NCEP Verification System (NVS) Documentation

Perry C. Shafran

The Forecast Verification System (FVS) is currently used to calculate many statistical products to determine the accuracy of models, to find biases, and to validate forecasts. FVS expects input to be of a certain format (see http://www.emc.ncep.noaa.gov/mmb/papers/brill/VSDBformat.txt for description of that format. This note will describe three codes that can be used to generate the format expected by FVS. These three codes are editbufr, prepfits, and gridtobs, collectively known as the NCEP Verification System (NVS).

Below is a flowchart that summarizes the flow of NVS, including all the inputs and outputs of each part of the code. (Chart courtesy of Guang Ping Lou, EMC). The rest of this document provides the details each of the codes, including the output and input files.

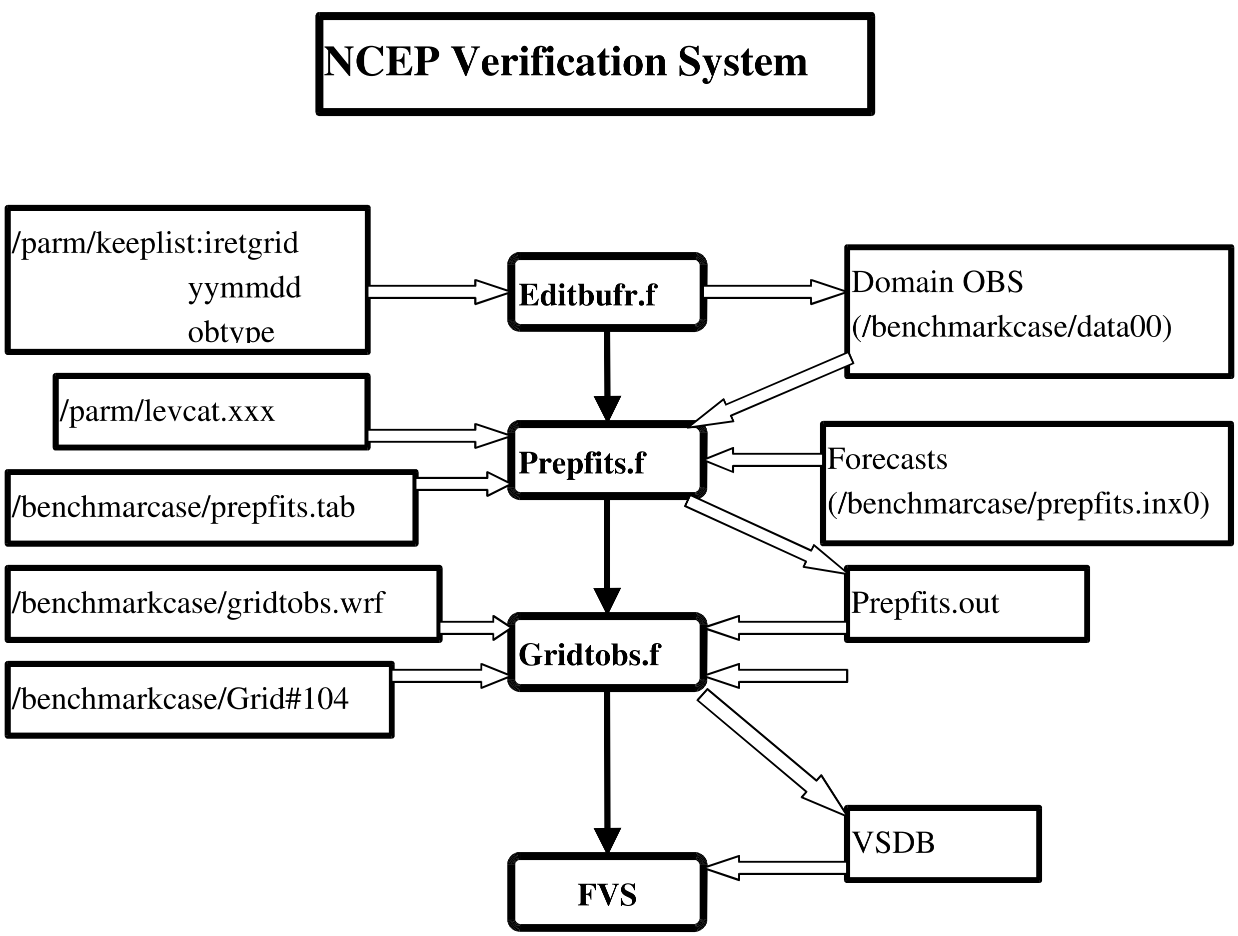

Script set up

The verification script is set up to verify all models, whether it be NAM, GFS, RUC, or even a parallel. That model name is determined at the top of the script and it put into a script variable, PLLN. For example, if PLLN is set to nam, then the other files referenced in the script have to be identified with the PLLN attached to it (e.g. keeplist.nam).

