NWP SAF

NWPSAF 1D-Var
GUI User Guide

Version 1.0

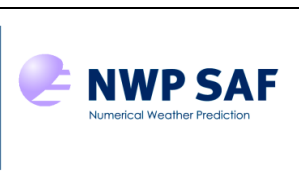22[nd] August 2014

NWP SAF 1D-Var GUI User Guide

This documentation was developed within the context of the EUMETSAT Satellite Application Facility on Numerical Weather Prediction (NWP SAF), under the Cooperation Agreement dated 29 June 2011, between EUMETSAT and the Met Office, UK, by one or more partners within the NWP SAF. The partners in the NWP SAF are the Met Office, ECMWF, KNMI and Météo France.

Change record			
Version	Date	Author/changed by	Remarks
0.1	21/02/14	P. Weston	First version, corresponds to NWPSAF 1D-Var v1.0 beta
0.2	21/05/14	P. Weston	Minor changes from beta test feedback
0.3	29/05/14	P. Weston	Updated with new header, new document ID, new copyright statement and changed font following comments from B. Conway
1.0	22/08/14	P. Weston	Version valid for NWPSAF 1D-Var v1.0

Table of Contents

1 INTRODUCTION

This document describes how to set up and run the NWPSAF 1D-Var graphical user interface (GUI). The GUI allows users to compile and run the NWPSAF 1D-Var code as well as plotting outputs from the retrievals. There are options within the GUI for users to customise their runs and various different plotting options to display a range of outputs.

It should be noted that the NWPSAF 1D-Var can be run from the command line without the need to use the GUI. The GUI represents extra functionality and some users may find it easier to compile and run the code via the GUI rather than the command line.

2 COMPILATION AND SETUP

The script used to run the GUI should perform the required compilation and setup automatically. If for any reason this doesn't work then prior to running the GUI in training mode then there are two routines which need compiling which are scale_oberr.f90 and scale_bgerr.f90. To do this run the make command in the GUI directory. You may need to change the Fortran compiler in the makefile to one which is available on your system (the default is ifort).

To run the 1D-Var relevant RTTOV coefficients for the instrument(s) you wish to run retrievals with must be in the 1D-Var top directory. The easiest way to do this is to link all of the coefficients for any instruments you wish to run with into one single directory. The location of this directory can then be set in the variable RTTOV_coeffs_dir within Control_GUI.sh, Control_test.sh and Control_training.sh or set within the GUI in general mode. If you cannot set up such a directory then you can comment out the lines which remove previous versions and link new versions of the RTTOV coefficients in each of the Control_* scripts and link in your own versions of the RTTOV coefficients for all the instruments you wish to run retrievals for. Other than this everything should run as expected without any further modification to the default code.

3 RUNNING THE GUI

To run the GUI simply run the following command in the GUI directory:

1dvargui

By default this will run the GUI in general mode which is intended for general users who want the option to customise their runs. This is intended to be the most frequently used mode of the GUI. See section 3.1 for a detailed description of how to run the GUI in this mode.

To run the GUI in test mode run the following command in the GUI directory:

1dvargui test

The test mode is intended to be run the first time you attempt to use the NWPSAF 1D-Var. It will test whether you have installed the code correctly and whether you can reproduce the results present in the Sample_Output directory.

To run the GUI in training mode run the following command in the GUI directory:

1dvargui training

The training mode is intended to help users to understand retrieval theory more clearly by using various case studies and extra customisation options. A version of the GUI in training mode will be used in the annual ECMWF/EUMETSAT NWP-SAF satellite data assimilation training course and by including this mode in the package it is intended that users who have attended the course can recreate the case studies, initially covered on the course, for further learning opportunities.

When running in all modes the standard output and errors from the GUI and the 1D-Var code itself will appear in the GUI_logfile file in the GUI directory. If the GUI fails to respond or there is any unexpected behaviour at any point you should look in the GUI_logfile to see what has gone wrong.

The following sections will describe running the GUI in each of these different modes in more detail.

3.1   General mode

After running the command to open the GUI in general mode:

1dvargui

The window in figure 1 should appear.

The initial interface allows several different options to be chosen before running the code. The available options are as follows:

1. Instrument
   Choose which instrument's observational data you wish to use. Available options in the drop-down menu are IASI, AIRS, CrIS, ATOVS, ATMS or SSMIS.

2. Retrieval levels
   Choose the number of model levels you want your retrieval to be valid on. Available options in the drop-down menu are 54, 51 or 43.

3. Truth profile
   If using simulated observations and background profiles then you may have the original truth profiles available to plot in addition to the observation, background and retrieval data. Tick this box if so and type the location of the truth profile in the box which appears as shown by figure 2.

