EPA/600/R-18/116

User's Manual: SSD Toolbox Version 1.0

Matthew Etterson

US Environmental Protection Agency
Office of Research and Development
Center for Computational Toxicology and Exposure
Great Lakes Toxicology and Ecology Division
Duluth, MN

March 2020

Page intentionally left blank

Contents

Glossary 5
Introduction & distributions supported 5
Normal, logistic, & triangular distributions 6
Gumbel distribution 6
Weibull distribution 6
$Burr_{III}$ distribution 7
Installation 7
Running the SSD Toolbox 8
Starting the program 8
Formatting and Importing Data 8
Fitting a distribution 13
Weighted Fits and censored data 14
Bodyweight scaling 15
Inferential Endpoints 15
The File menu 16
Import data 16
Load previous analysis 16
Save results 16
Clear table 16
Exit 16
The Plot menu 17
Show data 17
Interactive Plotting 17
The context menu 20
Percentiles (HCp) 21
Plot SSD 22
Parameter estimates 23
AIC table 23
Goodness-of-fit 24
Statistical estimation of lack of fit 25
Q-Q plots 26
Posterior Diagnostics 26

BIC Table. 26
Credible Intervals 27
Autocorrelation 27
Trace plots 28
Posterior distributions 28
Delete rows 29
Literature Cited 29

Glossary

Term	Definition
AIC	Akaike Information Criterion
ECDF	Empirical Cumulative Distribution Function
ECx	Concentration expected to cause an effect in x% of test subjects
GUI	Graphical User Interface
GoF	Goodness-of-fit
HCp	Concentration expected to be hazardous to p% of species tested
LC50	Concentration expected to be lethal to 50% of test subjects
LD50	Dose expected to be lethal to 50% of test subjects
MCMC	Markov Chain Monte Carlo
MCR	Matlab Compiler Runtime
SSD	Species Sensitivity Distribution

Introduction & distributions supported

The SSD Toolbox is a program for fitting species sensitivity distributions (SSD). It gathers together a variety of algorithms to support fitting and visualization of simple SSDs. The current version of the toolbox supports six distributions (normal, logistic, triangular, Gumbel, Weibull, and $Burr_{III}$). When any of the first four distributions are chosen, the data are first common-log transformed (log_{10}). When the Weibull or Burr distribution is chosen, the data are fit on their measurement scale. The toolbox also supports fitting distributions using four different methods (maximum likelihood, moment estimators, linearization, and the Metropolis-Hastings algorithm). However, not all fitting methods are available with all distributions (Table 1).

Table 1. Distributions available in the SSD Toolbox

Distribution	Transformation	Maximum Likelihood	Moment Estimators	Graphical Methods	Metropolis-Hastings
normal	log_{10}	Yes	Yes	Yes	Yes
logistic	log_{10}	Yes	Yes	Yes	Yes
triangular	log_{10}	Yes	Yes	Yes	Yes
Gumbel	log_{10}	Yes	Yes	Yes	Yes
Weibull	none	Yes	No	Yes	Yes
$Burr_{III}$	none	Yes	No	No	Yes

Choosing which distribution to fit and deciding between competing distributions is an important and difficult topic. A summary of the six distributions currently available is provided above in Table 1. In Table 2, some properties of the six distributions are given.

Table 2. Some properties of the distributions available in the SSD Toolbox

Distribution	Symmetric[1]	Number of Parameters
normal	yes	2
logistic	no	2
triangular	yes	2
Gumbel	no	2
Weibull	no	2
$Burr_{III}$	no	3

[1]Symmetry here is defined on the transformed scale (*i.e.*, following $\log_{10}$ transformation).

Normal, logistic, & triangular distributions

The normal, logistic, and triangular distributions are among the most commonly used distributions for SSDs. Each has two parameters. In general, differences in fit among these three distributions will be small and likely most apparent in the tails, with the triangular having a finite lower bound, compared to the other two. There are also asymmetric versions of the triangular distribution, with three parameters, but these are not considered here. An illustration of these three distributions fit to the Permethrin aquatic toxicity data can be seen below in Figure 1.

Figure 1. The normal, logistic and triangular distributions fit to a sample of Permethrin aquatic toxicity data using maximum likelihood

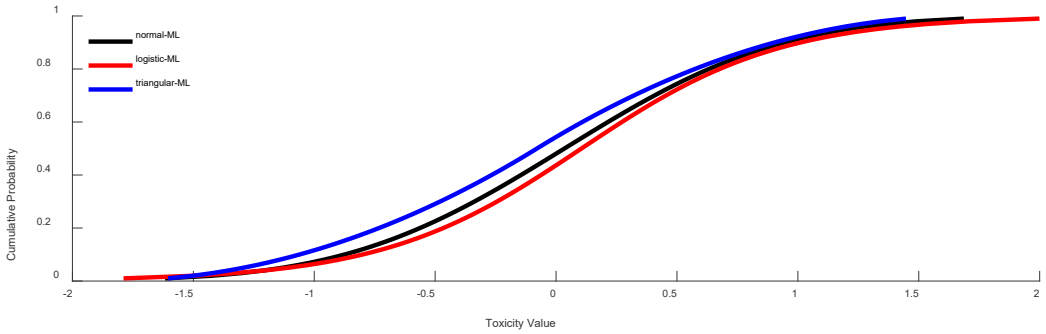

Gumbel distribution

The Gumbel distribution is an asymmetric two-parameter distribution with a heavier right tail than left tail (positive skew). It is also commonly referred to as an Extreme Value Distribution (Type I).

Weibull distribution

The Weibull distribution is an asymmetric two-parameter distribution that may have heavier right or left tail, depending on the value of its scale parameter. As implemented in the SSD Toolbox it is constrained to forms accommodating heavier left tails only (i.e., shape parameter k > 1, see Technical Manual). It is also commonly referred to as an Extreme Value Distribution (Type III).

$Burr_{III}$ distribution

The $Burr_{III}$ distribution is a three-parameter distribution that gives added flexibility in modeling asymmetry in the distribution of toxicity values. It is implemented in the software Burrlioz (https://research.csiro.au/software/burrlioz/). It has been reproduced here for comparative purposes only and does not contain several advantages offered by the Burrlioz implementation, including the limiting behavior to either the reciprocal Weibull or reciprocal Pareto distributions. In practice, this leaves the $Burr_{III}$ implementation here somewhat unstable, especially with small sample sizes, resulting in strong covariances among estimated parameters and sometimes poor optimization of parameter values. This is true for both the Maximum Likelihood and Metropolis-Hastings implementations in the SSD Toolbox. For this, and other reasons, it is not recommended here, but given its widespread use within the Burrlioz software it is provided for comparative purposes.

The preceding information is provided to help you understand and choose distributions *a priori*. The SSD Toolbox also provides considerable functionality for assessing the quality of an SSD *a posteriori* (i.e., after the distribution is fit to a data sample). The remainder of this manual is focused on the mechanics of using the software (which menu or button to use to get a particular result), but considerable guidance on assessing and fitting distributions is provided in the accompanying Technical Manual.

Installation

The SSD Toolbox is freely available on the USEPA Center for Computational Toxicology and Exposure website at:

https://www.epa.gov/chemical-research/species-sensitivity-distribution-ssd-toolbox

To get the latest version of the software, download the zip archive ("SSDToolbox.zip"). This archive contains five files:

1. SSDToolbox.TechnicalManual.March.2020.docx
2. SSDToolbox.UserManual.March.2020.docx
3. SSDToolbox.exe
4. ChlorpyrifosInvertLC50Data.xlsx
5. PermethrinAcuteData.xlsx

The User's Manual (this document) provides step by step instructions for use of the SSD Toolbox. The Technical Manual provides statistical and theoretical background for users wishing to gain a greater understanding of the functions provided in the Toolbox and on the background of SSD fitting. "SSDToolbox.exe" is the software executable (note it cannot be run without first installing the Matlab MCR, see below). "ChlorpyrifosInvertLC50Data.xlsx" are acute toxicity data for invertebrates exposed to Chlorpyrifos taken from USEPA 2016. "PermethrinAcuteData.xlsx" are Permethrin acute LC50 data for a diverse set of aquatic species

(both vertebrate and invertebrate) taken from Fojut et al. (2012). The examples in this User's Manual were generated using both of the above data sets.

The SSD Toolbox (SSDToolbox.exe) was programmed and compiled in Matlab 2018b (Mathworks 2018). To run the program you must first install the appropriate version of the Matlab Compiler Runtime (MCR). This can be downloaded free of charge from the Mathworks website:

http://www.mathworks.com/products/compiler/mcr

The correct version of the MCR for the SSD Toolbox is R2018b (9.5) for Windows 64 bit. NOTE: the MCR is very large (1.7 Gb) and can take some time to download, especially with a slow internet connection. The SSD toolbox executable may be placed in any directory on your computer.

Running the SSD Toolbox

Starting the program

The SSD Toolbox should launch by double-clicking "SSD Toolbox.exe". Note that it will fail to load if you have not previously installed the correct MCR (see Installation, above). It may also take some time to load, as the MCR must first be initialized.

Formatting and Importing Data

Data for analysis in the SSD Toolbox should be organized in columns and can be in Excel (*.xls, *.xlsx) or text (*.txt, *.csv) files. The first three columns of data should be Genus, Species, and Toxicity Value (*i.e.,* LD50, LC50, ECx, etc.). A fourth column may be used for birds when the you wish to standardize the toxicity values by body weight. In this case, the fourth column should be the mean weights of the tested animals. The Permethrin data file (PermethrinAcuteData.xlsx) is reproduced below in Table 2 (columns wrapped for economy of space).

Strictly speaking, there are no minimum sample sizes required by the SSD Toolbox, save for the requirement that sample size must be equal to or exceed the number of distribution parameters to be estimated (2 in most cases, except the $Burr_{III}$ distribution which has 3 estimated parameters). While it is understood that in most cases SSDs will be fit with small sample sizes, attempting to fit distributions to such limiting cases (sample size barely exceeding the number of estimated parameters) will almost certainly result in unreliable estimates of hazardous concentrations. Newman et al. (2002) recommended at least 40 data points (species), a criterion rarely met in SSD fitting and analysis.

Table 2. Permethrin aquatic toxicity data (from Fojut et al. 2012: Table 10)

Genus	species	LC50		Genus	species	LC50
Ceriodaphnia	*dubia*	0.25		*Danio*	*rerio*	2.5
Ceriodaphnia	*dubia*	0.652		*Daphnia*	*magna*	0.32
Ceriodaphnia	*dubia*	0.788		*Erimonax*	*monachus*	1.7
Ceriodaphnia	*dubia*	0.622		*Etheostoma*	*fonticola*	3.34
Ceriodaphnia	*dubia*	0.772		*Etheostoma*	*lepidum*	2.71
Ceriodaphnia	*dubia*	0.745		*Hyalella*	*azteca*	0.0211
Ceriodaphnia	*dubia*	0.858		*Ictalurus*	*punctatus*	5.4
Ceriodaphnia	*dubia*	0.571		*Notropis*	*mekistocholas*	4.16
Ceriodaphnia	*dubia*	0.58		*Oncorhynchus*	*apache*	1.71
Ceriodaphnia	*dubia*	0.609		*Oncorhynchus*	*clarki henshawi*	1.58
Ceriodaphnia	*dubia*	0.57		*Oncorhynchus*	*mykiss*	7
Ceriodaphnia	*dubia*	0.827		*Oreonectes*	*immunis*	0.21
Ceriodaphnia	*dubia*	0.585		*Pimephales*	*promelas*	9.38
Ceriodaphnia	*dubia*	0.849		*Procambarus*	*blandingi*	0.21
Ceriodaphnia	*dubia*	0.889		*Procloeon*	Sp.	0.0896
Ceriodaphnia	*dubia*	0.865		*Salmo*	*salar*	1.5
Chironomus	*dilutus*	0.189		*Xyrauchen*	*texanus*	5.95

Once data are appropriately formatted, they can be imported by choosing Import data from the File menu on the Graphical User Interface (GUI, Fig. 2).

