

ARM3 Release 4 Improvements

Forward by Ed Downs

We are pleased to release our latest ARM3 components for our Nirvana Club Members. The Genetic Algorithm and Neural Network components for Strategy Builder form the heart and soul of our valuable Artificial Intelligence technology, built over 10 years from thousands of experiments and many thousands of lines of software. While we know these tools appear somewhat complex on the surface, we encourage all members to attempt to use them to at least train on their own data.

The great thing about them is they are completely integrated into OmniTrader, making it possible for all members to experiment and further the state of the art with us. We are certain that the pruning capabilities introduced into both components this year will go a long way to helping all of us – Nirvana Research Staff and Members – build better solutions and in the end, make more money in our trading. Good luck! - Ed Downs

Neural Network – General Improvements

Copy/Paste of Inputs

A copy/paste feature was implemented within the AI components to ease the process of configuring new blocks.

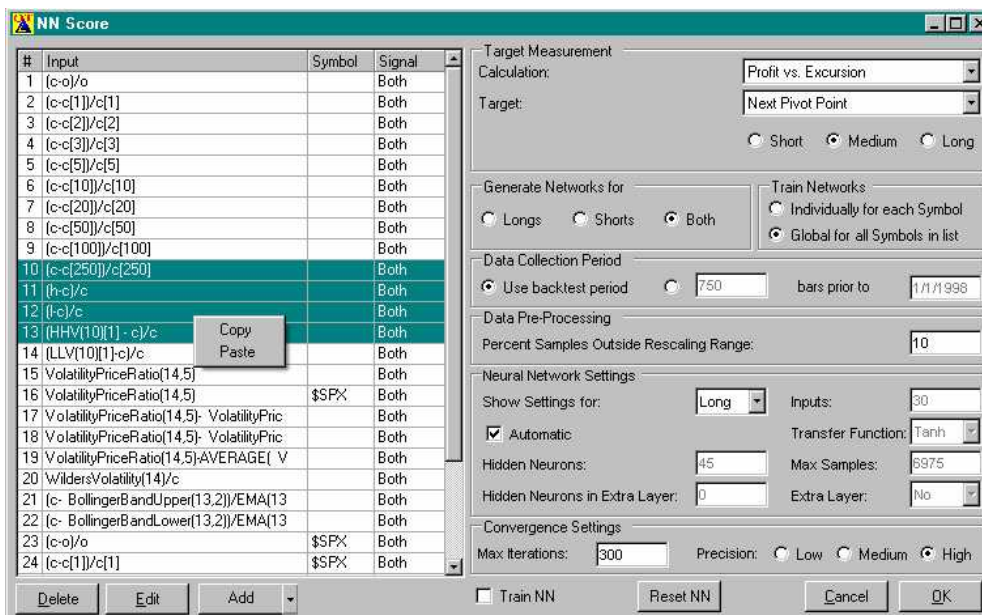


Figure 1. New features in the *NN Score* Settings form.

Within the *NN Score* Settings interface, select one or multiple inputs from the *Input* table and right-click to copy the highlighted inputs, which can then be pasted into any *NN Score* block across different strategies. Inputs can also be pasted in to any *GA Signals* block as genes, and vice versa (i.e. copied genes can be pasted as inputs).

Scaling Range

The user can now also specify the percentage of samples that falls outside the rescaling range, a variable used while pre-processing the collected data prior to training. To optimize the training process, the majority of all recorded inputs are rescaled into normal values (0 to 1), while the remaining samples are rescaled to fall outside this range (below 0 and above 1). This variable is now exposed in the *Data Pre -Processing* section of the UI. In the example of Figure 1 - with the value set to 10 – 80% of collected samples will be rescaled to values within the rescaling range, while 10% will take up values above and below the range, respectively.

Transfer Function Selection

An additional feature added to the Neural Network is the ability to choose the transfer function. All previous networks were training with hyperbolic tangent functions. A sigmoid function is now also available.

After training (while assigning scores to signals during analysis), the user can use the *Sample Validation* settings (Figure 2) to choose to ignore samples for which any of the input measurements are within a certain percentage of the minimum and maximum values recorded during data collection. These signals are assigned a score value of zero, independently of the network's actual output value. To assign a score to all signals, leave this setting set to 0.

Input Pruning for a Neural Network

After a Neural Network has been trained, it is now possible for us to gain insight into WHICH inputs are really contributing to the *NN Score* that is generated. This is done with the new Post-Training Analysis Tool. Figure 2 shows the main user interface for a *NN Score* block, that has already been trained.