4. Input files
   Rather than using the sample input files you can specify your own input files by checking the relevant tick box and then entering the location of your files in the relevant text box as shown by figures 3 and 4. If only using a subset of your own files then the remaining file locations will default to the sample files corresponding to the chosen instrument and number of levels which are provided with the package. If using all of your own files then the instrument and number of levels choices become obsolete and you can run the retrieval for any instrument and any number of levels.
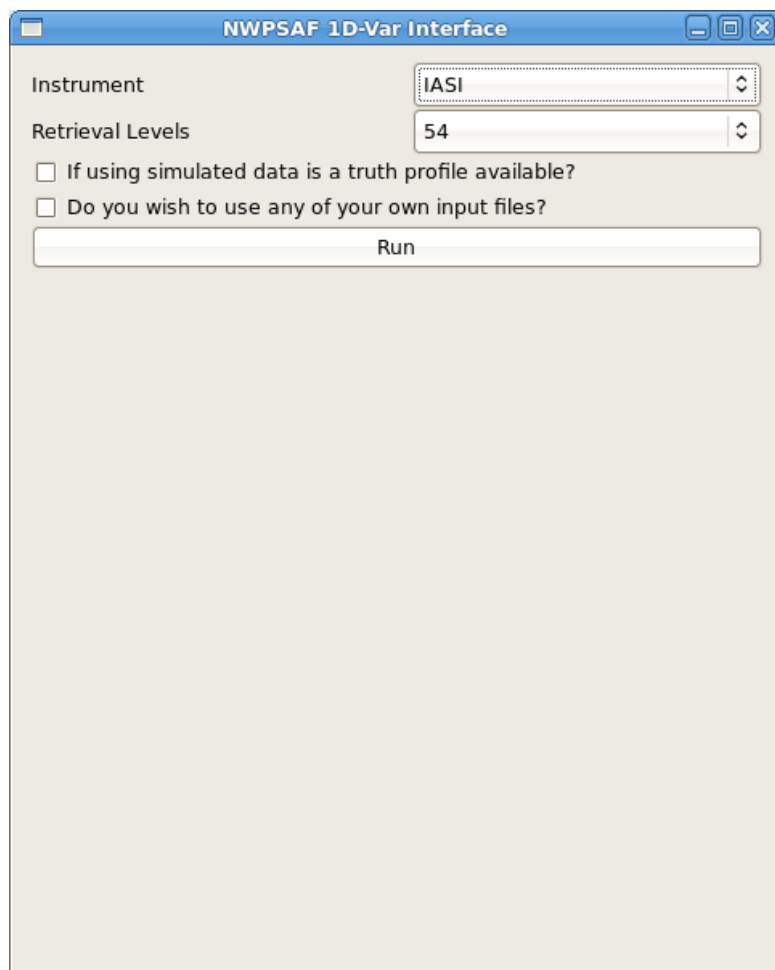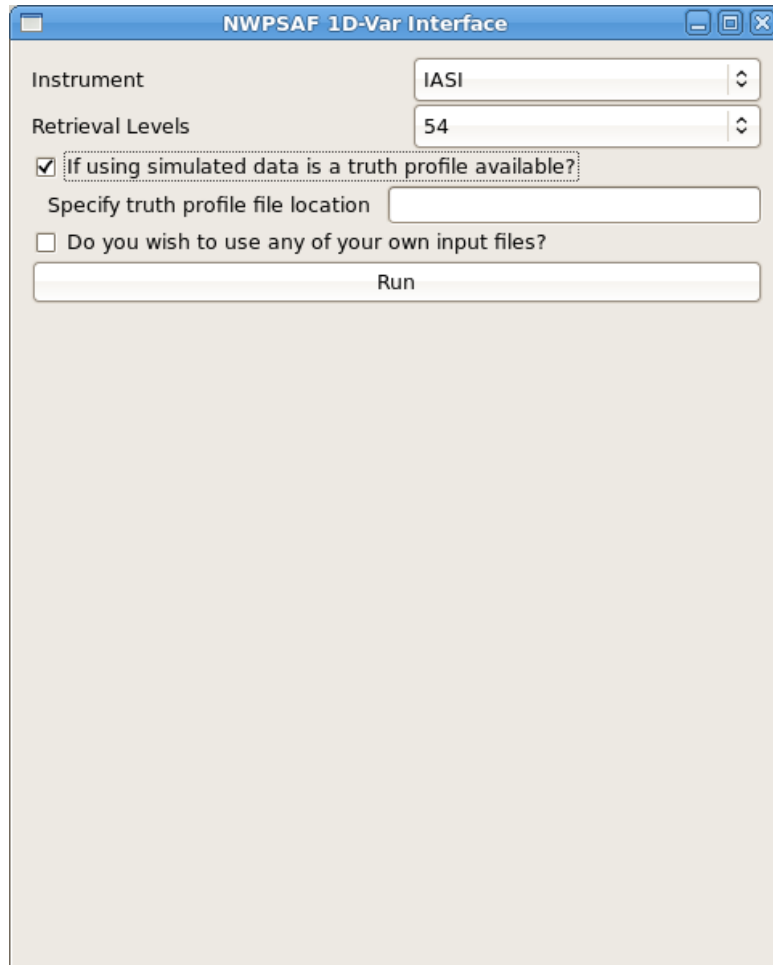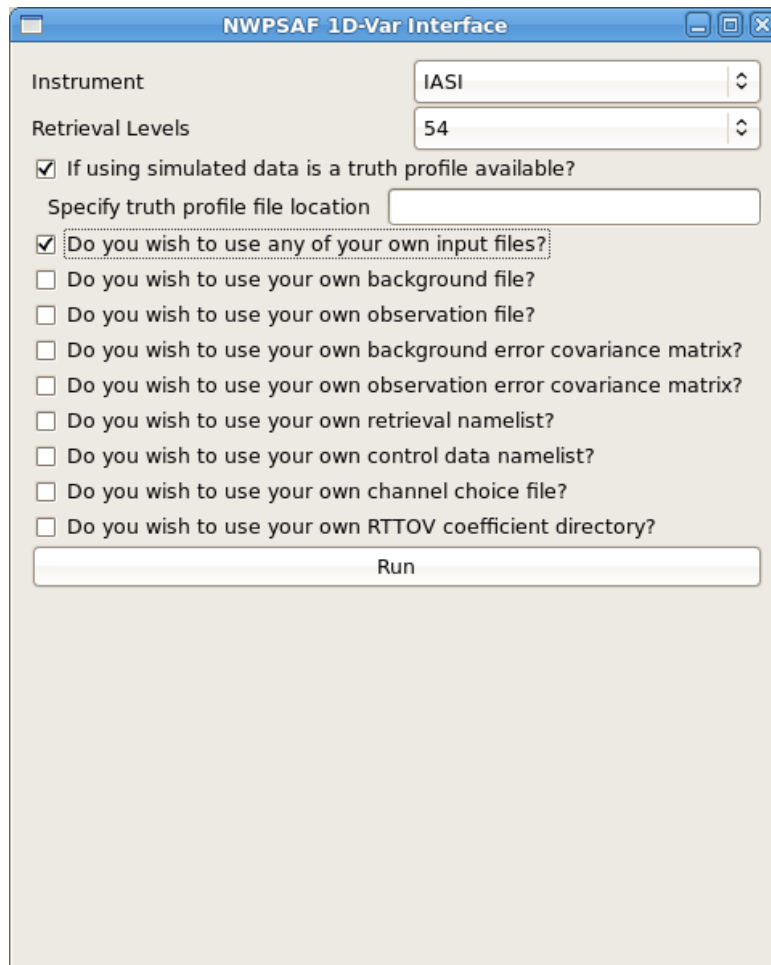
Figure 1: Initial NWPSAF 1D-Var GUI general mode window