Figure 2. The SSD Toolbox Graphical User Interface (GUI)

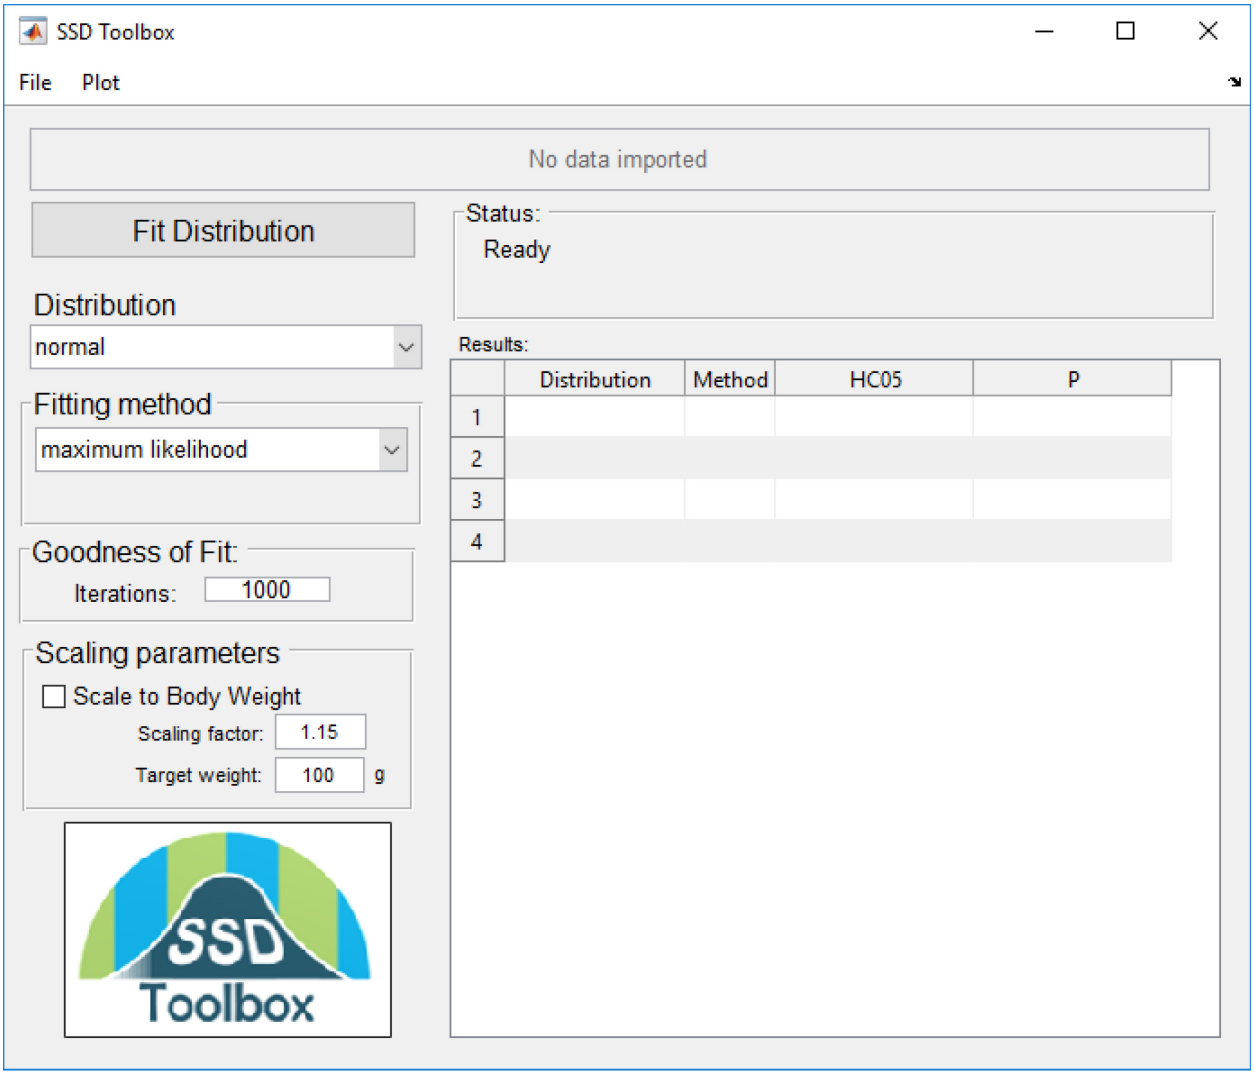

Whenever the Import Data option under the File menu is clicked, a warning will be displayed (Fig. 3) asking whether you wish to delete existing data and results. If you have fitted any distributions to an existing dataset, they will be deleted upon opening a new dataset. This warning is displayed regardless of whether you have previously fit any SSDs or whether you have previously imported data.

Figure 3. Warning

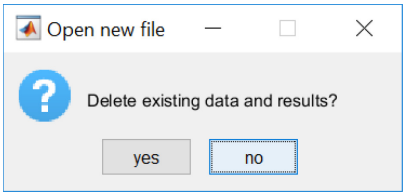

If you wish to perform analysis on bodyweight-standardized toxicity values, you must specify this prior to import using the checkbox on the main GUI (Fig. 2). In this case your data file must have bodyweights in column 4. On import, the bodyweights will be used to standardize the toxicity values using the scaling factor and target weight in the "Scaling parameters" frame. If you check the Scale to bodyweight option and your file does not contain numeric bodyweights in column 4 the program will generate an error warning (Fig. 4). If you choose bodyweight rescaling, all subsequent analysis and reporting will be done using the rescaled values.

Figure 4. Warning that bodyweight data were not properly formatted for scaling.

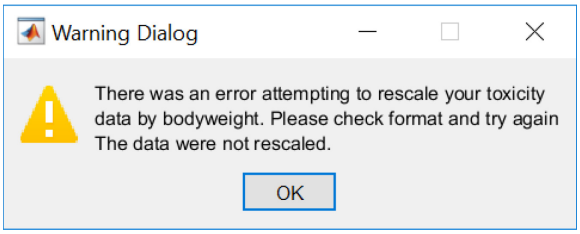

Once the file has been imported, the data will be displayed (Fig. 5). The number of test results is displayed above the table (34 in Fig. 5) and several choices allow you to generate histograms of the number of test results per species (Fig. 6), a toxicity histogram (Fig. 7), and the empirical cumulative distribution function (ECDF, Fig. 8). Values in the Toxicity value field are geometric means of all values for each species (prior to standardization if standardization is used). The "Replication" field gives the number of replicate toxicity values provided for each species. The ECDF field gives the value of the ECDF for each geometric mean toxicity value, sorted in descending order of empirical quantile. The Data table may be used to cross check the data against the original input file to make sure that all species are captured and that the number of endpoints per species is correct. It may also be used to decide on weight cutoffs for fitting graphical SSDs when indeterminate toxicity results are included (see Graphical Methods below for more detail).

When preparing data in Excel for import into the SSD Toolbox it is important to delete any trailing rows in the Excel file that contain formatting. This can happen inadvertently when editing a data file and deleting content using the "delete" key, which removes content, but not formatting in Excel. Matlab may interpret the residual formatting as data and import the empty cells. If this happens these empty cells will appear in the data table (Fig. 5) as the top row(s) and will have the toxicity value of NaN (not a number). To fix this problem, make sure that formatting is cleared from all cells immediately below the data you wish to import. Deleting entire rows is an easy way to ensure that these residual formats are removed.

Figure 5. Geometric means of the toxicity values by taxon in PermethrinAcuteData.xlsx (from Fojut et al. 2012: Table 10)

Data

Number of Test Results: 34

Plots

Sampling density | Tox Histogram | Toxicity ECDF

Data:

	Species	Toxicity Value	Replication	ECDF
1	Pimephales promelas	9.3800	1	0.9500
2	Oncorhynchus mykiss	7	1	0.9000
3	Xyrauchen texanus	5.9500	1	0.8500
4	Ictalurus punctatus	5.4000	1	0.8000
5	Notropis mekistocholas	4.1600	1	0.7500
6	Etheostoma fonticola	3.3400	1	0.7000
7	Etheostoma lepidum	2.7100	1	0.6500
8	Danio rerio	2.5000	1	0.6000
9	Oncorhynchus apache	1.7100	1	0.5500
10	Erimonax monachus	1.7000	1	0.5000
11	Oncorhynchus clarki henshawi	1.5800	1	0.4500
12	Salmo salar	1.5000	1	0.4000
13	Ceriodaphnia dubia	0.6641	16	0.3500
14	Daphnia magna	0.3200	1	0.3000
15	Procambarus blandingi	0.2100	1	0.2500
16	Oreonectes immunis	0.2100	1	0.2000
17	Chironomus dilutus	0.1890	1	0.1500
18	Procloeon Sp.	0.0896	1	0.1000
19	Hyalella azteca	0.0211	1	0.0500

Figure 6. Taxonomic sampling density for Permethrin data

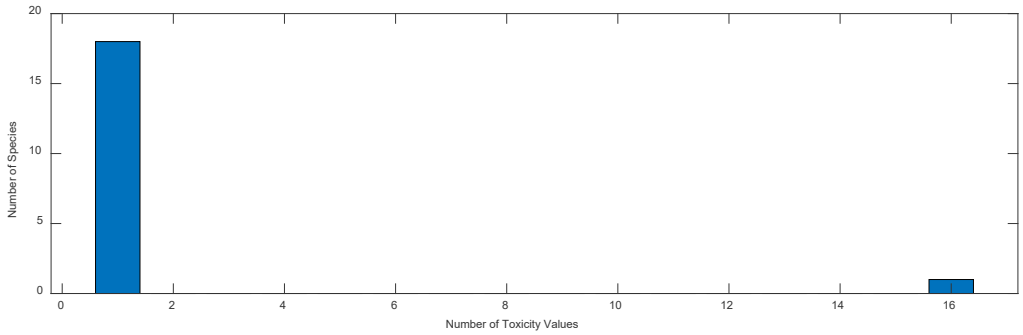

Figure 7. Log_{10} Toxicity Histogram for Permethrin data

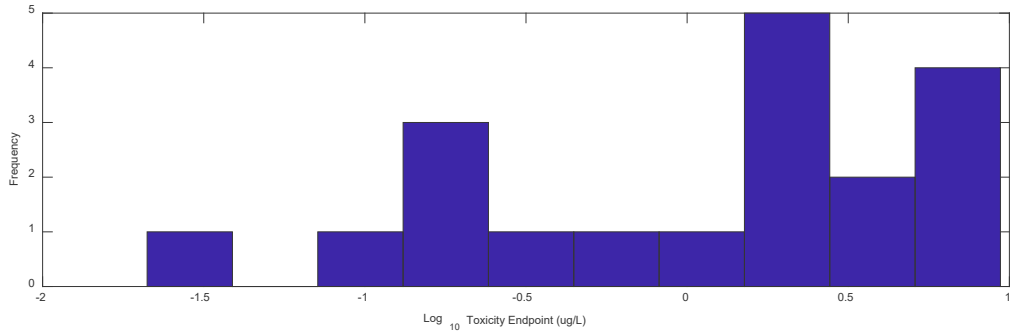

Figure 8. Empirical Cumulative Distribution Function (ECDF) for the Permethrin data

