IRSAViewer Help

Table of Contents

IRSA Viewer: Overview

The IRSA Viewer allows image visualization and catalog overlays, from off your local disk, the web, and/or other IRSA services.

You can initiate IRSA Viewer by loading the tool directly from the IRSA Viewer button on the front page, or by clicking on a link that says "to IRSA Viewer", or by clicking on a link in a list of search results that looks like this:  . If you have loaded the tool directly, the first time you see the IRSA Viewer, it does not have an image pre-loaded; if you have clicked a link or icon to launch it, the first time you see the IRSA Viewer, it already has an image loaded.

Besides the online help (also available as a PDF), note that there are also video tutorials, including quick start and longer AAS-demo-style overviews, available at the IRSA YouTube channel. Also see the IRSA Viewer playlists (one of which collects all the IRSA Viewer tutorials together), as well as the playlist of tutorials relevant for more than one archive.

Contents of page/chapter:

Terminology

The words in rectangles across the top are 'tabs.'

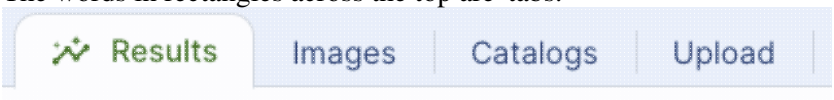

and you can add or remove tabs from this top level.

This icon in the upper left  pulls open a "drawer" from the left hand side which enables you to add or remove tabs from this top level (see below). It also can allow you to change the layout or appearance (dark or light mode) (see below).

When you have things loaded into IRSA Viewer, your browser window is divided into "**panes**", like "window panes." The contents of the panes depends on what you are doing with the tool, but could include an image pane, a catalog pane, and/or a plot pane. You can expand any of the window panes by clicking on the expand icon: 

Each of the three main kinds of 'panes' in the display has its own toolbox in its upper right corner which operates on things in that pane, and the basic functionality for each of these panes is covered elsewhere in this document:

- The Tables section covers (among other things) the tables toolbox 
- The Plots section covers (among other things) the plots toolbox 
- The Visualization section covers (among other things) the images toolbox 

The idea behind "**pinning**" is that you can retain a given item (a plot or table or image) within the tool. "Pinning" just means "hold on to this item within this tool." It doesn't mean "save this plot (or image or table) to disk", nor does it mean "download the data behind this"; it means "retain this item in this tool for now."

Tools Overview

There is a chapter on images in general, and then a separate chapter on the myriad visualization tools available for manipulating images.

There is a chapter on tables in general, and then separate chapters on both catalogs from IRSA and plots you can make from catalogs. You can also use VO services to conduct more complicated searches, including searches of other archives.

The images, catalogs, and plots are all interlinked such that clicking on one point in the plot highlights the same point in the catalog and the image.

Spectra are a special case of data that use images, plots, and tables, so there is a special chapter on just spectra.

Downloading data has its own page.

The Time Series Tool is somewhat of a separate tool within IRSA Viewer, but its documentation lives here, within IRSA Viewer.

Side Menu and Adding to the Tabs Menu

This icon  in the upper left pulls open a "drawer" from the left hand side; the top of it looks like this:

The highlighted bar ("Results" in this example) is the tab you have in the foreground on your main window.

You can use this side menu to add (or remove) blue tabs from the top of your IRSA Viewer interface. By default, Results, Images, Catalogs, and Upload are shown.

Images and Catalogs search IRSA holdings; if you want to search more IRSA data using different tabs, add the "Data Collections" or "VO TAP" tabs.

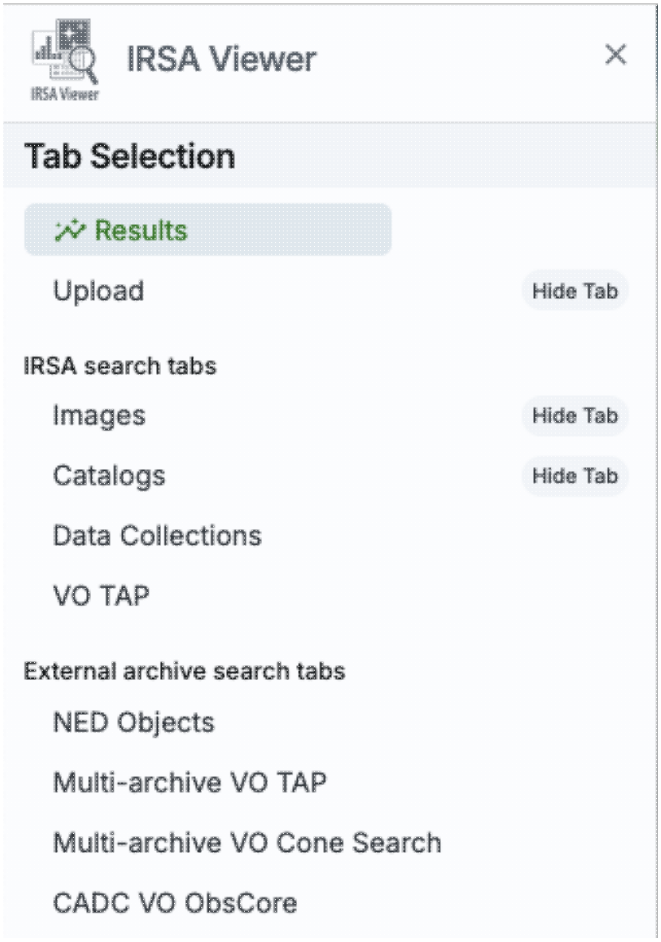

If you want to search other archives, add one of the tabs under "External archive search tabs": NED Objects, Multi-archive VO TAP (that is, a general TAP search), Multi-archive VO Cone Search, or a CADC VO ObsCore.

Click on the "Hide Tab" button to remove that corresponding tab.

Side Menu and Layout/Appearance

This icon  in the upper left pulls open a "drawer" from the left hand side; the bottom of it looks like this:

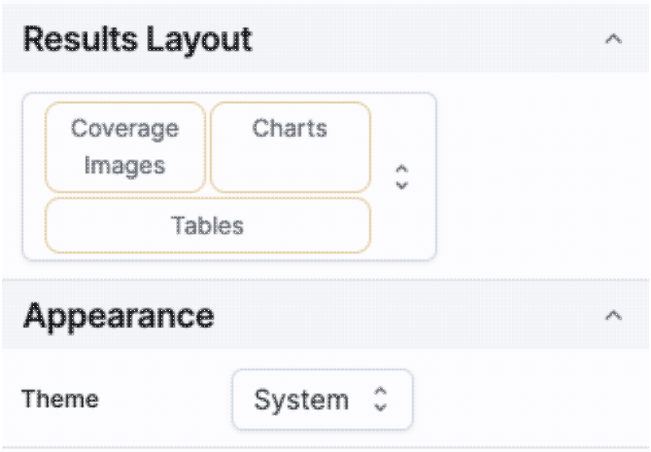

The first part of this is only available once you have some results shown in the tool. It controls how your results are displayed. By default, it looks like the stylized sketch as shown, with coverage and images on the upper left, charts (plots) in the upper right, and tables on the bottom. You have many more choices, though:

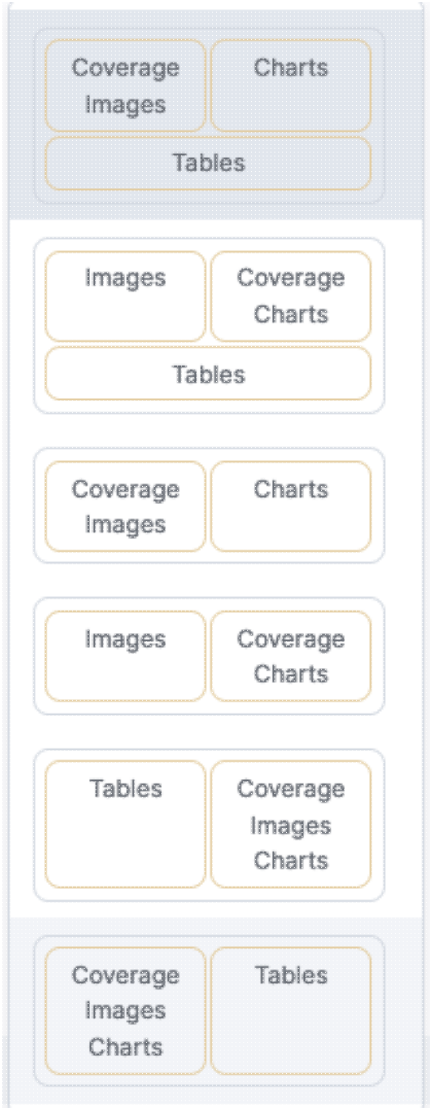

From this menu, you can control the layout of your Results tab -- the relative placement of your images, plots (charts), and tables. The darker one (the top one in this example) is the currently selected one.

The bottom of the side menu controls the appearance of the tool in your browser -- do you want it to run as light mode, dark mode, or respect whatever preferences you have set on your system? Try out the different modes; you may have a preference!

User Login

In the far upper right, there is a link to log in. IRSA Viewer can remember you when you return. See the user registration section for more information.

Getting More Help

The "Help" icon leads you into this online help. There are also context-sensitive help markers throughout the tools ( ). You can also download a PDF version of this manual; look at the top left of the help window.

(The PDF may be easier to search than the web pages; use your PDF reader's search function.)

You can submit questions to the IRSA Help Desk.

A set of frequently asked questions (FAQs) about IRSA Viewer is here.

The IRSA YouTube channel has lots of short videos about IRSA tools.

Found a bug? The known bugs and issues in this version of the IRSA Viewer are listed here. If you think you have found a bug, before reporting it, please check this list, and read this online IRSA Viewer help. It may be a "feature" we already know about. If you have found a new, real bug, then please do contact us via the IRSA Help Desk. Please include your operating system version and your browser software and version. If you can, please also include any specific error message you may have gotten. (NB: In our testing, copying shortcuts worked on Windows and Linux; the command-C did not always work on Macs, but selecting and clicking the right mouse button often did when command-C did not.)

IRSA Viewer: Quick Start for the Impatient

Yep, IRSA Viewer (and all its subsidiary tools: Time Series Tool, Data Collection Explorer, Catalog Search Tool, Data Discovery, IRTF, and Atlas) sure do look different now! We are very excited about this new look and feel for our tools. **All of the same functionality is still here, with some new features too.** We will be continuing to update the user experience (UX) to continue to make it easier to use our tools.

This chapter is aimed at people who have been familiar with the "old style" IRSA Viewer, but could also give new users a jump start on using the tool.

Contents of page/chapter:

You Did What Now?: Navigation Philosophy

The blue tabs across the top: 

are the main way to navigate through searches and results. Any searches you do will be from a blue tab, and all the results go into the results tab.

By default, you have tabs that read Results, Images, Catalogs, and Upload -- Images and Catalogs are both where you go to search IRSA images and catalogs, respectively.

Upload allows you to upload truly a stunningly wide variety of files -- images, catalogs, spectra, MOCs, and more. (This capability used to be buried in the old "catalogs" tab.)

Note that very few of these search options have a 'cancel' anymore -- these blue tabs persist in all views, so to cancel anything, just go to another tab. Example: I'm viewing "Results", I think I'd like a new image, I go to "Images", then think better of it. I can't 'cancel' or use the 'back' button on the browser; I just go click on "Results" to return to where I was before. This may cause a bit of adjustment!

When you add more kinds of searches (see below), you add new blue tabs.

The Background Monitor and the help pages are still accessible from links in the upper right of the page, as they were before, but they no longer look just like these main blue tabs.

Basic Searching

Searching IRSA images and catalogs are very, very similar to what they were before.

The Images tab drops you into what used to be the IRSA Viewer start screen, where you can search for FITS or HiPS images. The coverage image may be more prominent than you remember it being -- it was there before, but it may be more obvious now than it was before. The coverage image is basically a way for it (and you) to keep track of where you are working on the sky.

The Catalogs tab drops you into what used to be the "search IRSA catalogs" tab. This just searches catalogs held here at IRSA, such as 2MASS, IRAS, WISE, etc.

Basic Results

After doing an image and/or catalog search (and/or uploading images and/or catalogs), you have a screen that looks similar to the old IRSA Viewer interface.

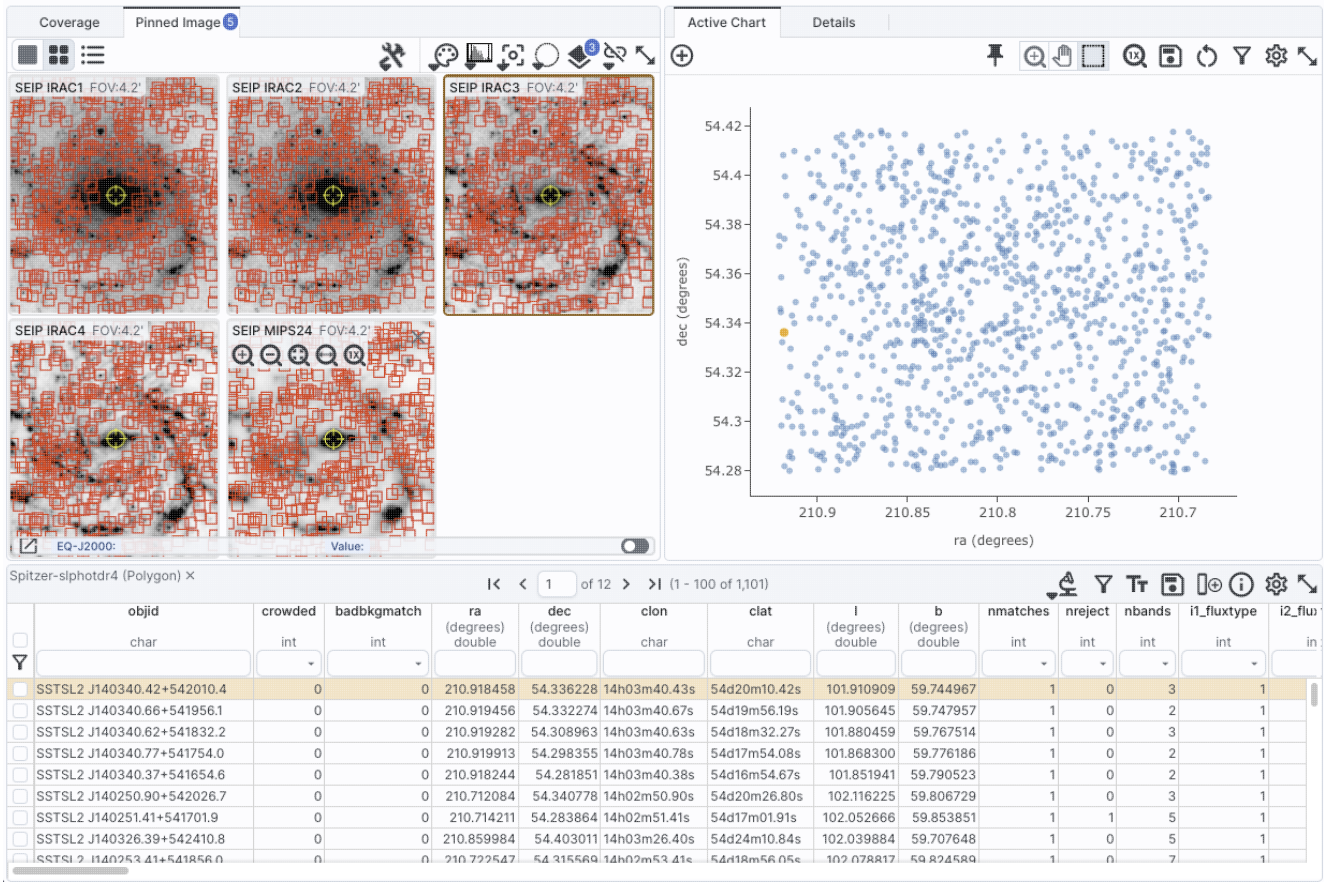

objid	crowded	badbkgmatch	ra (degrees)	dec (degrees)	clon	clat	l (degrees)	b (degrees)	nmatches	nreject	nbands	i1_fluxtype	i2_flu
char	int	int	double	double	char	char	double	double	int	int	int	int	in
SSTSL2 J140340.42+542010.4	0	0	210.918458	54.336228	14h03m40.43s	54d20m10.42s	101.910909	59.744967	1	0	3	1	
SSTSL2 J140340.66+541956.1	0	0	210.919456	54.332274	14h03m40.67s	54d19m56.19s	101.905645	59.747957	1	0	2	1	
SSTSL2 J140340.62+541832.2	0	0	210.919282	54.308963	14h03m40.63s	54d18m32.27s	101.880459	59.767514	1	0	3	1	
SSTSL2 J140340.77+541754.0	0	0	210.919913	54.298355	14h03m40.78s	54d17m54.08s	101.868300	59.776186	1	0	2	1	
SSTSL2 J140340.37+541654.6	0	0	210.918244	54.281851	14h03m40.38s	54d16m54.67s	101.851941	59.790523	1	0	2	1	
SSTSL2 J140250.90+542026.7	0	0	210.712084	54.340778	14h02m50.90s	54d20m26.80s	102.116225	59.806729	1	0	3	1	
SSTSL2 J140251.41+541701.9	0	0	210.714211	54.283864	14h02m51.41s	54d17m01.91s	102.052666	59.853851	1	1	5	1	
SSTSL2 J140326.39+542410.8	0	0	210.859984	54.403011	14h03m26.40s	54d24m10.84s	102.039884	59.707648	1	0	5	1	

From here, you can interact with the images, just like you could before, and interact with the tables and plots, just like you could before. "Pinning" is still a concept in use here -- that just means "hold onto this item within this tool". It doesn't mean "save to disk" (ours or yours), nor does it mean "download the data behind this"; it means "retain this item in this tool for now." You can pin images, tables, or plots, depending on the context. If you upload or extract spectra, then you can interact with them, just like you could before (though the newest features we have added are to the spectral capabilities!).

Downloading data works the same as it did before.

Side Menu: Adding Tabs

Now for the new stuff!

This icon  in the upper left pulls open a "drawer" from the left hand side.

You can use this side menu to add, or remove, or even just navigate among the blue tabs from the top of your IRSA Viewer interface. If you want to search more IRSA data using different tabs, add the "Data Collections" or "VO TAP" tabs. If you want to search other archives, add one of the tabs under "External archive search tabs." If you have too many tabs, click on the "Hide Tab" button to remove that corresponding tab.

Side Menu: Layout and Appearance

Yes, the oft-requested "Dark Mode" is finally here!!

The bottom of the side menu controls the appearance of the tool in your browser -- you can set it to run as light mode, dark mode, or respect whatever preferences you have set on your system.

The "Layout" portion of the side menu controls how the results tab displays the images, tables, and charts (plots). It gives you a stylized cartoon layout to pick from. (This control used to be in the upper right and was previously described tersely only in words, which may have been confusing.)

New Searching Options

OK, maybe not so much "new searching options" as "new ways of presenting some of the search options we had before, but maybe they were a little too buried or esoteric before."

"Data Collections" is a tab that works just like our Data Collection Explorer, which is (and was) a separate tool, but you can now also get results from that kind of a search right away into IRSA Viewer.

"VO TAP" is a TAP search that is limited to just IRSA services. (We had this before, too.)

NED Objects is a search of NED. (We had this before, too.)

Multi-archive VO TAP is, a general TAP search, which we had before but it was somewhat hidden. People complained, we heard you, and now this is back and more clearly accessible. You can use your own URL here, too.

Multi-archive VO Cone Search, which uses SCS. (We had this before; an oldie but a goodie.)

Finally, we now have something that is quite explicitly an VO ObsCore search, and it is limited to being just an ObsCore search, using CADC. Previously, to do this kind of a search, you had to go to the "Catalogs" tab, and it was confusing to be able to retrieve images, spectra, catalogs, and even services from a so-called 'catalog' search. So, now the ObsCore search is called out on its own as a separate entity, and the results of this search (by design) look an awful lot like a DCE search result page.

New Catalog Features

The newest, most prominent features to make an appearance in this tool all have to do with catalogs.

Parquet format

Parquet file format is a highly efficient, compressed, column-oriented format for tabular data that has been adopted by many recent wide area survey projects. It can enable faster searching for large tables. There are packages in astropy that handle parquet files. And now this tool can both read and write Parquet format!

Color swatches

If you have more than one catalog loaded, sometimes it is hard to remember which color in the image overlay corresponds to which catalog. Now, there is a little color swatch in the catalog header:

See this section for more details.

Hierarchical catalogs

It used to be the case that if you had >5000 sources, then the tool only *seemed* to be plotting all of the sources on the image(s); it really was only plotting a fraction of the sources. Now, it really is plotting all of the sources, but if there are >1000 sources, the tool will bin up the catalogs. In summary, the sky is broken up into sections, and the tool will show symbols with a number indicating the number of sources in that region. As you zoom in closer and closer, the tool will adjust those bins to smaller and smaller cells until it shows you individual sources. Go here for more information!

Questions, Feedback, Bugs

If you have questions, feedback, or bugs, please contact us via the IRSA Help Desk. If a bug, please include your operating system version, your browser software and version, instructions on how to reproduce it, and any specific error message you may have gotten.

IRSA Viewer: Upload

You can upload images and tables and even other kinds of files (like ds9 region files and MOC files!) to IRSA Viewer.

Contents of page/chapter:

Introduction

The "Upload" tab is one of the tabs that can appear on the top of the screen.

When you click on that tab, you get this:

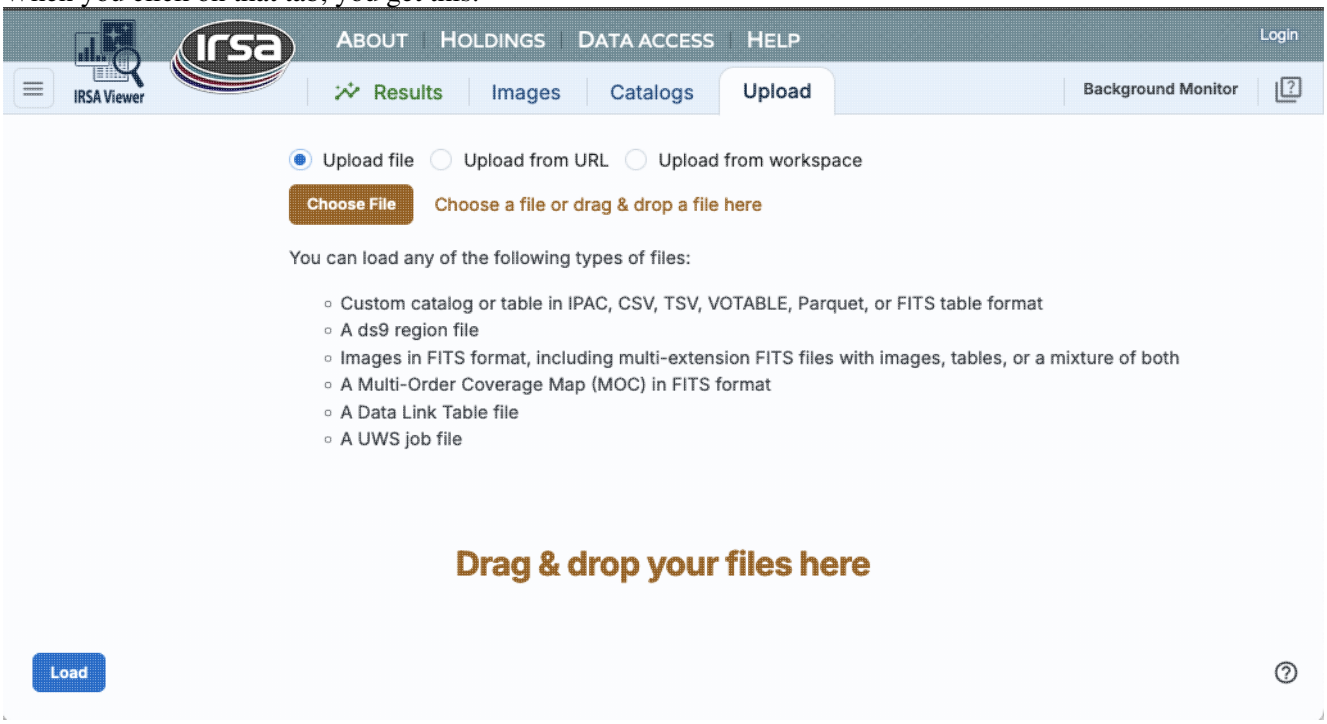

File Location

The file that you are uploading can come from your own disk, the web (type or paste in the URL), or from the IRSA Workspace. (Note that you need to be logged in to use the Workspace.)

Tips and Troubleshooting

- If you have multiple HDUs in your file, you can load a plane, and then come back to the the "Uploads" tab to pick another HDU without having to upload it again. Or, you can load all of them at once.

Catalogs

You may upload a file from disk (you can use the system browser to identify the file, or drag-and-drop them into this window), from the web via a URL, or from the IRSA Workspace.

The file can be in any of a number of formats, which we now briefly describe.

IPAC table format (*.tbl)

IPAC table format is plain text with a particular formatting. IRSA has a table reformatting and validation service which may be helpful, or you can download just about any catalog you find through IRSA, and mimic that format.

If you want it recognized as a catalog, your table file MUST have RA and Dec values, and unless it is specified, it assumes J2000. (See also "tips and troubleshooting", below.)

You can add a "SYMBOL" parameter to change the shape (X, SQUARE, CROSS, EMP_CROSS, DIAMOND, DOT) of catalog marks, e.g.:

```
\SYMBOL = X
```

You can add a "DEFAULT_COLOR" parameter to assign a CSS color name or a HEX value to catalog marks, e.g., either of these two:

```
\DEFAULT_COLOR = lightcyan
\DEFAULT_COLOR = #00FF00
```

You can find the CSS color code or the CSS color HEX values online.

Comma-separated values (CSV) format (*.csv)

CSV format is often accepted by spreadsheet programs, and most spreadsheet programs can output CSV. It is plain text, with values for each column separated by commas.

Tab-separated values (TSV) format (*.tsv)

TSV format is sometimes accepted (or generated) by spreadsheet programs, and sometimes in tabular data downloaded from journal articles from the 90s. It is plain text, with values for each column separated by tab characters.

VOTable (*.vot)

Virtual Observatory (VO) tables are a special case of XML tables. (All VO Tables are XML but not all XML are VO Tables.) It is a format developed by the International Virtual Observatory Alliance (more information here). XML tables initially look to the human eye like it might be HTML, but they are easier to parse with code. There are packages in astropy that handle VO Tables.

Parquet (*.parquet)

Parquet file format is a highly efficient, compressed, column-oriented format for tabular data that has been adopted by many recent wide area survey projects. It can enable faster searching for large tables. There are packages in astropy that handle parquet files.

FITS files (*.fits)

FITS tables can be loaded into this tool. Note that this tool is flexible enough to handle multiple header data units (HDUs), so that you can upload a FITS file that has both images and tables. You can specify which HDU you would like to load, and if you choose more than one image, it will give you a choice of loading them into individual frames or all into one frame. For more information on loading images this way, see below; there is another chapter entirely on images in this tool in general.

Nearly every file you load will result in a preview of the file you have uploaded. Here are some example previews of catalog uploads:

Uploading a CSV file:

Uploading an IPAC table file:

Uploading an xml file:

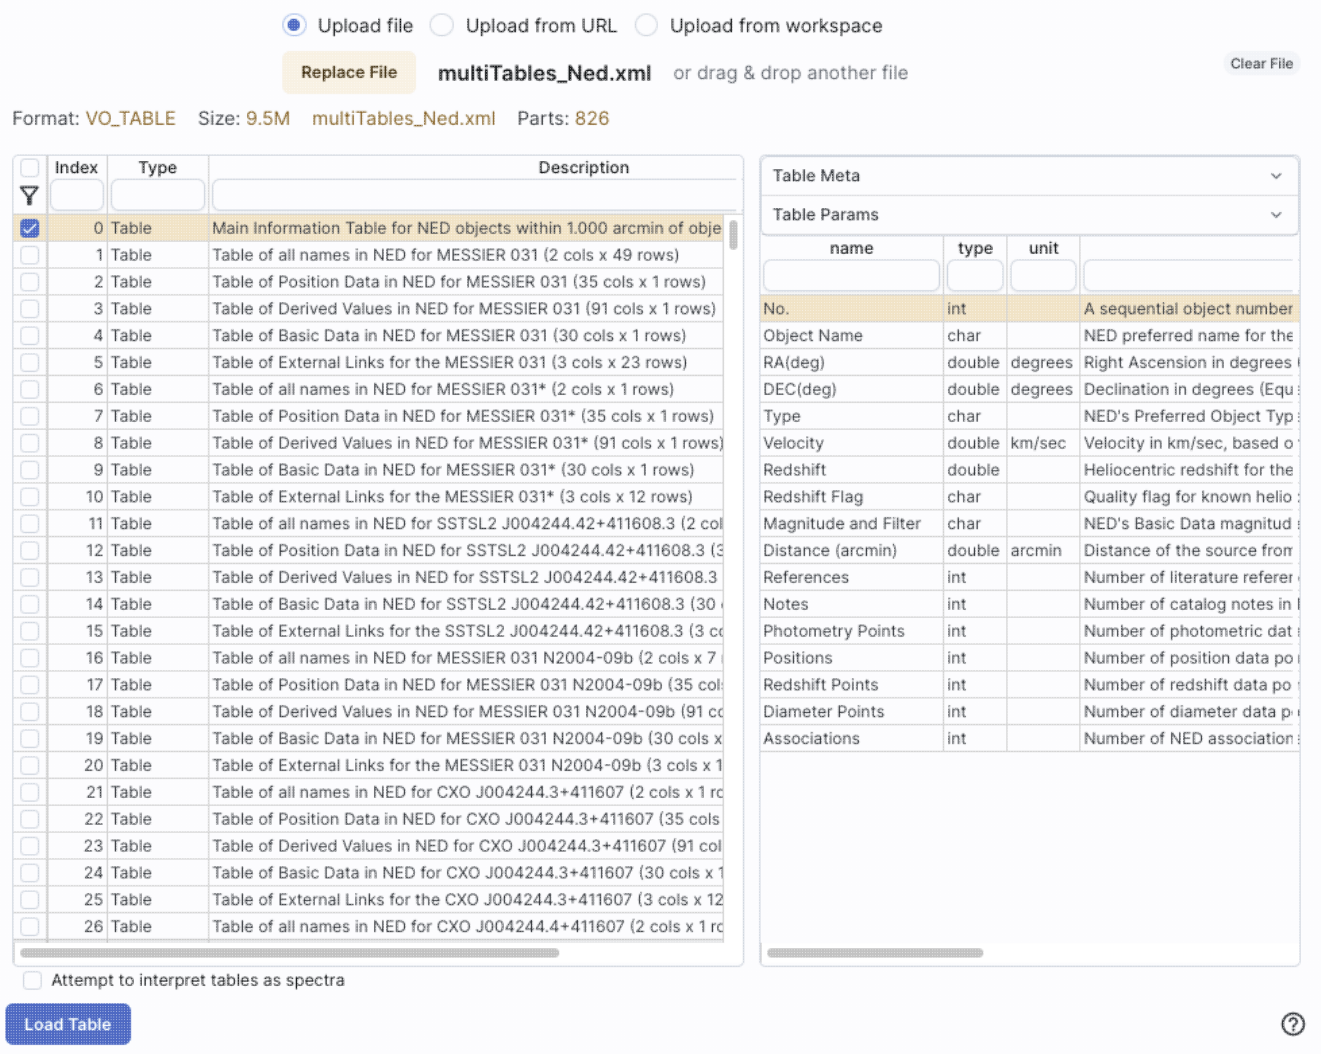

The tables are then shown and, if catalogs, interacted with in the same way as the other catalogs described here.

Tips and Troubleshooting

- If you would like to have your catalog overlaid on an image, it needs to have RA and Dec columns. If the tool doesn't seem to recognize your RA and Dec columns, check your formatting, or try headers of "ra" instead of "RA" and "dec" instead of "DEC" or "Dec". If you have columns like "_RA2000" and "_DE2000", it's going to be confused.
- Unless specified, the tool assumes any coordinates you give it are J2000.
- If there are no discernible positions in the uploaded file, it will still let you plot columns from the file after you've loaded it; it just can't overlay things on images in that case.

Region Files

DS9 is a popular program for visualizing FITS files. It uses a file format for storing image overlays called region files. IRSA Viewer can write and read DS9 region files. Usually, you would read regions files from the image toolbar, but you can also upload region files via this upload tab.

If you upload a region file here, the tool will just assume that you meant to upload it via the image toolbar and overlay it on your image(s). If you don't have an image loaded, it will warn you.

Tips and Troubleshooting

- If you want to overlay sources from a catalog, it is generally better to actually create a catalog (as above) and upload/overlay that, rather than a region file; you have more flexibility with what you can do, how you can interact with a catalog, than you can with a region file.

Images

You can upload single- or multi-plane FITS files into the tool, as well as multi-HDU FITS files. Again, it will give you a preview of what it thinks you are uploading. For multi-plane or multi-HDU files, you can select what portion(s) to upload.

Uploading a simple FITS file:

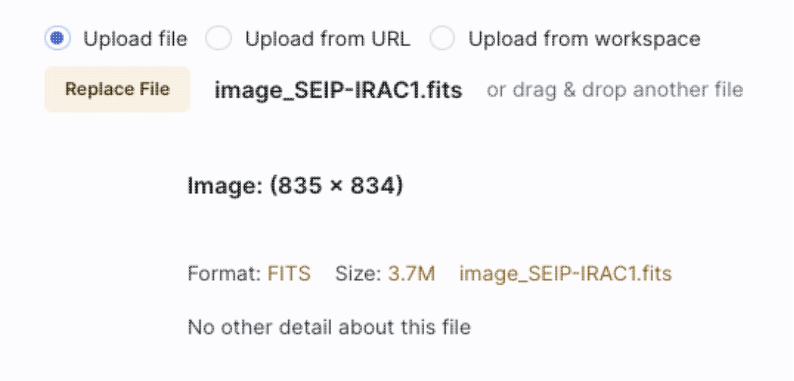

Uploading a FITS file with many image extensions, with just one selected:

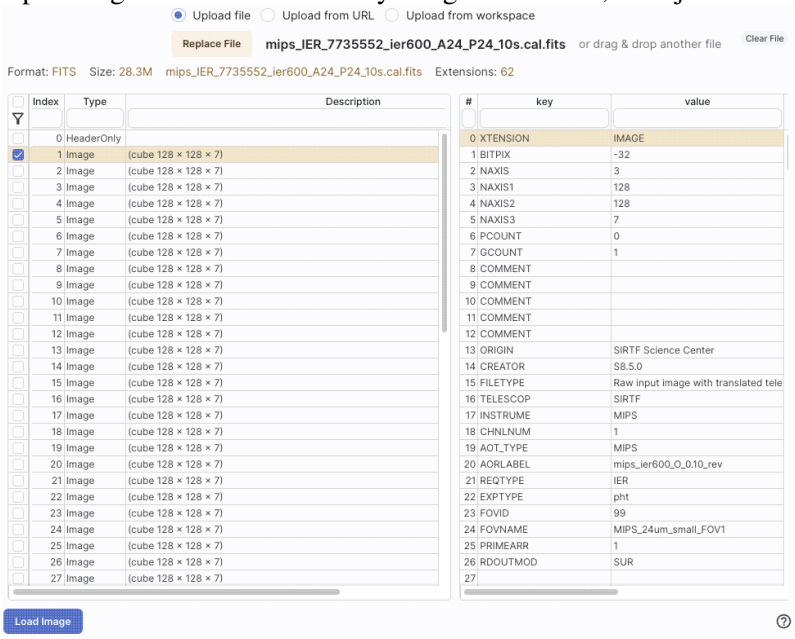

Uploading a FITS file with mixed extensions, with just one image plane selected:

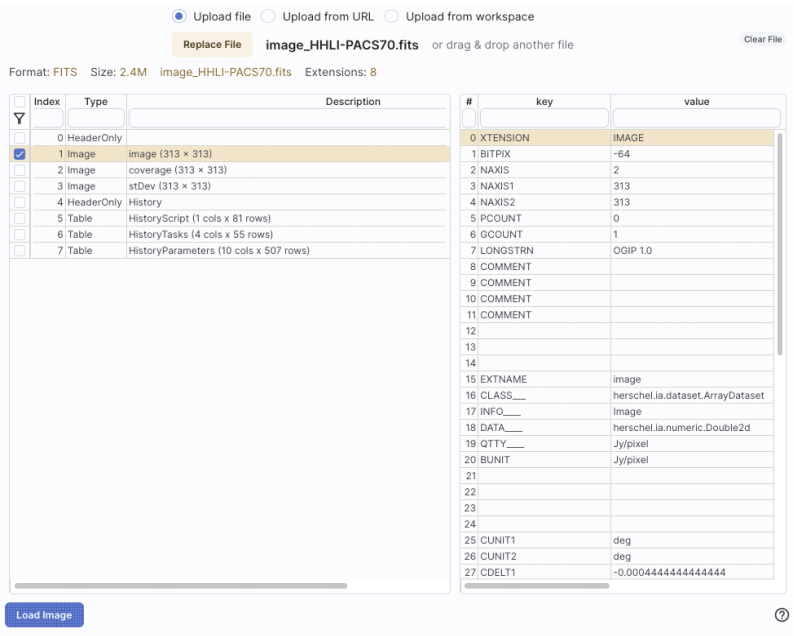

Uploading a multi-HDU FITS file containing a mixture of tables and images. (Note that in this case, all planes are selected and the lower left gives a choice for loading all the images into one window or one extension per window. Note also that it has the option of attempting to interpret tables as spectra.)

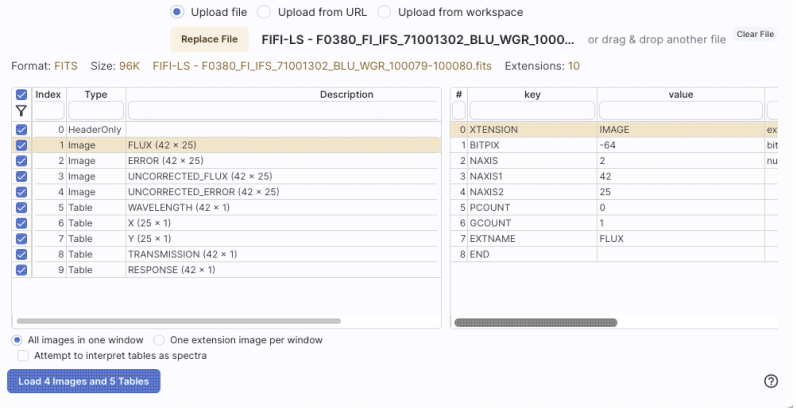

After you verify that the tool is reading your file correctly, and, if applicable, selected the HDU(s) you wish to load, click "Load" to load the file into the tool.

The images are then shown and interacted with in the same way as the other images described here.

Spectra

You can upload tables and tell the tool to attempt to load them in and treat them as spectra. Again, it will give you a preview of what it thinks you are uploading.

Uploading an IPAC tbl file that is a spectrum included, asking the tool to interpret the table as a spectrum:

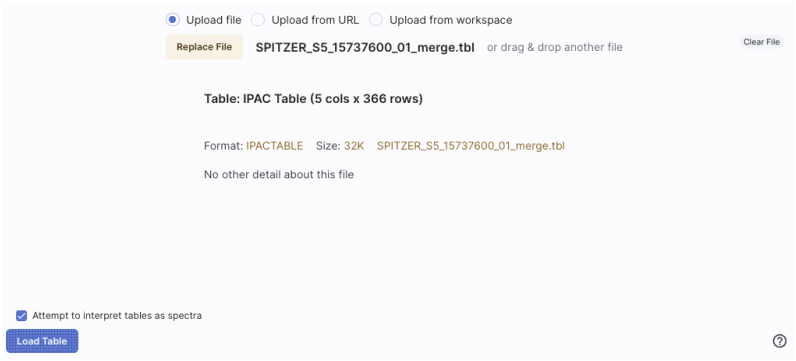

Uploading a FITS file that has a spectrum included, asking the tool to interpret the table as a spectrum:

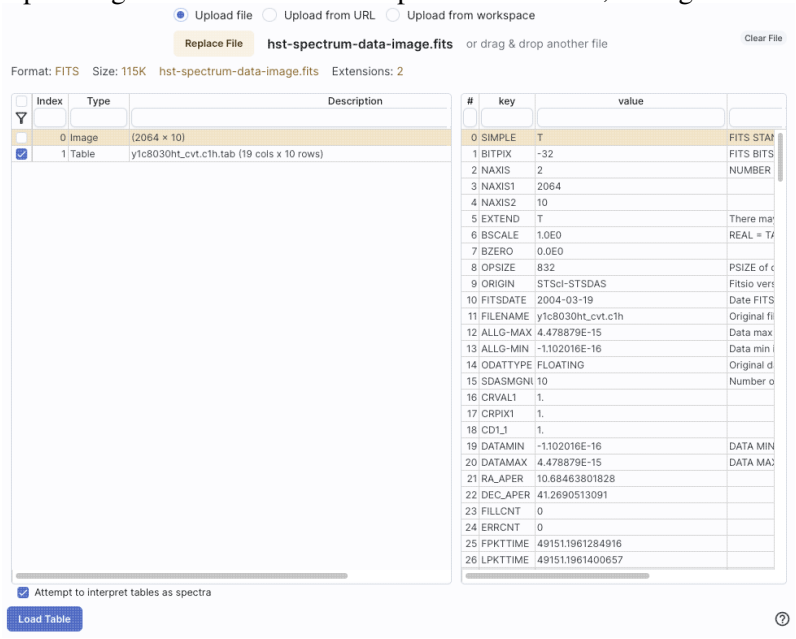

The spectra are then shown and interacted with in the same way as the other spectra described here.

MOC Files

Multi-order coverage map (MOC) files tell you where data exist (or don't exist). You can upload these kinds of files into this tool, and you can choose to view them as a table or as an overlay on HiPS files.

You can also load a MOC file from the HiPS/MOC menu once you have a HiPS image loaded.

A preview when uploading a MOC file :

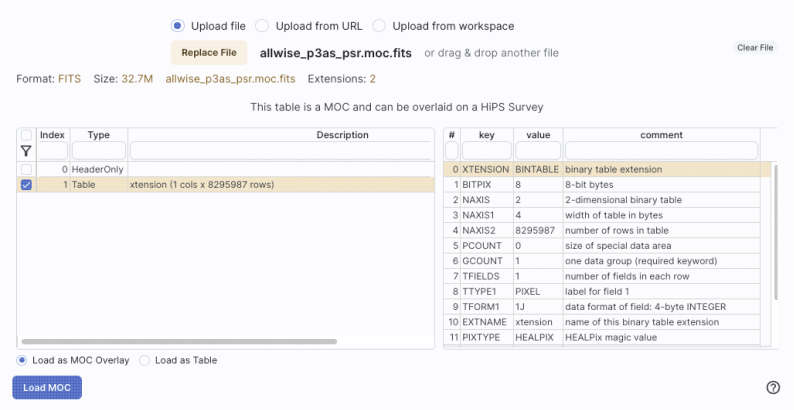

Data Link Files

A DataLink is a protocol developed by the International Virtual Observatory Alliance to specify more sophisticated linking of metadata and services to the data itself. You might use this kind of file to describe linkages to light curves or a light curve service from a single-epoch catalog. These kinds of files can also be loaded into this tool.

UWS Job Files

A Universal Worker Service (UWS) Pattern is a protocol developed by the International Virtual Observatory Alliance to manage asynchronous execution of jobs on a service. These kinds of files can also be loaded into this tool.

IRSA Viewer: Images

IRSA Viewer enables loading images as well as extensive interacting with images. This chapter covers loading images; visualization tools are covered in another chapter. IRSA Viewer can load images that are FITS and HiPS formats. Any catalogs you have loaded are overlaid on the images; see visualization chapter for more information.

Contents of page/chapter:

Searching for and Loading Images

When you first start IRSA Viewer, depending on how you get into it, an image may be pre-loaded. If the image is pre-loaded, you can start manipulating the image right away; see the Visualization section.

If the first screen you come to is an empty results page, then you need to decide what to load. To search for images, click on "images" at the top to begin an image search:

If you do not have an image pre-loaded, the default start position is in the position search for images. It is assuming by default that you want to load a FITS image from IRSA services, though you can also load a FITS image from disk or off another service on the web (see below). (Flexible image transport system, FITS, files are widely used in astronomy and are an easy way to store images.)

The search window looks like this:

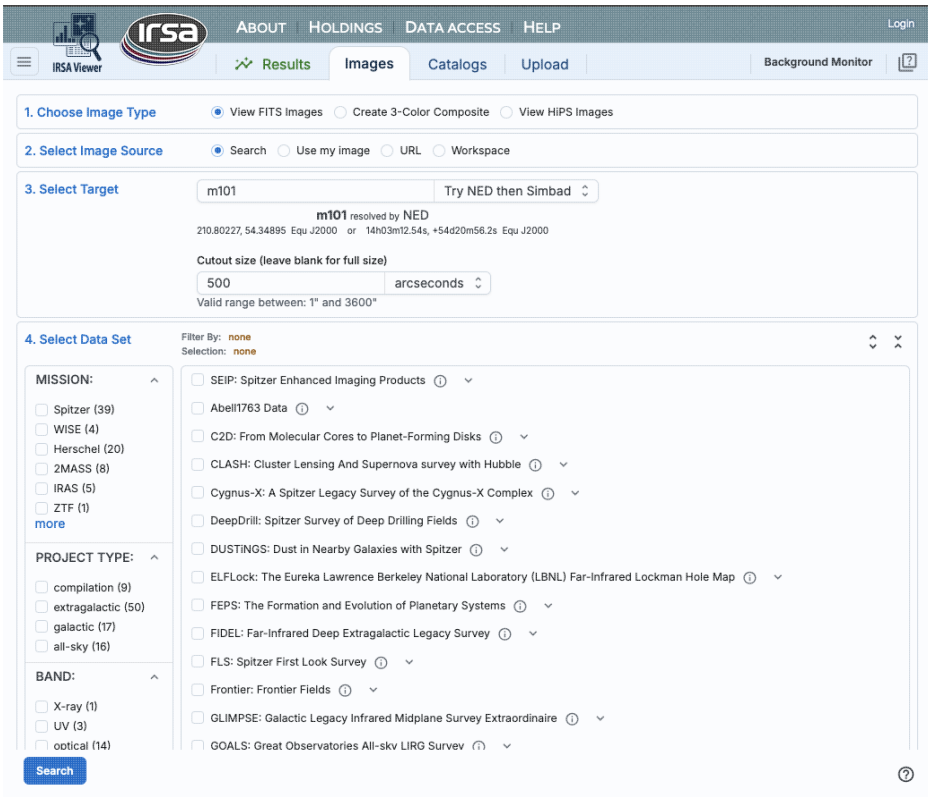

1. Choose Image Type

First, you select which images you want to load: FITS images (individually), FITS images that you load into a new three-color image (more on 3-color images below), or HiPS images (more on HiPS below; also see IVOA docs for more about what HiPS images are).

2. Select Image Source

Second, you select whether you want to pull an image from IRSA's archives ("Search"; see below), your own disk ("Use my image"), elsewhere on the web ("URL"), or the IRSA Workspace ("Workspace"). Note that to use the Workspace (reading from or writing to it), you'll need to log in.

In the cases other than "Search", nearly all of the additional options below this line vanish because they are no longer relevant. To **select an image off of your local disk**, select "Use my image", and then tell it where to find the image on your local disk. To load an **image from the web**, pick the "URL" option and enter the URL from which you want an image loaded. To load an **image from the IRSA Workspace**, pick the "Workspace" option and find the file you want to load.

If you would like to load an image from IRSA's archives, select "Search" and go on to these subsequent additional search parameters.

3. Select Target

Third, you select a target. You may enter a target name, and have either NED-then-Simbad or Simbad-then-NED resolve the target name into coordinates. Alternatively, you may enter coordinates directly. These coordinates can be in decimal degrees or in hh:mm:ss dd:mm:ss format, or Jhhmmss+ddmmss format. By default, it assumes you are working in J2000 coordinates; you can also specify galactic, ecliptic, or B1950 coordinates as follows:

- '46.53, -0.251 gal' means 46.53, -0.251 degrees in galactic coordinates
- '12.7, +4.3 ecl' means 12.7, +4.3 degrees in ecliptic coordinates
- '19h17m 11d58m b1950' means 19h17m 11d58m in B1950 coordinates

- ◊ a source name like 'J140320.67+542028.6' is parsed as 14h03m20.67s +54d20m28.6s.
- ◊ a source name like 'G102.0360+59.7715' is parsed as 102.0360 +59.7715 in galactic coordinates

Examples are given below the text entry box before you start typing in the box.

As you are completing a valid coordinate entry, it echoes back to you what it thinks you are entering. Look just below the box in which you are typing the coordinates to see it dynamically change.

Below the box where you enter the target, you can then specify the size of the images you want. You may enter the cutout size in arcseconds, arcminutes, or degrees; just change the drop-down option accordingly.

Tips and Troubleshooting

- ◊ Pick your units from the drop-down first, and then enter a number; if you enter a number and then select from the drop-down, it will convert your number from the old units to the new units.
- ◊ There are both upper and lower limits to your search size; it will tell you if you request something too big or too small. Note that these limits may be image-dependent; larger images may be available from certain surveys and smaller images may be available from other surveys.
- ◊ **If you want the whole image tile**, just leave the image size blank, then the closest image tile in its entirety will be returned.

4. Select Data Set

Fourth, you select the data set. There are myriad choices, which you can filter in various ways to allow you to find what you need. Statistically, any one spot on the sky will only be found in a few of these data sets, so it makes sense to weed down the list, at least a little bit.

On the left hand side, you can filter by :

Mission (or survey)
: Spitzer, WISE, Herschel, 2MASS, IRAS, ZTF, PTF, AKARI, DSS, SDSS, MSX, COSMOS, MUSYC, BLAST, IRTS, BOLOCAM.

Project Type
: Compilation (meaning, e.g., all the data available from a mission that was not an all-sky survey), extragalactic, galactic, all-sky.

Band
: X-ray, UV, optical, near-IR, mid-IR, far-IR, mm, radio. (Yes, you can find data from non-IR missions/surveys here, depending on what projects have delivered back to IRSA.)

The number in parentheses after each type is the number of data sets in that category.

To expand or contract the options below each of these broad categories, click on the black arrow on the right. To see all of the choices in each of these broad categories, click on the black arrow and then click "more" near the bottom of the list (after which you can click "less" to collapse it again). In this example, "Mission" is collapsed, "Project Type" is expanded, and "Band" is fully extended to reveal more options than are shown by default:

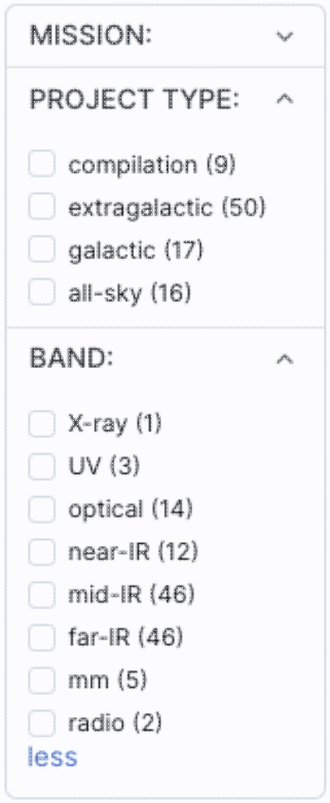

To select any of the options, click on the checkbox on the left. In response to your selections, two things happen. (1) checked items appear in highlights above this part of the page, next to "Filter By"; (2) the list of programs for selection on the right hand side changes. Here is an example where the list has been filtered down to just include WISE data:

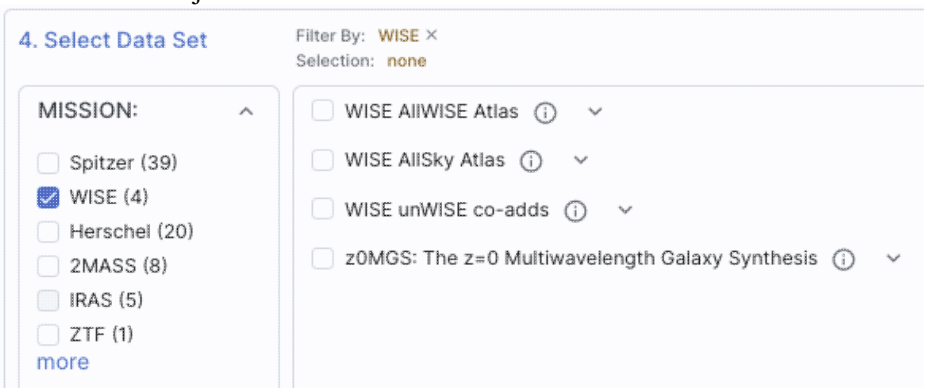

On the right hand side, you can select individual surveys and individual bands therein. To expand or contract the options below each of the categories, click on the black arrow on the right, next to the data set name. To select any of the waveband options, click on the checkbox on the left of the individual survey or individual bands. Here is an example showing all of the WISE/AllWISE data selected (via the tickbox next to "WISE AllWISE"), and just W2 from WISE/AllSky selected (via the tickbox next to "W2" in the expanded section under "WISE AllSky"). Note the "Selection" indication at the top.

