CRMSimulator User's Guide

The CRMSimulator is an easy-to-use implementation of the CRM dose-finding method. This document focuses on how to use the software, not the statistical method itself. A separate document gives some guidance for using the statistical method.

Text in **bold face** corresponds to text taken directly from the user interface.

The CRM Simulator has two tabs: one for specifying model parameters and one for running simulations.

Model parameters

Here is an example of the **Model parameter** tab:

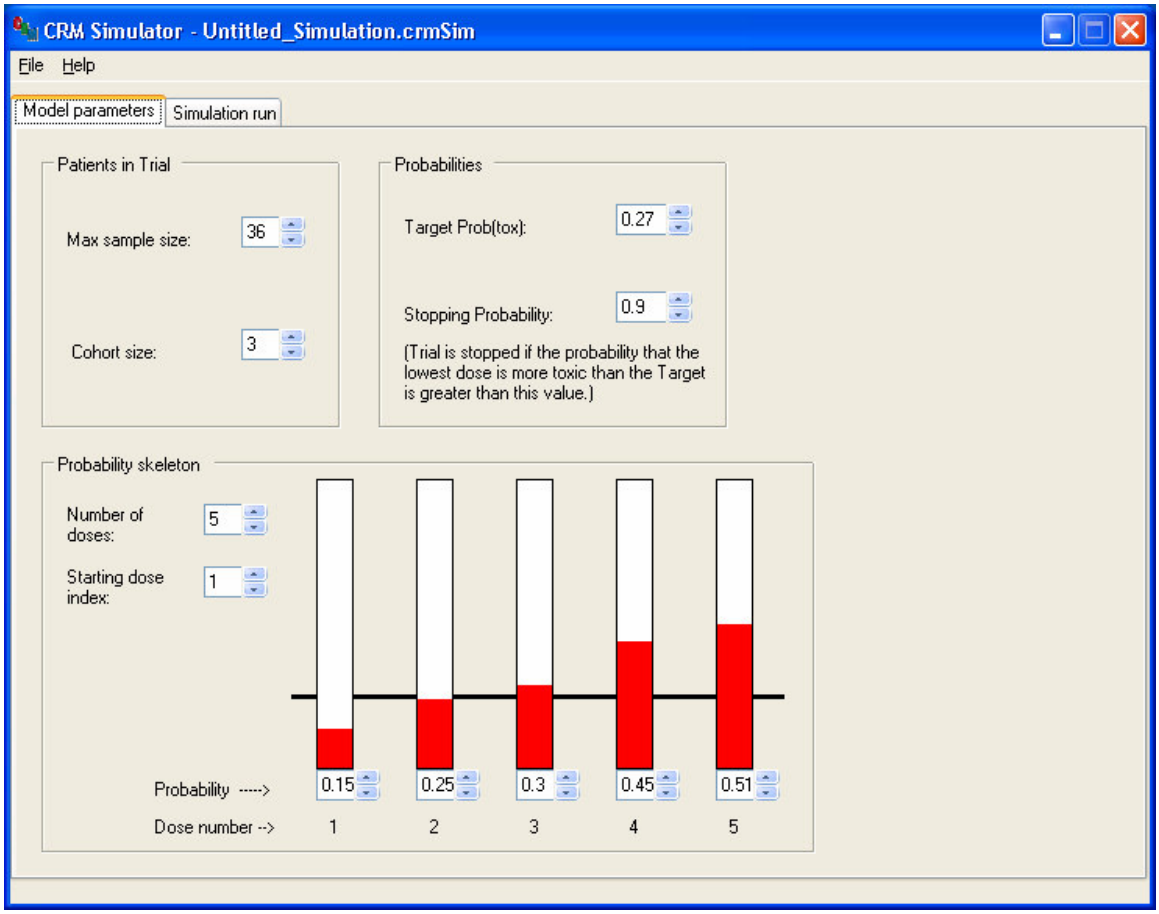

All numbers are entered into special edit boxes. One may simply type numbers into these boxes, or one may scroll through values by using the arrow keys or by clicking on the small up and down arrows beside the edit box.