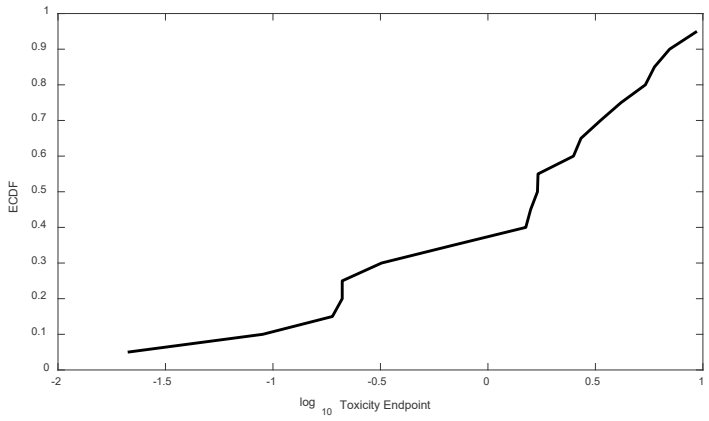

Fitting a distribution

Once data have been imported, fitting a distribution is done by simply clicking on the "Fit Distribution" button. You may choose which distribution to fit using the "Distribution" popup menu (Fig. 2). You may also choose among the four different fitting methods using the "Fitting Method" popup menu. The default fitting method is Maximum Likelihood. Much more information on choosing a fitting method is provided in the companion Technical Manual.

Below (Fig. 9), the SSD Toolbox is shown having fit all six distributions to the Permethrin data using maximum likelihood (ML). A limited amount of information is provided in the main table. Note that the P-value column (last column), corresponding to goodness-of-fit, is empty. To assess goodness-of-fit, and to obtain more information about the fitted distribution, you can use the context menu (see later section on the context menu).

Figure 9. Six distributions fit to the Permethrin data using maximum likelihood.

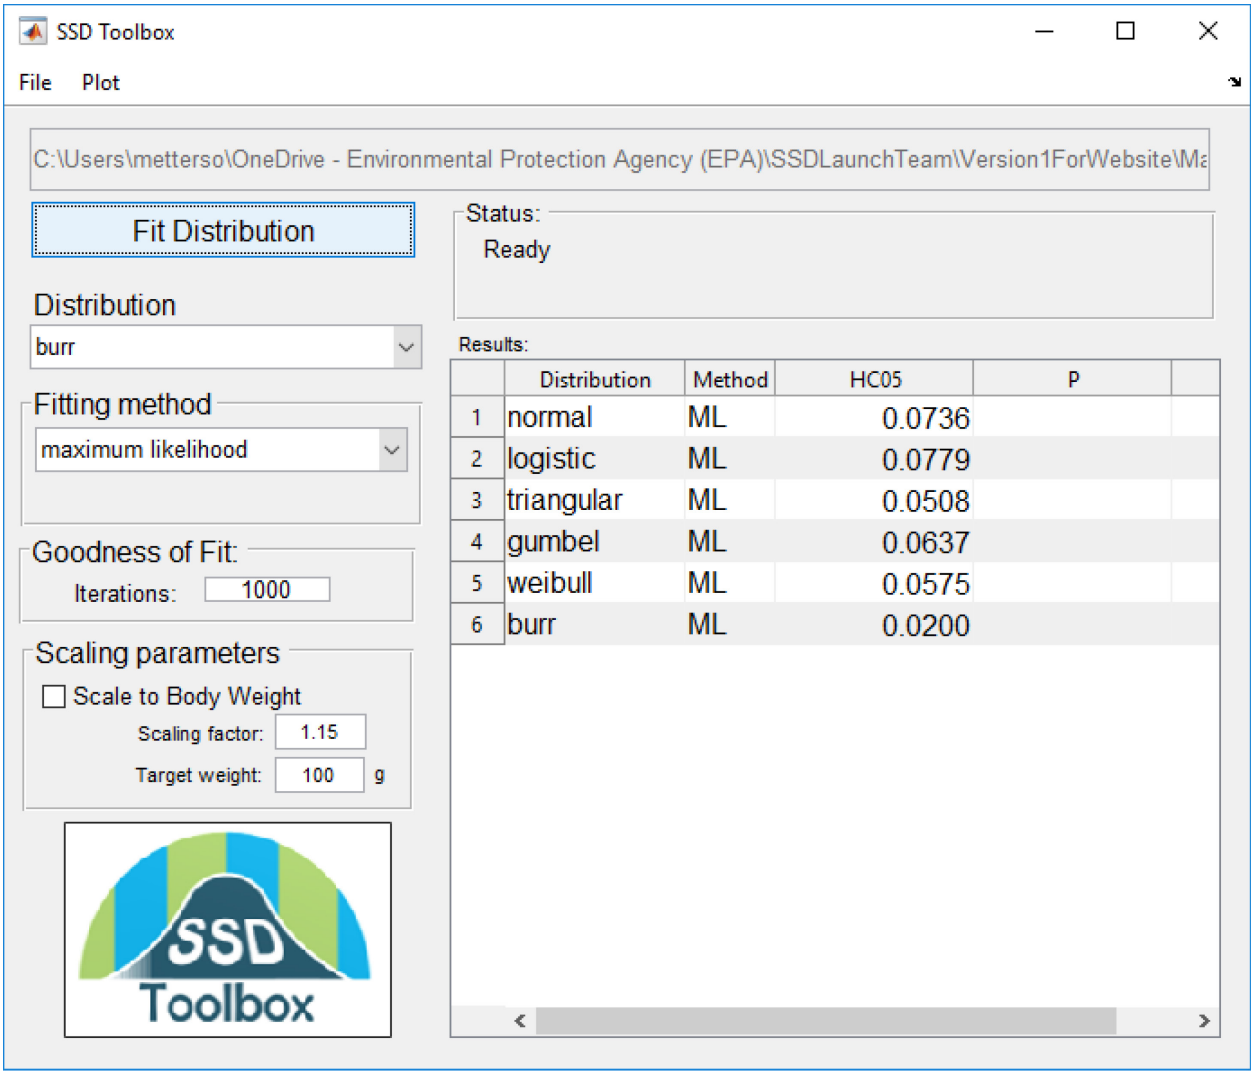

	Distribution	Method	HC05	P
1	normal	ML	0.0736	
2	logistic	ML	0.0779	
3	triangular	ML	0.0508	
4	gumbel	ML	0.0637	
5	weibull	ML	0.0575	
6	burr	ML	0.0200	

Weighted Fits and censored data

If you choose linearization to fit the SSD, you will have the option to weight the parameter estimates to the lower tail of the distribution. The weighted region is defined by the quantile cutoff box below the Fitting method popup (not visible unless using linearization). Choosing a value of 1 (default) gives equal weight to all data points. Choosing a value smaller than 1 will weight the distribution to those values that are below the quantile cutoff. For example, setting the quantile cutoff to 0.5 will use only the lower half of the toxicity values in the linearized regression (for details on how this process works see the Technical Manual). Figure 10 shows the effect of fitting a lognormal distribution to the Permethrin data with quantile cutoffs of 1, 0.5, and 0.25.

Figure 10. Lognormal SSDs fit to the Permethrin data using graphical methods with quantile cutoffs of 1, 0.5, and 0.25.

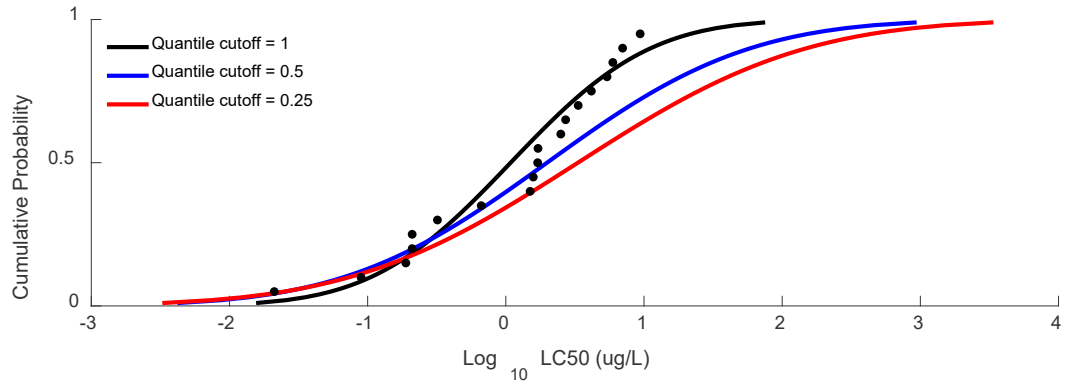

These weighted regressions can accommodate censored data by specifying the lower bound for the censored toxicity endpoints. The quantile cutoff must be lower than the lowest lower bound for censored data. Only linearization can accommodate censored data in this version of the SSD Toolbox.

Bodyweight scaling

For birds, when analyzing LD50 data, the SSD Toolbox allows you to rescale the toxicity values by body weight, using scaling parameters available for a subset of pesticides (Mineau *et al.* 1996). These scaling parameters are referred to herein as "Scaling factors," and the GUI is set to the default value (1.15). This rescaling must be done upon importing the data; if you skip this step on import you must go back and reimport the data, making sure that you have the correct target weight (default = 100 g) and scaling factor (default = 1.15) entered on the GUI. Rescaling is not typically done for fish, aquatic invertebrates, or plants.

Inferential Endpoints

It is assumed throughout that the desired inferential endpoint is an estimate of the HC05: the concentration expected to be hazardous to no more than 5% of tested species. This value is always reported on the main output window (Fig. 2). Regardless of distribution or of the method used to fit the distribution, the HC05 is estimated as the 5th percentile of the fitted distribution. This is given by the appropriate quantile function, which is unique to each distribution. The quantile functions for all distributions in the SSD Toolbox are provided in the Technical Manual. In cases where other quantiles are of interest (e.g., HC10, etc.), these can be obtained via the context menu described in detail below. In some cases, users may wish to use the lower confidence limits on hazardous concentrations as inferential endpoints, and these can also be obtained using the context menu, as described below.

An important consideration in making inference about hazardous concentrations is the distribution of available data across species. Strictly speaking, an assumption of all the methods considered below is that the data (test results) pertain to a random sample of species from the set of species for which the analysis is intended to apply. This assumption is always

violated; a relatively limited subset of taxa makes up the greater part of all toxicity tests. Therefore, in using SSDs to derive protection goals, one should consider the potential biases in the data set relative to the set of species for which the protection goal is intended to apply.

The File menu

The File menu has five options (Fig. 11), explained in more detail below.

