ToxFinder – Defining and Running Simulations

Tessella Project Number 3760

System User Guide

Tessella Support Services plc

e-mail: post@tessella.com **www.tessella.com**

Head Office: Abingdon

3 Vineyard Chambers, Abingdon, Oxfordshire OX14 3PX, England
Tel: +44 (0)1235 555511
Fax: +44 (0)1235 553301

Document Control

This document is under document control. All members of the distribution list will receive updated copies whenever alterations are made.

Contributors

Person	Role	Company	Contribution
Tom Parke	Author	Tessella	
Nick Battam	Reviewer	Tessella	

Document Approval

Person	Role	Company

Distribution

Person	Role	Company
Peter Thall, Peter Mueller		MD Anderson
Notes Projects Database 3760		Tessella

Revision History

Issue	Author	Date	Description
V1.R1.M0	Tom Parke	14-Oct-2004	1st version written for Alpha release v0.2.6
V1.R2.M0	Tom Parke	02-Nov-2004	Updated to use revised terminology.
V1.R3.M0	Tom Parke	19-Nov-04	Provisional release, updated following review by Peter Thall
V1.R4.M0	Nick Battam	03-Dec-04	Updated for release v1.0.0
V1.R4.M1	Tom Parke	07-June-05	Minor updates to 1.4 from Peter Thall

References

Ref.	Document	Date	Details and Version
[TMML]	"Dose Finding with Two Agents in Phase I Oncology Trials" Peter F. Thall, Randall E. Milikan, |Peter Mueller, Sang-Joon Lee	Biometrics 59:487-496, 2003	Original paper describing the dose-finding algorithm.

Table of Contents

1 Introduction

1.1 Purpose of this document

This document describes how to use the ToxFinder graphical user interface. It is intended for all end users of the system. Users are strongly encouraged to read the 2003 Biometrics paper [TMML] before using ToxFinder.

1.2 Scope of this document

This guide only describes using the ToxFinder graphical interface and the ToxFinder facilities accessible via that interface for running a trial; it does not cover all the statistical details of the model and method in TMML.

1.3 Context of this Issue

This issue documents version 1.0.0 of the system released to MD Anderson for review.

Note:

- **This release has not been approved by the author's of the [TMML] algorithm as a correct implementation of that algorithm.**

- **This release has been specifically prepared to allow the author's of the [TMML] algorithm to review the current state of the implementation.**

1.4 Definition of Terms

ToxFinder is a pair of computer programs that implement the TMML dose-finding algorithm, simulate and run clinical trials using the TMML design and provide charts and graphs of the results.

A Simulation Series is comprised of a 'Design' for a trial using TMML, a set of one or more 'Scenarios' – possible probabilities of toxicities – and one or more 'Variants' where some of the Design parameters can be varied to explore the effect of the variation on the operating characteristics of a Trial. A Simulation Series creates a number of Runs = (number of Scenarios) * (number of Variants).

A Design is set of parameters that define a trial design to be carried out using the TMML method. This consists of: the two dose ranges to be used, lambda - the relative cancer killing potential of the two drugs, the prior distribution of the parameters that characterise the probability of toxicity as a function of the dose pair of the two agents, the cohort size, the sample sizes of the two stages of the trial, the fixed line L1 used in the first stage of the trial and the first set of doses to use on L1.

A Scenario is a set of fixed values of the probability of toxicity as a function of the two agents' doses, $\mathbf{x} = (x_1, x_2)$. We denote this probability by $\pi(\underline{\mathbf{x}})$. As $\mathbf{x}$ varies over the two-dimensional domain of the dose pairs $\pi(\mathbf{x})$ forms a surface that may be illustrated graphically. ToxFinder does this, providing a graphical representation of each scenario specified by the user.

A Variant is a small modification to a Design – e.g. different prior, cohort size or study sample size.

A Run A run is a set of all the parameters required to simulate a trial; it is the combination of a Design and a Scenario. A Simulation Series consists of one or more runs. Each run is organised as a separate folder, all the parameters for a run are held in a single file called 'init.tmml' and all the outputs for the simulations of the run are held in the folder in '.csv' files. (Plain text files with one record per line, individual values separated by commas – this file format is readily imported into many other programs such as Excel, Access, SAS and Oracle). ToxFinder can have at most one run directory open at any one time and it is the files in this directory that provide the source data for the parameter values and visualisation.

Simulation A simulation is the result of probabilistically generating a single clinical trial using ToxFinder. ToxFinder creates data for subjects and their responses by randomly selecting from the relevant probability distributions of the scenario defined in the ToxFinder input file. A particular run may be simulated a number of times, indeed to accurately analyse the characteristics of a particular Design it must be simulated usually 1,000-10,000 times over a range of scenarios.

Trial A trial is a real clinical trial, where the user enters subjects' data and responses. The ToxFinder Algorithm is run to determine the doses to allocate and to analyse the results. This facility allows ToxFinder's adaptive allocation to be implemented to conduct small (single centre) trials.

2 Quick Guide

2.1 Starting ToxFinder

To start ToxFinder, double-click the ToxFinder icon that is installed on your computer desktop.

After ToxFinder has started, the main window will be displayed.

2.2 Main Window

The main window is displayed after ToxFinder has started and is present until ToxFinder exits. It is used to open and create designs and to access other parts of the application.

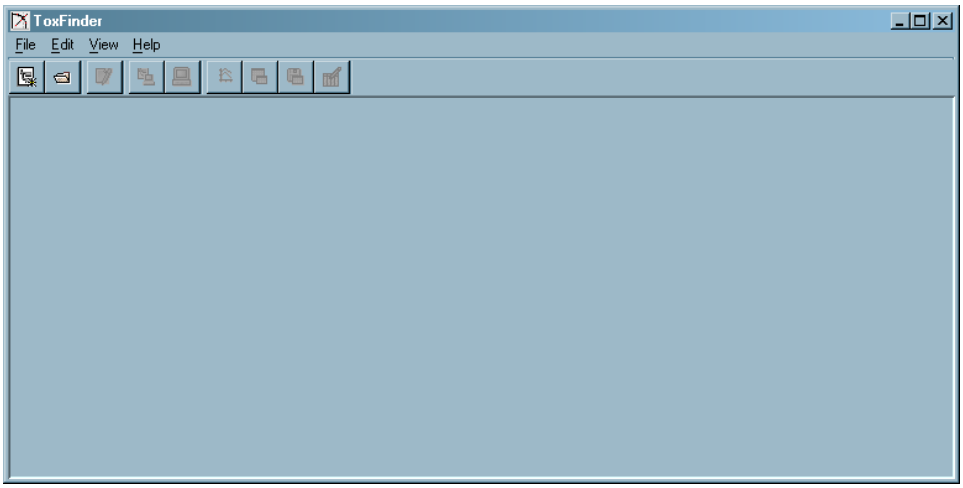

Initially the only enabled buttons are:

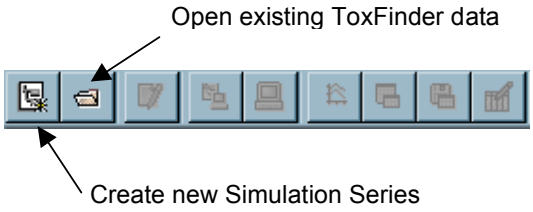

These commands are also available from the **File** menu (along with 'new trial').

2.3 Exiting ToxFinder

To exit ToxFinder, select **Exit** from the **File** menu of the main window. You will be prompted to save any unsaved changes that you have made to parameter values.

2.4 Displaying Existing Data

To display existing ToxFinder Data, select **Open** from the **File** menu, click the tool button or press the `ctrl-o` key. If there are currently any open parameter wizard, run manager, Desktop ToxFinder or visualisation windows, you will be notified that they will be closed before changing the current simulation series and asked if you wish to proceed.

Use the file browser to select the directory to open and then click the **Open** button. If you change your mind about opening a particular directory then click the **Cancel** button. As well as folders, the file browser displays the ToxFinder input ('init.tmml') and output files ('*.csv') in each directory to help you locate the correct one.

If the ToxFinder input files contain any invalid parameter values, then you will be notified via a message box. These values must be corrected in ToxFinder before ToxFinder can be run with this design.

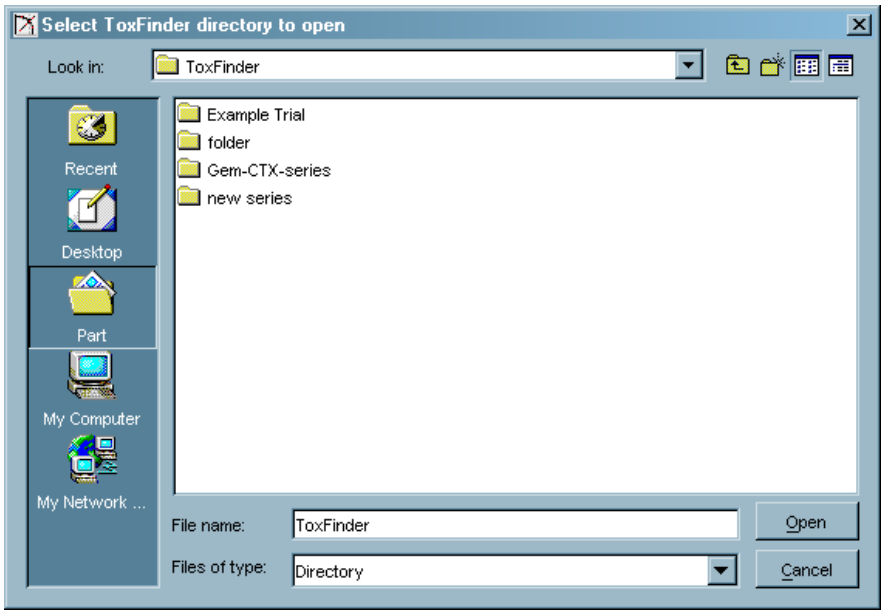

In the folder selected for opening ToxFinder will display all its folders that contain ToxFinder data:

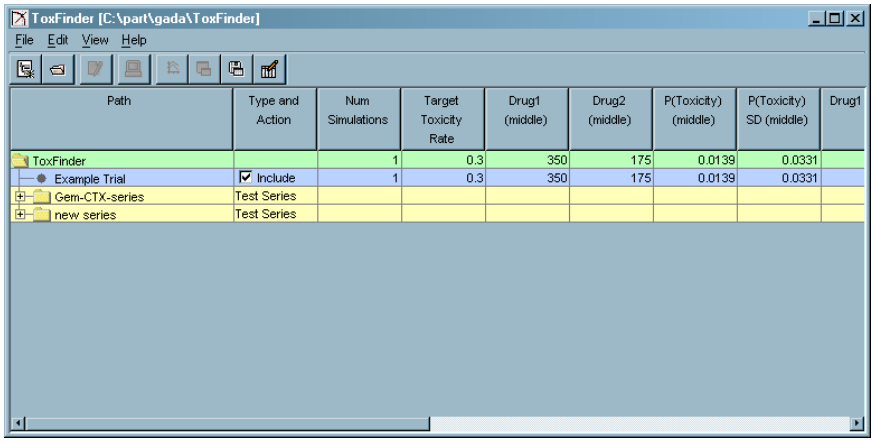

2.5 The Toolbar

Once a simulation series has been created, or a pre-existing one has been opened, the application's main panel becomes more interesting:

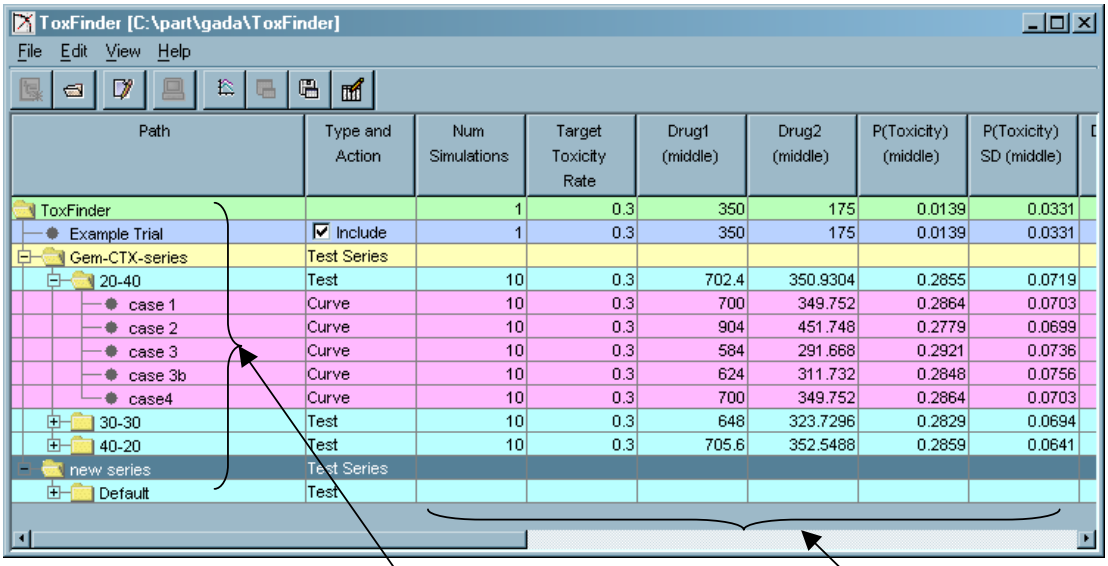

Directory containing the set of Simulation Series & Trials

Simulation Series and Trials

Summary data from simulations or actual trial

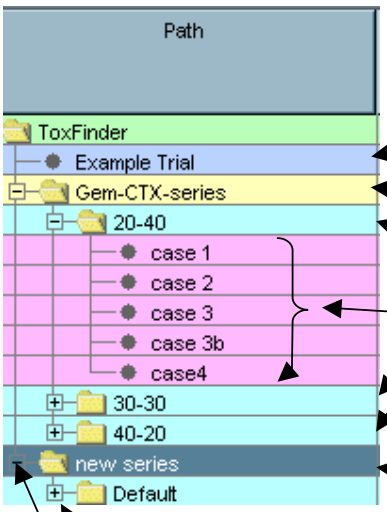

A Trial being run by ToxFinder

A Simulation Series

The design variants of the simulation series

The scenarios of design variant "20-40", all variants are simulated over the same scenarios.

The currently selected series

Levels can be expanded or hidden by clicking on the '+' or '-' boxes.

Select a simulation series or trial by clicking on its name – the line becomes highlighted. More of the ToxFinder Toolbar buttons become enabled:

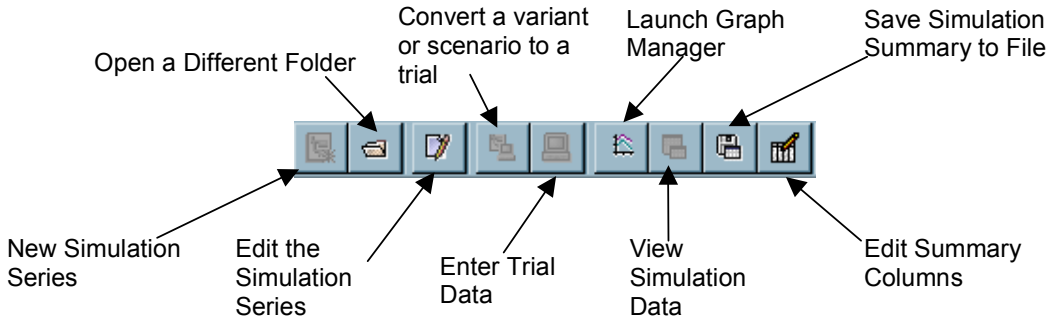

All of these operations are also available from the application's menus.

Right clicking on a Trial or Simulation Series selects it and opens a pop-up menu:

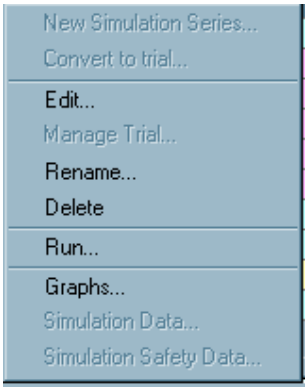

This gives short cuts to commands described elsewhere – it also allows the simulation series (the simulation series directory) to be renamed or deleted.

3 Design and Simulation of ToxFinder Trials

The main mechanism for designing and simulating trials in ToxFinder is the 'Simulation Series'.

A Simulation Series consists of:

1. a 'Design' (which provides the 'Default Parameters' for all of the different simulations),
2. a set of 'Scenarios' – probability surfaces describing the probability of toxicity over the range of combinations of the two agents under test
3. and a set of 'Variants' – variations on the underlying design in order to evaluate options in the trial design.

