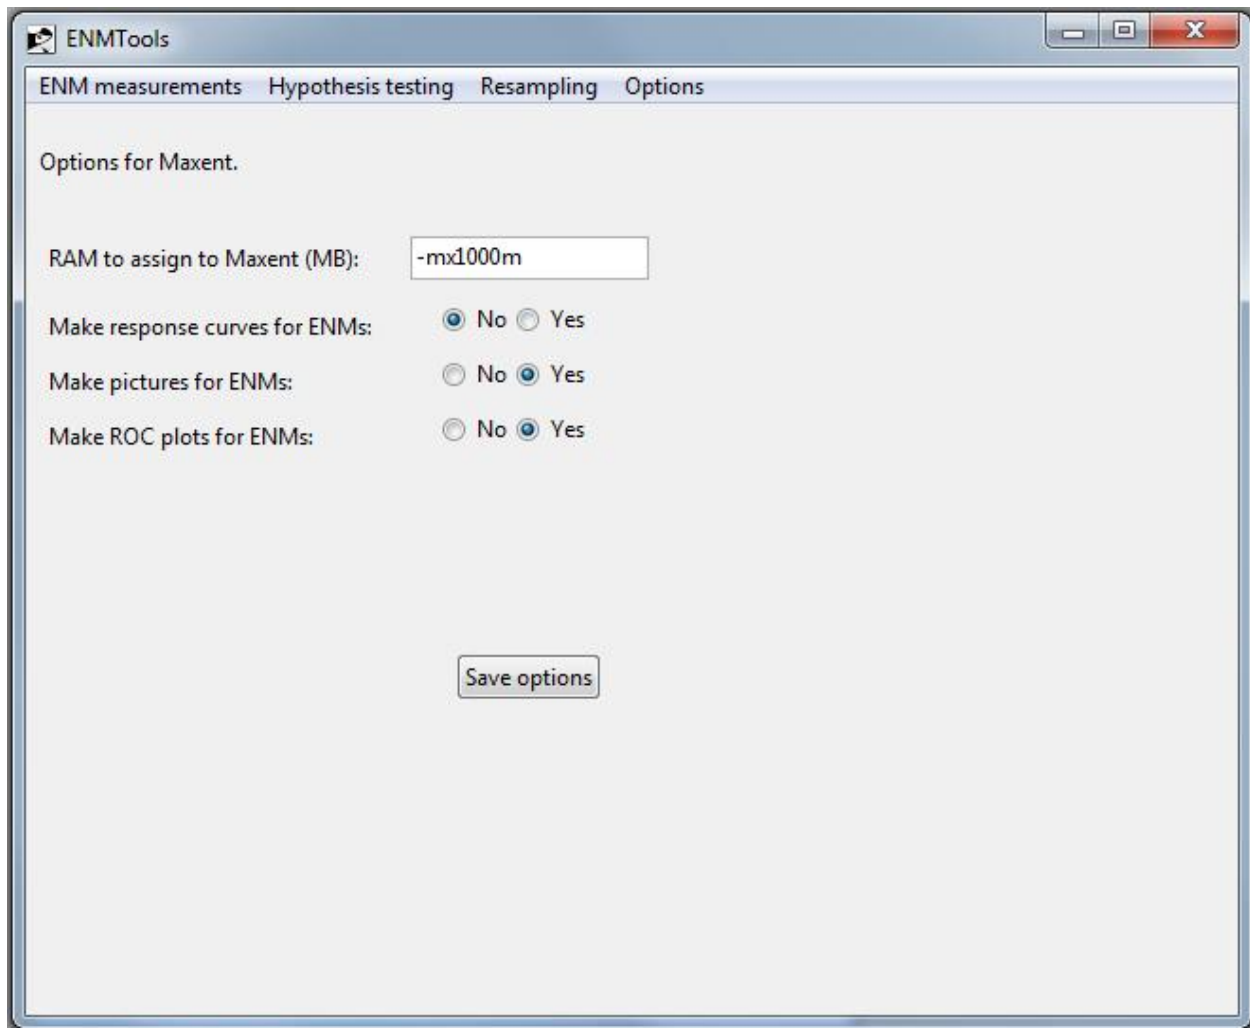
This page, also under the options menu, contains options that ENMTools will pass to Maxent for each run that it executes. At the present time, you are not able to manipulate the complete range of Maxent's options. We will probably enable more options as time goes on. All of currently implemented options should be fairly self-explanatory, with one exception: the "RAM to assign to Maxent" option requires an argument of the form –mx####m, where #### is the number of megabytes of RAM that you want Maxent to use. Make sure that it doesn't exceed the amount of memory available on your system!

At present, there is a bug on some 64 bit Windows systems that limits the amount of RAM that the Java virtual machine will allow you to allocate to Maxent. As far as we've been able to figure out right now, this isn't an ENMTools problem or a Maxent problem – it seems to be a Java problem. If anyone knows of a good solution, please contact Dan at dan.l.warren@gmail.com, we would appreciate the help.

II.b. The ENM Measurements Menu

Selecting the ENM measurements tab permits you to calculate two basic metrics from predicted habitat suitability scores generated by Maxent or some other niche modelling program: (1) quantification of niche overlap between ENMs generated from two or more species and (2) assessment of niche breadth. Conducting these analyses requires that ENMs and associated habitat suitability scores have already been generated by another software package.

II.b.i Measuring niche overlap
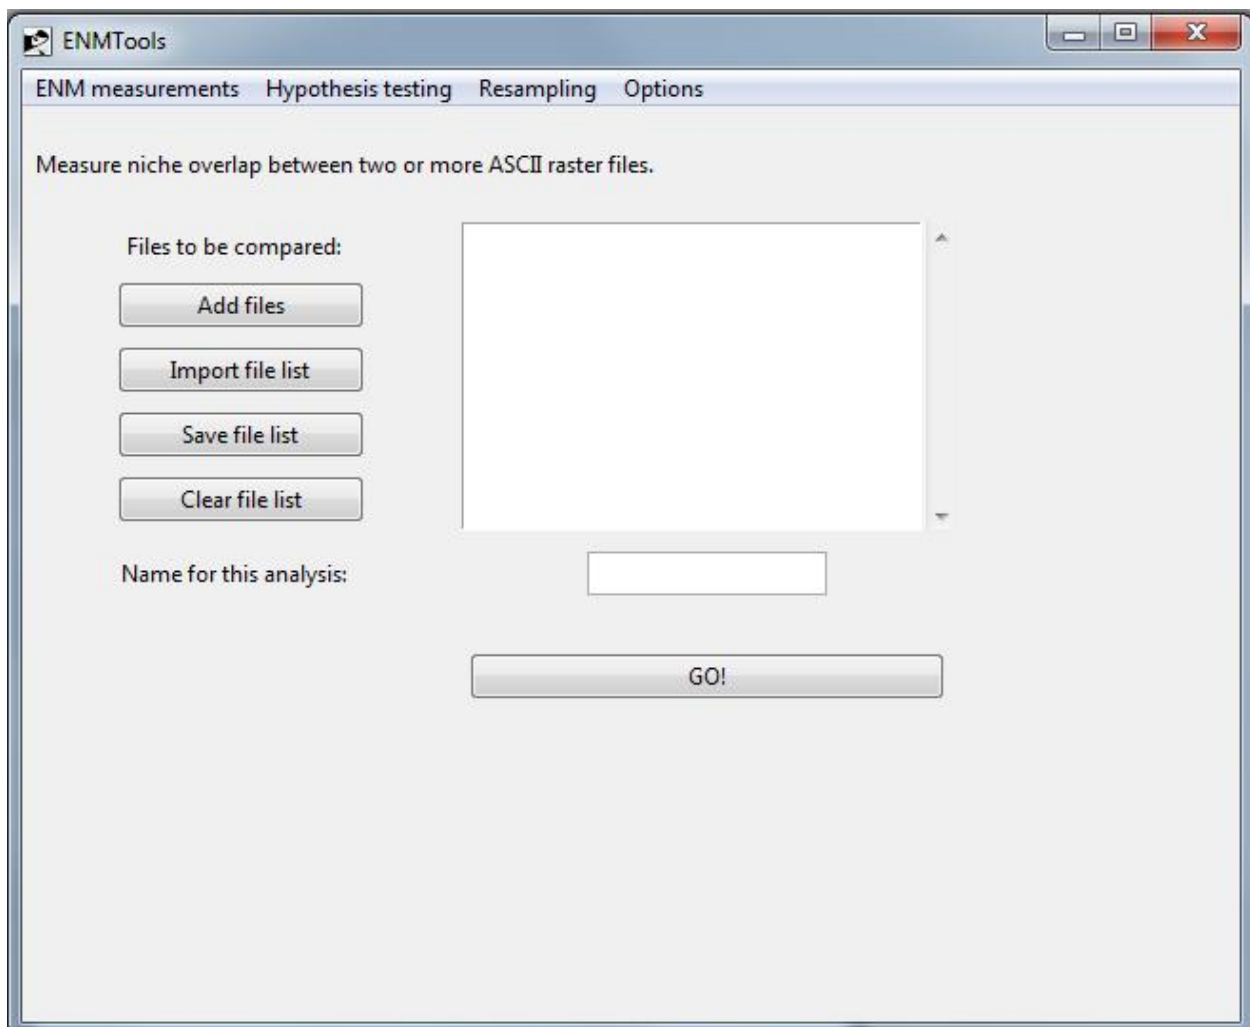
The overlap page is for measuring similarity between predictions of habitat suitability between one or more pairs of populations using methods introduced by Warren et al. (2008). The setup is fairly straightforward – you can add files, save and import lists of files, and name your analysis. The files to be added in this window should be ASCII files of predicted habitat suitability produced by Maxent or some