Import data imports data for analysis. Data files can be in MS Excel® (default), comma-delimited file format (*.csv), or text format (*.txt). This menu was explained in greater detail above under "Formatting and Importing Data."

Load previous analysis loads a file generated in a previous SSD fitting session. As with the "Import data" option, you will receive a warning that all previous data and analyses will be deleted (Fig. 3).

Save results saves results to a file that can be subsequently reimported for plot generation and further posterior analysis.

Clear table deletes existing results, but the imported data are still available for analysis.

Exit closes the SSD Toolbox.

Figure 11. The File menu.

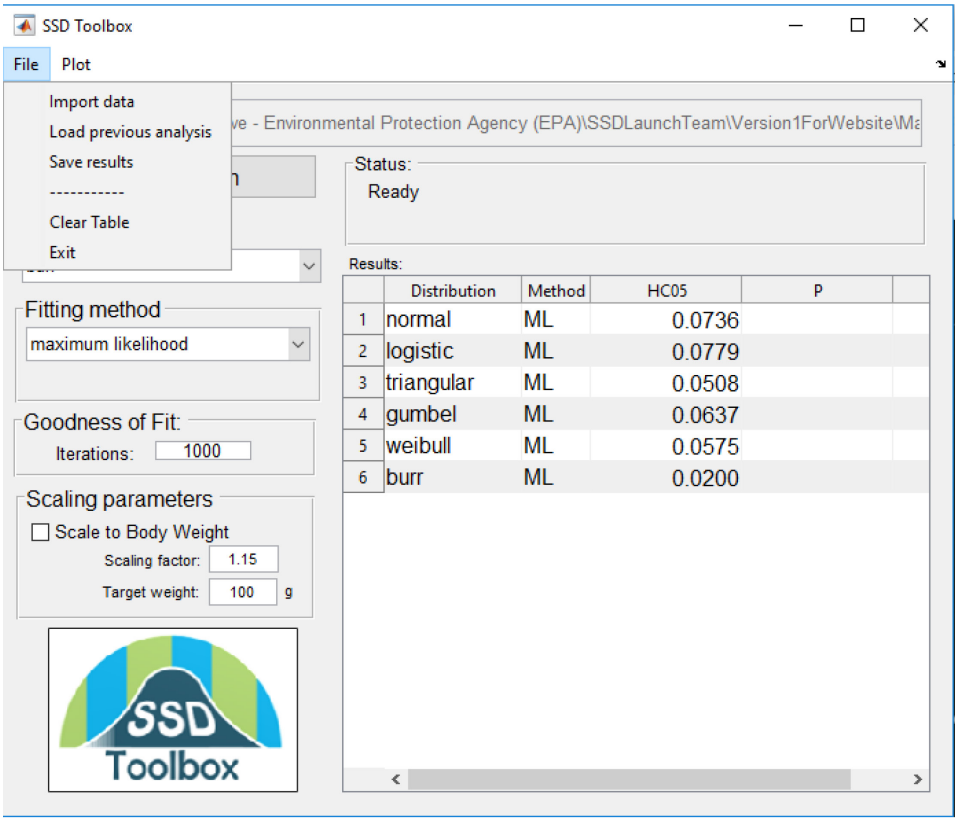

The Plot menu

Figure 12 shows the plot menu, which has two options.

Show data generates the data table (Figure 5, above)

Interactive Plotting opens the "Interactive Plotting" GUI, which allows you to choose elements of the SSD for plotting (Fig. 13).

Figure 12. The plot menu

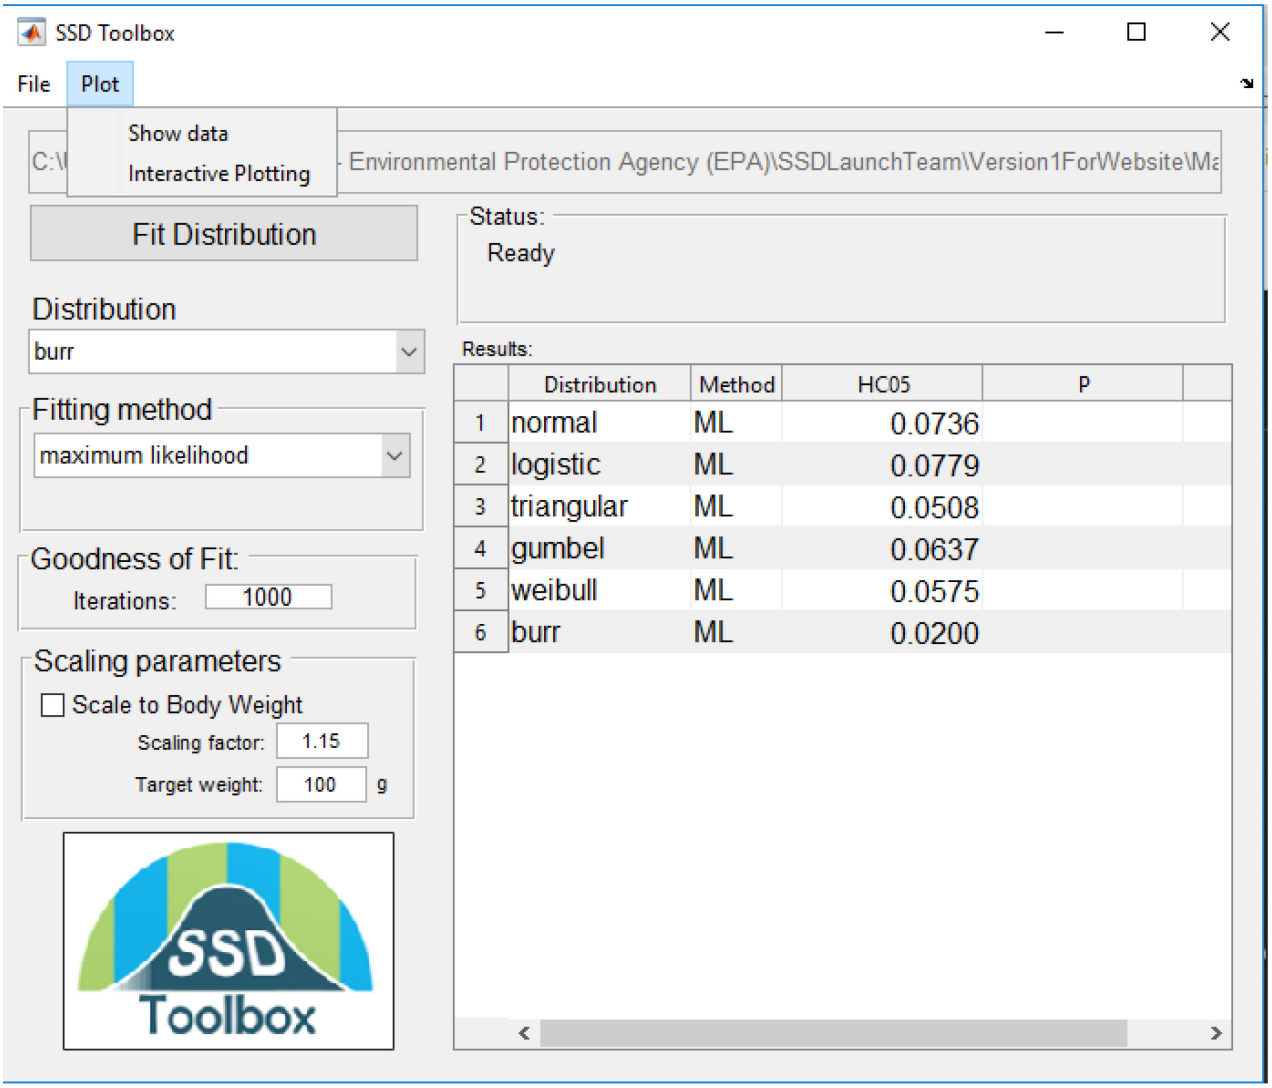

Figure 13. The Interactive Plotting GUI

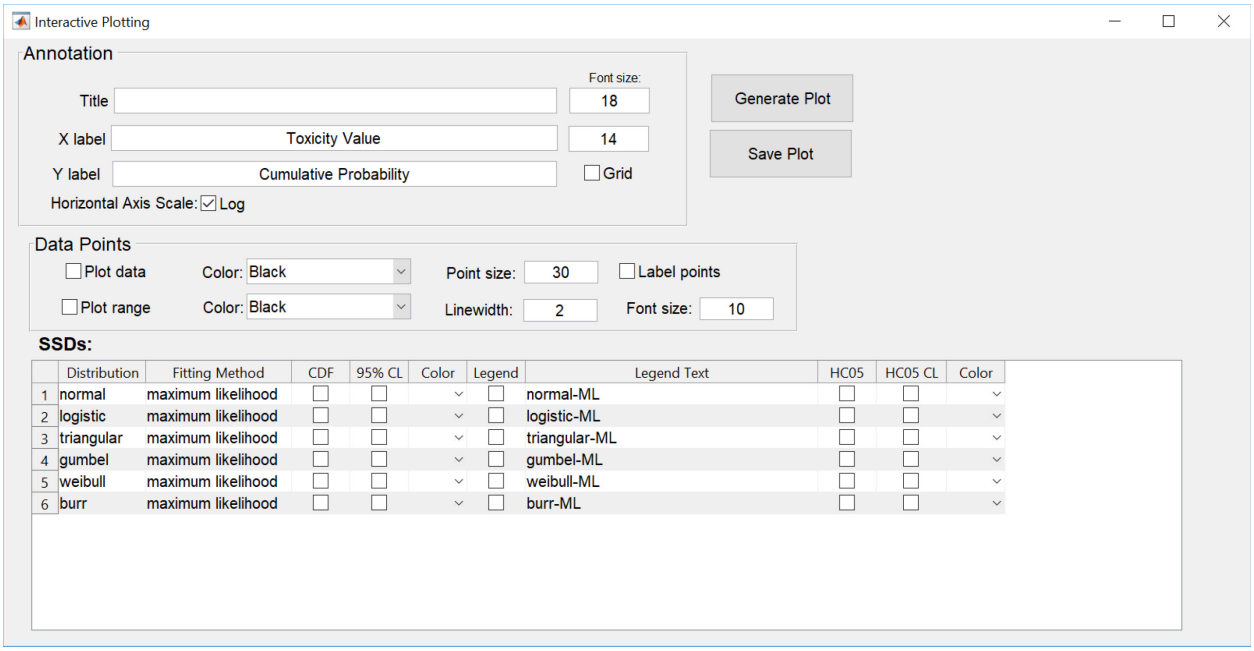

The “Interactive Plotting” GUI allows you to generate figures with specific desired elements. Checkboxes determine whether an element will be included in a plot. For most elements, the color can be chosen from a popup menu, and the default color is black. Point size and linewidth can also be specified for data ranges. Default axis titles are provided for the X-axis (“Toxicity Value”) and Y-axis (“Cumulative Probability”), but these can be edited. Font size can be changed for data point labels and for axis titles. There are too many permutations to include all possible plots in this guide, but a few minutes of experimentation will allow you to get a good feel for the possibilities. Two examples are provided below using the standardized sample data. The first example (Fig. 14 a, b) shows only the data points plotted with horizontal bars for the maximum and minimum values for each taxon. In this example the labeling for the toxicity scale was presented in natural toxicity units (ug/L) by unchecking the “Toxicity scale Log” option. The second example (Figure 15 a, b) shows a fitted log-normal SSD plotted with data points and 95% confidence limits.