Editbufr

The purpose of editbufr is to reduce the size of a BUFR file to only those observations that are required for verification. The input of editbufr is a file in PREPBUFR format at one specific time, and that is considered the verification time for a particular run.

Editbufr also requires an input file, called keeplist. The keeplist file has three sections that is defined by the user:

IRETGRID - GRID NUMBER OF THE RETENTION AREA

This defines the grid area that will be retained for verification. Ob locations outside of the user-defined grid area will be thrown out. For example, if you are mainly interested in continental US (CONUS) statistics, you can define grid 212, which is the 40-km Eta output grid. It does not matter that it is a 40-km grid; you could use any CONUS grid. If you wanted to verify only Alaska, this grid number would be 216, the Alaska output grid. Be sure to define the proper grid, otherwise the entire observation file might be thrown out.

YYMMDD - DATE OR TIME WINDOW INDICATOR

This defines the time, in hundredths of an hour, of the data to be kept. An observation's header, in BUFR, contains the time distance from the observation time, called DHR. A BUFR file contains observations up to 1.5 hours before and after the observation time. The number indicated here is the time distance that the user wants to keep; anything outside this time window is thrown out. Note that the number must be negative, but that does not indicate that only data, for example, 0.75 hr before the valid time is used. It means that data 0.75 hr before AND after the valid time is used. If you would like to simply keep a specific date's worth of data, this input would be the 2-digit year, month, and day (e.g. 060215 for 15 February 2006).

OBTYP - UP TO 30 OB TYPES TO BE RETAINED

This is a list of all ob types to be retained. The observation type can be found here: http://www.emc.ncep.noaa.gov/mmb/data_processing/prepbufr.doc/table_4.htm. The type indicates whether it is a raob, a surface land data set, and whether it is a mass or a wind observation. For example, the ob type 120 represents the upper-air rawinsonde observations. The user needs to list all the data types that needs to be verified against. Only such observations will be kept; all others will be thrown out.

The output of editbufr is another bufr file that contains only the observations that the user requested to be kept. This bufr file should be saved for use in the next code, which is prepfits.

Prepfits

Prepfits is the code that reads in the observation file from editbufr and also the model data, and interpolates the model data in space and in time to the ob location. It then writes out another new bufr file, the so-called "prepfits file", that will be used to calculate the statistics.

Prepfits requires several inputs. The first input is the bufr file created by editbufr. Another input is the levcat file. The levcat file looks like this, and it is read in via a namelist:

```
 &LEVCAT
 NUMLEV=19
 FIT=.T.,.T.,.F.,.F.,.T.,.T.,.T.,.F.,.F.,.F.
 /
```

The file provides two pieces of information. First, NUMLEV is the number of levels in the model output. Based on the number of levels, prepfits will be able to define the pressure at each of the levels. The second piece of information is the FIT variable. This defines what type of observation-level data the user wishes to verify. The levels are defined as

0 - surface
1 - mandatory level
2 - sig level T & moisture
3 - winds by pressure

4 - winds by height
5 - tropopause
6 - ANY single level
7 - auxiliary
8 & 9 - reserved

The user should define as .TRUE. or .FALSE. which observation-level data that are desired to be verified against.

Another input is the BUFR table. The BUFR table is the one that will be used to generate BUFR output. The input BUFR file, as modified by editbufr, already is encoded with its own BUFR table. The BUFR table identifies each BUFR variable used in the file and the groups they are contained within. Not much further will be said about the BUFR table here.

Prepfits is run at the verification time. That means that all forecasts that verify at that particular time should be input in a prepfits.in file. This file lists all the GRIB forecast files that are valid at that verification time. For example, at 00Z on a given day, all the forecasts to be verified are the 00Z analysis on that given day, the 6-hr forecast from 18Z the previous day, the 12-hr forecast from 12Z the previous day, the 24-hr forecast from 00Z the previous day, the 36-hr forecast from12Z two days ago, and so on. All of these forecasts are valid at the same time and should be listed.

The output of the prepfits is the prepfits file. This file is a BUFR file that contains, at each observation location, the observation and all the forecasts. This file needs to be retained for the final code in this process, gridtobs.

Gridtobs

Gridtobs (pronounced “grid to obs”) is the code that brings in the observation and forecasts on each ob location to write out “partial sums”. These partial sums are then used by other codes, most notably FVS, to calculate and plot out many different types of statistics. Like prepfits, gridtobs is also run at a single verification time. All the forecasts processed in a prepfits run are also processed in the proceeding gridtobs run.