Figure 2. *NN Score* main dialog box.

To access the NN Post-Training Analysis tool, click on the *Analyze* button. For a quick view of how each input affects the network’s output value, simply click on the *Analyze* button (Figure 3). The data recorded during data collection is used to calculate the average network’s output for each value a certain input. The *Network Range* plot in the lower right will display a histogram representing the minimum and maximum average score for each input.

This analysis must be run separately for the Long and Short network. The user can also choose to cut down on the *Number of Samples* used for this analysis. When several thousands of bars were collected prior to training, this analysis may take an unnecessary long time to complete. By using only a random portion of the collected data (1000, in Figure 3), this analysis will yield nearly equal results while saving a considerable amount of time.

When analysis on all inputs is complete, the user can automatically view which inputs never contribute to high or low score values, or those that result in a small score range (i.e. maximum score – minimum score), by using the scroll bars in the *Input Pruning* section. As any of the three settings are changed, the inputs on the left are grayed out if they don’t satisfy the constraints on the values for minimum score, maximum score and score range.

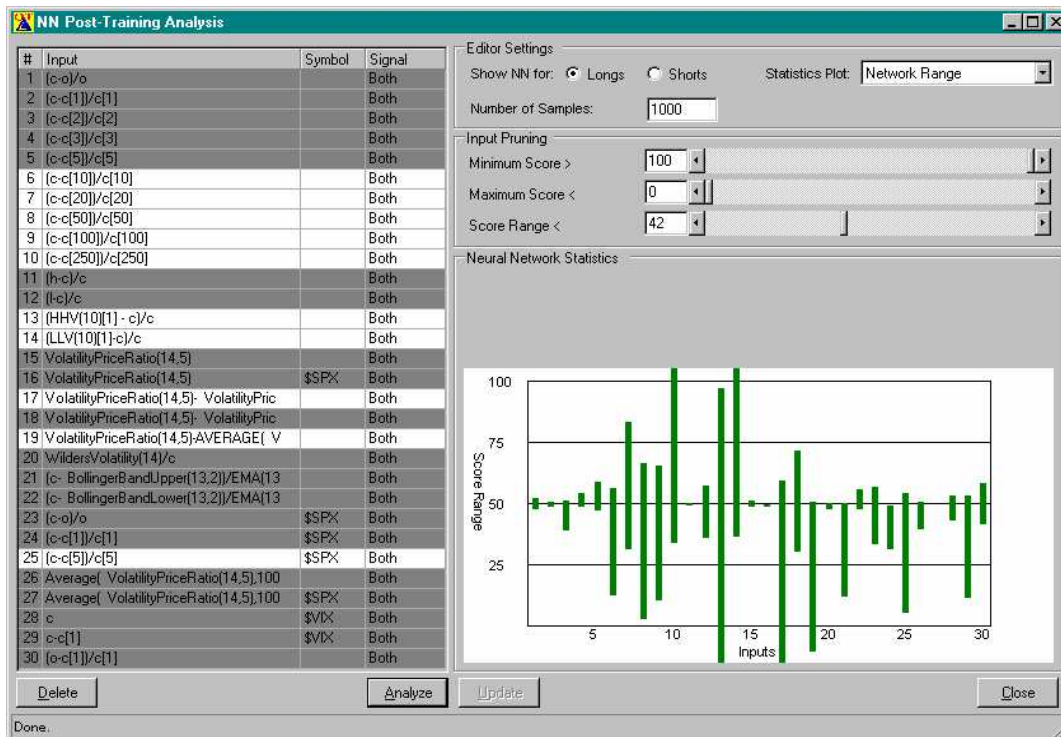


Figure 3. *NN Post-Training Analysis* UI.

A **sensitivity plot** for a specific input can also be shown by highlighting the desired input and switching the *Statistics Plot* setting to *Input Sensitivity* . A graph of the average score for each value of the selected input between its recorded minimum and maximum values is displayed (refer to Figure 4 for an example).