To generate a plot using the Interactive Plotting GUI, set the elements to their desired values and click the “Generate Plot” button. Plots thus generated can be saved by choosing “Save Plot”, which will save the figure to any of several common graphic file formats.

Figure 14a. Interactive Plotting GUI setup for generating a plot of the Permethrin data with taxon-specific ranges for toxicity values

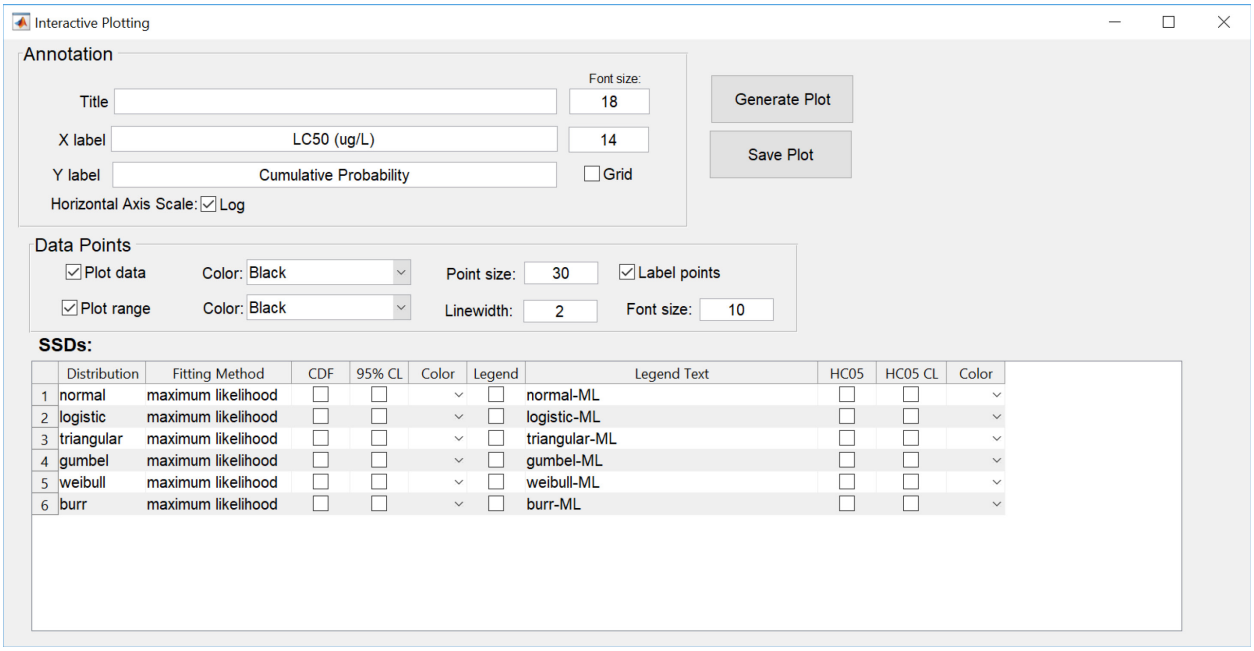

Figure 14b. Scatter plot of Permethrin aquatic toxicity data

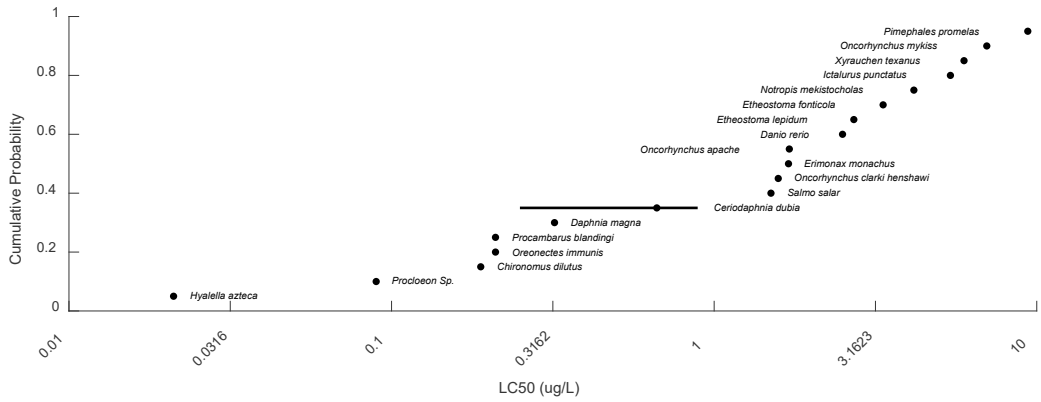

Figure 15a. Interactive Plotting GUI setup for generating a Weibull plot of Permethrin aquatic toxicity data

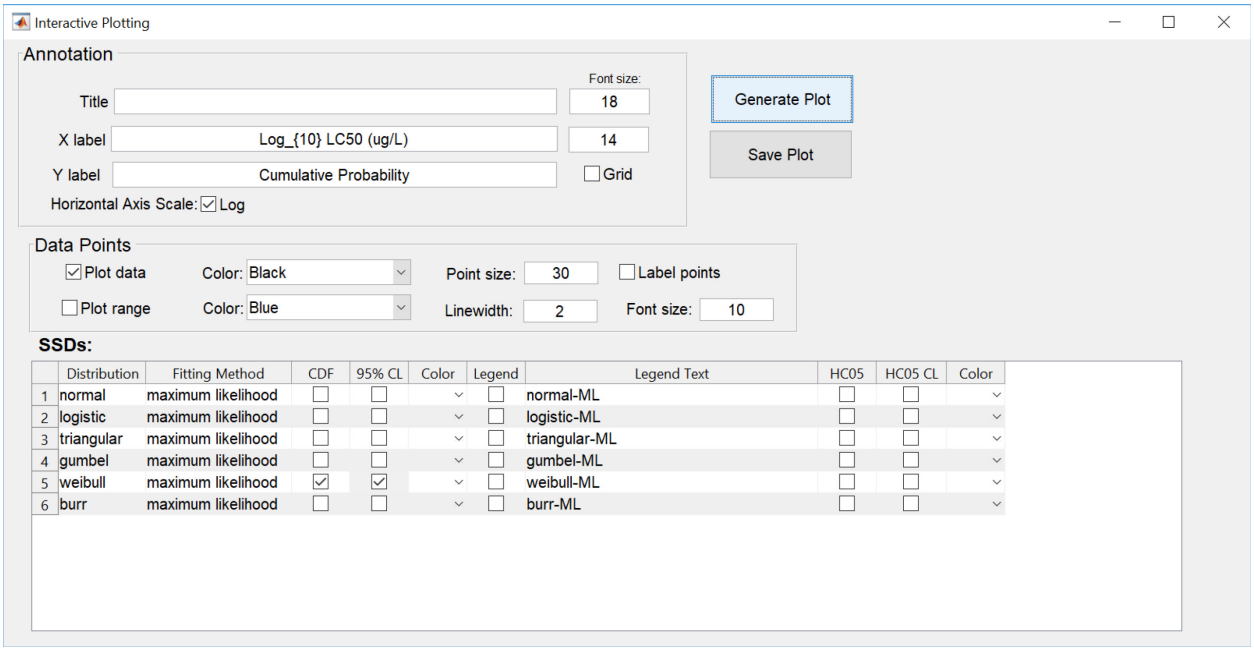

Figure 15b. Weibull SSD fit to Permethrin data

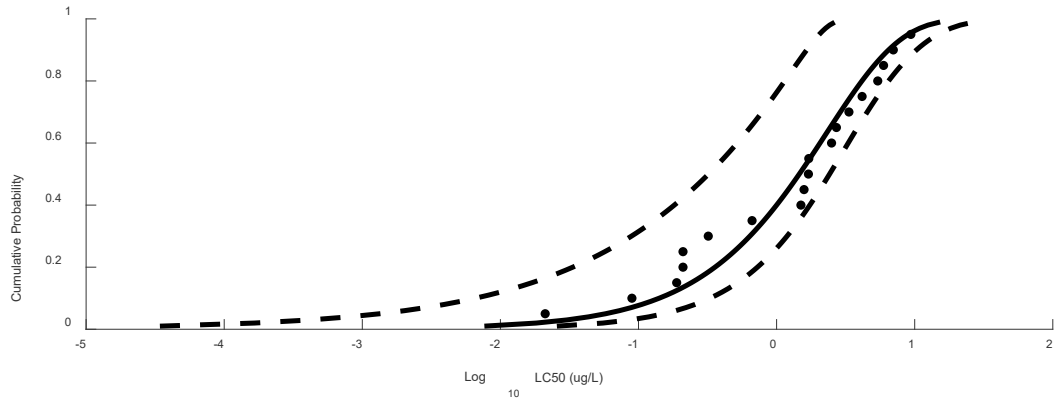

The context menu

The context menu is available by selecting a row in the main results table, and right-clicking (Fig. 16). The context menu has 6 choices (7 if Metropolis-Hastings was chosen as a fitting method), two of which (Goodness of fit and Posterior Statistics), themselves have submenus.