The result of gridtobs is the “VDB” (Verification Data Base) file. More will be discussed about the VDB file later in this document. For now, note that all the numbers that come out of gridtobs are domain-averaged over the grids selected in the gridtobs control file (discussed below). These numbers are called "partial sums", and these numbers are the domain averages of the forecast, the observation, the square of each, and the product of forecast times observation. From these a whole host of statistics can be calculated using FVS. For vector winds, there are U and V partial sums, with more items but the idea is similar.

Gridtobs has several inputs. First input is the prepfits file from prepfits. The second is the grid#104 file. The grid#104 file works in conjunction with another input, the regions file.

The regions file lists 30 regions for which the user can write out partial sums. The grid#104 file lists, for each gridpoint on NCEP grid 104, what region that point belongs to. All other regions are subsets of grid 104. These files do not have to be changed from run to run.

The main item controlled by the user is the gridtobs control file. This file sets precisely what you want to verify. Let's see an example of the control file:

```
V01   10                              Section 1
   1  NAM/218                                 2
   4  00                                      3
  06
  12
  18
   1  19                                      4
   1  ADPUPA                                  5
   7  G236                                    6
  G245
  G104/NPL
  G104/SPL
  G104/MDW
  72204
  72315
   2  SL1L2                                   7
  FHO
   6  Z                                       8
  T
  RH
```

```
PBL
CAPE 6 200 400 600 1000 1500 2000
VWND
 11  P1000                              9
P850
P700
P500
P400
P300
P250
P200
P150
P100
P50
```

Let's take this file one section at a time.

1) This line represents the version number of the verification system, which is always version 1.0 at NCEP. The "10" represents the version number of the input BUFR file to gridtobs. This line does not have to be changed.

All the following sections have a number before each section. That number denotes how many items follow in that list. For example, section 2 has a "1" before the "NAM/218", denoting that there is 1 item in that list. Section 3 has a "4" to denote that there are 4 items in the list to follow. Note that the first item is in the same line as the number of items in the list.

2) This is the grid that you are verifying. In this case we are verifying the NAM model on output grid 218. If you are verifying the NAM Grid 212, this line would read "NAM/212". If you are verifying the WRF Grid 216, this would read "WRF/216". If it were the GFS it would read "GFS/212". As a matter of fact you can put anything in this line to note what you are verifying, such as "NARR" for North American Regional Reanalysis, or "Model #1". The maximum number of models that can be verified currently is 6.
3) This lists all the forecast hours you are verifying. In this example, forecast hours 00, 06, 12, and 18 are being verified. That means that in this run of gridtobs, 4 different forecasts are being verified, all at this particular verification time. The maximum here is 20 forecast hours.

4) This is the list of verification dates being used. However, since gridtobs is generally run at one verification time, this line is ignored and the date is supplied by the prepfits file.
5) This is the ob type that is being verified. A maximum of 10 verification ob types can be verified. Unlike the keeplist file in editbufr, these are written out as they are in the prepfits file. In this case "ADPUPA", which represents all upper-air profile data, such as rawinsondes, pibals, or profilers, is being used. Other choices are below:
   - ADPSFC: Surface land data
   - AIRCFT: Aircraft data
   - AIRCAR: ACARS data
   - MSONET: Mesonets
   - ANYAIR: All aircraft data
   - SFCSHP: Ship and buoy data
   - ANYSFC: All surface data, including land and sea, and including the lowest level of upper-air profiles
   - ONLYSF: All surface data, not including the data from upper-air profiles
6) This lists the output domains that you want to verify on, up to a maximum of 16. These can be NCEP grids, as in the example above "G236". This means that the verification numbers are averaged over the G236 domain. Same with the G245 example.

These output domains can also be regions, as defined by a three-letter domain definition from the regions file. Each region must start with the "G104" since the domains are defined from the grid#104 file. In the example above, one item is "G104/NPL". That means that the partial sums will be averaged over the NPL grid as defined by the grid#104 file.

Finally, if you do not wish to average numbers over a domain, there is the option of posting the numbers for single stations. They are defined by their station ID, as above (i.e. 72204). These examples may or may not be real station ID's, only examples.

7) This is the statistic type. SL1L2 is used for scalar data, VL1L2 is used for vector data, and FHO is used for FHO (Forecast, Hit, Observation) statistic types. Note that just putting in the SL1L2 in this line will allow the user to verify both scalar and vector data, and the VL1L2 is not needed.
8) This is the list of variables to be verified, up to a maximum of 16 variables. In this list the user must list all scalar data first, then the vector data. In this case Z, T, RH, PBL, and CAPE are the scalar data being verified, and then the vector wind VWND. Here

is a list of what is frequently verified, although this list is by no means exhaustive:

- Z (for upper-air data)
- T
- RH
- Q (specific humidity)
- PBL
- CAPE (surface based)
- CIN (convective inhibition)
- LI (lifted index)
- PWO: Precipitable water, compared to integrated atmospheric water vapor in raob
- TROP: Tropopause level (in meters)
- VWND (vector wind)

The user may add or subtract variables but again note that all vector variables must be listed last.