To create a new simulation series, select either **New Simulation Series** from the **File** menu, or click the New Simulation Series button ().

A file browser will appear in which you should select the directory in which to save the new series:

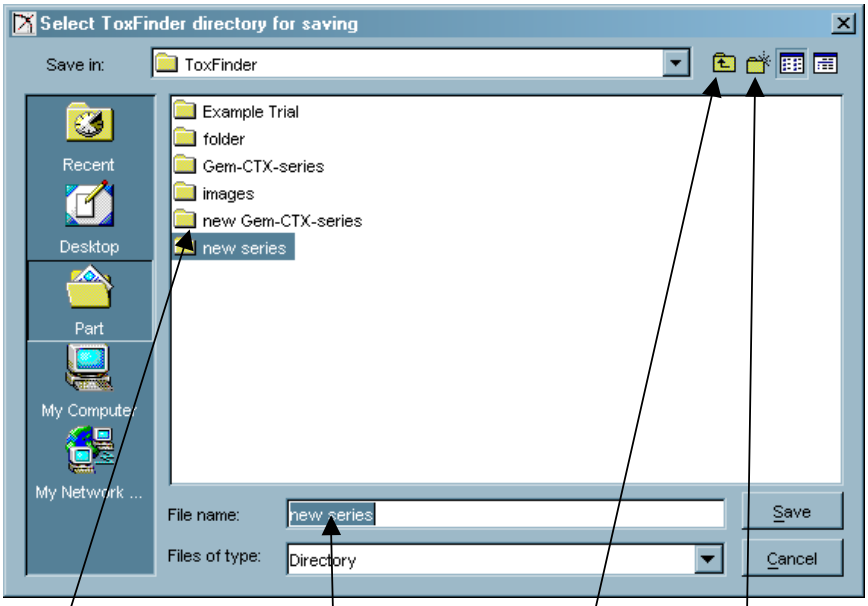

You should navigate to where you wish the new simulation series to be stored:

- double click on folders to move into a folder
- click on the 'parent folder' button to move up to the parent folder
- once in the folder that you want to act as the parent folder for the simulation series, enter the name of the new simulation series – this will be the folder for holding the simulation series data.
- click on the 'new folder' button if you wish to create a new parent folder – this will be useful for organising a number of related ToxFinder simulation series together:

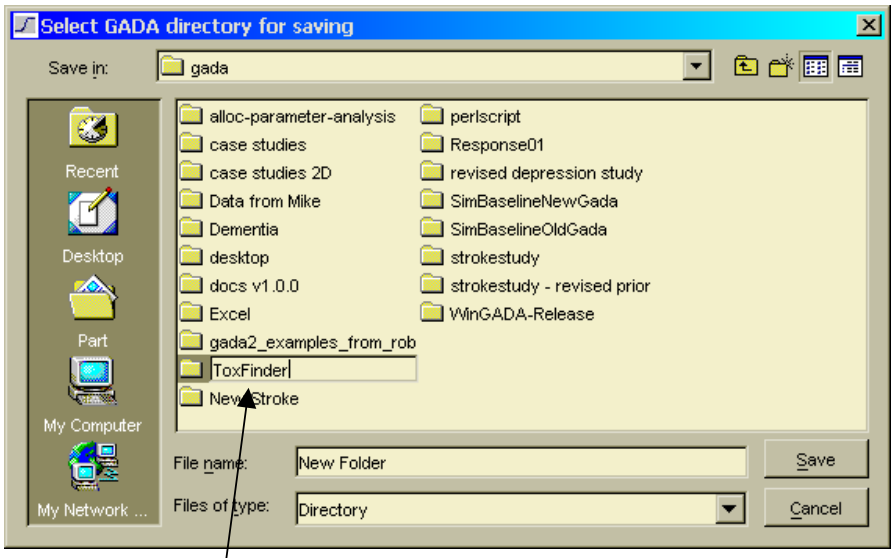

- if you select 'new folder' it will be called 'new folder' but the name will be selected and you can type a more meaningful one.
- you can then double click on that folder and then enter the name of a new directory for the simulation series.

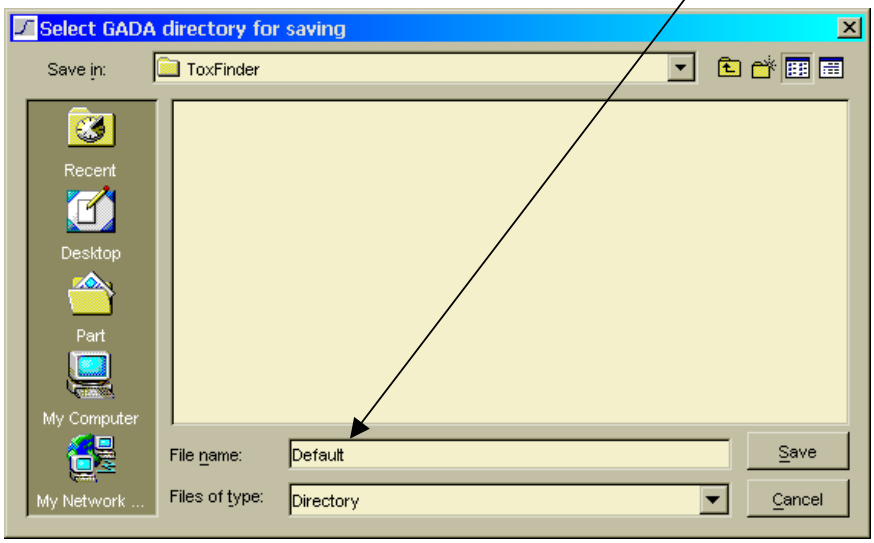

Click **Save** to save the new simulation series. The simulation series editor window will then open.

3.1 The Simulation Series Editor

- This panel allows you to select which aspect of the Simulation Series to define or edit. The step may be selected by clicking on it, or by using the **Previous** and **Next** buttons.
    - Series Definitions – enter the design of the trial.
    - Scenario Definitions – define the different toxicity probability surfaces to simulate from in the simulations.
    - Variants – define variations on the main design of the trial for comparison.
    - Simulation – displays a screen for running the simulations.

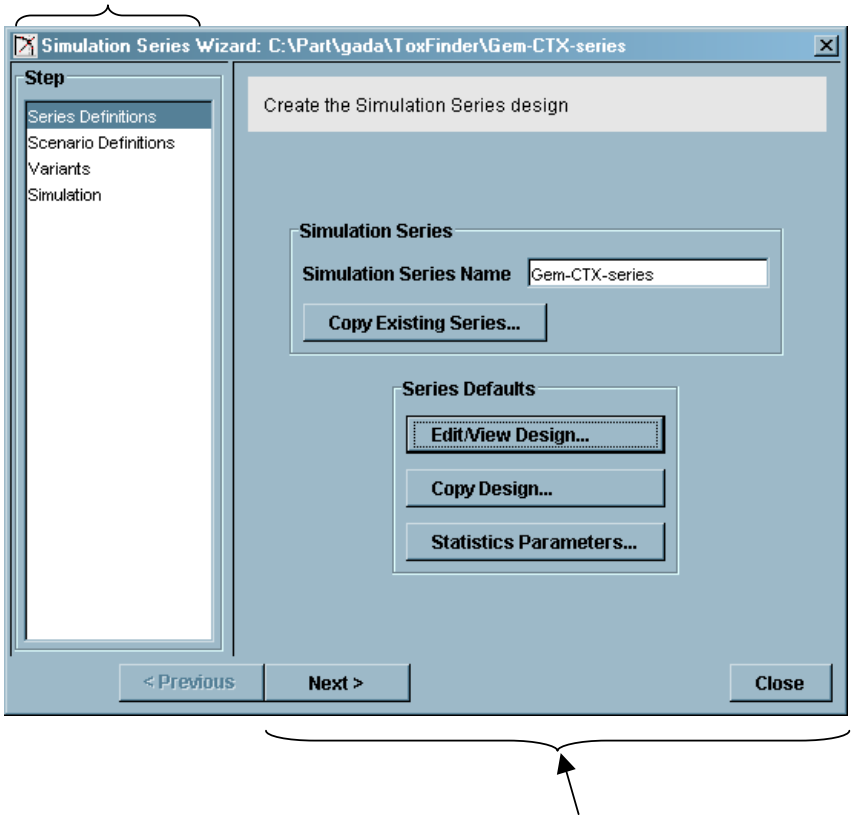

With **Series Definition** selected the rest of the screen looks like this. It allows the series name to be changed, the series design, scenarios and variants to be imported from an existing series (**Copy Existing Series**) or the Design to be setup from scratch (**Edit/View Design**) or imported from a particular run in another simulation series (**Copy Design**).

To design a simulation series from scratch, or modify a design that has been copied in, click on the **Edit/View Design** button.

3.2 Series Definitions - Editing The Design

We will now illustrate setting up a Trial design using the GEM-CTX example in the paper [TMML]. This uses a further sequence of screens where the sequence of steps is listed in the left-hand panel and the controls of the current step are shown on the right.

3.2.1 Defining the Drugs and Target Toxicity

Enter the details of the two drugs to be used:

- The name of the drug
- The units of the dose

Enter the relative Cancer Killing Potential of the two drugs, this is calculated as
λ.drug1 + drug2
this parameter sets the value of λ.

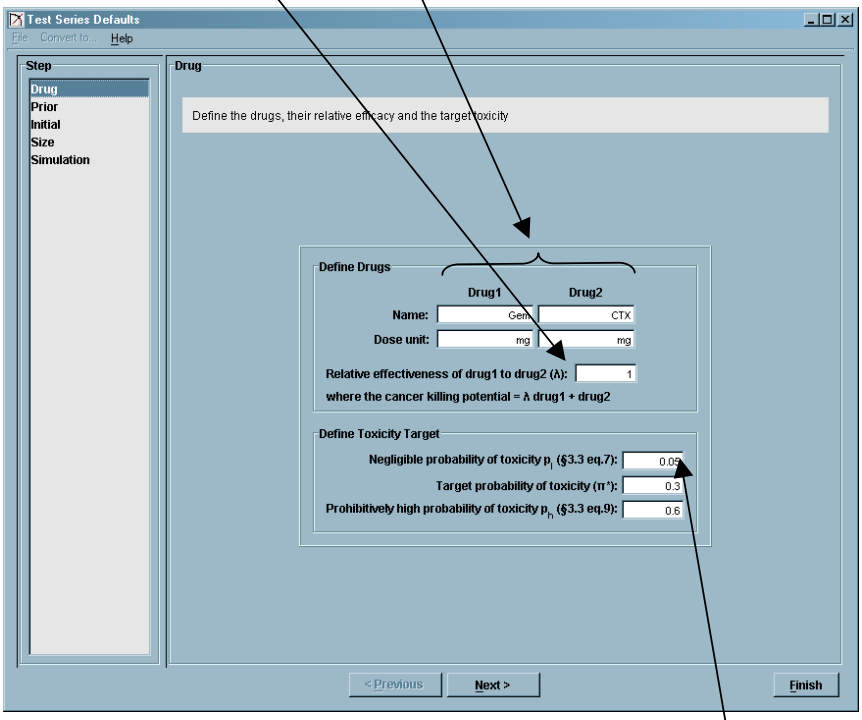

Enter 3 toxicity levels:

- A level of toxicity that is regarded as negligible.
- The level of toxicity that the study will attempt to target, i.e. it will attempt to identify and allocate dose combinations with this level of toxicity.
- A level of toxicity that is prohibitively high.

The prior probabilities of toxicity are specified in terms of the 'negligible' and 'prohibitive' toxicity levels, see the next screen. The two toxicity rates must be either side of the target toxicity rates. The negligible toxicity is used in equation (7) in [TMML], the high toxicity rate is used in equation (9) where the right hand side of the equation is set equal to h/(1-h).

3.2.2 Entering the Prior Toxicity Parameters

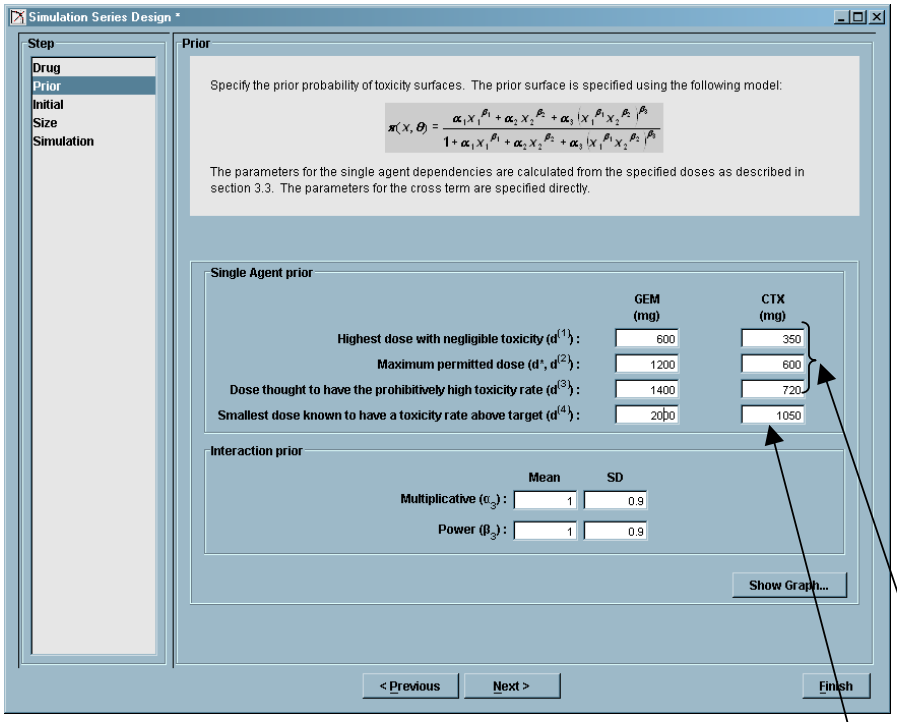

This screen displays the elicited the values that are used to set up the prior, and are best understood from section 3.3 in the paper [TMML].

- We specify, for each drug (on its own) the dose levels likely to correspond to the negligible '$d^{(1)}$', prohibitive toxicity '$d^{(3)}$', and the maximum strength of the drug than can be used on its own '$d^* = d^{(2)}$'.
- Finally, we ask for the dose level that is almost certainly above the target toxicity '$d^{(4)}$'. Although this is for the target toxicity level, that is below the prohibitive toxicity level, because of the much higher confidence level required for this figure typically: $d^{(4)} > d^{(3)}$, but this is not enforced, $d^{(4)} < d^{(3)}$ simply corresponds to a strong prior.

These specify the prior information about the toxicity of each drug considered by itself as a single agent.

For the interaction of the drugs the defaults are similar to those recommended in [TMML], adjusted for this particular implementation.

You can check the effect of the parameters entered for the prior graphically, by clicking **Show Graph**, which displays a contour plot of the resulting prior surface:

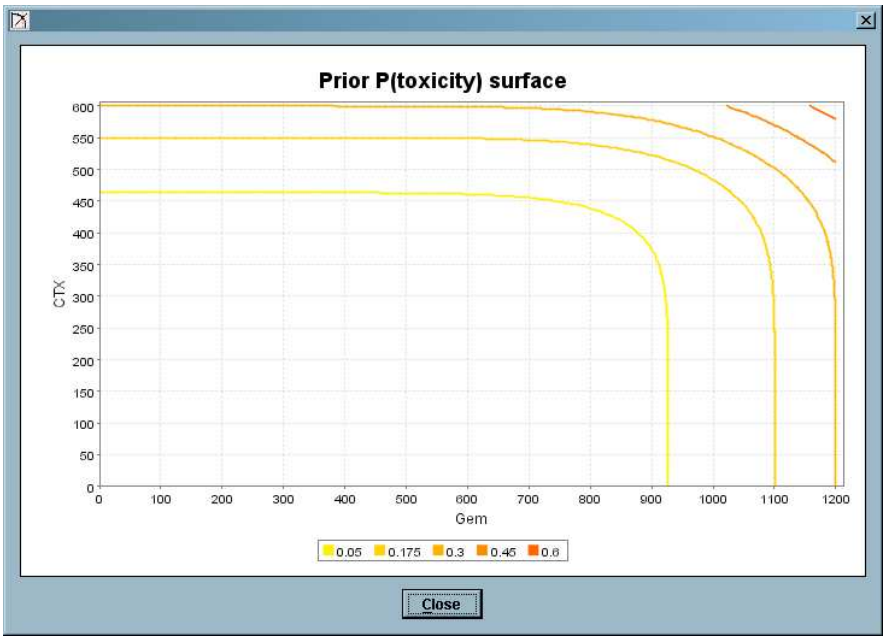

The contours plotted are for the 3 entered toxicity levels and two intermediate values.