other method. The program will automatically measure overlap using two different statistics – Schoener's D (Schoener 1968) and the *I* statistic (see Warren et al. 2008 for additional details).

All you have to do to measure overlap is load in a list of .asc files and provide a name for the summary output file. Every pairwise comparison will be made, so if you load in a list of ten files, you're making 45 comparisons. Execution time depends on the number of ASCII files and the resolution of the data.

The output of an overlap analysis is two files, each named using the analysis name you provided. These files can be opened using a text editor or spreadsheet program. Each one contains a table of pairwise overlaps using the statistic listed in the file name. Both *I* and *D* range from 0 (species have completely discordant ENMs) to 1 (species have identical ENMs).

II.b.ii Measuring niche breadth
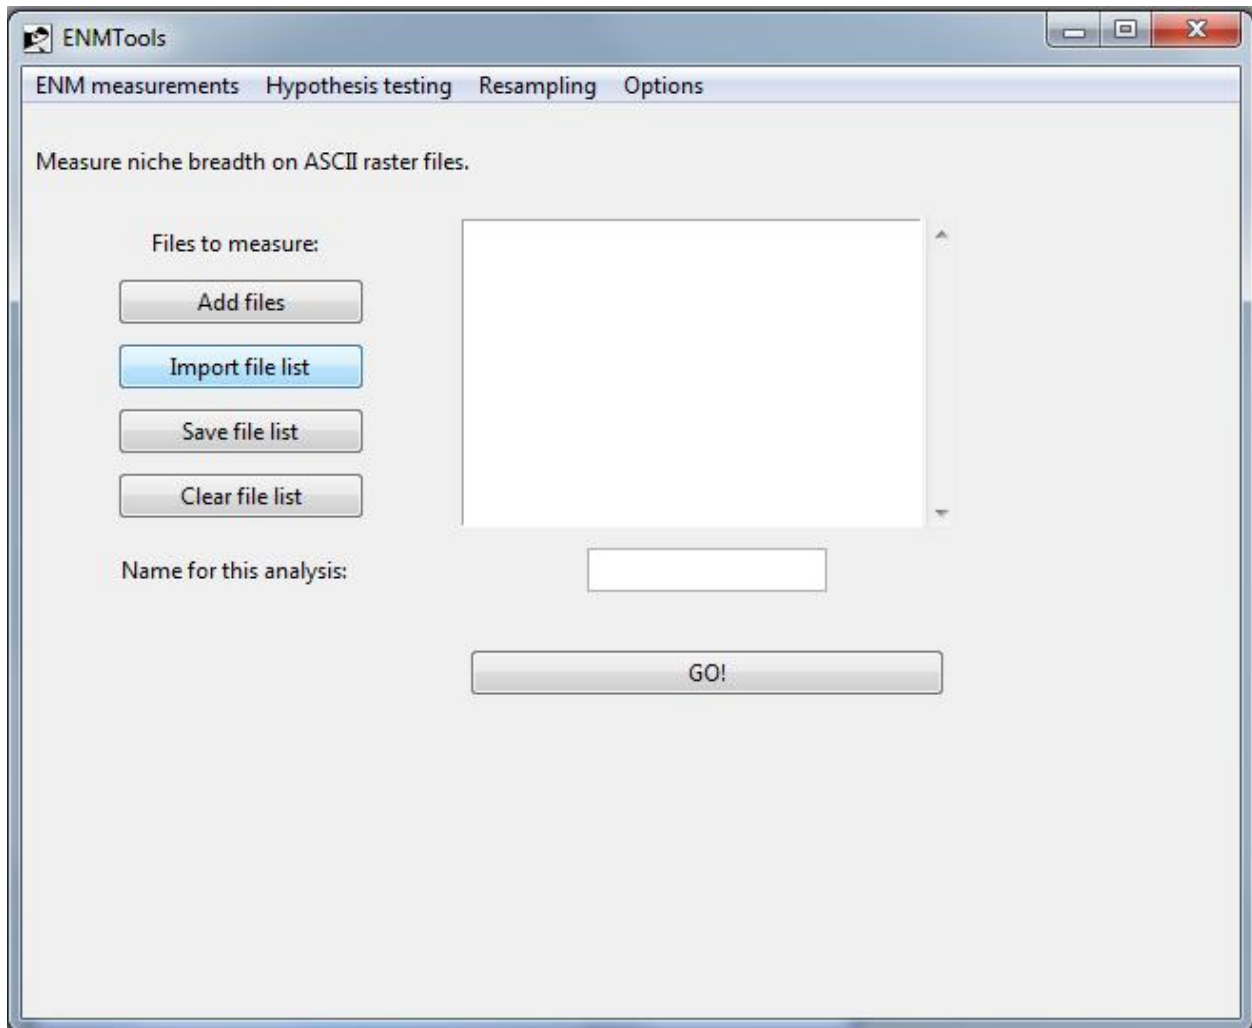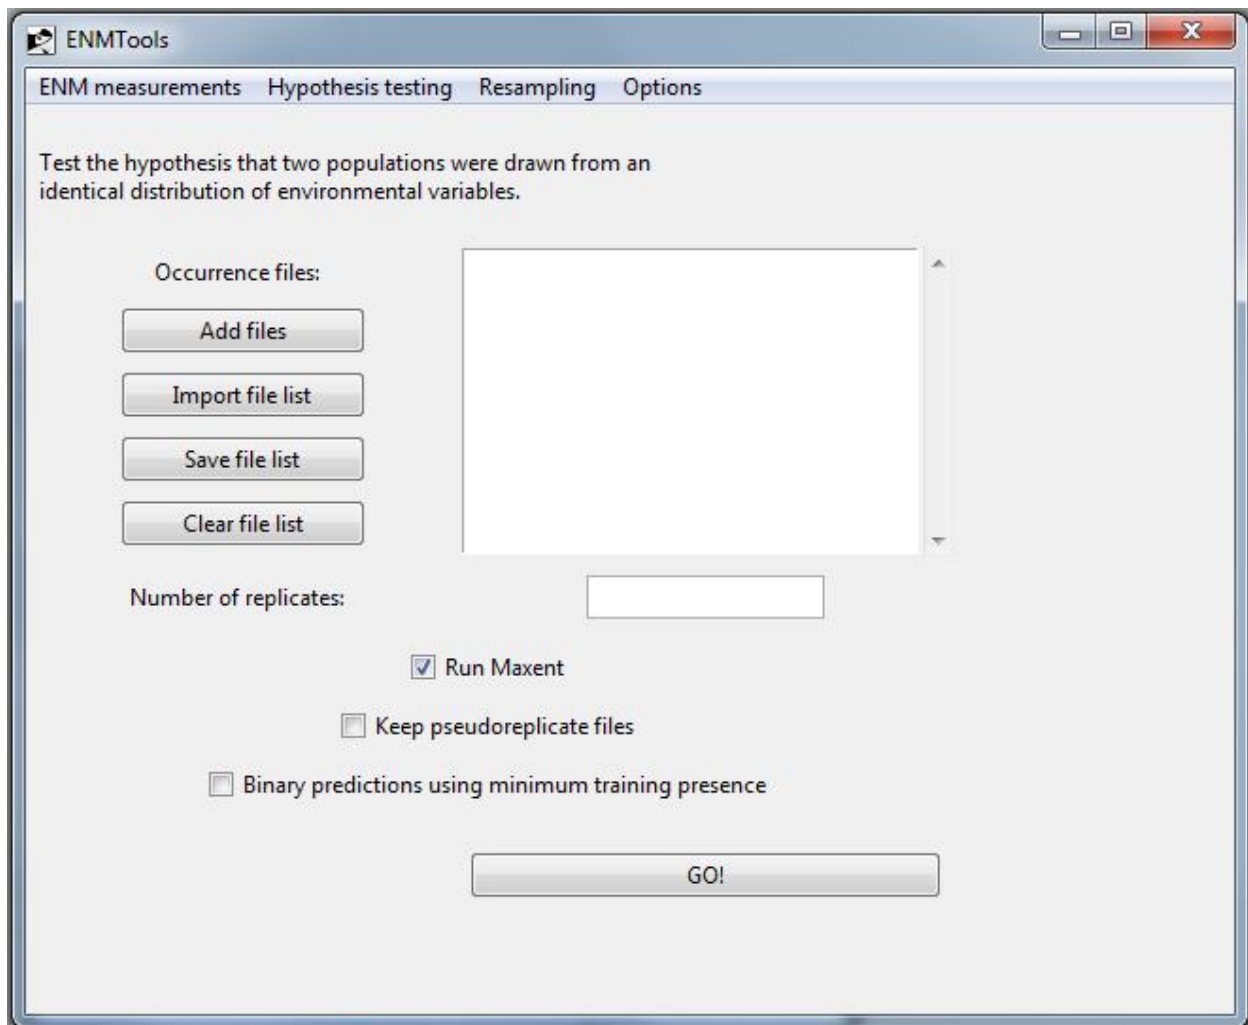
This page works almost exactly like the niche overlap page, except that it outputs a .csv file containing measurements of niche breath for each of the .asc files that you specify. The two niche breadth metrics are those of Levins (1968).