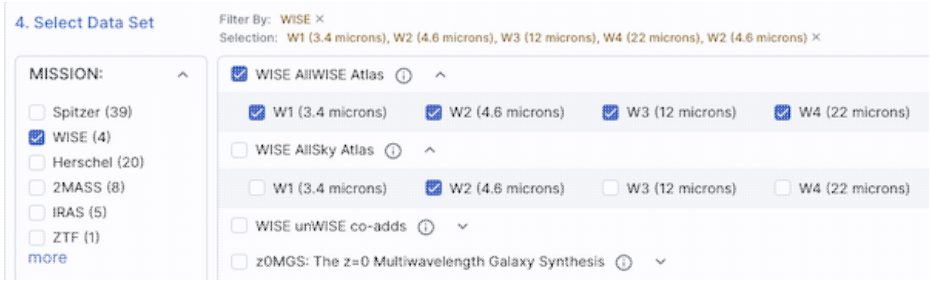

To find out more information about any given data set, click on the i in the circle  . This takes you to a master list of all data sets available in IRSA Viewer, from which you can obtain standard information about the data sets (mission, wavelengths, links to more information about the program or delivery, and more).

Search!

To actually initiate the search as specified, choose the "Search" button in the lower left.

Tips and Troubleshooting

- You don't have to select something on the left side before selecting something on the right side. If you know exactly what you want, just jump in and select things on the right and click "search."
- If you select something on the left side, it will limit your choices accordingly on the right side.
- 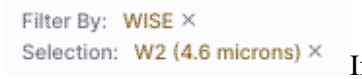

If you clear filters on the left side, it doesn't affect selections you've already made on the right. You must actually clear all the selections on the right to reset everything. The most efficient way to do this is by clicking the 'x' next to the selection summary at the top of this portion of the window -- the 'x' at the end of the "Filter By" line clears the filters on the left, and the 'x' at the end of the "Selection" line clears the selections.
-  If you want to expand all the choices on the right (to ease in band selection), click on the sets of arrows in the upper right of this part of the screen to show all options, or collapse them again.
- If you go back to add new images of the same target, be sure to uncheck the images you had previously selected, otherwise you will load in second copies of those images.
- To remove an image you have already loaded (or an error message from a data set that had no data covering your target), click on the small blue "x" in the upper right of the corresponding image tile.
- Most images that are returned will be cropped down to your requested size (if you entered a size). Some, however, cannot easily be cropped at the moment. Those are delivered full-size, and you can crop them down separately if you want.
- For more information on the data sets included in IRSA Viewer, see this list.
- IRSA Viewer will NOT return ALL the images for your specified combination of dataset+band+position; it returns the most centered science image it can find out of the bands it can access. This is usually but not always actually going to be the most useful to you, depending on what you're trying to do. Some programs delivered PSFs, mosaics created with alternate algorithms, etc. To find all the data at IRSA for a given target, you have several options.
    1. Enter the desired position in the big search box on the IRSA home page (this accesses a tool called Data Discovery).
    2. Navigate specifically to the page corresponding to your desired data set (e.g., see the pages at IRSA linked from this list) and explore the data there. Depending on the data, you may be able to wander through directories, download them all at once in a tarball, or search by position.
    3. Use the "Search Collections" option here within IRSA Viewer, or explicitly go to the separate tool called Data Collection Explorer. Note that as of this writing, not all of our data are yet

available in this tool.

Example of going outside of IRSA Viewer: I want data covering M16. M16, the Eagle Nebula, is essentially in the galactic plane. Thus, I don't need to search really any of the extragalactic data sets. I select on the left hand side, under project type, compilation, galactic, and all-sky. I am still left with a long list of programs, though it is much shorter than it was before. However, I can omit some programs from consideration based on their names -- programs studying Taurus or Orion are unlikely to have any data on M16. I can then individually select the dozen or so data sets likely to cover M16. I explore the images and decide that the Spitzer images from the GLIMPSE project are what I would like to explore in more detail. I can then navigate to the GLIMPSE page, read the documentation to find that they provided mosaics with two different pixel sizes, and click on the "DCE" link at the top of that page. I can search by position there to obtain the mosaics with different pixel sizes.

Much more detail about interacting with images can be found in the Visualization section.

Making 3-color Images

You can create 3-color images directly from the image search. Select "Create 3-Color Composite" from the top row of options. The rest of the window changes to look like the following:

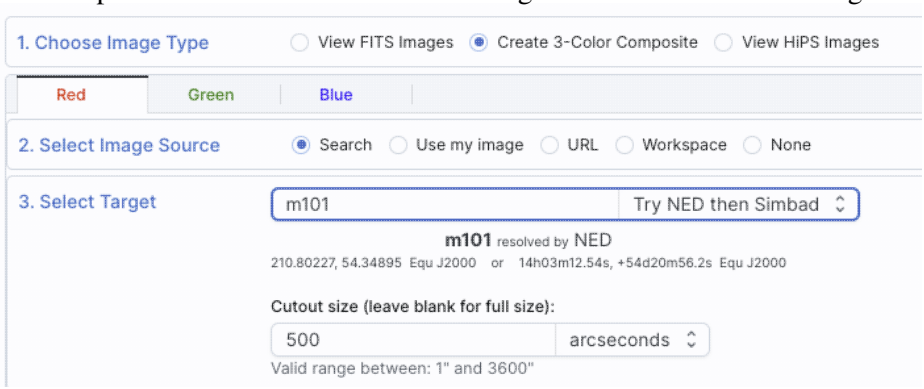

By default, you can select the red plane first; you then populate that color plane with all the same choices as you would have for a single channel image (as above). To set the additional color planes, click on "green" and then "blue" to populate those planes accordingly.

It assumes that you must want the same position for all three color planes.

Select your options individually for each color plane (red, green, blue), and click 'Search' in the lower left. To exit the search window (i.e., cancel) without creating a new 3-color image, click on any other tab at the top, e.g., "Results" returns you to the results you have already loaded into the tool.

To change the color stretch of each color plane individually, click on the "Color Stretch" icon in the toolbox on the top of the images pane; see the Visualization section. Much more detail about interacting with images can be found in the Visualization section.

Tips and Troubleshooting

- You load all three images at once, e.g., you do NOT pick red, click search, then go back and pick green, click search, then go back and pick blue, click search. Instead, click red and define what you want for

that image, then go to green and do the same, then go to blue and do the same. Don't click "search" until you have specified all three bands.
- The images will be downsampled to the resolution of the red image. If you, say, load an MSX image into the red plane, a WISE image into the green plane, and a 2MASS image into the blue plane, all of the images will have MSX-sized pixels. If you load a WISE image into the red and green planes, and a 2MASS image into the blue plane, the images will have WISE-sized pixels.

HiPS Images: General Information & Definitions

HiPS stands for hierarchical progressive surveys, and these kinds of images are multi-resolution HEALPix images (where HEALPix stands for Hierarchical Equal Area isoLatitude Pixelation). (Also see IVOA docs on HiPS.) In practice, what this means is that you can interact with images of a very large chunk of sky, and as you zoom, the pixel size changes dynamically. HiPS images are fundamentally different than FITS images, and as such, what you can do with the HiPS images are different than what you can do with the FITS images.

The whole point of HiPS images is to provide on-demand resolution changes. Zoom out, and it loads large pixels. Zoom in, and it loads smaller pixels. HiPS images are designed to cover large areas of sky efficiently. If you need to visualize many degrees, this is the image type to use.

There are HiPS images from all over the world available via the web; the complete list of HiPS images available from the images search page includes (once the "IRSA Featured" checkbox is unchecked) many HiPS images from CDS.

HiPS images typically have the color and stretch set by the person who originally made them. The color table can be remapped within this tool. But, you cannot, in general, change the stretch of HiPS images with as much flexibility as you can with FITS images. This is why there may be multiple versions of some data sets in the list of HiPS images.

HiPS maps typically come with a **Multi-Order Coverage map (MOC)**. A MOC is a format developed by the International Virtual Observatory Alliance to specify sky regions. In this context, a MOC tells you via a simple boolean yes/no, is there sky coverage from this data set in this region. You can overlay a MOC from one data set onto a completely different data set's HiPS image.

Searching for HiPS Images

Select "View HiPS Images" from the top row of options. The rest of the window changes to look like the following:

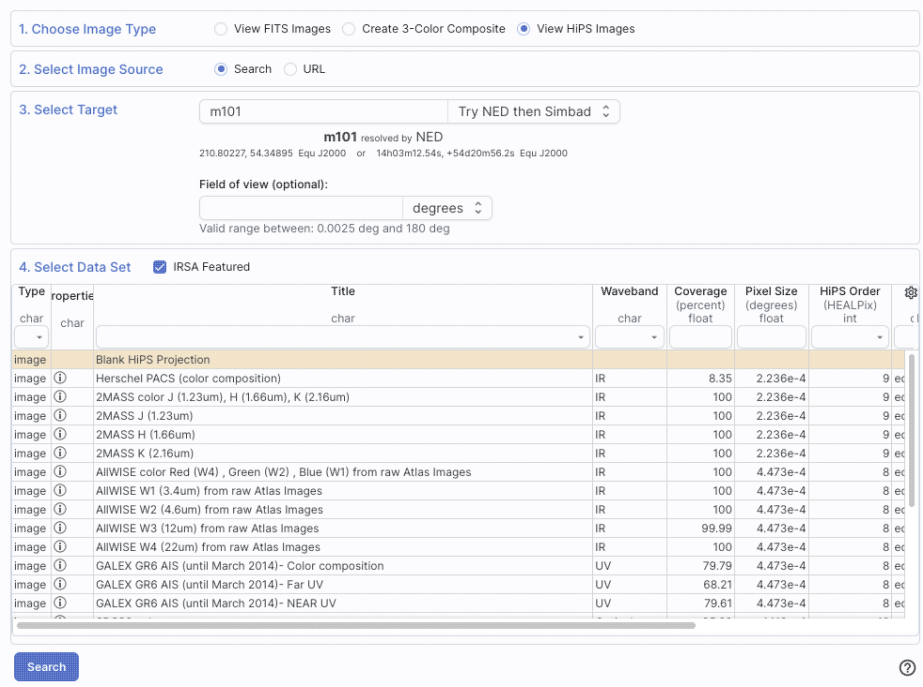

1. Choose Image Type

Image type should be "View HiPS Images".

2. Select Image Source

Second, you tell it whether you want to search among the HiPS images available to this tool, or give it a URL for it to construct a search. Either option retains the next section.

3. Select Target

Third, give it a target; this works as for FITS images above.

Note that here, in the context of HiPS images, the field of view (FOV) here has different limits than FITS images; it can be no smaller than 0.0025 deg, and can be up to 180 degrees. If you leave this blank, it will simply load the entire HiPS image for you.

4. Select Data Set

Fourth, select a data set. By default, the list of possible choices is limited (via the checkbox at the top of the table) to "IRSA Featured" choices, meaning data that IRSA users are most likely to want. If you uncheck this box, you will have a much larger list of HiPS image choices (from CDS ⧉) to pick from.

Column definitions. The table that appears when selecting a HiPS image has several columns:

◊ Type - image (all that is available in this context)
◊ Properties - links to more information for that image
◊ Title - descriptive words for the HiPS image
◊ Waveband - Approximate wavelength range
◊ Coverage - Approximate sky coverage (100% or less) -- HiPS images often carry with them something referred to as "MOC", or multi-order coverage. This number gives an indication of the sky coverage of the data.
◊ Pixel scale - At the highest order (zoomed in the most), this is the pixel size
◊ HiPS Order - HiPS order, e.g., how deep you can zoom
◊ Frame - coordinate sytem, e.g., equatorial, galactic, etc.
◊ Release date - Date that HiPS image became available
◊ Dataset IVOA ID - Unique (worldwide) identifier for the dataset

Search!

To actually initiate the search as specified, choose the "Search" button in the lower left.

Tips and Troubleshooting

- The table of HiPS choices is itself a Firefly table, like all the other tables in this tool, so you can sort/filter/etc. to locate the image you want to load.
- To learn more about any given image, click on the ⓘ in the second column of the table. Another window will spawn with basic information about that HiPS image.
- Note that color HiPS images are available and that the color stretch is set by the person making the HiPS image originally (so you can't really change it). You can change the color table (but not the stretch); see the Visualization section.
- A blank HiPS image (basically a blank canvas) is also available from the list of HiPS images.

Adding New Images

Additional images can be added at any time by clicking on the blue "Images" tab near the top.

You can use the same target as before, or change the target. To exit the search window (i.e., cancel) without getting a new image, click on any other tab at the top, e.g., "Results" returns you to the results you have already loaded into the tool.

Tips and Troubleshooting

- 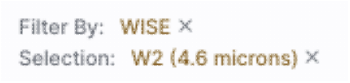

After you have done an initial search on a set of images, when you go back to the Images tab, *those same images are still selected.* If you just click 'search' again, you'll get a second copy of all of the images you initially selected. You must actually clear all the selections on the right to reset everything. The most efficient way to do this is by clicking the 'x' next to the selection summary at the top of this portion of the window -- the 'x' at the end of the "Filter By" line clears the filters on the left, and the 'x' at the end of the "Selection" line clears the selections.

Coverage Image

If you have launched IRSA Viewer by loading anything other than just images, it will also provide for you a "coverage image", which is basically a way for it (and you) to keep track of where you are working on the sky. This may be the most surprising when you are loading a catalog and therefore don't expect images to appear.

You may have a coverage image like this -- this is an example of a coverage image showing two different overlaid catalogs.

Note that in this case, the catalogs cover the whole sky, so the HiPS image is displayed in HiPS/Aitoff projection to show

the whole sky.

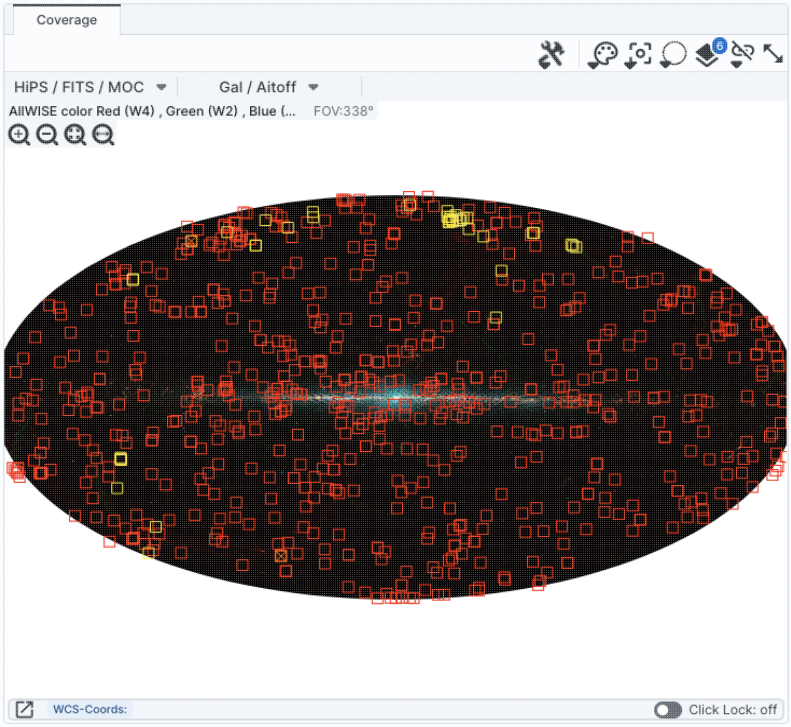

You can also have a coverage image like this, which shows a coverage image that has polygons demonstrating the coverage of each a list of loaded images (large polygons) and the locations of spectra (small squares). This case only has data over a relatively small region, so the HiPS image is zoomed in comparatively tightly on the relevant region.

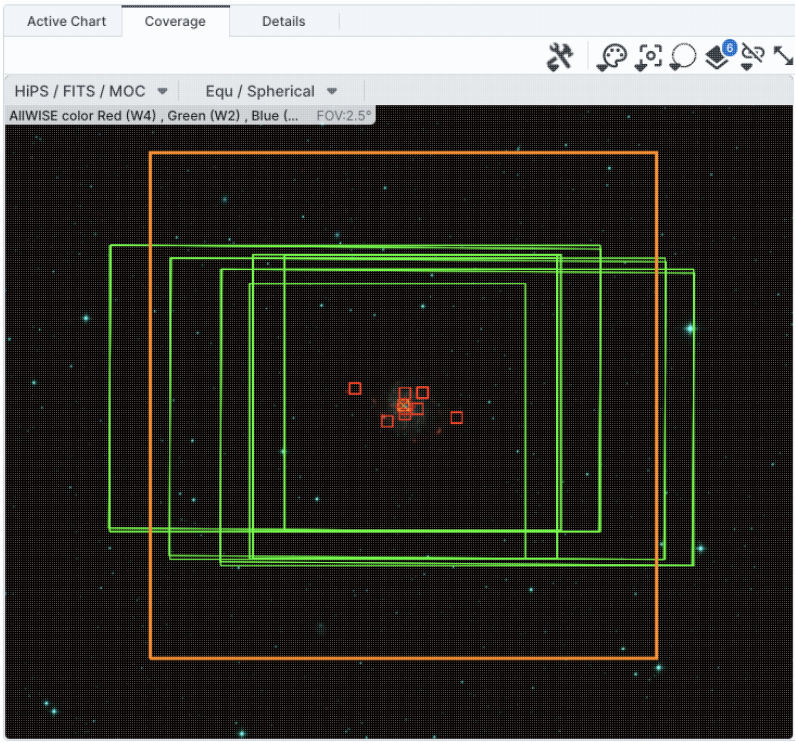

You might have a coverage image something like this, which shows a coverage image with a complex catalog overlaid. (This case still only has data over a relatively small region, so the HiPS image is zoomed in comparatively tightly on the

relevant region.) In this case, the catalog is more than 1000 sources, so the tool has shown you individual sources in the HEALPix cells where there are fewer sources than a given threshold, and where there are more sources, it simply shows you the cell and the total number of sources in that cell. The tool is rendering the catalogs in a hierarchical fashion, similar to how HiPS images work. You can control what this threshold is and how it renders the cells from the layers pop-up.

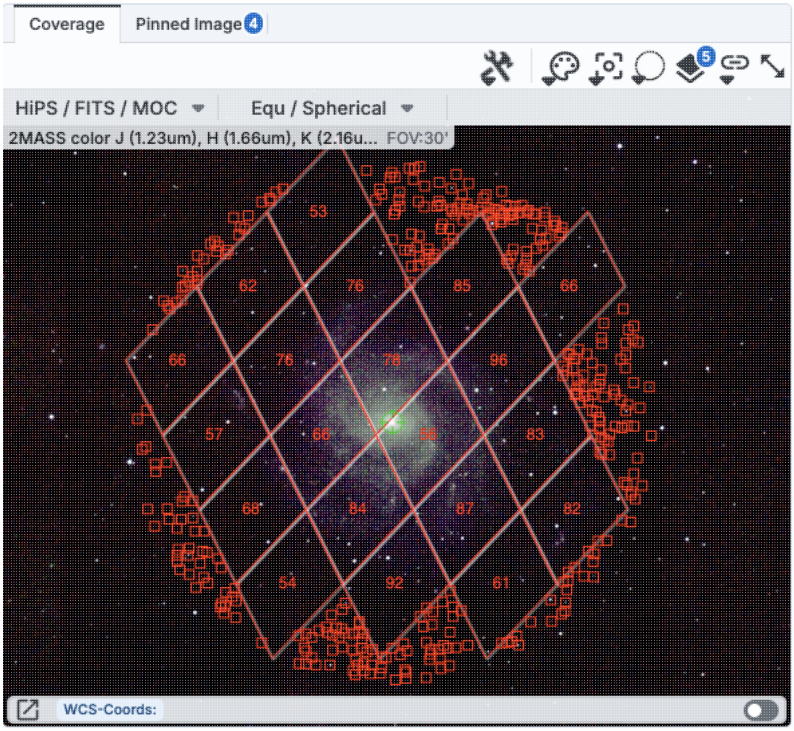

In all cases, you can interact with the coverage image in pretty much exactly the same way as you would any other image loaded into this tool; see the the visualization chapter for much more about those tools.

The thing that makes a coverage image a little bit different, however, is that it can automatically adapt, even beyond what a HiPS image can intrinsically do. The next subsection includes more details about how it can automatically change to accomodate your needs and zoom level.

Upper Left HiPS menus

In the upper left of the coverage image, there are two drop-down menus.

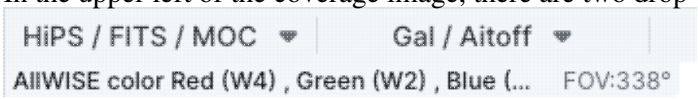

The first drop-down menu looks like this.

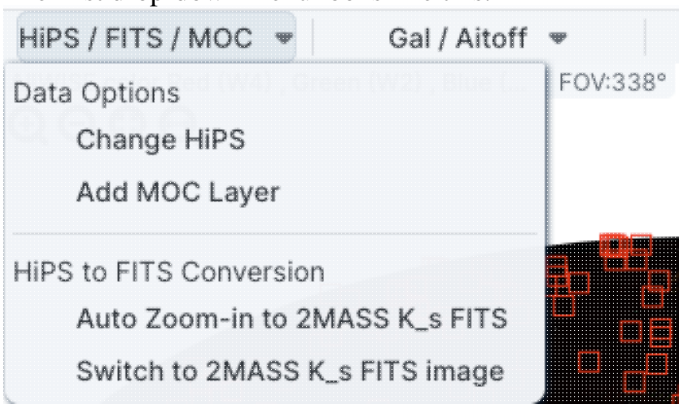

There are two sections here.

Under "Data Options", you can change what data are shown.

Change HiPS: Changing HiPS images

The choices made by any given creator of a HiPS image may result in any particular region being saturated or too faint to see. Color stretches are set by the creator of the HiPS map and cannot be changed; color tables can be somewhat changed by this tool, but that may be insufficient for your needs. If the HiPS map as shown does not suit your needs and you wish to change the HiPS image, click on the "HiPS/MOC" menu, and then click "Change HiPS". It brings up this pop-up, which resembles the HiPS search above:

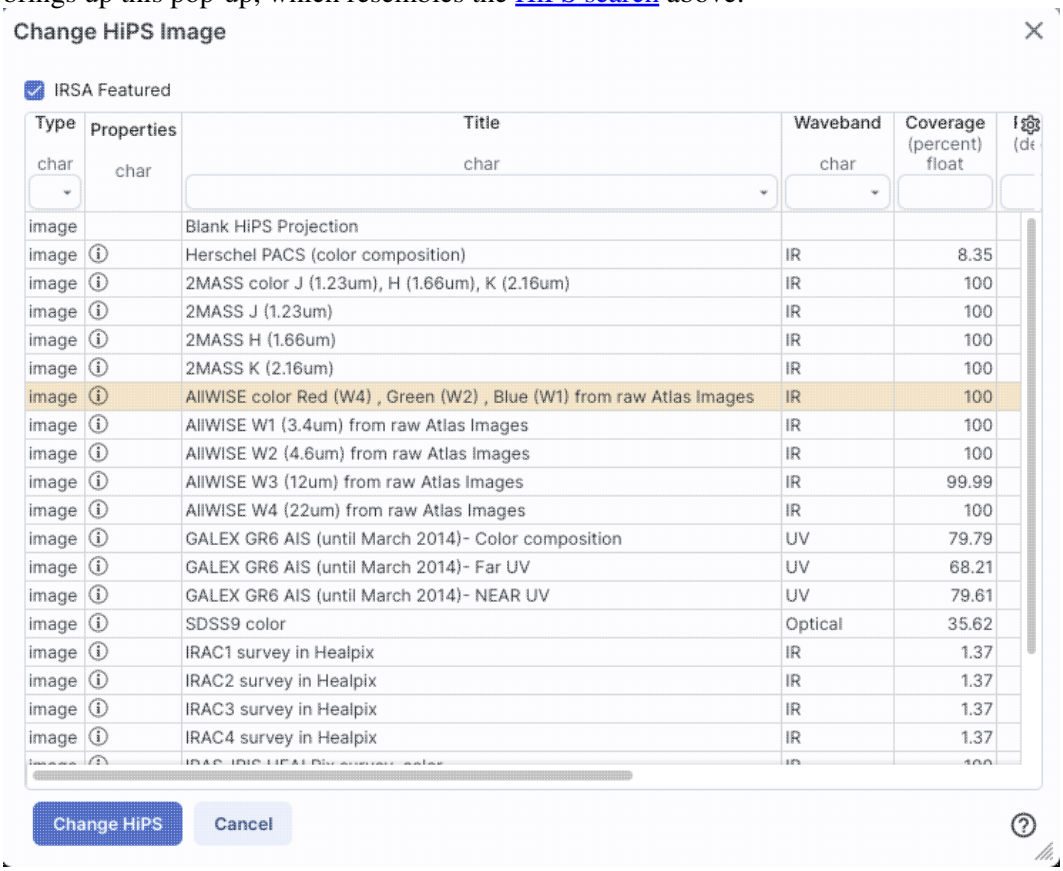

Type char	Properties char	Title char	Waveband char	Coverage (percent) float
image		Blank HiPS Projection		
image	ⓘ	Herschel PACS (color composition)	IR	8.35
image	ⓘ	2MASS color J (1.23um), H (1.66um), K (2.16um)	IR	100
image	ⓘ	2MASS J (1.23um)	IR	100
image	ⓘ	2MASS H (1.66um)	IR	100
image	ⓘ	2MASS K (2.16um)	IR	100
image	ⓘ	AllWISE color Red (W4) , Green (W2) , Blue (W1) from raw Atlas Images	IR	100
image	ⓘ	AllWISE W1 (3.4um) from raw Atlas Images	IR	100
image	ⓘ	AllWISE W2 (4.6um) from raw Atlas Images	IR	100
image	ⓘ	AllWISE W3 (12um) from raw Atlas Images	IR	99.99
image	ⓘ	AllWISE W4 (22um) from raw Atlas Images	IR	100
image	ⓘ	GALEX GR6 AIS (until March 2014)- Color composition	UV	79.79
image	ⓘ	GALEX GR6 AIS (until March 2014)- Far UV	UV	68.21
image	ⓘ	GALEX GR6 AIS (until March 2014)- NEAR UV	UV	79.61
image	ⓘ	SDSS9 color	Optical	35.62
image	ⓘ	IRAC1 survey in Healpix	IR	1.37
image	ⓘ	IRAC2 survey in Healpix	IR	1.37
image	ⓘ	IRAC3 survey in Healpix	IR	1.37
image	ⓘ	IRAC4 survey in Healpix	IR	1.37

Things to note:

- To select a new image, click on the row corresponding to the new HiPS image you want, and click "Change HiPS" on the bottom left.
- To cancel without selecting a new image, click "Cancel."
- This is an interactive table, so all the filtering and column manipulation tools apply here too. You can filter down the columns to find the image you want to use.
- To learn more about each HiPS map, click on the i with the circle in the second column. It will spawn another window with standardized information about the HiPS map.
- By default, it shows HiPS maps corresponding to IRSA data collections. To see a more comprehensive list, unclick the box marked "IRSA Featured."

Add MOC Layer: Adding a MOC Overlay

It is often useful to see what other data are available. Looking at a Herschel/PACS far-IR HiPS map? Overlay a Chandra MOC to see what X-ray data might be available to go with your IR data. Click on the "HiPS/MOC" menu, and then click "Add MOC Layer." It brings up this pop-up:

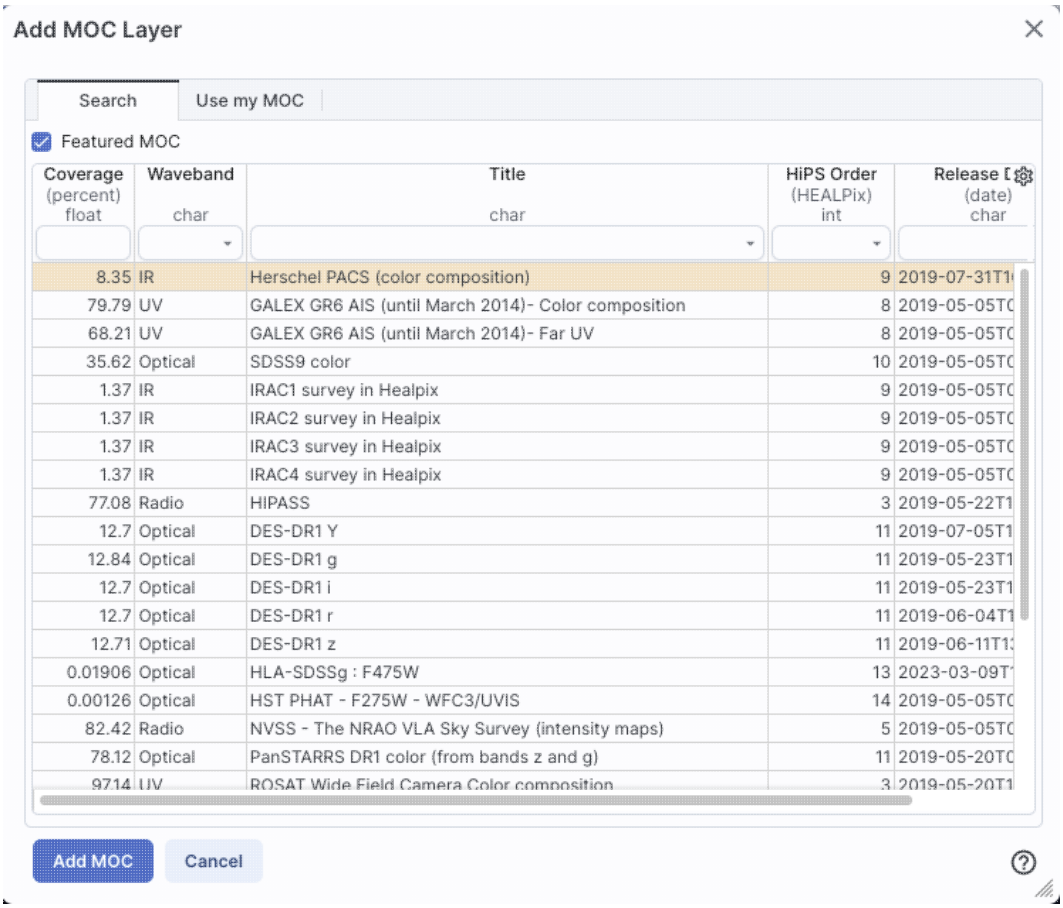

Things to note:

- Coverage (the first column) is the fraction of the sky covered by that MOC.
- To select a new MOC, click on the row corresponding to the new MOC you want, and click "Add MOC" on the bottom left.
- To cancel without selecting a new image, click "Cancel."
- This is an interactive table, so all the filtering and column manipulation tools apply here too. You can filter down the columns to find the image you want to use.
- By default, it shows MOCs likely to be most interesting for IRSA users. To see a more comprehensive list, unclick the box at the top left.
- If you have your own MOC FITS file, you may upload it via the "Use my MOC" tab on the top center.
- If you try to upload a MOC FITS file via the upload tab, it will behave as if you have uploaded it here.

Under "HiPS to FITS Conversion", you can control whether the viewer will automatically toggle between image types as needed.

By default, the coverage image is most likely a HiPS image. FITS images are best for small regions of the sky, and HiPS images are best for large regions of sky.

Auto Zoom-in to 2MASS K_s FITS:

If you select this, then when you zoom in very close to a target, it will automatically convert to a FITS image when you get close enough.

Switch to 2MASS K_s FITS image:

If you select this, then it will jump directly to a FITS image centered on the currently selected target.

Note that if you swap between HiPS and FITS and back again, it will include a region on the HiPS image that is the footprint of the FITS images you had just loaded. A label appears at the center of that footprint, which may be disconcerting if you are not zoomed out enough to see the region itself. Here is an example, zoomed out so it is more clear what is going on:

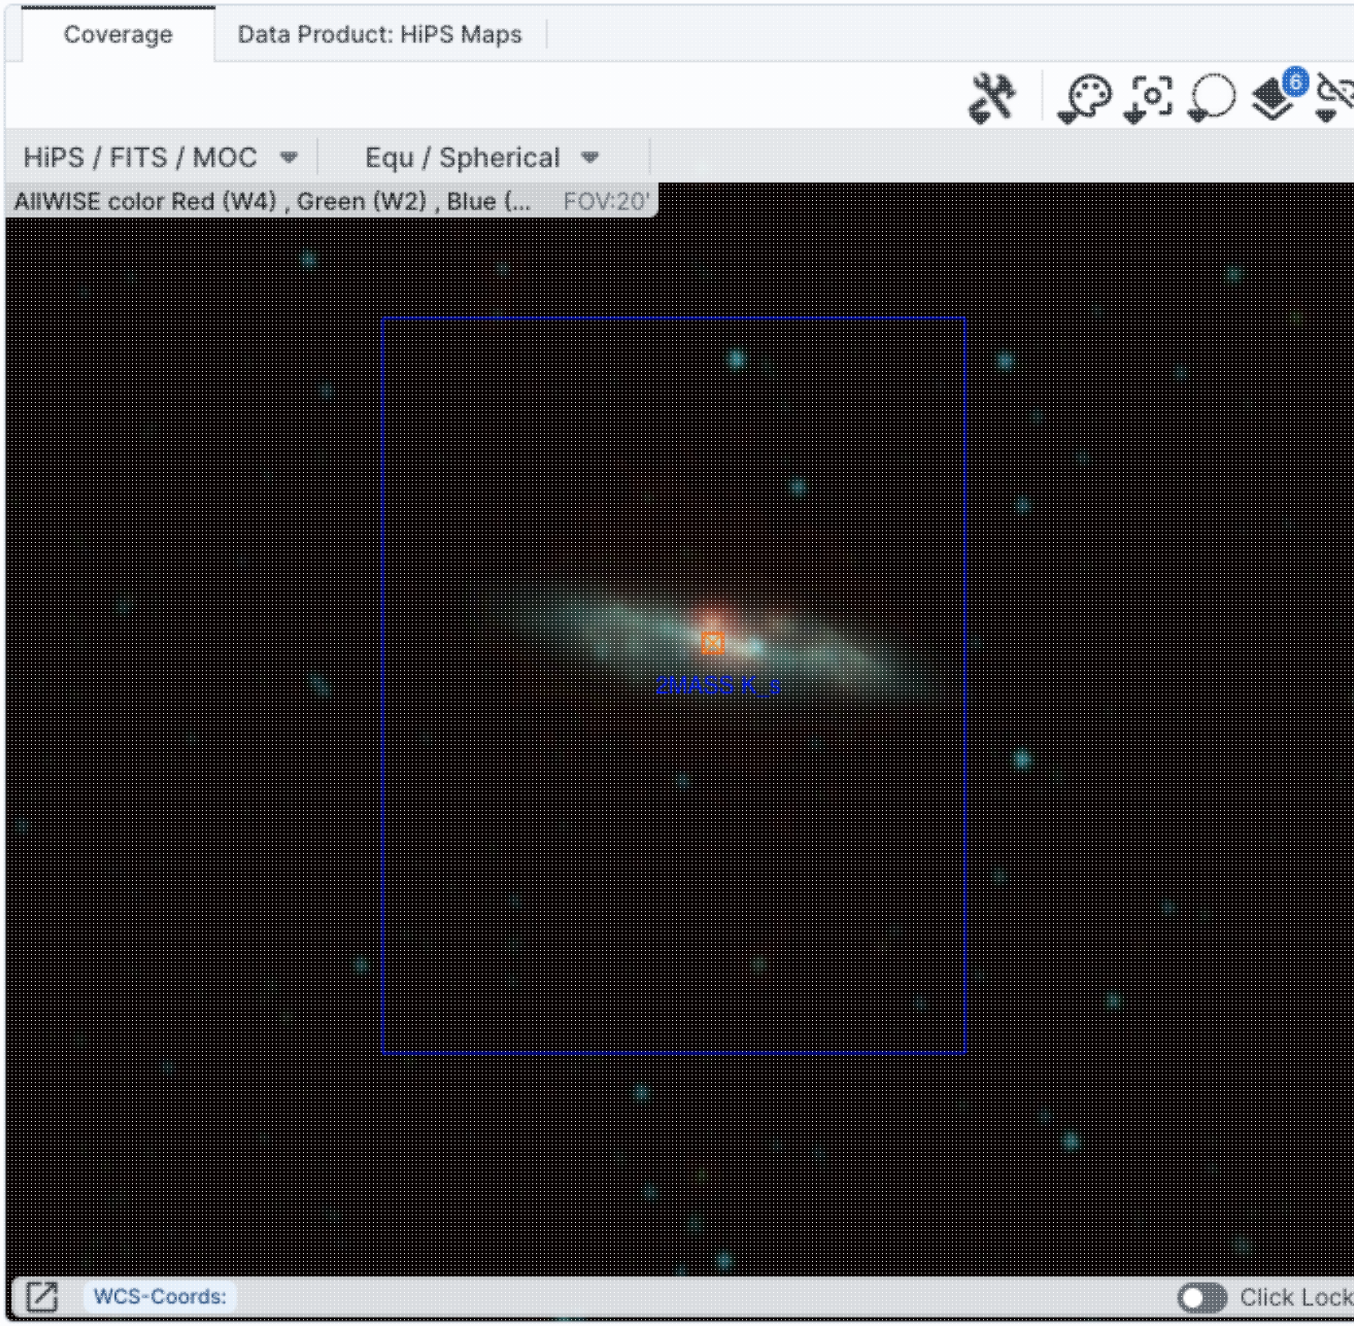

The second drop-down menu looks like this.

There are three sections in this menu.

Orientation

Under "Orientation", you can control whether the coordinates are in Galactic or Equatorial J2000 (RA/Dec). This can be used in conjunction with the image readout and/or the coordinate layer button (both described in the visualization section)

Center Galactic

If you select "Center Galactic", the HiPS image slews to put Galactic North up, the Galactic Center in the center of the field of view, and the Galactic Plane horizontally across your field of view. This is useful if you have been zooming or scrolling around to look at individual sources and need to bring it back to a familiar orientation quickly.

Projection

Under "Projection", you can control whether the display is in Spherical or Aitoff coordinates. Aitoff projection works better for all-sky displays.

IRSA Viewer: Visualization

At its core, IRSA Viewer is a way of visualizing and interacting with image files, both FITS and HiPS formats. You can also overlay catalogs. See the Catalogs section for more on catalogs. This section covers interacting with the images -- the visualization tools. The basics of searching for and loading images are in another section.

All of the interactive image visualization tools work the same basic way, and here we describe these basic options, in roughly the order in which you might encounter them in the window.

Contents of page/chapter:

FITS/HiPS Image Viewer

You can interactively explore the image with the mouse. Move your mouse over any image that is loaded into the viewer. Details about the image, specifically, the pixel beneath your mouse cursor, appear along the bottom left of the image window. Some information is updated in real time (coordinates); some information (pixel value) is updated when you stop moving your mouse for a second or two. The image can be interactively investigated in this fashion.

You can change the units of what is being read out, in terms of coordinates or pixel values.

If you click on the label of the coordinates, "EQ-J2000" in the screenshot example above, you get this pop-up, from which you can choose the coordinates from among:

- Equatorial (RA/Dec) J2000 in hh:mm:ss ddd:mm:ss format
- Equatorial (RA/Dec) J2000 in decimal degrees
- Galactic in decimal degrees
- Equatorial B1950
- Ecliptic J2000
- Ecliptic B1950
- FITS Image Pixel

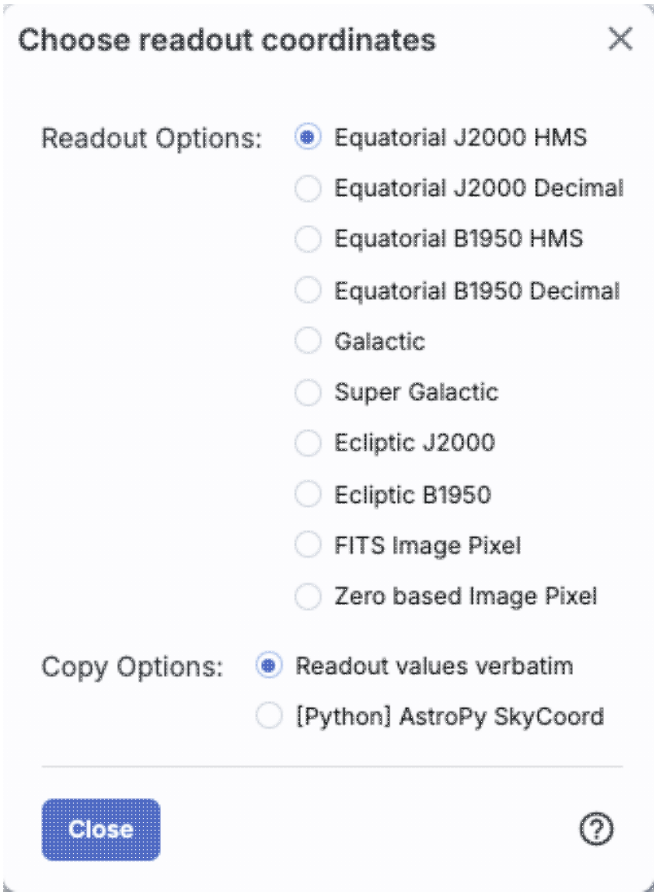

- Zero-based Image Pixel

If you click on the "click lock" toggle, the coordinates stop dynamically updating when you move your mouse, and they update only when you click on the image. When you do that, little clip boards appear next to each coordinate readout; clicking on those copy the position to your clipboard. From this pop-up window here as shown, you can control the format of the coordinates that are copied to your clipboard -- they can be as shown in the readout, or in the format that Python is expecting (for easy pasting into code).

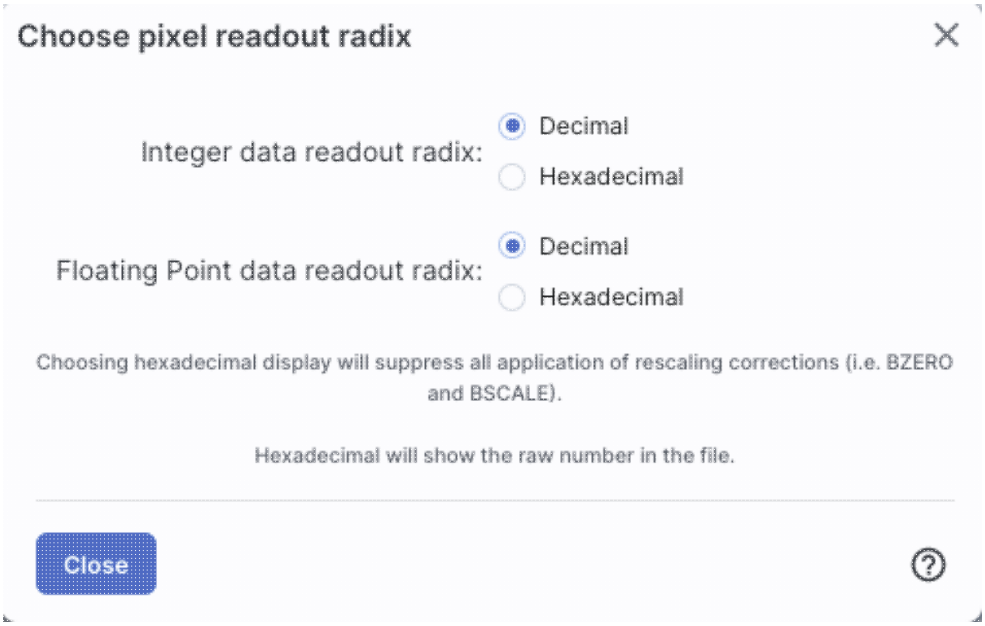

If you have a FITS image loaded, you have an additional readout. Click on the label of the readout, "Flux" in the tiny snippet of a screenshot example above, and you get this pop-up, from which you can choose the pixel readout from among:

- Integer data readout in decimal
- Integer data readout in hexadecimal
- Floating point data readout in decimal
- Floating point data readout in hexadecimal

If you choose the hexadecimal options, it will suppress all rescaling

corrections found in the header, like BZERO or BSCALE. It will just show you the raw binary number in the file. (For example, if the value in decimals is 5.13795757, the binary value in the file is 0x40a46a26; here is a conversion tool between decimal and hex.)

In the lower left of the images, if you click on this: , you get this pop-up. If you have a FITS image loaded, at the top of this pop-up, it shows the whole image; the orientation of the image is given with a compass rose. There is also a zoom-in of the image at the location under your cursor. Underneath that in the pop-up, whether you have a FITS or HiPS image loaded, you can get a readout of the pixel size, a readout of location on the image in two different coordinate systems, and a readout of the pixel value. You can change the units of those values by clicking on the name of the field: "Pixel Size", "EQ-J2000", "Image Pixel", and "Value". Each results in a pop-up, as above.

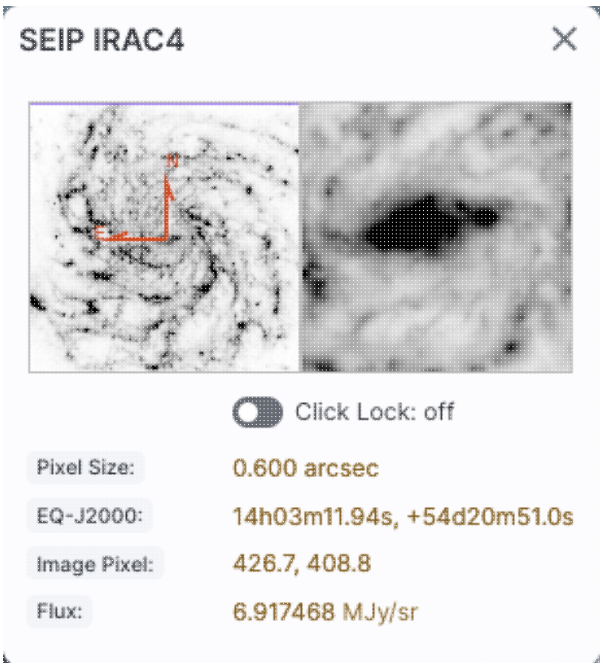

You can make the cursor 'stick' on a particular place on the image -- flip the "Click Lock: off" switch to "on" (either in the pop-up or in the lower right of the image window), and then click on the image at your desired location. When this is clicked, small "clipboards" appear near the position readout. Click on that icon to copy that position to your clipboard.

Image Information

The upper left corner of the loaded image has a label that tells you basic things about the image you are viewing. The telescope and/or instrument and/or channel and/or data release is first, in black. The field of view follows, in a paler font; this corresponds to the (horizontal) width of the image window. The currently selected image is outlined in brown; other images are outlined in grey.

Here are three examples of image labels. The first is from Spitzer Enhanced Imaging Products (SEIP), IRAC channel 4, which is 8 micron data, and the field of view is 12 arcmin. The second is from data delivered by the Local Volume Legacy (LVL) project, and it is from MIPS channel 3, which is 160 micron data; the field of view is 1.5 degrees. The third is also from the LVL

project, but it is far-ultraviolet (FUV) data, and the field of view is 32 arcseconds.

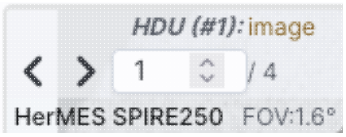

Images can have multiple planes; the arrows allow you to page through the planes. (This is from the HerMES project and is Herschel SPIRE 250 micron data.)

For HiPS images, the FOV is the angular size of the width of the HiPS viewer. Even if the image as displayed is smaller than the window, the FOV readout is the width of the window, not the image. If you shrink your browser screen, the FOV can get smaller because the viewer gets smaller. If you load more than one image, the FOV can get smaller because two viewers must fit in the same pane. As a result, the HiPS FOV requested in the search panel is approximate.

For FITS images, the FOV label on FITS images works analogously to the FOV label on HiPS images. If you zoom out, the FOV will increase even when the FITS image is entirely within the viewer. That's because the FOV is what the viewer can show you based on the pixel size. If you drag the image so that it is only partially seen through the viewer, the FOV will not change. For FITS images, the cutout size is not the same as the FOV.

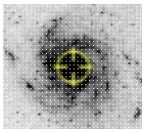
The target on which you searched is overlaid on the main image with a cross-hair marker, sometimes called a "reticle." You can remove this (or change its color) from the layers pop-up, described below.

Breaking Out of the Pane (and Going Back)

Panes: If you have both images and catalogs loaded into IRSA Viewer, the screen is broken up into panes - one for images, one for catalogs, and one for plots from the catalog. If you have more than one image loaded in, the image pane is further subdivided.

Make it big! For some purposes, it is useful to individually view just the table, or the images, or the plots, as large as possible. In any pane, this icon  appears in the upper right of the pane. Clicking on it will expand the pane into a larger window, as big as possible given your browser size.

Go back the way it was: 

The large "Close" arrow at the upper left is always available in the expanded views, and enables you to return back to the pane view.

Special case of images only: If you have only images loaded in, then the images are taking up all of your browser window, and it is already, by default, in this expanded mode. There's no 'close' arrow in the upper left since there is nothing else loaded in.

Removing things: To remove an image (or catalog) entirely, click on the small 'x' in the upper right of the image in the tiled view, or on the small 'x' in the corner of the catalog tab in the window pane view.

Note that you can also change the relative layout of the image, table, and plot panes from the side menu.

Also see the next section on image navigation.

Image Navigation

Single or Tiled Images

When you have many images loaded in, you can have icons like this: 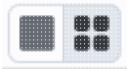 that portray (in icon form) the different views you can have of the images you have loaded. The first icon (the big square) denotes "show one image at a time." The second icon (the cluster of four squares) denotes "show smaller images of all the images I have loaded at once," e.g., tiled images. Whether the images (tiled or not) take up all the space or not depends on whether you are viewing in panes or in the full-screen mode (see immediately above on Breaking out of the pane).

Image List

Depending on what you have loaded, you may have an additional icon: 
Clicking on this icon brings up a list of the images you have loaded, with some additional information on each one. This list is a table like any other in this tool. There are two important things about it, though, that apply to this special case. (1) You can sort this table, and the images are sorted in response. (2) You can remove selected rows, or delete the failed images, with one click; click on the corresponding button on the lower right of this window.

The power of this table is best demonstrated by an example.

Example: Load the tool. Search on M101, with the default image size. Select all bands from the following data sets: SEIP, DUSTiNGS, LVL, MIPS_LG. Search. When the tool comes back, click on the image list icon above and obtain this pop-up:

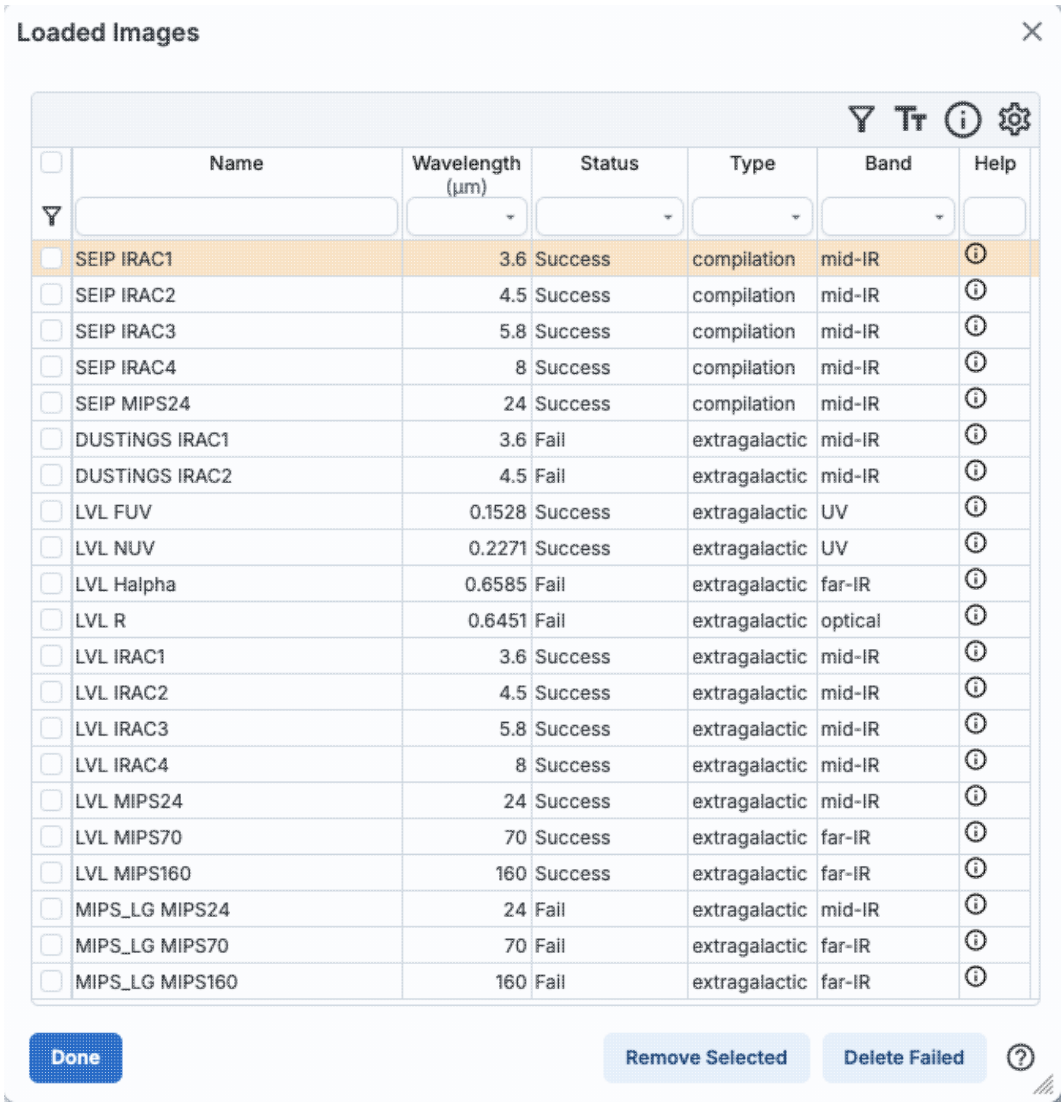

This table shows that it found images of M101 from SEIP, and some LVL bands, but nothing in DUSTiNGS or MIPS_LG. You can use the table to omit the failed images all at once. You can click on "Delete Failed" on the bottom right, or you can do the following: click the down arrow at the top of the "Status" column, tick the box next to "Success" and then click "filter." The failed images are removed all at once from the table and your display. Now, click the top of the "Wavelength" column. The table is sorted by wavelength, and the images in your display are now sorted that way as well.