Figure 2: If a truth profile is available then tick the relevant box

Figure 3: Ticking the box to specify your own input files opens up many more tick boxes that indicate which of the input files are to be user specified and which are to remain as the default versions
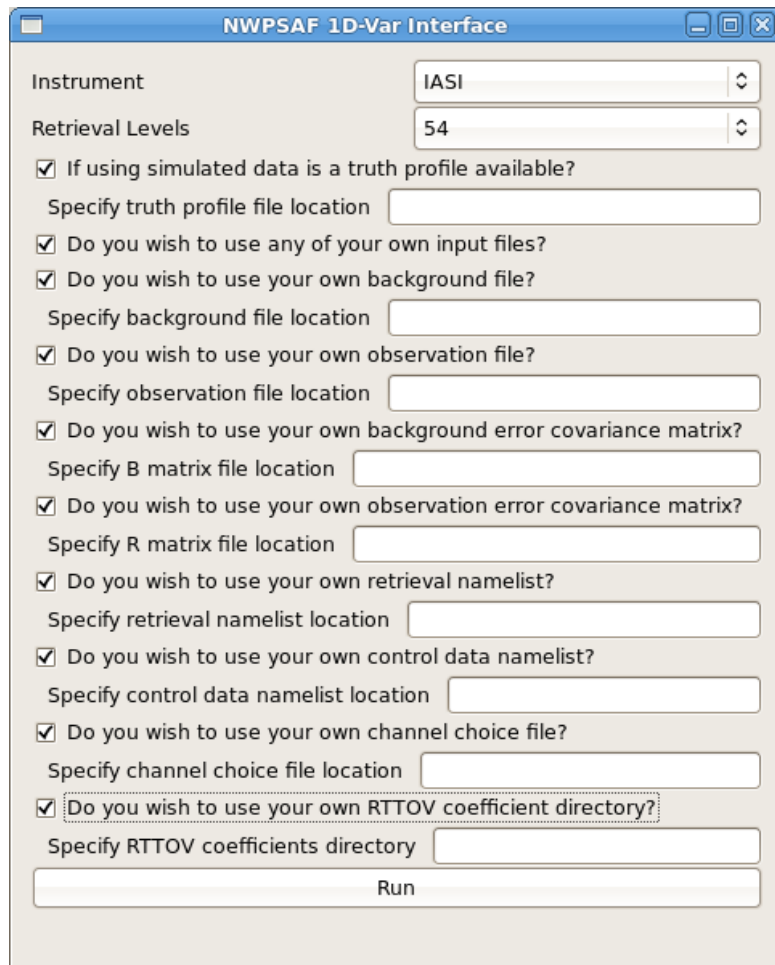
Figure 4: Showing the possible input files and directories that can be specified by the user

Once you have selected the required options for running the 1D-Var code click the "Run" button and the code will compile if not compiled already, your chosen files (either the ones specified or the default ones) will be copied into the required locations, RTTOV coefficients will be linked into the top directory and the code will run. After the code has run successfully another button will appear below the "Run" button labelled "Display plotting window" as shown by figure 5.

A description of how to use the plotting window interface in general mode can be found in section 4.1 below.

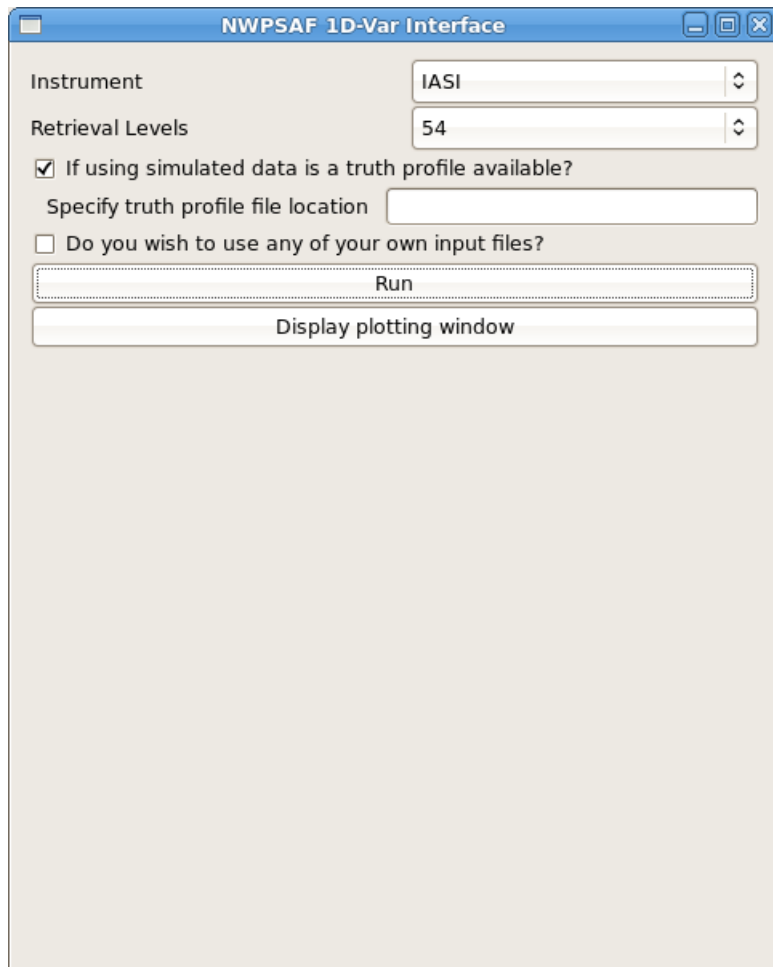Figure 5: The interface in general mode once the code has been run with the "Display plotting window" button displayed

3.2 Test mode

After running the command to open the GUI in test mode:

1dvargui test

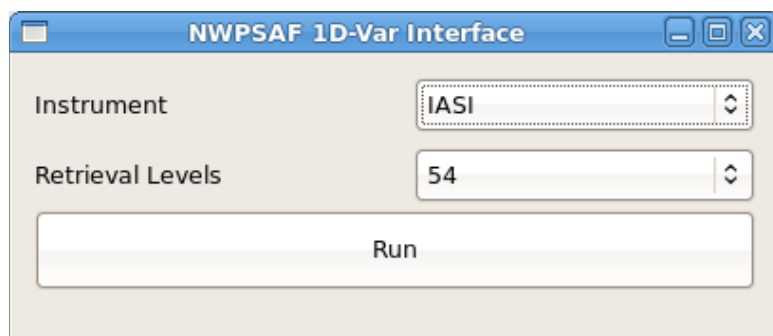The window in figure 6 should appear.

Figure 6: Initial NWPSAF 1D-Var GUI test mode window

The initial interface in test mode only allows two different options to be chosen: instrument and number of retrieval levels with the possible options being the same as in general mode as described in section 3.1.