Figure 16. The context menu

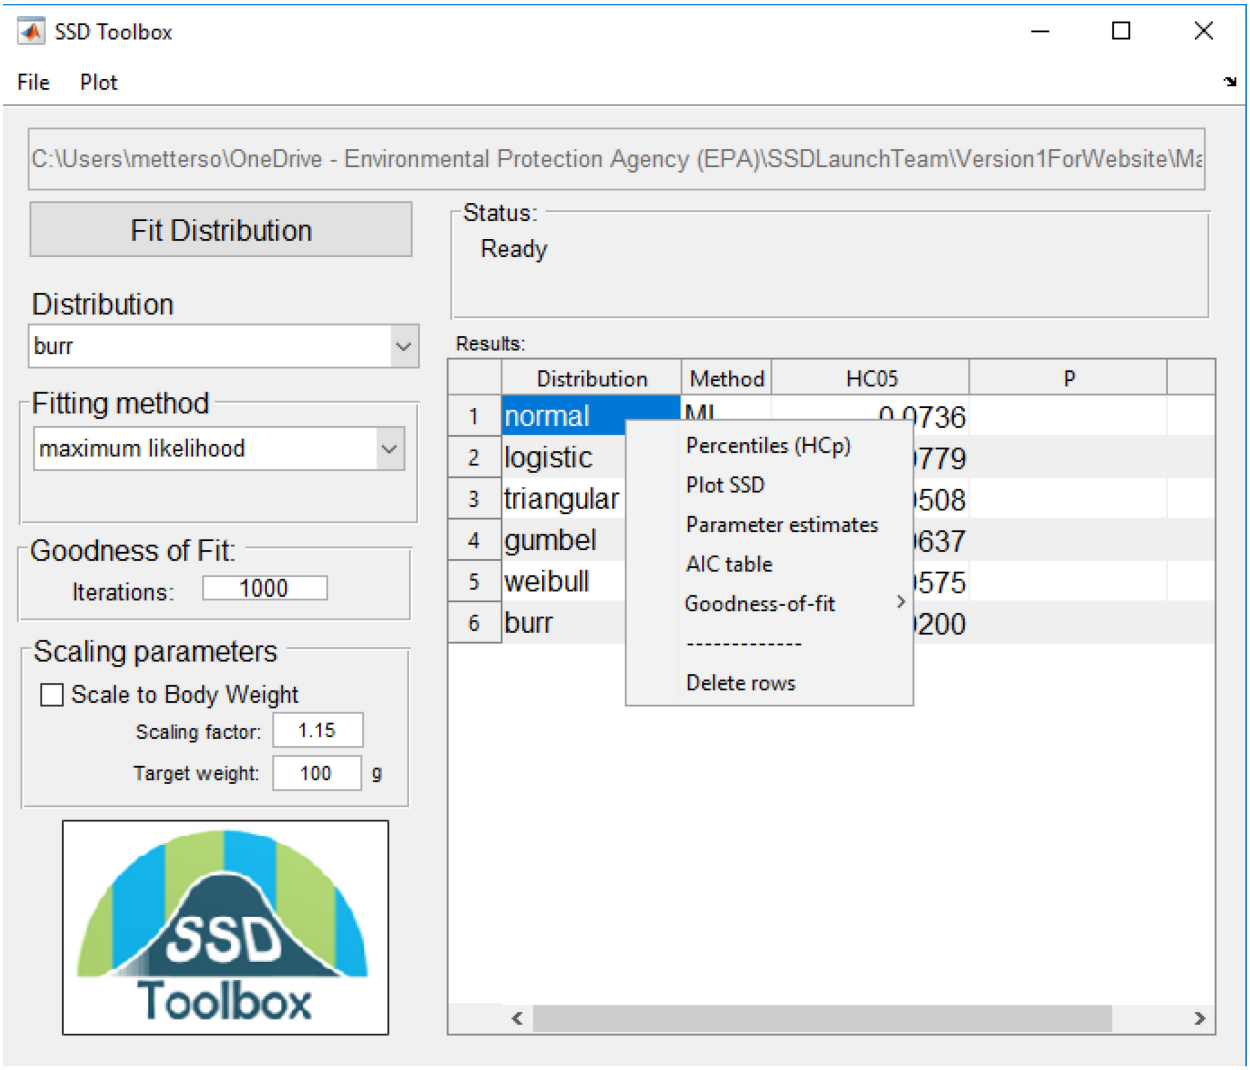

Percentiles (HCp) opens a window giving the percentiles (1 – 99) of the fitted distribution (Fig. 17). The "Variance method" button group allows specification of different variance methods. If goodness-of-fit has been run, bootstrap estimates of standard error (SE), coefficient of variation (CV) and lower and upper confidence limits around the estimated percentiles (Lower CL, Upper CL) will be displayed. If the distribution has been fit using maximum likelihood, then the same statistics will be available but estimated using the Hessian matrix. If the distribution has been fit using Metropolis Hastings then the bootstrap choice will display standard errors and credible intervals for the quantiles from the posterior distribution of parameter values. In all cases, these values can be copied to the clipboard and pasted into a spreadsheet application if desired.

Figure 17. The HCpTable for a Weibull distribution fit to the Permethrin data using maximum likelihood, displaying variance estimate from the Hessian matrix. Row-labels are values representing “p.”

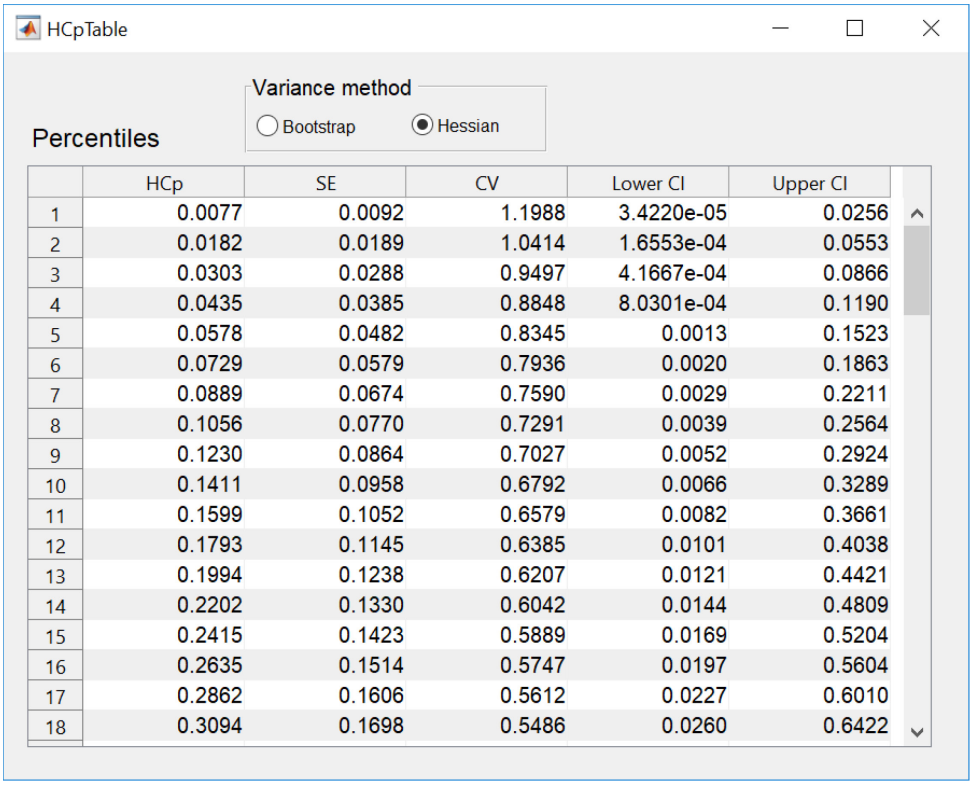

	HCp	SE	CV	Lower CI	Upper CI
1	0.0077	0.0092	1.1988	3.4220e-05	0.0256
2	0.0182	0.0189	1.0414	1.6553e-04	0.0553
3	0.0303	0.0288	0.9497	4.1667e-04	0.0866
4	0.0435	0.0385	0.8848	8.0301e-04	0.1190
5	0.0578	0.0482	0.8345	0.0013	0.1523
6	0.0729	0.0579	0.7936	0.0020	0.1863
7	0.0889	0.0674	0.7590	0.0029	0.2211
8	0.1056	0.0770	0.7291	0.0039	0.2564
9	0.1230	0.0864	0.7027	0.0052	0.2924
10	0.1411	0.0958	0.6792	0.0066	0.3289
11	0.1599	0.1052	0.6579	0.0082	0.3661
12	0.1793	0.1145	0.6385	0.0101	0.4038
13	0.1994	0.1238	0.6207	0.0121	0.4421
14	0.2202	0.1330	0.6042	0.0144	0.4809
15	0.2415	0.1423	0.5889	0.0169	0.5204
16	0.2635	0.1514	0.5747	0.0197	0.5604
17	0.2862	0.1606	0.5612	0.0227	0.6010
18	0.3094	0.1698	0.5486	0.0260	0.6422

The default behavior of the HCp Table is to show bootstrap estimates of the variances around the percentiles of the distribution. Bootstrap estimates are not available until the Goodness-of-fit routine is run (see below). If the distribution was fit using maximum likelihood, then the Hessian statistics are also available by clicking the “Hessian” radio button.

Plot SSD generates a quick plot of the SSD with data points and data ranges (Fig. 18).

Figure 18. Plot of Weibull SSD fit to Permethrin data using maximum likelihood. Horizontal blue lines indicate the range of toxicity values. Red points are geometric means for taxa with multiple estimates. Black points are single estimates.

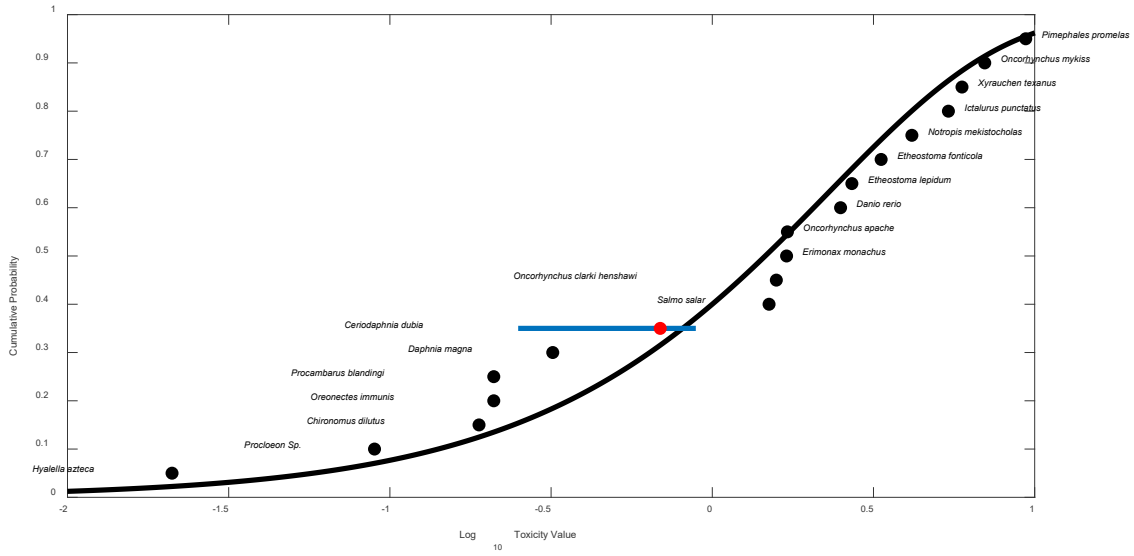

Parameter estimates opens a window giving the parameters of the fitted distribution (Fig. 19). For normal, logistic, triangular, and Gumbel distributions, the parameters will be for log-transformed data. If goodness-of-fit has been run, then the bootstrap standard errors will also be displayed. If the distribution was fit using maximum likelihood the standard error estimated using the Hessian matrix will be displayed. If the distribution was fit using Metropolis Hastings then the standard error of the estimate and credible interval are calculated from the posterior distribution of parameter values.

Figure 19. Parameter estimates for a Weibull distribution fit to the Permethrin aquatic toxicity data using maximum likelihood and displaying the standard error and confidence limits generated from the Hessian matrix.

Output

Parameter Estimates:

	Estimate	SE Hessian	LCL Hessian	UCL Hessian
lambda	2.2994	0.6876	0.9517	3.6472
k	0.8063	0.1516	0.5092	1.1034

AIC table opens a window with an AIC_c table created using all distributions fit using maximum likelihood (Fig. 20). The AIC table shows standard AIC statistics (corrected for sample size, AIC_c), as well as the HCp and standard error (from the Hessian matrix) of the HCp. Above the table,

several edit boxes provide the model-averaged HCp, the standard error and coefficient of variation of the model-averaged HCp. The model-averaged statistics are generated using the AIC_c weights in the Weight column. To change the percentile of interest, you may type a different percentile in the top edit box. A detailed mathematical description of how these model-averaged values are calculated is given in the Technical Manual.