Notice that the prior in this illustration corresponds to a case in which the target toxicity probability occurs towards the upper end of the permitted dose range for each drug individually, and there is a slight positive interaction.

3.2.3 Defining The First Stage

On this screen the parameters of stage 1 of the study, where the dose combinations are restricted to a fixed line segment, ‘L1’, are defined.

- The line segment L1 is defined by its end points, the minimum and maximum values for Drug1 and Drug2. This maximum dose should not be greater than the 'single drug acceptable dose', ‘d*’, the answer to question 2 in the prior elicitation section [TMML].

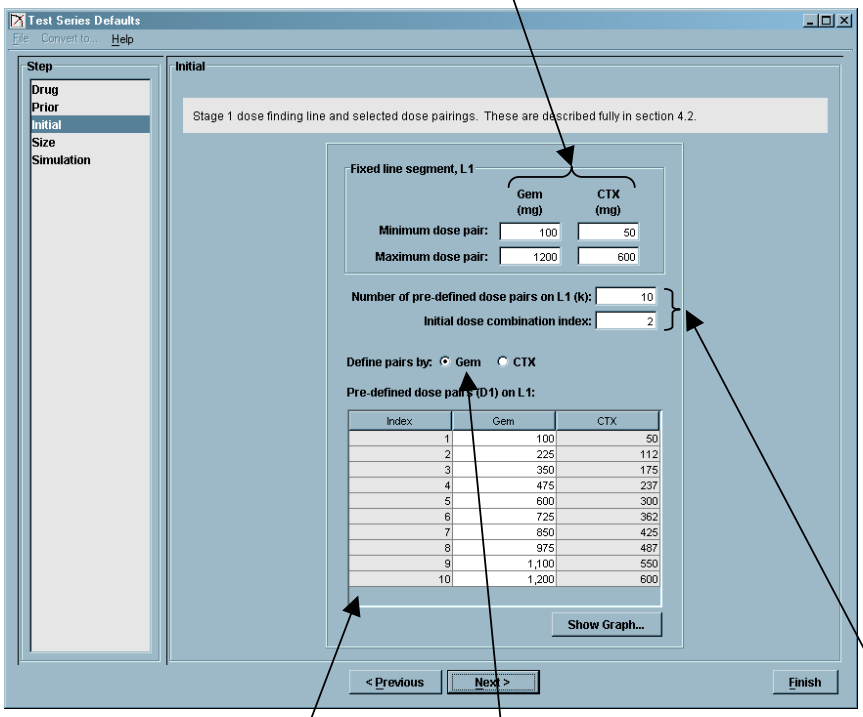

- The number of pre-specified dose combinations along the line segment L1 can be specified. This controls the number of rows in the panel below, and which will be the first dose pair to be allocated.
- As the dose combinations at this stage are constrained to lie in the line segment L1, each combination can be specified by a single value, here select whether you will use values of Drug1 or Drug2 to specify the initial combinations.
- The specified dose pairings (10 in this illustration) are then entered, by entering the values for the Gem half of the pair.

Clicking **Show Graph** displays a chart showing the defined line segment and dose combinations:

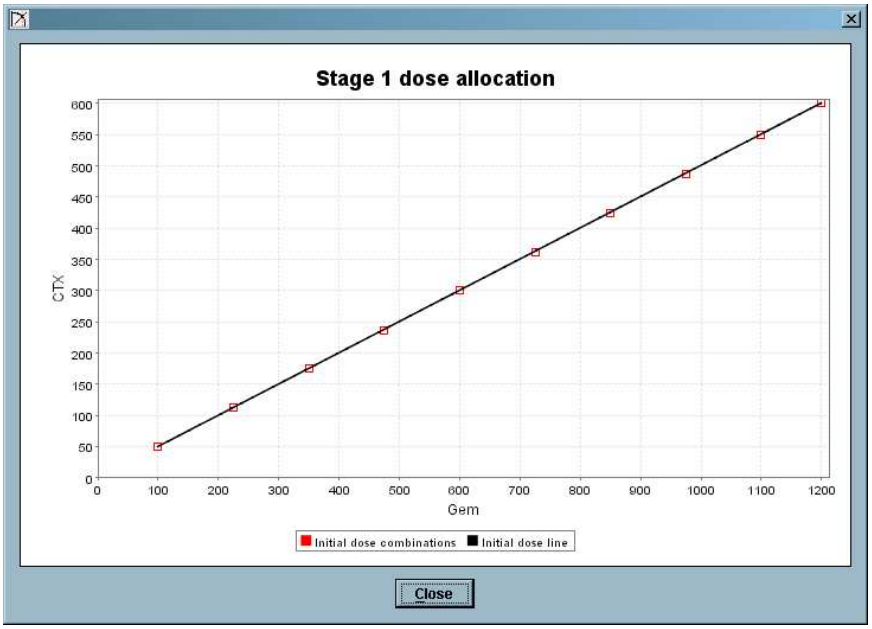

3.2.4 Specifying the Size of the Study

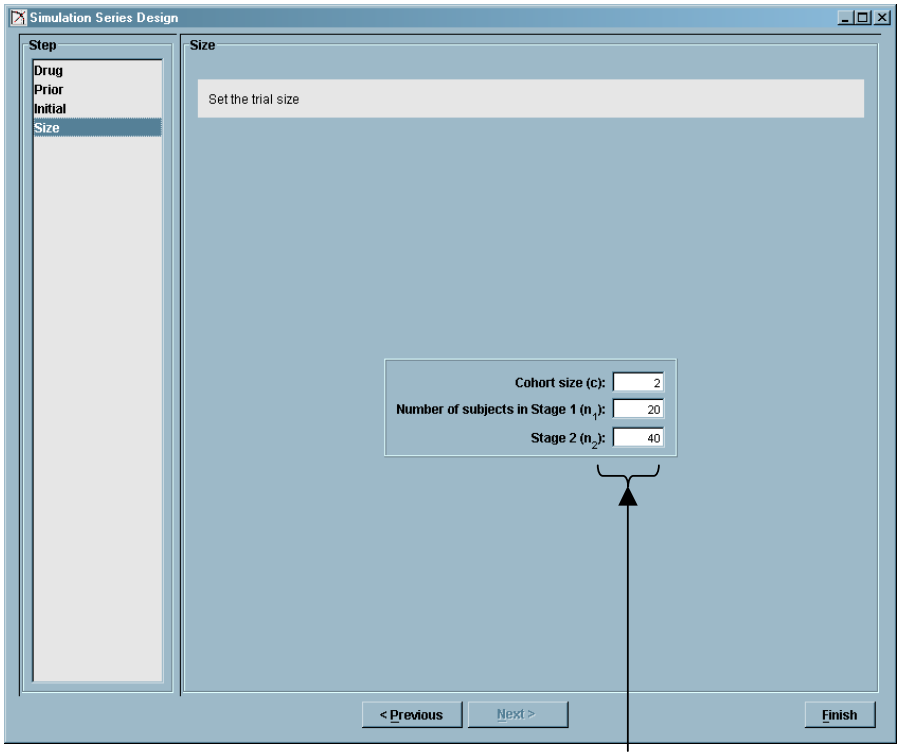

On this the final parameter screen we simply set the size of the cohorts, and the number of subjects (which must be a multiple of the cohort size) in each of the 2 stages.

Clicking on **Finish** saves the parameters.

3.3 Scenario Definitions – The Probability Surfaces to be Simulated

The second step in the overall definition of the Simulation Series is to define some example toxicity surfaces to simulate from to assess the design's performance.

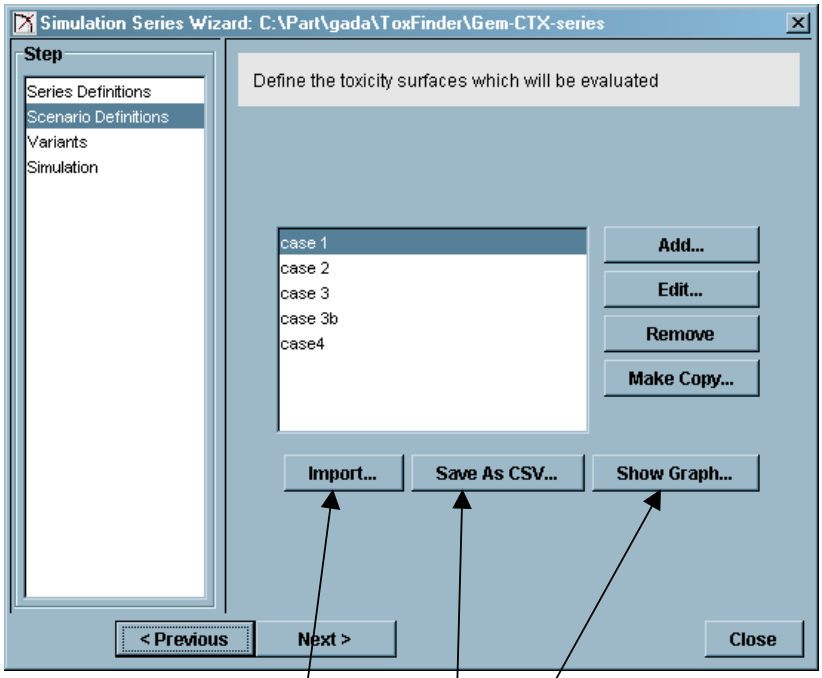

The main buttons allow curves to be imported from other simulation series, and for the scenario data to be saved as CSV files (e.g. for importation into Excel, S-Plus or other graphing tool). ToxFinder itself can also display contour plots of the surfaces.

The second set of buttons are self explanatory – they allow a new scenario to be added to the simulation series, or the scenarios already defined to be edited, removed or copied. Before using these last three, ensure the scenario you want the operation to apply to is selected in the list of scenarios.

3.3.1 Defining a Toxicity Probability Surface

Pressing **Add** or **Edit** displays a screen for defining ToxFinder probability surfaces.

- The grid
- The Number of points to display in the grid.

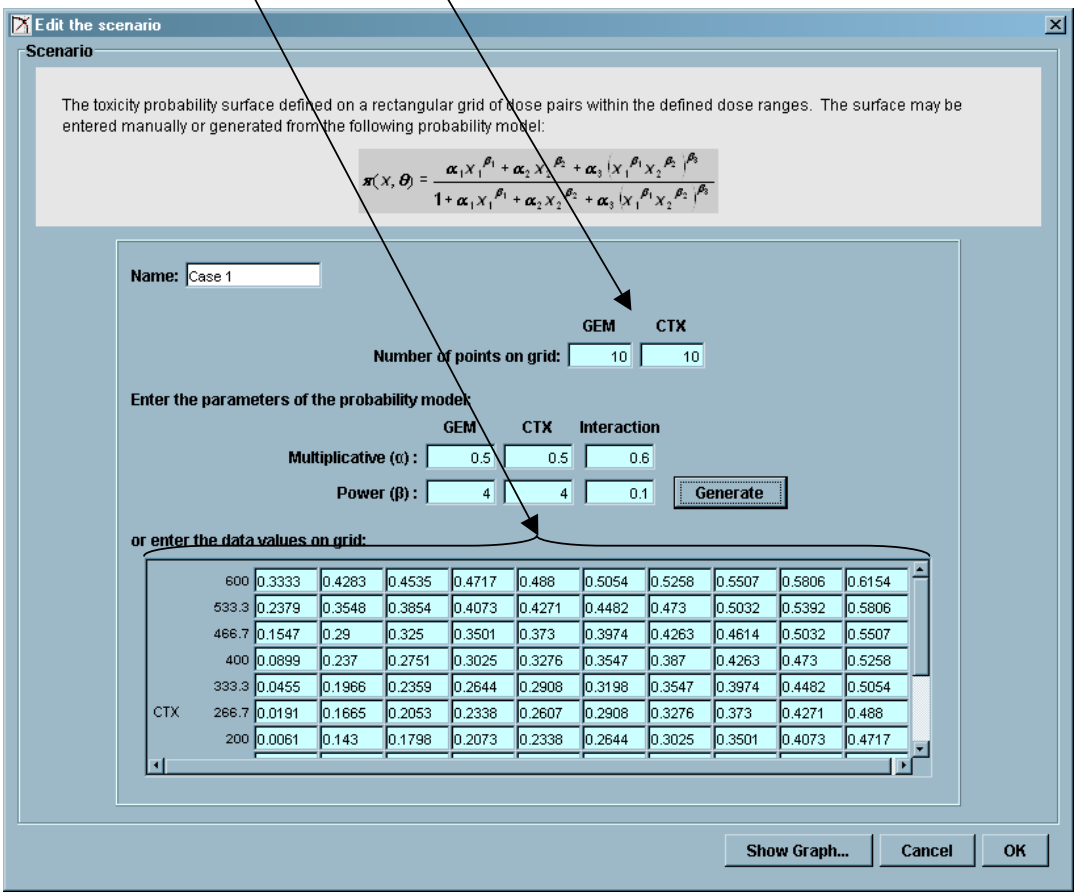

There are two methods for specifying the toxicity to be simulated.

1. The user may enter values in the grid, and scale or add an offset to all the values in the grid.
2. The user may enter 6 values for the parameters of the equation used in the model to specify a surface. After entering the parameters, clicking on **Generate** causes the toxicity probability surface to be created, and the results are shown in the grid.

The two methods may be combined. An initial surface may be generated using the parametric model, and then individual values in the grid may be modified by hand.

3.3.2 Example Toxicity Probability Surfaces Generated from the Parametric Model

	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	1	1	0	1	1	0
Power	4	4	1	8	8	1
	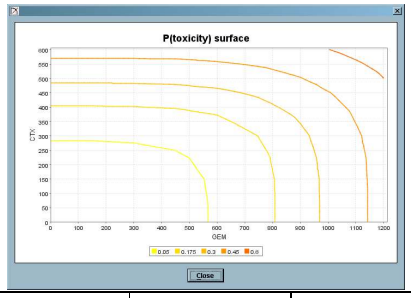			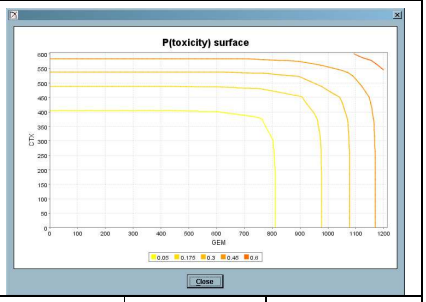		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	1	1	0	1	1	0
Power	1	1	1	2	2	1
	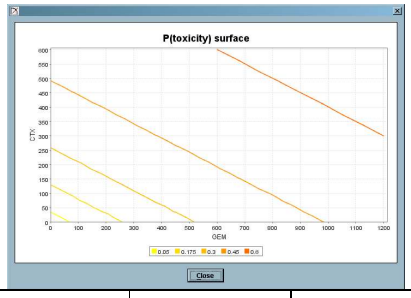			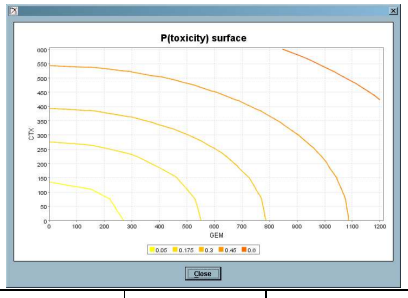		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	1	1	0	1	1	0
Power	0.25	0.25	1	0.5	0.5	1
	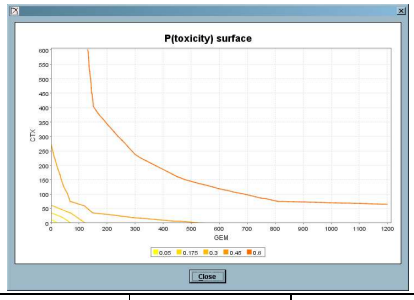			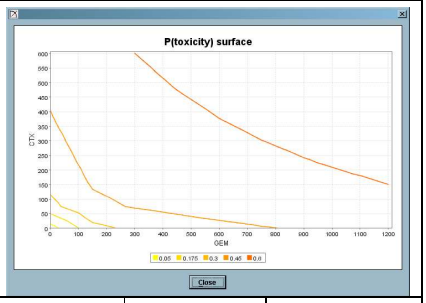		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	1	1	0	1	1	0
Power	0.25	8	1	0.5	4	1
	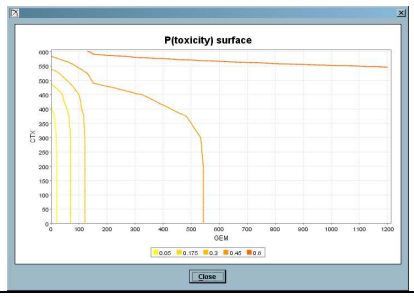			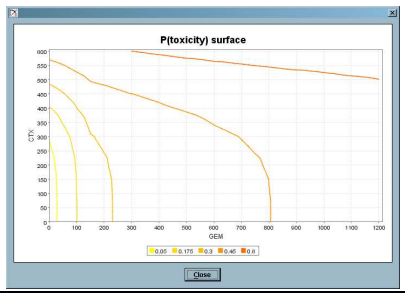		