II.c. The Hypothesis Testing Menu

Testing a wide range of comparative hypothesis is ENMTools's raison d'etre. Hypotheses that may be tested with ENMTools include:

1. **The niche identity test.** This test is used to ask whether ENMs generated from two or more species are more different than expected if they are drawn from the same underlying distribution.

2. **The background similarity test.** This test is used to ask whether ENMs drawn from populations with partially or entirely non-overlapping distributions are any more different from one another than expected by random chance.

3. **Range-breaking tests.** This suite of tests is used to ask whether biogeographic boundaries correspond with significant environmental variation.
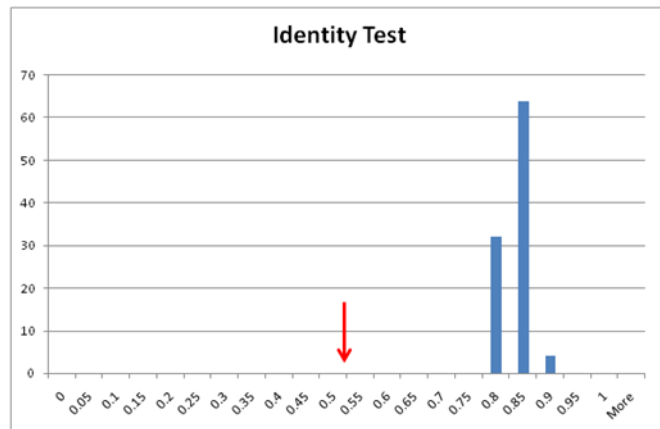
II.c.i. Identity tests

ENMTools

ENM measurements   Hypothesis testing   Resampling   Options

Test the hypothesis that two populations were drawn from an
identical distribution of environmental variables.

Occurrence files:

Add files

Import file list

Save file list

Clear file list

Number of replicates:

☑ Run Maxent

☐ Keep pseudoreplicate files

☐ Binary predictions using minimum training presence

GO!

The niche identity test allows the user to test whether the habitat suitability scores generated by ENM models from two species exhibit statistically significant ecological differences. It does this by pooling empirical occurrence points and randomizing their identities to product two new samples with the same numbers of observations as the empirical data. See Warren et al. (2008) for details.

To the niche identity test you need to use the Add, Import, Save and Clear buttons to select a set of .csv occurrence files for analysis (NOTE: unlike analyses of ENM similarity, the files used for hypotheses testing are .csv files of occurrence points, not .asc files of habitat suitability scores). Once the files with the desired occurrence points have been selected, the program then conduct pairwise identity tests for every pair of occurrence points. Although we've found that it is often easier to simply maintain different .csv occurrence files for each species or population you intend to analyze, hypothesis testing in ENMTools can also accommodate a single .csv with multiple different labels for occurrence points. If a file with multiple different occurrence labels is used, the points with the different labels will be treated as distinct datasets and every possible pairwise comparison will be conducted by default (e.g., if you have one file with species A and B and another with species C and D, you will be doing comparisons A-C, A-D, B-C, and B-D as well as A-B and C-D).

Hypothesis testing is conducted by default on the set of layers that was previously specified in the Options tab. It is important to note that the test of niche identity (and other hypothesis tests conducted in ENMTools) can take a long time, depending on the size and resolution of the spatial data. The output from the niche identity test includes a pair of .csv files with niche overlap scores from each pseudoreplicate. These scores represent the expected degree of niche overlap when samples are drawn from the same distribution (i.e., the pooled sample of occurrence points from two populations). By comparing the overlap between ENM models generated from the actual data for each species (obtained using the overlap tab under the ENM measures menu) to the null distribution obtained using the identity test tab, it is possible to ask whether ENMs produced by two populations or species are statistically significantly different. In the example to the right, the red arrow indicates the measured overlap between species and the histogram illustrates the distribution of overlaps from pseudoreplicates. It is clear from this example that the ENMs built based on the actual occurrences of the two species are more different than expected by chance. This outcome amounts to a rejection of the hypothesis of niche identity.

Options for running the Identity test.

Run Maxent. By default, the "Run Maxent" dialogue box should have a checkmark in it. This simply means that . By de-selecting "Run Maxent", you can tell ENMTools to simply generate input files for pseudoreplicates, but not analyze those files in Maxent. These input files could then be analyzed using other methods of ENM construction.

Keep pseudoreplicate files. The button marked "keep pseudoreplicate files" tells ENMTools whether or not it should delete the files Maxent generates for each pseudoreplicate ENM. When you first begin analyses of a given dataset it will be useful to individually inspect results of each pseudoreplicate to ensure proper behavior of ENMTools and Maxent. However, saving all of the data associated with each pseudoreplicate for a large analysis can quickly eat up all your hard-drive space.

Binary predictions using minimum training presence. This button is used to conduct analyses on simple binary predictions that involve predicted presence or absence, rather than some quantitative measure of habitat suitability. Because Maxent produces quantitative measure of suitability, the use of a binary presence/absence prediction requires simplification of the original output from Maxent. The simplest approach to doing this is to assign a threshold value to the Maxent predicted suitability output scores that corresponds with the minimum training presence (i.e., the lowest predicted suitability score that corresponds with a known occurrence). When this option is used in ENMTools, the identity test is conducted on grids that consist strictly of 0s (species assumed to be absent due to predicted habitat

suitability score lower than the minimum training theshold) and 1s (species assumed to be present due to predicted habitat suitability score higher than minimum training threshold). Some caution should be applied in interpreting the results from such binary predictions: in addition to the general concerns expressed about comparing thresholded niche models in appendix S1 of Warren et al. (2008), there are potentially conceptual problems with comparing the overlap between thresholded models for two species to the overlap seen between thresholded models for subsets of their pooled occurrences.