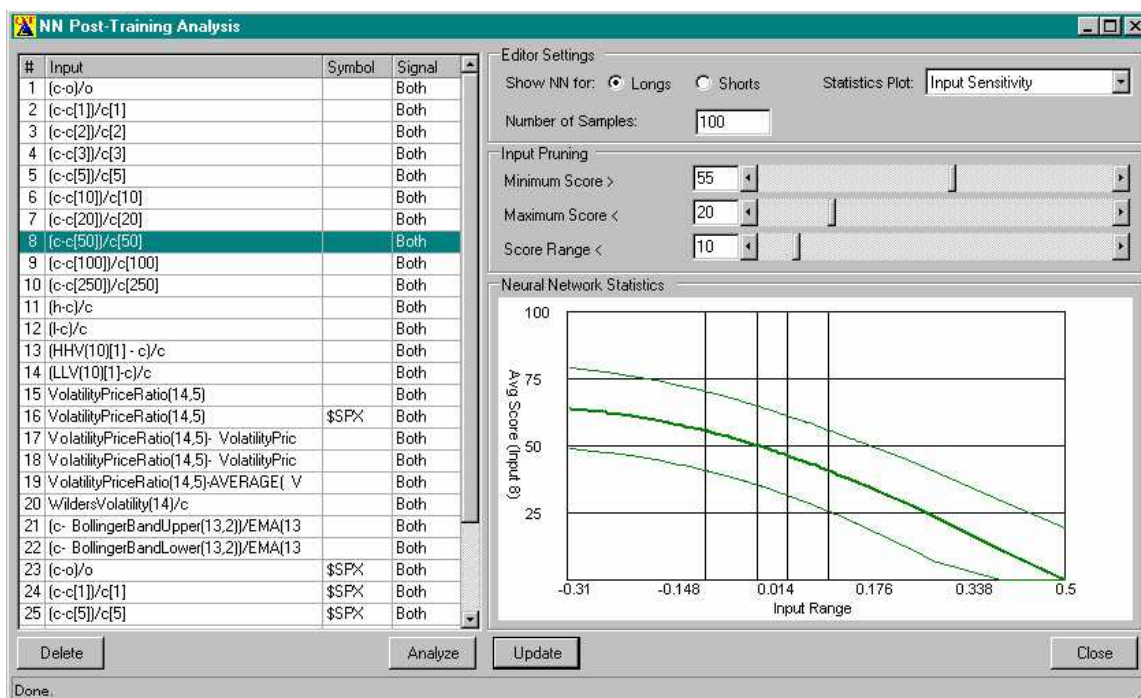


Figure 4. Example of *Input Sensitivity* plot.

The thinner lines around the (thicker) average score represent a standard deviation measurement: the distance between these lines is one standard deviation wide. This graph also features a vertical grid that provides information on the distribution of recorded input values within its minimum and maximum; the vertical grid lines that separate the graph in 5 vertical slices are spaced in a way that results in each slice containing the same amount (20%) of the collected samples used for NN Post-Training Analysis. To show the sensitivity plot for a different input, highlight the desired input and click on *Update*.

Hitting the *Delete* button will remove the grayed-out inputs and reset the network. The user can then close this window and return to the NN Settings interface to add more inputs, if desired, and retrain the network.

What are we looking for?

When reviewing the Sensitivity Plot, we are looking for an upward or downward-sloping curve related to an input. That is, the left side (or right side) is lower (or higher) than the other. This indicates that, as the input changes, the scores in the Network change and thus, the input is being used. Inputs which have a flat curve are probably not contributing much to the overall effectiveness of the Network.

OK. I found my inputs that appear to be ineffective. Now what?

First, you want to save the Strategy that has the Neural Network so you can get back to it, because the step of deleting inputs is irreversible. The Neural Network **MUST BE RETRAINED** after you delete any input. After making a copy, bring up the Strategy and delete the inputs that seem to be ineffective and retrain the Strategy. You should see a negligible effect in performance in your profile.

Now that you have cleared out some inputs that are not useful, you can add more! Think about which measurements or systems might be predictive, create an input based on them (measurement or signal from a system) and retrain. We have discovered improved results almost immediately from executing this process. It is our hope that members will do the same thing, and share their discoveries with the group on the Nirvana Club web site.

GA Signals Knowledge Base Editor

The GA Knowledge Base Editor is similar to the NN Post-Analysis tool, in that it is possible to gauge the effectiveness of various inputs. However, the process is different because the GA is based on RULES. It is possible to delete inputs (called “genes” in the GA) and retain all rules that do not use the given genes. This makes for a very robust editing capability, and allows us to build and improve our Knowledge Bases over time by sequential and selective input reduction.