Models only using the interaction component:

	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0	0	1	0	0	1
Power	1	1	4	1	1	8
	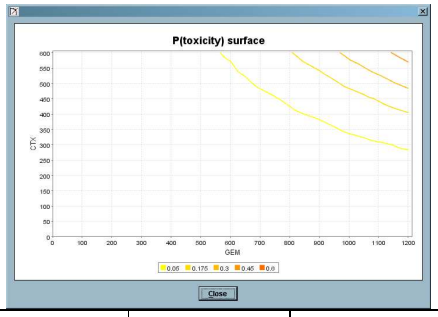			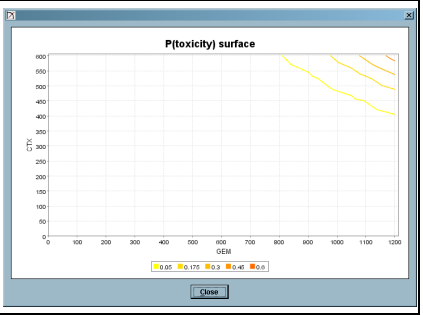		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0	0	1	0	0	1
Power	1	1	1	1	1	2
	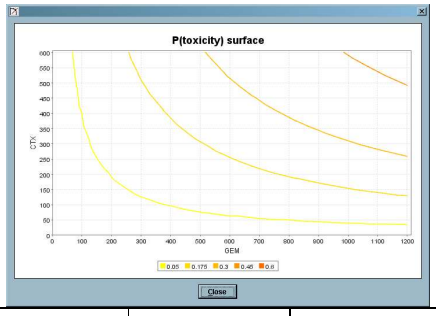			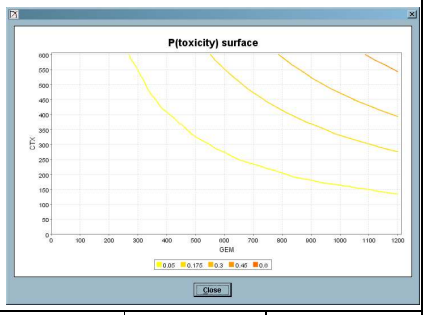		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0	0	1	0	0	1
Power	1	1	0.25	1	1	0.5
	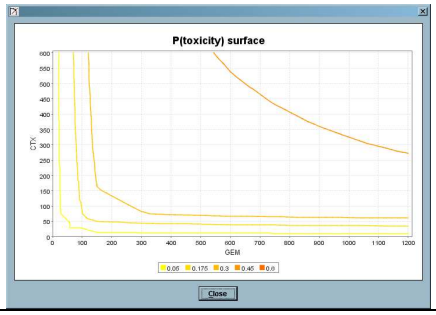			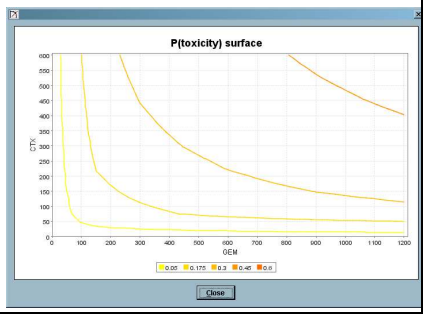		

Example models using combined single agent and interaction components:

	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0.5	0.5	0.5	0.5	0.5	0.5
Power	4	4	0.25	1	1	8
	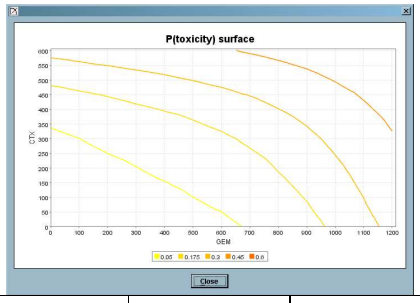			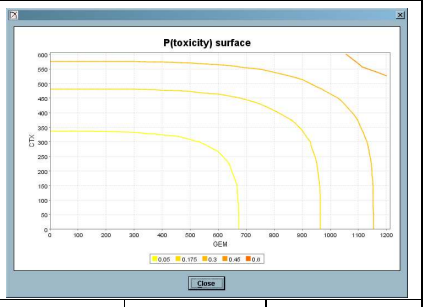		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0.5	0.5	0.5	0.5	0.5	0.5
Power	1	1	0.25	1	1	8
	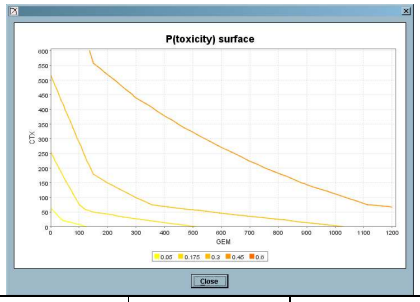			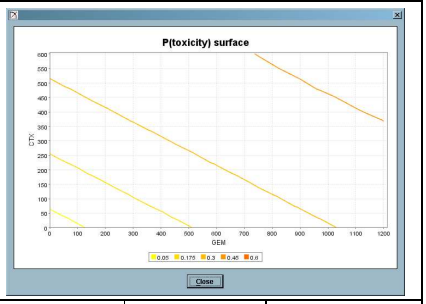		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0.5	0.5	0.5	0.5	0.5	0.5
Power	0.25	0.25	0.25	0.25	0.25	8
	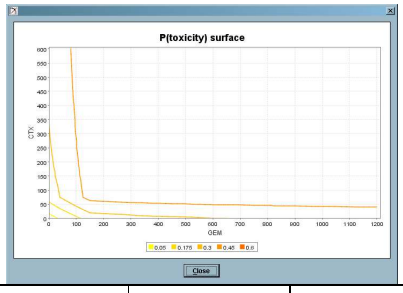			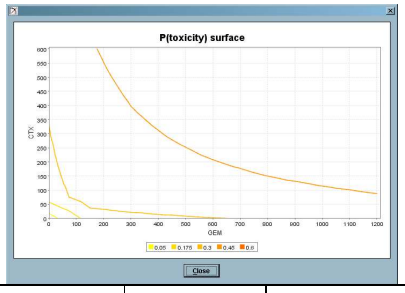		
	Drug 1	Drug 2	Interaction	Drug 1	Drug 2	Interaction
Multiplicative	0.5	0.5	0.5	0.5	0.5	0.5
Power	0.25	8	0.25	0.25	8	8
	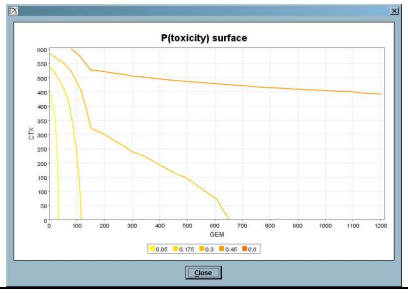			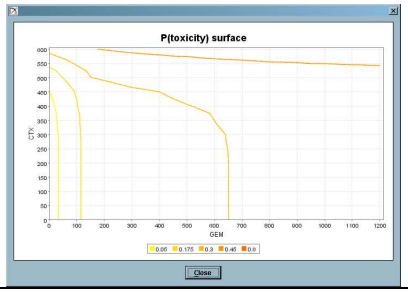		

3.4 Defining Variants

Variants are variations on the default parameters set in the 'Series Definitions' or trial design. In this illustration three versions of the trial design with stage 1 and stage 2 sample sizes (n1,n2) = (20,40), (30,30), and (40,20) are specified. At least one variant is required – but this can simply be a default with no changes.

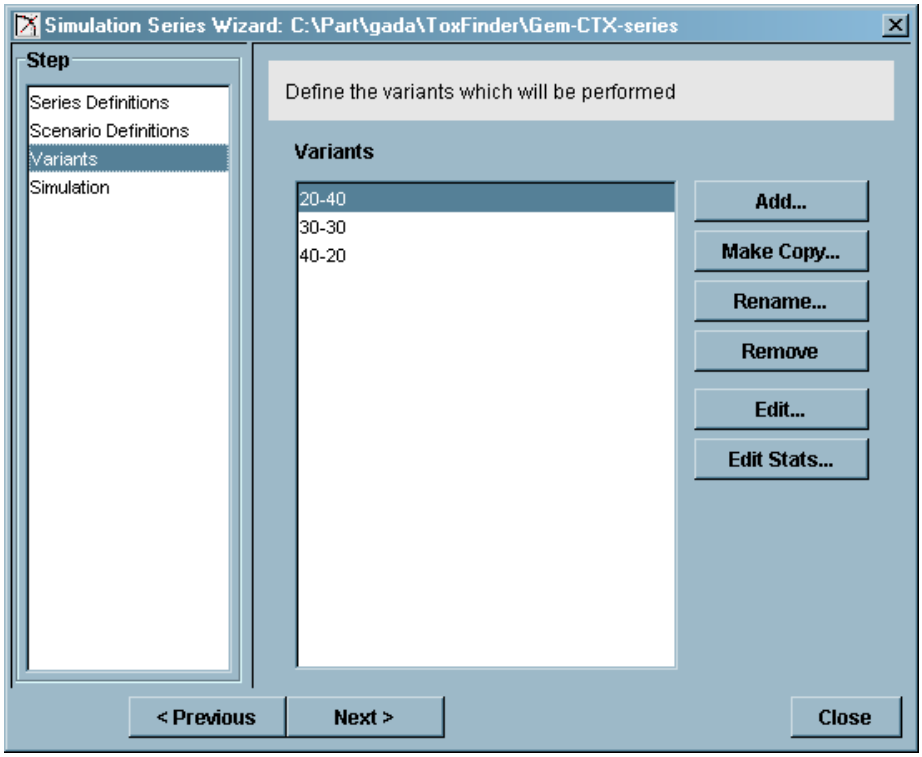

3.4.1 Editing Variants

Editing the variants presents the same screens as used in defining the series, all pre-populated with the design parameter values entered in the series definition. There is no restriction on which parameters, or how many parameters can be changed. It is left to the user to judge what is sensible and what is not.

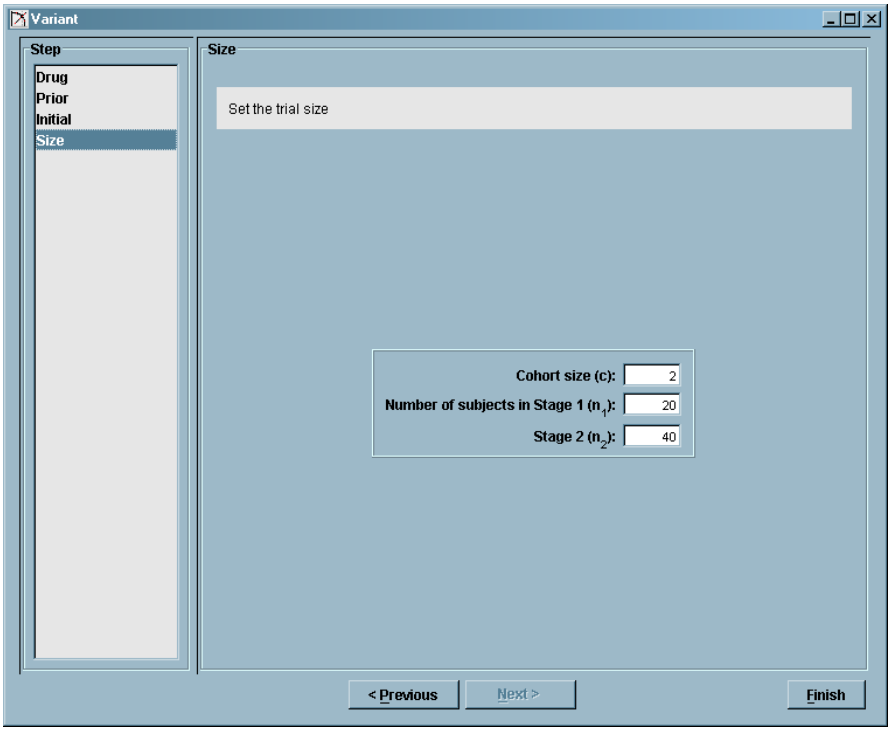

4 The Statistics Parameters

In both the series definition and the variant definition stage it is possible to change the 'statistics parameters'. These parameters are for statisticians who wish to investigate the effects of various choices in the underlying model.

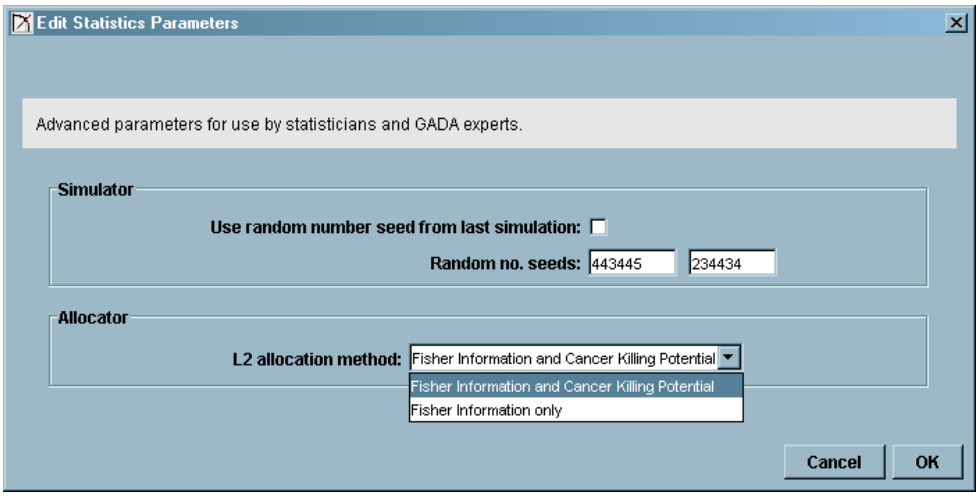

This allows the random number seed to be specified and (possibly for this version only) the L2 allocation method to be set. There are 2 L2 allocation methods that can be selected:

1. Fisher Information and Cancer Killing Potential (as per [TMML] but using the 2nd derivative Fisher Information Matrix)
2. Fisher Information only

If you make changes on this screen it is necessary to use the **File > Save** command from the menu to save them.

5 Simulating Trials using ToxFinder

The implementation of the TMML algorithm can run on the simulation series from the last step in the Simulation Series editor. These all cause the Simulator screen to be displayed:

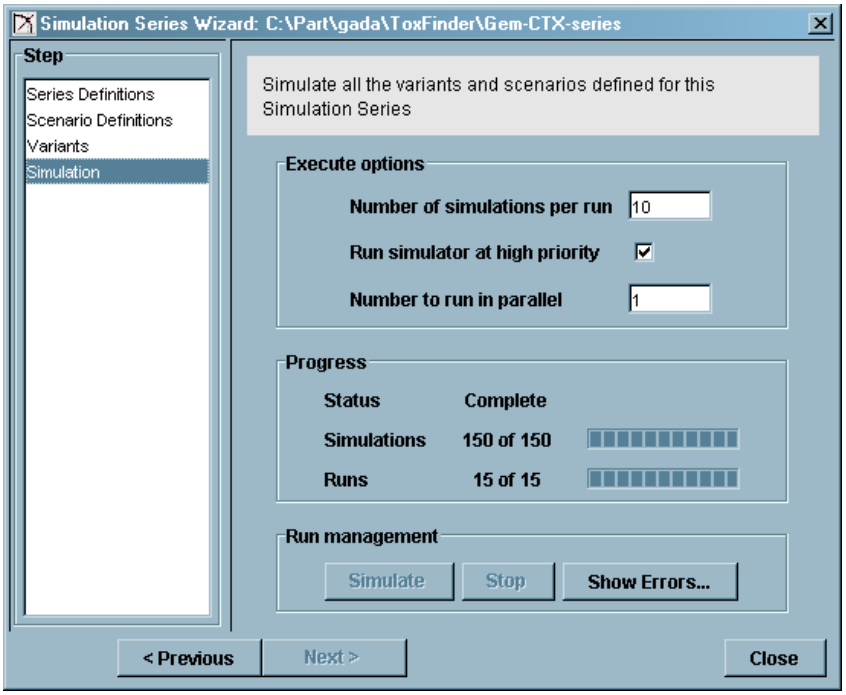

The 3 parameters to set control:

- the number of trials to simulate each Run,
- whether to run them at high priority – this causes the simulations to run faster but makes it difficult to use your computer for anything else at the same time. If this option is left unchecked then the simulations are run in the background and your computer gives priority to whatever else the user is going while they run.
- The 'number to run in parallel' should normally be left at 1, however on very new PC's with CPU's with 'Hyper-Threaded' architectures, these CPUs can effectively run two programs simultaneously with no loss of performance to either one. On these PCs it is worth setting this parameter to '2'. To tell if your PC has this kind of CPU, right click on the Windows task bar (normally at the bottom of the screen). Select 'Task manager' from the pop-up menu and select the 'Performance' tab on the Windows Task Manager dialog. If you see two charts of 'CPU usage history' then you have a Hyper-threaded CPU, otherwise not.
- Once you have clicked 'Simulate' the progress bars will update as the simulations run.