One word of caution – be sure that you measure your actual overlap using ENMs that were built using the same data and same suitability measure as your pseudoreplicates! Comparing *D* to *I* makes little sense, and comparing overlap of logistic Maxent ENMs to overlap of raw or cumulative ENMs makes even less.

II.c.ii Background tests
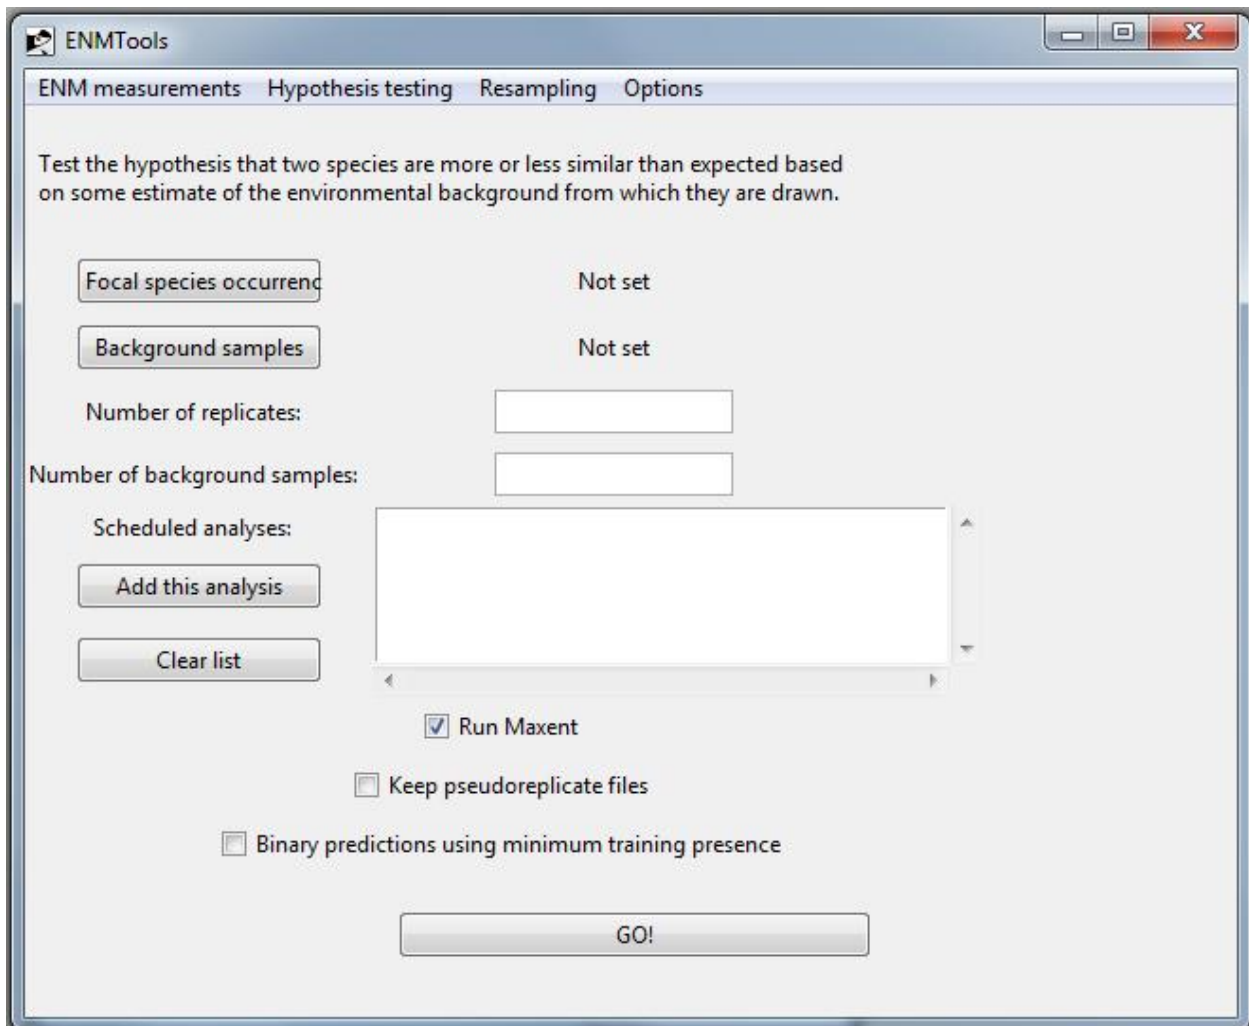
This test uses randomization to determine whether two species are more or less similar than expected based on the differences in the environmental background in which they occur. The identity test above is a very strict condition for assessing ecological similarity – identity is expected to be rejected only when species tolerate the exact same set of environmental conditions *and* have the same suite of environmental conditions available to them. Although the former condition may hold for some species pairs, the latter is unlikely to hold for allopatric species.

To conduct each background test, you need two files. The first file is a .csv file with occurrences for a single species – the focal species for this particular test. The second file is either a .csv file of points (lat-long) from the designated "background" region or an ASCII raster file to use as a mask to generate random background points. If a mask file is provided, any cell designated "-9999" is interpreted as being outside the study area, and any other value is interpreted as being inside. In either case, generating the background file requires a bit of knowledge of a GIS program like ArcGIS. Perhaps the simplest approach is to draw a polygon around the background area in ArcMap and export the region inside this polygon as a new layer that will be used as the mask in ENMTools.

As an example – say you want to compare a pair of species (A and B) to see whether they were more or less similar than expected based on the environmental conditions available to them. You would compare the species occurrences for A to the background for species B, using either a set of points or an ASCII mask for the background area. The number of replicates determines the statistical resolution of the test (100 replicates is enough to get a nominal resolution of 0.01, but because of sampling error this is only an approximation of the true P value). The number of background points you use should be the number of points you have for ***species B*** (if you have 25 points for A and 50 for B, you compare the actual occurrences for A to 50 randomly chosen background points from the area of species B). Once these settings have been made, they are added to the queue by pressing "add this analysis". You would then want to make the same comparison in the opposite direction – compare the points for B to the background for A. When all analyses have been added, press the "Go!" button.

The output from these analyses looks like that of the identity test, and is analyzed similarly, with one exception; a pair of species may be either more or less similar to each other than expected based on the suite of habitats available to them, so this is most often treated as a two-tailed test.

One conceptual issue that arises with this test is the definition and justification of what we consider "background". In some cases (e.g., making comparisons between island endemics or between other allopatric taxa separated by geographical barriers) the appropriate geographic regions are fairly clear. Other situations are not as simple. The results of this test can vary based on the definition of the background areas, so users are strongly encouraged to either (1) have a clear biological justification for defining background regions or (2) conduct sensitivity tests by specifying alternative background regions.

These analyses can be very time-consuming, depending mostly on the size and resolution of the spatial data.