Once you have chosen the required options click the "Run" button to run the code and, as in general mode, a "Display plotting window" button should appear as in figure 7.
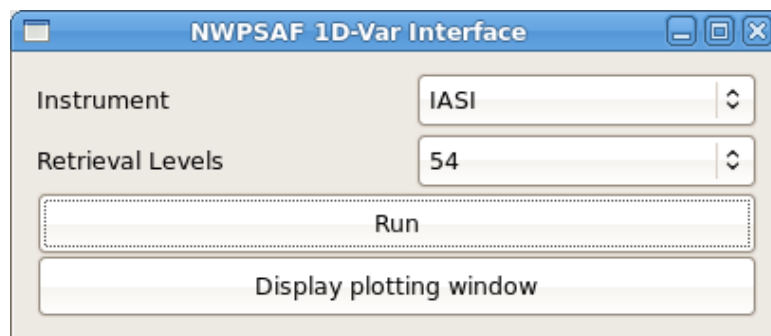

Figure 7: The interface in test mode once the code has been run

In addition to this any differences between the output just generated and the output within the relevant Sample_Output directory will be shown in the terminal which you launched the GUI from. If the code has been set up correctly then there shouldn't be any differences (or very small <0.1% differences if using a compiler other than the default, ifort).

The plotting window gives further options to help visualise any differences and this functionality will be described in section 4.2.

3.3 Training mode

After running the command to open the GUI in training mode:

1dvargui training

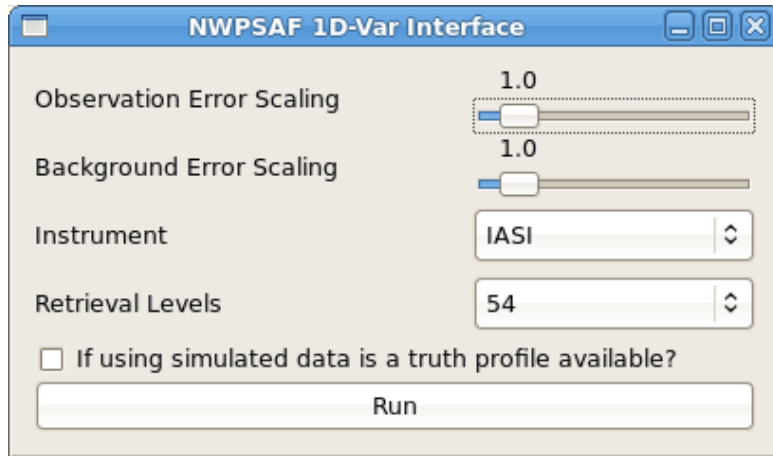The window in figure 8 should appear.

Figure 8: Initial NWPSAF 1D-Var GUI training mode window

In addition to the instrument, retrieval levels and truth profile options from general mode there are two extra options available in training mode which are:

1. Observation error scaling
   A slider which can multiply the observation errors used in the retrieval by a factor ranging from 0.1 to 10.0. This can be used to discover the effects of down or up-weighting the observations on the retrieval.

2. Background error scaling
   As the observation error scaling but for background error.

Once the options have been chosen click the "Run" button to run the code and as in the other modes the "Display plotting window" button should appear as shown in figure 9.
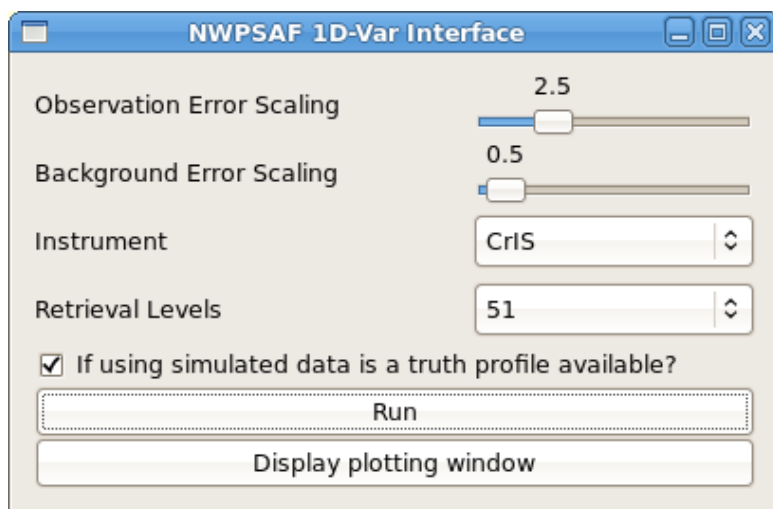

Figure 9: The interface in training mode once the code has been run

4   DISPLAYING PLOTS

Clicking on the 'Display plotting window' button should open a new window as shown in figure 10 for the general and training modes and figure 18 for test mode.

4.1   General and training modes

In general and training modes there are four drop-down menus as shown in figure 10 which allow the production of a variety of plots which will be covered in detail below.
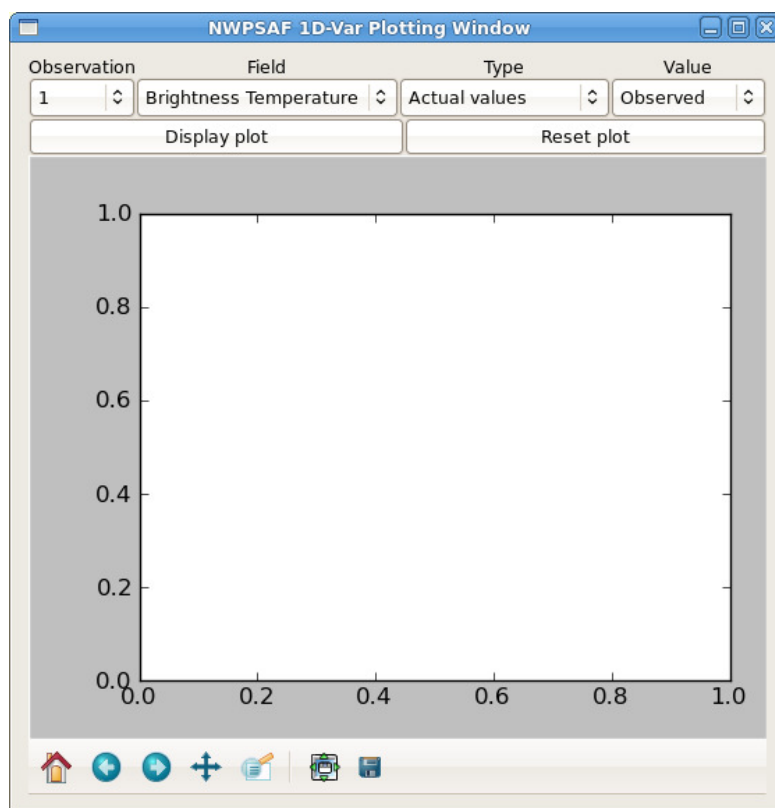

Figure 10: Initial NWPSAF 1D-Var plotting window in general and training modes

The available options are:

1. Observation
   This drop-down menu lets you plot the data from any of the observations you have run retrievals for. The observations will be numbered from 1 to n, where n is the total number of observation locations in the specified observation file. In training mode n will always be equal to 1.