Figure 20. AIC_c table for six distributions fit to the Permethrin data

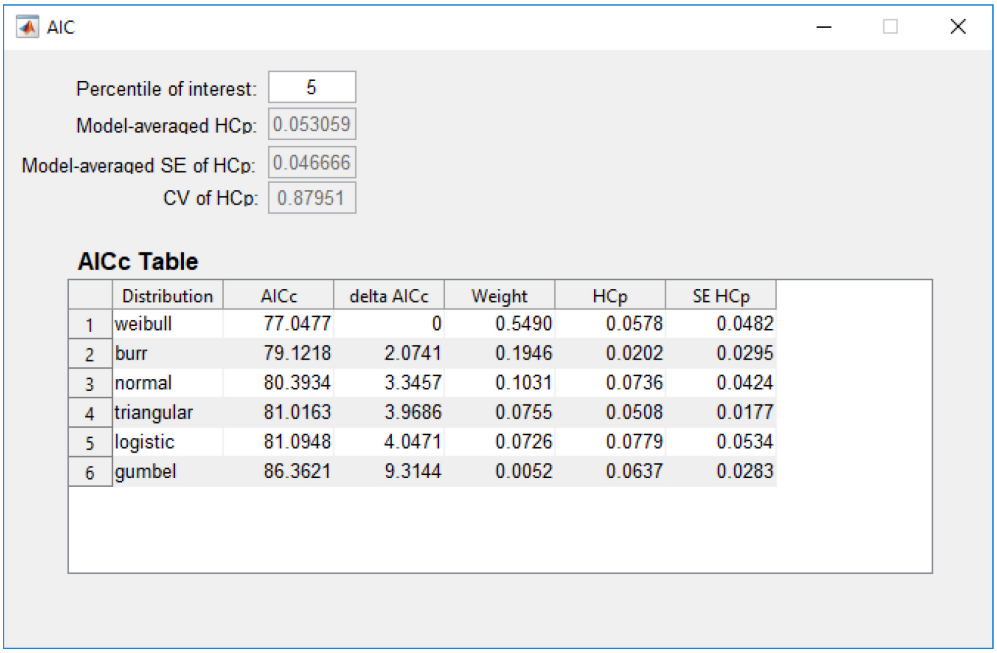

	Distribution	AICc	delta AICc	Weight	HCp	SE HCp
1	weibull	77.0477	0	0.5490	0.0578	0.0482
2	burr	79.1218	2.0741	0.1946	0.0202	0.0295
3	normal	80.3934	3.3457	0.1031	0.0736	0.0424
4	triangular	81.0163	3.9686	0.0755	0.0508	0.0177
5	logistic	81.0948	4.0471	0.0726	0.0779	0.0534
6	gumbel	86.3621	9.3144	0.0052	0.0637	0.0283

The AIC table menu option will generate a warning if you have not fit at least two distributions using maximum likelihood (Fig. 21)

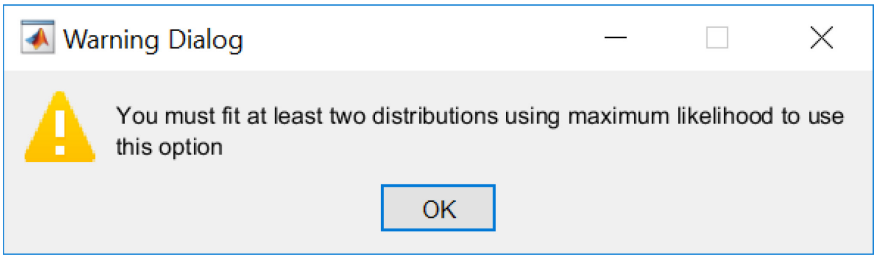

Goodness-of-fit allows you to run a parametric bootstrap goodness-of-fit test. This option has sub-menus allowing you to run the test only on desired distributions (selected rows in the table) or on all distributions. Once goodness-of-fit has been run, a P-value will be displayed in the right-most column. A low P-value indicates lack-of-fit. A third sub-menu generates Quantile (Q-Q) plots for visual inspection of fit.

Statistical estimation of lack of fit

To increase speed for the goodness-of-fit test, the SSD Toolbox takes advantage of parallel processing capabilities in computers with multi-core processors. This requires Matlab to load the parallel processing engine the first time such a test is run. Thus, there may be a delay of several seconds or more the first time you run a goodness-of-fit test in an SSD Toolbox session. On computers without multicore processors bootstrap goodness-of-fit estimation may be quite slow, taking 30 minutes or more, depending on the number of iterations specified. This is especially true for distributions fit using maximum likelihood. Note also, that you may get a warning about the windows firewall blocking some features of "SSDToolbox.exe" or "Ctfxlauncher.exe" (Fig. 22). If so, click "cancel" or "run anyway" and the program should run fine. If operating on a personal computer, you may see an alternative message that either of the two programs are requesting permission to circumvent the windows firewall. It is fine to deny permission. The program should still run fine.

Figure 22. Windows firewall warning on use of the goodness-of-fit algorithm

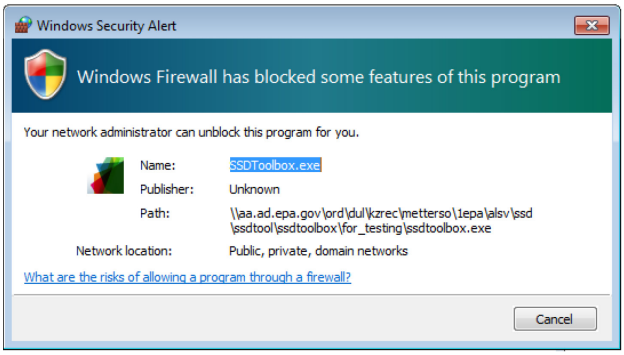

As noted above, when a distribution is first fit, the column corresponding to the goodness-of-fit P-value is empty. This is because the parametric bootstrap goodness-of-fit algorithm implemented in the SSD Toolbox can take a few minutes to run, especially for distributions fit using maximum likelihood. To assess goodness-of-fit, choose the number of bootstrap samples desired (1,000 is the default) and then choose "Goodness of fit" from the context menu and select the desired choice for assessing fit of selected distributions or all distributions. This starts the test running.

While the goodness-of-fit test is running, the SSD Toolbox is also using the bootstrap iterations to assess sampling variance around the parameters of the fitted distribution and around the quantiles (the HCp). This makes bootstrap estimates of the standard errors, coefficients of variation, and confidence limits for the percentiles available for the assessed distribution. These can be viewed in the context menu under "Percentiles (HCp)". If bootstrap estimates of

goodness-of-fit have not yet been run, then these fields will be blank when the Percentiles window is opened.

Q-Q plots

Quantile, or Q-Q, plots show the predicted quantiles from a fitted distribution against the empirical quantiles from the empirical cumulative distribution function. The closer these are to a straight line, the better the fit of the distribution to the data. Q-Q plots are useful for diagnosing deviations from fit in specific areas of the distribution (e.g., in the region of the HC05) and for finding outliers. Figure 23 gives a sample Q-Q plot for the Permethrin aquatic toxicity data.

Figure 23. Quantile plots for a. Weibull (best distribution by AIC), and b. Gumbel (worst distribution by AIC) for Permethrin aquatic toxicity data (see also Fig. 20 above)

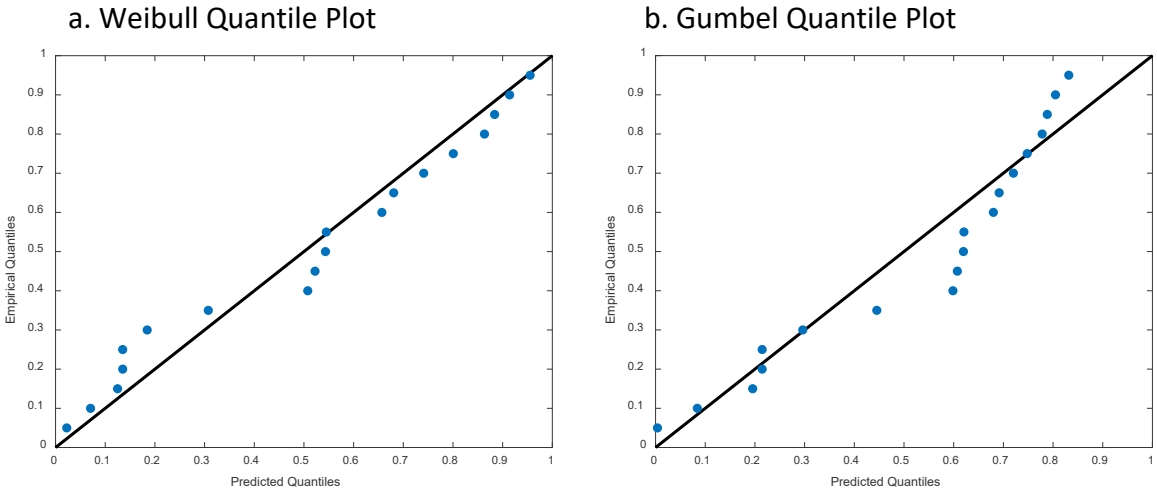

Posterior Diagnostics will only be available as a context menu choice when an SSD has been fit using the Metropolis-Hastings algorithm. This menu choice has four sub-menus. Additional guidance on interpreting these statistics is provided in the Technical Manual.

BIC Table. This option, like the AIC Table option for SSDs fit using maximum likelihood will generate a model selection table using Bayesian Information Criterion instead of Akaike's Information Criterion.

Figure 24. Distributions fit to Permethrin data with Bayesian *p*-values.