II.c.iii Range Breaking
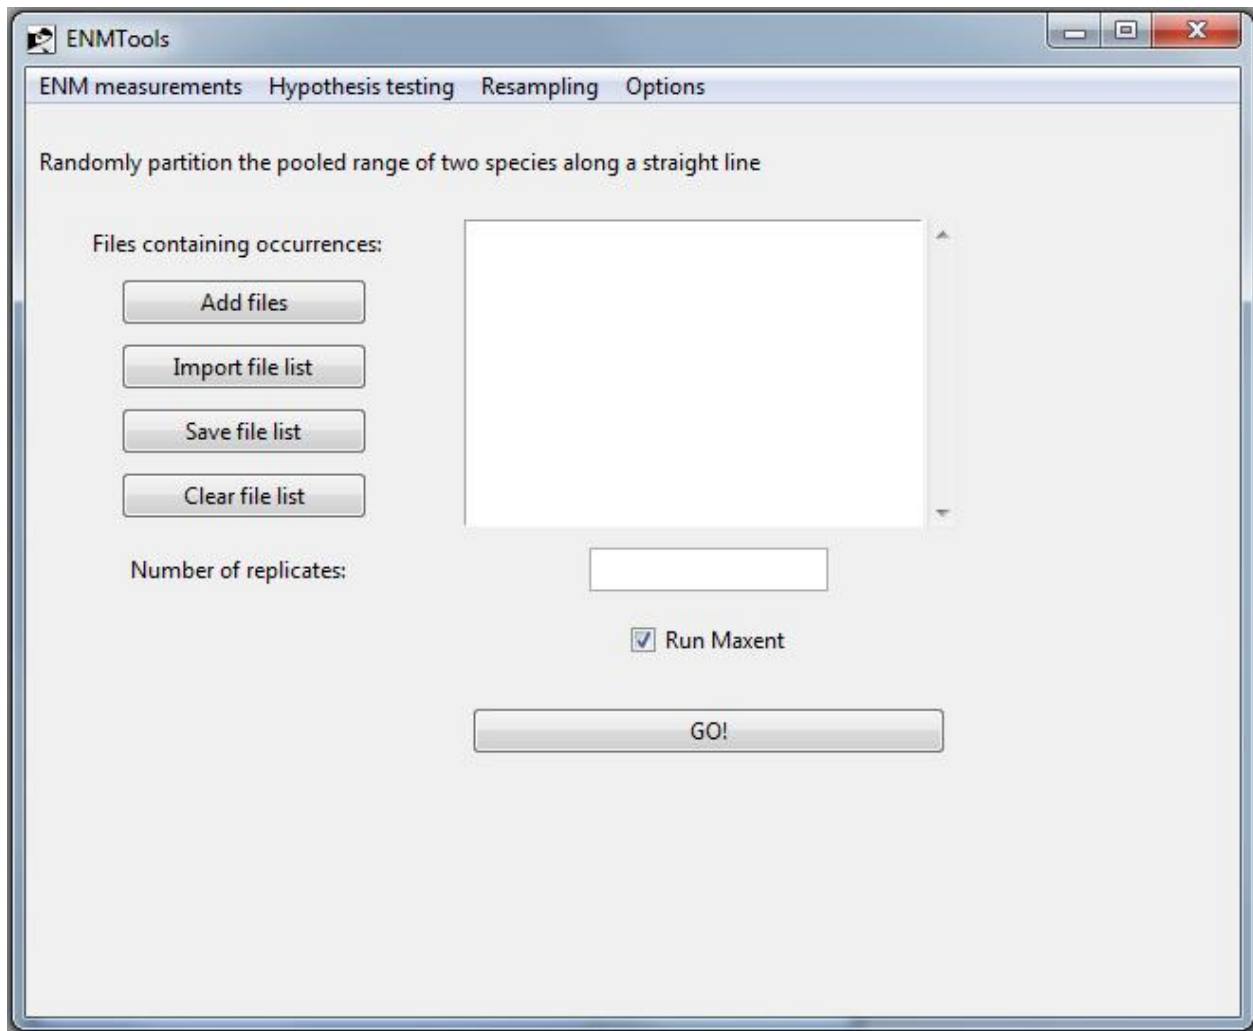
These range-breaking tests are used to ask whether geographic boundaries between species or populations are associated with significant environmental variation (Glor and Warren, *in prep*). Range-breaking tests come in three main flavors: linear, blob, and ribbon. The linear and blob tests are two versions of a test that permit one to ask whether the geographic regions occupied by two species are more environmentally different than expected by chance. The ribbon test, meanwhile, is designed to test whether the ranges of two species are divided by a region that is relatively unsuitable to one or both forms. To demonstrate what's happening with the range breaking tests, take the following examples:

Linear Range-breaking
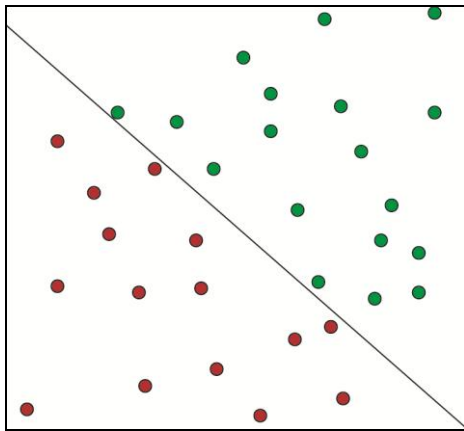
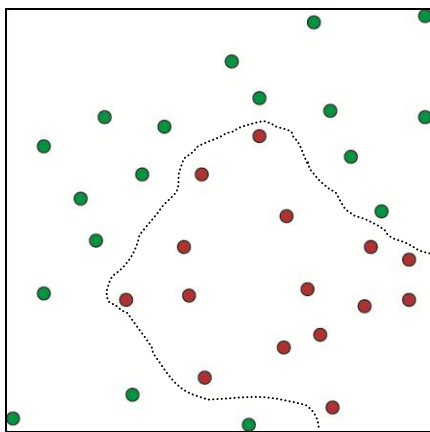**Actual**   **Replicate 1**
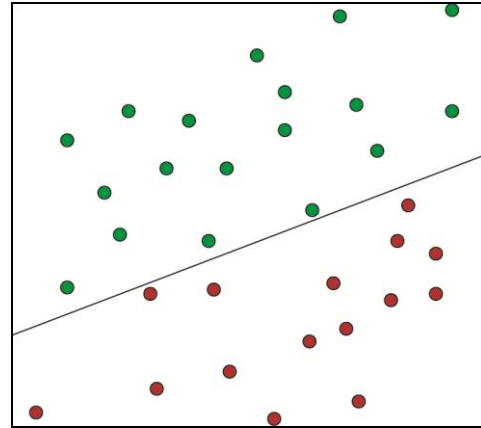
In this example, we have two species (red and green) that are separated by a linear boundary, and we wish to test the hypothesis that this geographic boundary between these species corresponds with a significant environmental boundary, with significance defined as more different than expected by chance. To test this hypothesis, ENMTools randomizes the location of the boundary itself. ENMTools does this by pooling the occurrences for the two species, randomly selecting a slope for the linear boundary, and then choosing an intercept such that the new boundary splits the pooled occurrences into artificial species with the same sample sizes as in our original data set (in this case 18 green and 15 red). ENMTools then construct ENMs for the two newly generated species and measures overlap between them. By doing this many times, ENMTools is able to construct a distribution of expected overlaps across a randomly drawn linear barrier that splits populations into the sample sizes present in our actual data. By comparing the niche overlap measurement from the empirical data to those obtained after randomly drawning the barrier we can ask whether the real barrier is partitioning habitat into more environmentally distinct regions than expected by chance.

There are a few things to be cautious about with the linear range breaking analysis. First, randomly drawing linear barriers like this can become problematic when the pooled set of occurrences is much longer in one spatial direction than in others. This is due to the limited number of ways there are to

bisect very elongate ranges – in the extreme case where all occurrences fall along a single line, there are at *most* two ways to split the pooled set of occurrences. This clearly leads to a lack of independence of pseudoreplicates. The same sort of problem can occur when sample sizes for the two populations are highly asymmetric – when this is the case, occurrence points near the center of the set of pooled ranges will preferentially be assigned to the species with the most occurrences in the original data set, once again leading to a lack of independence of pseudoreplicates. Finally, having many replicate occurrence points for either of the two species (i.e., when a larger portion of one or both of their occurrences are repeated records for the same geographic coordinates) can also lead to a lack of independence of pseudoreplicates, as it restricts the number of ways that the pooled occurrences can be bisected. To avoid a lack of independence among pseudoreplicates it is necessary to carefully examine resulting output and delete duplicate pseudoreplicates (i.e., pseudoreplicates that result in identical partitions).

Blob Range-breaking

The blob method of range breaking is in most respects similar to the linear analysis.


