FREQUENTLY ASKED QUESTIONS:

A lot of questions can get answered by inspecting/searching the COMMENTS section at the top of the IRI Fortran programs (in particular irisub.for and irifun.for).

Are there plans in the future to provide the source code in languages other than FORTRAN?
GITHUB.com is a good place to search for IRI code versions in different languages. Check also links on the AVAILABILITY section of irimodel.org.

Getting NaN outputs from IRI:
This could be a compiler issue. Most compilers will initialize all the variables to zero if they do not already have a value (e.g., f77, f90 on macOS). If the compiler does not initialize everything automatically it most likely will have flags that can be set to achieve this, e.g.,
Simply Fortran: -ffpe-summary=none -finitlocal-zero -fno-automatic
gfortran: -finit-local-zeros -fno-automatic

I am not sure which one of the solar indices in the IG_RZ.DAT and APF107.DAT files are definitive and which are predicted?
All indices in the APF107.DAT file are definitive except for the last two values in each line which are the 81-day average of the solar 10.7cm radio flux F10.7 (F5.1) and the 365-day average of F10.7 centered on the date of interest (F5.1). At start and end of the index file the 81-day and 365-day averages are calculated taking only the available indices, e.g. for the first date the 81-day average is only over 40 F10.7 values and over 41 values on the 2nd date. In the IG_RZ.DAT (also ASCII) the release date is given in the first line in the form Month, Day, Year. Next is a blank line and than Start_Month, Start_Year, End_Month, End_Year for the indices given in the file. The indices provided are the 12-month running means therefore the last definitive is the one six month before the given release date, so for a release date of Jan 2013 the first predicted index is the one for July 2012.

What is the format of the APF107.DAT file and which parameters are included?
The format of the file is explained in subroutine APF in IRIFUN.FOR: year(I3), month(I3), day(I3), 3-hour Ap indices for the UT intervals (0-3),)3-6),)6-9), ..,)18-21),)21-24(in an array of dimension 8 (8I3), daily Ap (I3), -11(I3), F10.7 radio flux for the day (F5.1), 81-day average of F10.7 radio flux (F5.1), 365-day average of F10.7 centered on the date of interest (F5.1). At start and end of the index file the 81-day and 365-day averages are calculated taking only the available indices, e.g. for the first date the 81-day average is only over 40 F10.7 values and over 41 values on the 2nd date.

What is the format of the IG_RZ.DAT file?

The indices file IG_RZ.DAT contains the 12-month running mean of the ionosphere global (IG) index (IG12) and of the sunspot number (Rz12). IG_RZ.DAT is structured as follows (values are separated by comma) line by line:

month,day,year of the last update of this file,
a blank line
start month, start year, end month, end year,
a blank line
the IG12 index for the month before the start month of start year (needed for interpolation)
the IG12 indices for the rest of the start year,
the twelve IG12 indices for the year following the start year
.. and so on until the year before the end year,
the IG12 indices for the end year from January to the end month+1 (needed for interpolation)
a blank line
the Rz12 index for the month before the start month of start year (needed for interpolation)
the Rz12 indices for the rest of the start year,
the twelve Rz12 indices for the year following the start year
.. and so on until the year before the end year,
the Rz12 indices for the end year from January to the end month+1 (needed for interpolation)

Have there been efforts to assimilate data into IRI or to update IRI with measured data?

There are quite a number of different mathematical methods for data assimilation and model updating. Here are a few references:

Assimilation: (1) Yue Xinan, William S. Schreiner, Ying-Hwa Kuo, Douglas C. Hunt, Wenbin Wang, Stanley C. Solomon, Alan G. Burns, Dieter Bilitza, Jann-Yenq Liu, Weixing Wan, Jens Wickert, Global 3-D Ionospheric Electron Density Reanalysis based on Multi-Source Data Assimilation, J. Geophys. Res., 117, A09325, doi:10.1029/2012JA017968, 2012.

(2) Galkin, I. A., B. W. Reinisch, X. Huang, and D. Bilitza (2012), Assimilation of GIRO data into a real-time IRI, Radio Sci., 47, RS0L07, doi:10.1029/2011RS004952.

(3) Schmidt, M., D. Bilitza, C.K. Shum, and C. Zeilhofer, Regional 4-D modeling of the ionospheric electron density, Adv. Space Res. 42, #4, 782–790, doi:10.1016/j.asr.2007.02.050, 2008

Updating: (1) Hernandez-Pajares M., J. Juan, J. Sanz, and D. Bilitza, Combining GPS measurements and IRI model values for Space Weather specification, Adv. Space Res. 29, #6, 949-958, 2002.

(2) Bilitza, D., S. Bhardwaj, and C. Koblinsky, Improved IRI predictions for the GEOSAT time period, Adv. Space. Res. 20, #9, 1755-1760, 1997.

Why do I have to specify a height even if I generate a height profile of an IRI parameter?

This question relates to the IRIweb online interface as well as the IRI_WEB subroutine in IRISUB.FOR. In both cases you can decide what type of profile you want to generate. This can be a latitudinal, longitudinal, diurnal, etc profile of an IRI parameter. If you select a latitudinal profile you than have to select the begin, end, and stepsize for the latitude interval. You still have to also enter a latitude in the earlier part of the IRIWeb interface or the IRI_WEB subroutine arguments, but this value will be ignored. The same hold true for any other type of profile, e.g. altitude.