The "help" links in the far right of this table take you to a master list of all data sets available in IRSA Viewer, from which you can obtain standard information about the data sets (mission, wavelengths, links to more information about the program or delivery, and more).

Paging through single image views

If you have many images loaded in and click on the single big square to view one image at a time, you are provided with navigation aids in the upper right, like this:

The arrows allow you to scroll through your list of images, sorted as specified in the image table. The filled dot in the list of circles shows you where your currently displayed image is in the set of loaded images. The "auto play" tick box on the left triggers automatic scrolling through each of the loaded images.

Scrolling through multi-image views

If you have several images loaded in and click on the cluster of four squares to view them all at once, you will also have this as a choice: 

If you toggle this on: 

then each image tile becomes bigger, and you can use your mouse to scroll through the collection of images. If you are on a Mac, your scrollbar may be hidden until you try to scroll.

Tips and Troubleshooting

◊ If your mouse is in a currently active (selected) image (that is, highlighted in brown), then your image will zoom rather than scroll. Just move your mouse over to another image, and then your window will scroll rather than zoom. Or, find your scrollbar.

Removing things

To remove an image (or anything else removable) entirely, click on the small 'x' in the upper right of the image in the tiled view. Closing the upper-right image leaves your mouse on or near the x for the next image that fills that corner, allowing multiple images to be closed with minimal mouse movement.

Image Toolbar (FITS and HiPS)

The image toolbox is always present as a row of tools associated with the images you have loaded, or that are loaded on your behalf. They're generally located above the images you have loaded, with most of them on the upper right. Letting your mouse hover over any of these icons will result in a "tool tip" that appears in order to remind you what the icon does. Most items apply equally to FITS and HiPS images, but some only apply to FITS images.

This is the image toolbox when you have clicked on a FITS image you have loaded:


And, this is the image toolbox when you have clicked on a HiPS image you have loaded:


The two toolbars are different, but if the same icon appears, it has the same effect on the image. Many of the icons have a downward pointing black triangle, which means that there are additional options in a drop-down menu that appear when you click on the icon.

We now discuss each icon in the order in which they appear.

 *Tools drop down*

The choices here look like this:

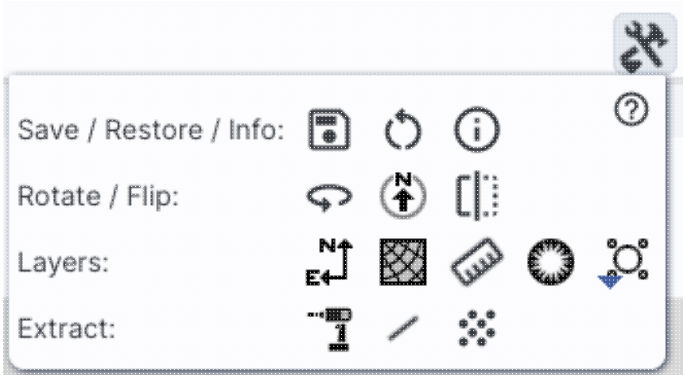

jump to save
jump to rotate
jump to layers
jump to extract

Saving the image

The diskette icon will allow you to save the current image. You can save files to your local disk or to the IRSA Workspace. Note that **you** control where the file is saved on your disk through your browser; your browser may be configured to store all downloads in a particular location on your disk.

If the current image is a FITS file, you can save it as a FITS or PNG or regions file to your local disk. If it is a HiPS file, your only choices are PNG or regions file. Saved FITS images will not save the color stretches or overlays; it will just save the underlying FITS image. Saved PNG files WILL include any overlays or annotations you have placed on the image, but will not include the underlying FITS image. Saved regions files will not save the underlying image, but will just save the overlays as a DS9 Regions file. See the DS9 website for more information on the syntax of these DS9 region files.

Note that you can save the original or a cropped version of a FITS file; see the "select region" icon below to crop, then click on the save icon. Be sure to save the cropped FITS image (see annotated figure). This feature is not available for HiPS images.

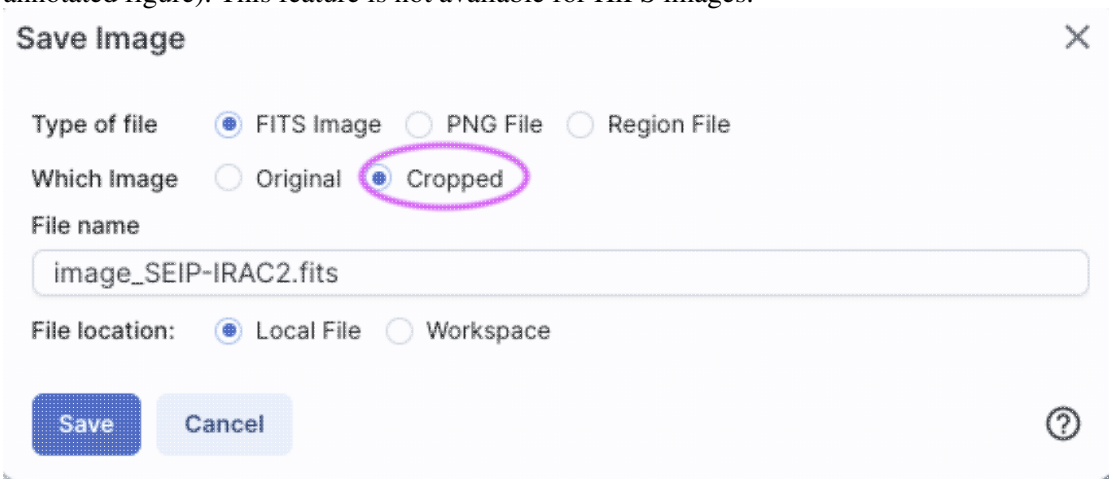

Note that if you overlay a large catalog on an image, then turn around and save a regions file from the catalog overlay, the full catalog may not be saved to the regions file. If you have >5,000 sources, it's entirely likely that not every source will be overlaid on the image (because of hierarchical catalogs display), and thus will not be in the regions file. If you want to save your entire catalog as a regions file, **save the catalog from the table pane**.

The saved PNG is the same size as it is on your screen. If you want a big version, make the desired image big on your screen (view one-at-a-time; see here) before saving the PNG.

You can't save HiPS images from within IRSA's tool. To download your own copy, you will have to track down the original source of the image.

Restoring everything to the defaults

If you've played around a lot with the image, you may want to undo everything you've done. Click this button to restore everything to their original default values. Some layers may persist; remove them via the layers icon.

Viewing the image header

This icon displays a pop-up window with information about the image. If a FITS image is selected, it will show the FITS header of the image; if a HiPS image is selected, it will show the HiPS properties of the image. These are Firefly tables like all the other tables in this tool, so they are sortable and filterable, etc. If you click on the columns in the pop-up, it will sort the keywords alphabetically by that column. This is useful for finding individual keywords in particularly densely populated FITS headers. Click the header again to sort in reverse-alphabetical order, and a third time to return to the default order. Below are examples of an original and sorted FITS header. To make this window go away, click on the 'x' in the upper right of the pop-up, or click "close" on the bottom left.

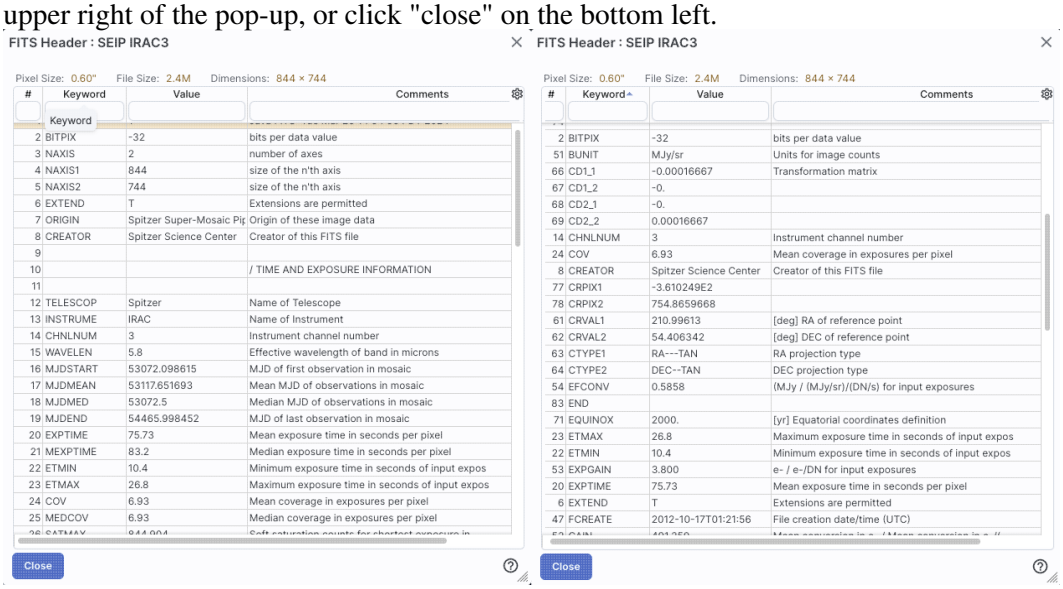

For comparison, an example of the HiPS properties window is here:

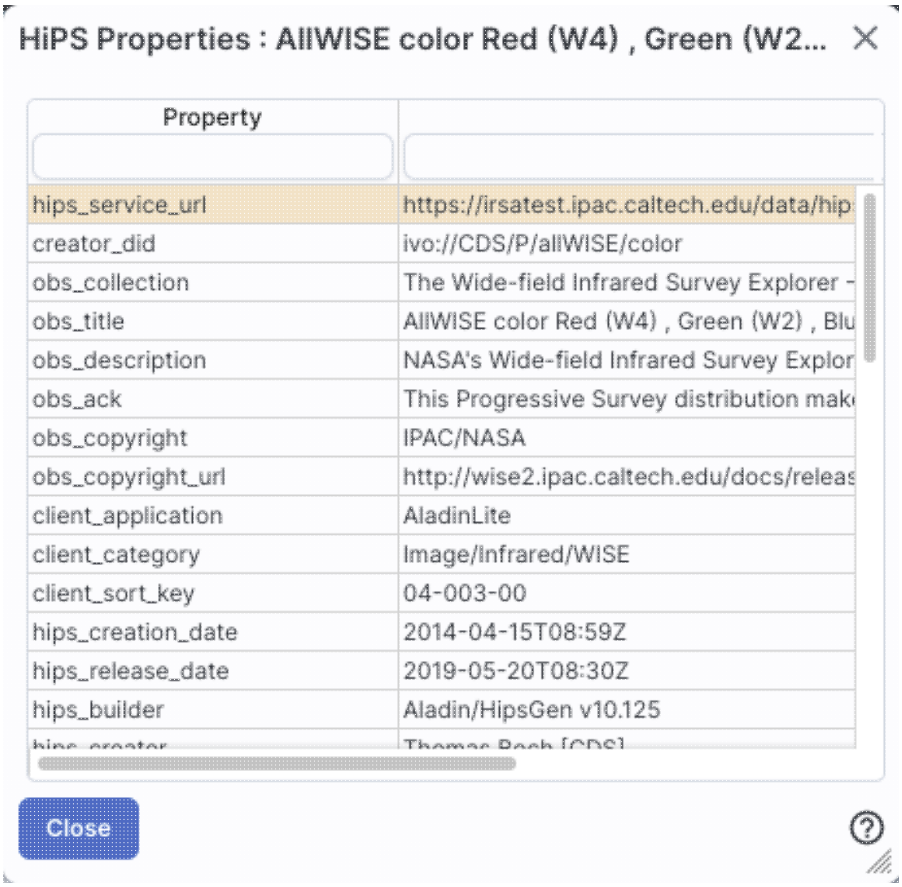

Rotating the image so that North is up

Images retrieved from IRSA Viewer are frequently (but not always!) already oriented such that North is up, or close to it. However, when interactively investigating images, or loading images from other sources, you could find yourself in a situation where North is not necessarily up. Clicking this icon will orient the selected image so that North is up. (This option is only available for FITS, not HiPS, images.)

Flipping the image on the y-axis

Clicking on this icon flips the image on the y-axis. (This option is only available for FITS, not HiPS, images.)

Add a compass rose

When you click this icon, arrows appear on the image showing which direction is North and which is East. Clicking on this icon a second time removes this compass rose. (You can also remove this layer via the layers icon, described below.)

Add a coordinate grid

Click on this icon to overlay a coordinate grid on the image. (Only available for FITS images, not HiPS, but see information on HiPS grid in the WCS section.) Click it again to remove it. Customize the units of the grid (to, e.g., Galactic coordinates) via the "layers" icon (described below).

Measuring a distance

When you click this icon, at first, nothing seems to happen. However, you can now click and drag to draw a line on the image, and the length of the line is displayed (in the middle of the line). The units for the measured distance (and the color of the overlay) can be changed from the "layers" icon (described below). You can calculate the difference in RA and Dec separately via the layers icon as well; find the layer associated with the distance measurement and tick the "offset calculation" box. When it displays the offset calculation, it will give you the angle in degrees in one corner, and the length of the line segment in the RA and Dec directions, in the units you have specified. When you are done with the distance tool, you can click on the

that appears next to the image toolbar, or click on this icon a second time to remove the distance tool. (You can also remove this layer via the layers icon.)

Read in a DS9 Regions file

When you click this icon, you get a pop-up window from which you can read in a DS9 regions file from your local disk. See the DS9 website for more information on the syntax of these DS9 region files. The supported regions are text, circle, box, polygon, line, and annulus. To make this window go away without doing anything, click on the 'x' in the upper right of the pop-up.

Tips and Troubleshooting: If you overlay a list of sources you created in ds9 regions format from your disk, it will only be overlaid on the current image, not all of the images you have loaded. If you want to have it overlaid on all the images you have loaded, create a catalog from your source list and overlay it as a catalog. Then it will appear on all of the images you have loaded, provided that the positions overlap on the sky.

Put a marker on the image

When you click this icon, a drop-down menu appears with several possible options:

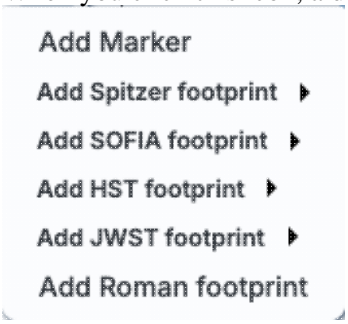

The first overlay choice (simply called 'marker') is a red circle. Initially, it appears in the center of the images, and is meant to be moved to wherever you first click in the image. It looks like this:

. The dash-dot line around it means that it is 'active', so you can move (click and drag the marker) or resize it (click and drag the dash-dot boundary). You can change the color of the marker (and change the label) via the "layers" icon (described below). You can also remove this layer via the layers icon. There are several additional options in the drop-down, enough that they have their own section below.

Drill down through the image

If your FITS image has multiple planes or HDUs, especially if each plane or HDU represents a different wavelength, it can be useful to "drill" down through the image cube at a given position on the sky. This tool allows you to do just that. When activated, this tool extracts the data at the place your mouse clicks down through the cube. For more information on saving the information, see the extraction section below.


Draw a line in the image

When this tool is activated, you can draw a line in your FITS image with your mouse, and it will extract for you the pixel values along that line. If you have more than one image loaded and visible, you can shift-click in another image to see the same line in another image. For more information on saving the information, see the extraction section below.


Make points in the image

When this tool is activated, you can click in your FITS image with your mouse, and it will extract for you the pixel values at the location of your click, creating a catalog for you as you click. If you have more than one image loaded and visible, you can shift-click in another image to extract points from another image. For more information on saving the information, see the extraction section below.


Zoom

When your mouse is in an image, these options appear in the upper left of the image.


Zooming in or out

Clicking on these magnifying glass icons zooms in or out of the image. The readout of the net effect of your zooming on the displayed field of view (FOV) appears at the top left of each image.

If you click zoom in or out rapidly, a pop-up window appears to allow you to more rapidly select the zoom level (field of view) you want. Select the desired level, or click on the 'x' in the upper right to make the window go away. Here is an example:

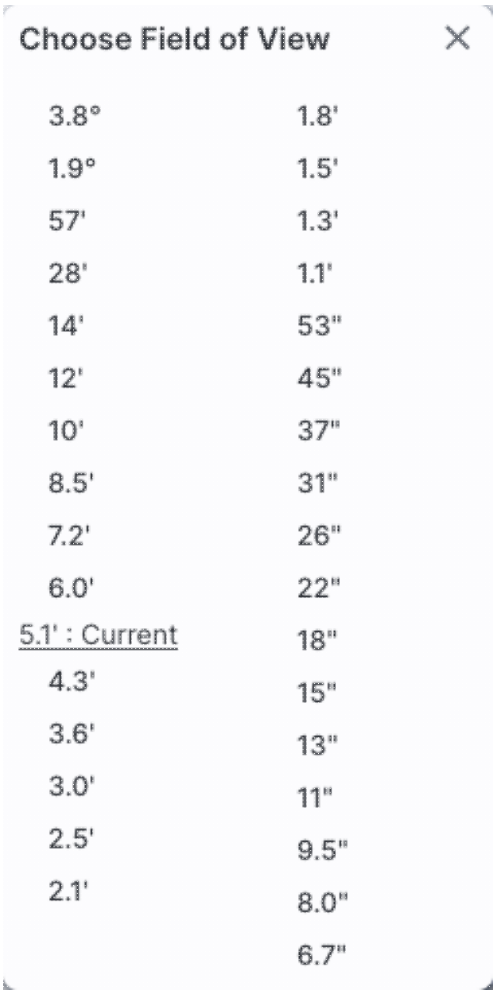

You can alternatively zoom using the mouse wheel (or drag forward and backward on a touchpad or magic mouse).

Note that there is a maximum (or minimum) allowed zoom level, and they are different for FITS and HiPS images. A notification will appear when you have reached the maximum (or minimum) allowed zoom level for a given image. To enlarge images more (or less) than that, please repeat your search to obtain new images with smaller (or larger) spatial extent. HiPS images are specifically designed for large areas, so if you need a big area, use HiPS. If you want to zoom in close enough to see individual original pixels, your best choice is FITS.

See also the section in the Images chapter on changing coverage images, specifically that on automatic transitions while zooming.

Fit image to screen or fill screen

These two icons are designed to maximize the available space in your browser window. The first one automatically picks a zoom level such that the image entirely fits within the available space (which could be your whole browser window, or just the portion of it where that image is loaded). The second one automatically picks a zoom level such that the image fills as much of the available space as possible (e.g., it is zoomed such that short axis of the window is filled with the image, whether that short axis is left-right or up-down).

By default, the images that are returned are frequently but not always centered on your search target. Clicking on these icons let you see the whole image that is returned, whether or not it is centered on your target.

This is available for both FITS and HiPS images, though note that FITS images retrieved from IRSA using this tool are typically square, and HiPS images cover the sky, so fitting the image to the screen might not be what you want to do.

Zooming to a 1-to-1 size

Clicking this icon will zoom the image such that one pixel in the image is one pixel on your screen. This option is only available with FITS images; HiPS images by their nature have pixels of varying sizes, so this button would have no meaning in this case.

Color table drop down

This icon enables you to change the color table of the displayed image. (This option is available for FITS and HiPS images.) When you click the button, a drop-down menu appears.

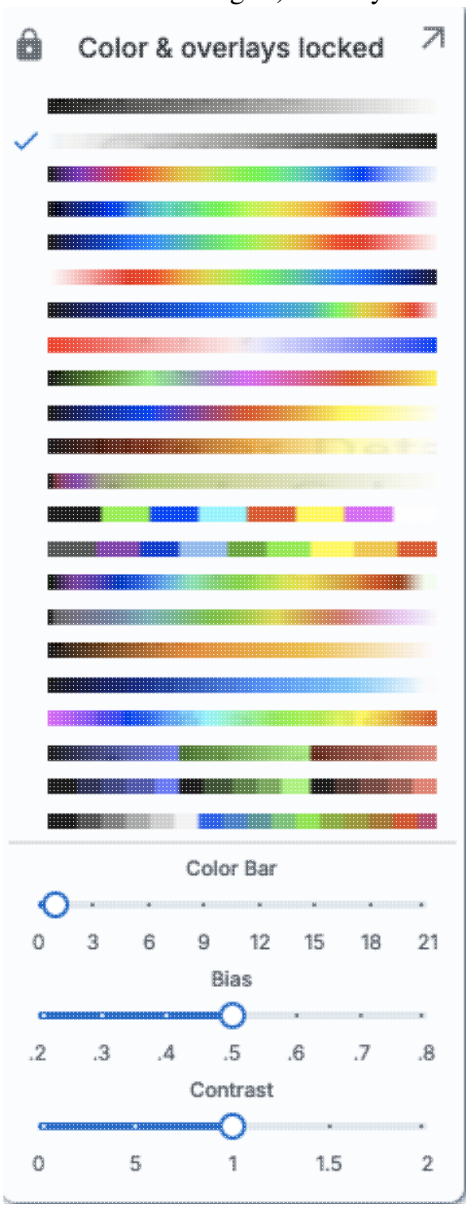

◊ The top of the menu either says "Color and overlays locked" or "Color and overlays unlocked" -- by default, all of the (FITS) images that you have loaded are locked together for color and overlays. What that means is if you change the color table (via this menu), then the color table for all the (FITS) images are changed. (Or, if you add a layer to one image, then the layer is added to all the images; see below.) If you don't want this to happen, select "Color and overlays locked" to unlock it. Select the text again to lock it again.

◊ The arrow in the upper right creates a pop-up window out of this drop-down menu so that you can leave the choices up while settling on the best option.

◊ The next portion of the menu has a wide variety of color table choices. Select your new color table from the options shown.

◊ Alternatively, you can use the "color bar" slider to move among the color tables by number. These numbers correspond to the color bar number used in the Python implementation of the Firefly tools.

◊ Below the color bar slider, there are sliders controlling the bias and contrast. Click or drag the slider to change the image display.

Color stretch drop down

This icon enables you to change the color stretch of the displayed image. (This option is only available for FITS, not HiPS, images.) Because this is complicated, for much more information, please see below.

Re-center the image drop down

Clicking this icon produces a drop-down menu:

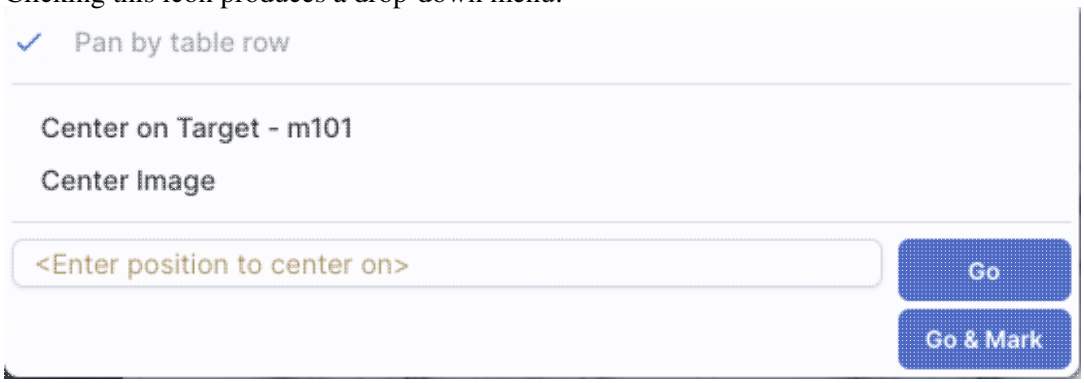

By default, "Pan by table row" is turned on (checked), but, depending on how you have loaded your images, or whether you have catalogs loaded, it may not seem to do very much. However, if you have a catalog loaded and are zoomed in on your images, as you scroll through your catalog, the FITS image underneath will move as needed when you have selected an observation in a different part of the sky than was originally shown.

Example of "Pan by table row" functionality: Load the tool. Search on M101, and ask for all SEIP channels (4 IRAC, 1 MIPS). When it returns, search on catalogs. Select a WISE/AllWISE catalog, and ask it for a polygon search covering the image that is loaded. When the catalog loads, go up to the image toolbox and be sure to select Align and Lock by WCS (see WCS Alignment below for more information). Sort the catalog by RA by clicking on the top of the RA column. Note that the images are all now slewed to center the first object (the furthest east) in the catalog. Click on a source in the plot near the west edge of the image. Note that the images are all now slewed to center that newly selected source.

Other choices are to center on the target of the observation, center the image in the window, or center on a target of your choice. For the last of those, you can simply center on that target, or center and leave a marker on the image at that location.

That last option may or may not appear, depending on what you have been doing before getting to this screen. If it can, it gives you a choice to center on recent positions. Move your mouse over to the arrow to select from a list.

Selecting a region drop down

When you click this icon, you can select a region of the image for further actions. Because this is complicated, for much more information, please see below.

Image Layers: Viewing/Changing the Layers on the Image

Every time you add something new to the image, you add a 'layer' to the image. This is complex, so please see below for much more information.

Lock/unlock images

You may have this "lock images" icon appear in your toolbar; it will appear as the first icon if they are

locked and the second icon if they are unlocked. The main purpose of this icon is to lock all the images you have loaded for zooming, scrolling, etc. You need to specify how it locks and for how long. Clicking it produces this drop-down menu:

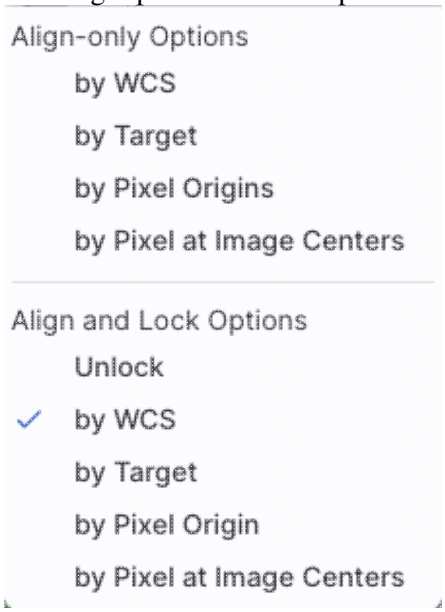

The first set of options aligns the images only once, temporarily; the second set of options makes the alignment persist ("lock") when you move the images (that is, when you move one, they all move). You can align by the images' WCS (e.g., RA and Dec), by the target, by the pixels according to the origin of the coordinate system in the image header, or by the pixel at the image center. The most common choice is likely the WCS align and lock. You can align FITS and HiPS images to each other. This is discussed in more detail in the WCS section.

 *Pin image*

This icon does not always appear. If you do have this icon, you can "pin" the image. Within this tool, "pinning" just means "hold on to this item within this tool." It doesn't mean "save this data to disk", nor does it mean "download all the data that goes with this"; it means "retain this item in this tool for now." Think of it as if you have a metaphorical bulletin board behind your computer monitor and you want to put an image you see on that bulletin board temporarily (with a pushpin!) to remember it while you continue to look through other images. In many cases, the tool will treat images as already pinned for you. In other cases, you have to actually pin them before it will create a pinned images tab.

 *Getting help*

Clicking on this icon takes you to this help page, usually to a spot in the help relevant to whatever you're currently doing.

To remove an image (or anything else removable) entirely, click on the small 'x' in the upper right of the image in the tiled view, or on the small 'x' in the corner of the image (or catalog) tab in the window pane view.

Color Stretches

 This icon enables you to change the color stretch of the displayed image. (This option is only available for FITS, not HiPS, images.) When you click the button, a drop-down menu appears with a variety of choices. You

can choose from a set of pre-selected options:

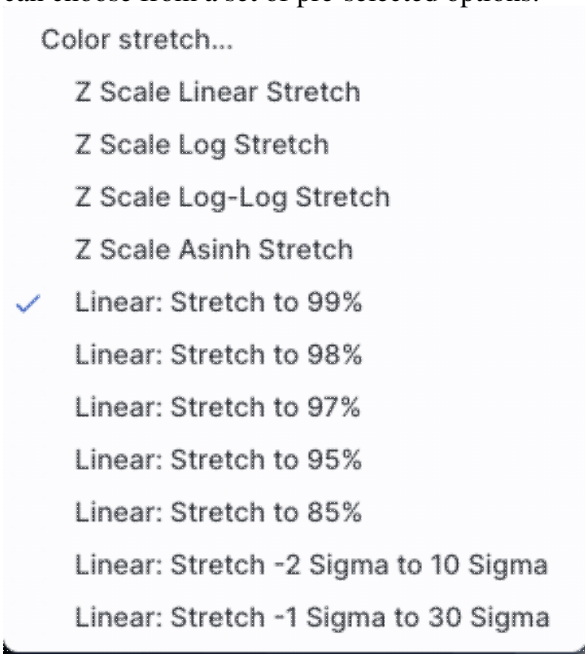

If you pick the first one, "color stretch", you can customize the stretch. A pop-up window appears with a histogram of the values in the image, and you can change the stretch type and range.

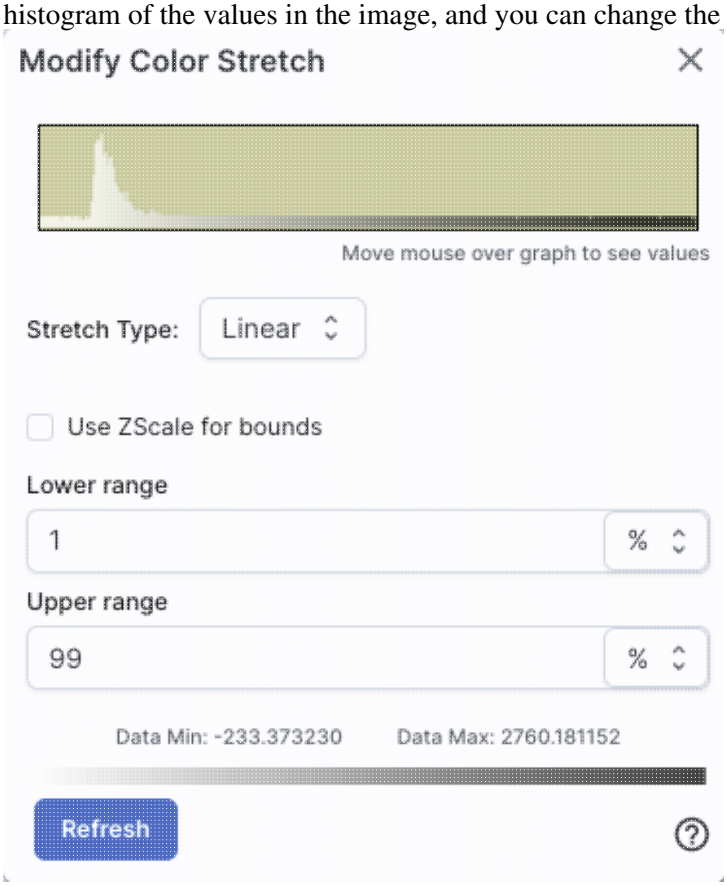

If you pick a color stretch from the pre-defined options, the pop-up window reflects this change. If you change the color stretch in the pop-up window, the drop-down menu changes correspondingly.

Example: Display the pop-up for color stretch. From the main drop-down, pick 'Linear stretch to 99%'. Go back to the color stretch pop-up. Note that it has filled out the stretch type and ranges to reflect the current

choice. Then -- either with the pop-up window still up or not -- go back and pick a different pre-defined stretch from the standard options. Note that the values in the pop-up change to reflect this current choice. From the pop-up, pick a different stretch type -- try "histogram equalization." Select "refresh" to update the images. Go back to the drop-down menu. The last 7 items have changed to be based on histogram equalization, as opposed to the "linear" default.

If you have a 3-color image, you can change the stretch in each color plane separately; select the tab at the top accordingly for red, green, or blue. By default, it stretches each band independently, and you can set the parameters in the stretch pop-up accordingly.

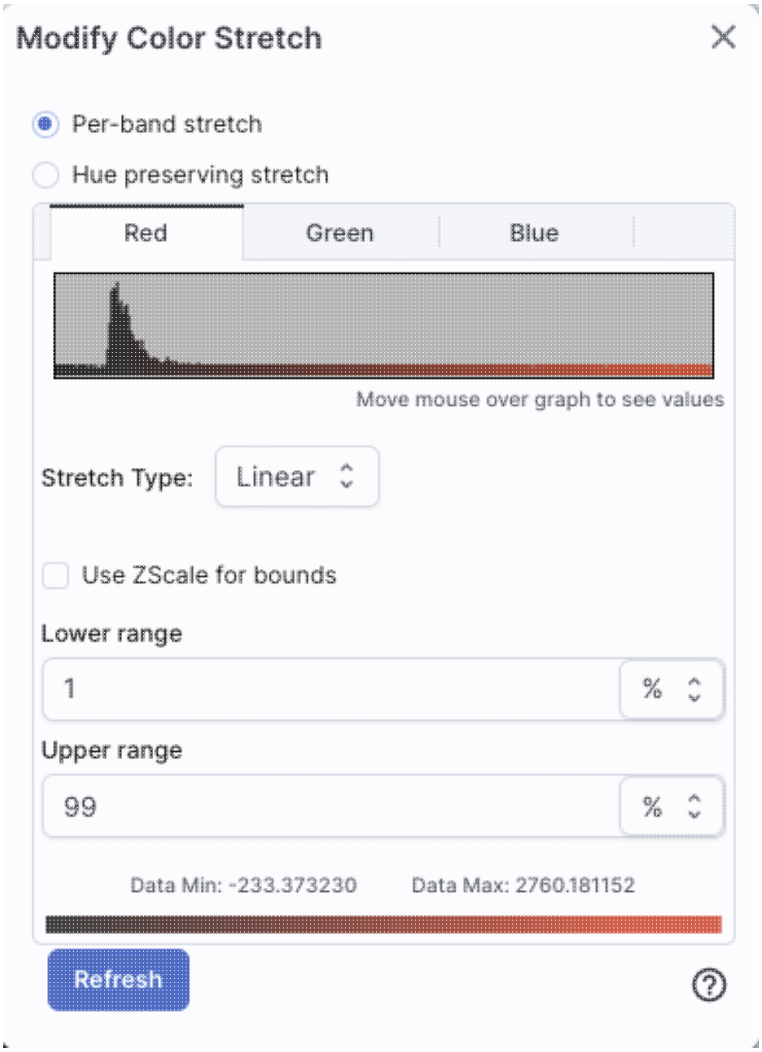

As described in Lupton et al. (2004), a different algorithm may be useful for creating 3-band color images. Select "Hue preserving stretch" to invoke this option. This stretch should be a brightness-independent color-preserving asinh stretch, though in practical terms, it seems to work best for optical images.

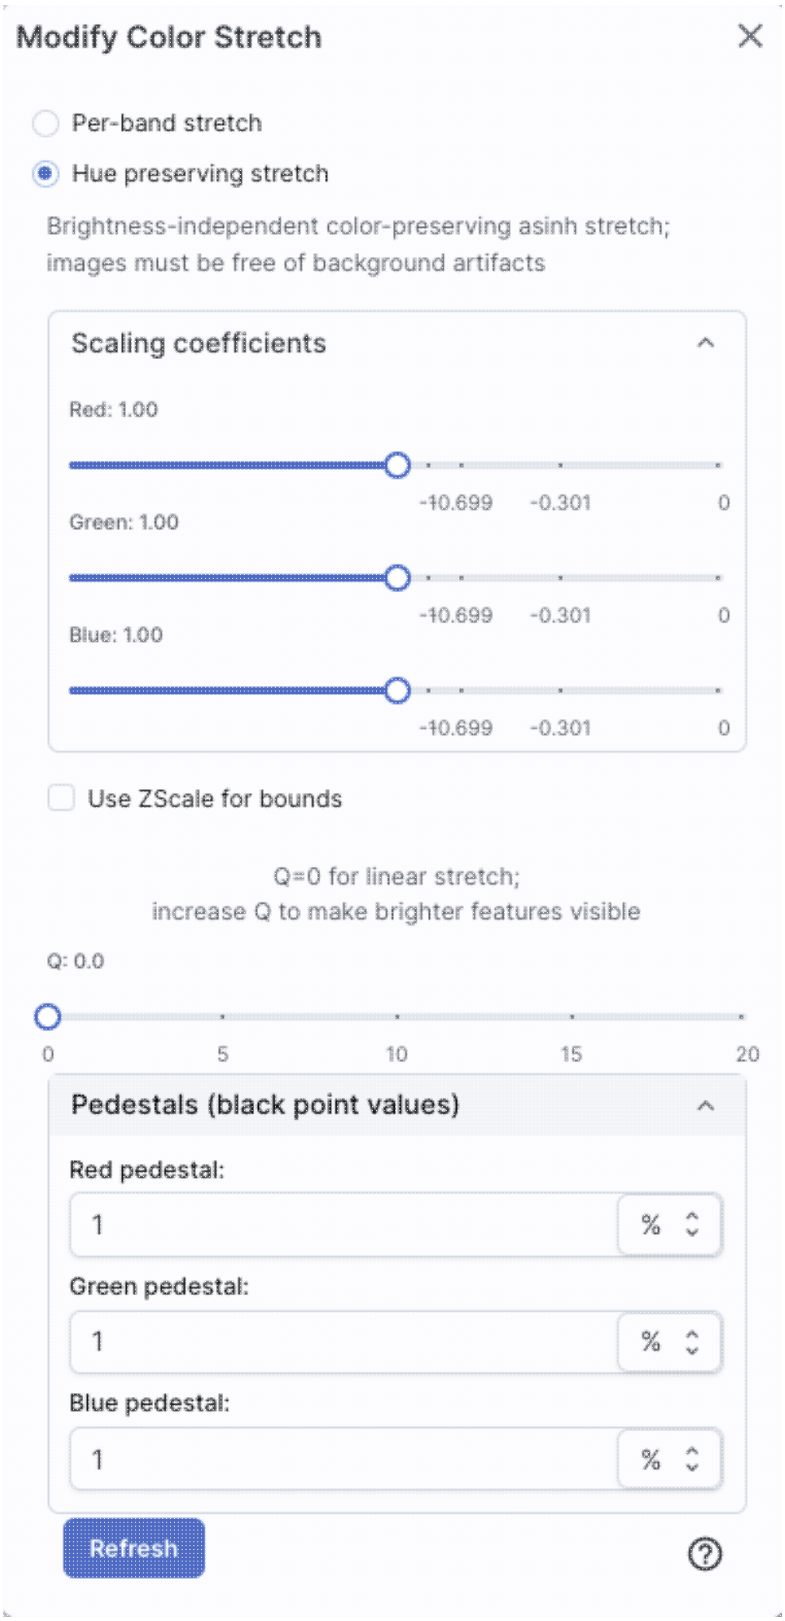

It may be useful to scale individual channels; sliders allow you to do so. The Q parameter has another slider. For a linear stretch, Q=0; increase Q to change what features are emphasized. Pedestal values can also be set to allow the level assigned to "black" to change.

Viewing/changing the layers on the image

Every time you add something new to the image, you add a 'layer' to the image. You can have an image with a lot of annotations on it.

The number that appears circled in blue over the layers icon tells you at any given time how many layers you have on the currently selected image (the image outlined in brown).

If you click this layers icon, you will get a pop-up window with a list of all the layers you have on top of the image. Here (on the right) is an example of a well-populated layers pop-up; in real life, this is scrollable to see several more layers). From this pop-up, you can:

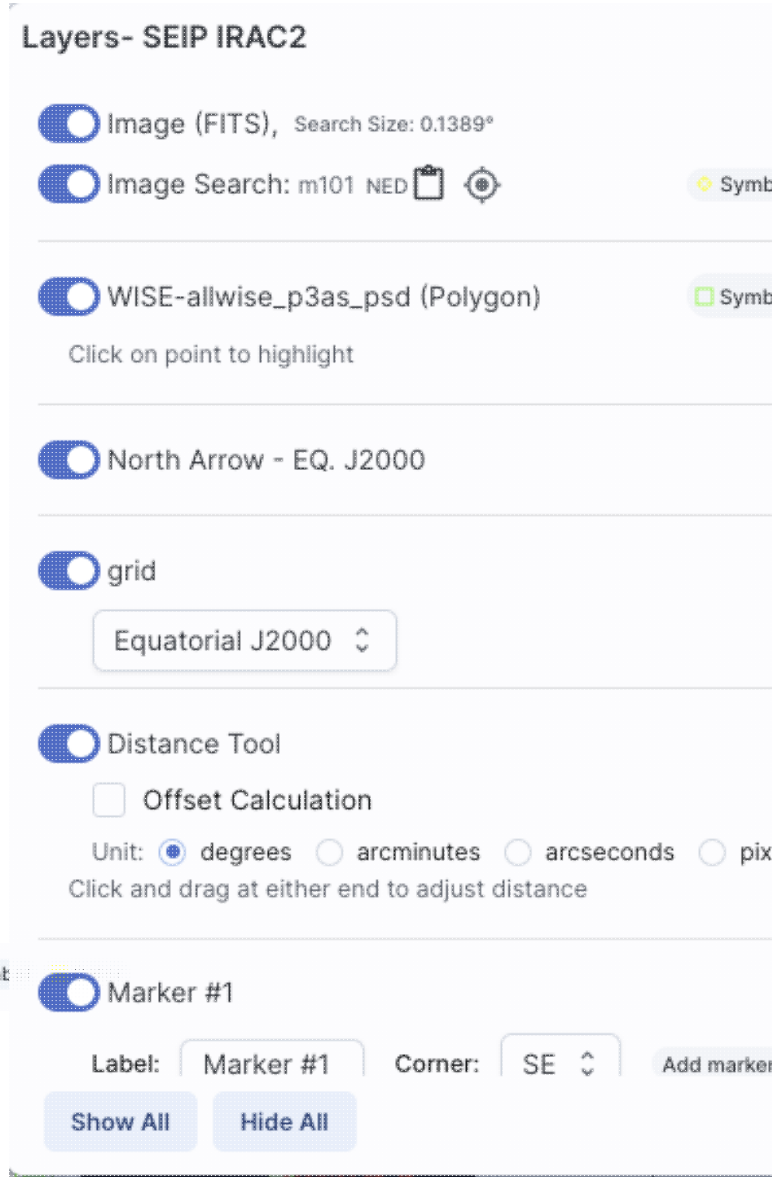

- turn layers off and on (click on the switch on the left of the corresponding row);
- remove layers entirely (click on the 'x' on the right of the corresponding row);
- change colors of overlays (see below);
- change symbol shapes and sizes (for overlaid catalogs), including hierarchical catalog settings;
- change annnotations (for markers);
- or change units (for the coordinate grid or the distance tool).

To add entirely new layers, though, you need to go to other options within the toolbar.

You can "show all" or "hide all" with the buttons on the lower left of the pop-up window. To make this pop-up window go away, click on the 'x' in the upper right of the pop-up.

Note the target description:

This reminds you of the target on which you searched -- here, it was M101, where the coordinates were resolved by NED (as opposed to Simbad). The two icons next in that row indicate, respectively, "copy this location to the clipboard" and "center image on this position."

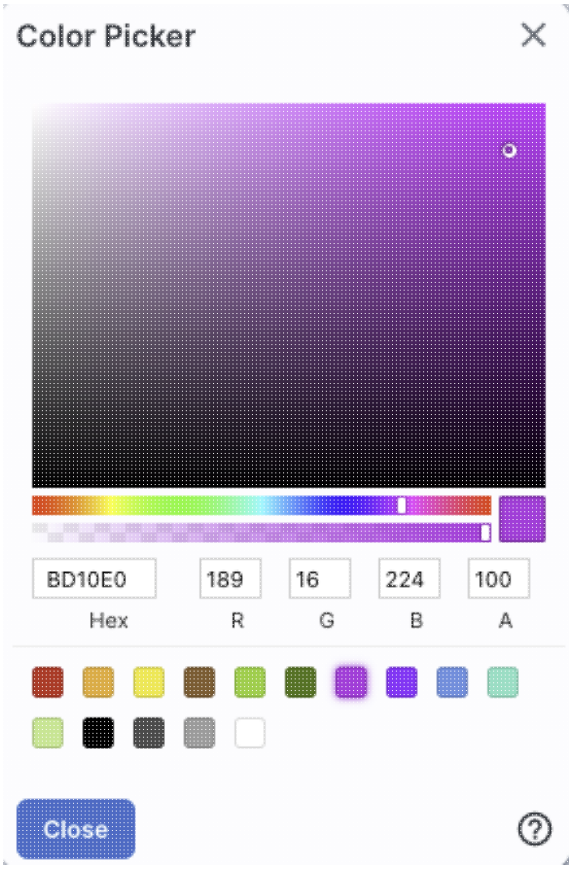

Where it's possible to change colors of a layer, click on the 'colors' link to be taken to a new pop-up from which you can select a new color.

From here, you can click on your desired color in the top colorful box. Immediately below that box, you can change the color and saturation of the top box so that you can select from a different range of colors. Below that, you can enter numerical hex codes or RGBA values (where the value for RGB is between 0 and 255, and A is in units of percent, e.g., 50 = 50%). Finally, you can also select from a pre-defined set of 15 colors by clicking on any of the small boxes. Note that the numerical codes update as you select different colors. Your choices are implemented as soon as you select them. Click 'Close' to close the window, or click 'x' in the upper right.

If you have a catalog loaded into the tool, you can also obtain this pop-up by clicking on the color swatch in the heading of the catalog tab.

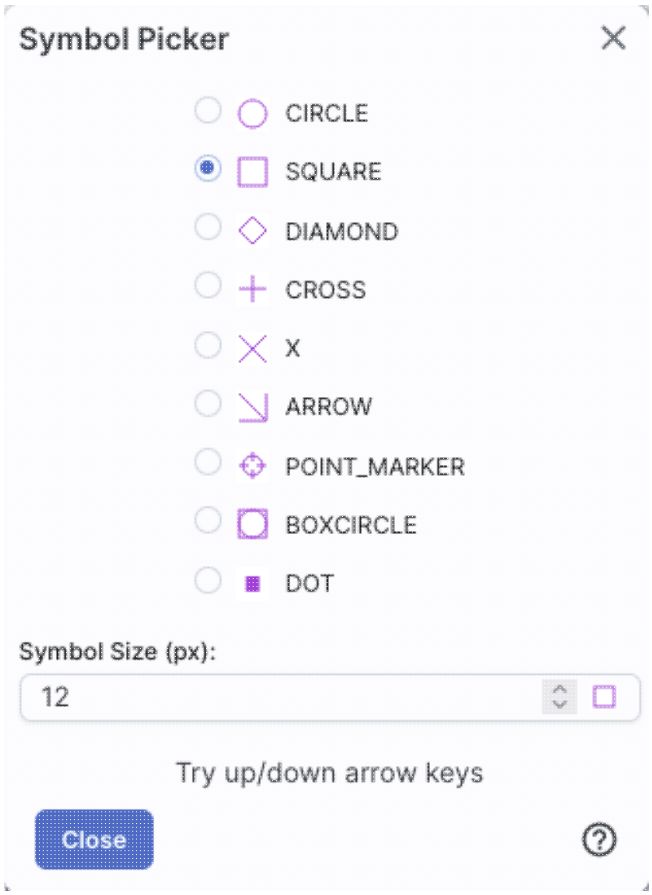

For catalogs or the search target, you can also select the symbol shape and size. To adjust the size, type in the symbol size in pixels or use the up/down arrow keys to change the size by one pixel at a time. Your choices are implemented as soon as you select them. Click 'Close' to close the window, or click 'x' in the upper right.

Note that if you load both FITS and HiPS images at the same time, it can include a marked layer on the HiPS image that is the footprint of the FITS images you have loaded. A label appears at the center of that footprint, which may be disconcerting if you are not zoomed out enough to see the region itself. Here is an example, zoomed out so it is clear what is going on:

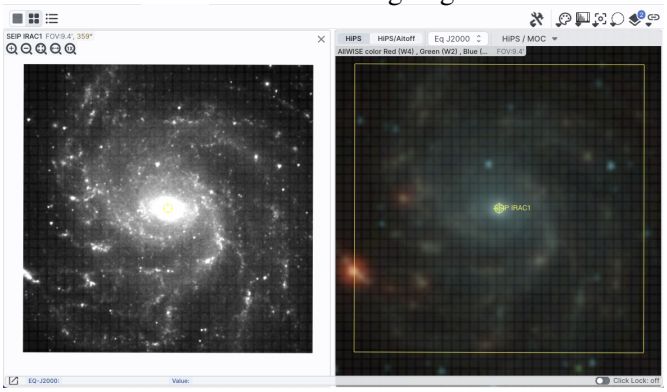

Once you have loaded a HiPS image into IRSA Viewer, if you select the HiPS image and click on the layers icon ( ), you will have new, HiPS-specific choices in the layers:

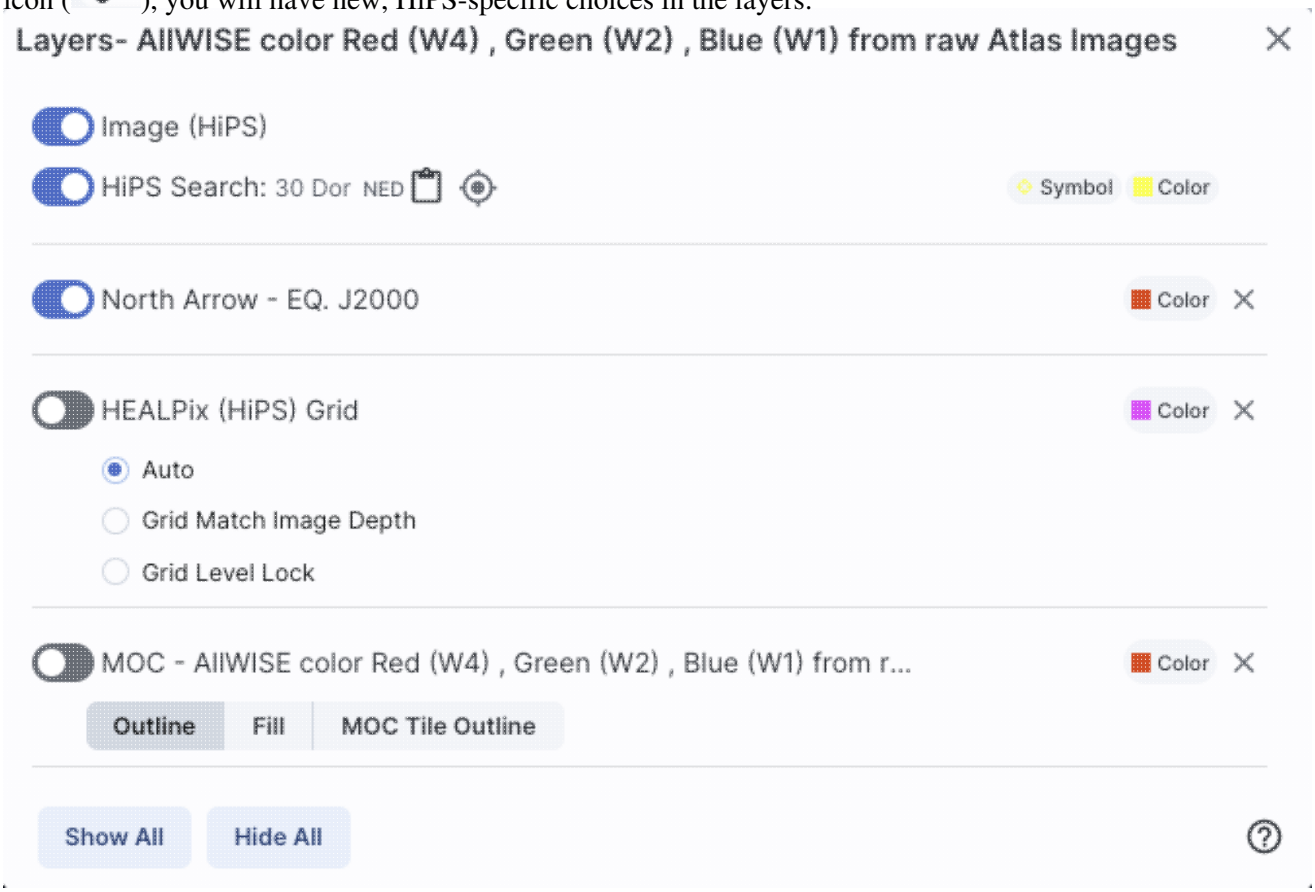

HEALPix (HiPS) Grid

To turn on these choices, toggle the switch to the left of "HEALPix (HiPS) Grid". (See here for more information on HiPS images in general.)

Auto: This option overlays a position grid, with the tile numbers marked in the center of each box. As you continue to zoom in, when smaller tiles are needed, they are drawn, with the new tile numbers

marked. You may not zoom beyond HiPS Norder level 14 tiles. The numbers after the "/" is in the "NESTED" (as opposed to RING or NUNIQ) numbering system; see the IVOA standards document for more information.

Grid Match Image Depth: If you select this option, the grid will adjust to a new level when you zoom in and a new level of HiPS image both exists and is used for the display.

Grid Level Lock: Selecting this option yields an additional numerical drop-down menu. The higher number you pick, the smaller the grid boxes are that are drawn. When this option is selected, the boxes stay the same size regardless of how zoomed-in on the image you are.

HiPS MOC

To turn on these choices, toggle the switch to the left of "MOC".

(See here for more information on MOCs in general.) A MOC tells you via a simple boolean yes/no, is there sky coverage from this data set in this region. The choices here are:

- ◊ Outline - an attempt to outline the entire region covered by the data; it still sometimes struggles near the edges of coverage, so zoom in to get a better sense of the coverage edges.
- ◊ Fill - filled regions, where you can control the opacity of the overlay by going to the color picker; you control the opacity by changing the number above the "A".
- ◊ MOC Tile Outline - individual tile outlines, where the tiles are set by the MOC tiles themselves (as opposed to tiles created by the mosaic tiles that make up the data set).