Actual   **Replicate 1**
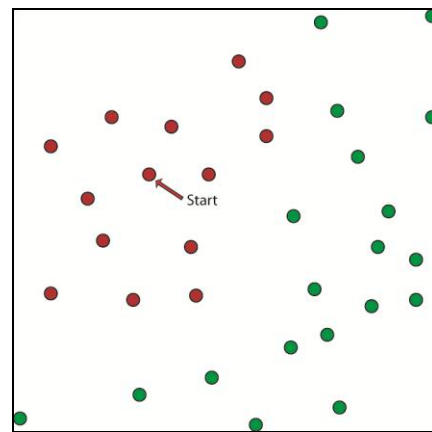
In this example we have two populations that are allopatric, but the barrier between them does not appear to be linear. Instead of comparing the overlap between them to a distribution of overlaps produced by randomly drawn linear barriers, we will compare them to a distribution of overlaps produced by randomly drawing pseudoreplicate populations with non-linear ranges (a.k.a. "blobs"). As with linear range breaking, the analysis begins by pooling the occurrence records for the two populations of interest. We then randomly select a single occurrence point (labeled "start" in the figure above). From this starting point, we continue adding the nearest occurrence points until we construct a data set of the same size as the less numerous of our two species. This effectively draws the smallest possible polygon around our starting point that includes the required number of points. The remainder of the comparison proceeds in the same manner as the linear analysis.

Although the blob method of generating spatial partitions works well in situations where ranges are elongate and sample sizes are asymmetric, having a large number of repeated occurrences can still lead to reduced independence of pseudoreplicates. Again, we recommend deletion of duplicate pseudoreplicates prior to statistical inference from blob range-breaking analyses.
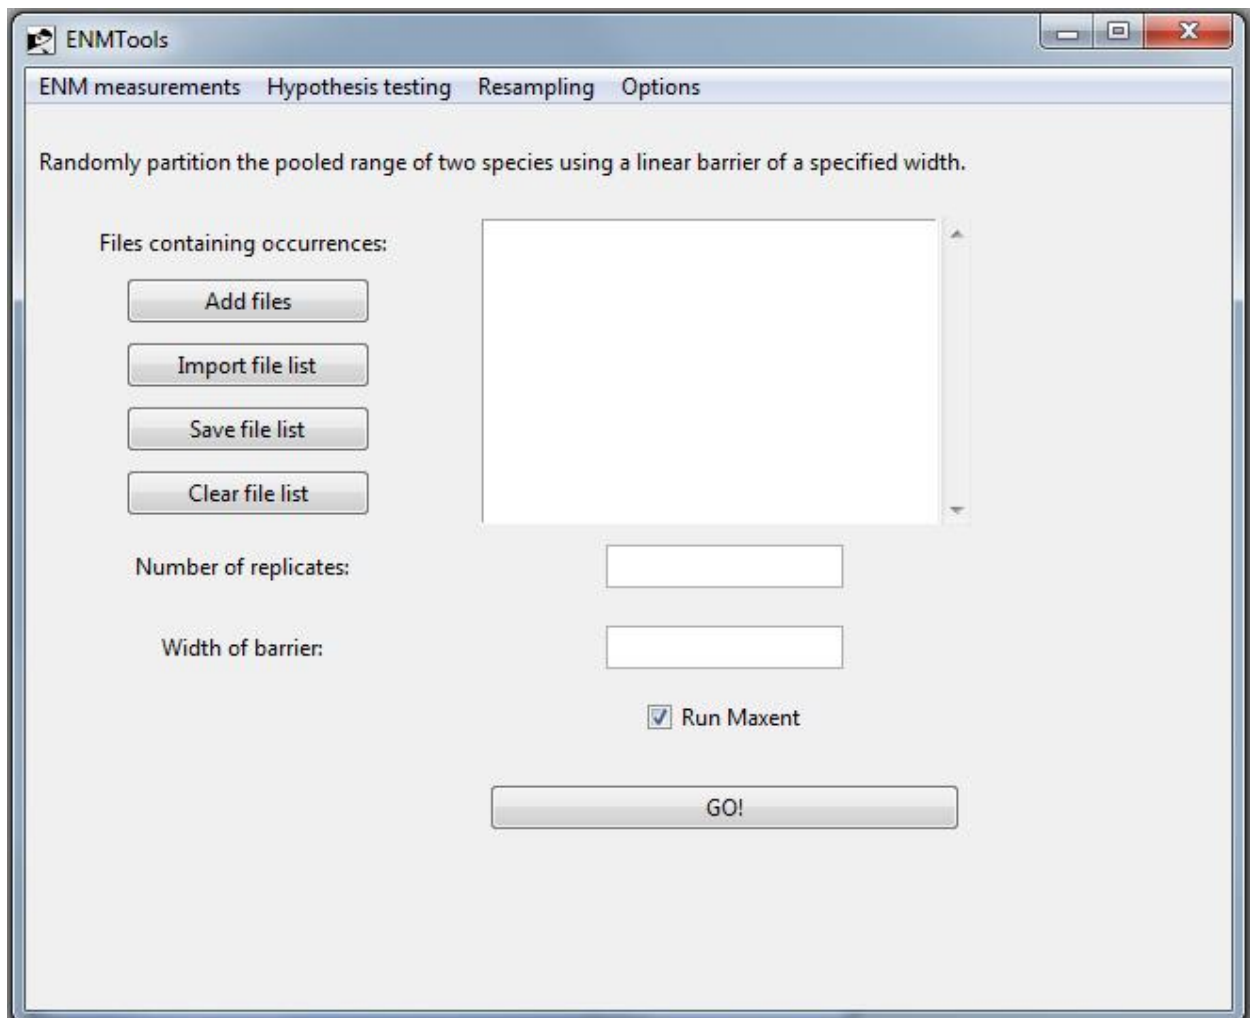
Ribbon

| Actual | Partitioned | Replicate 1 |
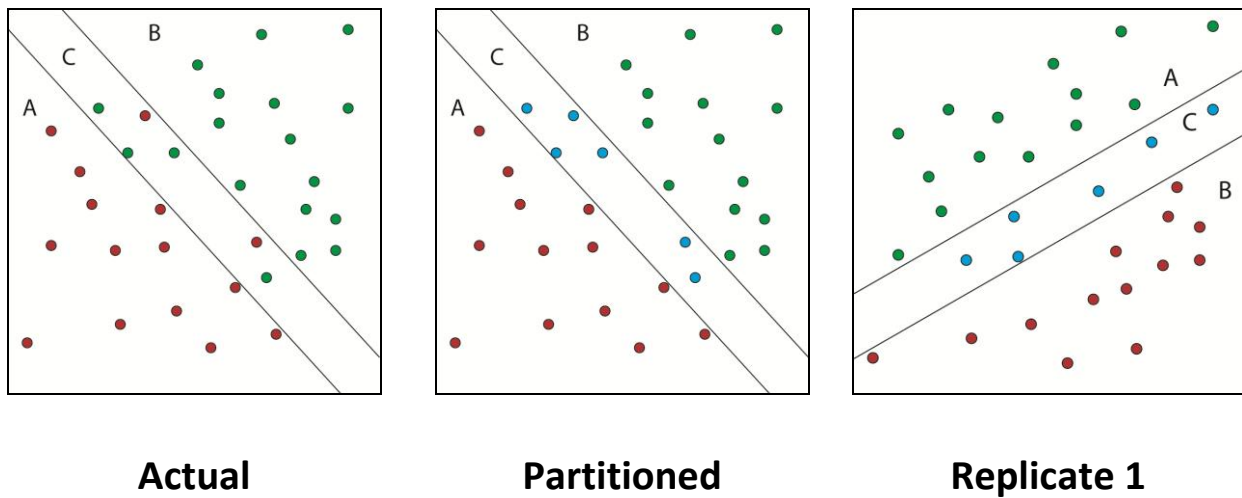
The ribbon range breaking analysis is used to test the hypothesis that two populations occurring in similar, highly suitable habitats are separated by a band of less suitable habitat of a known width (e.g., a band of desert habitat separating two forested regions). In the "Actual" distribution in the above example, the red species is primarily in area A and the green species primarily in area B. Area C is a region where the species' geographic distributions overlap, but is hypothesized to be marginal habitat for both. In this analysis, we pool the samples in the region of overlap, and measure overlap between the two allopatric flanking regions (A vs. B) and each flanking region and the ribbon of unsuitable habitat (A vs. C and B vs. C). Finally, we can measure the overlap between the pooled occurrences of the two species on the outside of the putative barrier and those on the inside (A+B vs. C).