6 Data

6.1 Summary Data

After running simulations the main window can show the summary data for the different scenarios and design variants simulated. The summary data falls into two groups – overall data and then data specific to each of the 3 finally selected dose pairs at the target toxicity probability the 'left' pair using more of Drug 2, the 'middle' pair lying on the diagonal line L1 and the 'right' pair using more of Drug 1.

It is possible to edit the selection of columns displayed using the column editor by clicking on the 'column editor' button in the toolbar (), or selecting **Change Columns** from the **Run** menu. See below for a description of the column editor.

It is also possible to save the contents of the all the columns to a 'comma separated values' (CSV) text file that can be easily imported into a number of other applications such as Excel, Access, SAS, S-Plus etc.. by clicking on the 'Save summary data to file' button on the toolbar (), or selecting **Export Table** from the **File** menu.

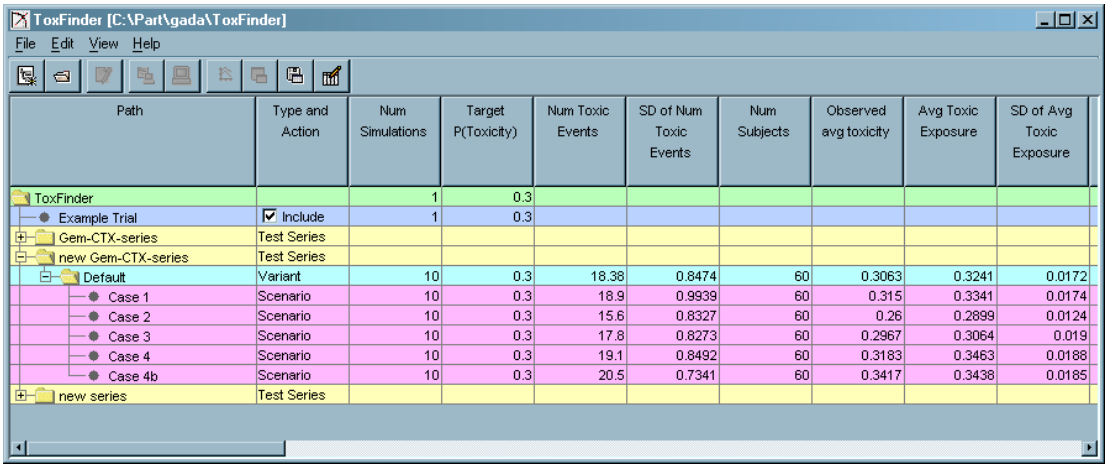

Path	Type and Action	Num Simulations	Target P(Toxicity)	Num Toxic Events	SD of Num Toxic Events	Num Subjects	Observed avg toxicity	Avg Toxic Exposure	SD of Avg Toxic Exposure
ToxFinder		1	0.3						
Example Trial	Include	1	0.3						
Gem-CTX-series	Test Series								
new Gem-CTX-series	Test Series								
Default	Variant	10	0.3	18.38	0.8474	60	0.3063	0.3241	0.0172
Case 1	Scenario	10	0.3	18.9	0.9939	60	0.315	0.3341	0.0174
Case 2	Scenario	10	0.3	15.6	0.8327	60	0.26	0.2899	0.0124
Case 3	Scenario	10	0.3	17.8	0.8273	60	0.2967	0.3064	0.019
Case 4	Scenario	10	0.3	19.1	0.8492	60	0.3183	0.3463	0.0188
Case 4b	Scenario	10	0.3	20.5	0.7341	60	0.3417	0.3438	0.0185
new series	Test Series								

These columns are:

Num Simulations	The number of trials simulated for each scenario, only 10 have been run in the example above – this is a wholly inadequate number to be able to draw reliable conclusions about the operating characteristics of the design.
Target P(Toxicity)	This is the 'acceptable' level of toxicity 'π*' in the design.
Num Toxic Events	The mean number of subjects, per simulated trial, who were simulated to have a toxic event.
SD of Num Toxic Events	The SD of the result: "Num Toxic Events" over the set of simulated trials.
Num Subjects	The number of subjects in each simulated trial.
Observed avg. toxicity	'Num Toxic Events' / 'Num Subjects'.
Avg Toxic Exposure	The mean, over the simulated trials, of the average toxicity probability of the dose combinations allocated to the subjects in the trial. This is, given the scenario being simulated, the average risk of experiencing toxicity during the trial.
SD of Avg Toxic Exposure	The SD of the result: "Avg Toxic Exposure" over the set of simulated trials.

For each of the 3 final recommended dose pairs (left/middle/right) the following summary data is presented (here the data for the middle dose pair is shown):

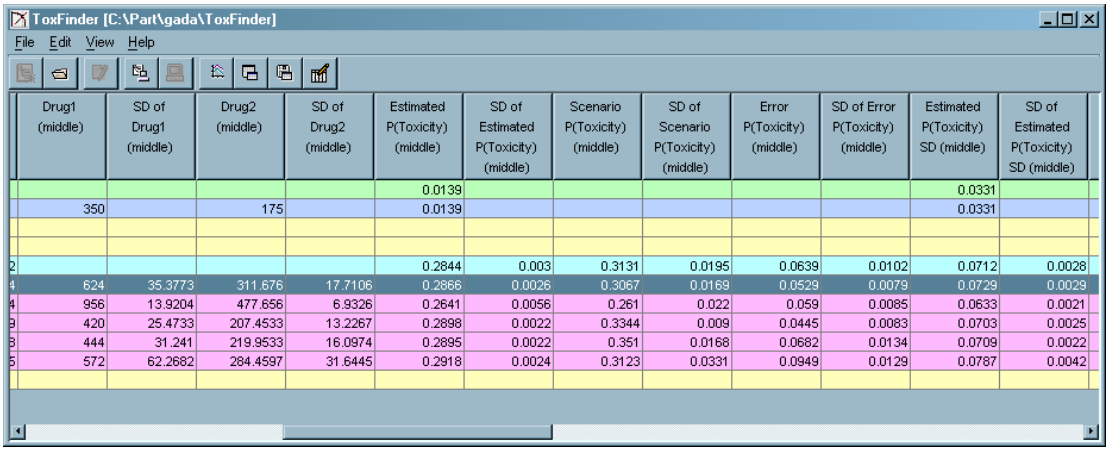

Drug1 (middle)	SD of Drug1 (middle)	Drug2 (middle)	SD of Drug2 (middle)	Estimated P(Toxicity) (middle)	SD of Estimated P(Toxicity) (middle)	Scenario P(Toxicity) (middle)	SD of Scenario P(Toxicity) (middle)	Error P(Toxicity) (middle)	SD of Error P(Toxicity) (middle)	Estimated P(Toxicity) SD (middle)	SD of Estimated P(Toxicity) SD (middle)
				0.0139						0.0331	
350		175		0.0139						0.0331	
2				0.2844	0.003	0.3131	0.0195	0.0639	0.0102	0.0712	0.0028
4 624	35.3773	311.676	17.7106	0.2866	0.0026	0.3067	0.0169	0.0529	0.0079	0.0729	0.0029
4 956	13.9204	477.656	6.9326	0.2641	0.0056	0.261	0.022	0.059	0.0085	0.0633	0.0021
9 420	25.4733	207.4533	13.2267	0.2898	0.0022	0.3344	0.009	0.0445	0.0083	0.0703	0.0025
8 444	31.241	219.9533	16.0974	0.2895	0.0022	0.351	0.0168	0.0682	0.0134	0.0709	0.0022
6 572	62.2682	284.4597	31.6445	0.2918	0.0024	0.3123	0.0331	0.0949	0.0129	0.0787	0.0042

Drug 1	The mean strength of the dose of Drug 1 in the (middle) selected combination over the simulated trials.
SD of Drug 1	The SD of the result: “Drug 1” over the set of simulated trials.
Drug 2	The mean strength of the dose of Drug 2 in the (middle) selected combination over the simulated trials.
SD of Drug 2	The SD of the result: “Drug 2” over the set of simulated trials.
Estimated P(Toxicity)	The mean of the estimated toxicities ‘$\overline{\pi}(x)$’ at the (middle) selected combination over the simulated trials.
SD of estimated P(Toxicity)	The SD of the result: “Estimated P(Toxicity)” over the set of simulated trials.
Scenario P(Toxicity)	The mean of the scenario toxicity surface at the (middle) selected combination, ‘true $\pi(x)$’ over the set of simulated trials.
SD of Scenario P(Toxicity)	The SD of the result: “Scenario P(Toxicity)” over the set of simulated trials.
Error P(Toxicity)	The mean absolute error between the estimated toxicity and the scenario toxicity at the (middle) selected combination over the set of simulated trials.
SD of Error P(Toxicity)	The SD of the result: “Error P(Toxicity)” over the set of simulated trials.
Estimated P(Toxicity) SD	The mean estimate of the SD of the result $[\pi(x)\|Z_n]$ at the (middle) selected combination over the set of simulated trials.
SD of Estimated P(Toxicity) SD	The SD of the result: “Estimated P(Toxicity) SD” over the set of simulated trials.

6.2 Modifying the Columns Displayed

For any of the data summary screens it is possible to configure the columns displayed and their order:

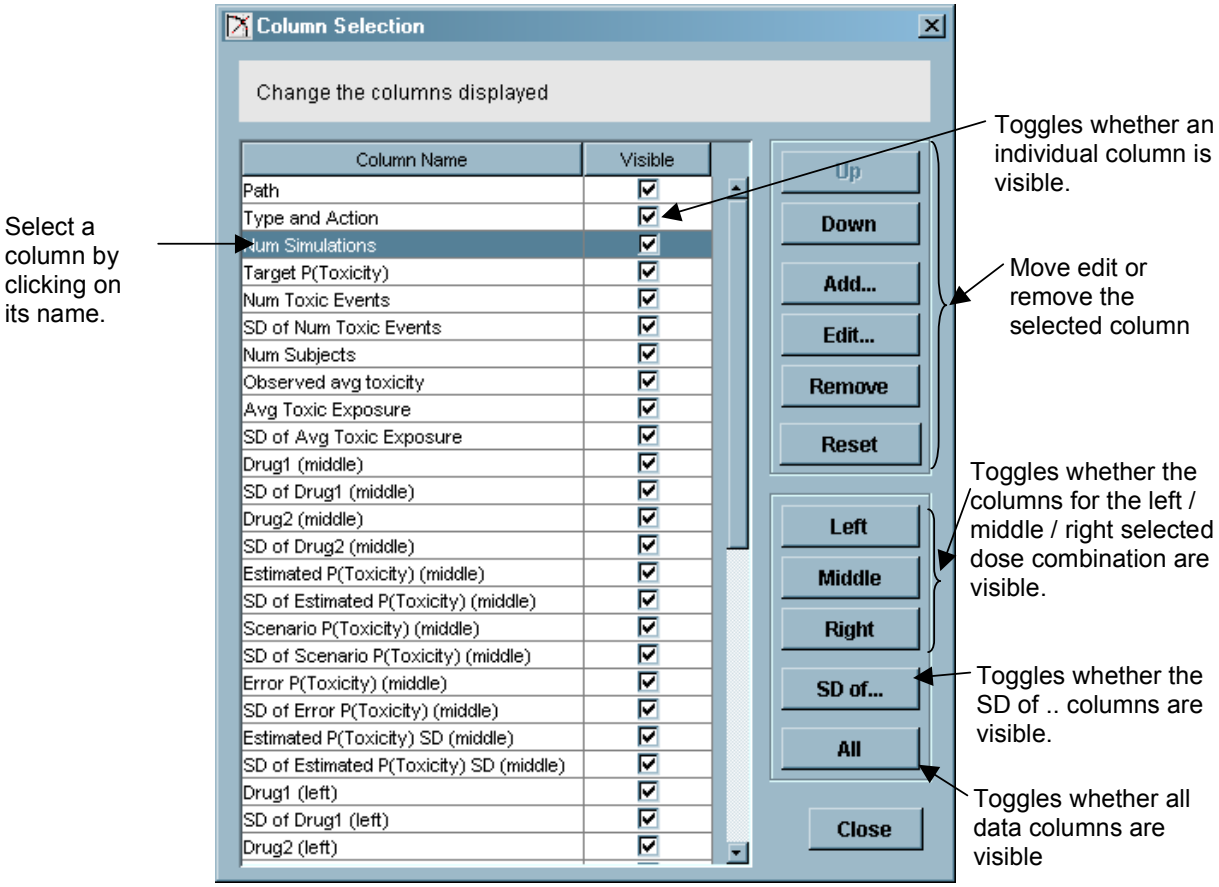

6.3 Editing or Adding a New Column

If you select to edit or add a new column the following dialog is displayed:

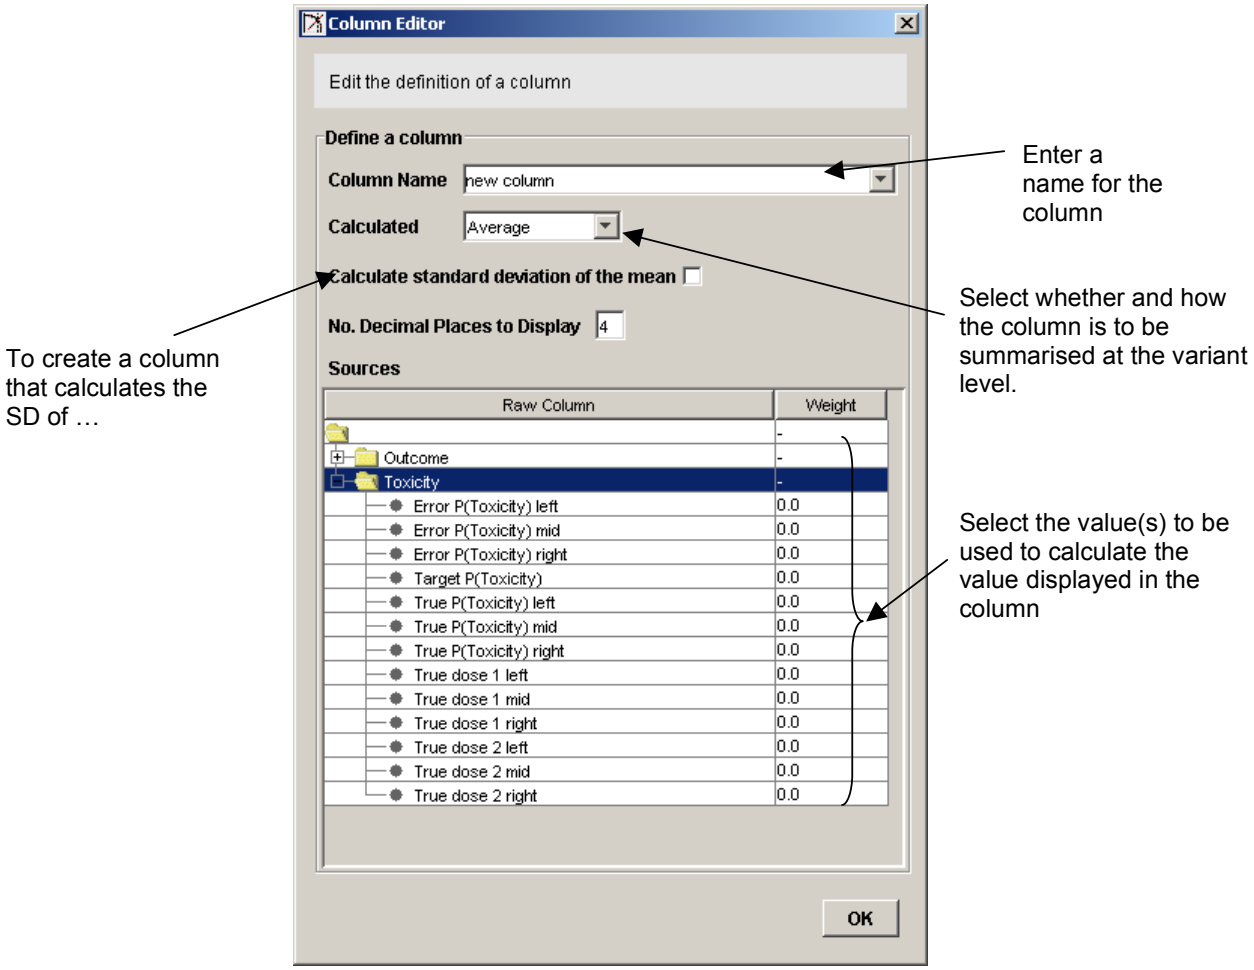

This dialog has couple of features that warrant further explanation:

- The values to be displayed per run, can be defined as the weighted sum of a number of values. To display a single particular value simply set its weight to '1'. To scale it (e.g. by 100) enter the value to scale it by. To sum a set of values set all their weights to 1, to average a set of values set their weights to the appropriate decimal fraction for the number of values being averaged.
- The values can be summarised (or not) at the Design Variant level – either as average, ranking or not at all (e.g. for recommended dose values). For instance variants could be ranked by their mean number of toxic events.