Tips and Troubleshooting:

- The entire concept of a MOC is built upon the "tiles" that are inherent to the HiPS concept. As a result, those tiles are imprinted on how the MOC is rendered, especially near edges or corners of coverage. Strange behavior may result; you can always zoom in to get a better sense of the coverage. For authoritative information, download the actual data for the region you are concerned about.
- For the "fill" option for a MOC, depending on how you display a MOC, you may see two shades of color in the MOC. It is important to note, though, that the information it is displaying does not include depth of coverage, merely boolean "is there data there or not." Why is it displaying shading? Well, it's rounding. For example, a given WISE MOC might be generated at order 13. At this order, there are 805,306,368 HEALPixels on the sky, and they are about 26 arcseconds across. When zoomed out far, there is no point in trying to show each of these pixels, so the application "rounds up" the MOC to an order in which there are roughly 100-200 displayed HEALPixel polygons horizontally across the image. When it does this, it flags the rounded up polygons with the paler color. So the boundaries of a coverage region in the MOC all get a paler color. If you zoom in far enough on a MOC, the two-tone colors go away.
- You might be here in the documentation looking for more information on hierarchical catalog settings; this is also related to HEALPix.

World Coordinate System (WCS) Alignment

As described above, there is a way to lock/unlock images to each other for position matching. This section describes the image locking in more detail.

When aligning images, you can specify how the images align and for how long. Clicking the lock images icon produces this drop-down menu:

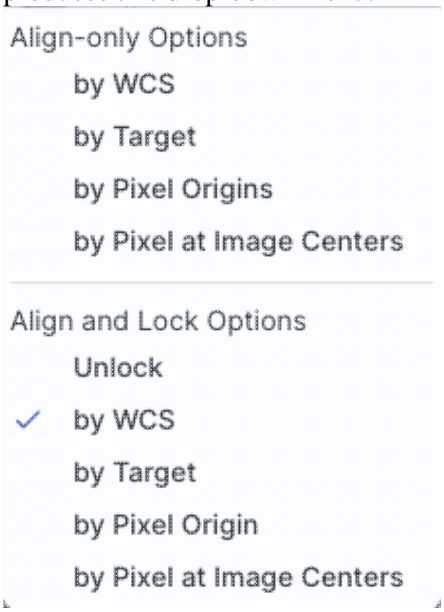

The first set of options aligns the images only once; the second set of options makes the alignment persist ("lock") when you move (zoom, etc.) the images.

You can align by the images' WCS (world coordinate system, e.g., RA and Dec), by the target (align by target on the screen, regardless of position in the sky), by the pixels according to the origin of the coordinate system in the image header, or by the pixel at the image center. The most common choice is likely the WCS align and lock.

Here are examples of different alignments, in this order: align by WCS, by pixel origin, and by pixel at image centers.

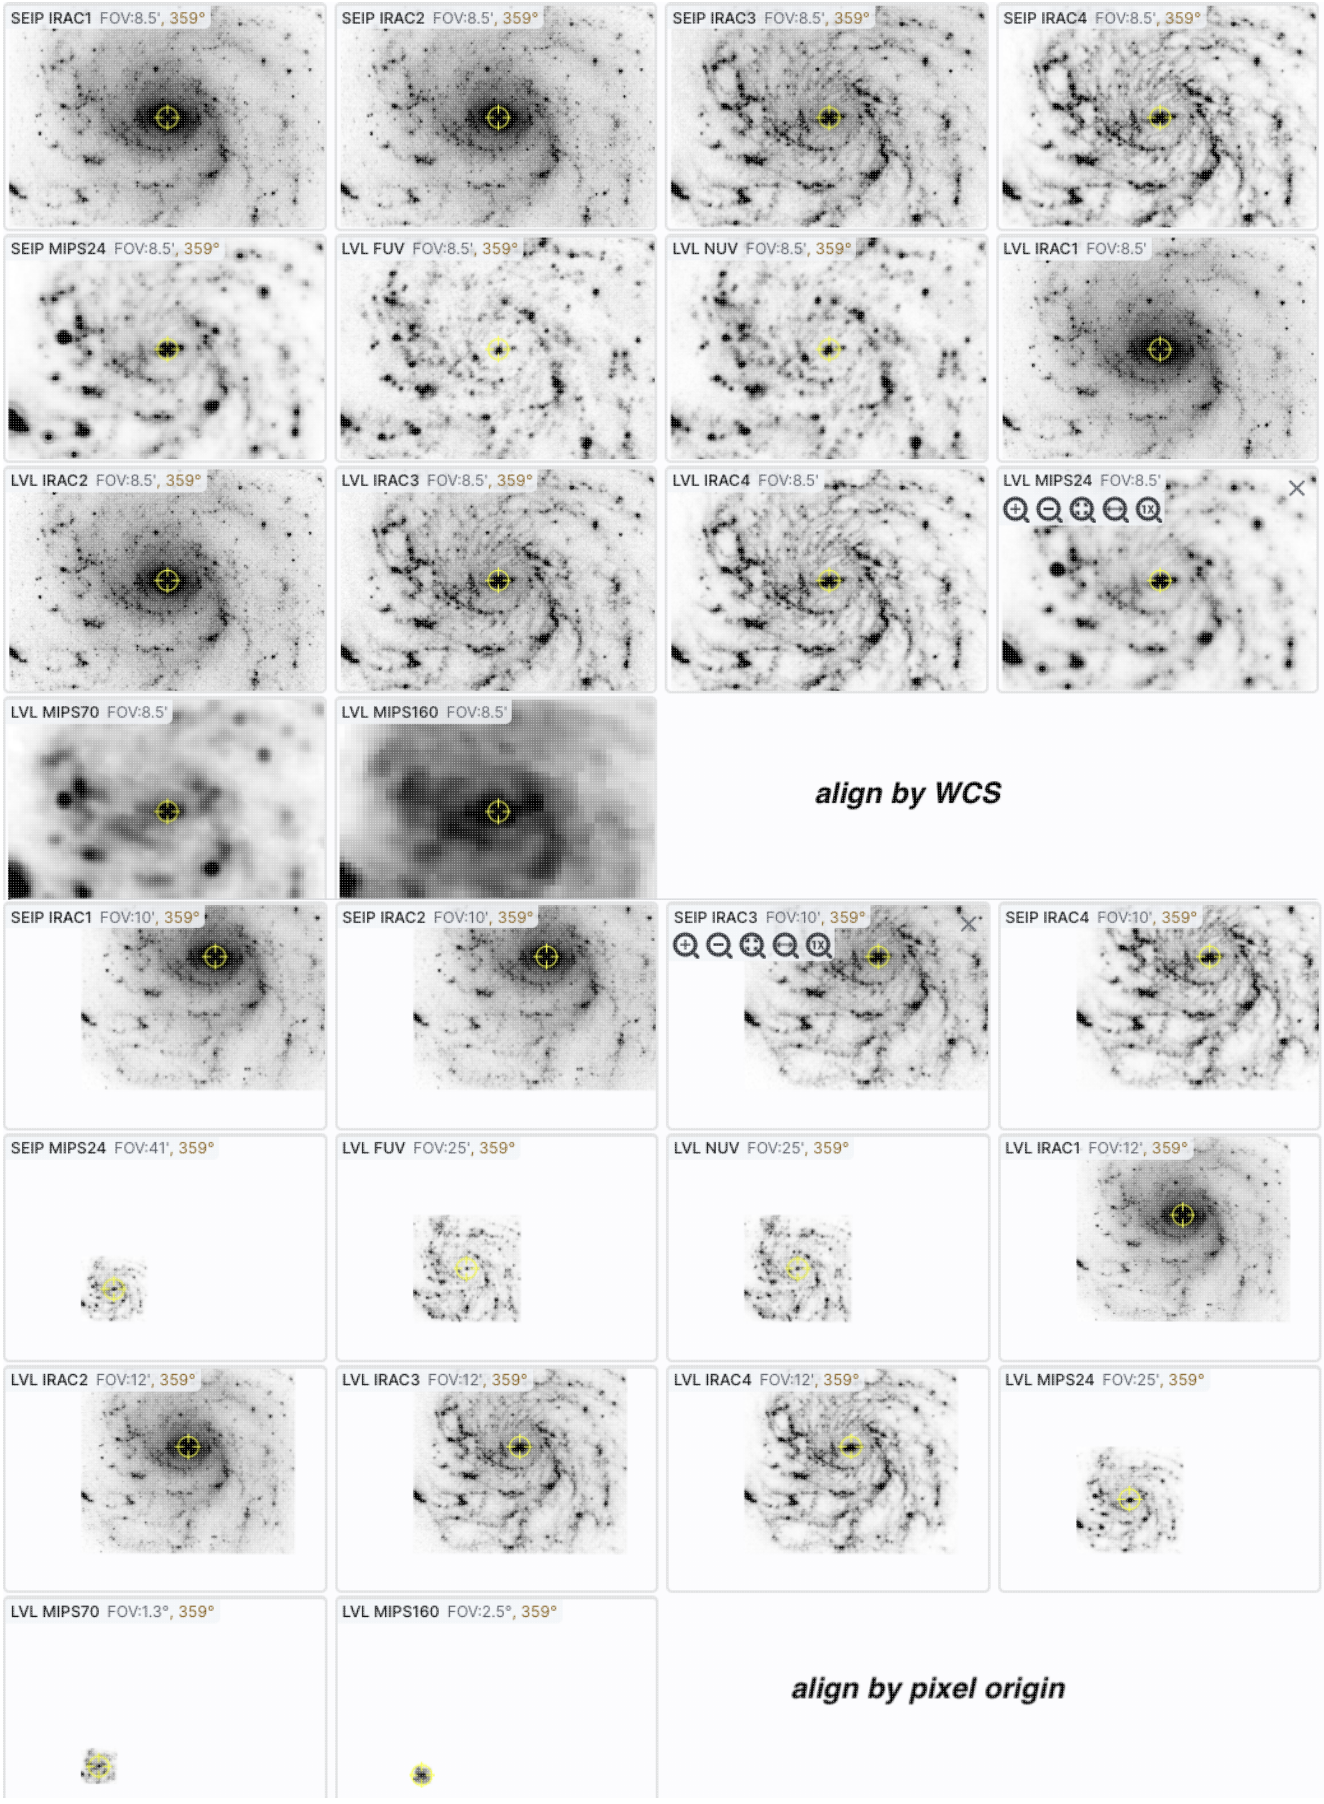
SEIP IRAC1 FOV:8.5', 359°
SEIP IRAC2 FOV:8.5', 359°
SEIP IRAC3 FOV:8.5', 359°
SEIP IRAC4 FOV:8.5', 359°
SEIP MIPS24 FOV:8.5', 359°
LVL FUV FOV:8.5', 359°
LVL NUV FOV:8.5', 359°
LVL IRAC1 FOV:8.5'
LVL IRAC2 FOV:8.5', 359°
LVL IRAC3 FOV:8.5', 359°
LVL IRAC4 FOV:8.5'
LVL MIPS24 FOV:8.5'
LVL MIPS70 FOV:8.5'
LVL MIPS160 FOV:8.5'
align by WCS
SEIP IRAC1 FOV:10', 359°
SEIP IRAC2 FOV:10', 359°
SEIP IRAC3 FOV:10', 359°
SEIP IRAC4 FOV:10', 359°
SEIP MIPS24 FOV:41', 359°
LVL FUV FOV:25', 359°
LVL NUV FOV:25', 359°
LVL IRAC1 FOV:12', 359°
LVL IRAC2 FOV:12', 359°
LVL IRAC3 FOV:12', 359°
LVL IRAC4 FOV:12', 359°
LVL MIPS24 FOV:25', 359°
LVL MIPS70 FOV:1.3°, 359°
LVL MIPS160 FOV:2.5°, 359°
align by pixel origin

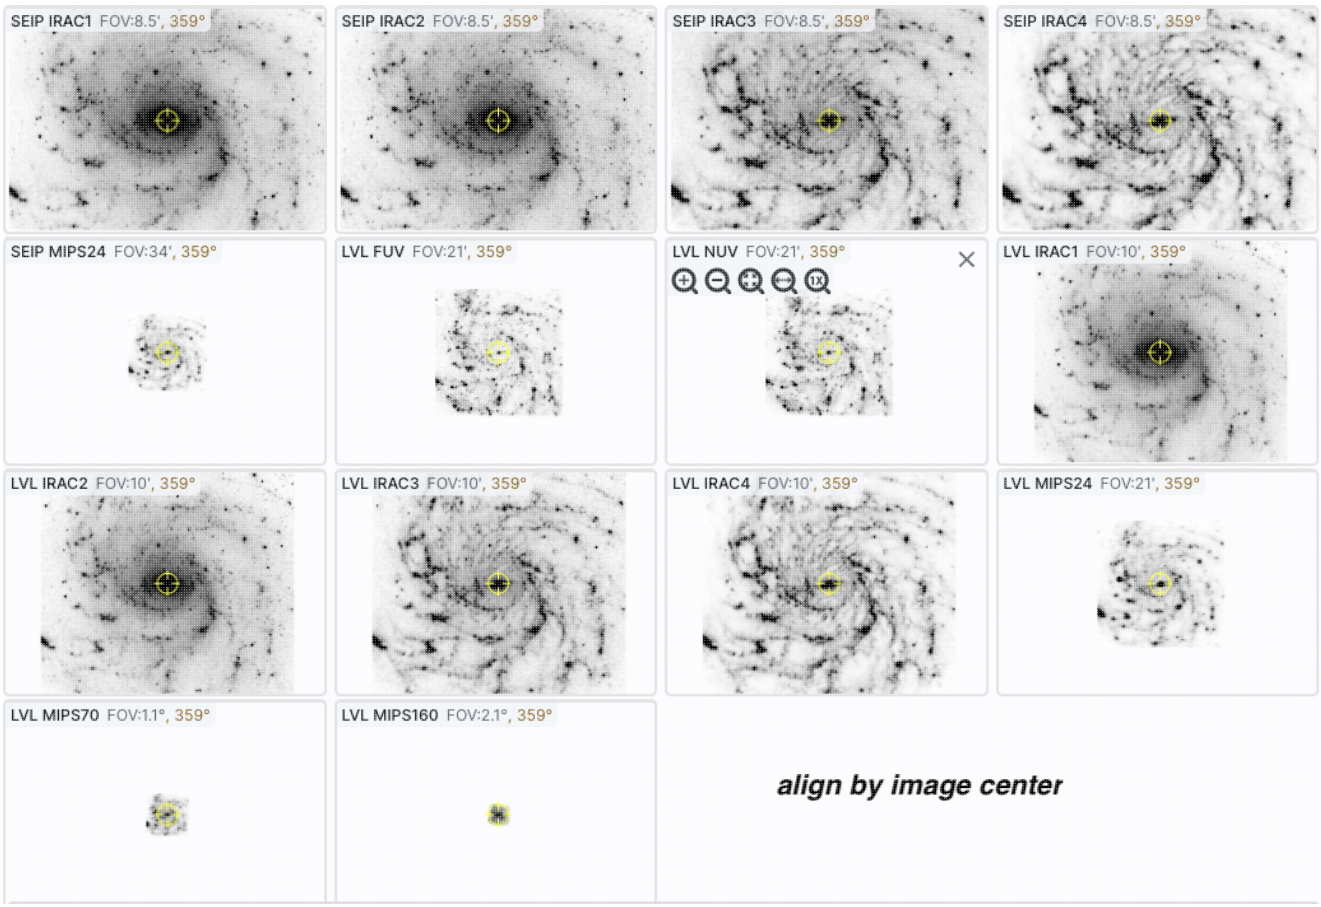

Note that aligning by WCS puts North up, and aligned so that each image has the same angular scale.

In contrast, here is an align by target - six different galaxies, but the target used for each image is in the center of each image tile.

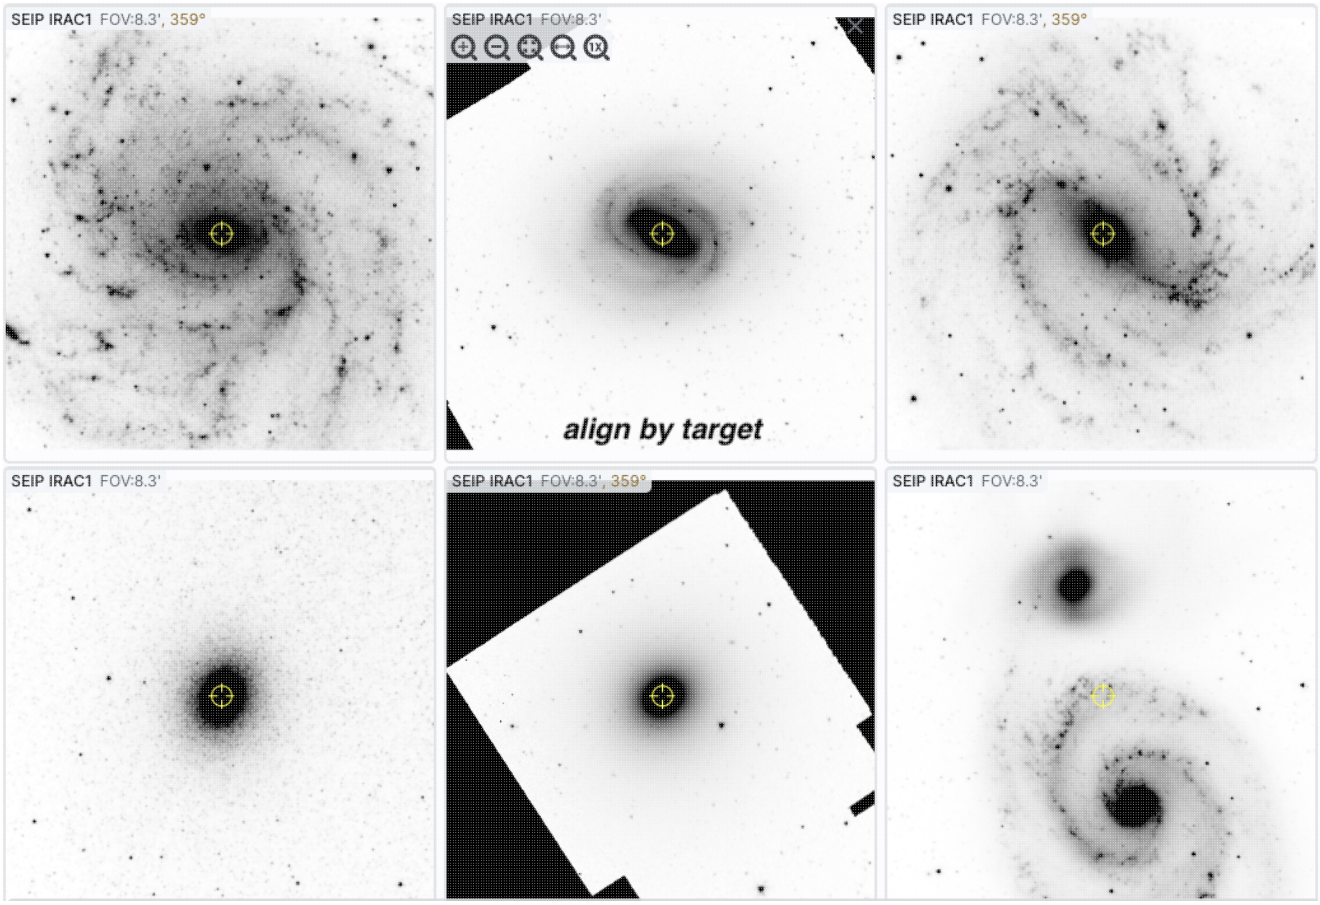

Aligning images by position on the sky is likely to be the most common use of locking. You can align FITS and HiPS images to each other.

Extraction Tools

Several tools allow you to extract information from images or image planes, but only for FITS (not HiPS) files.

-  -- Extract down through image planes
-  -- Extract a line from the image
-  -- Extract points from the image

All three of these follow the same basic structure --

1. Intitiate extraction mode
2. Set aperture
3. Try extraction; repeat if desired
4. Pin (retain) extraction if desired
5. Download (as table or chart) if desired
6. Repeat extraction if desired

7. End extraction mode

The drill pokes down through multiple planes, the line moves across pixels in a plane, and the points extracts points from a plane.

Here, we cover the basic approach, with specifics of each tool integrated as we go along.

Intitiate extraction mode. When you click on one of these icons, you enter into the extraction mode. Text appears next to the image toolbar to remind you that you are in this mode: End Extraction When you are done, to end this mode, click on this "end extraction."

When starting out, the pop-up window that you get depends on the tool you pick.

For the drill:

for the line:

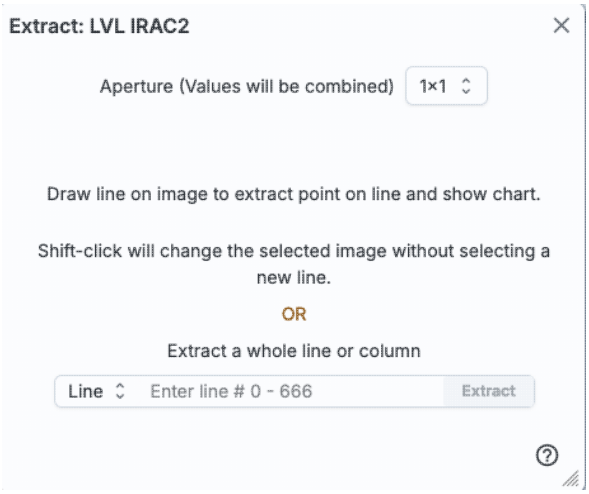

and for the points:

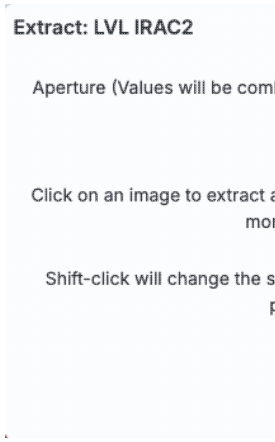

Set aperture. In all three cases, the top center of the extract pop-up window has a drop-down from which you can select the aperture value. For the drill and points, you can choose, in pixels, 1x1, 3x3, 5x5, or 7x7. Values can be summed or averaged. For the line, it's a little different. The apertures are 1x1, 1x3, 1x5, or 1x7, and you control whether the points are summed or averaged. You can have it extract along a line that you draw with your mouse on the image (useful for examining brightness profiles, e.g., across spiral arms of a galaxy), or you can have it extract an entire line or column of the image that you specify (useful in the cases where the 'image' is a file where the first row is an extracted spectrum, the second row is the error, the third row is a mask, etc.).

Try extraction. From this point, you can click on your image, or click and drag for the line tool. The pop-up then contains a plot of your extraction.

For the drill:

for the line:

and for the points:

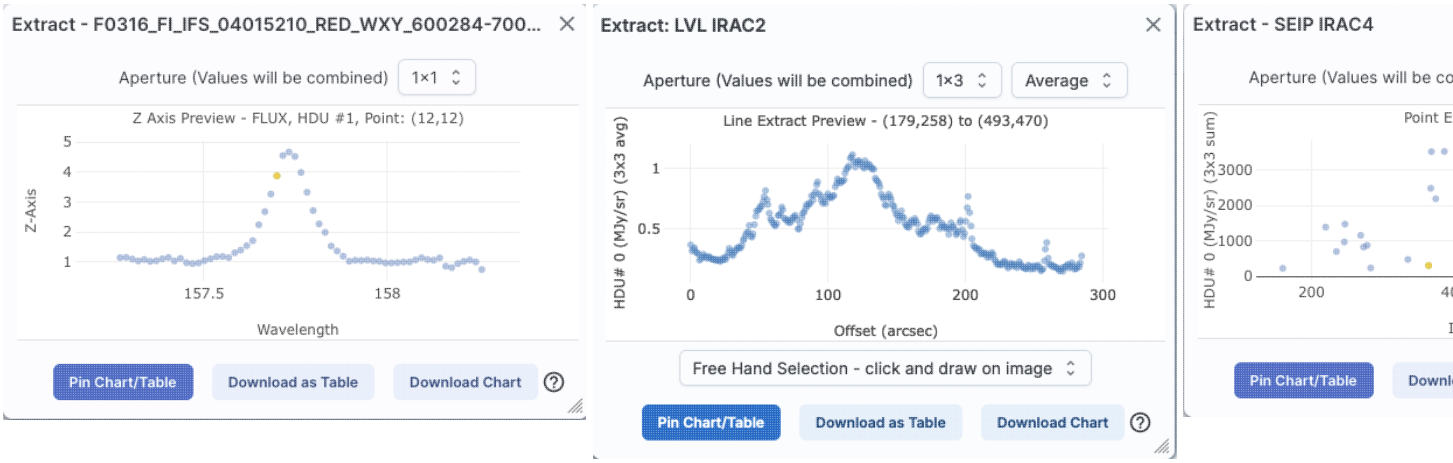

In each of these cases, it does the best that it can to label the axes accordingly. The cube that is used for the drill example here follows the spectral data model so that it is recognized as a spectrum, and it plots against wavelength. The line extraction is a 1x3 average and plots the offset in arcseconds from the initial mouse click, given the WCS information in the header of the image. The point extraction is a 3x3 sum, and plotted in image x coordinates, but can be changed to show image y coordinates.

Note that for the line, if you have more than one image loaded and visible, you can shift-click on a new image to see the same line on a new image. Similarly, for the points, you can shift-click to change images without extracting points. For the line extraction, if you want to change at this point to extraction along a line or column, use the drop-down menu at the bottom of the pop-up (shown here as "free hand selection").

Pin extraction. Once you have an extraction that you like, you can retain the extraction for further analysis. "*Pin chart/table*" extracts the information as a table, just like any of the other tables in this tool, with an accompanying plot. You can then manipulate the table/plot just like any other table or plot in this tool. If the tool recognizes the extraction as a spectrum, you may have additional capabilities.

Once you pin or save your extraction, the tool leaves a "footprint" of your extraction on the image so that you can remember what the extraction was. **NOTE THAT it is not interpolating** across fractional pixels here. It is averaging if you have asked it to average, but particularly if your pixels are large, if you draw a line that is diagonally across pixels, it will be immediately obvious that it's not interpolating. This line gets rendered as these pixels:

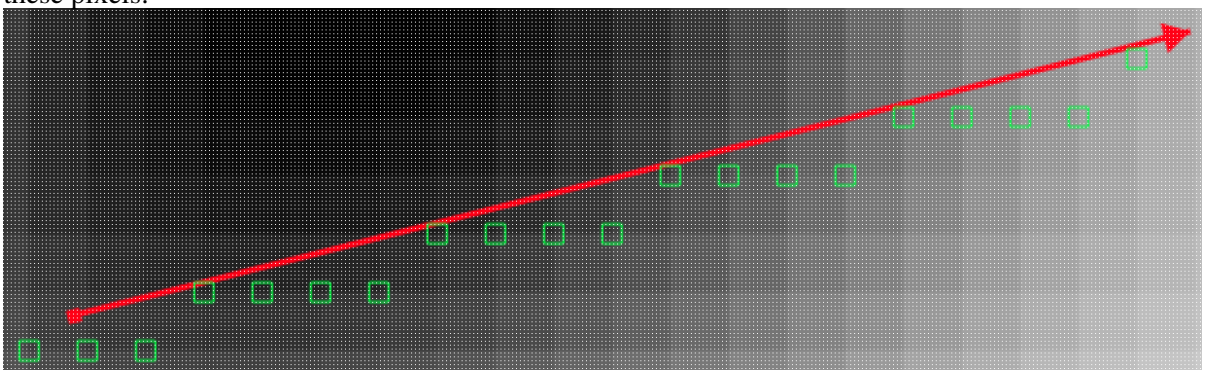

The point appears on the image at the lower left corner of the relevant pixel.

You can pin as many different extractions as you want. Each one will result in new tabs with the corresponding table at the bottom of the screen. There are navigation aids within the tables section that may help.

Download extraction. You can download the extraction as a table or plot without pinning it. *Download as Table* saves the table to your local disk with all the same options as a regular table. *Download Chart* saves the plot as shown, as a png file.

After pinning an extraction, you also have the extraction appearing as a table in the tables section of your window and as a plot in the plots section of your window. As with any table in this tool, you can save the table by clicking on the diskette icon in the table. You can choose from a variety of formats; see the tables chapter for more information. Similarly, as with any plot in this tool, you can save the plot by clicking on the diskette icon in the plot. See the plots chapter for more information.

Repeat extraction. As long as the extraction pop-up is still open, you can continue to click points or draw lines to make additional extractions.

End extraction mode. 

When you are done, to end the extraction, click on "end extraction" to end the extraction. Alternatively, just click on the 'x' in the upper right corner of the extraction pop-up.

Region Selection

 When you click this icon, you can select a region of the image, from which then you can do a whole host of things to the image and to the catalog you may have overlaid upon it.

First, from the drop-down, you are given a choice of a rectangular selection or an elliptical selection:

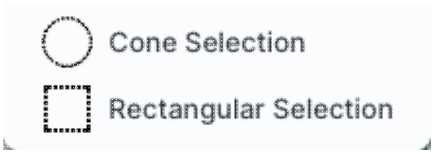

After you make that choice, you can click and drag in the image, selecting a box or a circle on the image. For the circle, your click is the center of the circle and the drag is the radius of the circle; for the box, your first click is one corner of the rectangle, and the drag is the opposite corner of the rectangle. This region can be resized by grabbing and dragging the corners of the box or the pixels delineating the corners of a box around your code. If you need to move the image underneath, you can hold the shift key and click and drag.

When you have selected a region of the image, additional icons appear above the image, and exactly which icons you see is a function of whether you are working on a FITS or HiPS image, and whether or not you have a catalog overlaid:  These icons allow you to do several things:

 *Crop the image*
(FITS only) Crop the image to the selected region. Then you can save the cropped FITS image via the save icon described above.

Note that, if you have a rotated FITS image such that a crop would have to bisect pixels, it will show you the region that encompasses your selection. If you crop at that point, then, it will crop in image space (such that pixels are not bisected). See the figure below -- in the original image, north is up. This has been rotated 45 degrees. The selected region is in white. The yellow dash-dot line is the crop in pixel space that encompasses the selected region.

Select sources (and *cancel selection*)
(Only if a catalog is overlaid) Select the catalog sources overlaid on the image within the region. Selecting highlights the sources in the list and plot with a different color row or symbol. Once there are selections made, the second icon appears to give you an option to cancel the selection.

Filter sources
(Only if a catalog is overlaid) Filter the overlaid catalog down to the sources within the enclosed area. When you choose to impose a filter via this selection mechanism; the filters icon changes above the catalog to indicate that there is a filter applied ( ). To clear the filters, click on the cancel filters icon (which also appears after you impose filters): . There is much more on filters in the Tables section.

Zoom the image
Zoom the image to fit the selected area into your field of view.

Recenter the image
Recenter the image on the selected area.

Obtain statistics
(FITS only) Obtain statistics from the image on the region. The statistics option results in a pop-up that looks something like this:

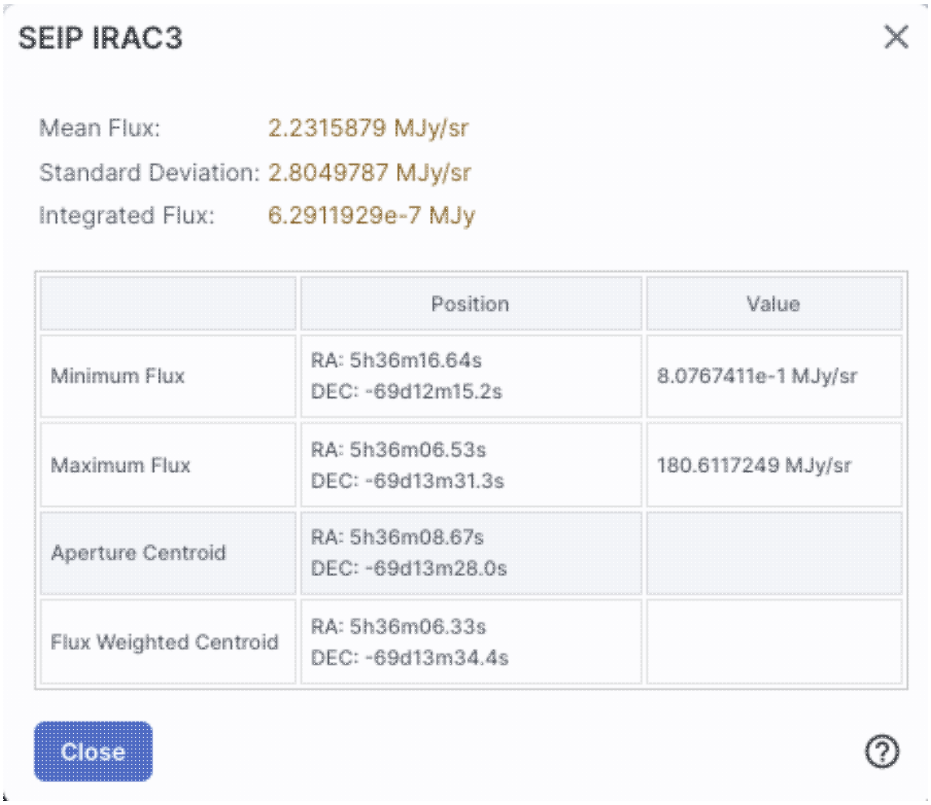

Note that it calculates the location of the minimum and maximum fluxes, and the aperture and flux-weighted centroids; the flux values given are in the same units as the FITS file. If you put your mouse over the row of the table in the pop-up, that location appears as an 'x' on the image.

Search

This tool implements a new search, an "action", on the region you have selected. It results in this drop-down (right). where this example is based on a region centered on 84.036131, -69.224431, J2000 decimal degrees, over a 4-cornered polygon. (You can also use the region tool to define a cone; this example happends to be a rectangle.) From this drop-down, you can launch:

- ◊ A TAP polygon search over this region (more information about TAP searches)
- ◊ A NED cone search at this position with a radius attempting to correspond to this polygon (more information about NED searches); results loaded into this tool.
- ◊ A Simbad cone search at this position with a radius attempting to correspond to this polygon; results loaded into this tool.

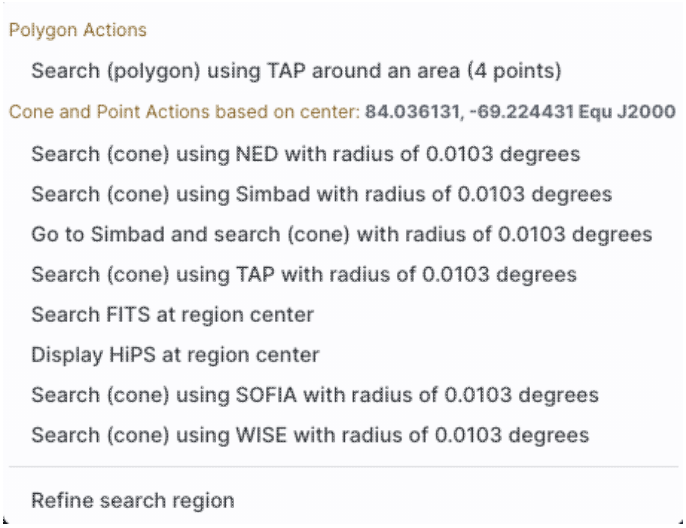

- A Simbad cone search at this position with a radius attempting to correspond to this polygon, but launch another browser window or tab at Simbad with the results.
- A TAP cone search at this position with a radius attempting to correspond to this polygon (more information about TAP searches); results loaded into this tool.
- A FITS image search at this position (at IRSA, via this tool); results loaded into this tool.
- A HiPS image search at this position (via this tool); results loaded into this tool.
- A cone search of the SOFIA archive at this position with a radius attempting to correspond to this polygon; results loaded into the SOFIA tool.
- A cone search of the WISE Image Service at this position with a radius attempting to correspond to this polygon; results loaded into the WISE tool.
- Refine the search region

The last option brings up another pop-up window (similar to this) that allows you to refine the search region iteratively by choosing a cone or polygon, setting the center, and setting the cone size or polygon vertices.

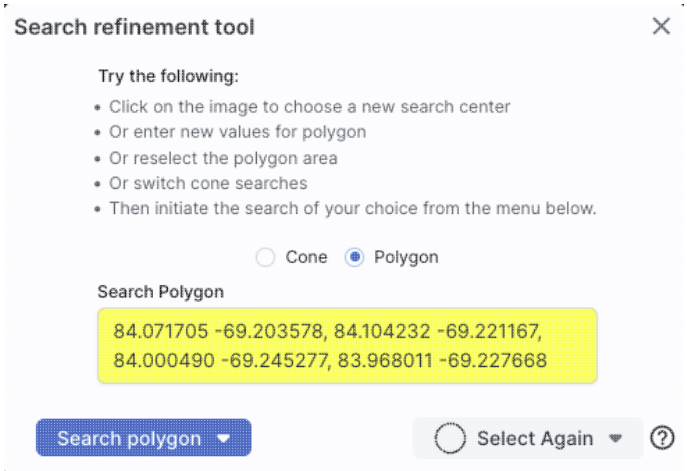

From here, you can change the kind of search, refine the positions, launch searches from your refined position (blue button on lower left), and select from the image again (drop-down on the lower right).

When you are working with the selection tool, this: End Select appears next to the image toolbar to remind you that you are in that mode. If you are refining positions for a search, it will be End Search Marker. Either way, when you are done with the selection tool, if your other actions don't turn off the selection tool, click on that text to turn it off.

Tips and Troubleshooting

- The "region selection" tool also appears in the slightly different context of interactive target refinement in several other places in this tool, where it works in a a very similar fashion.

Footprints

The marker icon () has a drop-down menu with several possible options:

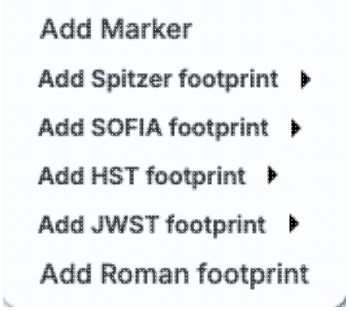

Any of the options with an arrow on the right can expand to additional subsidiary choices, e.g.,:

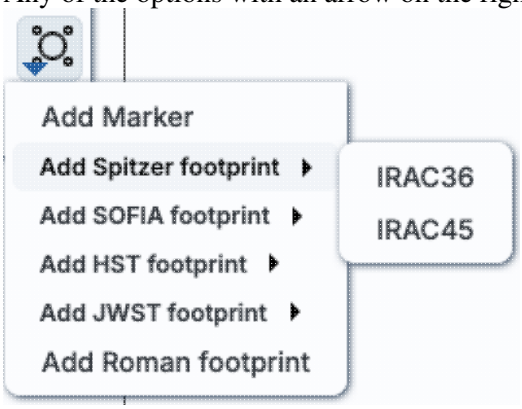

We now describe these various footprints here.

For each of these choices, the markers appear initially in the center of the loaded images. The first mouse click you make in any of the images will move the marker to that location.

Each of these marker choices, when overlaid and/or selected as 'active', has a dot-dash square around it. If it is asymmetrical (most of them are), it has an additional "appendage" and a red plus at the center of the footprint:

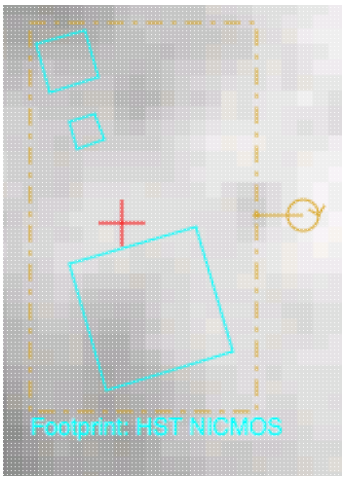

These so-called "handles" allow you to resize and/or rotate the marker, depending on the nature of the marker. These handles only appear when the marker is selected as active; if you wait a few seconds, they vanish.

Tips and Troubleshooting

- Some of these footprints are *large*. If you have a small image, some of these footprints will be larger than your image. Zoom out to see it, or find a larger image to use. If you overlay, say, a Nancy Grace Roman Space Telescope (formerly WFIRST) footprint on a 2MASS Atlas FITS image, you may need to zoom out a considerable amount before you can see the Roman footprint. You will see the center indicator of the marker before you will see the Roman footprint itself.
- You can add multiple copies of the same marker using the layers pop-up (described generally above). From the layers pop-up, there is a link right under the 'angle' option that says "Add another [marker type]" -- click on that to get an additional marker of the same type. You can also add a label to the marker from the layers pop-up, or change its color.
- If you have many footprints on the same image, you may have trouble grabbing and moving footprints lower in the stack of layers on the image. For example, overlay footprint 1, then footprint 2, and you might have a hard time grabbing and rotating footprint 1 after footprint 2 has been added. The only workaround here is to use the layers pop-up (described generally above) to temporarily hide footprint 2, then move footprint 1, then restore footprint 2.
- If you have images of very different resolutions loaded (e.g., IRAS and really anything else), sometimes it struggles to render the marker on each image. You may need to place markers on one image at a time. (Unclick the "lock color & overlays" option to place markers one image at a time.)

The first overlay choice (simply called 'marker') is a red circle.

The remaining markers are all footprints from various telescopes: Spitzer, SOFIA, HST, JWST, and Roman. HST, JWST and Roman are derived from information provided via MAST (see http://gsss.stsci.edu/webservices/footprints/help.html ⧉.) For Roman in particular, they are pre-launch values.

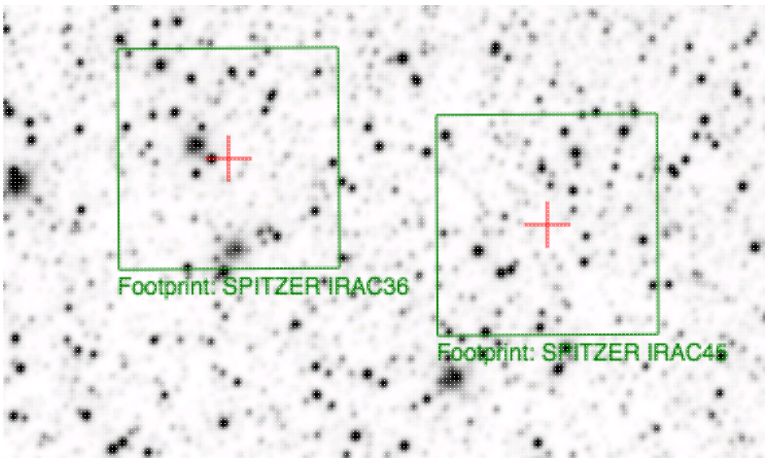

Spitzer/IRAC 3.6 and 4.5 micron footprints. These two footprints are placed separately from each other. The footprint can be moved or rotated. Click and drag the center of the footprint. A circle appears with four small circles ("handles") around it. Grab and drag the small circles to rotate it, or drag the big circle to move it. Change the color, delete, or add more copies of the IRAC footprints from the layers pop-up.

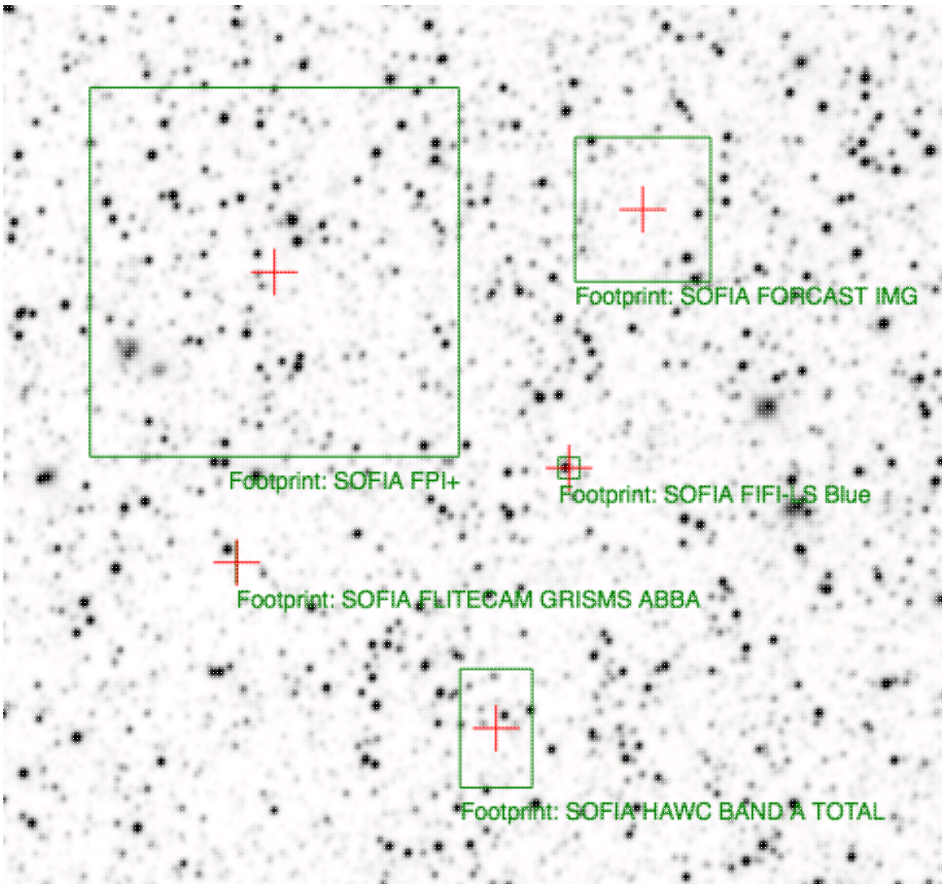

SOFIA footprints. Several different SOFIA footprints are available; the graphic here shows a selection of them. The available footprints (all of which are placed separately) are:

- FIFI-LS
  - Blue (50-120 microns)
  - Red (110-200 microns)
- FLITECAM
  - Imaging

    - Grism ABBA
    - Grism AB
- FORCAST
    - Imaging
    - Grism a
    - Grism b
- FPI+
- HAWC+
    - 53 microns (Band A), Total Intensity
    - 53 microns (Band A), Polarization
    - 89 microns (Band C), Total Intensity
    - 89 microns (Band C), Polarization
    - 154 microns (Band D), Total Intensity
    - 154 microns (Band D), Polarization
    - 214 microns (Band E), Total Intensity
    - 214 microns (Band E), Polarization

Change the color, delete, or add more copies of the SOFIA footprints from the layers pop-up.

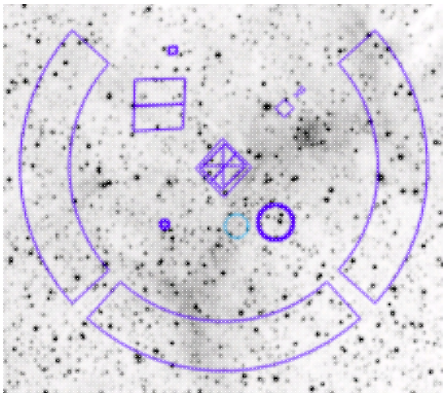

HST footprints. You can overlay the whole focal plane footprint, shown here, or individual instrument footprints (NICMOS, WFPC2, ACS/WFC, ACS/HRC, ACS/SBC, WFC3/UVIS, and WFC3/IR). Consult the HST documentation for specifics on which apertures are which. The footprint can be moved or rotated. Click and drag the center of the footprint. A circle appears with four small circles ("handles") around it. Grab and drag the small circles to rotate it, or drag the big circle to move it. **Note that** if you overlay the footprint on a very small image, nothing will appear to have happened. You need at least a 45 arcmin image to comfortably see the footprint. Change the color, delete, or add more copies of the HST footprints from the layers pop-up.

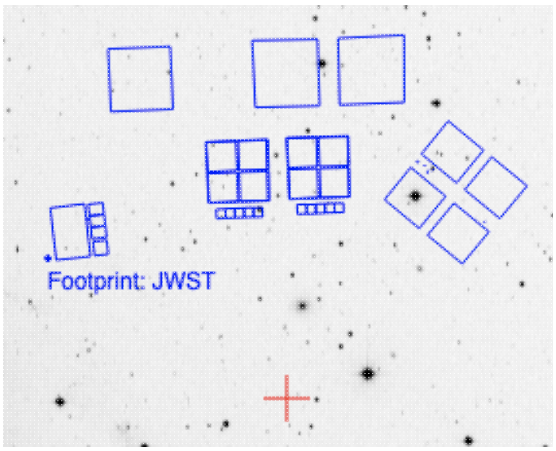

JWST footprints. You can overlay the whole focal plane footprint, shown here, or individual instrument footprints (FGS, MIRI, NIRCAM, NIS, and NIRSPEC). **Note that** if you overlay the footprint on a very small image, nothing will appear to have happened. You need at least a 30 arcmin image to comfortably see the entire JWST focal plane. Please consult the JWST documentation for details about the footprints. In all cases, if the footprint is 'active', a circle near the middle of the footprint will appear with four small circles ("handles") around it. Grab and drag the small circles to rotate it, or drag the big circle to move it. Change the color, delete, or add more copies of the footprints from the layers pop-up.

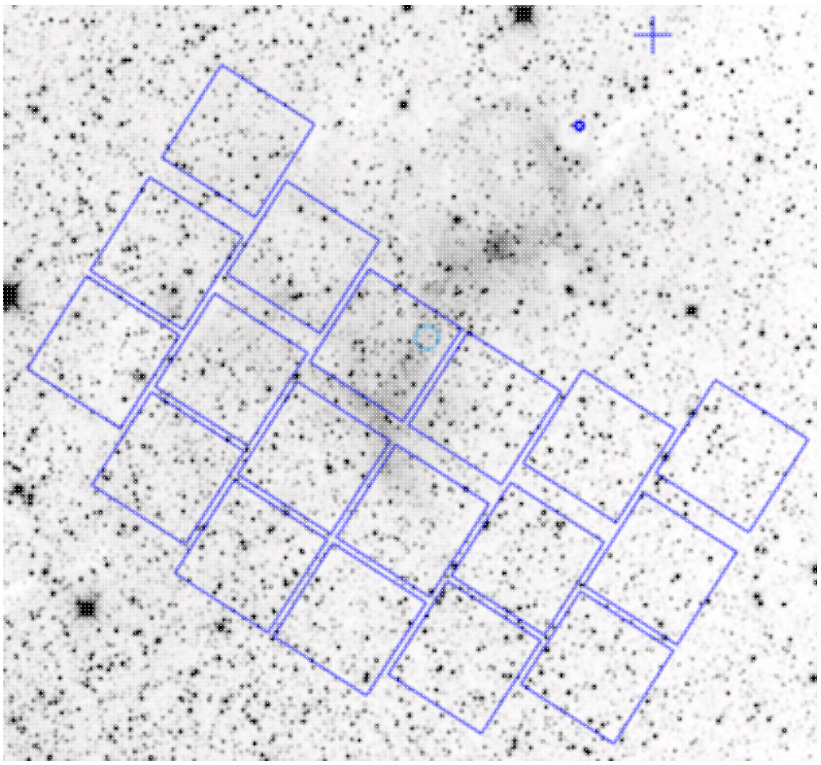

Nancy Grace Roman Space Telescope focal plane footprint. As above, the footprint can be moved or rotated. Click and drag the boresight (the cross hairs), which appears by default to the upper right of the array of squares. A circle appears, centered on the boresight, with four small circles ("handles") around it. Grab and drag the small circles to rotate it, or drag the big circle to move it. **Note that** if you overlay the footprint on a very small image, nothing will appear to have happened. You need at least a 60 arcmin image to comfortably see the footprint, and even then you will probably have to click and drag to see the entire footprint. Consult the Roman documentation for specifics on the apertures. Change the color, delete, or add more copies of the Roman footprint from the layers pop-up.

IRSA Viewer: Tables

All of the tables in IRSA Viewer (whether they are catalogs, or a list of the loaded images, or the contents of a FITS header) are interactive tables, and they have the same basic properties, discussed in this section. The specific broad case of catalogs is in another section.

Contents of page/chapter:

Table Header

These interactive tables -- called Firefly tables, after the software that is running here -- all have the same functionality, regardless of the contents of the table. If it looks like this sort of table, you can work with it whether it is a catalog, an image header, a list of images, etc.

If you have loaded a table into a tab and it doesn't take up the whole screen, to see more of the window, grab the divider between the window panes and slide it up/down or over as needed, or use the expand arrow icon ( ) to enlarge the window pane to take up the whole window.

The table is shown exactly as it has been provided to the tool, with all columns as defined by the creator of the table. To understand what each column is, please see the documentation associated with that catalog or table.

The tab (and table) name itself likely indicates its origin. To remove the tab, it is likely that you can click on the "X" on the tab.

At the top of the table, there can be several symbols:

which we now describe, going from left to right along the top of the catalog tab.

Table navigation
The first thing to notice is that (typically) only the first 100 rows of the retrieved catalog (or table) are displayed in the table. In the example here, there are 319 sources that were retrieved as a result of the search, grouped into 4 'pages.' The left/right black arrows plus the page number allow you to navigate among these 'pages' of 100 sources each. Note that the entire set of results (not just the 100 rows you are currently viewing) can be sorted by clicking on any column's name.


Table Actions: Searches
This drop-down has choices to launch new searches, and as such, it has a separate section below.

Filter
Filters are complex and powerful enough that they are covered in a separate section below.

Table as text
Clicking on this changes the table display into a fixed-width text display. The icon then changes to click this again to return to the default table view.

Save table
This is how you can save the table. It has a separate section below.

Add a column
This icon adds a new column to the table. This has a separate section below.