This illustration shows the GA Signals Knowledge Base Editor. Click the **Edit KBs** button on the main GA interface to launch this editor.

The screenshot displays the GA Signals Knowledge Base Editor interface. It features several configuration panels and a data visualization section.

Knowledge Base Lookup Panel:

- Minimum Requirements:
 - Basic | Advanced
 - Profit per Trade: 5
 - Hit Rate: 80
 - APR: [Empty]
 - Hits: 10
 - Max Std. Deviation: 0.05

Editor Settings Panel:

- Show KB for: Longs Shorts
- Statistics Plot: Gene Usage/Fitness

Gene Pruning Panel:

- Remove: Selected Gene All Genes Below
- 10 % Usage
- % Fitness

Knowledge Base Statistics Table:

Statistics	Current KB's	Affected Rules	Remaining Rules
Number of Rules	179	135	44
Average Fitness (%)	0.01154	0.01183	0.01067
Average Hits	12.49	12.53	12.39
Average Standard Deviation (%)	0.01782	0.01924	0.01347
Average Profit Per Trade (%)	6.58	6.62	6.44
Average APR (%)	301.83	303.37	297.13
Average Hit Rate (%)	86.69	85.57	90.12
Average Signal To Noise Ratio	0.01154	0.01183	0.01067
Average Profit vs. Excursion	4.8	4.82	4.76
Average Signal vs. Wilders Risk	0.0125	0.01283	0.01149
Average Bars In Trade	7.72	7.69	7.81

Gene List Table:

#	Gene	Symbol	Signal
1	ADX-B(14, 25)		Both
2	BND-C(13, 6)		Both
3	BOL-C(13, 2)		Both
4	BOL-T(14, 20, 30, 2.4)		Both
5	CCI-C(13)		Both
6	CCI-D(13, 0, 0)		Both
7	CCI-FP(80)		Both
8	CCI-P(13)		Both
9	CHA-D(0, 0)		Both
10	CHA-P(2000)		Both
11	DMI-C(14)		Both
12	KBA-C(13, 2.5, 7)		Both
13	MAC-D(12, 26, 0, 0)		Both
14	MAC-M(9, 12, 26)		Both
15	MFR-B(13)		Both
16	MFR-C(13, 20, 80)		Both
17	MFR-D(13, 0, 0)		Both
18	MDM-P(13)		Both
19	MV2-C(9, 50)		Both
20	WLR-P(13, 20, 80)		Both
21	WLR-D(13, 0, 0)		Both
22	WLR-C(13, 20, 80)		Both
23	WLR-P(13, 20, 80)		Both

Gene Usage and Fitness Chart:

The chart displays Usage (left Y-axis, 0% to 28%) and Fitness (right Y-axis, 0.01 to 0.04) for 40 genes. The X-axis represents the Gene number (1 to 40). Usage is shown as green bars, and Fitness is shown as a line graph.

GA Signals Knowledge Base Editor

Any changes you make in the Knowledge Base Editor dialog box will also affect the settings in the Main GA Signal dialog box. All of your changes to this dialog box are permanent. Once you make a change in the Knowledge Base Editor and click the Close button, your changes are saved.

In addition to the Canned Settings and Knowledge Base Lookup sections explained in the Main GA Signal interface section, the GA Signals Knowledge Base Editor has the following options:

Editor Settings

The Editor Settings section specifies which statistics to display in the editor window. You can modify the following options to view a variety of information available in the editor.

- **Show KB for Longs or Short**—Select Longs to display the Knowledge Base statistics for Longs. Select Shorts to display the Knowledge Base statistics for Shorts.
- **Statistics Plot**—Specifies which statistics plot to display in the Knowledge Base Statistics graph. You can display the statistics for gene usage/fitness as well as the number of iterations for each gene or for the bin distribution for each gene in the Knowledge Base.

Gene Pruning

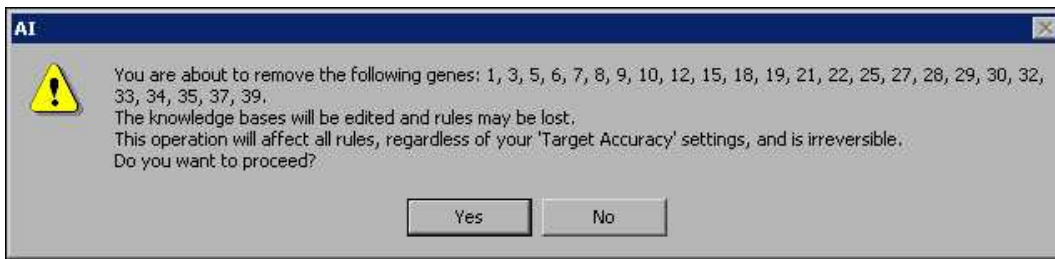
The gene pruning section allows you to remove a specific gene or to specify criteria that all genes must meet in order to be included in the Knowledge Base. For example, this illustration shows that all genes below a 10% usage will be removed from the Knowledge Base.