2. Field
   This drop-down menu lets you control which field you wish to plot whether it be in observation space or model space. The available options are brightness temperature, temperature and humidity.

3. Type
    This drop-down menu allows you to control what type of values of the chosen field are plotted. The available options are actual values, normalised values, increments and errors (if temperature or humidity is the chosen field).

4. Value
    This drop-down menu allows you to control what values are plotted and the options available depend on the choice of field or type. If the chosen field is brightness temperature and the chosen type is actual values then the available options are observed, background and retrieved. If the chosen type is normalised then the available options are normalised observed, normalised background and normalised retrieved (which also plots the other two values relative to the chosen normalised value). If the chosen type is increments then the available options are observed – background, observed – retrieval and retrieval – background. If the chosen field is not brightness temperature then there are no observed values available but there may be a truth profile available depending on whether this option has been triggered in the original interface. Hence, if the chosen type is actual values then the available options are background, retrieval and truth. If the chosen type is normalised then the available options are normalised background, normalised retrieval and normalised truth. If the chosen type is increments then the available options are retrieval – background, retrieval – truth and background – truth. Finally, if the chosen type is errors then the available options are background and retrieval.

Once the options have been chosen the "Display plot" button can be pressed and a plot should appear in the window relating to the chosen options e.g. a plot of actual observed brightness temperatures is shown in figure 11. If you want to plot more than one value on the same axis then you can choose this value via the drop-down menu and then press the "Display plot" button again and another line should appear on the plot e.g. lines representing the background and retrieved data are shown in figure 12 in addition to the observed data just shown in figure 11. The plot can be reset to a blank plot either by pressing the 'Reset plot' button or by changing either the field or type options because it doesn't make sense to plot different types or fields on the same axis.

Figures 13 to 16 show various examples of the resulting plots when choosing different options in the interface:

- Figure 13 shows the observed and retrieved brightness temperatures in relation to the normalised background brightness temperatures using the "Brightness Temperature" field option and "Normalised values" type option.
- Figure 14 shows an example where the various brightness temperature increments are plotted using the "Increments" type option.

- Figure 15 shows an example where the background, retrieval and true temperature profiles are plotted using the "Temperature" field option and "Actual values" type option.
- Figure 16 shows an example where the background and retrieval humidity errors are plotted using the "Humidity" field option and "Errors" type option.
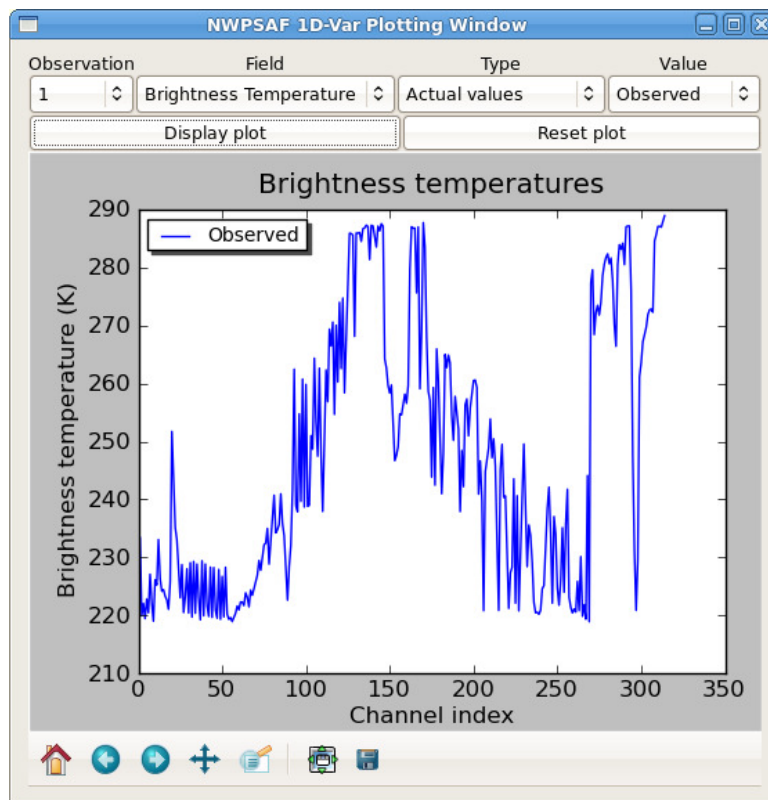

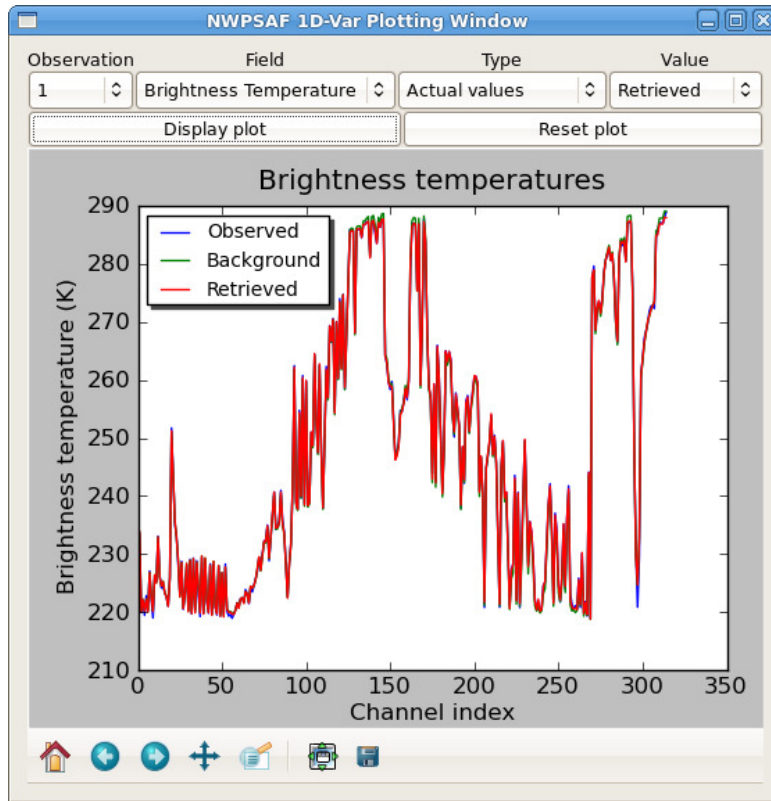Figure 11: A plot of observed brightness temperatures