Info link
You may or may not see this icon. This is an "information" button and, if it exists, it may provide additional information about the table. It could have information about the job that was used to retrieve it:

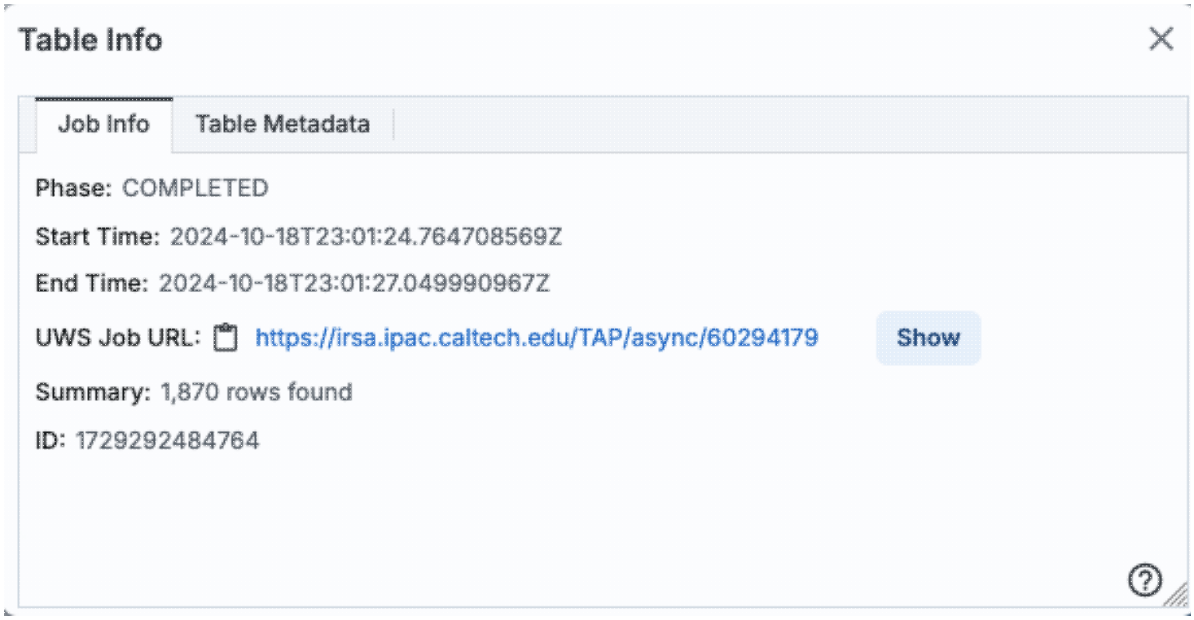

where the direct link to the job is given there (and can be copied by clicking on the clipboard, ready to be pasted into a helpdesk query, for example), with a job id as shown. It could also just have information about the table metadata:

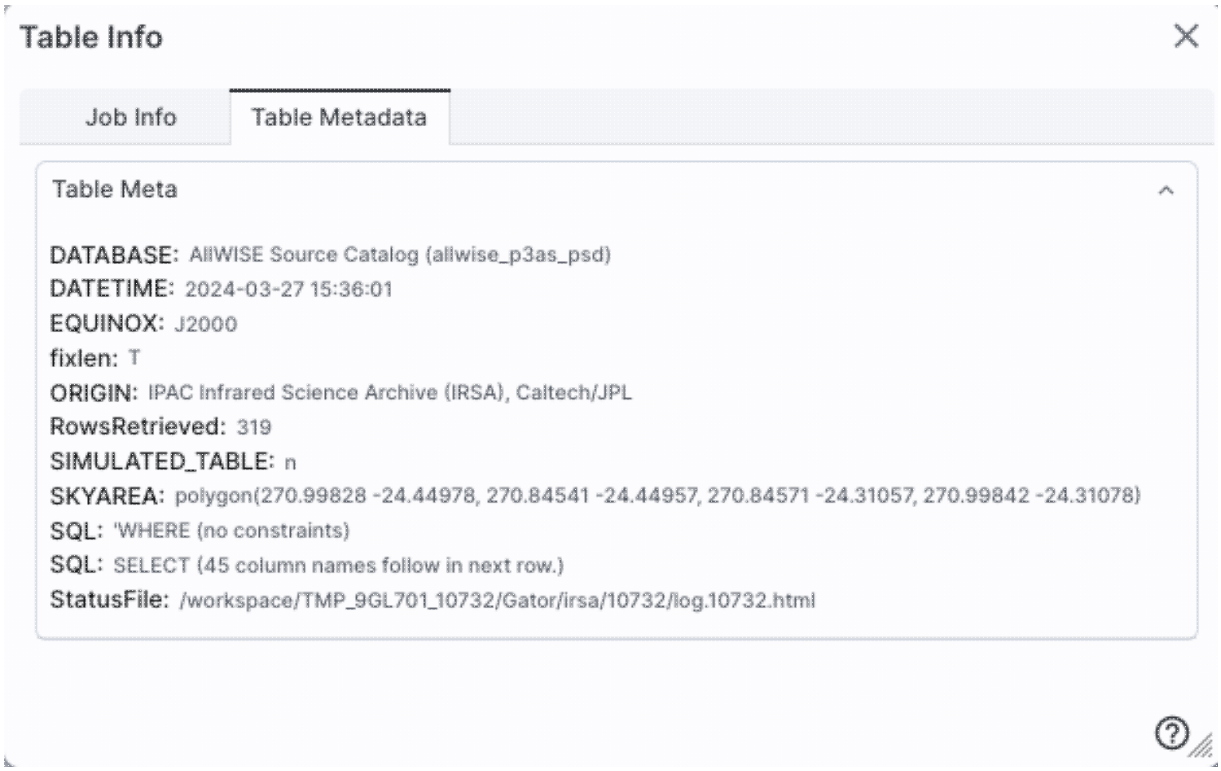

where the information about this table includes information about the query that produced it.

Row details

You may or may not see this icon. This is how you get more information about the currently-selected row. It has a separate section below.

Table options

Clicking on this icon brings up options for the table, e.g., how many rows are displayed per page, which columns are shown, metadata about each column if available, whether units and data types are shown at the top of the column, shown here:

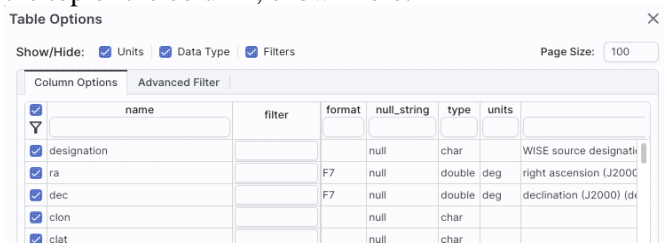

By default, it is often but not always the case that all columns are shown. To show or hide columns, select the tickbox in each given row. The default page size is 100 rows. Note that expanding the page size to numbers much greater than 100 may result in a substantial performance degradation (e.g., your browser will appear to freeze or not appear to be doing anything while it manages and renders the large table). See the Filters section below for more things to do from the table options pop-up.

Expand

Clicking on this expands the catalog window pane to take up the entire browser window. To return to the prior view, click on "Close" in the upper left.

Help

The last option on the top of the catalog tab may be a context-sensitive help marker, which should bring you to this online help.

Table Columns

Depending on what you did to display a table, the columns that are shown may be in easily-human-readable form, or may reflect column names used within the individual catalog. Please consult the detailed documentation associated with your specific table if the headers are not clear to you.

The table is shown exactly as it appears in the corresponding database (or as it appeared on your disk), with all columns as defined for that catalog. To understand what each column is, please see the documentation associated with that catalog. (For IRSA catalogs, this documentation is available via navigating through the IRSA website.)

Clicking on the column names sorts the table by that column; clicking once sorts in ascending order, clicking a second time sorts in descending order, and clicking a third time returns the table to the original order. Small arrows appear next to the column names to remind you if the column is sorted in ascending or descending order. When you do a single-position search on catalogs, depending on how you do it, it could be that two new columns are appended to the end of the catalog as it is returned to you. These columns are:

- dist - the distance between the source in question and the location you specified
- angle - the position angle between the target position you requested and the object it found (degrees E of N)

When you do a multi-position search on catalogs, you could have three new columns prepended to the catalog as it is returned to you. These columns are :

- cntr_01 - the target position you requested
- dist_x - the distance between the target position you requested and the object it found
- pang_x - the position angle between the target position you requested and the object it found (degrees E of N)

These additional columns can help you assess if the target it found is the target that should be matched to the position you requested.

It could be that, when you do a multi-position search on catalogs, you have an option for "one-to-one matching". If that is selected, the line (and only one line) of output is included for each line of input. It chooses the closest source within the radius you specify, or if there is no match, it adds a line indicating no match.

Tips and Troubleshooting

- You can hide or display columns; click on the gears () to get to the table options, and tick the box corresponding to the row you want to hide or show.
- If you are constructing and loading your own catalog where you don't specify the formatting, the tool tries to guess the formatting of the column based on the first row's values. Therefore, if you have a value in the first row that happens to be a string like "null" where the rest of the column isn't like that, or happens to have only one decimal place where the rest of the column has 2 or more, it will guess incorrectly and format the other rows just like the first one. Try editing your catalog file to put the most general case as the first row.

Adding Columns

This icon  allows you to add a new column to your catalog. When you click on it, it brings up this pop-up window:

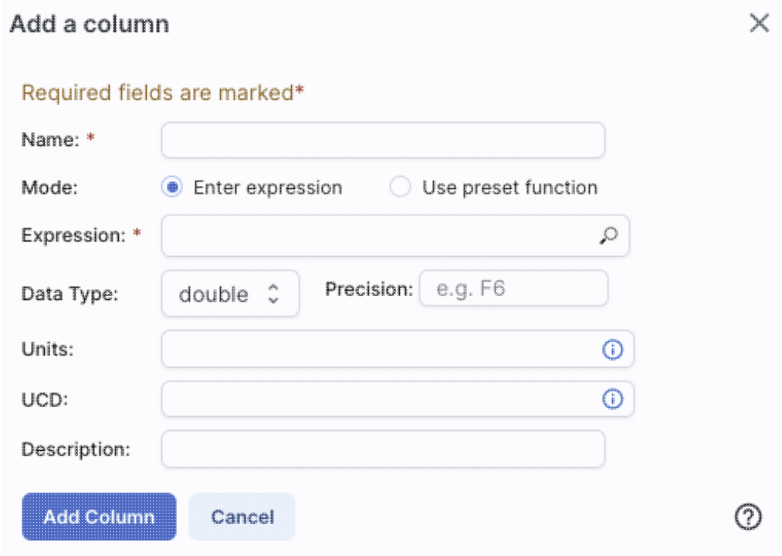

This window asks for:

- *Name* of the column (required) -- it cannot have special characters like a minus sign or a percent symbol; you can only use letters, numbers, and underscores.
- *Mode* "Enter expression" or "Use preset function" -- options shown here correspond to "Enter expression" options; the preset function options are included in the last bullet here.
- *Expression* (required) -- using basic SQL operators, you can manipulate columns to create the new column. (See more on this below.)
- *Data Type* -- specify if your new column is a double precision floating point ("double"), a long integer ("long"), or a string ("char").
- *Precision* -- if you have selected "double" for data type, select how many decimal places your new column should display. For example, if you want the numbers to display as 1.23, enter "F2".
- *Units* -- specify the units of your new column. For more information, see IVOA documentation
- *UCD*, or unified content descriptor -- for VO compliance, add this for your new column. For more information, see IVOA documentation
- *Description* -- add a description for your new column.
- *Select a preset* -- if you select "Use preset function", you can choose (a) "set filtered rows to 'true' and the rest to 'false'", (b) "set selected rows to 'true' and the rest to 'false'", or (c) "number rows in current sort order". These options are useful for tagging items you have selected in myriad ways (from plots, images, complicated filters), or ordered in complicated ways. For example, if you have constructed a complicated filter, then you can create a column that is true for the selected rows; if you cancel the complicated filter, you can then easily recreate the complicated filter by simply filtering on your newly created column.

In order to construct the expression for your new column, your input should follow the syntax of an SQL expression. If you click on the magnifying glass next to the form input, you get a pop-up window that can help you construct an expression; click "apply" to apply the expression.

You need to use the column names exactly as they appear in your catalog. Supported operators are: +, -, *, /, =, >, <, >=, <=, !=, LIKE, IN, IS NULL, IS NOT NULL. You may use functions as well; for a list of all available

functions, see here. Some examples include:

- "w3mpro" - "w4mpro"
- sqrt(power("w3sigmpro",2) + power("w4sigmpro",2))
- ("ra"-82.0158188)*cos(radians("dec"))
- "phot_g_mean_mag"-(5*log10(1000/"parallax") - 5)

Tips and Troubleshooting

- When you create a new column that is calculated from other columns, it is created statically. That is, it is not dynamically updated like a spreadsheet, but calculated once and left alone after that.
- When you create a new column, the header of the new column is red to let you know that the column is not present in the original catalog.
- When you save the catalog, the header of the saved catalog indicates that you have added a column. When you load the catalog back into the tool, the header of the new column is still red.
- You can edit or delete columns after you have created them; click on the gears ( ) to get to the table options, and then click on the edit icon to bring up a dialog box to edit or delete the column.
- You can hide columns; click on the gears ( ) to get to the table options, and tick the box corresponding to the row you want to hide or show.
- If you create a new column that turns a floating point column into an integer with the "FLOOR()" function, you need to be sure to set the resulting column type to "long". For example, if you have data covering several days or years, and you have a column that is a floating-point MJD, you can convert it into an integer, e.g., via FLOOR(mjd) for days or FLOOR(mjd/365.24) for years, then you can use the drop-down filter menu for the new column to quickly compare different time ranges. But, the new column must be an integer (e.g., "long") in order for this to work properly.
- If you need to, say, take the square root of a column that occasionally has a negative number, and you want it to attempt to handle this in a physically reasonable manner, you can construct expressions like this for a column named 'col' that has some positive and some negative numbers:
  if("col">=0,sqrt("col"),-sqrt(-"col"))

Table Filters

Filters are a *very* powerful way of exploring the table full of search results. Click on this icon in order to start the process of adding filters. A text entry box appears above each of the current catalog columns, with a small version of the filter icon corresponding to that row on the far left. You can type operators and values in these boxes -- hit return or tab after typing or click in another box to implement the filter.

Example: From a catalog, show only those sources with declination above a certain value (say, 31 degrees), type "> 31" in the box at the top of the "dec" column. Or, if you have retrieved a WISE catalog and would like to only view the objects with a W1 (3.4 micron) profile-fitted magnitude less than 6 magnitudes, in the box at the top of the 'w1mpro' column, type "< 6" in the form.

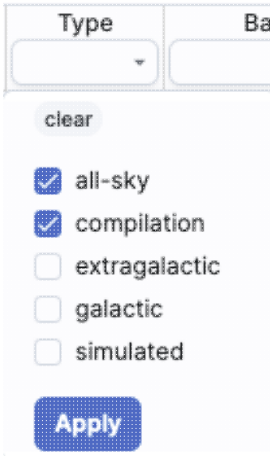

For columns (fields) with a limited set of choices, on the right edge of the text entry box, an arrow appears; click on it to get a drop-down from which you can select the available choices. To implement the filter, make the choices, and click "Apply" when you are done. Click "clear" in the top of the drop-down menu to remove that filter.

After you impose a filter, then the number of rows in the table is restricted according to the rules you have specified, and the "filters" icon on the top right of the catalog pane has changed to remind you that there has been a filter applied, in this case four filters:  . To clear the filters, click on the cancel filters icon (which also appears after you impose filters):  .

Filters can be used in combination. Note that the filters between columns are logically "AND"ed together, but filters within the same column can be logically "AND"ed or "OR"ed together; examples are below.

The available logical operators are :

- = which means 'equal to' (exactly!), e.g., the parameter on which you are querying (the column headers as shown) is exactly equal to this value you are specifying.
- > which means 'greater than'
- < which mean 'less than'
- != which means 'not equal to' (exactly!)
- >= which means 'greater than or equal to'
- <= which means 'less than or equal to'
- IN which means 'included within this list', e.g., the parameter on which you are querying is included within the list you are specifying (if the column filter is free-form text, type "value1,value2" and it will give you rows that have value1 or value2).
- LIKE which means 'resembles the text that is entered', e.g., the text resembles the text that you type in the box.
- IS which effectively is the same as =
- IS NOT which effectively is the same as !=

Examples:

- Retain rows for which a certain parameter is not an empty string: !="
- Retain rows for which a certain parameter is not NULL and is larger than 1.234: > 1.234 and IS NOT NULL
- Retain rows that have values between -0.5 and 1.25: > -05 and < 1.25
- Retain rows with a parameter greater than one value or exactly not equal to another value: > 12345 or != 3000
- Retain rows with a parameter equal to one of the values in a list: IN a,b,c,d

You can also interactively impose filters from plots from a catalog. Moreover, all the same operators that are available for making plots can be applied in filters. In both cases, see the plots section for more information.

You may also be able to select rows one at a time via the far left column and then filter that down. Example: Retrieve a catalog of any sort. Select rows by ticking the box on the far left, say, every other row out of the first 12. Click on the filter icon on the top of the column. The filter is imposed to only include the 6 rows you selected.

If you click on the table options icon ( ), you get a pop-up that includes a place to filter columns. Here is an example of the table columns for a WISE catalog that has been filtered to just have SNR>10 for each of the four bands:

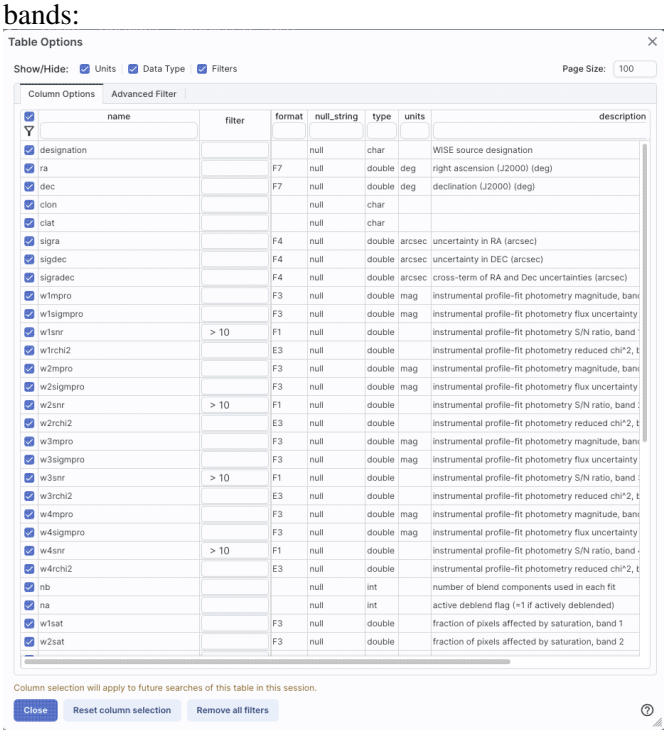

You can type in constraints in the filter box in much the same way as you can from the catalog itself; note that the column description is included here, which may make this way of setting filters more useful when working with a new (to you) catalog.

The second tab is the advanced filter interface:

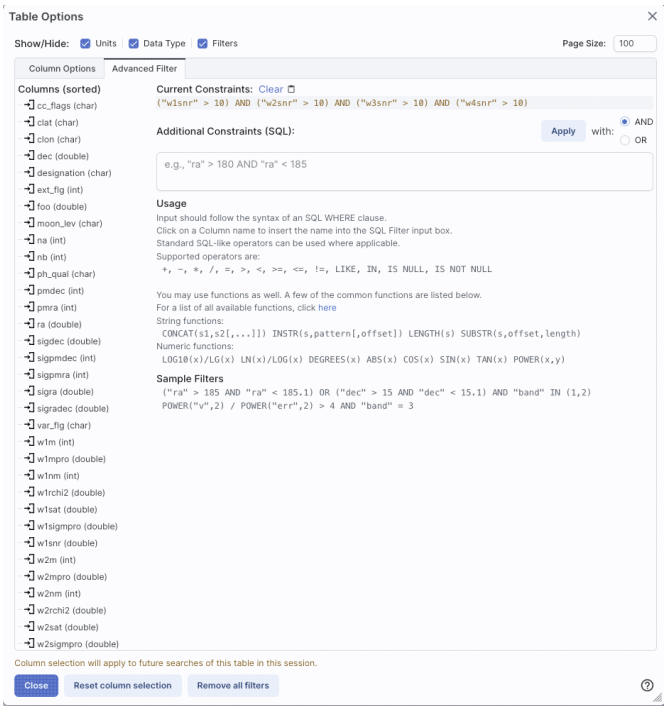

Here, the columns are listed alphabetically on the left, the constraints you have imposed are in the "current constraints" box, and you can apply additional SQL constraints via the free-form text box. Hints for syntax are given below the entry box. This window can be resized such that you can see the whole set of hints and imposed filters. Note that in this interface, column headers must be enclosed within double quotes.

Cancelling filters

After you impose a filter, then the "filters" icon on the top right of the catalog pane has changed to remind you that there has been a filter applied, in this case just one filter:  To clear all the filters at once, click on the cancel filters icon (which also appears after you impose filters):  .

Tips and Troubleshooting

- If you impose logically inconsistent restrictions such as "exposure_time > 160" and "exposure_time < 100" (">160;<100" in the filter box for a column called "exposure_time"), no data will result, because no data are (can be) both less than 160 seconds long and greater than 100 seconds long at the same time. However, "exposure_time > 160 or < 100" works just fine (">160 or <100" in the filter box for the column "exposure-time").
- If you impose nonsensical filters (like using a letter where a number should be, like "w1snr < f") then it will let you know that something has gone really wrong, and let you go back to fix it.
- However, a filter like "ph_qual < f" could be valid -- if the column with which you are working is a string, then a string is a valid filter. It is case-insensitive. For the ph_qual column in the 2MASS catalog, the values are always three letters, such as AAA or ABA or CUU. A filter like "ph_qual < f" will operate as if you have alphabetized the list. Any string that starts with A comes before F and so will be retained. If you do "ph_qual < BBB" then AAA will be left in, but so will "BAU", because alphabetically, BAU precedes BBB.
- If you impose filters from a plot, it can manifest as several filters on the catalog, e.g., one for each side of the square you have drawn on the plot. If you want to remove, say, just one of the four filters (rather than all of them by cancelling all filters), you can do so from the table options pop-up.

- If you want to copy all of the "current constraints", even if the entire field is not visible to you, you can click on the clipboard icon to copy the entire string, and then paste it into another field or application to see what it is.
- If you are choosing filters from a list of terms, cancelling those filters might not work the way it cavalierly seems like it should. If a down arrow appears next to the filter box, then a list of options you can select is available. Tick the boxes you want, and click "Apply" to apply the filter. Now, if you want to change the filter, click the down arrow, select different options, and click "Apply" again to re-impose the new filter. To remove the filter, you have to treat it like you would when applying a modified second filter -- unselect the choices, or hit 'clear', and **then hit Apply again**. If you deselect the choices and then click elsewhere in the window without hitting clear, your actions are interpreted as 'cancel without doing anything' as opposed to 'impose the new filter I just set' (which is 'cancel all filters'). When you are resetting the filter to be 'select nothing', it is treating that as a new filter setting, so you need to set up that filter and click 'Apply' in order for it to understand.
- If you impose filters from a plot, image, or table, you can cancel them from a plot, image, or table. That is, if you impose filters from a plot, and different additional filters from the image, when you click the 'cancel filters' icon from the table, *all* the filters are canceled at once. If you have filters imposed from multiple places, clicking on 'cancel filters' doesn't cancel just the filters imposed from that place; it cancels all of them. If you want to be able to reconstruct a complicated set of filters, though, you can add a column to your table that is one of the preset functions -- set it to true if the row is filtered. Then you can cancel all your filters at any time but you can easily recover the filtered data by filtering on that one new column.

Table Actions

 This icon implements a new search, an 'action', on the currently selected row of the table. When you click on it, it reveals a drop-down:

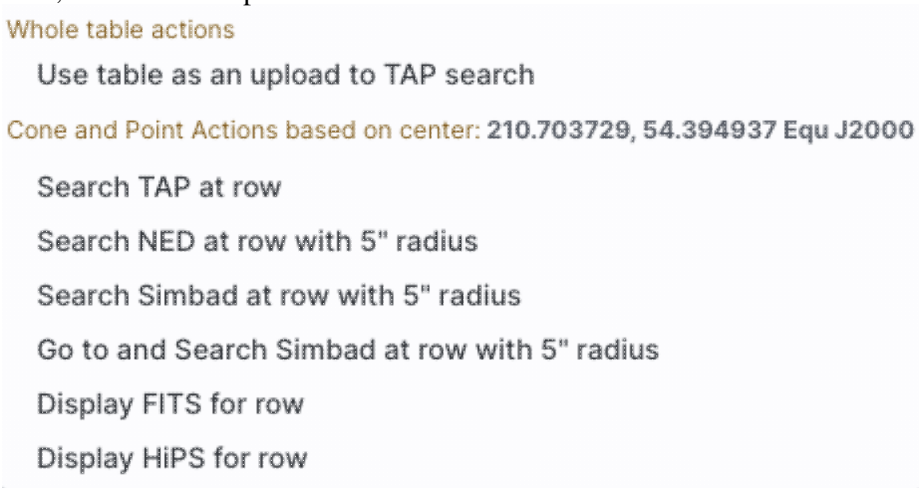

For this example, the selected row is 210.703729, 54.394937 in J2000 decimal degrees. From this drop-down, you can:

- Send the entire table to a TAP search
- Launch a TAP cone search at this position (more information about TAP searches)
- Launch a NED search at this position with a 5 arcsecond radius (more information about NED searches)
- Launch a Simbad search at this position with a 5 arcsecond radius, and put the results in a table here in the tool
- Launch a Simbad search at this position with a 5 arcsecond radius, but start another browser window or tab at Simbad with the results

- Launch a FITS search at this position (at IRSA, via this tool)
- Launch a HiPS search at this position (via this tool)

Any of these searches (except the search that launches another browser window with the Simbad results) places the search results into this same tool.

Row Details

 This icon is not always available. When it is available, when you click on it, a new pop-up window appears with information about the row you have selected:

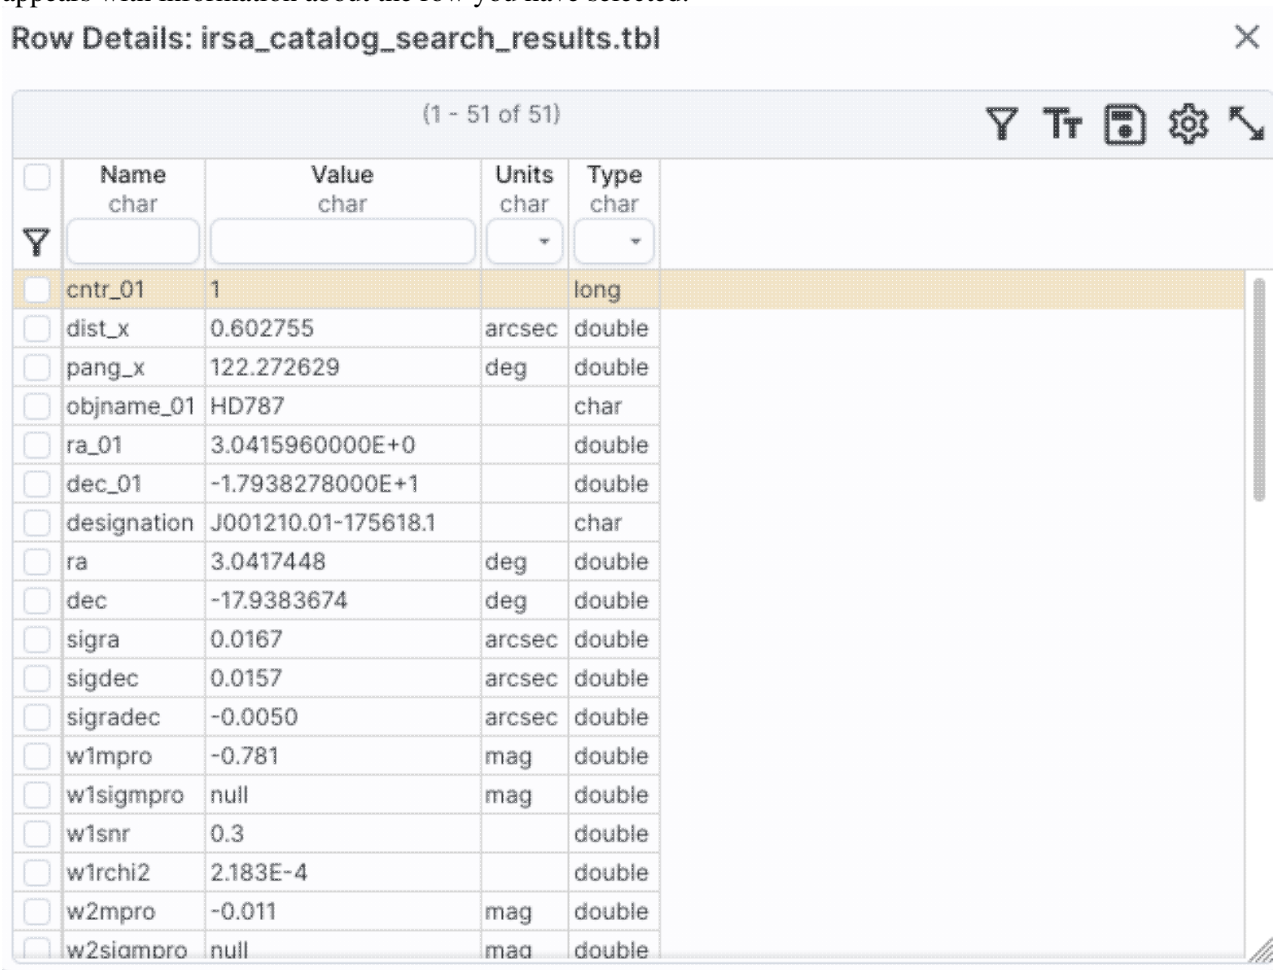

Name char	Value char	Units char	Type char
cntr_01	1		long
dist_x	0.602755	arcsec	double
pang_x	122.272629	deg	double
objname_01	HD787		char
ra_01	3.0415960000E+0		double
dec_01	-1.7938278000E+1		double
designation	J001210.01-175618.1		char
ra	3.0417448	deg	double
dec	-17.9383674	deg	double
sigra	0.0167	arcsec	double
sigdec	0.0157	arcsec	double
sigradec	-0.0050	arcsec	double
w1mpro	-0.781	mag	double
w1sigmpro	null	mag	double
w1snr	0.3		double
w1rchi2	2.183E-4		double
w2mpro	-0.011	mag	double
w2sigmpro	null	mag	double

In some tools, like IRSA Viewer, this content appears as additional tab elsewhere in the tool (not in the table pane, but often viewable at the same time as the table itself), as an additional tab called "Details."

In either incarnation, this information is sometimes called a "property sheet."

This table consists of each of the columns of the retrieved catalog with additional information about each field where available. (Not every catalog may have this information available.) For additional information, please consult the full documentation that accompanies the catalog.

Note that if you leave the pop-up or tab open as you select different rows in your main table (or catalog), it is dynamically updated.

Tips and Troubleshooting

- The property sheet is a more expanded, vertical view of the information shown in a row of a catalog, along with documentation of the catalog columns. If you think of the main table (catalog) view as a single row that you have to scroll left and right to see in its entirety, this view is sort of an orthogonal view, where you have all of the same contents of that row but shown as its own table, with the full header description, if available, and you can scroll up and down to see the entire contents (as opposed to left-right). *This is sometimes a more user-friendly way to view any given row.*
- Because you can sort/filter the data in the property sheet, you can restrict what values are shown. Those filters are respected as you page through your main table.
- When changing rows in the main table, the property sheet scrolls to preserve the visibility of whatever row in the details tab is highlighted. If you scroll down in the property sheet *without changing the highlight*, when you change rows in the main table, because the first row in any table is always highlighted by default, the property sheet will scroll back to the top.

Table Cells

Some cell values may be too long for the cell space. If that is the case, an ellipsis ("...") will appear in the cell as you mouse over it:

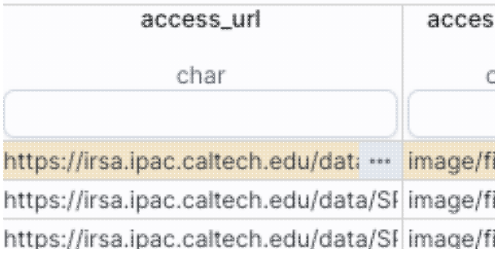

If you click on the ellipsis, you get a drop-down menu:

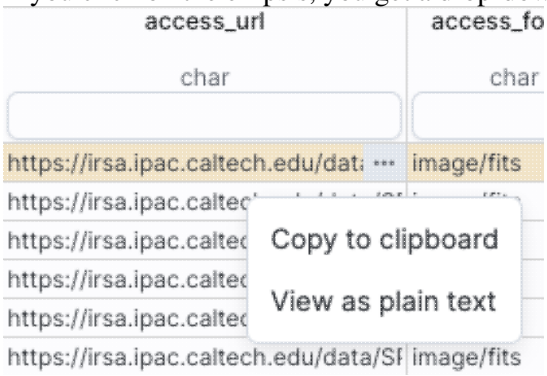

from which you can choose to copy the cell value or view it in a pop-up window.

If you view it in a pop-up window, it will attempt to format it in a readable fashion:

which, for most cells, will be functionally indistinguishable from no special formatting. It will only matter for very complicated cells.

Saving Tables

Click on the diskette icon ( ), if available, to save the table.

You obtain this pop-up:

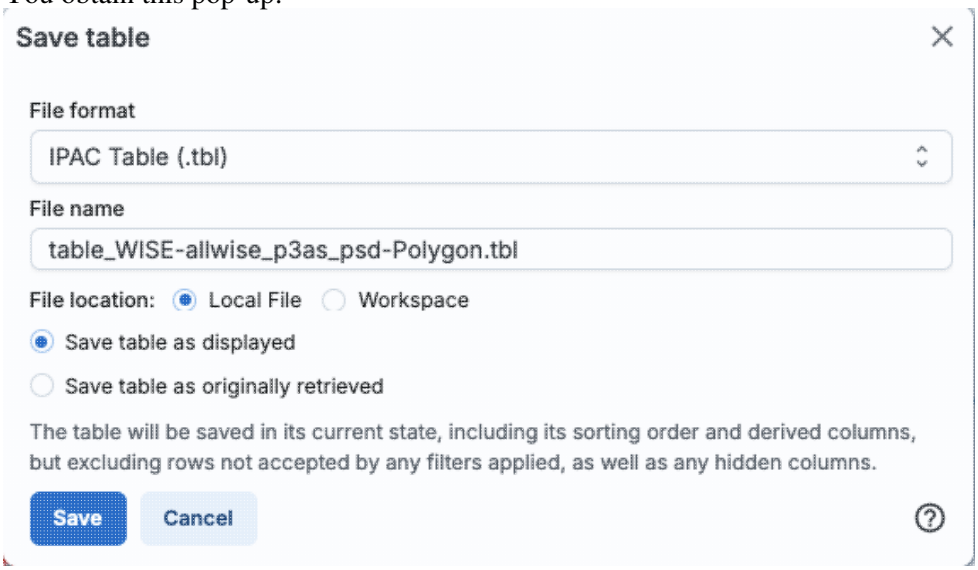

You have several choices to make. In order, they are:

File format

You can save the table in a variety of formats:

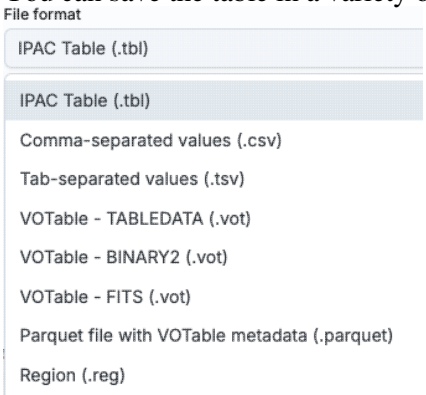

It defaults to saving it as an IPAC table file, which is basically ASCII text with headers explaining the type of data in each column, separated by vertical bars.

Other formats include comma-separated values (csv, suitable for, e.g., Excel), tab-separated values (tsv), and three different versions of VO tables. You can save the file in parquet file format, which is a highly efficient, compressed, column-oriented format for tabular data that has been adopted by many recent wide area survey projects. You can also save the file in DS9 Regions file format. The advantage of saving it here as a regions file (as opposed to from the visualization) is that this way, the entire catalog is guaranteed to be saved.

File name

The tool tries to make a guess at a sensible filename. Feel free to change it to something that makes sense to you.

File location

You may save your file to a local file (on your disk) or, if you are logged in, in the IRSA Workspace.

Modifications to the table

Depending on what you have been doing to the table at this point, you may have filtered or added columns. If you want to save the table as it is currently displayed, with all filters as imposed and any columns hidden, and any added columns as shown, select "Save table as displayed." If you want to save the original table, with all rows and the original columns intact, choose "Save table as originally retrieved."

Table Navigation

At any time, you can move among tables by just clicking on the tab name. But, with the ability to extract data from images comes the ability to rather quickly drown in tables. There is a way to navigate among a lot of table tabs that have accumulated.

This is probably best explained via an example. Here, I have tabs showing that I have extracted several lines, drills, and points from my images:

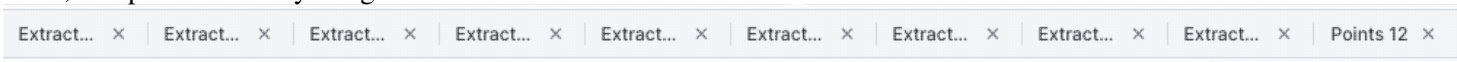

On the far right of the collection of tabs, at the top right of the table pane, I have a downward pointing arrow. When I click on this, I get a drop-down:

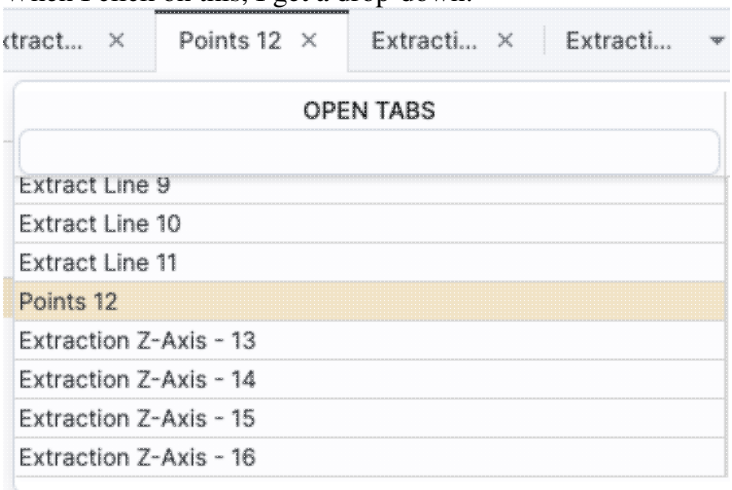

This drop-down shows that the tab marked "Points 12" is in the foreground (it is shaded yellow) and this drop-down has a filter box at the top. That filter box at the top works just like the filters discussed above, so if I want to find the tabs with the extracted lines, I can type "line" in the box, and it will filter down the list, leaving only those tabs with "line" in the name:

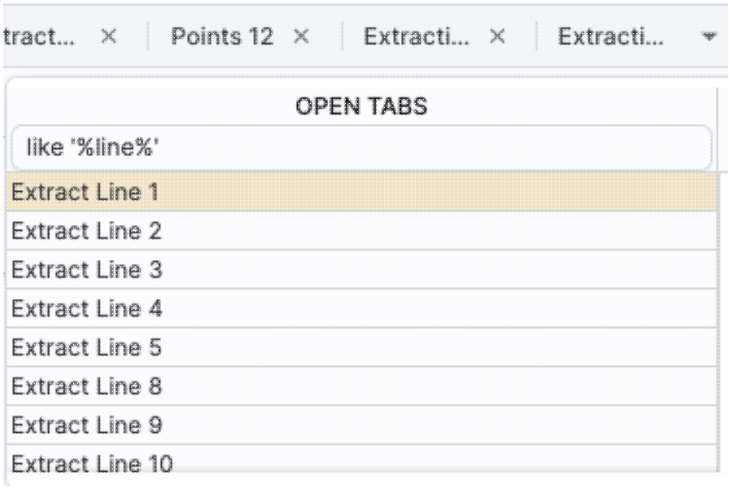

Then, from there, I can select the tab I want to bring to the foreground. (When I select a tab to bring to the foreground, that tab is also in the foreground for plotting.)

In this fashion, I can navigate easily among many tabs that are open at the same time, even if I have so many tabs that their headers are not completely shown.

IRSA Viewer: Catalogs

Catalogs are a special case of tables; the basic functionality of tables is covered in the Tables section. You can load a wide variety of catalogs to load for overlaying on your visualized data, but this section focuses on IRSA catalogs. If you don't have an image loaded, the tool will pick a "coverage image" for you and overlay the catalog on that image. It will also make a plot for you, which you can change.

Contents of page/chapter:

Introduction

There are several different ways to get catalogs into IRSA Viewer. This chapter focuses on IRSA catalogs.

When you click on the "Catalogs" tab at the top of the IRSA Viewer window, you are dropped into an IRSA Catalogs search, which is what we describe here. There are other kinds of searches you can also do that will load catalogs (or other tables) into the tool.

IRSA Catalogs -- Searching for catalogs at IRSA

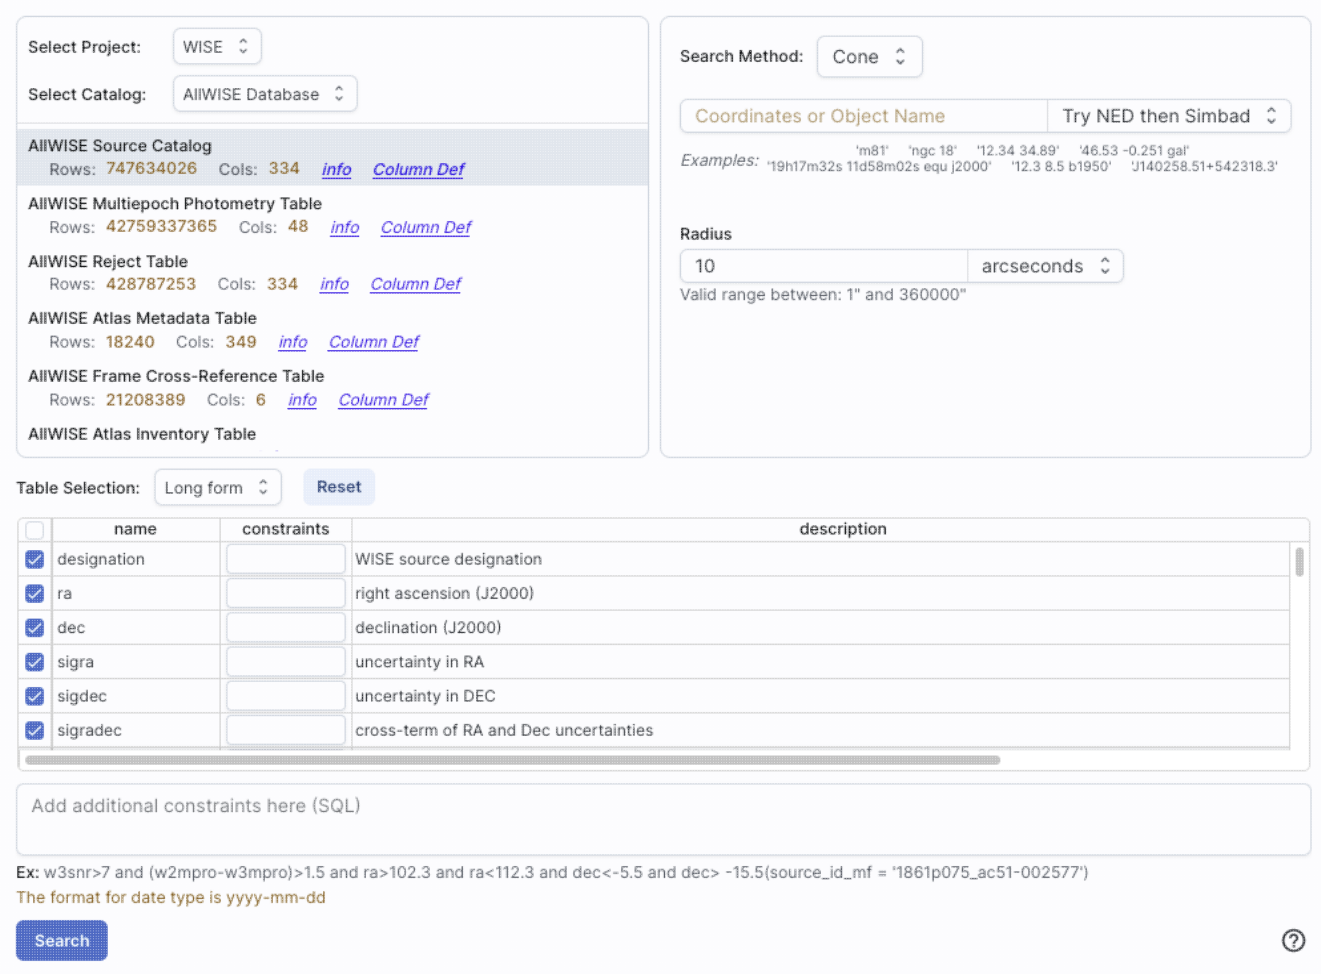

The upper left quadrant of this window is where you specify which catalog you want to search. To change catalogs, first select the "project" under which they are housed at IRSA, such as 2MASS, IRAS, WISE, MSX, etc. The available choices underneath that change according to the project you have selected. A short description is provided for each of the catalogs, with links for more information (including definitions of the sometimes cryptic column names); an example is here:

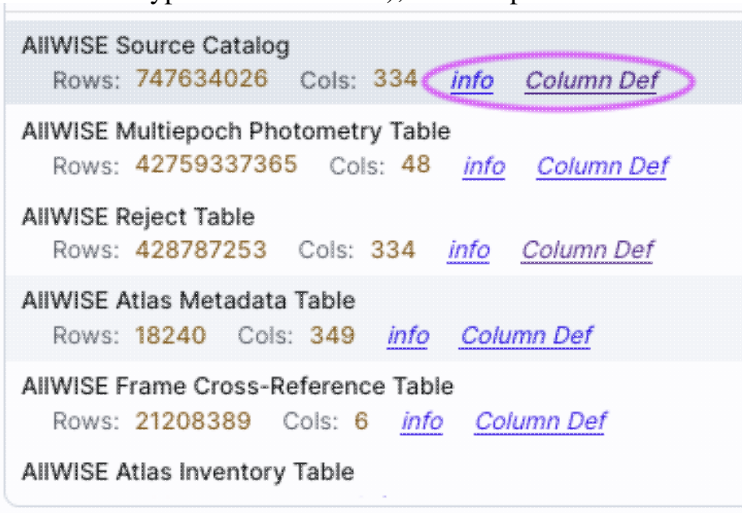

The upper right quadrant of this window is where you specify the target (the position is sometimes pre-filled

with its best guess as to what you want) and the search method (cone, elliptical, box, polygon, multi-object, all-sky), and the parameters that go with that search method (e.g., the radius of the cone). The parameters for each of these searches change dynamically as you select search options, as follows:

Caution:

Pick your units from the drop-down first, and then enter a number; if you enter a number and then select from the drop-down, it will convert your number from the old units to the new units. There are both upper and lower limits to your search radius; it will tell you if you request something too big or too small. Note that these limits are catalog-dependent.

Cone search:

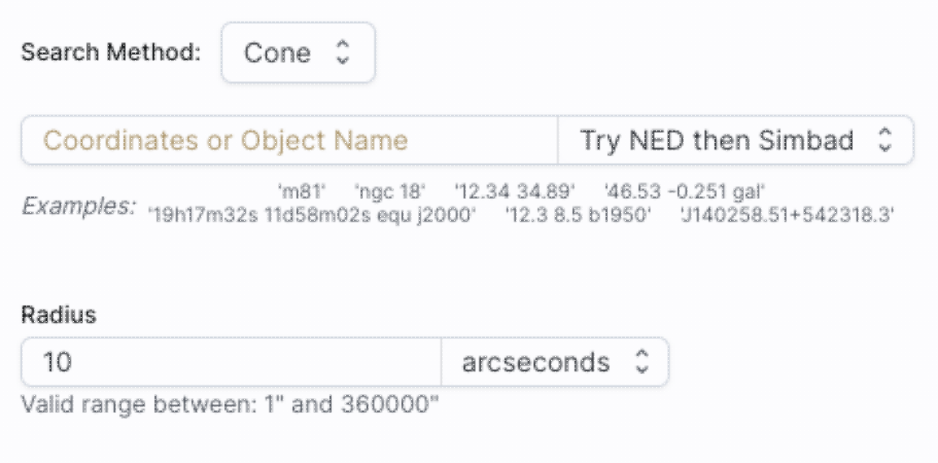

You can put in a position, but sometimes it attempts to guess a position, based on prior searches. You specify the cone radius; the default is 10 arcsec.

Elliptical search:

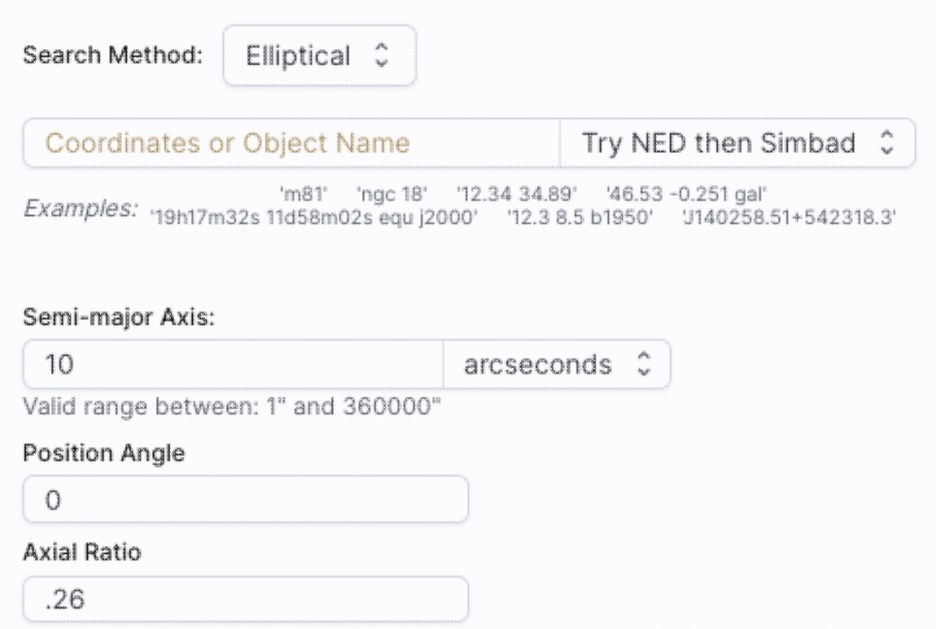

You can put in a position, but sometimes it attempts to guess a position, based on prior searches. You specify the search ellipse's semi-major axis, position ratio, and axial ratio. Defaults are as shown.

Box search:

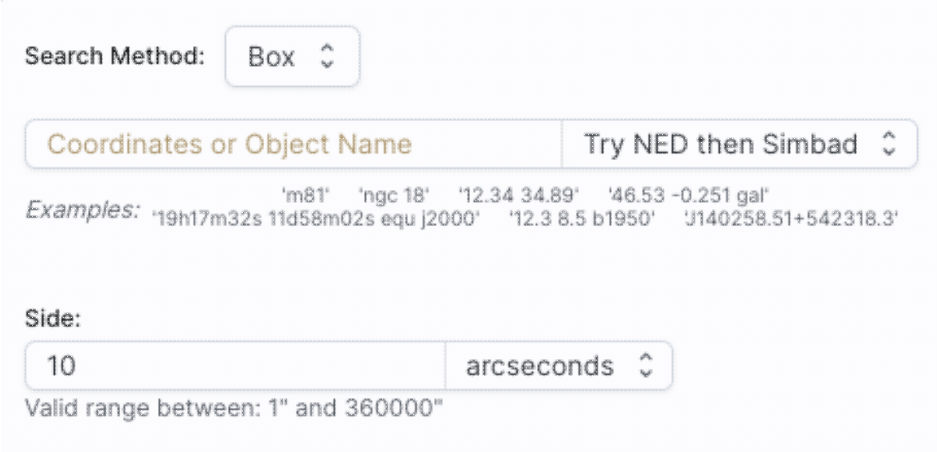

You can put in a position, but sometimes it attempts to guess a position, based on prior searches. You specify the box's length on a side; default is as shown.