6.4 Simulation Data

In addition to the summary data displayed in the ToxFinder main window, at the individual run level it is also possible to see simulation results data and safety data. These can be accessed by right clicking on the Run's line in the main ToxFinder window and selecting '**Simulation Data …**' or **Simulation Safety Data …**' or by selecting the Run's line and accessing these commands from the **View** menu.

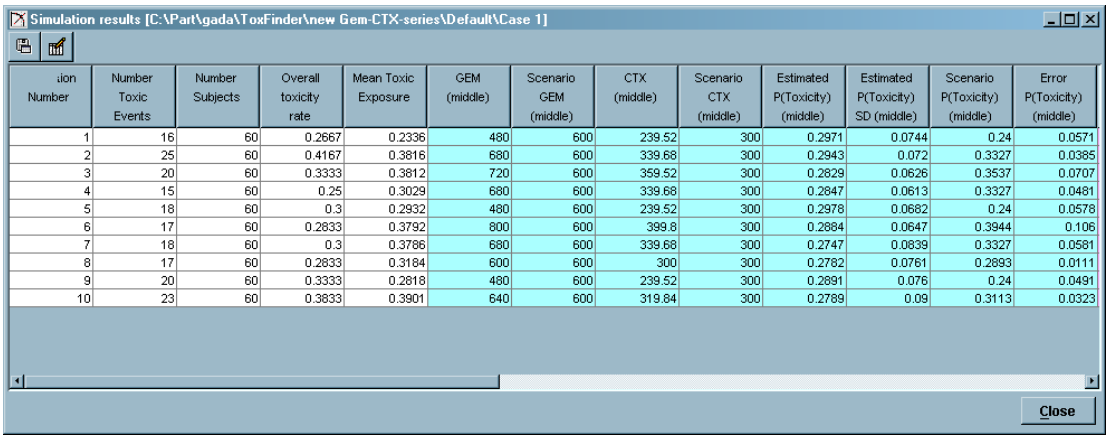

ion Number	Number Toxic Events	Number Subjects	Overall toxicity rate	Mean Toxic Exposure	GEM (middle)	Scenario GEM (middle)	CTX (middle)	Scenario CTX (middle)	Estimated P(Toxicity) (middle)	Estimated P(Toxicity) SD (middle)	Scenario P(Toxicity) (middle)	Error P(Toxicity) (middle)
1	16	60	0.2667	0.2336	480	600	239.52	300	0.2971	0.0744	0.24	0.0571
2	25	60	0.4167	0.3816	680	600	339.68	300	0.2943	0.072	0.3327	0.0385
3	20	60	0.3333	0.3812	720	600	359.52	300	0.2829	0.0626	0.3537	0.0707
4	15	60	0.25	0.3029	680	600	339.68	300	0.2847	0.0613	0.3327	0.0481
5	18	60	0.3	0.2932	480	600	239.52	300	0.2978	0.0682	0.24	0.0578
6	17	60	0.2833	0.3792	800	600	399.8	300	0.2884	0.0647	0.3944	0.106
7	18	60	0.3	0.3786	680	600	339.68	300	0.2747	0.0839	0.3327	0.0581
8	17	60	0.2833	0.3184	600	600	300	300	0.2782	0.0761	0.2893	0.0111
9	20	60	0.3333	0.2818	480	600	239.52	300	0.2891	0.076	0.24	0.0491
10	23	60	0.3833	0.3901	640	600	319.84	300	0.2789	0.09	0.3113	0.0323

The columns are the same as for the summary data, except they show the results per simulated trial, not the means, and hence there are no SD's of the mean's either.

6.5 Simulation Safety Data

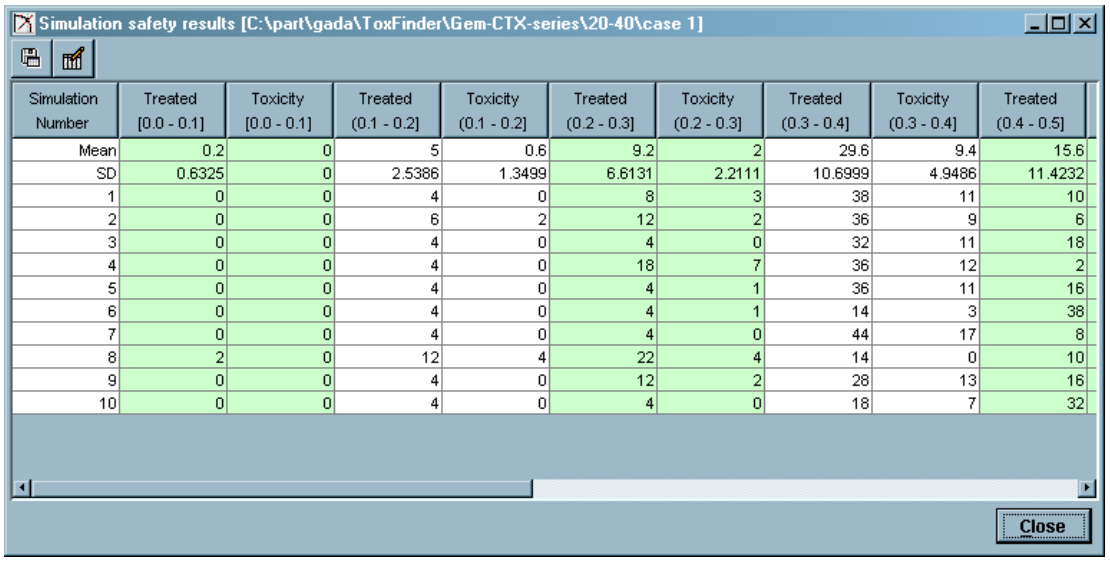

Simulation Number	Treated [0.0 - 0.1]	Toxicity [0.0 - 0.1]	Treated (0.1 - 0.2]	Toxicity (0.1 - 0.2]	Treated (0.2 - 0.3]	Toxicity (0.2 - 0.3]	Treated (0.3 - 0.4]	Toxicity (0.3 - 0.4]	Treated (0.4 - 0.5]
Mean	0.2	0	5	0.6	9.2	2	29.6	9.4	15.6
SD	0.6325	0	2.5386	1.3499	6.6131	2.2111	10.6999	4.9486	11.4232
1	0	0	4	0	8	3	38	11	10
2	0	0	6	2	12	2	36	9	6
3	0	0	4	0	4	0	32	11	18
4	0	0	4	0	18	7	36	12	2
5	0	0	4	0	4	1	36	11	16
6	0	0	4	0	4	1	14	3	38
7	0	0	4	0	4	0	44	17	8
8	2	0	12	4	22	4	14	0	10
9	0	0	4	0	12	2	28	13	16
10	0	0	4	0	4	0	18	7	32

The simulation safety table shows the number of subjects allocated to each decile of the toxicity range and the number of toxic events experienced by those subjects. The data shows the mean, SD and individual values for each simulated trial.

7 Graphs and Charts

Once simulations have been run, it is possible to see charts and graphs of the results.

If you select a Simulation Series, Variant or individual Run -

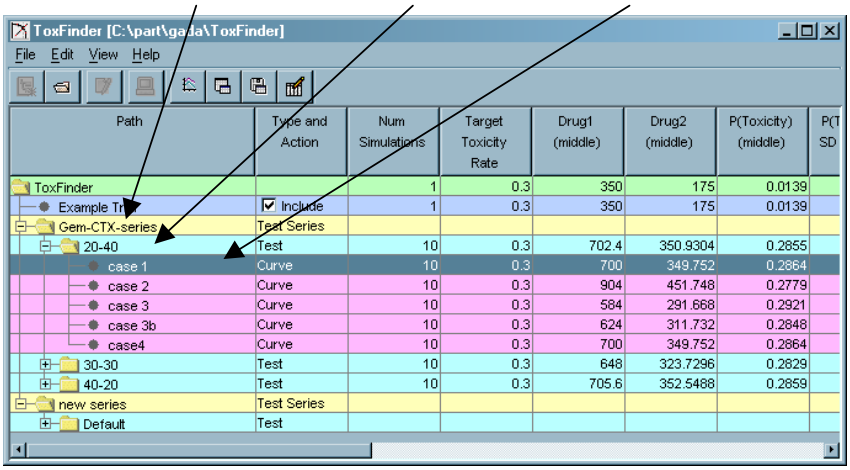

You can then view the graph manager by clicking on the graph manager button on the tool bar (), selecting **Graph Manager** from the **View** menu or right clicking on the Simulation Series, Variant or Run line in the ToxFinder main window and selecting **Graphs** from the pop-up menu.

7.1 Histograms for Series and Variants

If the Graph Manager is opened on a Simulation Series or a Variant, then you can select to display the data in any of the columns as a histogram.

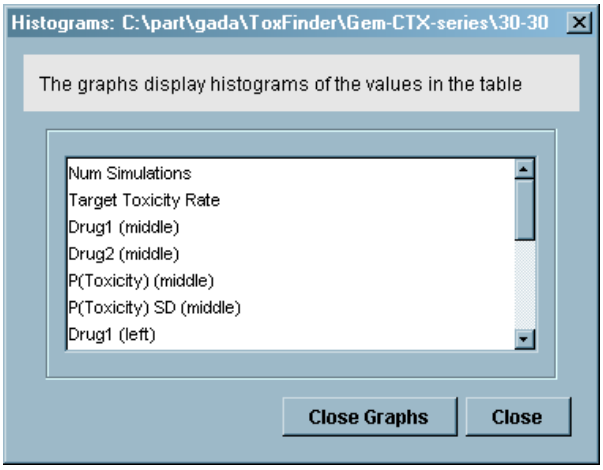

For a Simulation Series the histogram is plotted with a column per Design Variant in the series:

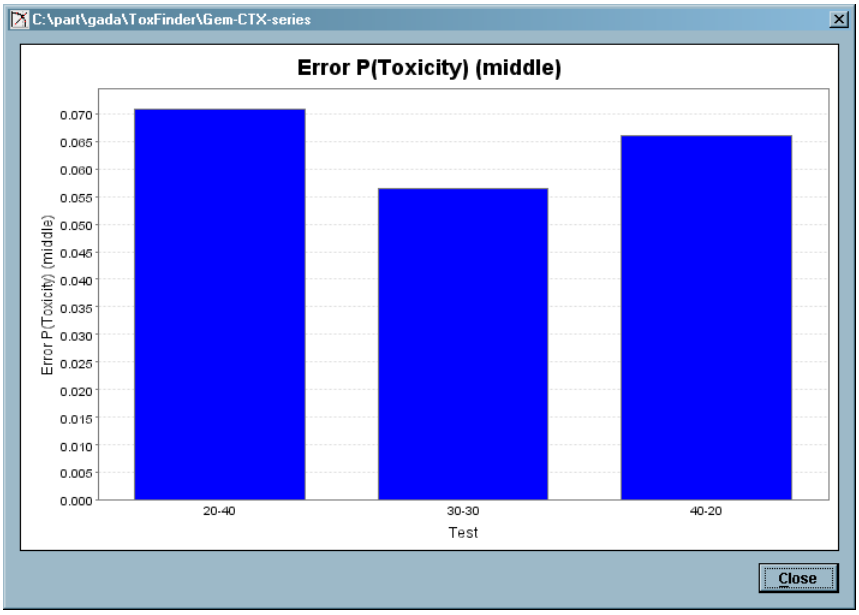

For an individual variant, the histogram is plotted with a column per Scenario:

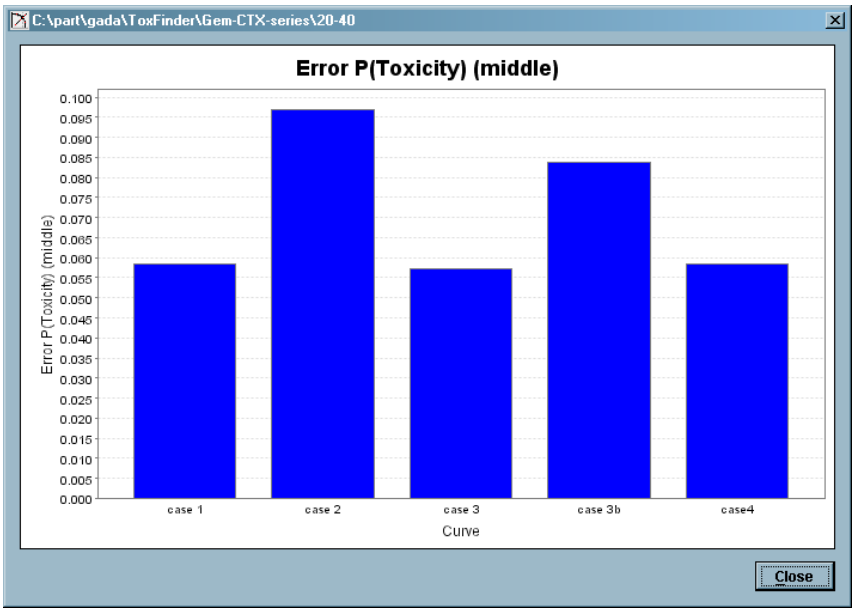

7.2 Charts for Runs

For an individual Run the charts available are quite different:

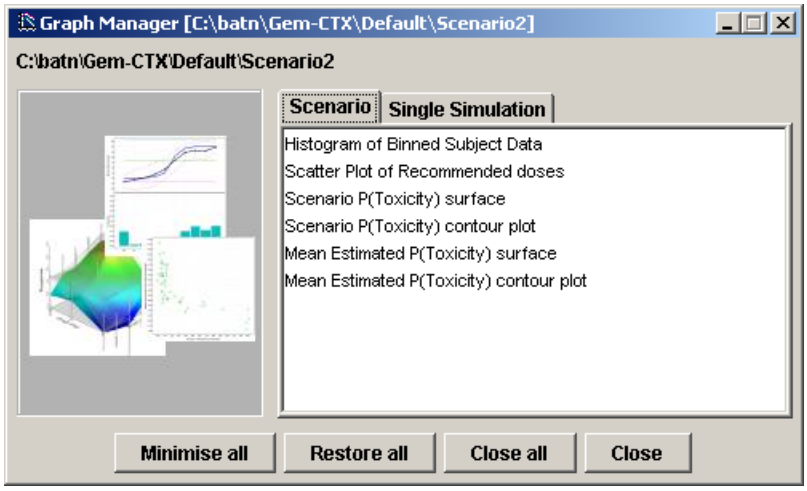

All charts have a standard set of buttons in the top left corner:

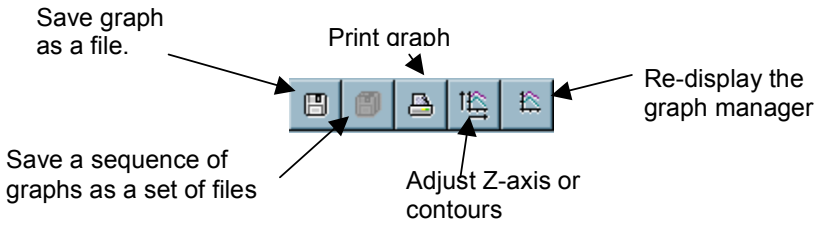

All 3D surface plots have a standard set of controls in the top right corner:

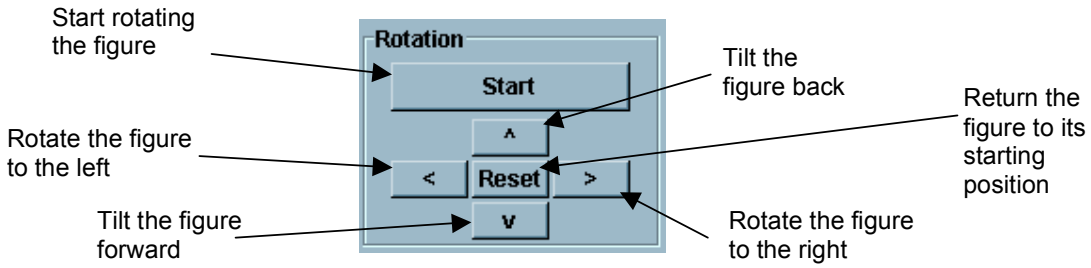

Graphs for individual simulations have a 'selector' control at the bottom:

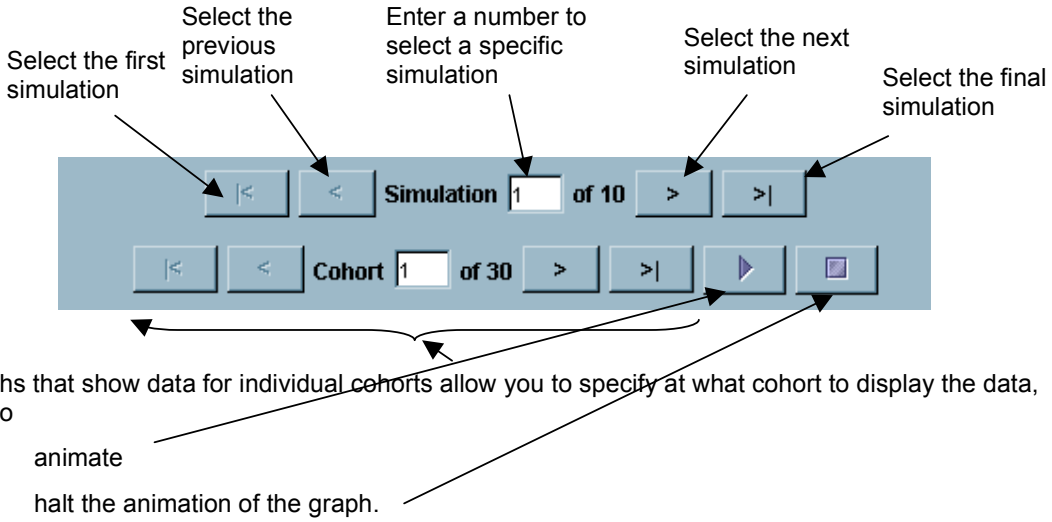

Graphs that show data for individual cohorts allow you to specify at what cohort to display the data, and to

- animate
- halt the animation of the graph.

7.2.1 Histogram of Binned Subject Data

The Histogram of Binned Subject Data shows the number of subjects allocated to dose combinations with different levels of toxicity, and the number that experienced toxicity.

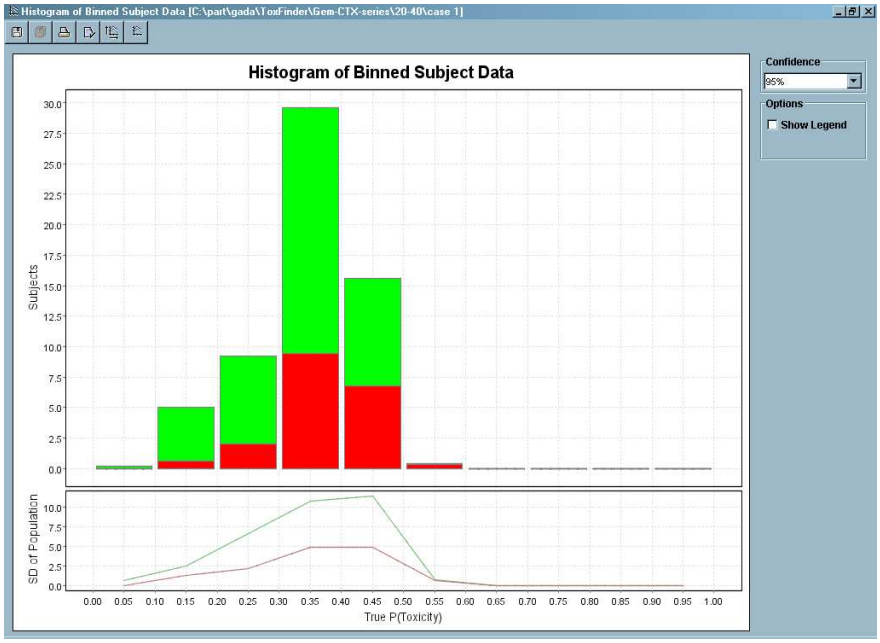

Below the histogram the line plot shows the variation (as a standard deviation) of the population totals across the simulations.

7.2.2 Scatter Plot of Recommended Doses

This plot shows the finally selected right, middle and left dose combinations of all the simulations in the Run, superimposed on the P(Toxicity) contours used in the scenario for the high, target and negligible levels of toxicity.

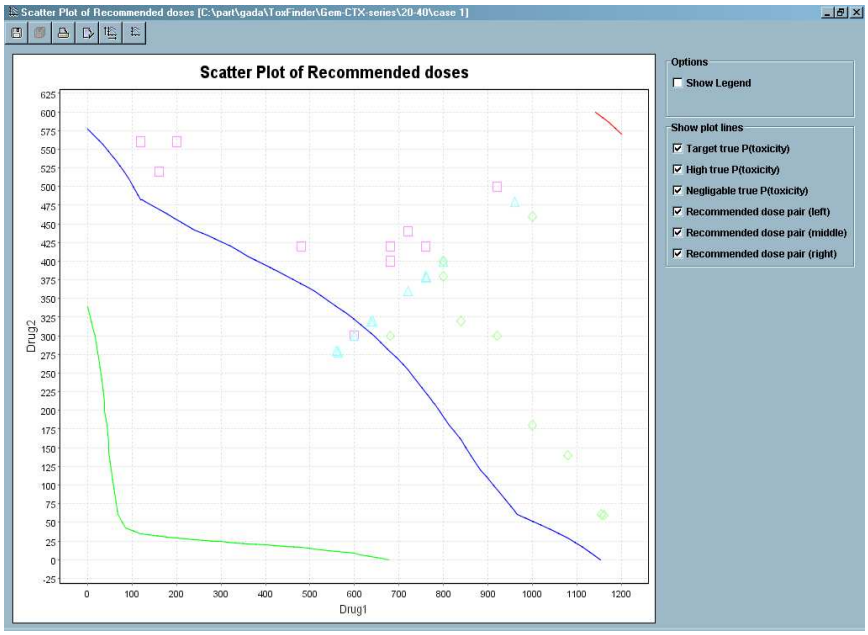

7.2.3 Scenario P(Toxicity) Surface & Contour Plot

The scenario P(Toxicity) surface – the 'true' one that the model is trying to estimate, can be displayed as a 3D surface plot:

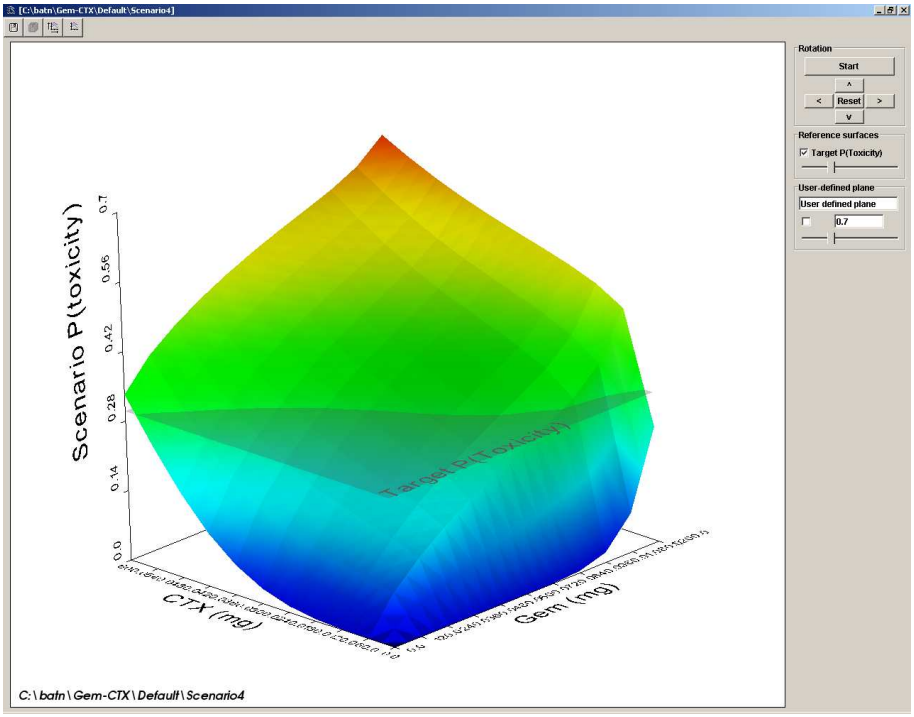

or as a 2D contour plot:

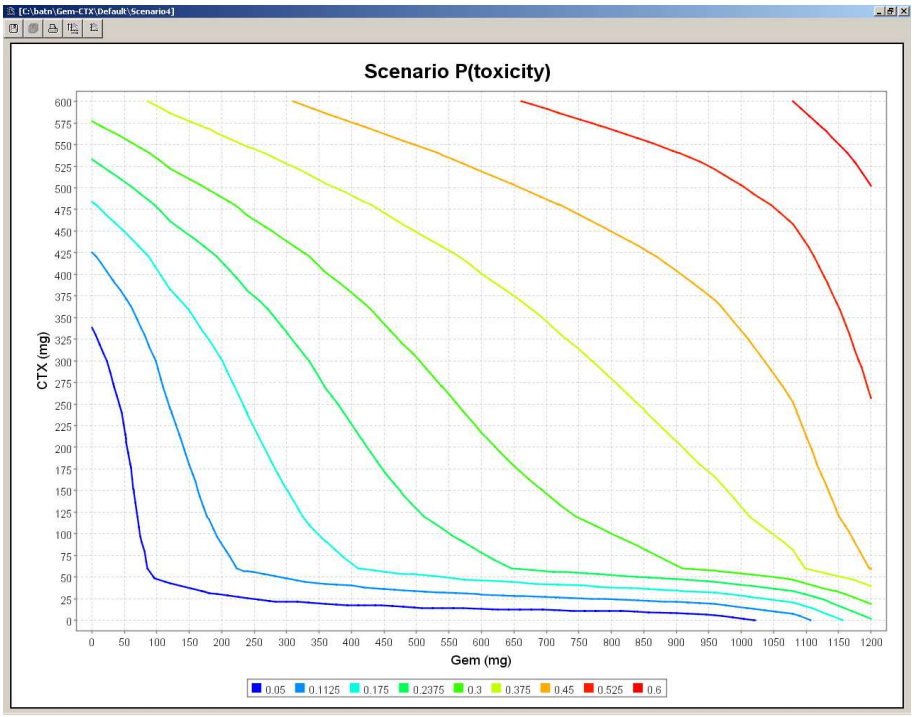

7.2.4 Mean Estimated P(Toxicity) Surface & Contour Plot

The Mean Estimated Toxicity Surface shows the average toxicity surface found over all the simulations with error bars calculated from the distribution of the simulations. This surface can some time to compute the first time it is displayed, depending on the number of simulations to be averaged.

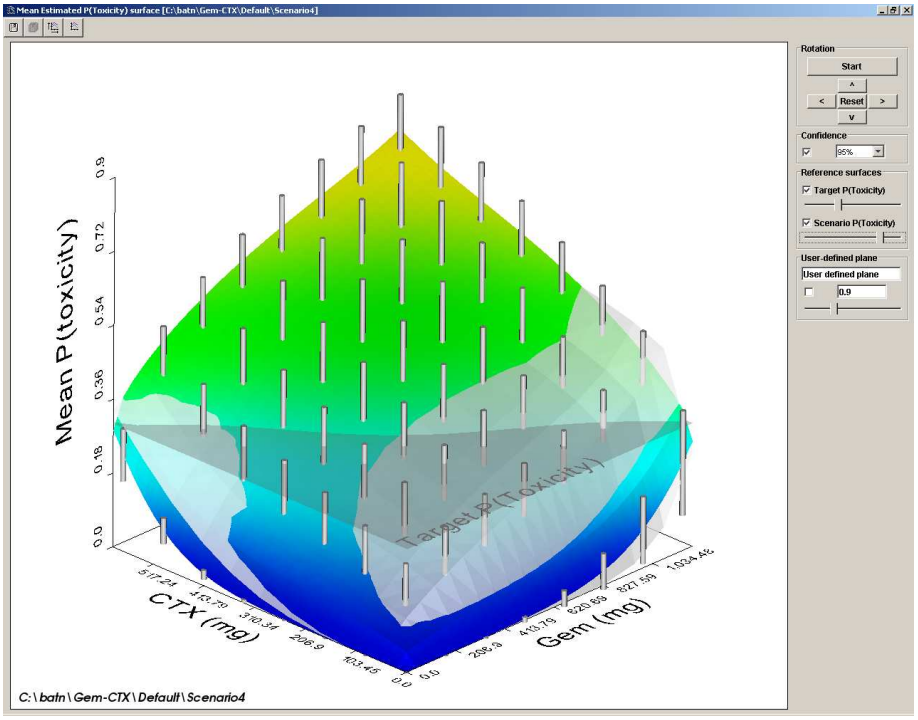

or as a 2D contour plot:

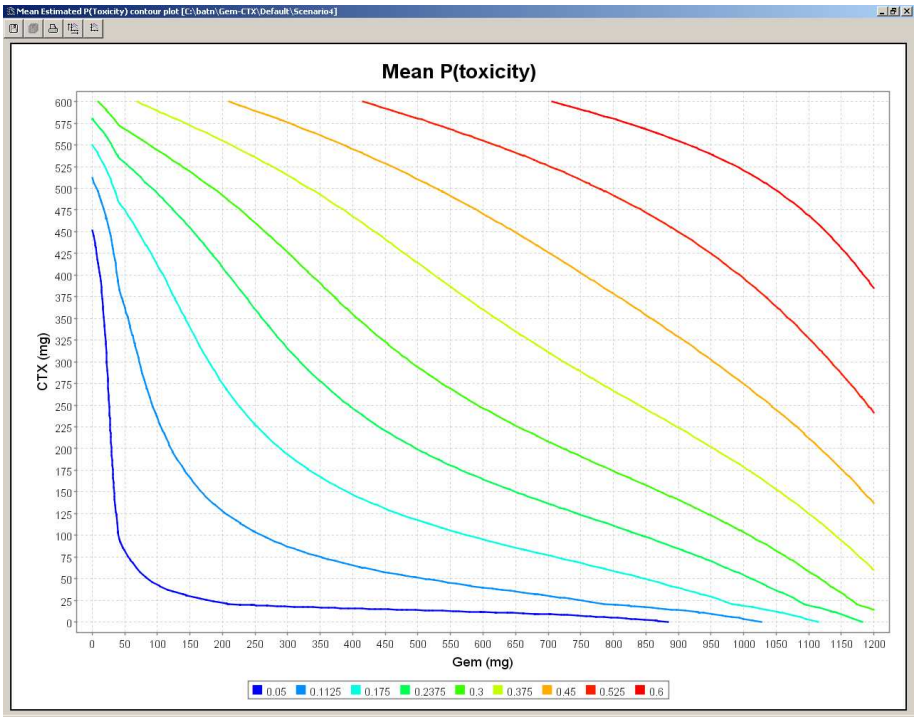

7.3 Charts for individual simulations

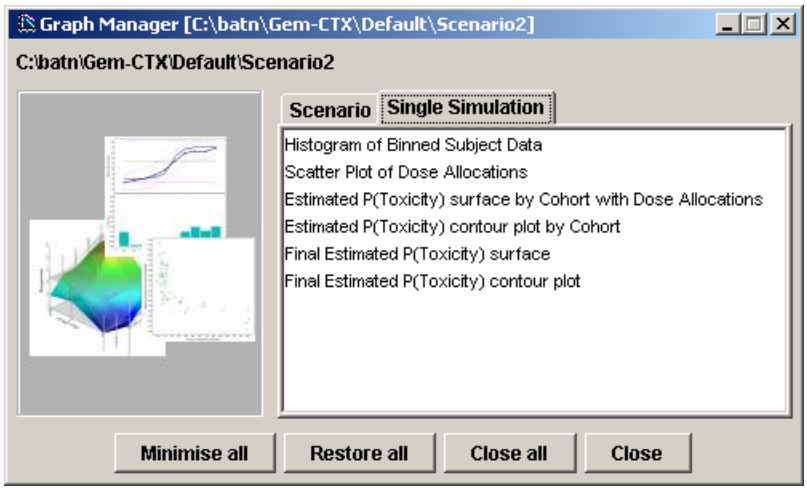

7.3.1 Histogram of Binned Subject Data

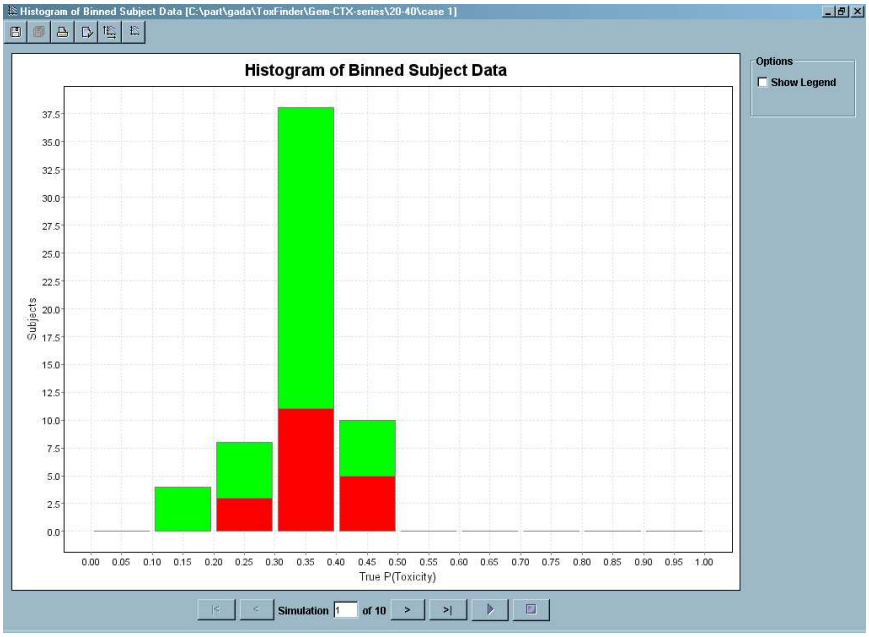

Displays the histogram plot of the number of patients allocated to the different levels of toxicity, and the number of toxic events observed, for each simulation.

7.3.2 Scatter Plot of Dose Allocation

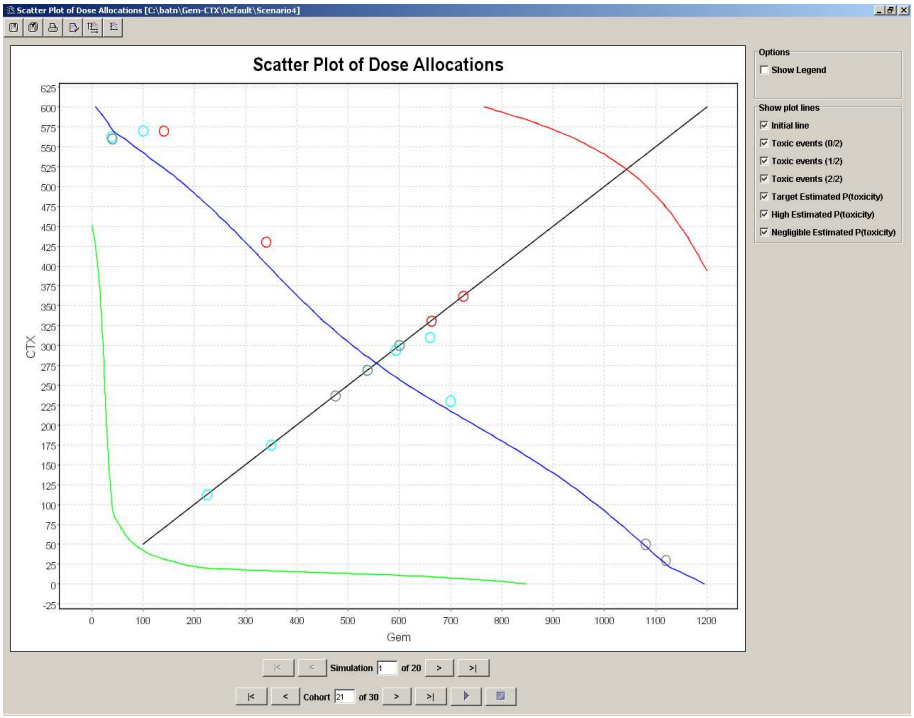

The scatter plot shows each cohort allocated in a simulation and the number of toxic events for each cohort. The circles showing the cohorts are coloured from blue to red indicating the proportion of the cohort experiencing toxicity.

Circles for cohorts with exactly the same dose combination are dithered to make them more visible.

The current estimated negligible, target and high toxicity contours are also displayed.

7.3.3 Estimated P(Toxicity) Surface by Cohort

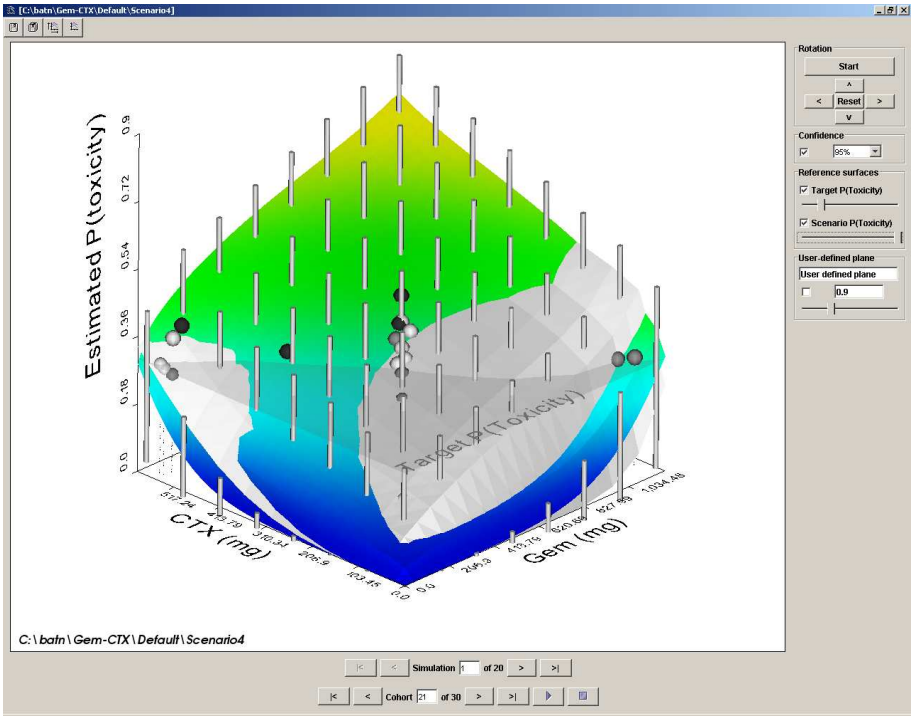

The toxicity surface by cohort graph shows the scatter plot of cohorts (each grey sphere shows a cohort) superimposed on a 3D image of the fitted surface (coloured), the surface being simulated from (grey surface) and the plane at the target level of toxicity. The sphere's showing the cohorts dose combinations are coloured from white to dark grey indicating the proportion of the cohort experiencing toxicity. **The cohort sphere's position on the toxicity axis does not correspond to their toxicity**. It was felt that this would make the sphere's position on the drug axes two difficult to determine and hence colour coding is used instead.

Spheres for cohorts with exactly the same dose combination are dithered to make them more visible.

The estimated mean toxicity surface through the simulated trial can also be viewed as a contour plot:

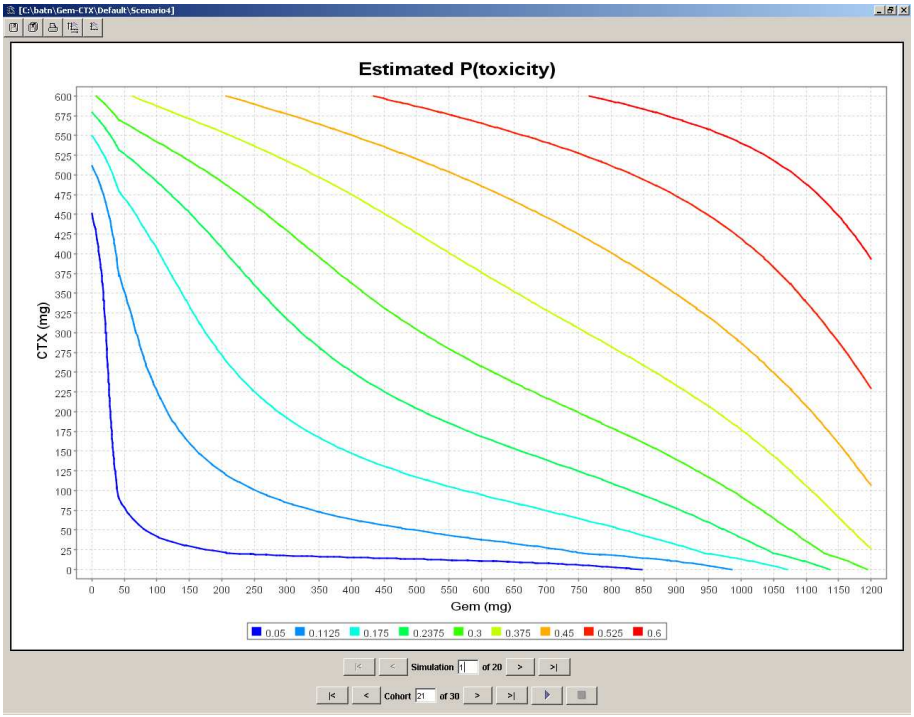

7.3.4 Final Estimated Toxicity Surface and Contour Plot

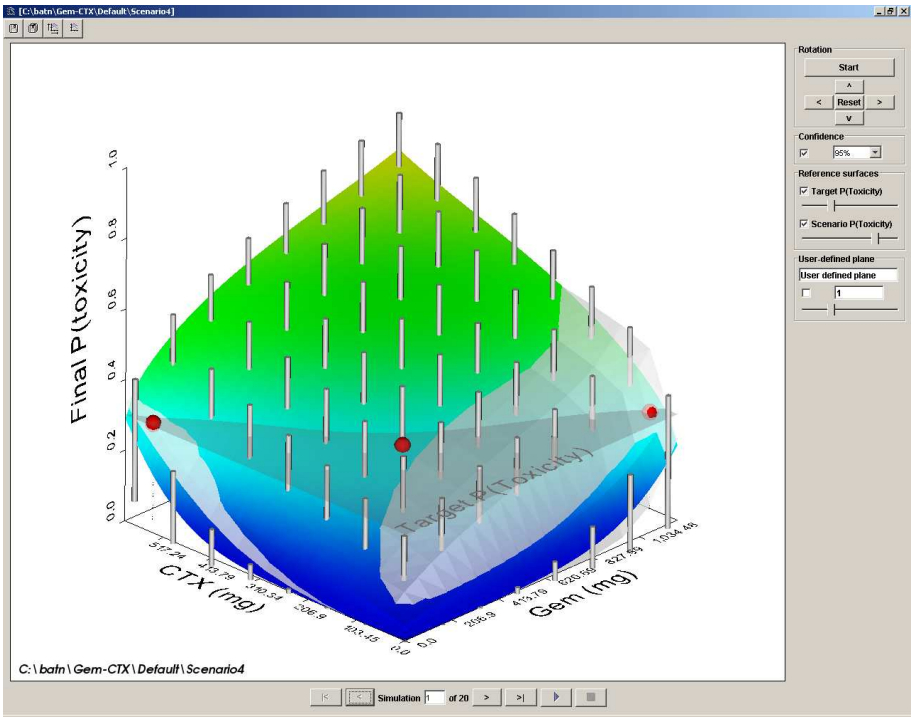

The Toxicity surface graph shows the final toxicity surface (coloured) and also the three selected dose combinations (‘left’, ‘right’ and ’middle’ in [TMML]).

The final estimated mean toxicity of a simulation can also be displayed as a contour plot:

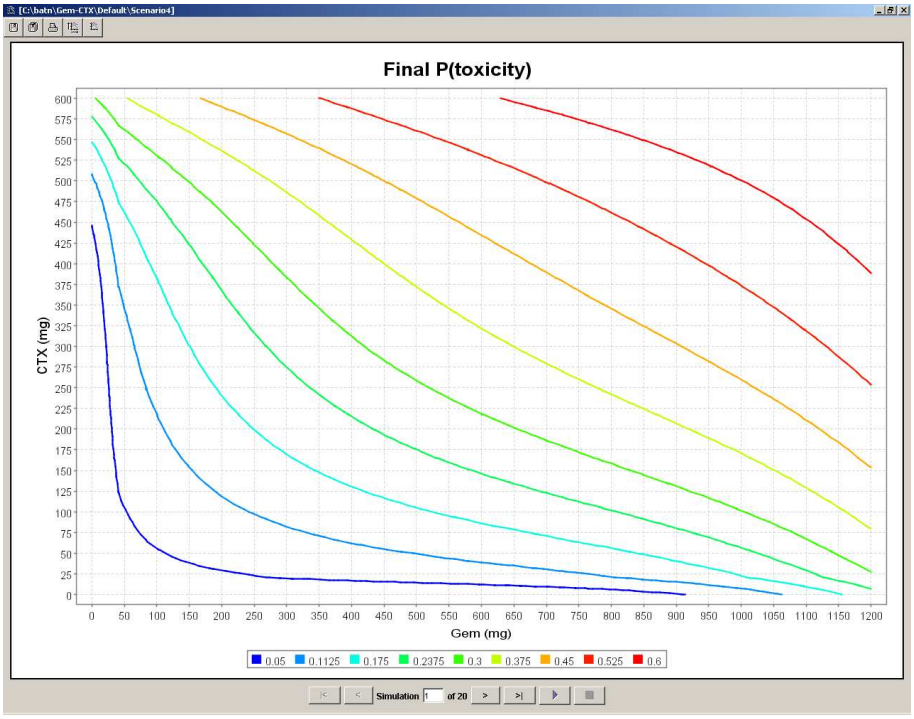

8 Instructions for evaluating the model implementation

8.1 Dose allocation rules:

So the different options can be compared there is an option on the statistics page to switch between them, this is how the 2 options work:

Fisher Information & CKP	Each half of the L2 contour is searched for the point which maximises cancer killing potential (CKP) and the point which maximises Fisher Information (log(det(Hessian))) Cohorts are then allocated to a dose halfway between these two. This is the default and should be how the paper was intended to work.
Fisher Information only	The cohort is allocated to the point on L2 which maximises Fisher Information. This is provided to make it easier to see what the Fisher information is doing.

8.2 Debug output:

In order to see in detail what the implementation is doing (and help determine that the implementations are indeed correct), diagnostic output from the model is available. This is controlled by the parameter "verbose".

This parameter is not accessible through the user interface but can be manually edited in the 'init.tmml' file – e.g. "verbose 1"

"verbose" takes a single integer value:

0 = No debug output - the default

1 = Writes debug.csv - see below

2 = Write information on how the chosen metric (FI, etc) varies along L2 to the file 'out.txt'

3 = Writes fisher.csv which dumps a grid of log(det(FI)) after each pair of cohorts

debug.csv contains pictorial representations of the grid used to find L2. Each cell has the format A;BC where:

A = Number of toxicities (only if B>0)

B = Number of patients (only if >0)

C = A letter code.

L and R are points along the L2 contour

D is where L2 crosses L1,

K is the maximum of CKP

F is the maximum FI (or random, or "farness" depending on the allocation method)

A is where the dose is allocated

Letters overwrite each other, so if K and F are at the same dose, both will be hidden by A.