Figure 12: A plot of observed, background and retrieved brightness temperatures
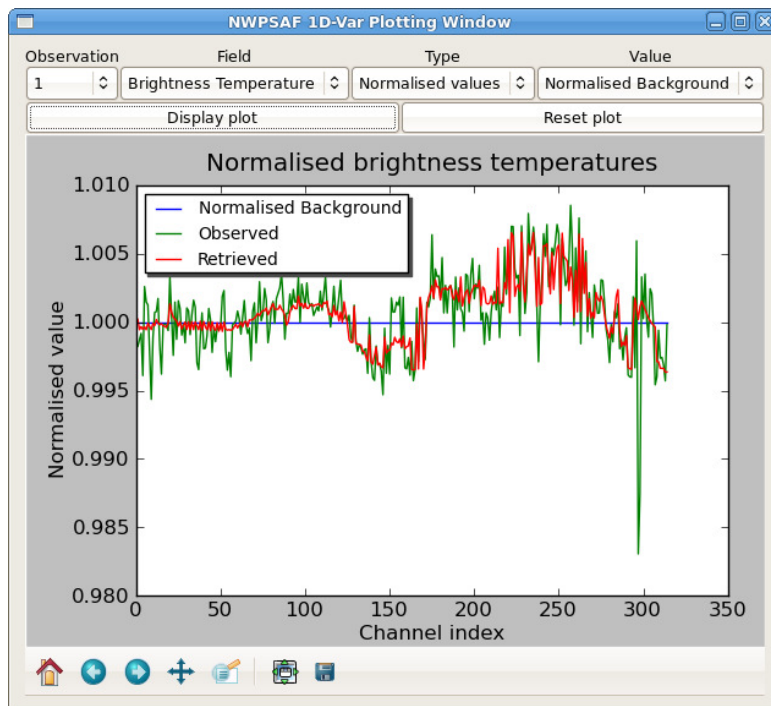

Figure 13: A plot of observed and retrieved brightness temperatures normalised by the background brightness temperatures
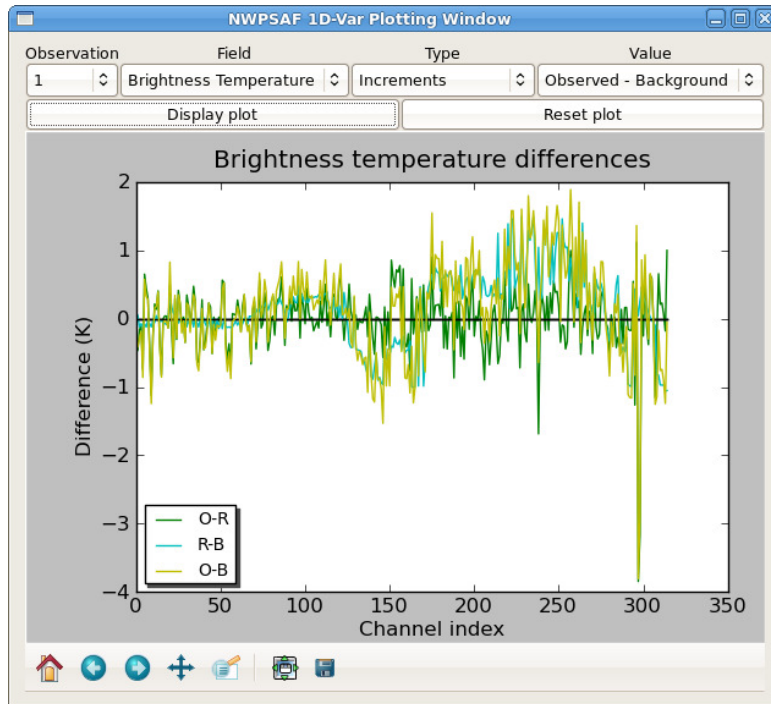
Figure 14: A plot of the observed – retrieved, retrieved – background and observed – background brightness temperature increments
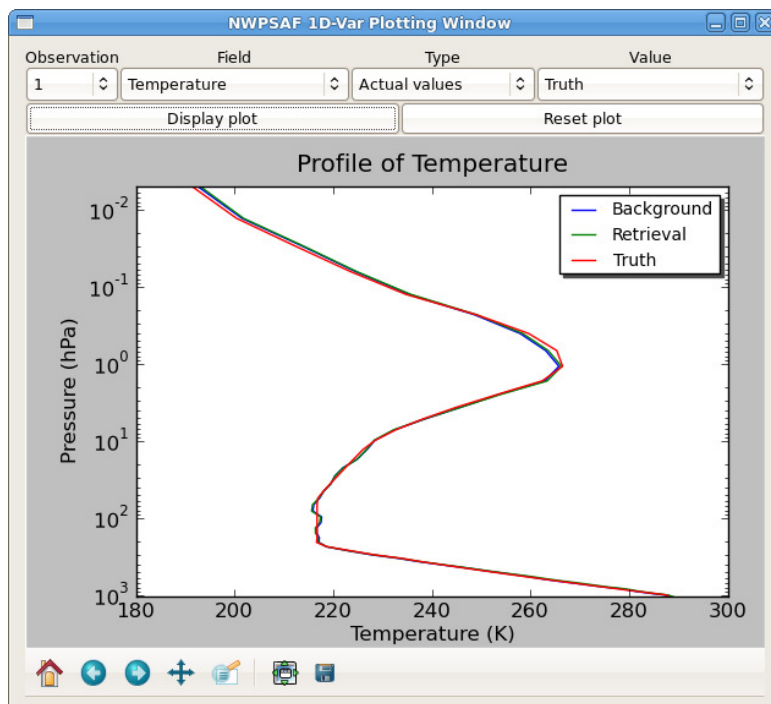