The maximum number of patients in the trial is specified in the **Maximum sample size** field. The maximum sample size must be an integer multiple of the **Cohort size** value.

The cohort size may be 1, 2, 3, or 4.

The CRM method tries to find the dose with posterior probability of toxicity closest to a specified target probability of toxicity. This target is labeled **Target Prob(tox)**. This value is graphically represented as a black horizontal bar across the probability skeleton values.

The trial will stop if the posterior probability that the lowest dose is more toxic that the target is sufficiently high. This threshold value is labeled **Stopping Probability**. Typically this value is around 0.9 or higher. If it is much lower, the method will stop frequently. If it is too high, the method will not stop when it should. One may set the stopping probability as high as 1, effectively turning off the stopping rule. Of course this would typically be unethical in an actual clinical trial.

The probability skeleton is graphically represented by "thermometer" controls. There is a thermometer for each dose: these controls appear and disappear as one changes the **Number of doses** value. The height of each red bar corresponds to a skeleton value, the prior mean probability of toxicity at that dose.

The maximum number of doses is eight.

(continued on next page)

Simulation run

Here is an example of the **Simulation run** tab.

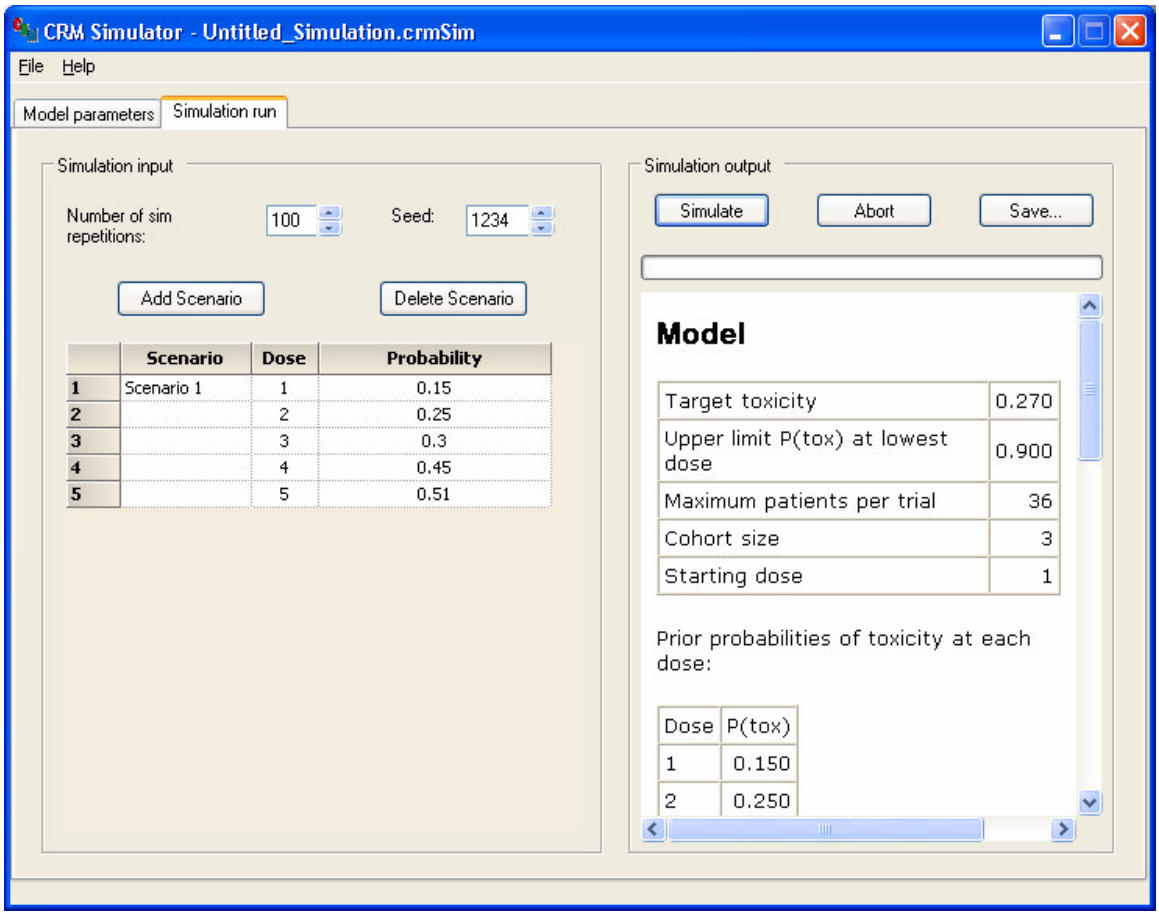

One may add a scenario by clicking on the **Add Scenario** button. One may delete a scenario by clicking on some row of that scenario then clicking the **Delete Scenario** button.

Number of sim repetitions denotes the number of times each scenario will be run. One can typically get a good idea of the characteristics of a design from running 100 repetitions while experimenting with parameters. Such runs will complete quickly. Once a design has settled, people often report the operating characteristics based on 1000 repetitions.

Seed is the random number generator seed. Results will be reproducible if the same seed is used.

To run a set of simulations, click the **Simulate** button**.** Simulations may be stopped by clicking the **Abort** button. The results of a simulation may be saved by clicking the **Save** button.