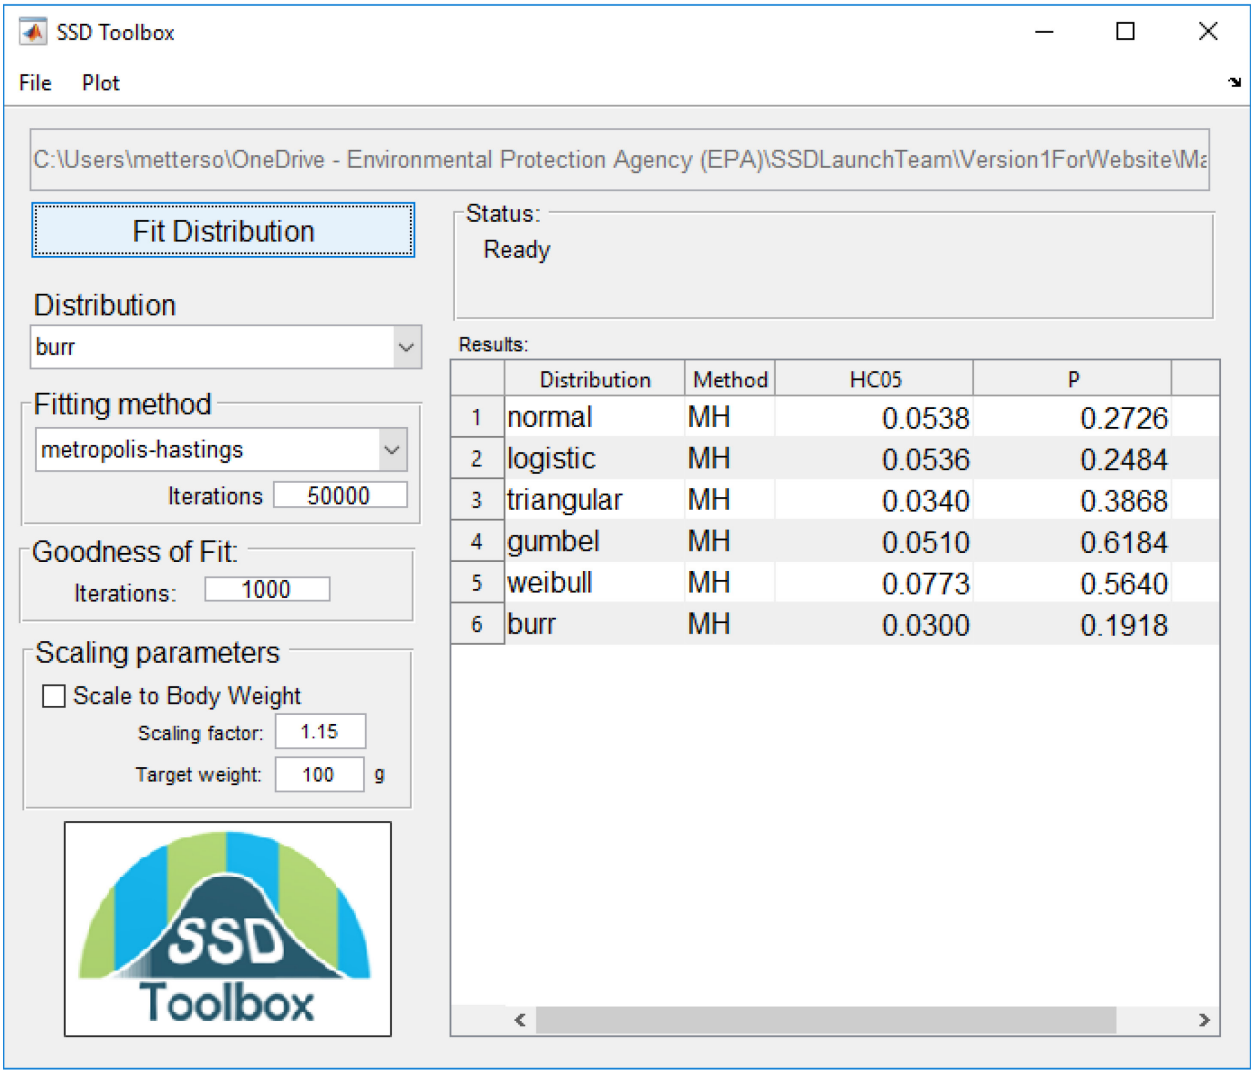

Credible Intervals After running the Metropolis-Hastings sampling routine, the program summarizes the simulated posterior distribution. During this time, it also calculates the Bayesian credible interval around parameter estimates and around the quantiles of the distribution. These will be subsequently displayed when "parameter Estimates" or Percentiles options, respectively, are chosen in the context menu.

Three additional graphical options are available to help diagnose the fit of Bayesian SSDs fit using the SSD toolbox. These are autocorrelation plots, trace plots, and posterior density plots, each of which may be helpful in understanding the quality of the fitted SSD. These plots are described briefly below. More information is provided in the Technical Manual.

Autocorrelation plots show the serial autocorrelation between sequential parameter values in the Markov chain. Typically, the autocorrelation should decline to zero quickly with increasing

lag-length. If not, this is an indicator of problems in the MCMC run. Note, this algorithm can take a minute or two to run. Figure 25 shows autocorrelation plots for normal distribution parameters.

Figure 25. Parameter autocorrelation plots for log-normal distribution fit to the Permethrin data

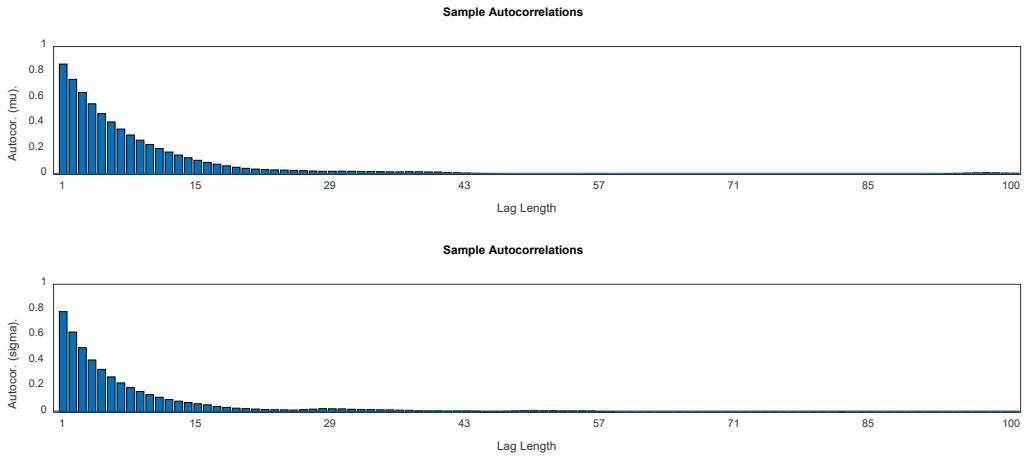

Trace plots show the sequential values sampled by the MCMC algorithm (Fig. 26). These are useful for judging the adequacy of sampling by the MCMC algorithm. Ideally these should show relatively little patterning other than frequent jumps about a central tendency.

Figure 26. Trace plots for log-normal distribution parameters for the Permethrin data.

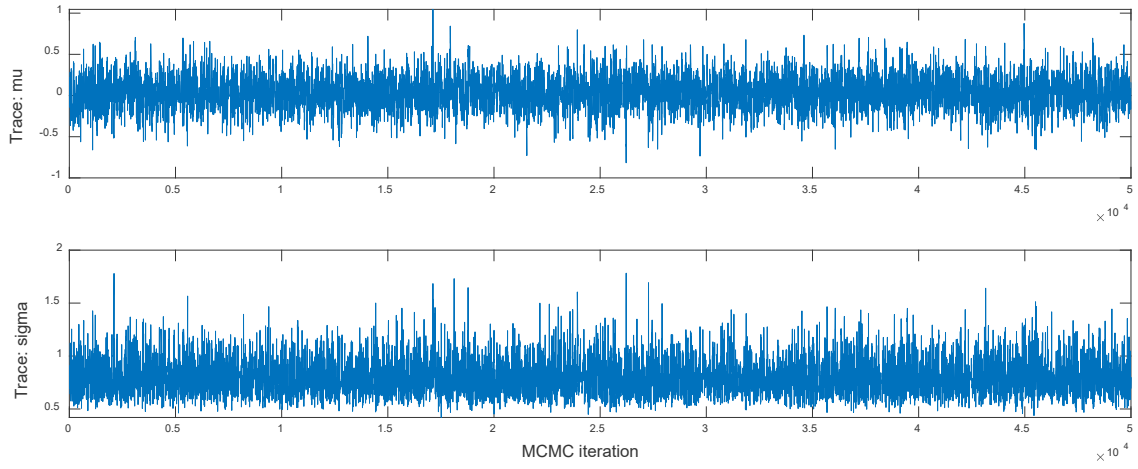

Posterior distributions plots the posterior distributions of the parameters of the fitted distribution. The value of statistics generated from Bayesian fits are the averages of these posterior distributions (the marginal distributions). Figure 27 shows the posterior marginal distributions (on the diagonal) and joint distributions (off-diagonal) of parameter estimates. Ideally the marginal distributions should be

unimodal and relatively smooth-edged. The joint distributions should be round or blocky, ideally not revealing covariation among sampled values from the posterior.

Figure 27. Posterior parameter distributions for log-normal distribution fit to the Permethrin data.

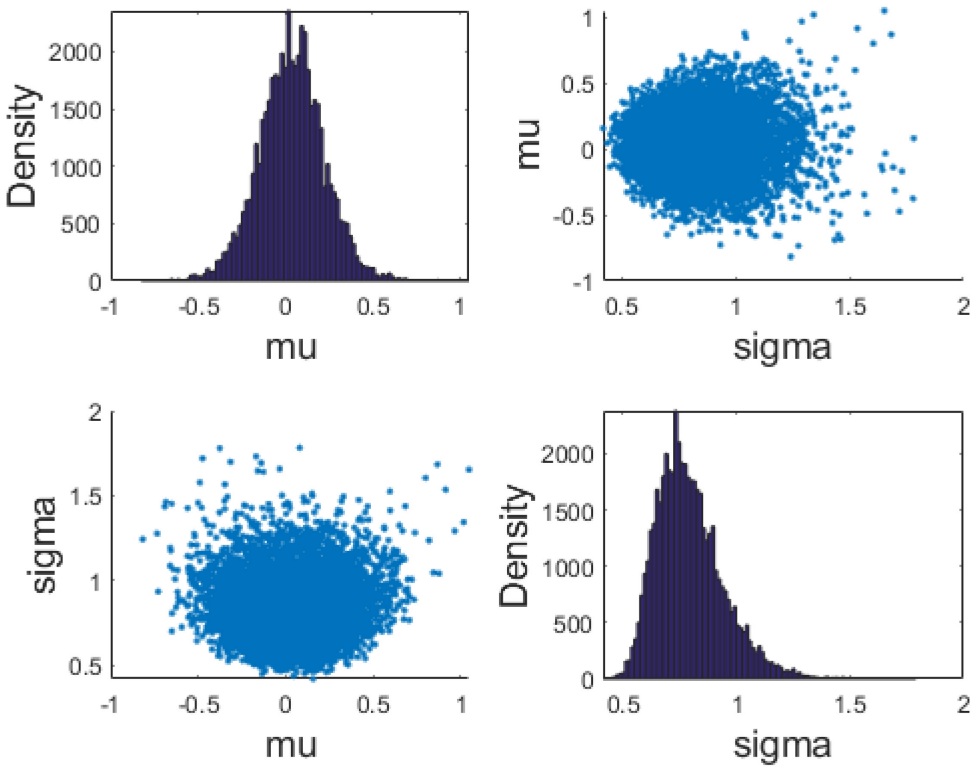

Delete rows. This menu option deletes all selected rows of the table. They cannot be recovered.

Literature Cited

Fojut, T.L, A.J. Palumbo, and R.S. Tjeerdema. 2012. Aquatic life water quality criteria derived via the UC Davis method: II. Pyrethroid Insecticides. Reviews of Environmental Contamination and Toxicology 2016:51-104.

Link, W.A. and R.J. Barker. 2010. Bayesian Inference With Ecological Applications. Elsevier, Academic Press, London, UK.

Mathworks 2017. Matlab. Mathworks. Natick, MA.

Mineau, P., Collins, B.T., and A. Baril. 1996. On the use of scaling factors to improve interspecies extrapolation of acute toxicity in birds. Regulatory Toxicology and Pharmacology, 24: 24-29.

Newman, M.C., D.R. Ownby, L.C.A. Mézin, D.C. Powell, T.R.L. Christensen, S.B. Lerberg, B.-A. Anderson and T.V. Padma. 2002. Species sensitivity distributions in ecological risk assessment: Analysis of distributional assumptions, alternate bootstrap techniques, and

estimation of adequate number of species. Pp. 119-132 in L. Posthuma, G.W. Suter, and T.P. Traas (Eds.) Species Sensitivity Distributions in Ecotoxicology. Lewis. Boca Raton, FL. USA.

[USEPA] US Environmental Protection Agency. 2016. Biological evaluation chapters for Chlorpyrifos ESA assessment. [cited 2019 July 3]. https://www.epa.gov/endangered-species/biological-evaluation-chapters-Chlorpyrifos-esa-assessment