In order to estimate the distribution of expected overlaps for these four measures, ENMTools first picks a slope and intercept that divide our pooled set of occurrences into samples of the size of the two species in our actual data, just as with the linear range-breaking test. Following that, we expand the linear barrier outwards until it is the same width as the band of putative unsuitable habitat that we are studying (i.e., if the desert habitat in the above example is 10 km wide, we construct pseudoreplicates with a 10 km wide ribbon). Because our hypothesis deals with the amount of environmental heterogeneity expected for a band of habitat of a certain size, we keep the width of the barrier (rather

than sample size) constant. This means that sample sizes for A, B, and C may change slightly for different pseudoreplicates, so the ribbon analysis is best applied in situations where sample sizes for both species are high enough that minor deviations in sample size will have little effect on ENMs.

The ribbon range breaking page is identical to the linear and blob range breaking pages, with the exception of the field "width of barrier". The width provided here should be in the same units as the "cellsize" argument in your ASCII environmental layer files (e.g., if your ASCII files are in arc seconds, the barrier width should be also).

II.d. Resampling

II.d.i. Jackknife/Bootstrap
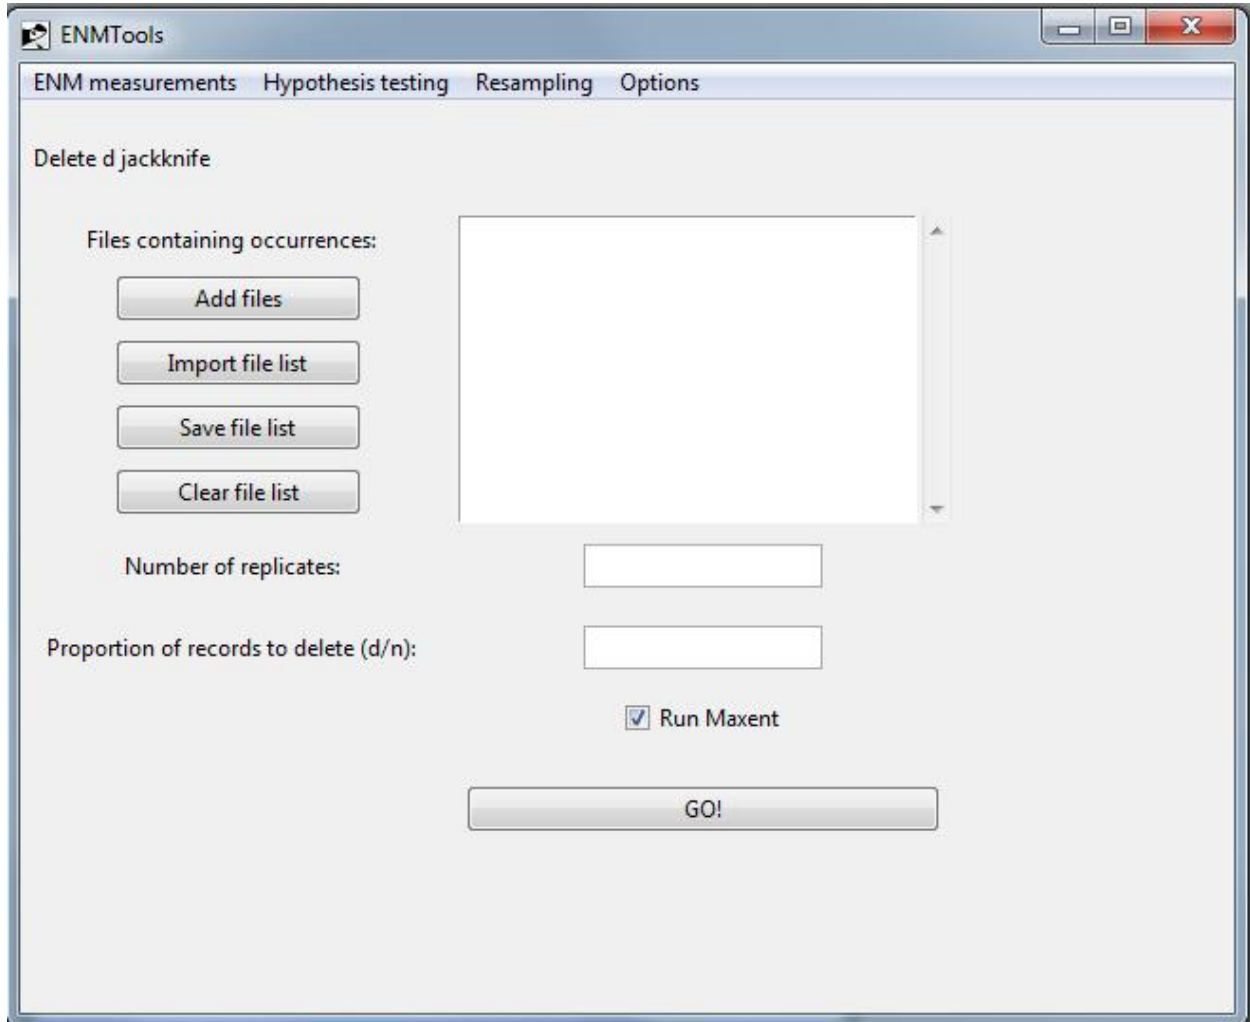
At present, ENMTools offers four methods of randomly resampling data: (1) nonparametric bootstrap, (2) delete D jackknife, (3) delete 1 jackknife (k-fold cross-validation), and (4) retain X jackknife[1].

---

[1] When these resampling analyses were initially written into ENMTools, Maxent didn't allow resampling via its own GUI. Now that Maxent offers similar functionality, the resampling tools in ENMTools may be largely obsolete before they are ever used publicly: most users will likely prefer to conduct their resampling in Maxent due to the availability of confidence intervals on marginal suitability functions. Nevertheless, ENMTools's resampling functions may be useful in some instances, particularly when resampled data sets for use with other ENM construction methods than Maxent are desired.

1. **Nonparametric bootstrap** – This method creates replicates from the original data set by resampling with replacement. The user simply selects a list of files and the number of replicates to execute.

2. **Delete D jackknife** – This method builds replicate data sets by randomly deleting a portion of the data (specified by D) and constructs ENMs using the remainder. For this analysis the user selects a list of files to analyze, the number of replicates to construct and analyze, and the portion of records ($0 < D < 1$) to delete each time.

3. **Delete 1 jackknife** – Also known as k-fold cross validation, this method constructs N pseudoreplicates for a data set of size N, each of which is missing one of the points in the original data set. The user simply selects a list of files, as the number of replicates is determined from the sample size.

4. **Retain X jackknife** – This method builds replicate data sets by deleting all but X occurrences, where X is an integer. Unlike the delete *D* jackknife, this method will keep sample size for pseudoreplicates constant across species even if the sample sizes for the original data sets differ. For example, if X is set to 20 and you analyze two species, one with a sample size of 100 and one with a sample size of 200, replicates for both species will be constructed using 20 points.   For this analysis the user selects a list of files to analyze, the number of replicates to construct and analyze, and the *number* of records ($0 < X < N$) to retain each time.

All of these methods calculate summary statistics (mean and variance) for all runs once Maxent analyses are complete. Due to autocorrelation between pseudoreplicates, a variance inflation factor is applied to the delete *D,* retain *X*, and delete 1 jackknife so that the variance estimates will more accurately represent the true uncertainty in habitat suitability (Efron and Tibshirani 1994, p. 149). Although this procedure is standard for jackknife resampling, there is reason to suspect that it generally overinflates estimates of variance in ENMs when sample sizes are small and *D* is large (Warren and Phillips, *in prep*).