Simulation results appear in the pane to the right. Note that one may grab the lower right corner of the application and drag it to make the application larger, exposing more of the output at one time.

Here is a sample output of a simulation run.

Model

Target toxicity	0.270
Upper limit P(tox) at lowest dose	0.900
Maximum patients per trial	36
Cohort size	3
Starting dose	1

Prior probabilities of toxicity at each dose:

Dose	P(tox)
1	0.150
2	0.250
3	0.300
4	0.450

Simulation Settings

Number of repetitions	100
Random number generator seed	12345

Scenarios and Results

Scenario 1

Dose	True P(tox)	P(selection)	# Subjects treated
Stop	NA	0.00	NA
1	0.150	0.13	10.62
2	0.250	0.49	14.52
3	0.300	0.34	8.91
4	0.450	0.03	1.62
Toxicities per trial: 8.00			

First the model parameters are repeated so that it is clear which design the simulation results are associated with.

The **P(selection)** column of the output reports the selection probability of each dose. The **# Subjects treats** column reports on average how many patients were given each dose in each trial. The **Toxicities per trial** field at the bottom of the table gives the average number of toxicities each time the trial was run under the given scenario.

The selection probability for **Stop** reports the probability of the method stopping a trial early under the given scenario. The **True P(tox)** and **Times given** columns are marked **NA** (not applicable) because they do not apply to the stopping decision.

(continued on next page)

Saving and opening files

Program input (model parameters and scenarios) and output (simulation results) are saved in a single HTML file from the File -> Save menu.

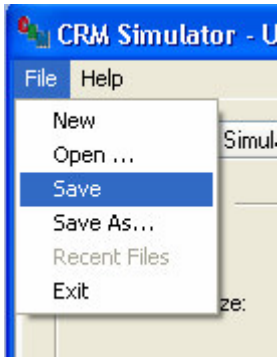

This HTML file can be viewed and printed from any web browser.

If this file is opened via the File -> Open menu, the parameters are re-imported and the simulation results appear in the output window.

Validation and errors

If a validation rule is violated, the software displays a flashing red dot with an exclamation point inside. For example, the maximum cohort size for this application is four. The following shows the result of entering a cohort size of five:

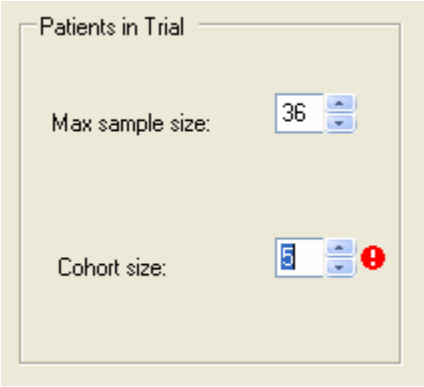

Hovering the mouse over the red dot pops up a small box explaining the reason for the warning.