Polygon search:

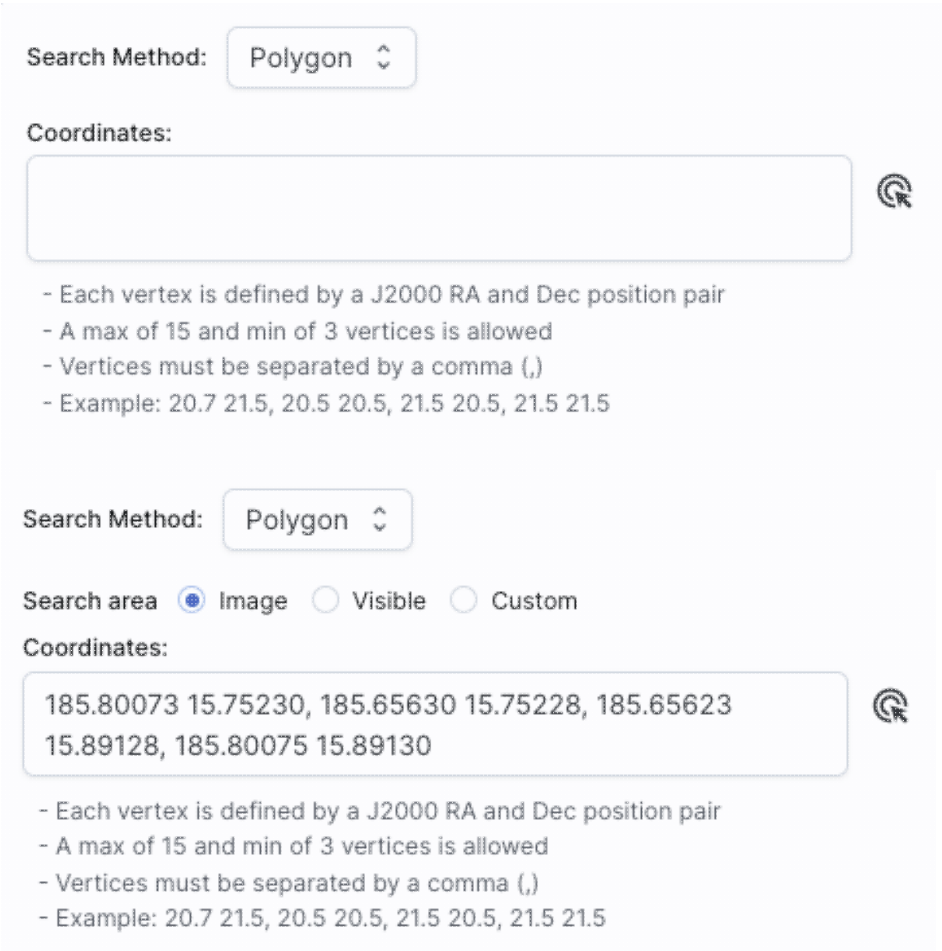

For this, note that it no longer has a single target location. It will sometimes try to pre-fill the vertices of the position it thinks you want, based on prior searches. If you have images loaded, it will give you choices based on the current image -- you can select whether you want the catalog request to match the entire area of the image you have selected ("image"), or just the portion of the image you can see in the current view ("visible"), or your own ("custom") area. (However, note that if you have selected a HiPS image before searching, you are limited to a maximum of 5 degrees.) The list of vertices in the coordinates box are in decimal RA and Dec in degrees. You must enter at least 3 and at most 15

vertices, separated by a comma. Note that, for overlaying catalogs on HiPS images, you cannot select "image", because HiPS images are generally very, very large, so this would result in too many points being returned. There is a maximum of 5 degrees imposed on catalog searches to match HiPS images.

If you select a rectangular region of your image and then select a polygon catalog search, you will have a fourth radio button above, "selection", which matches the corners of your selected image region.

If you select the "bullseye" icon on the right ( ), you get a pop-up with a way to interactively select your target; this works just like this interactive target refinement (go there for more details) :

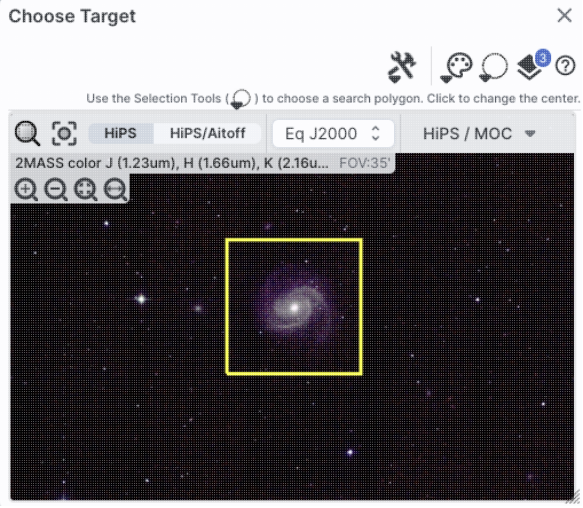

Multi-Object search:

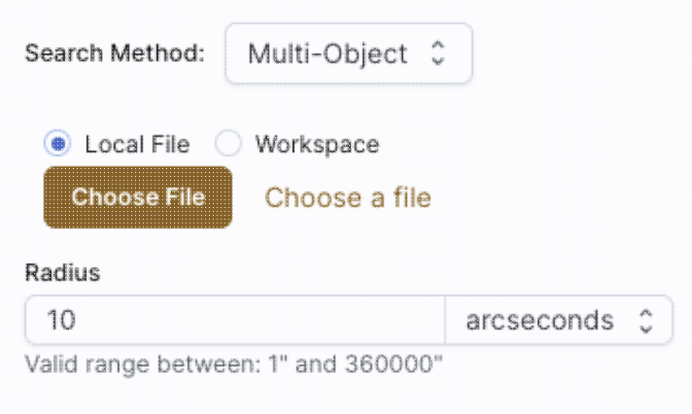

For a multi-object search, it can't guess what position you want. You need to upload a file (from your disk or the IRSA Workspace) in IPAC table format , which is a varietal of plain text. (IRSA has a table validator which may be helpful.) Note that you also have to specify the radius over which to search for each of the targets in your list.

When you do a multi-position search on catalogs, three new columns are added to the catalog as it is returned to you. These columns are :

◊ cntr_01 - the target position you requested
◊ dist_x - the distance between the target position you requested and the object it found
◊ pang_x - the position angle between the target position you requested and the object it found

These additional columns can help you assess if the target(s) it found is the target that should be matched to the position you requested.

All-sky search:

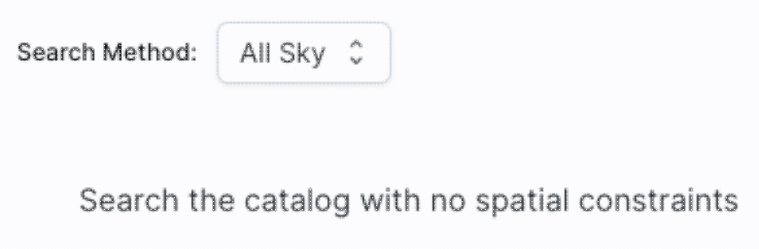

Because this is an all-sky search, it does not have a single target entry box. In order to constrain this search, you need to impose constraints on the bottom of the screen (see below).

The bottom of this window allows you to set restrictions on specific columns. It gives you a list of all the available column names in the corresponding catalog. (Most catalogs have identical "standard" and "long form" selections, but some have more columns available in "long form".) From here, you can choose what to display (tickboxes on the left), and filter what is returned ("constraints" column). For example, only return objects with values in column y that are greater than x. If you add more than one restriction, they are combined logically using an "AND" operators; be careful, because you can thus restrict data such that none of the catalog meets your criteria.

Click on "Search" to initiate the search. It will load the catalog into a tab of its own. The objects will also be overlaid on any images you have loaded, and a default x-y plot will be shown. (For more on the x-y plots, see Plots section.) All of these representations are interlinked -- clicking on a row in the table shows it on the image and in the plot, and clicking on an object in the image shows it in the table and in the plot, and clicking on an object in the plot shows it in the table and on the image.

To close the catalog search window without searching for a catalog, click on "Cancel".

Tips and Troubleshooting

- If the catalog search is successful quickly, it will promptly return the results in a tab of its own.
- The search may take a long time to return, especially if you have asked for a large catalog, and you may think that nothing has happened, but be patient and eventually it will return a tab.
- Use large search radii with caution! Be sure you understand how many sources you are likely to retrieve. Searches that retrieve more rows will take longer. Searches that retrieve tens of thousands of rows will take quite a while.
- If you want to impose additional constraints on the catalog during your initial search, you can do so in the lower half of the screen (e.g., SNR > n in some band, or an SQL command), you can place constraints at this point. However, be advised that it is easy to combine constraints such that no sources are retrieved!
- If you overlay a large catalog, the tool will show cells with a number enclosed indicating the number of sources in that region. As you zoom in closer and closer, the tool will adjust those bins to smaller and smaller cells until it shows you individual sources. Go here for more information!
- If you overlay a large catalog, then turn around and save a regions file from the catalog overlay, then the entire catalog may not be saved. To get a large catalog saved as a regions file, save it from the table.
- If you have "pan by table row" turned on (see Visualization chapter), then it may be disconcerting to have the images "jump" right after the catalog loads. It is centering the selected catalog object (the first one upon catalog loading) in the viewer. If you don't like this, turn off "pan by table row".
- By default, it may show you fewer columns than are available in the full catalog. By selecting "long form" (above the list of columns), you can access the full range of available columns. In some cases, there are literally hundreds of columns that you can access!
- If you start searching from a HiPS image, you are limited to a 5 degree search radius.

The search results are then shown in a Firefly table and you can interact with it.

Interacting with Catalogs

When you load a catalog, the tool may create a table, a plot, and/or, if your catalog has position information (e.g., RA and Dec), it overlays the catalog on an image. Tables, plots, and overlays on images are all interlinked and interactive.

Catalogs are a special case of tables; the basic functionality of tables is covered in the Tables section. You can sort and filter the table.

Plots are also covered in a different section. You can make scatter plots, heat maps, and histograms. You can plot columns from your catalog, including simple mathematical manipulations of catalog columns.

If the catalog has positions included, the catalog will also be overlaid on the loaded image(s). The Visualization section includes information about that. Each catalog that you load is overlaid on the image using different, customizable symbols and colors.

When you have catalogs loaded into the tool, the header of the catalogs has the name of the catalog and a color swatch:

This color swatch corresponds to the symbol color that is used in the image overlays. You can change the color by clicking on the color swatch in the header, or by navigating to the layers in the image pane. See the color picker section of the visualization chapter for more information.

Tips and Troubleshooting

- Large catalogs will be displayed hierarchically! See next section.
- If you save the overlays from an image as a regions file, you may not get your complete catalog, especially if it is a large catalog (see next section!). However, you can save the full contents of a single catalog as a regions file using the "save" (diskette) icon in the table toolbar, instead of the image toobar.
- The "color swatches" may not appear immediately. To make loading faster, sometimes the colors don't load until they are actually needed. If you are in a situation where no images are visible, then no color swatches may appear until you ask the tool to show you an image (like the coverage image), and then the color swatches will appear.

Hierarchical Catalog Display

If one has a large catalog loaded into the tool overlaid on top of lots of images the possibility exists that the computer or the network could be overwhelmed trying to render all the points on all the images. Historically we dealt with this by "thinning out" the catalog and not showing all the points. However, there is a better solution, which is now employed here!

For catalogs below about 1000 points, the tool will show the individual points on the image.

For catalogs above that threshold, the tool will bin up the catalogs based on HEALPix pixels (see HiPS section here for more links). In summary, the sky is broken up into sections, and the tool will show symbols with a number indicating the number of sources in that region. Then, when you zoom in, it will dynamically adapt to

show you smaller and smaller cells until it shows you all the individual sources.

From the layers icon (see visualization chapter), you can bring up many display options. Below are examples of what is displayed, the options seen in the layers, and additional options. The same catalog and zoom level and minimum group size are used for each view. The "Min Group" option here is 50, so if there are cells with fewer than 50 sources, then the individual sources are shown, and if there are more than 50 sources, then the cell is shown with a number inside corresponding to the number of sources from the catalog. (See below for additional information.)

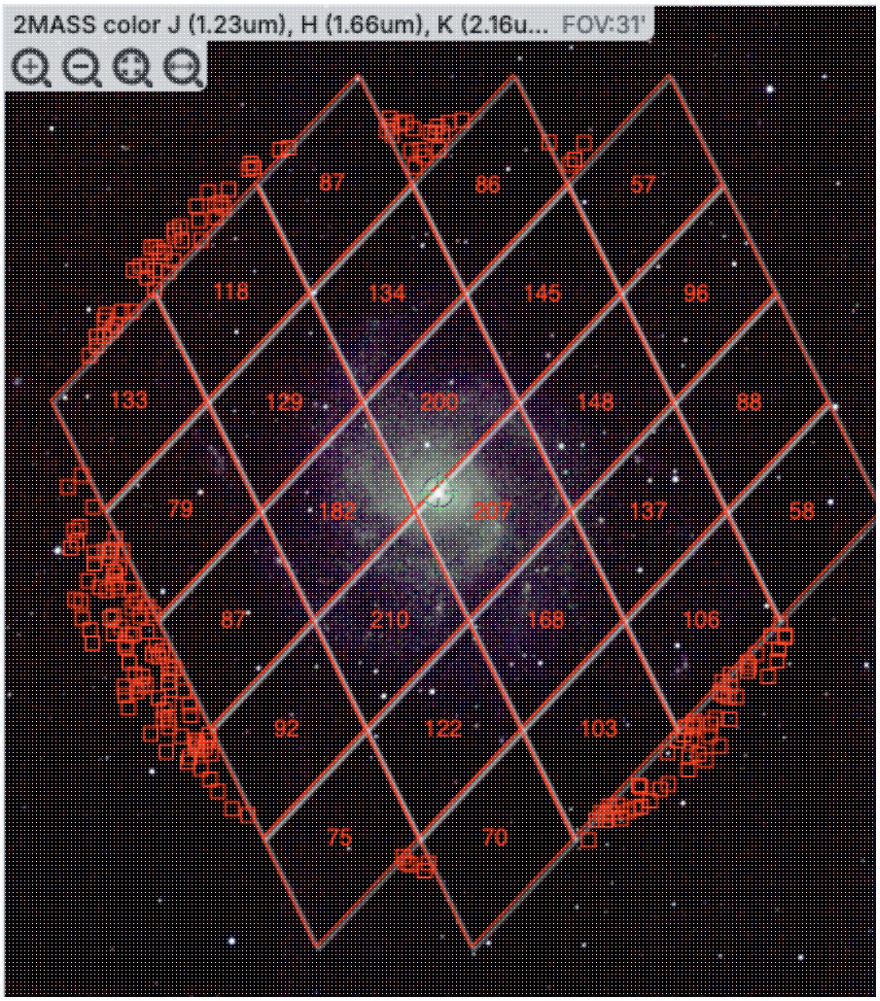

In this view, the 'cells' used are the cells explicit the size of the cells is very clear. In the top row across the top have fewer than 50 sources (so the individual sources are shown), then the next row 57 sources respectively.

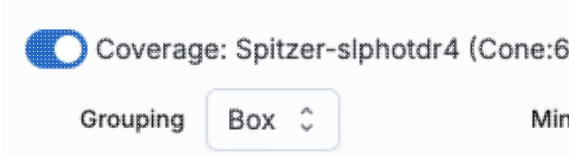

In this view, the 'cells' are shown by circles encl sizes are the same as in the prior screenshot, but obvious to new users.

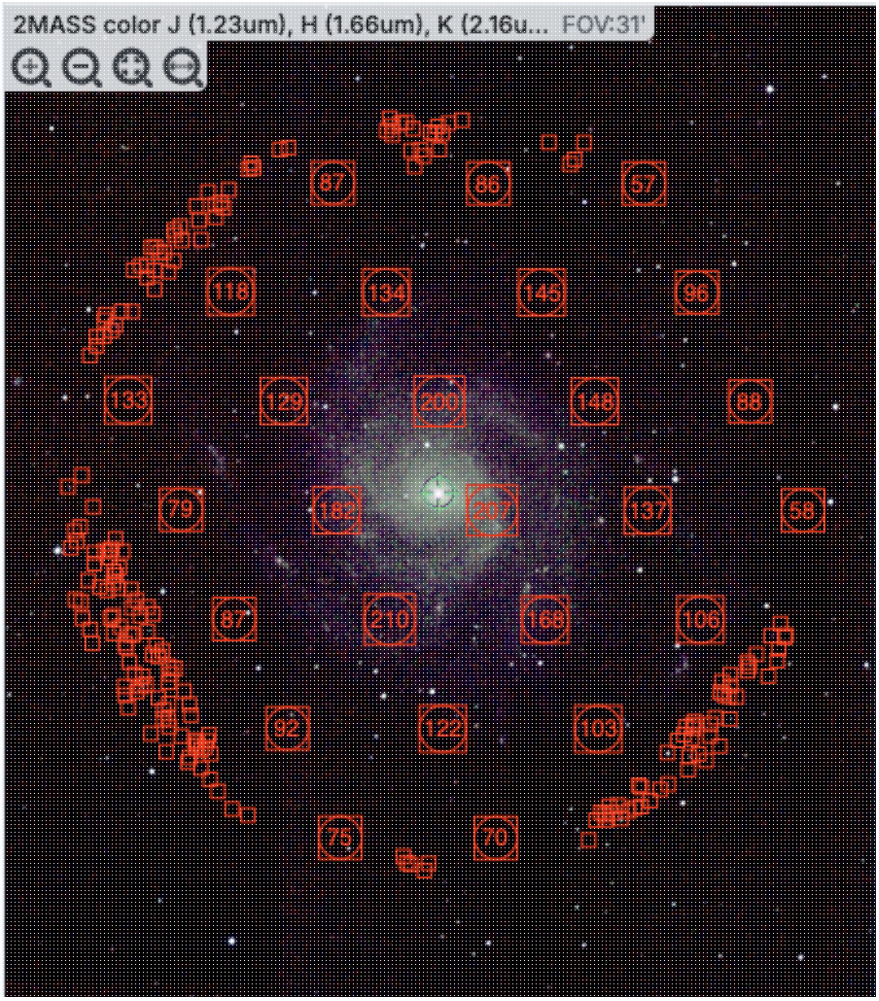

In this view, the 'cells' are shown by ellipses sho
cell sizes are the same as in the prior screenshot,
less obvious to new users. It may be more obvio
groups of points.

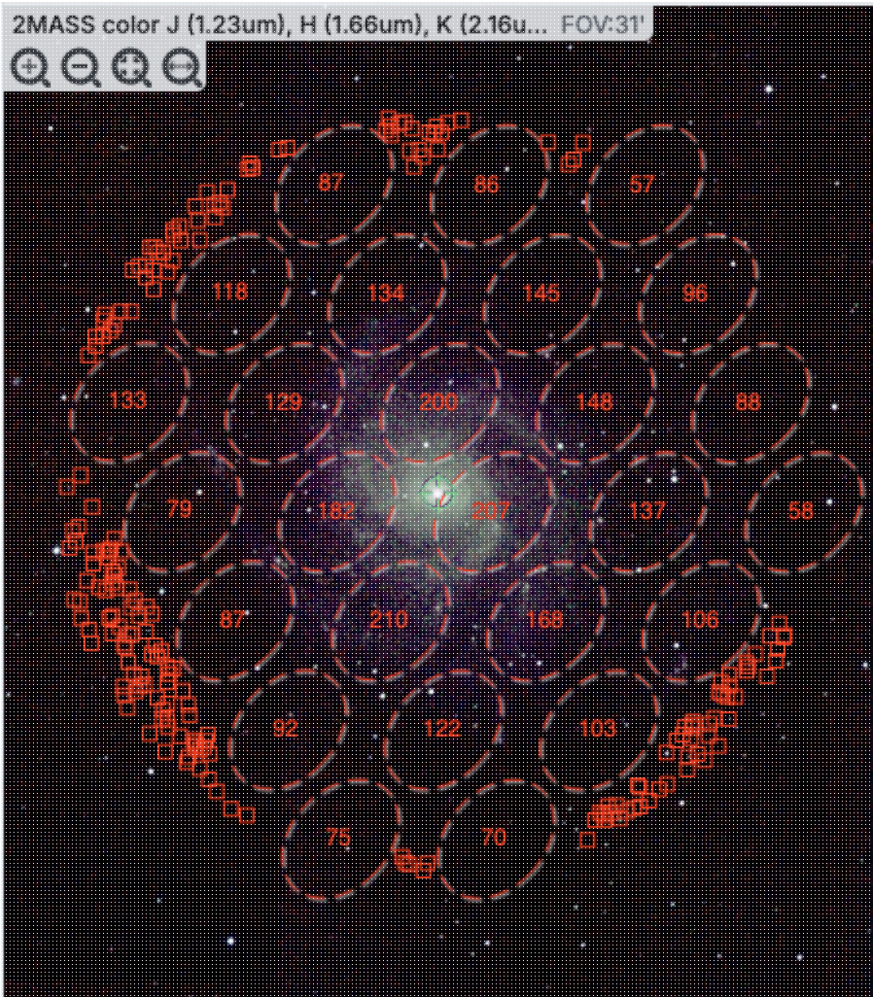

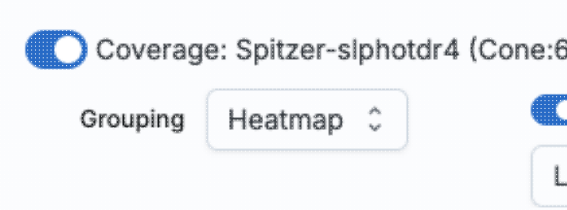

Finally, in this view, the 'cells' are again shown a
color of the cells corresponds to the number of s
"Linear", "Linear Compressed", or "Log Stretch
the color range by changing the color using the
which you can also change the transparency. Thi
how many sources are in each cell, but makes it
though you can change the transparency of this
still can make seeing the image challenging in so

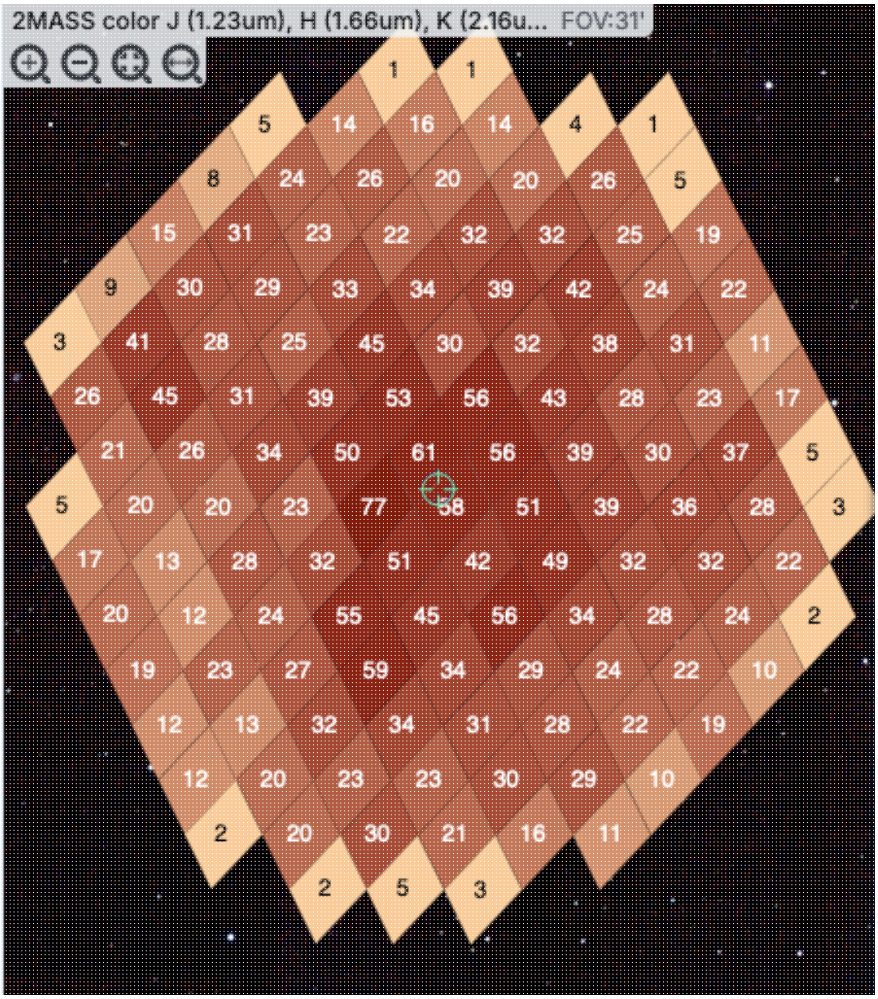

Tips and Troubleshooting

- For all of these renditions, when you zoom in close enough, it will dynamically adapt and show you individual sources when you zoom in. (That is, it no longer decimates the overlaid catalog, which is what it used to do.)
- For all of these renditions, if you click on a cell, it will display all of the sources in the cell. You can click on many cells in a row and it will continue to display all the sources it can until it reaches the point at which it thinks performance will suffer, at which point it will turn some of the points back into cells.
- If you want to have more of your catalog shown as individual sources, pick a smaller "min group" number.
- If you have more than one catalog loaded, the numbers within the cells (and in some cases the cell indicators themselves) will be offset slightly so that you can see them.

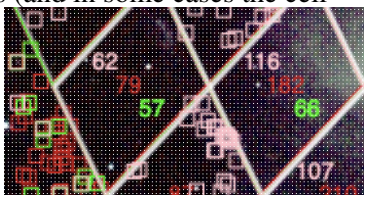

- If you have a catalog that includes sources from all over the sky, it very well may just give you box groupings, and may not allow you to change that view until you zoom in.
- If you have cells where only 1/4 of a cell is populated, it automatically renders a smaller cell, so if you have a sparsely populated but still large catalog, the size of the display will always be "small" size cells.

- If you are looking at many footprints from, say, a complex, and long ObsCore search, if you have more than 30,000 footprints, it may not be able to render all of the outlines of all of those images. It may render the centers of all of those images as if it were a catalog, in which case you will encounter these kinds of hierarchical catalog display options.

Details Tab

If you load a catalog from IRSA, you will likely have an additional tab on the right hand side, under the plot, called "Details." This additional tab is sometimes called a "property sheet." This tab is, itself, another Firefly table, and consists of each of the columns of the retrieved catalog with additional information about each field where available. (Not every catalog may have this information available.) This information can be used to learn more about each of the columns in retrieved. For additional information, please consult the full documentation that accompanies the catalog.

Tips and Troubleshooting

- The property sheet is a more expanded, vertical view of the information shown in a row of a catalog, along with documentation of the catalog columns. Because you can sort/filter the data in the property sheet, you can restrict what values are shown. Those filters are respected as you page through a catalog. So, for example (see screenshot below), you can pull up the property sheet, filter it down to only show the profile-fitted magnitudes and errors by filtering on "mpro", and then step through the values in the catalog and inspecting the brightnesses as shown in the property sheet for each source.
- When changing rows in the main table, the property sheet/details tab scrolls to preserve the visibility of whatever row in the details tab is highlighted. If you scroll down in the property sheet *without changing the highlight*, when you change rows in the main table, because the first row in any table is always highlighted by default, the property sheet will scroll back to the top.

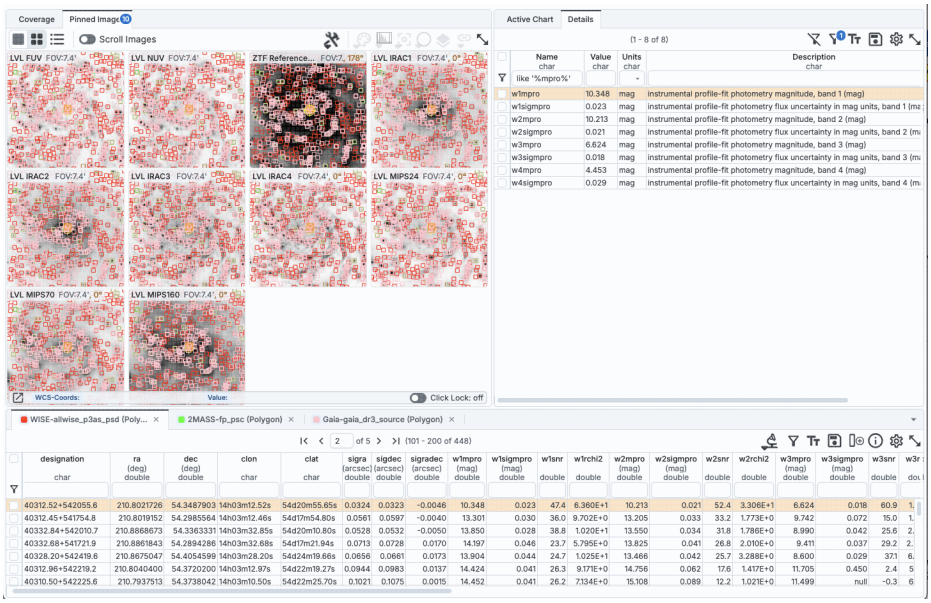

IRSA Viewer: Other Searches

There are several other searches you can add to IRSA Viewer using the side menu to add to the tabs at the top. This chapter covers most of these choices. Nearly all of them are VO searches that retrieve some sort of tables, many of which are catalogs. For the results of any of these searches, if the tool recognizes positions in a catalog, it will overlay the catalog on images and make plots.

Contents of page/chapter:

Introduction & Terminology

There is a lot of terminology in this chapter to understand.

- VO = Virtual Observatory
- TAP = table access protocol. TAP services enable complex queries of tables using ADQL.
- ADQL = astronomical data. query language
- ObsCore = core components of Observation Data Model
- Data model = A standard logical structure for a type of dataset; more flexible and general than a data format.
- ObsTAP = TAP service that serves ObsCore tables
- SCS = simple cone search.

There are myriad places on the web to learn more about TAP queries and ADQL, as well as all the rest of the VO standards and protocols. We just provide a brief overview here in the context of this tool.

IRSA Viewer can help you interactively create ADQL which then you can copy and use in your own code elsewhere.

By using TAP and ObsTAP queries, you can use IRSA services to talk to other archives that also comply with these standards, world-wide.

The first part of this page talks about interactive target refinement and some constraints that are common to more than one of the searches discussed here, and then specifics of particular searches follow after that.

Data Collections is another tab and possible search, but it has its own help page.

Interactive Target Refinement

Whenever you see this icon  in IRSA Viewer, you can click on it to bring up a window to **interactively refine your target selection via clicking on a HiPS map**. Here, we are using a TAP search to demonstrate this process, but you can find this kind of target refinement in several places in IRSA Viewer.

When you click on the icon  you bring up a window:

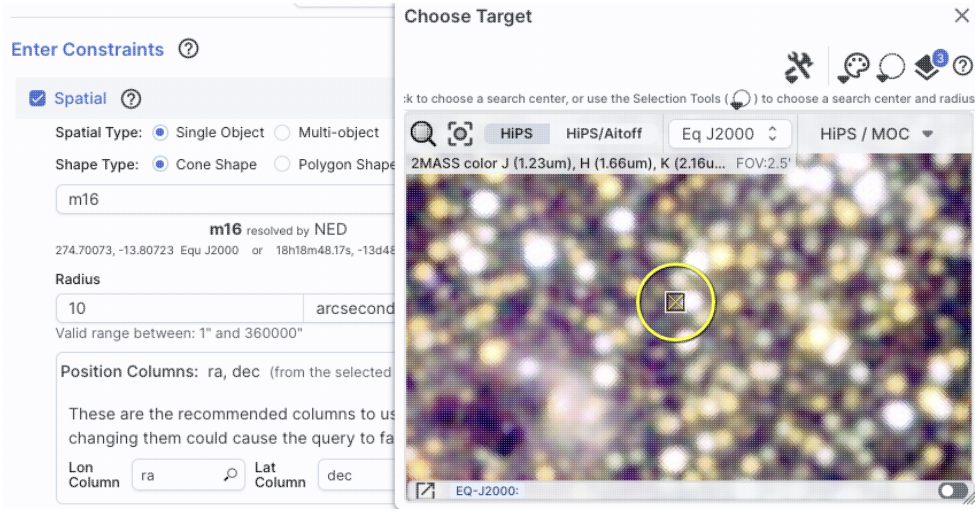

If you have entered a target already, the window arrives already centered on the target. If not, it is centered on the galactic center, zoomed out. If you have entered a cone search radius already, then the circle drawn on the image is that cone size. You can manipulate this image with the same basic tools as in the visualization tools.

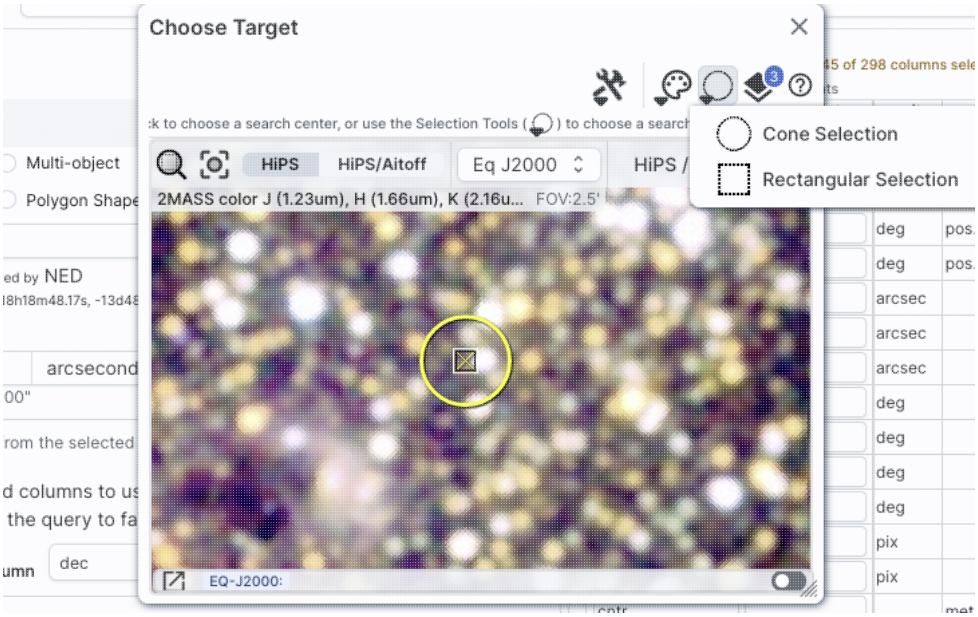

To change the search region interactively, choose the selection tools and draw a shape on the image.

Note that if you have selected a cone search on the left, no matter what you select on the right, it will give you a cone search. If you change the cone position or radius in the yellow boxes after you change the selection, it will update the region in the image.

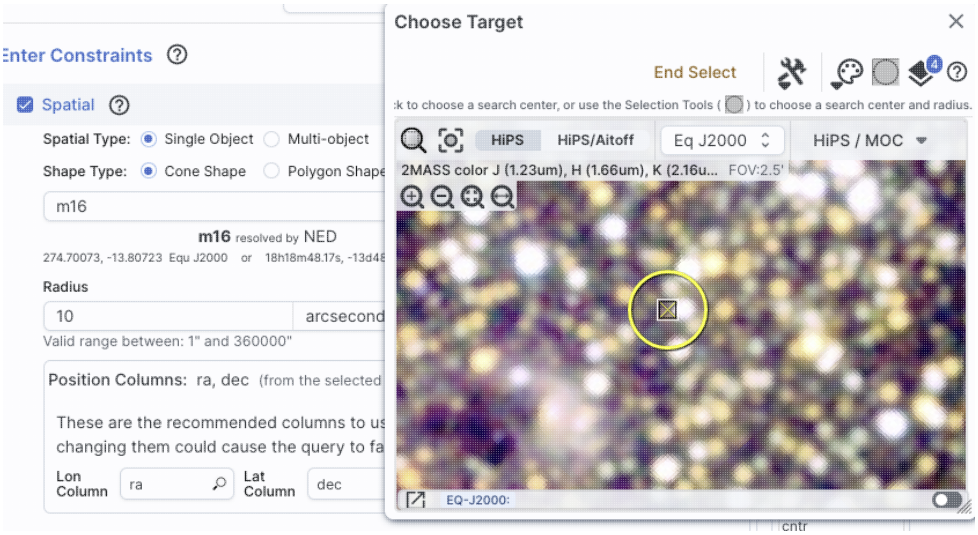

If you want to quit out of the selection without changing, click on "end selection" (the brown text near the top of the image).

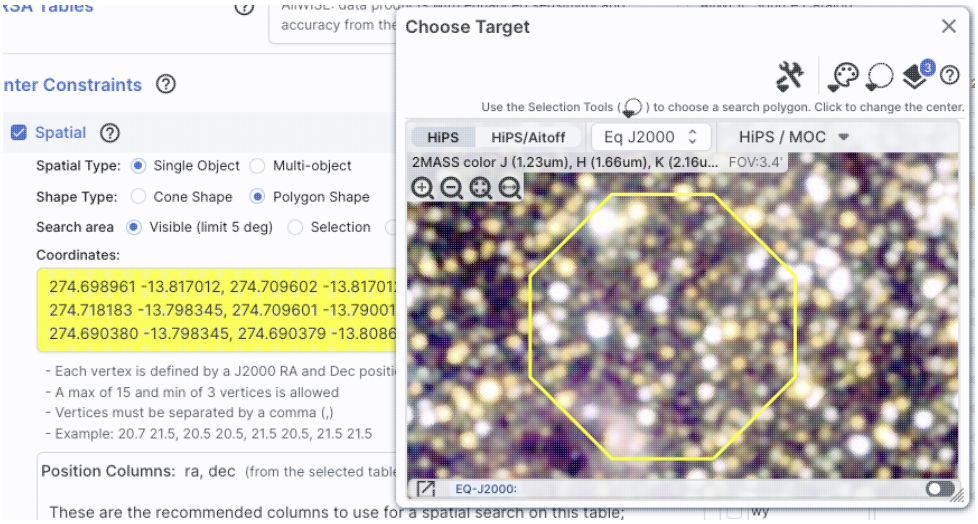

If you select polygon on the left, and you use the selection tool for "cone selection" on the right, you will get a spherical polygon (a polygon where the line segments are on a sphere).

When you are done with this pop-up window, click on the 'x' in the upper right of the window. Then you can continue with whatever you were doing before you started to refine your target parameters.

VO TAP Searches: More information about constraints

You can have several different ways of constraining your search depending on the options you have selected before the "Enter Constraints" section, and the options depend on what kind of service is available at the TAP service you have selected. If the options do not appear initially, click on the downward arrow to "unfold" the options.

Enter Constraints: Spatial

This part of the interface allows you to specify the details of a spatial search. You need to specify both what kind of search you want to do and which columns of the catalog are to be used for coordinates.

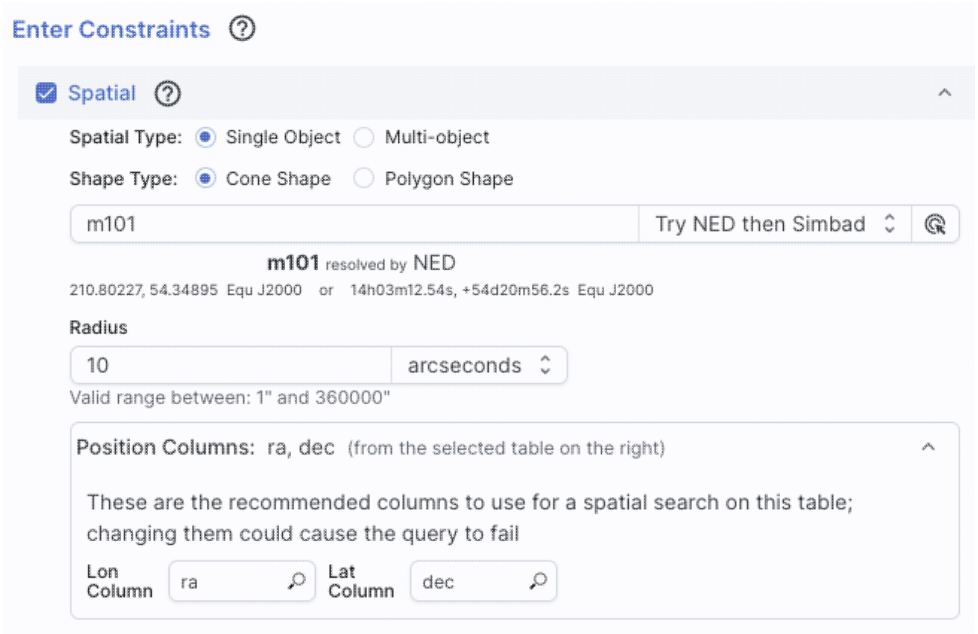

This is what it looks like when you do a **single target cone search**; note that you have the same name resolution options as in any other search here.

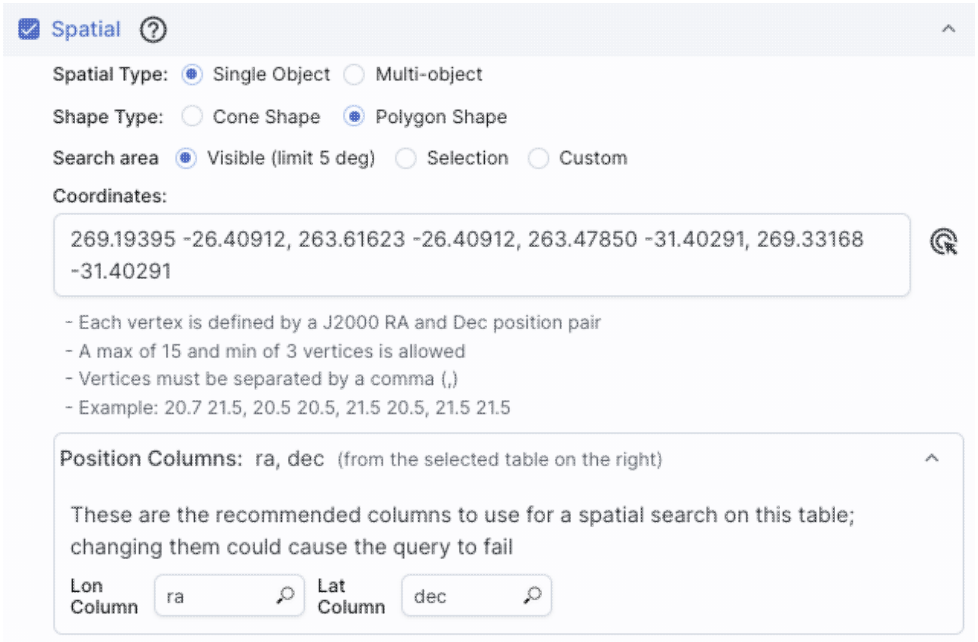

And, this is what it looks like when you do a **single target polygon search**. The search areas here (visible, selection, and custom) are the same as when you do a polygon search on catalog -- that is, you can select whether you want the catalog request to match the entire area of the image you have selected ("image"), or just the portion of the image you can see in the current view ("visible"), or your own ("custom") area. The list of vertices in the coordinates box are in decimal RA and Dec in degrees. You must enter at least 3 and at most 15

vertices, separated by a comma.

You can also click on this icon to interactively refine your search position.

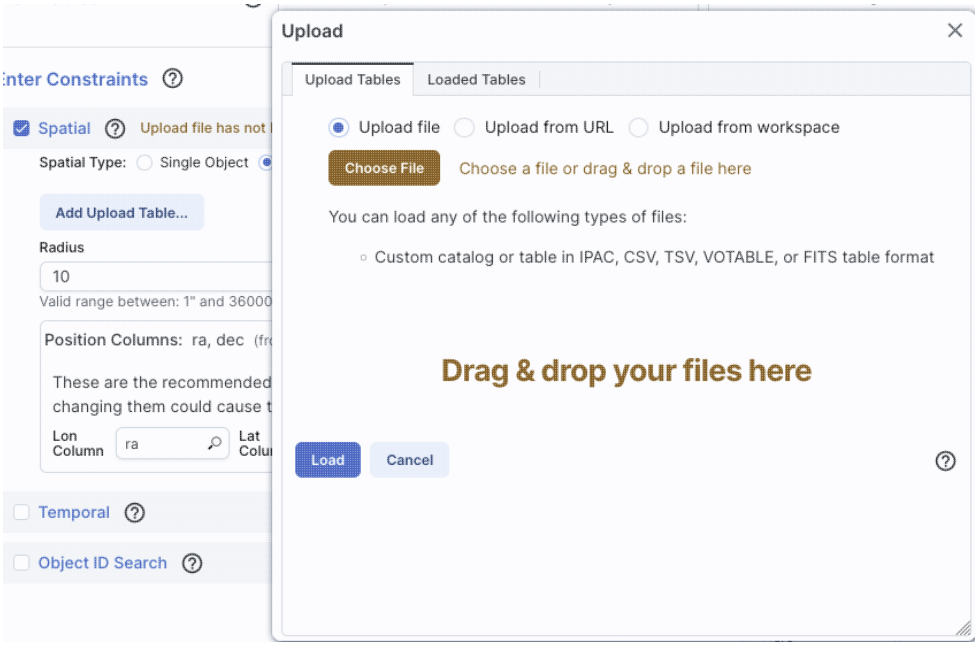

If you want to perform a **multi-target search**, click on "multi-object", which automatically brings up this pop-up, from which you can load a table from disk ("Upload tables" tab) or select one of the tables you have already loaded into the tool (click on the "Loaded tables" tab). Your uploaded catalog has to follow all the same rules as normal catalogs from disk.

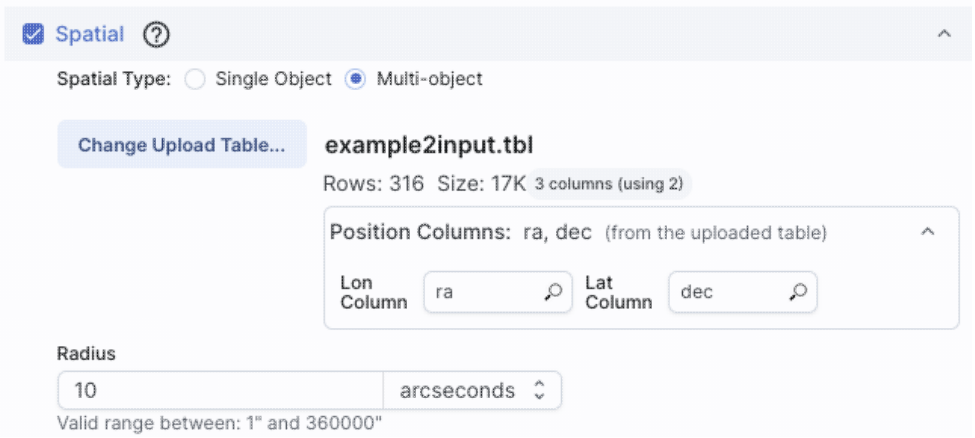

After you find your file with your listed positions and upload it, the tool attempts to guess which two columns are the position columns. In this example, it has (correctly) guessed that the position columns are "ra" and "dec". If it guesses wrong, or can't figure it out, you can help it along by clicking

on the down-arrow to 'expand' that part of the panel and selecting the two coordinate columns to use.

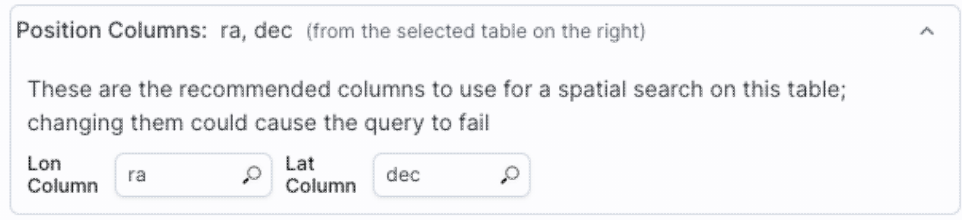

Regardless of what configuration you use, the last thing to check is which columns the tool has assumed are the position columns in the catalog to be matched to your position, region, or list of positions. Again, it attempts to make an educated guess as to the right columns, but if it guesses wrong, you can help it along by clicking on the down-arrow to 'expand' that part of the panel and selecting the two coordinate columns to use.

Enter Constraints: Temporal

This part of the interface allows you to specify which column of the catalog is to be used for timing, and allows you to constrain the date in two different systems.

This is what the panel looks like initially, where you specify the column in the catalog you are searching with the time and then the dates. If you don't remember what the column is in the catalog, click on the magnifying glass to get a pop-up with a list of all of the columns.

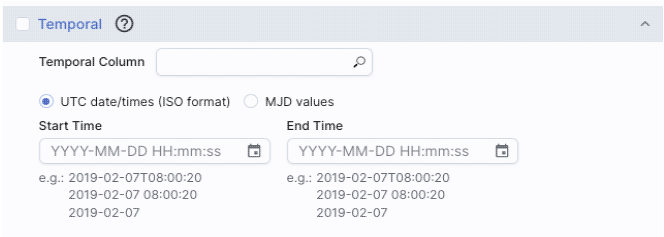

For the dates and times, if you click on the calendar icon at the far right of the entry box, you get a pop-up from which you can specify the date and time, shown here.

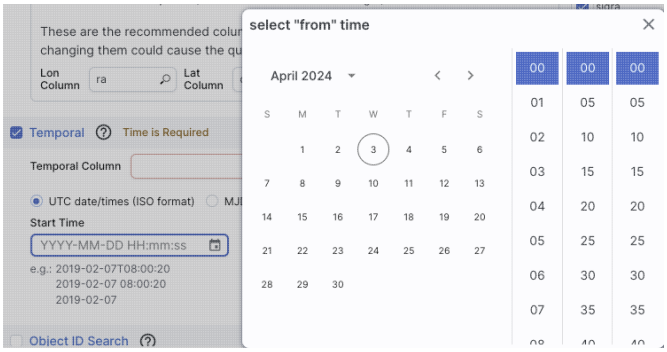

If you would like to work in MJD instead of ISO dates, select the "MJD" radio button. Note that it echoes below the box what it thinks you've entered in two different systems (UTC and MJD) to verify what you have entered.

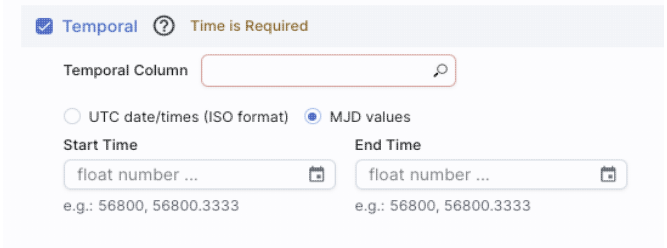

Enter Constraints: Object ID

This part of the interface allows you to match object IDs.

This is what the panel looks like initially:

This is what the panel looks like after you have selected your uploaded list of IDs (in this case, a file called "gaiaids.tbl", which consists of an IPAC table file that is just the list of Gaia IDs, in a column called "gaiaid"), and it is being matched against the Gaia DR3 main catalog, where the relevant catalog is "source_id".

Tips and Troubleshooting

◊ The names have to match exactly, so in the case of this example, since the Gaia column source_id is a long, the IPAC table file must also cast the Gaia ID as a long.

◊ If the catalog to which you are matching is not indexed by the name column you are using, the search may take a long time.

◊ If you are doing, say, an object ID search, you need to turn off the position search, otherwise it doesn't understand what you want it to do.

VO ObsCore Searches: More information about constraints

These are several additional ways of constraining your search depending on the options you have selected before the "Enter Constraints" section. These options appear if you have selected an ObsCore search. If all of these options do not appear initially, click on the downward arrow to "unfold" the options.

Enter Constraints: Observation Type and Source

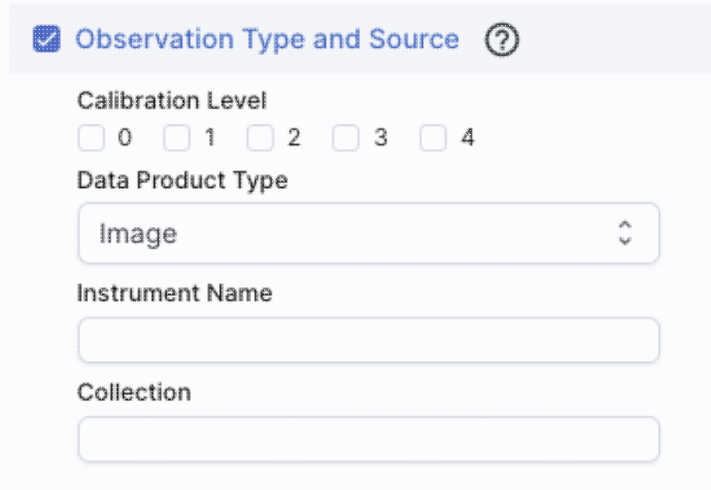

This panel provides a way to constrain the:

◊ *Calibration level* - 0 is the least processed and 4 is the most processed, and not all instruments provide all levels
◊ *Data product type* - image, cube, spectrum, SED, time series, visibility, event, measurement, or none specified. You can select more than one of these at a time by using a right mouse click.
◊ *Instrument name* - must match exactly, wild cards not accepted
◊ *Collection* - must match exactly, wild cards not accepted

Enter Constraints: Location

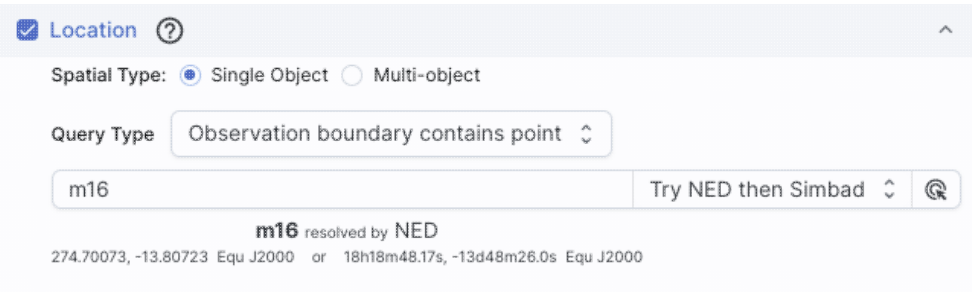

This panel provides a way to constrain the location of your search. Here, it is a single object search, which works just like it does above, including the interactive target refinement via the bullsye icon. You can also upload a list of targets by selecting "multi-object" -- it brings up the same pop-up as above, from which you can load a table from disk ("Upload tables" tab) or select one of the

tables you have already loaded into the tool (click on the "Loaded tables" tab). Your uploaded catalog has to follow all the same rules as normal catalogs from disk.

You can specify via the drop-down the type of your query: "observation boundary contains point", "observation boundary contains shape", "observation boundary is contained by shape", "observation boundary intersects shape", and "central point (s_ra, s_dec) is contained by shape." The latter refers to the columns "s_ra" and "s_dec" in the ObsTAP table.