Here is a list of surface data that can be verified:

- SLP: Sea-level pressure
- T: 2-m temperature
- TMAX: Max 12-hr temperature (use only at 00Z)
- TMIN: Min 12-hr temperature (use only at 12Z)
- DPT: Dew point
- VIS: Horizontal visibility
- TCLD: Total cloud fraction
- GUST: Max wind gust
- HEAT: Heat index
- CHILL: Wind chill
- PWO: Precipitable water (compared to the GPSIPW observations)
- CLDBT: Height of cloud base in meters
- VWND: 10-m vector wind

Note that in the example one of the above examples there is a series of numbers after the variable name:

CAPE 6 200 400 600 1000 1500 2000

This means that in this example, CAPE will not only be verified using the SL1L2 statistics, but also FHO statistics. The first number following the name of the variable denotes the number of thresholds that FHO stats will be calculated for. Then the thresholds are listed. Make sure that the number of thresholds listed is equal to the number indicated first in the list, otherwise the code will not run.

Note that not every variable needs to have a list of thresholds. If you do not need to verify a variable using FHO statistics, simply do not list any threshold values.

9) Finally, this is the list of levels to verify, up to a maximum of 19. In this example, there is a list of pressure levels. P1000 represents the 1000-mb level, P500 represents the 500-mb level, and so on. The "P" must be included to indicate that pressure-level data is to be verified. If you are verifying surface data, then only one level is needed: SFC (no other indicator). SFC will verify all the surface data, including the 2-m temperature and 10-m winds.

The user can place multiple sections into one control file. It is generally best to list the upper-air portion separately from the surface portion, or the aircraft, etc. If you want to verify more than 16 domains, you might want another list in the control file, similar to above. Please refer to the list above for the limits of number of variables, models, etc. The user can list them one after another, with each section starting with the "V01 10" line. Note that only one control file can be input into any particular run of gridtobs.

Verification for FVS users

Most people use the Forecast Verification System (FVS) to calculate statistics and plot them out. For more specific information on FVS, please consult http://www.emc.ncep.noaa.gov/mmb/papers/brill/FVShelpfile.txt or email Keith Brill, of the NCEP's Hydrometeorological Prediction Center at Keith.Brill@noaa.gov. The following is only meant to set up your files for use in FVS.

A particular run of the NCEP Verification System (NVS) produces a "VDB" file, which is only valid at one forecast hour. It is best to concatenate all the files for one day into a "VSDB" file. FVS will search for VSDB files, not VDB files. The VSDB file contains the entire day's worth of verification.

For example, a particular run of NVS produces a file called nam_2006092400.vdb. This is the verification valid at 00Z on 24 September 2006. You might also have files called nam_2006092403.vdb, nam_2006092406.vdb, etc. Concatenate all these files together to produce a file called nam_20060924.vsdb. This command works best:

```
cat nam_${DATE}*.vdb > nam_${DATE}.vsdb
```

where $DATE=yyyymmdd.

Your files must be organized in a directory so that FVS will be able to search for them. FVS requires that three things must be the same: the model name in Section 2 of the gridtobs control file above (which in turn is written out inside of the VSDB file), the model name in the file name, and the model name in the directory it is stored in. For example, if your model is NAM, you would have NAM/218 (as an example, the grid number does not matter) in your Section 2 of the gridtobs control file. Your filename would include the model nam (as in the example directly above), and the directory the files are stored in is called nam, like this:

```
$VSDB_storage_directory/nam/nam_20060920.vsdb
$VSDB_storage_directory/nam/nam_20060921.vsdb
$VSDB_storage_directory/nam/nam_20060922.vsdb
.
.
.
```

FVS will search for the files in $VSDB_storage_directory as long as you set the search environment variable to do so. The following command does so:

```
setenv VSDB_DATA $VSDB_storage_directory
```

You may have several directories in the $VSDB_storage_directory (such as gfs, wrf, namy, etc.) as long as the files inside those directories follow the conventions discussed above.

You may run FVS on either the IBM supercomputer cluster or the workstation system. In either case, you need to have all your environmental variables set to the location of the FVS codes, scripts, etc, that you need to link. Please email me at Perry.Shafran@noaa.gov for help in setting up to either system.