Why do I get discontinuities in the height-derivative of electron density at the E-valley top?

The E-region valley electron density profile is modeled by a 5^{th} order polynomial that is defined by the following constraints: (1) Takes the value NmE at hmE and has derivative=0 there; (2) Produces valley with given WIDTH, so density at hmE+WIDTH is NmE; (3) Reaches the lowest density at height h_deep with derivative=0 at that point; (4) Valley has the DEPTH=(NmE-N(h_deep))/NmE; (5) Derivative at valley-top divided by NmE is DLH. Coefficients A(2:5) are determined from given WIDTH, DEPTH, HD=h_deep-hmE, and DLH. The transition at the valley top is continuous in value but not in derivative, because the DLH is given as input and not determined from the F1-region profile function.

Why do I get discontinuities at sunrise and/or sunset in the diurnal variation of electron density in the E-valley region?

Because of the very deep night valleys we had to switch to a logarithmic representation at nighttime. Please also see answer to "Why do I get discontinuities in the height-derivative of electron density at the E-valley top".

Why do I get NaN for electron density?

In some rare cases, when computing the electron density, a NaN can be obtained. It was found that in the computation of the magnetic inclination DIP= ASIN(BDOWN/ BABS) in the igrf_dip subroutine of the igrf.for file the argument BDOWN/BABS of ASIN can be slightly larger than 1.0, thus generating the NaN. This has to do with the compiler used and the system on which the program is being run, because the same rare case did not case problems for other configurations. A simple fix is to change the statement

```
DIP=ASIN(BDOWN/BABS)
```

into the following four statements

```
argasin=BDOWN/BABS
if(argasin.gt.1.0) argasin=1.0
if(argasin.lt.-1.0) argasin=-1.0
DIP=ASIN(argasin)
```

Problem with blank COMMON being initialized twice:

Some compilers complain about the blank COMMON being initialized in two subroutines (SHAMDB0D and SHAB1D of IRIFUN.FOR) . Here is a step-by-step solution to the error:

Open irifun.for:

Replace lines 5388-5389, 5560-5561, and 5734-5735 with:
" COMMON /bnew/BINT,BEXT,RE,TZERO,IFIT,IB,KINT,LINT,KEXT,"
" * LEXT,KMAX,FN"
Replace lines 5392-5394 and 5566-5568 with:
" THETA = 180."
" ICEN = 0"
" IREF = 0"
Add the following subroutine just above SUBROUTINE SHAMDBOD on line 5356:
" SUBROUTINE INITB ()
" PARAMETER (IBO=0,JBO=1,KDIM=6,LDIM=4,L=-1)
" DIMENSION FN(0:KDIM,0:KDIM), CONST(0:KDIM,0:KDIM)
" DIMENSION BINT(0:KDIM,0:KDIM,1-IBO-JBO:LDIM),
" * BEXT(0:KDIM,0:KDIM,1-IBO-JBO:LDIM)
" COMMON /bnew/BINT,BEXT,RE,TZERO,IFIT,IB,KINT,LINT,KEXT,
" & LEXT,KMAX,FN
" DATA RE,TZERO,IFIT,IB,KINT,LINT,KEXT,LEXT
" * /6371.2,1.0, -1, 2, 6, 4, 0, -1/
" RETURN
" END
Save and close irifun.for and open irisub.for
Add the following line after line 1293:
" CALL INITB()"
Save and close irisub.for
Run compilation/link command in Fortran

Why do I get WARNINGS when compiling IRIDREG.FOR?
The large BLOCKDATA statement used for the FIRI model for the D-and E-region causes some compilers problems. The WARNINGS mostly are due to the fact that such large BLOCKDATA statements slow down the program and take lots of memory.

Why is there a sharp increase in the H+ density at 300km
In IRI different models are used for the ion composition above 300km (O+, H+, He+) and below 300km (O+, O2+, NO+) separating the parts dominated by light ions and by molecular ions, respectively. The H+ density at 300km is very small but below 300km it is zero, and thus there is a jump in H+ density at 300km. One way to avoid this jump is to extrapolate the H+ density profile from above 300km to the region below best using an exponentially decreasing function.

Representation of the D-region electron density in IRI
The CCMC online version provides access to two options: FIRI model (**FT-2001**; Friedrich and Torkar, J. Geophys. Res., 106/A10, 21409-21418, 2001) that is based on Friedrich's compilation of reliable rocket data. The older **IRI-95** model is based on a much smaller selection of typical rocket profiles (Mechtly and Bilitza, Report IPW-WB1, Freiburg, Germany). **IRI-95** is the Default. NOTE:
In the IRI code irisub.for a third option is given for the D-region electron density. This is the model developed by Danilov et al. (1991, 1995) which includes options to account for

the effects of stratospheric warming and D-region winter anomaly. Their model is based on a separate set of rocket measurements. The output values are stored in a special section of the IRI output file and the iritest.for test program illustrates how to access these values.
The IRI-95 D-region profile is normalized to the E-peak, the profiles of the other two options are not!