Enter Constraints: Timing

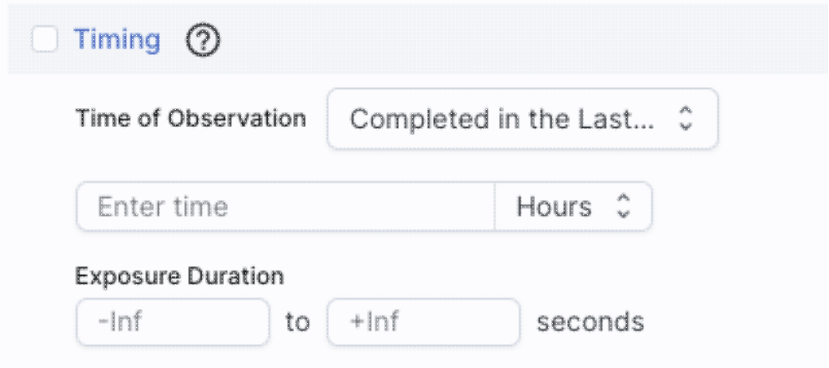

This panel provides a way to constrain the observation time of your search. This is the default option, where you want data completed in the last x hours (or other unit of time).

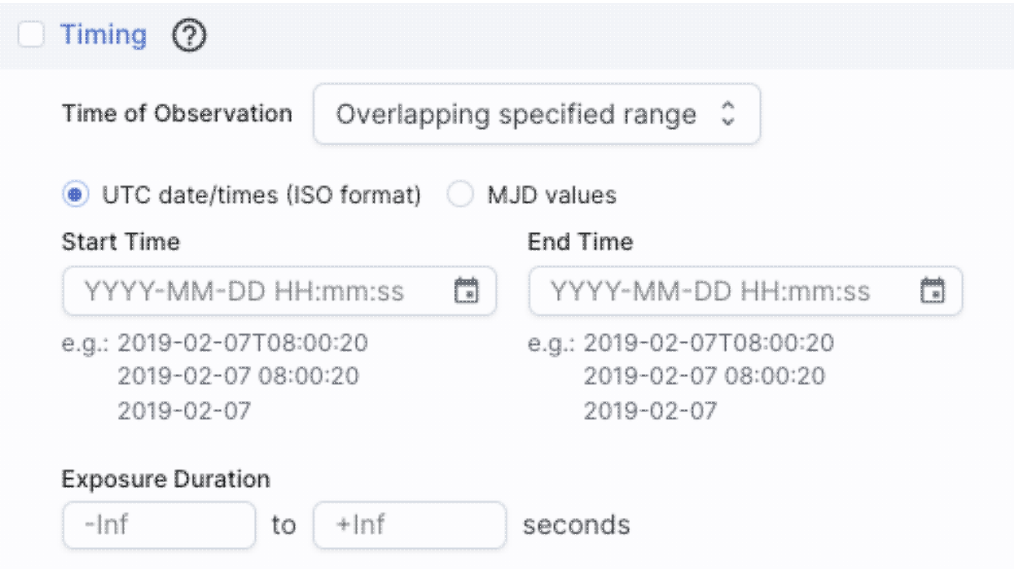

This is the alternate option, where you want data overlapping a specified date range, where you can specify UTC or MJD times.

Enter Constraints: Spectral Coverage

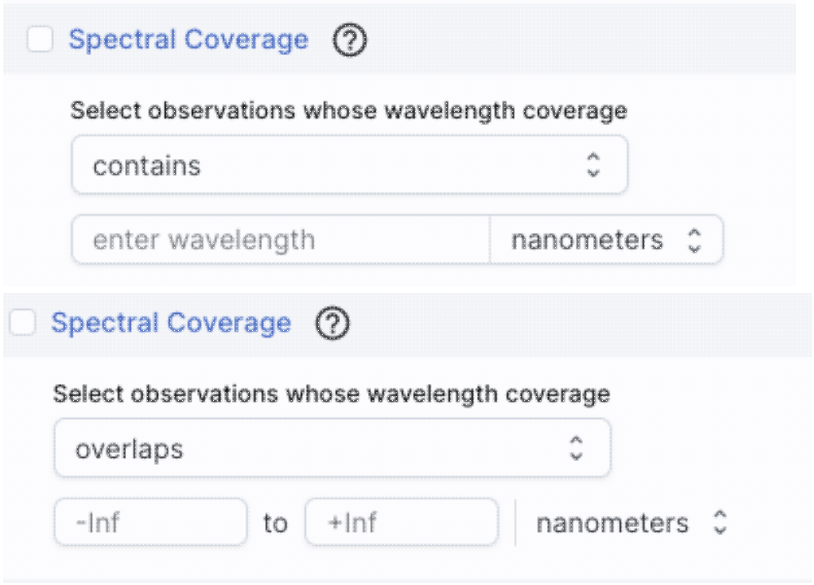

This panel provides a way to constrain the spectral coverage of your search. This is the default option, where you want data containing a given wavelength.

This is the alternate option, where you want data overlapping a specified wavelength range.

IRSA TAP Search

To see this tab as a choice on the top, you may need to select it from the side menu. (You may also find yourself having landed here from a general TAP VO search or a CADC ObsCore search; follow those links for more information.)

This is what the IRSA VO TAP search screen looks like by default:

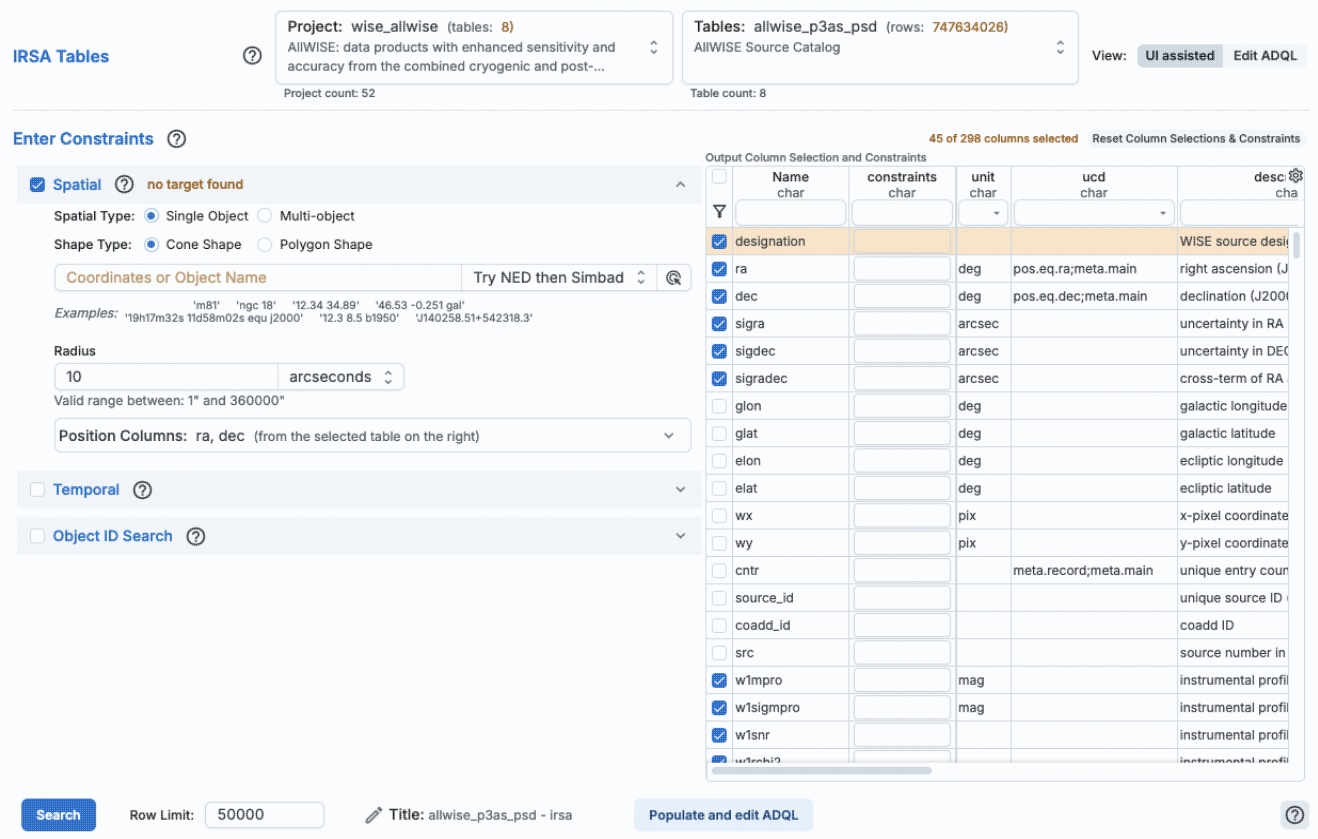

Just do it: a quick start

Select Table: It comes up ready to search on *IRSA Tables*. You first need to select the "project" (sometimes called "Table Collection" or "Schema" in other contexts). Then, having selected that, the drop-down menu on the right changes to reflect the tables available under that schema.

Then **Enter Constraints**: On the **left**, you can impose a variety of constraints. In addition to selecting the tickbox indicating that you wish to impose a particular kind of constraints, you need to specify which columns should be used for those constraints. More information on these constraints is above. On the **right** is a list of the columns in the selected table, with tickboxes to indicate which columns will be returned. You can also set constraints on the columns from here, following the same filter rules as for any tables here. Above this section of the screen, there is an indication of which columns are selected (e.g., 45 of 298 columns). You can reset the column selection via the button here as well.

Then to **actually do it**, click "Search."

Getting more out of it: Taking advantage of additional options

Selecting a Query Type

On the far right of the top row, there is a slider or button:

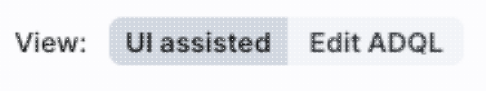

By default, it is set to "UI-assisted", as opposed to "Edit ADQL". Especially when starting out, UI-assisted is easier. By using the UI assisted" option, you can select pre-defined options and have the interface construct the query in ADQL. Alternatively, if you are already fluent in ADQL, you can select the second option, "Edit ADQL", to construct even more complex queries.

After populating the search parameters using the UI, you can click the button on the bottom, "Populate and edit ADQL" -- this takes the parameters you have entered, creates the ADQL, and launches the "Edit ADQL (advanced)" interface.

Advanced ADQL

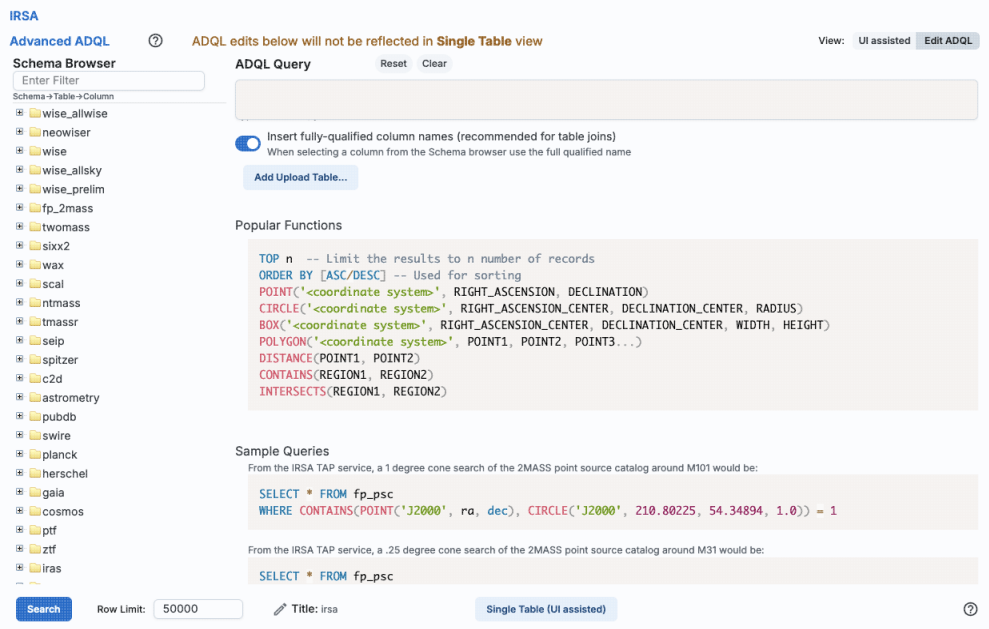

You can get to this screen by selecting "Edit ADQL (advanced)" in step two, or by clicking on "Populate and edit ADQL" after filling out the UI.

You can select the schema from the left side of the screen. Each of the schemas can expand into viable tables and then columns within each table via clicking on the "+" to the left of the folder icon. Click on a column name to have it appear at the location of your cursor in the ADQL query box on the right. If you have the tickbox checked on the right that says "Insert fully-qualified column names", clicking on the column name inserts fully-qualified column names at your cursor location in the box.

You can type the ADQL directly into the box. If you configured a search on the "UI assisted" page, this box is already pre-filled with the ADQL version of your search, and you can proceed to edit it further.

Examples of useful functions and queries are given on the lower right of this window; you may need to scroll down.

Tips and Troubleshooting

- All the drop-downs near the top are roughly searchable. This is very useful if you are trying to find a particular table or service that you can't quite recall. Click in the drop-down area, type the first letter of what you are looking for, and it will jump to the first instance of a string starting with that letter. Hit that letter again, and it goes to the second instance of a string starting with that letter.
- Not every table available via this interface even has RA/Dec or, if it does, it may not be searchable via positions. If you have selected a table that doesn't have positions, it will not yield results if you try to search by position.
- There is a maximum number of returned rows at the bottom. If you are anticipating more than this number of rows, increase this number!
- You can set the title of the returned table (e.g., what it will use in the tab that it shows in the tool) by clicking on the 'pencil' icon at the bottom in the center next to "Title" to change the title.
- If you arrive at the advanced ADQL page from the "Populate and edit ADQL" button, this is a one-way trip -- any changes you make to the ADQL here are NOT transmitted automatically back to the "UI Assisted" query page.

If you choose to use the GUI, you can work within it to set the constraints you desire at the bottom of the screen; see VO TAP: More About Constraints for all of the information about setting constraints.

The result of an IRSA TAP search is a catalog that can be interacted with like any other catalog in this tool.

General VO TAP Search

To see this tab as a choice on the top, you may need to select it from the side menu. Or, you may land in this tab from a table action.

When you first go to this tab, you will see this near the top of your screen:

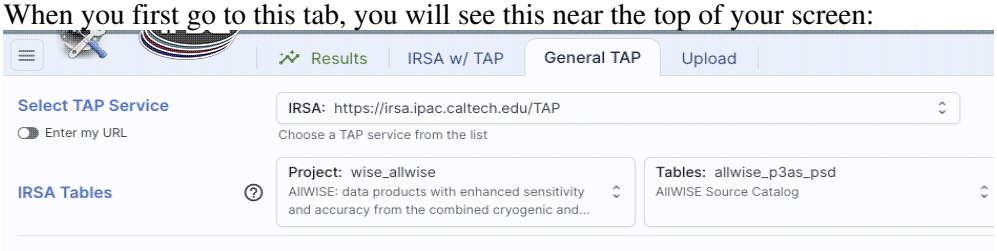

At the top, you now have a choice of which TAP service you want to use, and it defaults to IRSA's. You can select your favorite from the list, or use the toggle on the left to enter your own custom URL. It will remember

your custom URL in the same session, so you can select it again later. If you want to hide this top row after setting it (to, say, regain screen real estate), look for this on the far right:

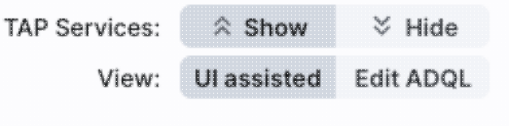

the "TAP Services:" button (show/hide) will reveal or conceal this top row.

The rest of this search screen is basically identical to that which you get from the IRSA VO TAP Search screen, even if you pick a TAP service other than IRSA's (with a few exceptions, including if it's ObsCore; see below).

For VizieR's services in particular, because there are so many tables, the tool will give you a slightly different interface under the "Tables" section of the window. Here is the default Vizier choices as of this writing:

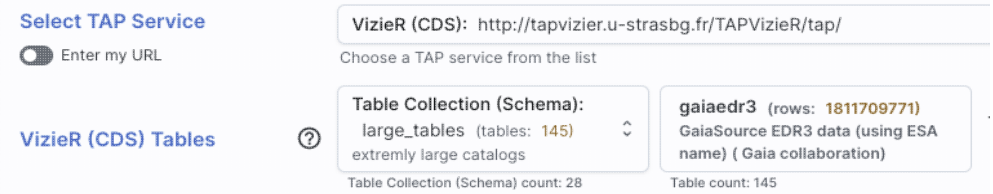

Note that it tells you how many tables and rows are available. Switching to tables associated with journal articles, far more tables are available:

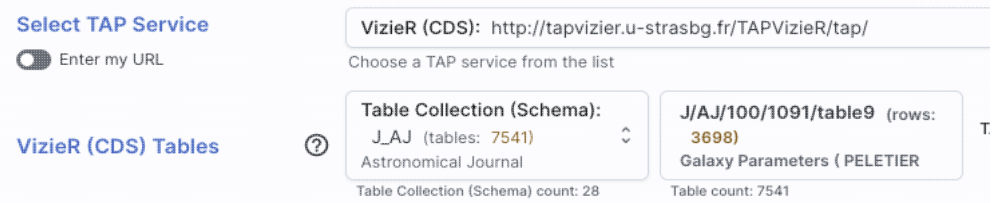

Now, if you click on the second tier menu (J/AJ/100/1091/table9), you get a pop-up, which is another Firefly table like any other in this tool, so it's searchable and sortable:

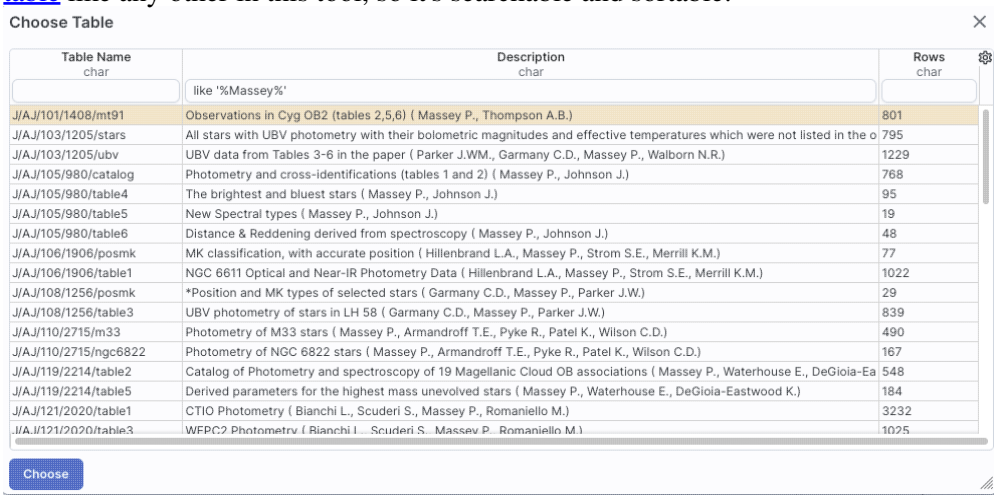

which makes it trivially easy to find tables in which you are interested, such as those associated with papers by Massey, as shown. Select the table that you want to search, and then continue to specify the rest of your search, just as described above in the IRSA VO TAP Search screen discussion.

The result of any VO TAP search is a catalog that can be interacted with like any other catalog in this tool.

NED Objects -- Searching for NED objects

(NED= NASA Extragalactic Database.)

To see this tab as a choice on the top, you may need to select it from the side menu.

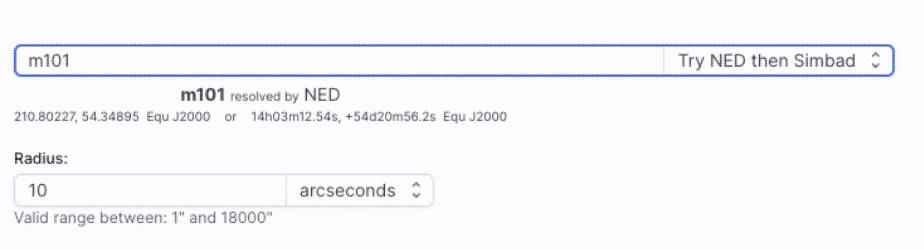

As for the other catalog searches, the tool may pre-fill the target position with its best guess of the coordinates of the target with which you have been working. You can use an object name in place of coordinates. Note that although NED is used for name resolution, the actual search is then performed using coordinates, as opposed to name. In this case, you are limited to a cone search, so the next option is the cone search radius. Pick your units from the drop-down first, and then enter a number; if you enter a number and then select from the drop-down, it will convert your number from the old units to the new units. There are both upper and lower limits to your search radius; it will tell you if you request something too big or too small.

From the NED results, you have one-click access to the fully detailed information from NED on any object returned from such a search. The search results will generally include a column "Details", though you may have to scroll to the right to see it. Clicking on a link in this column takes you directly to the full NED information display for the selected object in a new window.

The result of any NED search is a table that behaves like any other table in this tool.

VO SCS -- VO Simple Cone Search

(SCS = simple cone search.)

To see this tab as a choice on the top, you may need to select it from the side menu.

As for the other searches, the tool may try to pre-fill the target position with the coordinates of the target with which you have been working. In this case, you are limited to a cone search, so the next option is the cone search radius. As usual, pick your units from the drop-down first, and then enter a number; if you enter a

number and then select from the drop-down, it will convert your number from the old units to the new units. There are both upper and lower limits to your search radius; it will tell you if you request something too big or too small.

If you know your VO URL already, you can jump down to the Cone Search URL box and type or paste your URL into the box and hit search.

More commonly, however, users do not know *a priori* which URL to use. Click on "Find Astronomical Data Resources" to be dropped into a VO search. Find the URL corresponding to the catalog you want, copy it, and go back and paste it in the URL box. The URL should not have the RA and Dec in it; the tool will add your RA and Dec as listed to the URL in the right syntax. Click on "Search" to initiate the search.

Example

Load the tool. Search on IC1396. Go to the catalogs tab. Choose "VO Catalog." It wants the root URL for a cone search. Click on "Find Astronomical Data Resources", which goes here. Search on IPHAS. Get this page. Look for the complete catalog release (not just one associated with one specific study). The name of the catalog goes here. Hit the [+] to expand it. There is one URL listed there, under "available endpoints for the standard interface." Copy that URL and paste it into the search form. The IRSA tool will append your coordinates and radius and return you a table.

Tips and Troubleshooting

- Note that searching the VO means that you are using resources not specifically housed at IRSA, so servers may be down, or timeouts set, or limits on numbers of returned sources, etc., that are beyond our control. In most cases the solution is to specify as precise a search as possible. The URL you enter into the box in the search panel must be a Cone Search base URL (not containing RA and Dec parameters, which are inserted into the URL by the tool in response to the search parameters you give it).
- The master list of registries is here. You can also search the registries directly via that link (as opposed to via the IRSA tools).

The result of any VO SCS search is a catalog that can be interacted with like any other catalog in this tool.

CADC ObsCore Search

An ObsCore search is technically a subcategory of a TAP search, but it is a special subcategory in that it can return images, spectra, catalogs, and more, or even links to services. As such, it is a wholly different section than the General TAP section above.

There are ObsCore servers all over the world, but the default server for this release of IRSA Viewer is the CADC ObsCore service. To see this tab as a choice on the top, you may need to select it from the side menu.

This is what you see when you go to this tab:

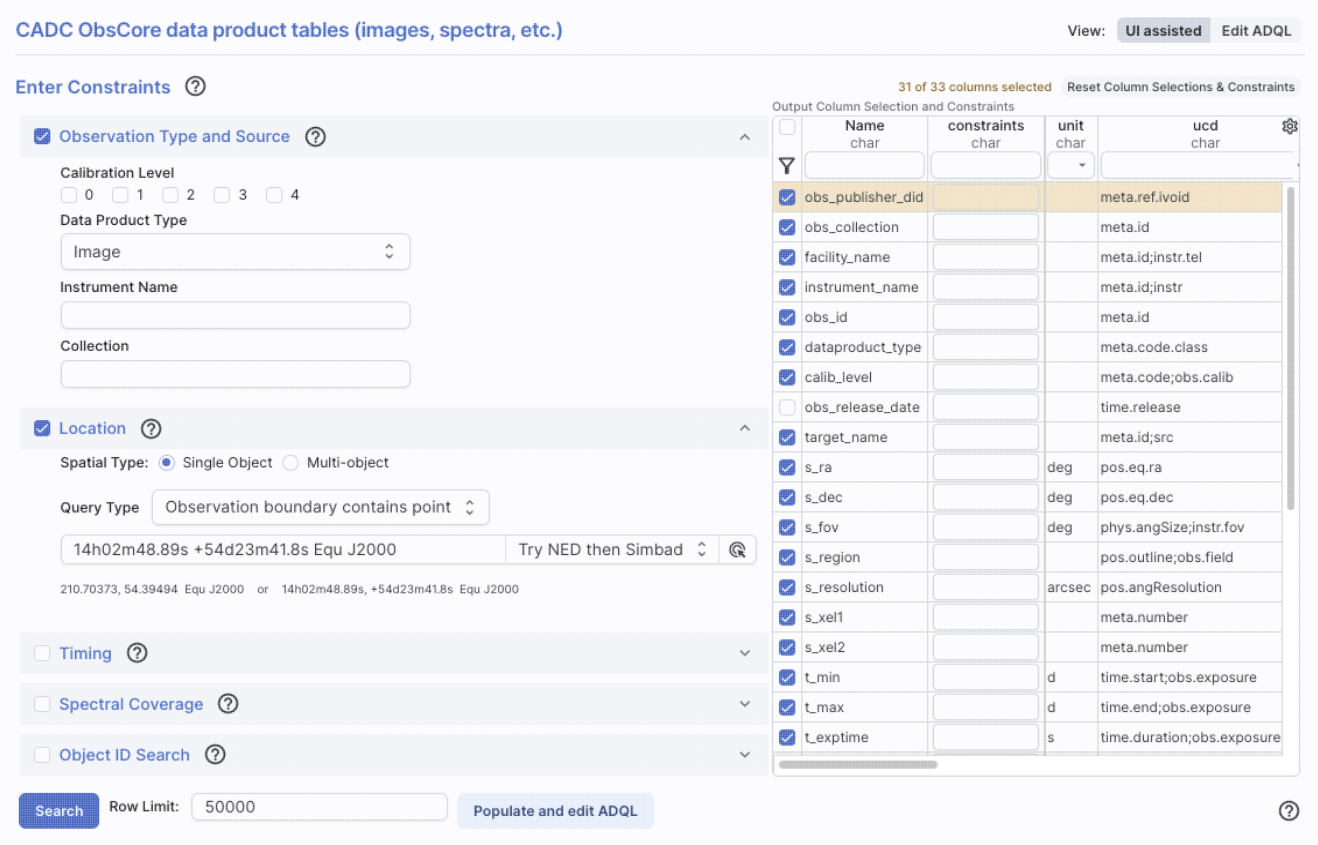

The "UI Assisted"/"Edit ADQL" switch at the top right works just like it does above -- use the UI to construct a query or dive right into the ADQL yourself.

Then in the remaining part of the screen, impose the constraints you want -- see the constraints section above.

At the bottom of the screen, you can "Populate and edit ADQL" if you want to work with the ADQL directly, or just search straight away after setting your search parameters.

Here are results of a basically unconstrained search on M16:

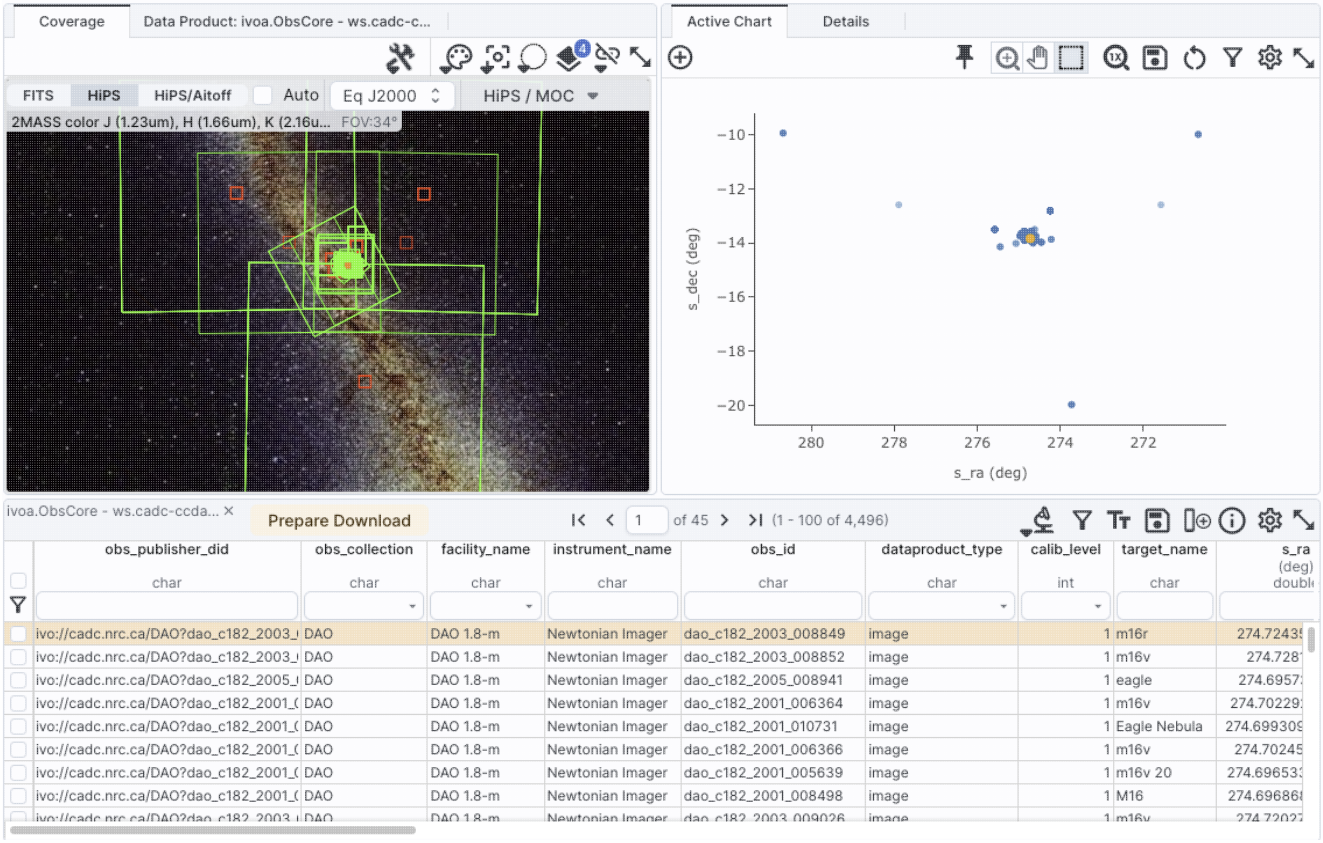

The coverage image on the left shows the polygons of coverage of the observations it found, and the plot on the right is the (relatively uninformative) plot of the positions associated with the observations. The table on the bottom as a list of the observations that it found consistent with the search parameters. This table is like any other table in this tool, so it can be sorted, filtered, etc. Data can be selected for download via the tickboxes on the left. The other tab on the upper left ("Data Product") shows a preview, when possible, of the data product corresponding to the highlighted row in the table. This interface is very similar to that for the "Data Collections" tab.

Tips and Troubleshooting:

- If you want to do an ObsCore search with a service other than CADC, go to "Multi-archive VO TAP", select your service, and if that service supports ObsCore, then an ObsCore switch appears. Toggle that if so desired. 

The result of any ObsCore search is a table that is a list of observations, and that table can be interacted with like any other table in this tool. However, it is a table of observations and/or services, so it yields much more than just a row in a catalog; it can give you images, spectra, and more.

Because ObsCore searches can produce search results that are services, when the tool encounters services, it gives you different choices. In the upper left of the search results, you may see a 'more' option, which expands like this:

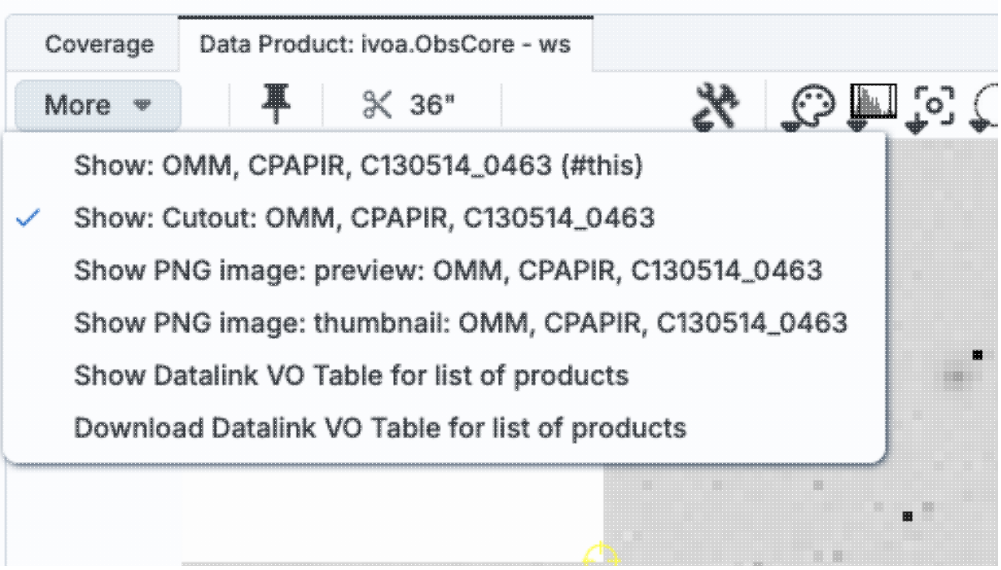

As shown here, the 'cutout' option is selected, and additional choices appear above the image. The scissors icon allows you to choose the size of a cutout centered on your target from the selected image product.

IRSA Viewer: Data Collections

Contents of page/chapter:

Introduction

The Data Collections tab is one of the searches you can add to IRSA Viewer using the side menu to add to the tabs at the top. (It is also essentially the same as the separate, standalone Data Collection Explorer tool!) The information we cover here is for the Data Collections tab within IRSA Viewer.

In order to learn more about any enhanced data product, the particular values that are returned, how the data were reduced, its strengths and weaknesses, etc., please see the documentation corresponding to the products.

Searching

When you first load the Data Collections tab, you see this screen:

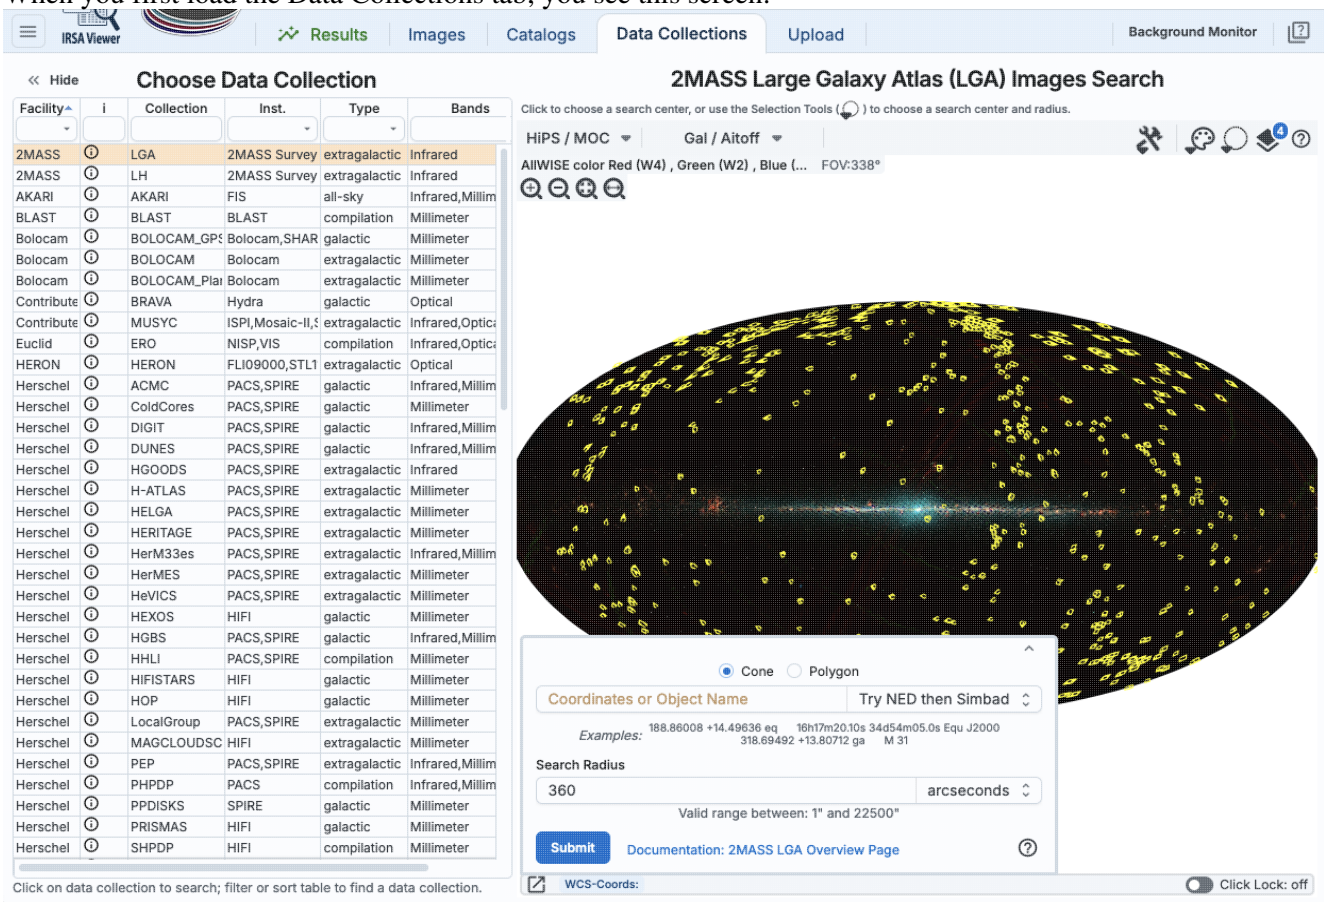

On the left, you see a table entitled, "Choose Data Collection." This table is like any other table here in that it is searchable, sortable, filterable, etc. Each row here corresponds to a data set that is currently available within this tool.

If you select a row in this table by clicking on it, then the contents of the image on the right changes in response. If the dataset is all-sky (like AKARI or IRAS/ISSA), then no polygons appear on the right. If the dataset covers just portions of the sky, then the polygons on the right indicate the sky coverage of the dataset. In the example above, the 2MASS LGA is selected on the left, and the small, yellow polygons on the right show the footprints of the individual images that make up the LGA survey.

Tips and Troubleshooting:

- The HiPS image as well as the MOC are customized to each data set, so it may take a few seconds to update. For some data sets with many small footprints (e.g., z0MGS), it may take longer to render the MOC than others.
- If you know exactly what you want, you don't have to wait for the MOC to render before you search -- just type in your target and 'Search'!
- If no MOC is shown on the HiPS image, the data set is all-sky (or there is a bug!)... if it is an all-sky data set, it should say "Covers whole sky" in the title of the search portion of the screen on the right.
- The i with a circle ( ) in the table is a link to more information about the data collection. You can also click on the "Documentation..." link in the search box on the right.
- The image on the right is a HiPS image which has the MOC for the data set overlaid. As described in the discussion of the layers icon, specifically that on the HiPS MOC, you can change how the MOC is rendered; here it is attempting to draw the MOC as outlines, and its accuracy will be improved the closer you zoom in.
- If you want more screen 'real estate' to work with, click on "<< Hide" in the upper left to collapse the Data Collection table and maximize the space with the image and survey overlay.
- The search box can be collapsed (click on the 'disclosure arrow' in the upper right of the search box), or made translucent (click anywhere off the search box/on the image). If you put in some search information and then collapse the search box, a summary of the information you have entered is still shown when the box is collapsed.

You can navigate around on the image on the right just like you can on any other image here -- pan, zoom, etc.

To select a target, you can type in a name and have NED or Simbad resolve it into coordinates, type in coordinates in any of a variety of units, or click on the image to select a target based on your mouse position. This is just like entering a target anywhere else in the tool, or interactive target refinement elsewhere in the tool.

Enter a search radius, and click "Search" to start the search.

Results

A typical results page looks like this:

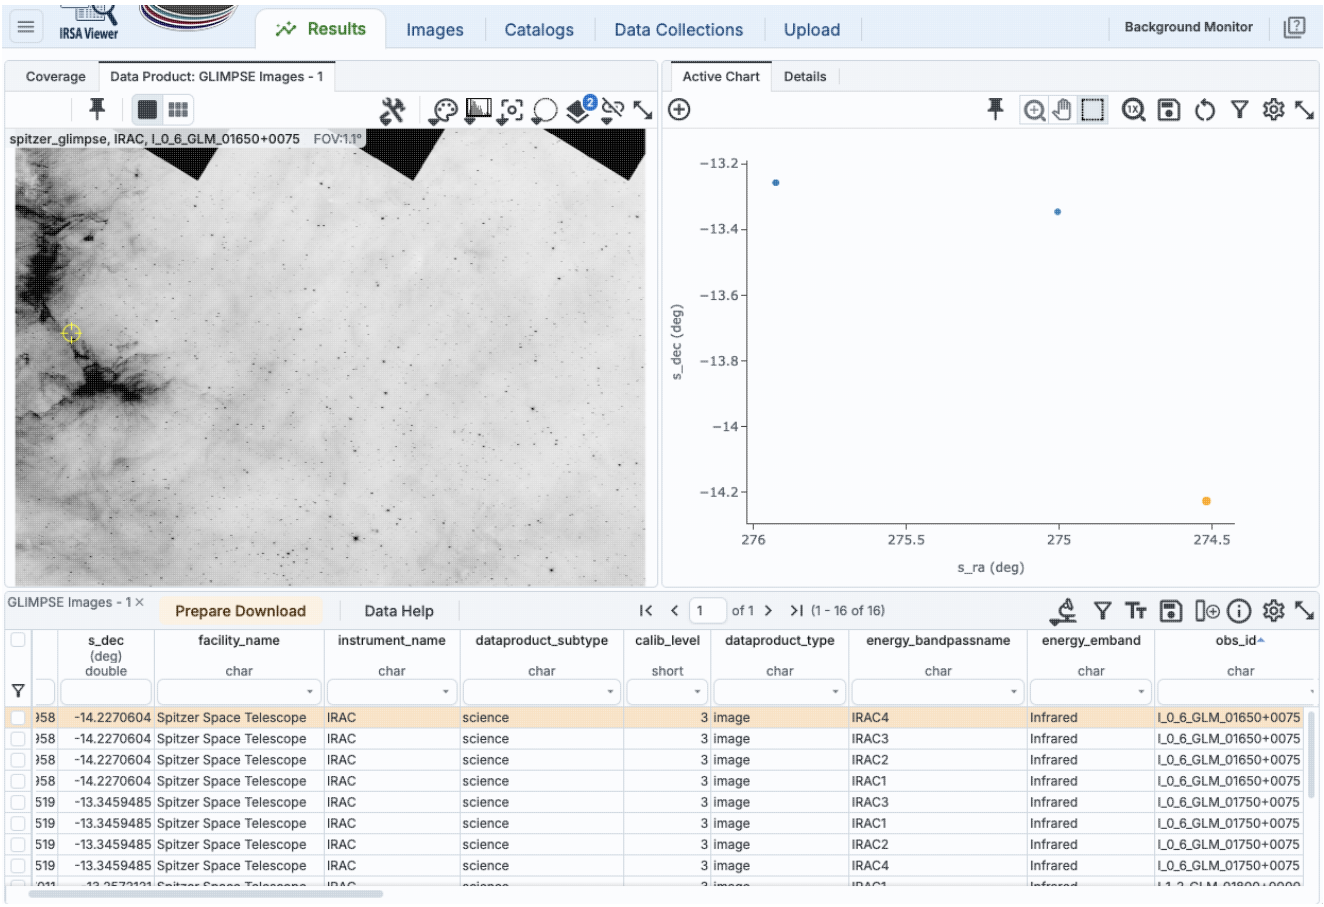

The table at the bottom is a list of the observations that it found consistent with the specified search. This table is searchable and sortable just like all the other tables here. The "Data Help" button at the top left of the data table will take you to more information about the data set. The image (or plot in some sitautions) in the upper left is a preview of the selected row in the table. Pin the image to hold onto it if you want.

The plot on the right is often initially relatively uninteresting, because it has just the positions of the data products it has retrieved. It's a plot like any other in this tool, so you can make it plot other things.

Tips and Troubleshooting:

- The coverage image, which is by default on the left, under the window showing data product previews, will show polygons corresponding to the data products' coverage. However, if there are too many polygons to manage, the tool may fall back to showing just the positions of the data products it has retrieved. (Whether that is the central point or the lower left corner of the image depends on the data product itself.)

Downloading Data

Here, you have a Prepare Download button, so select the data as described on the Download page, and click that button.

IRSA Viewer: Plots

Plots (sometimes called charts) can be made from Tables. Plotting is covered in this section. The Tables section discusses tables more generally, and the specific case of loading catalogs is in another section. If your table has RA and Dec in it, the Visualization section covers how the catalog can be overlaid on images. Note that spectra are in a different section entirely.

Contents of page/chapter:

Default Plot

By default, after a table has loaded, a plot appears in the browser window.

To obtain a full-screen view of your plot, click on the expand icon in the upper right of the window pane when your mouse is in the window:  . To return to the prior view, click the "Close" arrow in the upper left.

The plotting tool, by default, starts with RA and Dec plotted if it can find RA and Dec in the corresponding table. Note that it does so following astronomical convention -- RA increases to the left. If the catalog does not have RA and Dec, it plots the first two numerical columns it finds.

Plot Format: A First Look

If you have loaded a catalog with many (> 5,000) points, you may have an RA/Dec plot that looks something like the one on the left here. If you have loaded a catalog with few (< 5,000) points, you will have an RA/Dec plot that looks more like the one on the right here.

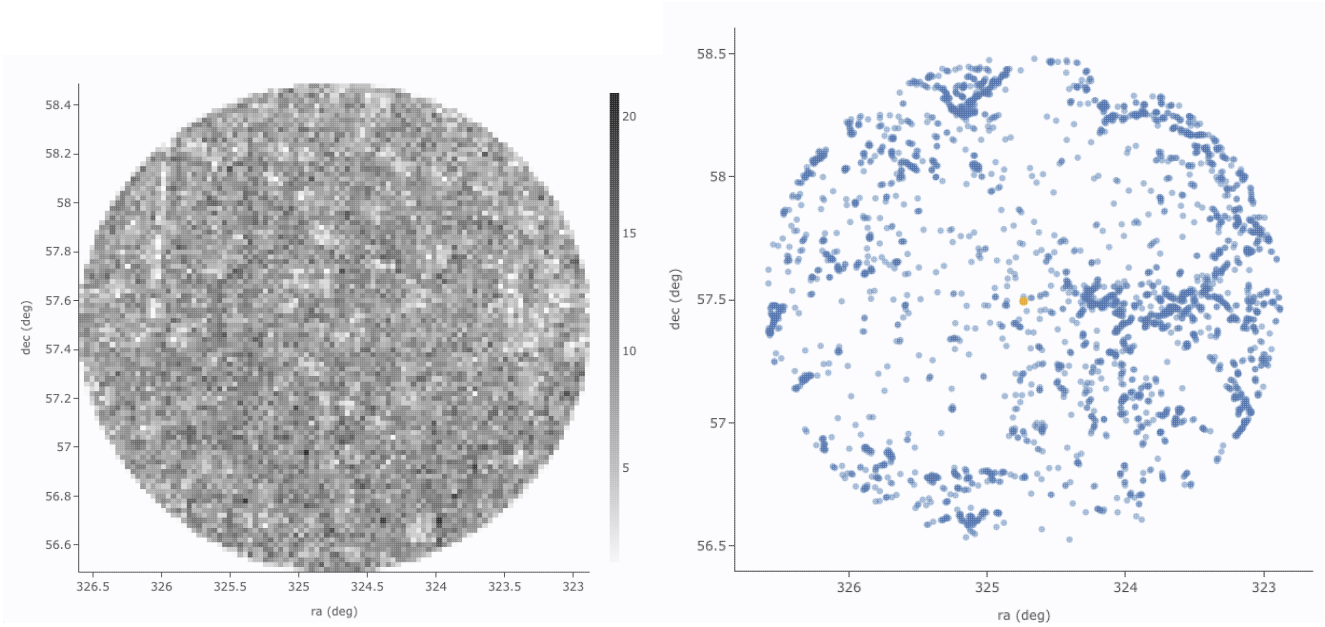

The difference between them is that, for larger catalogs (left), the plot is binned -- more points are encompassed in a black tile and fewer points are encompassed in a white tile. In the context of this tool, this is called a **heatmap**. The shades of grey correspond to how many points are encompassed in each 'cell', with the density scale given on the right hand side of the plot. For smaller catalogs (right), each individual point is shown as a blue dot. In the context of this tool, this is called a **scatter plot**. Note that even when individual points are shown, where the points overlap, the color is darker.

In either case, letting your mouse hover over a point tells you the values of the point under your cursor, and (if binned) how many points are represented:

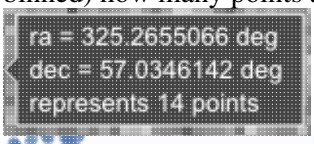

for binned plots, and

for just one point.

Clicking (in an unbinned plot) highlights that point, and it stays highlighted, though you must keep your mouse on the point in order to see the information about it.

The reason the tool makes a heatmap for large catalogs is to more fairly represent the point density -- and to make the plotting faster. In these cases, though, it will not give you the option to overplot errors (see below). If you have a heatmap and want a scatter plot *by default*, you need to filter or otherwise restrict the catalog to have fewer points (see below). You can change the bin size and shading via the plot options pop-up (more on this below). Or, you can force it to make a scatter plot anyway (see below).

Plot Navigation

The top of the plot window has a row of icons something like this:


which we now describe.

Add new plot

You may or may not have this icon. Clicking on this icon adds a new plot. This has a separate section below.

Pin plot

This icon may not always appear. Clicking on this icon pins the plot. This has a separate section below.

Show table

This icon does not always appear. Clicking on this icon pulls to the foreground the table that generated the plot that is currently in the foreground. This is related to pinning, which has a separate section below.

Combine chart

This icon does not always appear. Clicking on this icon attempts to combine plots; it has a separate section below.

Plot mode

This trio of icons controls the plot interaction 'mode'. By default, you are in 'selection' mode, as seen here -- the last icon is darker, like a pushed-in button. To activate the other modes, click on the other icons, and they become darker or "pushed in."

Zoom mode

When this mode is active, when you click and drag in the plot, the plot is zoomed to the region you have selected. Even when this mode isn't active, you can also zoom using your scroll feature on your mouse. To return to the original view, click on  .

Pan mode

When this mode is active, when you click and drag in the plot, it moves around in response to where you drag. To return to the original view, click on  .

Select mode

When this mode is active, when you click and drag in the plot, you are given additional options at the top of the plot :  The checkmark means "select" and the funnel means "filter." The difference is that filtering (temporarily) limits what is shown in the plot, catalog, and image (see general information on filters), and selecting just highlights the points enclosed within your selection. To cancel either one, click on cancel filters  or cancel selection  .

Re-scale plot

Return to the view that optimizes the range of x and y to show the currently displayed points.

Tips and Troubleshooting: Did you accidently zoom in the plot with your magic mouse or touchpad? Click on this icon to reset the plot.

Save plot

Save the plot. It will save as a png file, wherever your browser is configured to save files. The saved png is the same size as it is on your screen. If you want a big version, make the desired plot big on your screen (expand the view to take up as much space as possible) before saving the png.

Undo

Restore everything to the defaults. If you've played a lot with the plot, you may want to undo everything you've done. Click this icon to restore everything back to the defaults.

 *Filter from plot*

Pull up interactive filters. This button brings up filters for the displayed catalog in an interface like all the other tables here, except you don't see the values in the catalog themselves; you can enter filters here in the same way you can everywhere else in this tool (see general information on filters).

 *Configure plot*

Click on this icon to change what is plotted (much more on this below).

 *Expand plot*

Click on this icon to make the plot take up the whole browser window. To return to the prior view, click the "Close" arrow in the upper left.

 Help

This icon may not appear, but if it does, it is a context-sensitive help marker, which should bring you to this online help.

Plot Linking: Plots are linked to catalog and image(s)

If you move your mouse over any of the points in the plot, you will get a pop-up telling you the values corresponding to the point under your cursor. For scatter plots, if you click on any of the points, the object(s) corresponding to that point will be highlighted in the overlays in the images shown, and highlighted in the catalog table. This works the other way too -- click on a row in the catalog, or an object in the images, and the object will be highlighted in the plot or the catalog or the image.

Changing What is Plotted

To change what is plotted, click on the gear icon in the upper right of the plot window pane:  .
Configuration options then appear; the options are a little different depending on whether the points are binned or not. This section describes how to change what is plotted, i.e., the "Modify Trace" option at the top of both of these pop-ups. The overplotting option (and, for that matter, adding plots) are covered in more detail below.

This is the configuration window for a binned (a.k.a. heatmap, or greyscale) plot. By default, the "chart options" may be hidden; to reveal them, click on the name "Chart Options" or the disclosure arrow on the right. To hide them again, click on the disclosure arrow on the right.