Figure 15: A plot of the background, retrieval and true temperature profiles
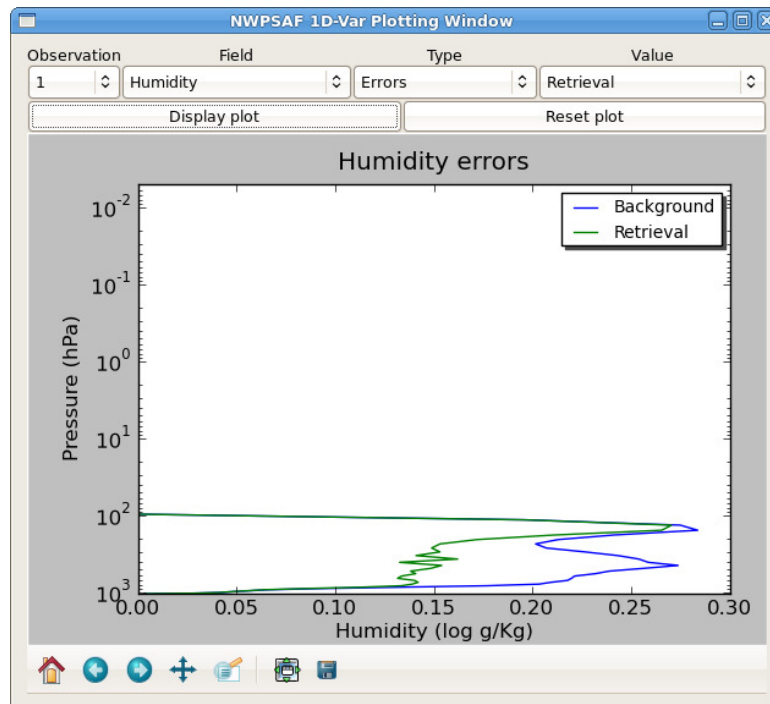
17

Figure 16: A plot of the background and retrieval humidity errors

In training mode it is also possible to compare plots coming from different run time options such as what happens when you up or down-weight the observations or background. This can be done by:

- Plotting the first set of values in the plotting window
- Leaving that window open and going back into the initial 1D-Var interface window
- Using the sliders to modify the error scalings
- Rerunning the code by clicking the 'Run' button and waiting for it to finish
- Then coming back to the plotting window and clicking the 'Display plot' button without changing any of the options

An example of the result of doing this is shown in figure 17 where retrieval increments resulting from different observation and background error scalings are used in the retrievals. The green line uses the default errors, the cyan line uses smaller background errors and the yellow line uses smaller observation errors.
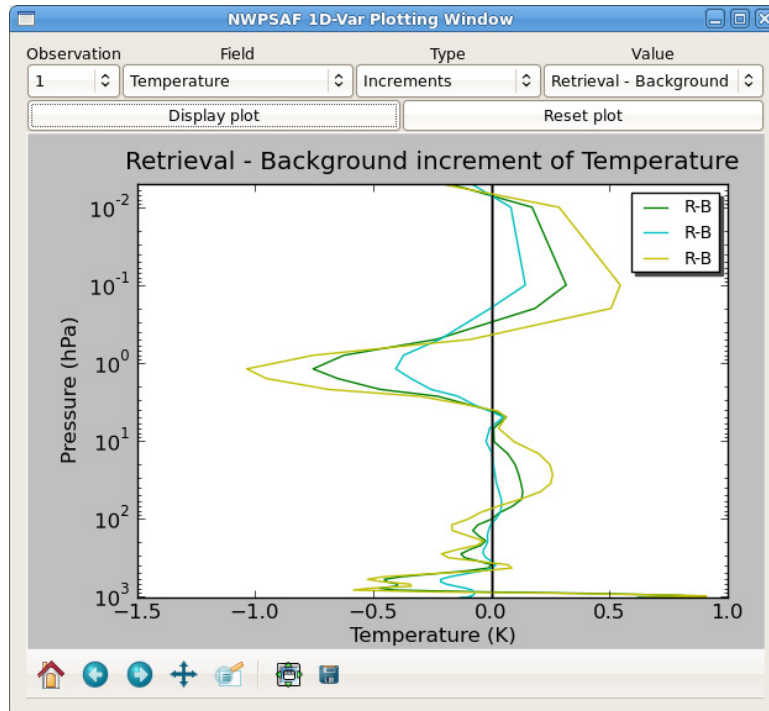
Figure 17: Results of up (yellow) and down-weighting (cyan) observations on the retrieval increments. The initial result with no scaling is shown in green.

4.2 Test mode

In test mode there is only one drop-down menu in the plotting window as shown in figure 18. This drop-down menu is the same field menu as in the general and training modes and has the available options of brightness temperature, temperature and humidity.

Choosing one of these options and clicking the "Display plot" button will plot the difference between the results of the retrieval you just ran with the sample output supplied with the package. If you have made no changes to the package, used the same version of RTTOV and RTTOV coefficients as are specified in the readme.txt file and have installed the package correctly these differences should be zero. Example plots are shown for brightness temperature in figure 19 and temperature in figure 20.

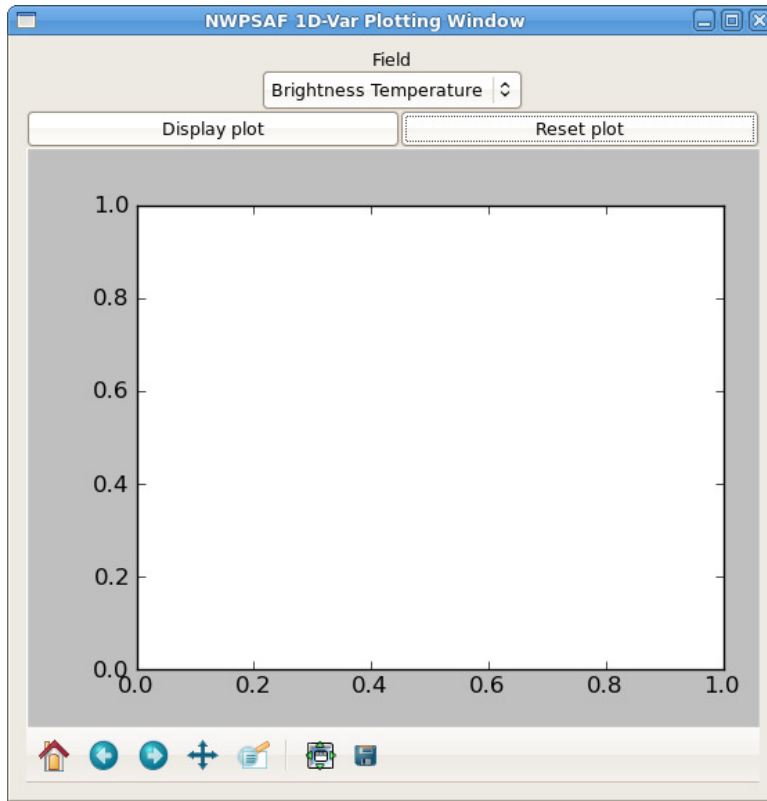|	NWPSAF 1D-Var GUI User Guide	Doc ID : NWPSAF-MO-UD-030 Version : 1.0 Date : 22/08/2014