Gene Pruning Options in GA Knowledge Base Editor

The % Usage option specifies the percentage of rules that must have a value other than Don't Care. The % Fitness option specifies the percentage of rules that must meet the specified Fitness value.

To see how the values you specify in the Gene Pruning section effect the Knowledge Base, click the Delete button in the Knowledge Base Editor. For example, if you set the % Usage value at 10 and click the Delete button, you will see the following AI dialog box appear.



Artificial Intelligence Gene Deletion Dialog Box

Knowledge Base Statistics

The Knowledge Base Statistics section displays the average results of the settings you have specified for Profit per Trade, Hit Rate, APR, Hits, and Maximum Standard Deviation. For more information about setting these requirements, refer to the *Error! Reference source not found.* section in this chapter.

The Knowledge Base Statistics section gives you the statistics for all current Knowledge Bases and gives you the statistics on the number of rules affected and remaining after your criteria is applied. The Update button below the grid will refresh the statistics in the grid after you have made changes to the settings.

Below the Knowledge Base Statistics grid is a graph showing the statistics plot you specified in the Editor Settings section. If you are viewing the gene usage and fitness results, you will notice that the fitness results are displayed in black and the usage results are displayed in green. The fitness value is the average fitness of all the rules that use the gene.

Gene List

The Gene List is a list of the genes currently in use. You can use the options below the Gene List to Delete, Add, and Disable genes used in the Knowledge Base. When you delete a gene, it is permanently removed from the Knowledge Base. When you click the Add button, you can choose to add a System, Measurement, or Boolean gene.

The Disable Gene option turns the gene off for signal generation but does not remove the gene from the Knowledge Base. The Disable Gene option is a good tool for testing. If you want to edit the genes and/or parameters, click the Settings button on the Main GA interface, modify the options on the GA Signals dialog box, and retrain the GA.

When you click the Add button to add a gene, you are given the following three options to choose from:

- **Add System**—Opens the Add System dialog box where you can select any of the existing systems or system measurements, as well as edit parameters, tolerance value, and applicable symbol. You can use any of the systems included with OmniTrader or you can use a system that you have created using OmniTrader Professional. When you add a system, you will see the Add System dialog box.

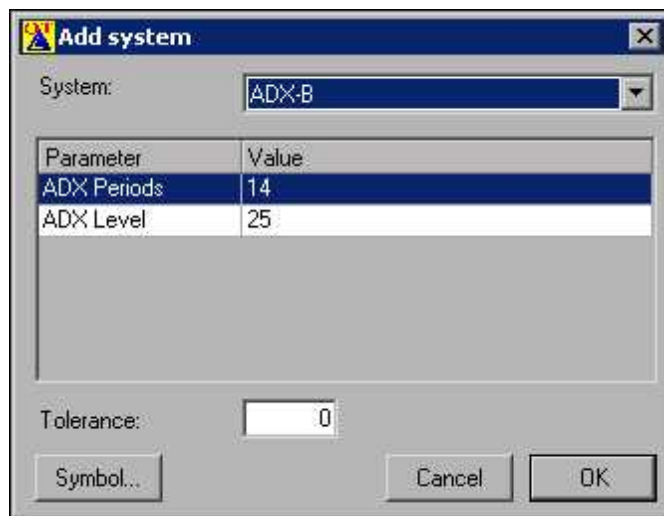


Figure 76. Add System Dialog Box

From this dialog box, you can set the parameters of the system, set the tolerance, and specify to run the system on a particular symbol, index, sector, group, or subgroup.

- **Add Measurement**—Opens the Formula Builder where you can create a formula that returns a continuous measurement. For more information, refer to the *Using the Formula Builder* section of Chapter 7, *Programmable Filters* in the *OmniTrader User Guide* .
- **Add Boolean**—Opens the Formula Builder dialog box where you can create a formula that returns a Boolean value.