Ideally, all of these methods should produce identical estimates of the uncertainty in suitability scores that is due to incomplete sampling (sampling variance). However, there are reasons why these methods may not be precisely equivalent when applied to ENMs. The nonparametric bootstrap allows individual occurrence points to appear in a single pseudoreplicate multiple times. Given that many ENM construction efforts begin with trimming duplicate occurrence points, some users have expressed concern that re-introducing these duplicates may produce unreliable estimates of variance. If this is the case, one of the jackknife methods may be preferable.

However, jackknife methods necessarily produce pseudoreplicate data sets that are smaller than the original data set, which may alter feature selection and other aspects of the modeling process. This may result in features being selected (or not) during jackknife runs that might not be preferred if the full data set was used. This could affect jackknife's ability to inform variable selection, and may be partially responsible for inflated estimates of variance under the conditions mentioned above. The delete 1 jackknife is unlikely to exhibit this problem to any great degree. A systematic comparison of nonparametric bootstrap and delete *D* jackknife methods as applied to ENMs is currently underway (Warren and Phillips, *in prep*).

The retain X jackknife is primarily useful when comparing measurements on pseudoreplicates between species that have different sample sizes – e.g., making comparisons of niche breadth with jackknife estimates of uncertainty. One might generally expect that estimates of niche breadth would demonstrate some correlation with sample size, particularly when sample sizes are small. For this reason, it makes more sense to compare species using identically-sized subsets of their occurrence points.

II.d.ii. Spatial cross-validation
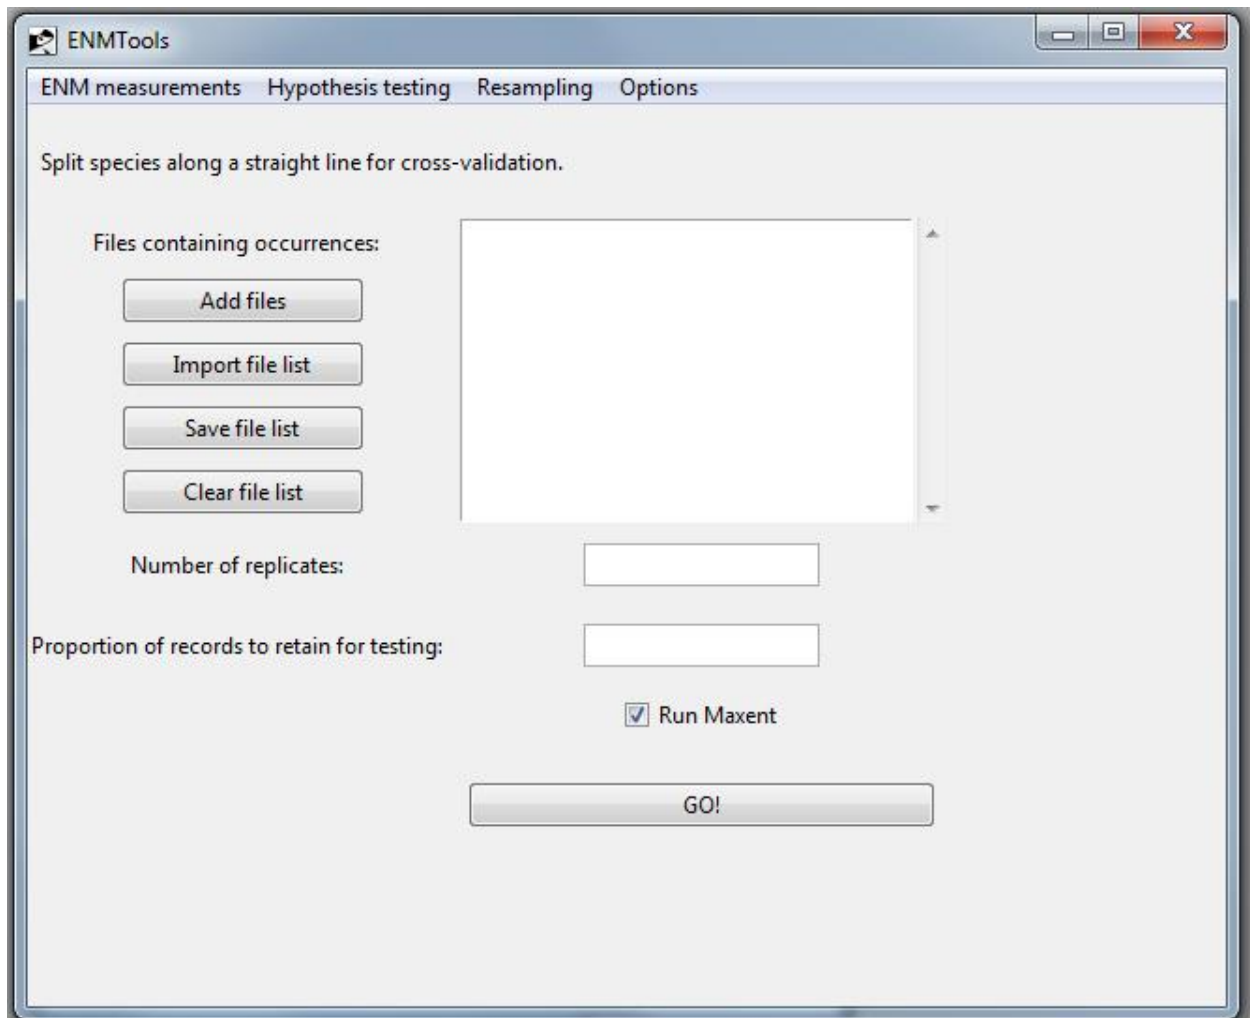
These two analyses use the range-breaking methods described above (section III-c) to randomly partition the data for a single species into two non-overlapping data sets. The program then uses one of the partitions as training data and the other as test data for a Maxent run. These analyses can be useful in studies of sample selection bias, and *may* be relevant to studies of model transferability in some situations (but see Phillips 2008).

Users simply specify a list of files to analyze, the proportion of records to withhold as test data, and the number of replicates to perform. The caveats mentioned in section III-c about the possible non-independence of pseudoreplicate data sets in certain situations also apply to these tests, and must be taken into consideration when using them.

III. Citing ENMTools

We kindly request that users who publish results obtained from ENMTools cite the paper that introduced ENMTools's core methods (Warren et al. 2008), and the application note that formally introduces the program (Warren et al., In press):

Warren, D.L., R. E. Glor, and M. Turelli. 2008. Environmental niche equivalency versus conservatism: quantitative approaches to niche evolution. Evolution 62:2868-2883.

Warren, D.L., R.E. Glor, and M. Turelli. In press…

Users who are analyzing data obtained from Maxent for ENM measurement analyses in ENMTools, or using any of ENMTools's hypothesis testing functions should also cite the paper that introduces Maxent:

Phillips, S. J., R. P. Anderson, and R. E. Schapire. 2006. Maximum entropy modeling of species geographic distributions. Ecological Modelling 190:231-259.

ENMTools users who make use of the range-breaking functions, are also asked to cite the paper that introduced these methods (Glor and Warren, In press):

Glor, R. E. and D. L. Warren. In Press.

Thanks in advance for using and citing ENMTools!

IV. Literature Cited in this Manual

Efron, B., and R .J. Tibshirani. 1994. An Introduction to the Bootstrap. Chapman & Hall/CRC. New York, USA.

Levins, R. 1968. Evolution In Changing Environments. Monographs in Population Biology, volume 2. Princeton University Press, Princeton, New Jersey, USA.

Phillips, S. J., R. P. Anderson, and R. E. Schapire. 2006. Maximum entropy modeling of species geographic distributions. Ecological Modelling 190:231-259.

Phillips, S. J. 2008. Transferability, sample selection bias and background data in presence-only modelling: a response to Peterson et al. (2007). Ecography 31:272-278.

Schoener, T. W. 1968. Anolis lizards of Bimini: resource partitioning in a complex fauna. Ecology 49:704-726.

Warren, D. L., R. E. Glor, and M. Turelli. 2008. Environmental niche equivalency versus conservatism: quantitative approaches to niche evolution. Evolution 62:2868-2883.