Figure 18: Initial NWPSAF 1D-Var plotting window in test mode
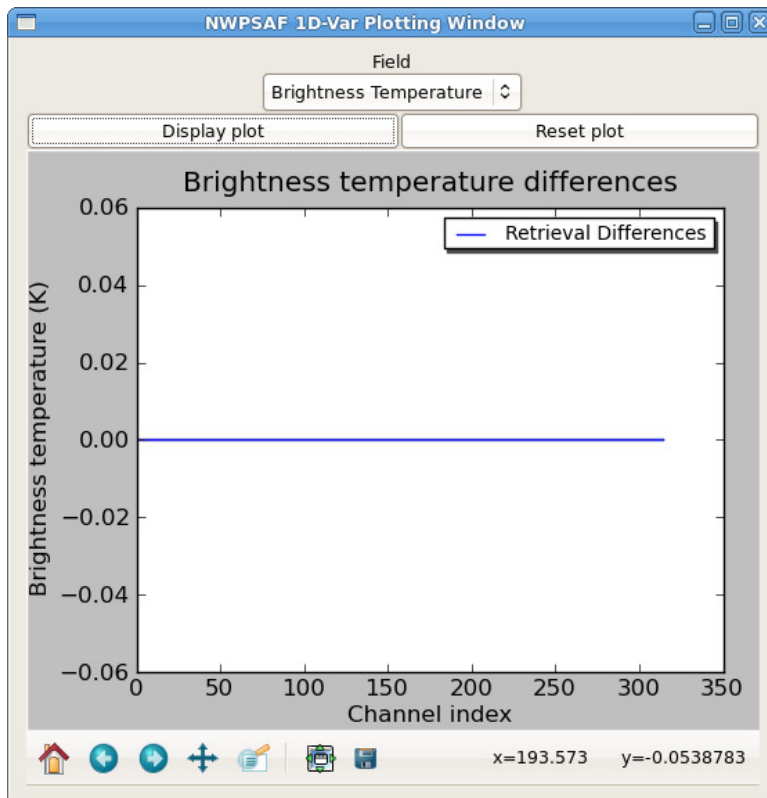

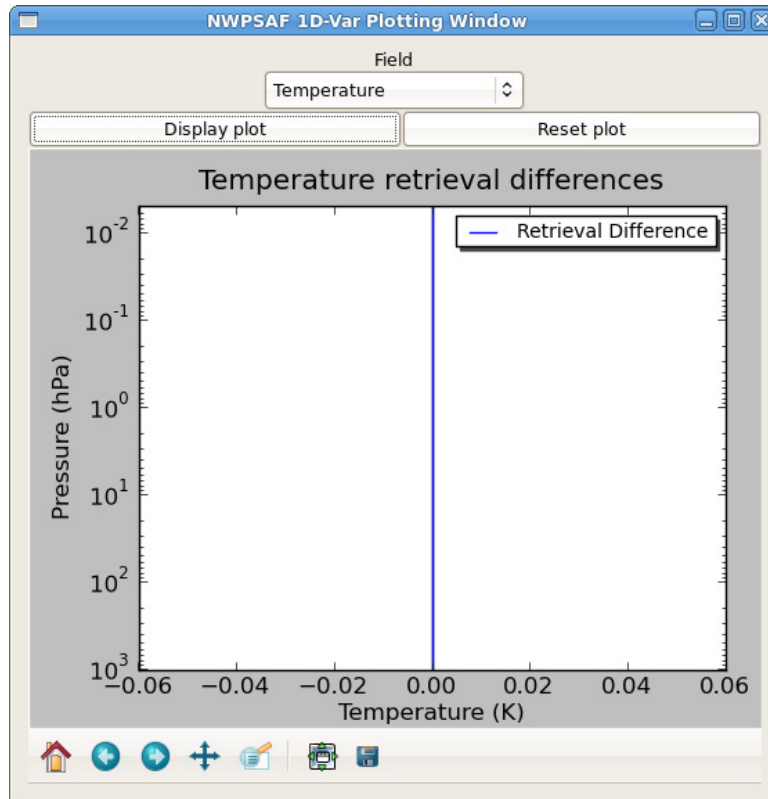Figure 19: A plot of retrieved brightness temperature differences

Figure 20: A plot of retrieved temperature differences

5   LIMITATIONS

In the current version customisation is limited to specifying different input file locations. If you wish to further customise your runs by, for instance, running cloudy retrievals or using an emissivity atlas then there are currently no options within the GUI for doing this and these changes must be made manually as in versions of the Met Office 1D-Var. This sort of functionality is planned to be added in future releases of the code.

When the scaling applied to the background and observation errors has been modified in training mode the retrieval errors plotted correspond to an optimal retrieval i.e. assuming that the true background and observation errors used to create the sample simulated data matches the assumed background and observation errors. If scalings other than 1 are used then the retrieval will not be optimal and the retrieval errors plotted will not be accurate. The optimal retrieval error ($\mathbf{A}_{opt}$) is calculated using equation (1) whereas the true retrieval error ($\mathbf{A}_{true}$) is calculated using equation (2).

$$\mathbf{A}_{opt} = (\mathbf{H}^T \mathbf{R}_a^{-1} \mathbf{H} + \mathbf{B}_a^{-1})^{-1} \qquad (1)$$

$$\mathbf{A}_{true} = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{B}_{true}(\mathbf{I} - \mathbf{K}\mathbf{H})^T + \mathbf{K}\mathbf{R}_{true}\mathbf{K}^T \qquad (2)$$

$$where \ \mathbf{K} = \mathbf{B}_a \mathbf{H}^T (\mathbf{H}\mathbf{B}_a\mathbf{H}^T + \mathbf{R}_a)^{-1} \qquad (3)$$

In the equations $\mathbf{B}_a$ and $\mathbf{R}_a$ are the assumed background and observation error covariances and $\mathbf{B}_{true}$ and $\mathbf{R}_{true}$ are the true background and observation error covariances which are consistent with the simulated background profile and observations. It can be shown by simplifying equation (2) that if $\mathbf{B}_a = \mathbf{B}_{true}$ and $\mathbf{R}_a = \mathbf{R}_{true}$ then $\mathbf{A}_{opt} = \mathbf{A}_{true}$.

This issue will be addressed in future releases of the code.