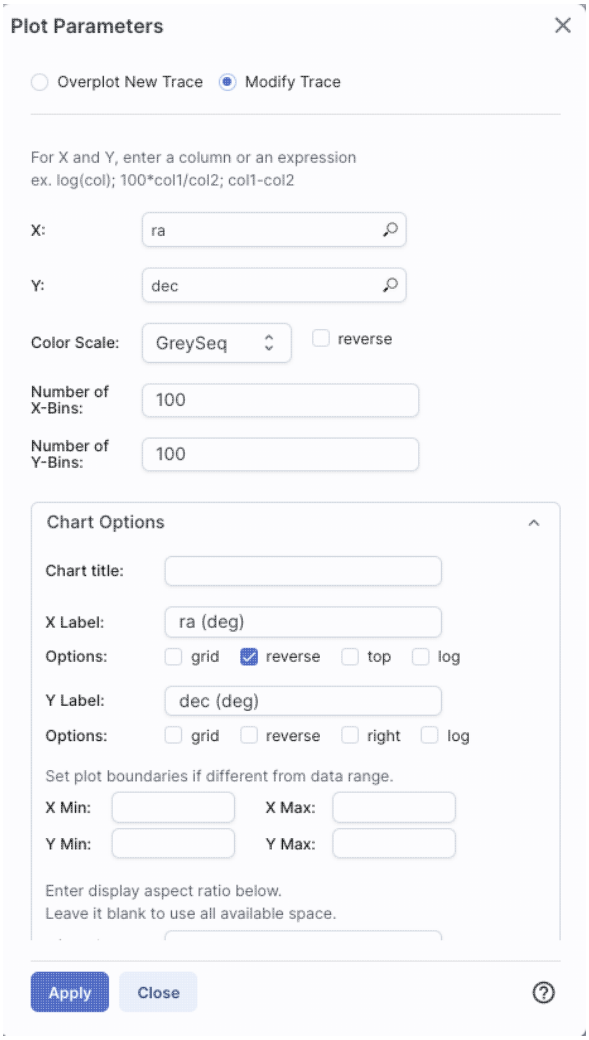

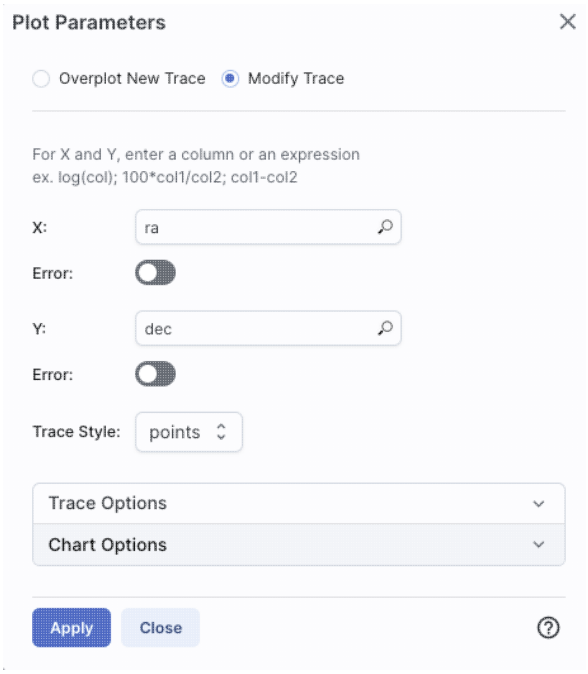

The configuration window for a plot that shows individual points, once fully extended, is much longer (and scrollable), and so is shown here in two parts. Both the "Trace Options" and "Chart Options" may be hidden by

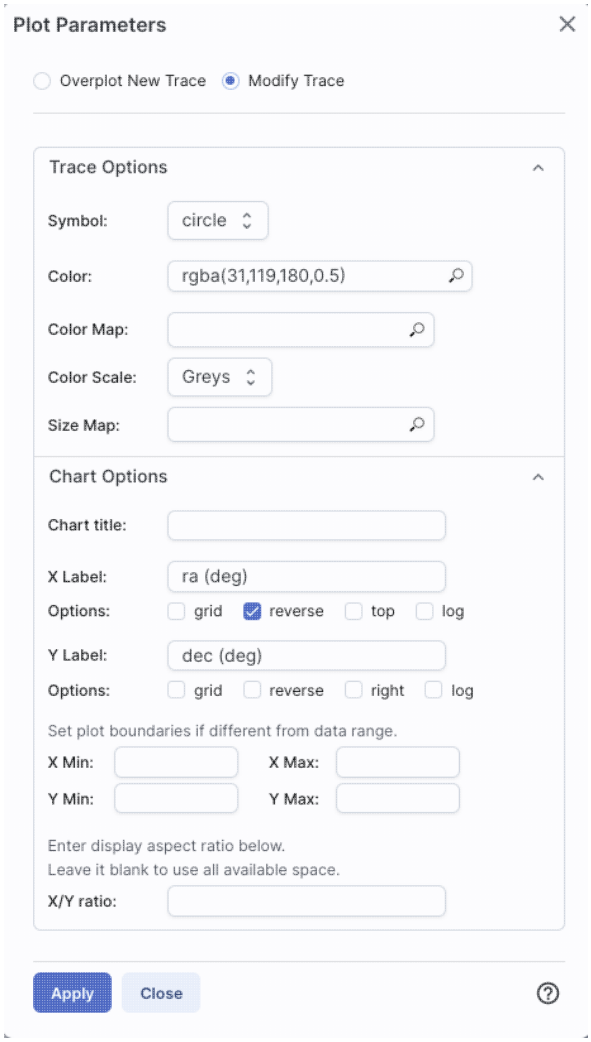

default; to reveal them, click on the name or the disclosure arrow on the right. To hide them again, click on the disclosure arrow on the right.

Options found in both kinds of plots

In either case, you can **specify what should be plotted on each axis**. The magnifying glass is a link that brings up a table that lists all of the available columns in the catalog. Alternatively, you can just start typing, and viable options appear below the box. Whatever you put in the box must match the column name as shown in the catalog *exactly*.

Click on the black triangle to reveal additional options.

In both of the examples above, RA is plotted on the x-axis. It has pulled the column name for the label; in this table, the column is "ra" rather than "RA", and it is case-sensitive. It has copied over the units ("deg") from the catalog, and plotted the x-axis increasing to the left as per astronomical convention. You can change what column is plotted, and whether or not errors are shown. Under "**Chart Options**", you can specify:

◊ title of the plot;
◊ labels on the x and y-axis;
◊ whether or not there is a grid shown;
◊ whether or not the axis is reversed (as for ra in the examples above);
◊ whether the x-axis is on the top or bottom and the y-axis is on the left or right;
◊ whether or not the axis is logarithmic;

- ◊ the maximum and minimum values of the plot range;
- ◊ the aspect ratio of the plot (e.g., square or rectangular).

By default, the boundaries of the plot are set to encompass the full data range. Here you can change the boundaries to specific numbers. (This can also be set via filtering from the plot; see below.)

You can enter **simple mathematical relations** in these boxes too, such as (for a WISE catalog) "w1mpro-w4mpro" to put [W1]-[W4] on one axis. Supported operators:

- ◊ +,-,*,/
- ◊ abs(x), acos(x), asin(x), atan(x), atan2(x), ceil(x), cos(x), exp(x), floor(x), lg(x), ln(x), log10(x), log(x), power(x,y), round(x), sin(x), sqrt(x), tan(x)
- ◊ degree(x) and radians(x) are also supported -- these are the same functions as in ADQL and convert radians to degrees or degrees to radians. For small astrometric offsets, you could make a scatterplot of dec2-dec1 vs. (ra2-ra1)*cos(radians(dec1)) instead of typing cos(dec1*pi()/180). (NB: pi() is also a supported function you can use, instead of typing 3.14159.)
- ◊ Non-alphanumeric column names (e.g., those with - or + or similar characters) should be quoted in expressions.

Click "Apply" to apply, and "Close" to return to the plot without making changes. (For the latter, you can also click the 'x' in the upper right.)

Options found only in binned plots

(Plots are binned by default if there are > 5,000 points in the catalog.) From the pop-up, you can control the color table that is used (greyscale is the default; there are many other choices in the drop-down menu), as well as the number of bins in the x and y directions. The default value for the number of bins is 100 in both directions.

Options found only in plots showing individual points

You can add errors. Toggle the error switch, and then additional choices appear. From there, you can select symmetric or asymmetric errors, and then you can specify an error as either an existing column in the catalog, or calculated from a column in the catalog.

Under "**Trace Style**," you can control whether the points are shown as individual points, connected points, or just lines connecting the points.

Under **Trace Options**, you have many choices.

- ◊ Choose the **symbol type**: circle (default), open circle, square, open square, diamond, open diamond, cross, x, upward-pointing triangle, hexagon, or star.
- ◊ Choose the **color**. By default, the point color is a mid-range blue that is darker where more points. This is specified by the rgba vector shown in the example here (31, 119, 180, 50) where the last number is in units of fraction of 1, so 0.5=50% in this example. Click on the magnifying glass to bring up a color picker window:

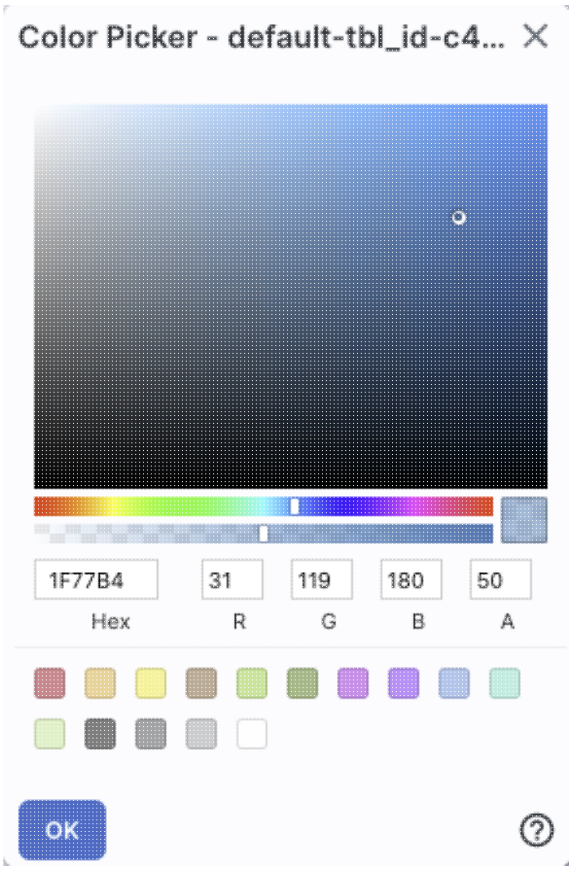

From here, you can click on your desired color in the top colorful box. Immediately below that box, you can change the color and saturation of the top box so that you can select from a different range of colors. Below that, you can enter numerical hex codes or RGBA values (where the value for RGB is between 0 and 255, and A is in units of percent, e.g., 50 = 50%). Finally, you can also select from a pre-defined set of 15 colors by clicking on any of the small boxes. Note that the numerical codes update as you select different colors. Click "OK" to implement your color choice, or click 'x' in the upper right to close the window without changing the color.

Tips and Troubleshooting: Don't like the transparency feature of the points that makes them darker when there are more points? Set the last value of the vector (A) to 1. Don't like the blue? Pick a diferent color entirely. Want the faintest point to be brighter than it is by default? Set the last element of the color vector ("A") to be 0.7 or 0.8.

◊ Choose the **color map**. By default, all of the points are the same color, but darker where there are more points. You can change this such that the color scale of the points is tied to a column value, such as w1snr (WISE-1 signal-to-noise ratio) in a WISE catalog. If you select this option, you can also change the color scale to any of many different options (see the drop-down). Simple mathematical relations (as above) are also permitted in this box.

◊ Choose the **size map**. By default, all of the points are the same size. You can change this such that the color scale of the points is tied to a column value, such as w1snr (WISE-1 signal-to-noise ratio) in a WISE catalog. Simple mathematical relations (as above) are also permitted in this box.

Example: Load a large WISE catalog. Plot w1snr (WISE-1 signal-to-noise ratio) vs. w1mpro (WISE-1 profile fitted magnitude). It defaults to a heatmap. Change the labels, making the y-axis label "WISE-1 SNR" rather than the more cryptic column header "w1snr". Change the x-axis label to "[W1]." Change the greyscale to

yellow-green-blue ("YlGnBu") to make it easier to see the lowest-populated bins. Depending on your catalog, you may need to adjust the ranges. Obtain this plot:

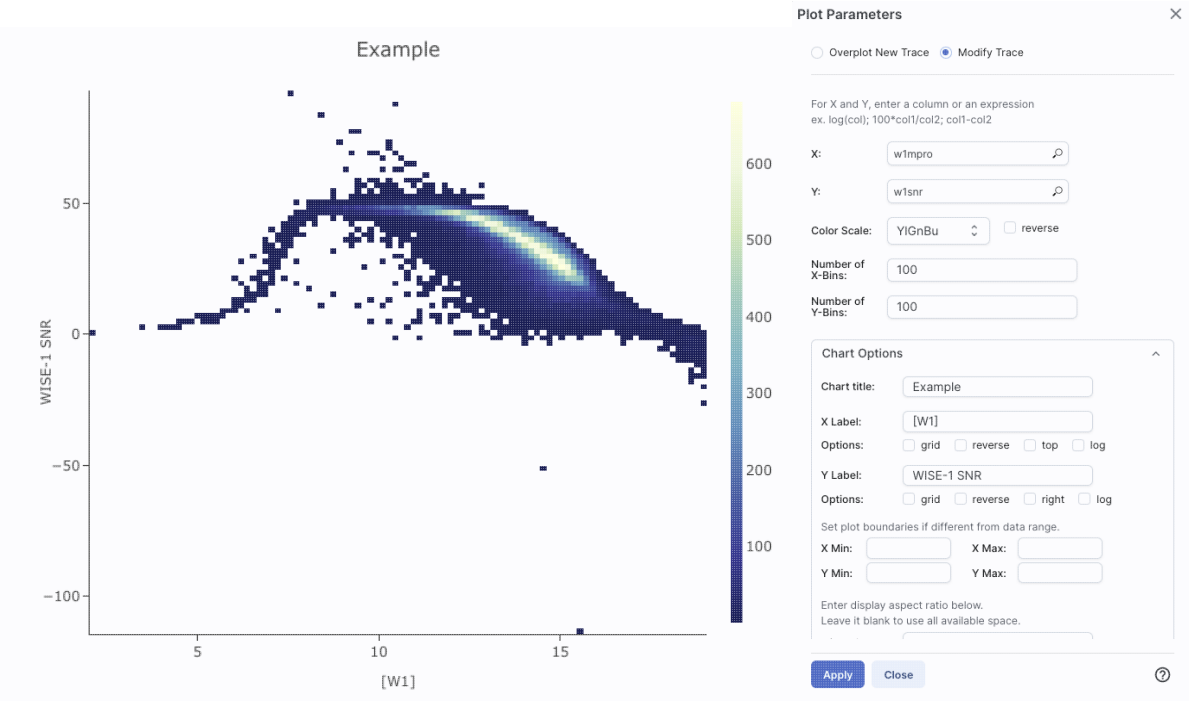

Example: Load either a smaller WISE catalog, or the same large WISE catalog, but filter it down such that w1snr, w2snr, and w3snr are all greater than 10, which limits the number of points to be <5,000. Plot w1snr vs. w1mpro. It shows the points individually. Change the labels. Change the point color map to scale with w2mpro (WISE-2 profile fitted magnitude). Change the point size map to scale with w4snr (WISE-4 signal-to-noise). Obtain this plot:

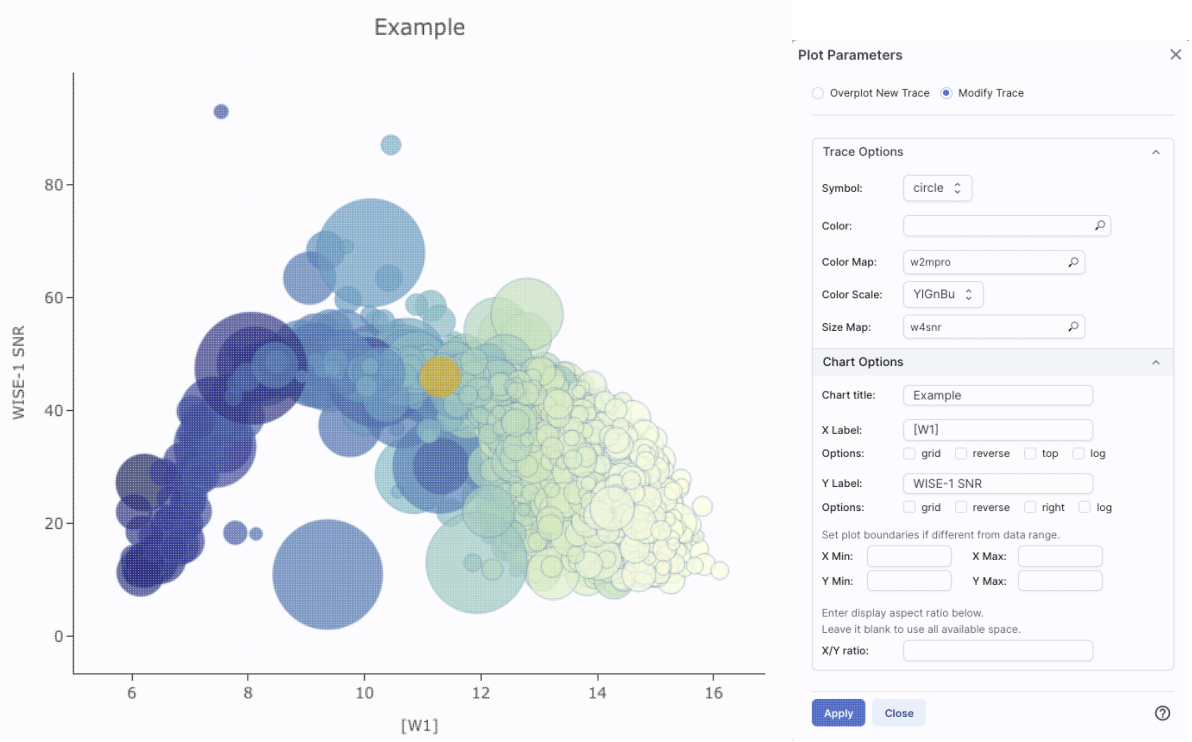

Plotting Manipulated Columns

You can choose a single column to plot against another column, as above. However, you can also do simple mathematical manipulations.

For example, if you have loaded a WISE catalog, you can plot [W1]-[W2] vs. [W3]-[W4]. In terms of the names of the columns in the database, this is w1mpro-w2mpro vs. w3mpro-w4mpro.

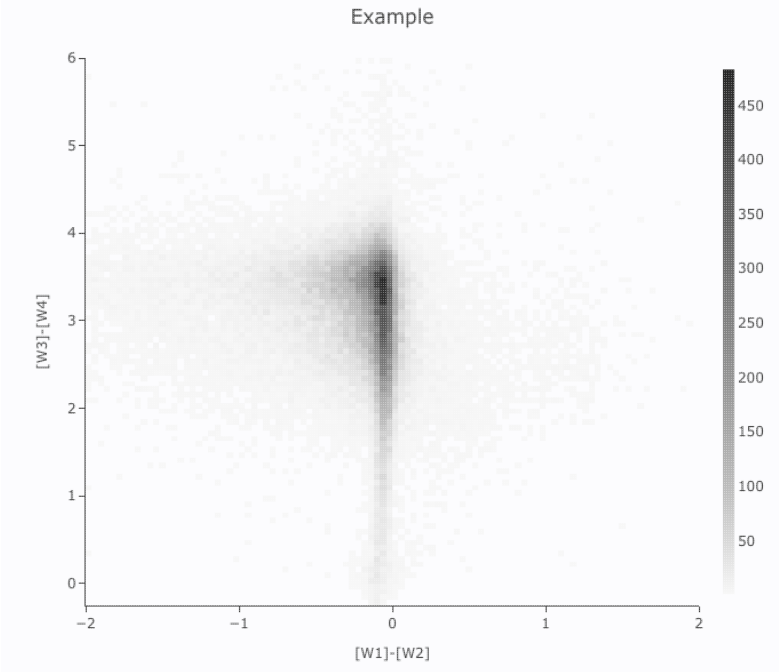

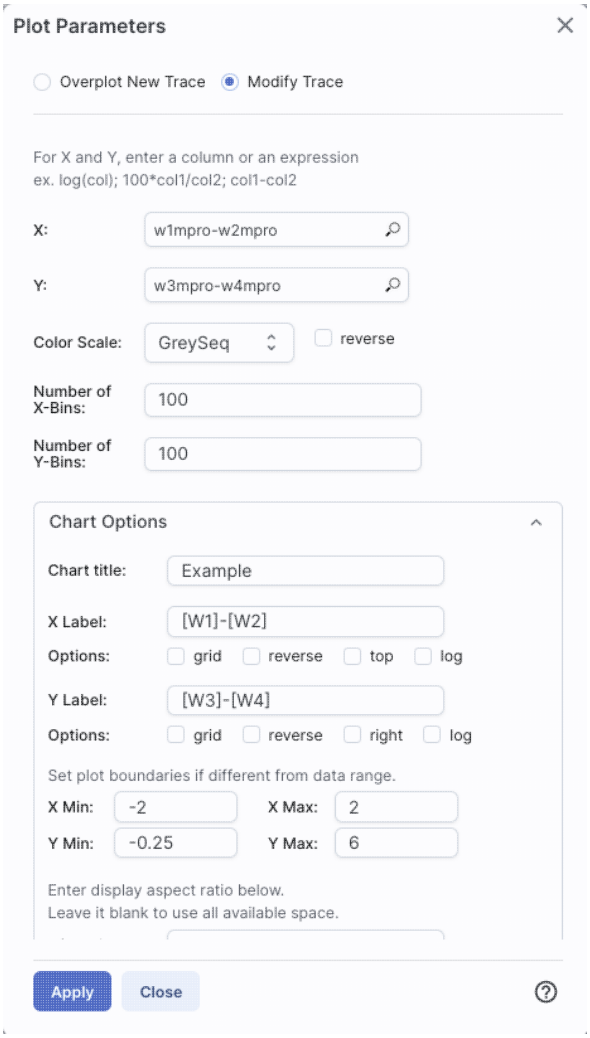

If you have few enough points that the plot is not binned, you can add errors that you calculate. Here, the expression for the x-axis errors is sqrt(power(w1sigmpro,2)+power(w2sigmpro,2)) and for the y-axis errors, it is sqrt(power(w3sigmpro,2)+power(w4sigmpro,2)) -- that is, the errors for the individual photometric points added in quadrature.

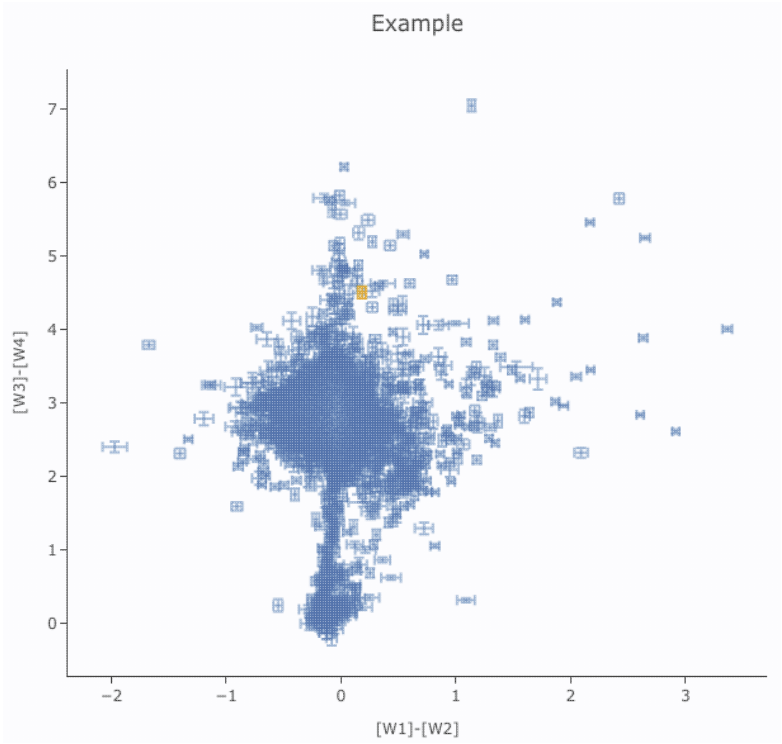

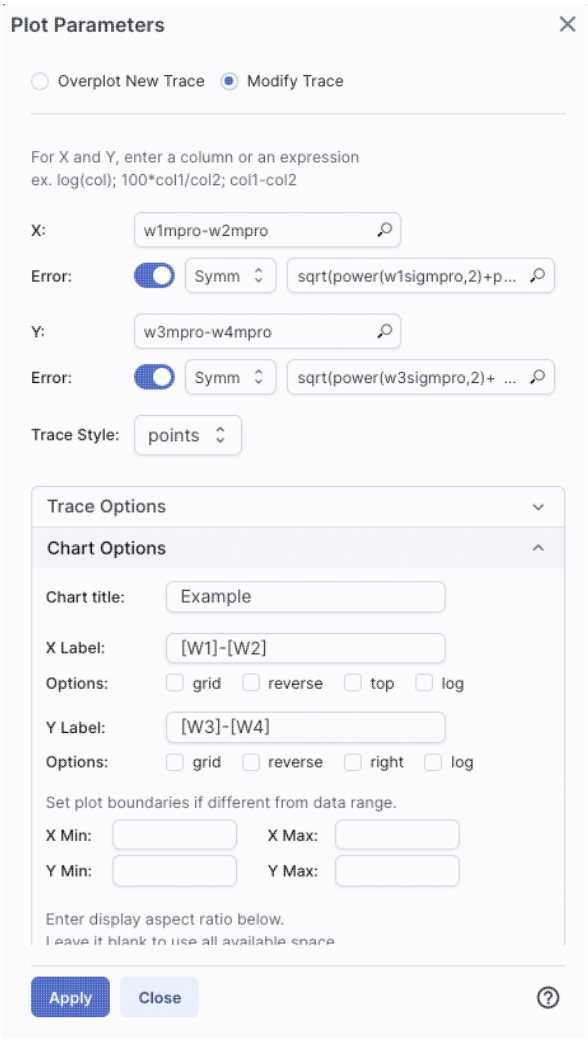

Restricting What is Plotted (from the plot)

You can also restrict what data are plotted in any of several different ways.

You can filter the catalog from the table itself (discussed in another section).

You can set axis limits on the plot itself from the plot options pop-up (discussed above).

However, and perhaps more powerfully, you can set limits from the plot itself using a rubber band zoom. Click on the select icon in the plot  Then, click and drag in a sub-region of the plot. New icons appear:  If you click on the funnel icon, only those data points that pass the filter are shown in the plot, in the table, and/or overlaid on the image(s). (This is the behavior of 'filter', as opposed to 'select'; the former restricts what is shown, the latter just highlights the points.) For more on filters, see the filtering discussion in the tables section.

Example: Obtain a WISE catalog of a star-forming region, say IC1396. Filter down the catalog to only have detections at all four WISE bands. (Limits have undefined errors, so ask the catalog to filter down such that w1sigmpro>0, w2sigmpro>0, w3sigmpro>0, and w4sigmpro>0). Plot w1mpro-w4mpro on the x-axis, and

w1mpro on the y-axis. Reverse the y-axis to put bright objects at the top. Click and drag in the plot to select the bright and red objects, and filter them down to get a subset of bright and red sources. For clarity, the screenshot here has the sources selected, not filtered.

Overplotting

At the top of the pop-up that you get when you click on the gears, you have two radio buttons:

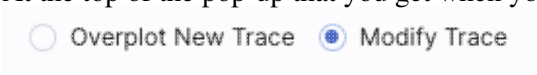

They are "Overplot New Trace" and "Modify Trace." Modifying traces (plots) has been covered above; in this section, we will cover overplotting. This is sometimes called "multi-trace," meaning that more than one thing is plotted.

When you select "Overplot New Trace," you get a new interface that is very similar to the original interface where you selected what to plot:

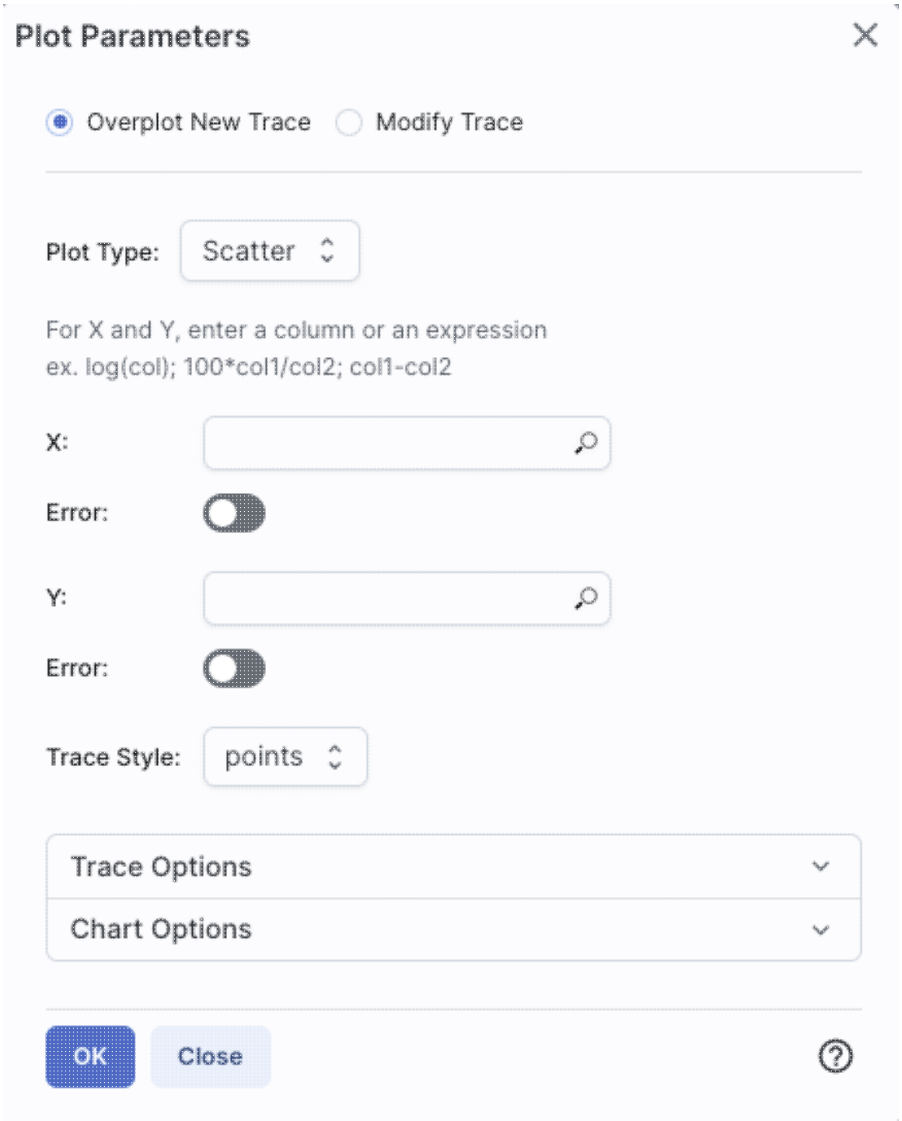

As before, you need to :

- select a plot type (scatter, heatmap, histogram);
- tell it what column(s) (and and manipulations thereof) you want for x, y, and associated errors;
- select the trace style (points, connected points, lines);
- set any additional trace options;
- set any additional chart options.

The best way to explain how to use this feature is probab
example. We have a plot of [W1] vs. [W1-W4] from abo
add on top of it a plot of [W2] vs. [W1-W4]. Click on the
bring up the pop-up. Select "Overplot New Trace." Enter
"w1mpro-w4mpro" for x and "w2mpro" for y. Expand "C
Options." Note that it has preserved the overall chart title
before, but has erased the X and Y labels (and lost the re
the y axis) because the overplot could literally be anythin
need not be the same columns or even the same units as
already plotted. Type them in again. Here is the configur

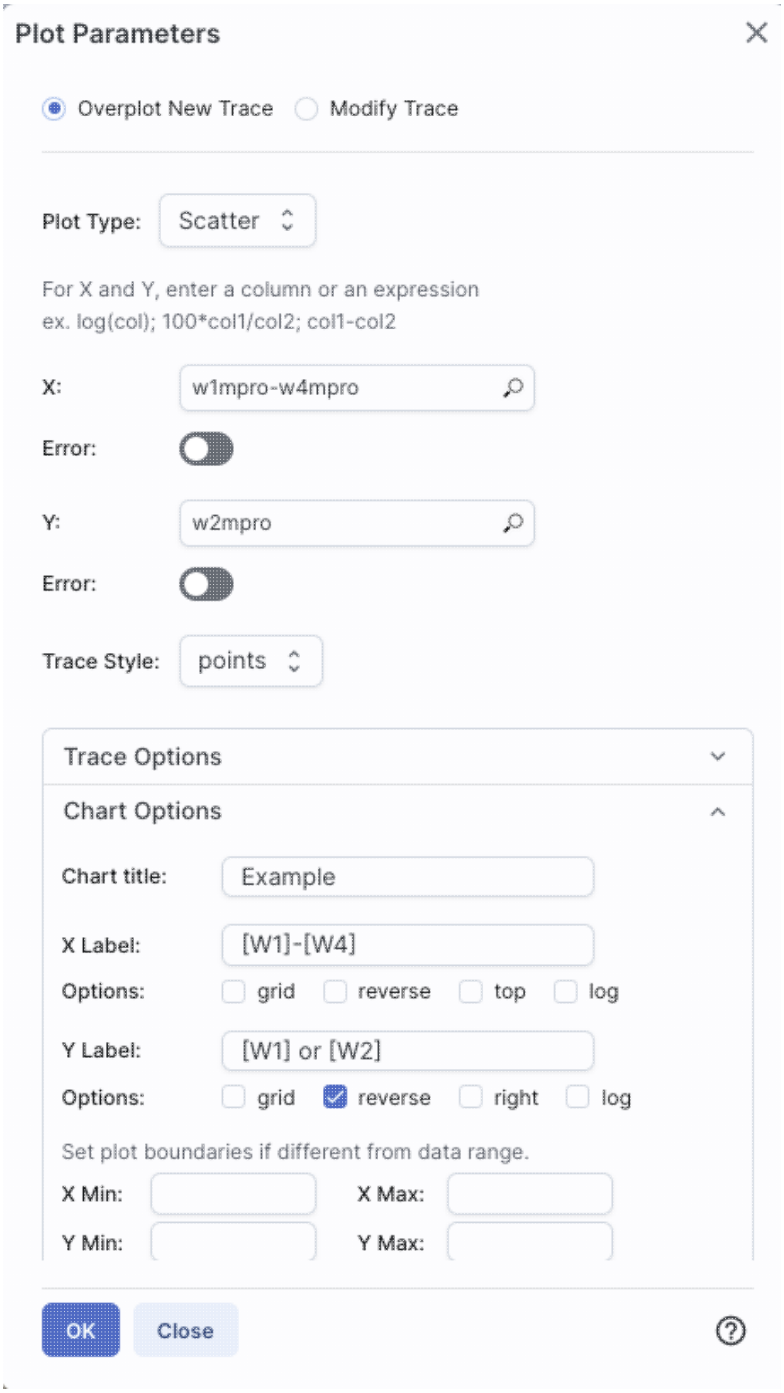

window right before clicking "ok", and the resultant plot

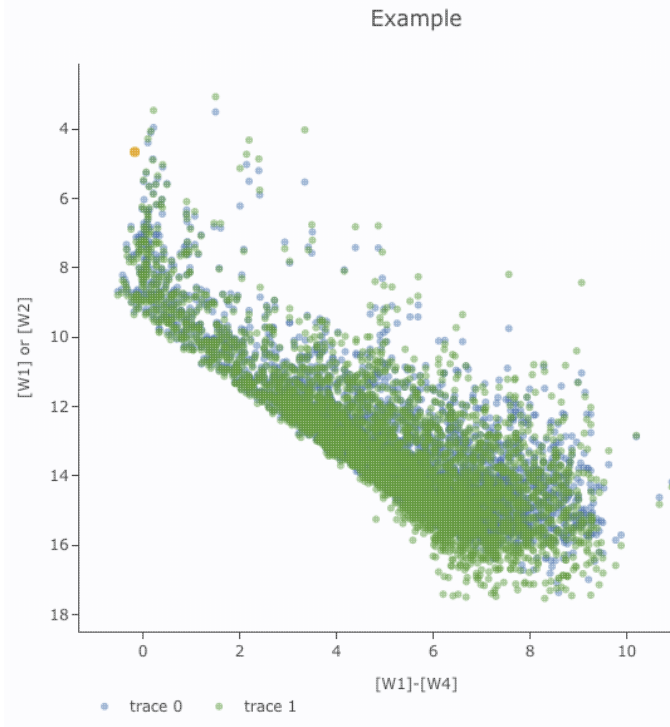

After you add the overplot, if you click on the gears agai
that the choices at the top of the window have changed. Y
add another overplotted trace, modify a trace, or remove
trace. Each trace that you add is a new 'layer' on the plot.
drop-down menu near the top of the window controls wh
is 'active' for setting the x, y, errors, trace style, name, sy
color, etc. there is now a drop-down menu at the top of th
There is a legend on the plot specifying which color corr
to which trace. In this example, the plot above has appea
a blue and green color scheme, which may be too hard to

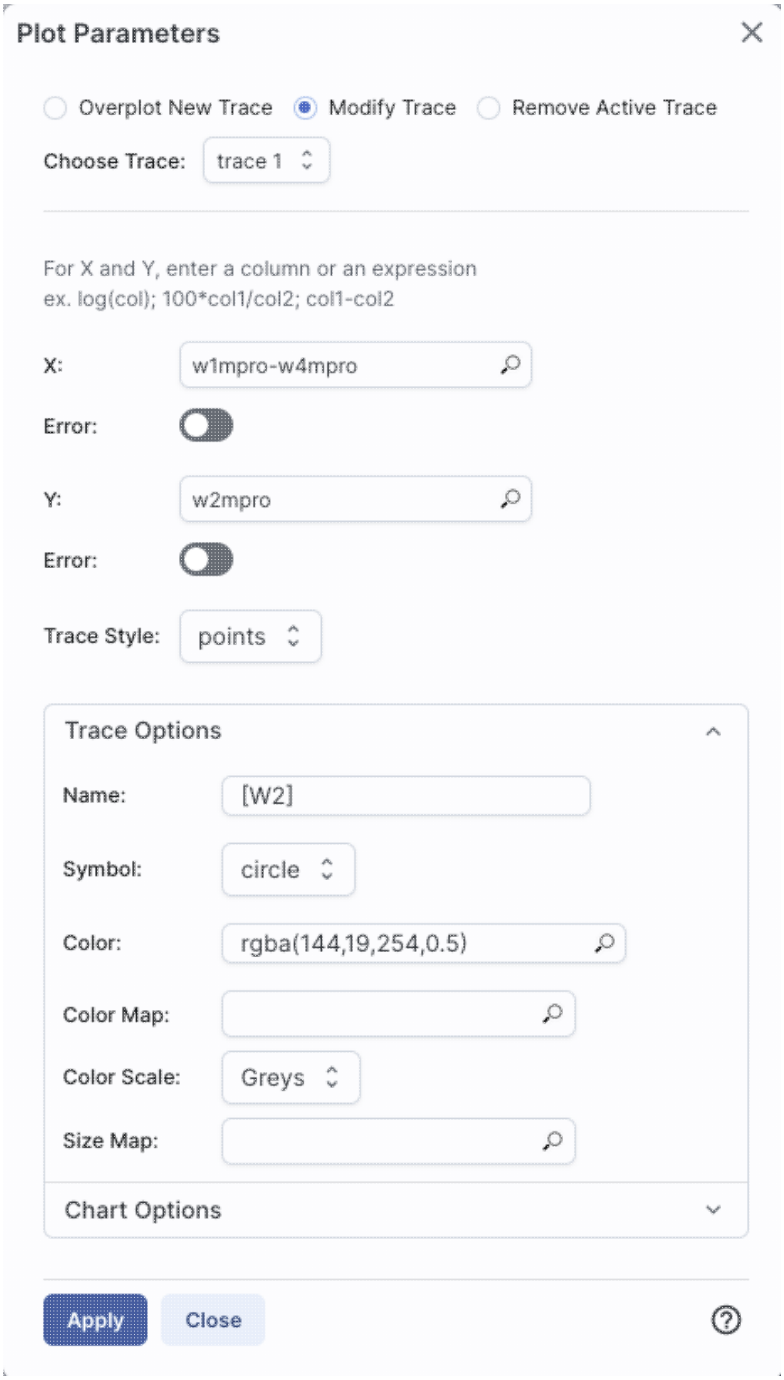

differentiate. To change the new points' color, click on th
ensure "Modify Trace" is selected, select "trace 1" (as op
"trace 0", the first one you loaded), go down and expand
"Trace Options" and pick a different color. You can also
the legend name from "Trace 1" to, in this case, "[W2]".
"apply" to apply the changes to the plot. Note that once y
change the trace name, the relevant drop-down menus in
pop-up window and the legends on the plot update accor

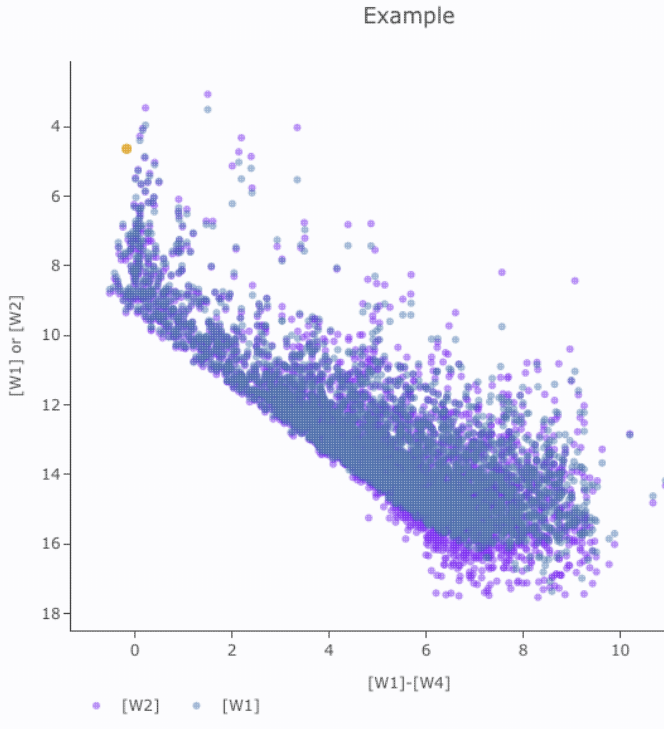

Note that the pop-up spawned by clicking the gears now has an additional option at the top: "Add New Chart", "Overplot New Trace", "Modify Trace", and "Remove Active Trace." From here, you can modify a trace you have already plotted (as described above), overplot another trace (also as described above), or remove the selected trace:

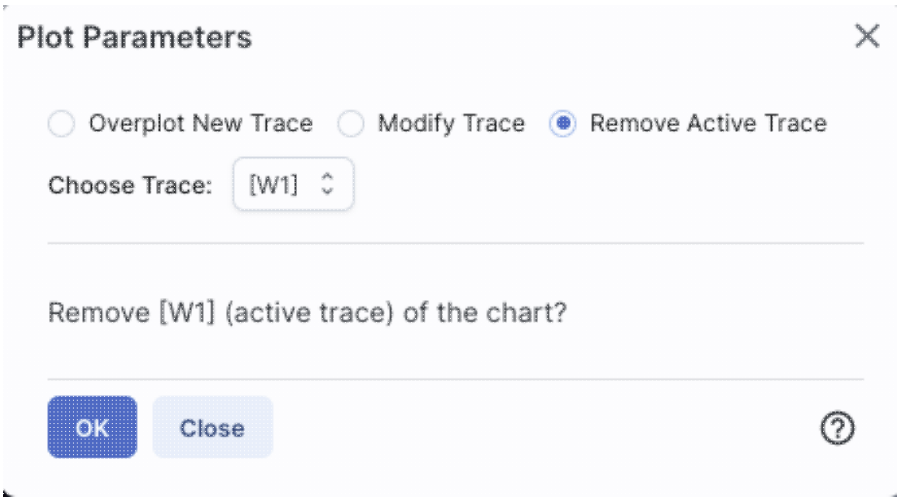

Tips and Troubleshooting

- Right now, the overplotting only works from the same catalog -- that is, you cannot plot [W1] vs. [W1]-[W4] from one catalog and overplot [W1] vs. [W1]-[W4] from another catalog. (We enthusiastically await this capability too.)
- You can easily get yourself into a physically nonsensical situation, say, by overplotting a histogram onto a scatter plot. If you find yourself in a hopeless mess, click on the "undo" icon to reset everything and try again.
- When you have more than one thing (trace) plotted, double click on the legend to bring that trace to the foreground and temporarily hide the other traces.
- You can overplot a scatter plot on top of a heatmap if you really want to!

The context where this feature really shines is in plotting multiple spectral orders. In that case, it makes complete sense to plot many things from the same file (and only things from the same file) on the same plot. However, spectra are sufficiently different and complicated that all of that information is collected into a chapter about spectra.

Adding Plots

Clicking on this icon brings up a dialog from which you can choose to make another scatter plot (left below), a heatmap (center below), or a histogram (right below):

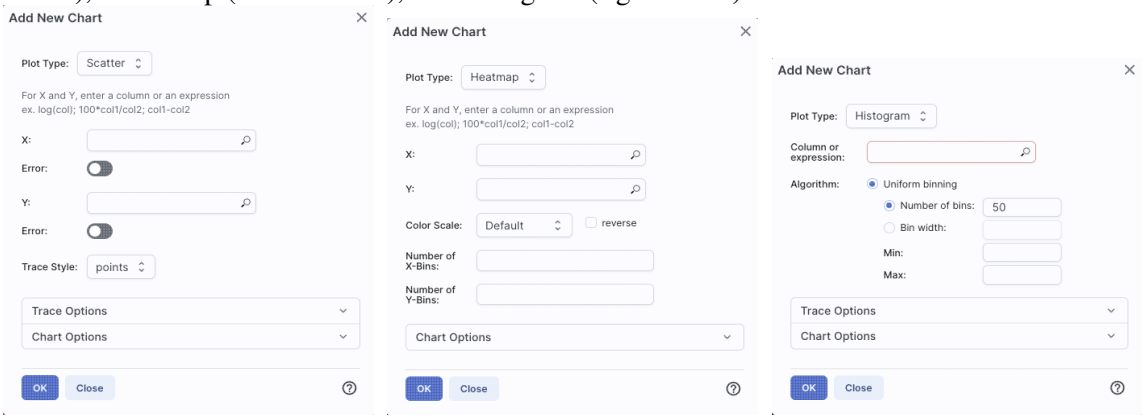

The options for these plots here are very similar to what is described above. You can specify which columns to plot or manipulate and plot, specify labels, etc.

Scatter plots allow you to choose points, connected points, or lines; you can add errors to each point.

Heatmap plots are binned scatter plots; you can choose what color scale and how many bins to use.

Histogram plots allow you to choose how many bins or the bin width. Note that, if you provide a minimum number, the binning starts at the minimum value you provide, and may exceed the maximum you entered in order to fit in a whole bin.

You can change what is plotted after plotting by clicking on the gears, as described above.

You can have many plots up at the same time.

You can view multiple plots all at once or one at a time by clicking on the corresponding icons above the plots (just as when you have multiple images loaded).  The single box means "one at a time", the set of four boxes means "all the plots at once". If you are viewing one at a time and have more than one plot loaded, you will see the ">" and "<" signs (as in the image here), and you can scroll among the plots by clicking on these arrows (just as when you have multiple images loaded).

Tips and Troubleshooting

- Note that many plots of a large catalog may make your browser run slowly.
- You can force the tool to make a scatter plot even if you have a catalog of >5,000 points; it just might make your browser slow down. You can even overplot a scatter plot on top of a heatmap if you really want to!
- To remove a plot, click on the 'x' in the upper right corner of the plot.

Pinning Plots

The idea behind "pinning plots" (or "pinning charts") is that you can retain a plot. Within this tool, "pinning" just means "hold on to this item within this tool." It doesn't mean "save this plot to disk", nor does it mean "download the data behind this"; it means "retain this item in this tool for now." Think of it as if you have a metaphorical bulletin board behind your computer monitor and you want to put a plot you make on that bulletin board temporarily (with a pushpin!) while you continue to work on other plots or other catalogs. For example, you could make the same color-magnitiude diagram from two catalogs of two different regions, pin both, and then compare them side-by-side.

When you make a plot that you want to retain, click on the "pin chart" icon  and it will make a copy of that chart and keep it in your plot area, even as you continue to make more plots and work with more catalogs.

It saves your pinned plot in a different tab than your current plot. To change between them, just click on the corresponding tab. Here, there are three pinned charts, but the current view is the active chart:

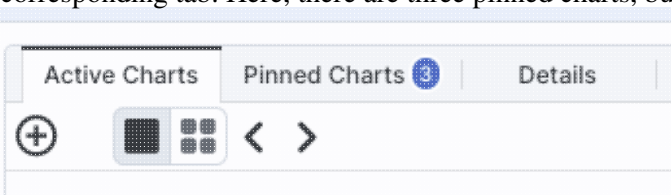

Tips and Troubleshooting

- When you pin a plot, and then filter its parent catalog, the plot updates correspondingly, because it is still linked to its parent catalog.
- When you have multiple catalogs (or tables of any sort) loaded, if you want to bring to the foreground the catalog tab corresponding to a pinned plot, click on "show table"  , and it will bring the corresponding catalog to the foreground of your tables pane.
- If you click on a source in the image (or catalog or plot), the same source is highlighted in the catalog (or image or plot).
- To change a pinned plot, click on the plot to bring it to the foreground, then click on the gears in the the corresponding plot area, and change what is plotted as described above.
- You can view the plots one a time or in a grid by clicking on the icon for one-at-a-time (big square) or tiled (4 small squares; if you are viewing them one at a time, you can use the ">" and "<" arrows to scroll through the list): 
- There is a maxium of 12 charts that you can pin at one time.
- When many plots are pinned, and you view many plots at once, the margins on the plots will shrink down (effectively vanish) to make the data more easily visible. To view the plots with axes, view the plots one at a time.
- To remove a pinned plot, click on the blue 'x' in the upper right corner of the plot.
- If you remove a catalog by clicking on the 'x' in the corresponding catalog tab, the pinned plot will be removed as well.

Combining Plots

 When you have more than one plot pinned, you have an additional icon that can appear -- it means "Combine Chart".

This option only appears if you have pinned at least two plots, and it will only let you combine plots if it recognizes that you have spectra loaded. These can be spectra you extracted or have loaded in from a file.

Because combining is only possible for spectra at this time, the information on how to combine plots can be found in the chapter on spectra.

Examples of Catalog Plots

Example: Plotting [W1] vs. [W1]-[W4] in a star-forming region

For this example, we are trying to find young stars in a star-forming region. We are searching in the WISE AllWISE catalog. Stars without circumstellar dust should be at a variety of W1 brightnesses, but all have [W1]-[W4]~0. Background galaxies should be faint and red. Stars with circumstellar dust (e.g., young stars) should be bright and red. Here, we will make a plot, identify a bright and red object in the plot, and find where it is in the WISE images.

1. Launch the IRSA Viewer tool. Click on "Images" at the top to search on images. For "Select Image Type", choose "View FITS Images", and for "Image Source", select "Search." Enter target as IC1396. Use Simbad's interpretation of the coordinates. Select image size of 1 degree (pick the units first, then enter the number). Select WISE AllWISE Atlas, all four bands. Click 'Search.'
2. Click on "Catalogs" at the top to search on catalogs. Select Project=WISE. Select the AllWISE database, and the AllWISE Source Catalog. Leave the cone search method and target as what it comes set up with as the default - it is set to the target we just used for our search. However, make sure to

change the units of the drop-down to "deg" and enter 0.5 degree in the cone radius to match the images. Leave all the defaults in the Table Selection section. Click "Search." You may have to wait a bit for the catalog to be returned, because there are ~17,000 sources. If you choose to put the catalog search in the background, load it when it is ready.

3. By default, the catalog is overplotted on the images in the upper left, plotted as an x-y plot on the upper right, and loaded as a table on the bottom. Note that the plot on the upper right is a greyscale plot, which means that the plot shows a binned "heatmap" where darker colors mean more points. Only the first 100 are shown in the catalog on the bottom, though all are loaded, and you can page through the list, 100 at a time. Note that selecting a source in the image makes its corresponding row in the catalog change color. It doesn't highlight in the plot, because there are too many points in the plot.
4. Filter down the catalog to only have detections. By default, you should have a place to put in filters at the top of each column, but if you don't, go to the the upper right of the catalog pane, and click on the funnel icon  to add filters to the top of each column in the catalog. In this AllWISE catalog, limits have null errors. To limit the catalog to just those with detections, in the top of the "w1sigmpro" (WISE-1 profile fittted magnitude error) column, type ">0" (without the quotes). Repeat for w2sigmpro, w3sigmpro, and w4sigmpro. After this process, you should have ~3,000 sources, enough fewer that the plot now has individual blue points (it is no longer a heatmap). Now the individual sources are also **all** shown on the images.
5. The plot comes up with an RA/Dec plot by default. Click on the gears  icon in the upper left of the plot window to change what is plotted.
6. Ensure "Modify Trace" is selected in the pop-up window (by default, it should be selected). Enter in the x box: w1mpro-w4mpro. This is WISE-1 profile-fitted measurement in magnitudes minus WISE-4 profile-fitted measurement in magnitudes,or [W1]-[W4].
7. Enter in the y box: w1mpro. This is the [W1] mag.
8. Click on "Chart Options", and enter "[W1]-[W4] (mag)" for the x-label, and "[W1] (mag)" for the y-label. Click on "reverse" for the y-axis to make the brighter objects appear at the top of the plot. Check "grid" to make the grid appear for both axes.
9. Click "Apply."
10. Obtain this plot (configuration also provided for reference):

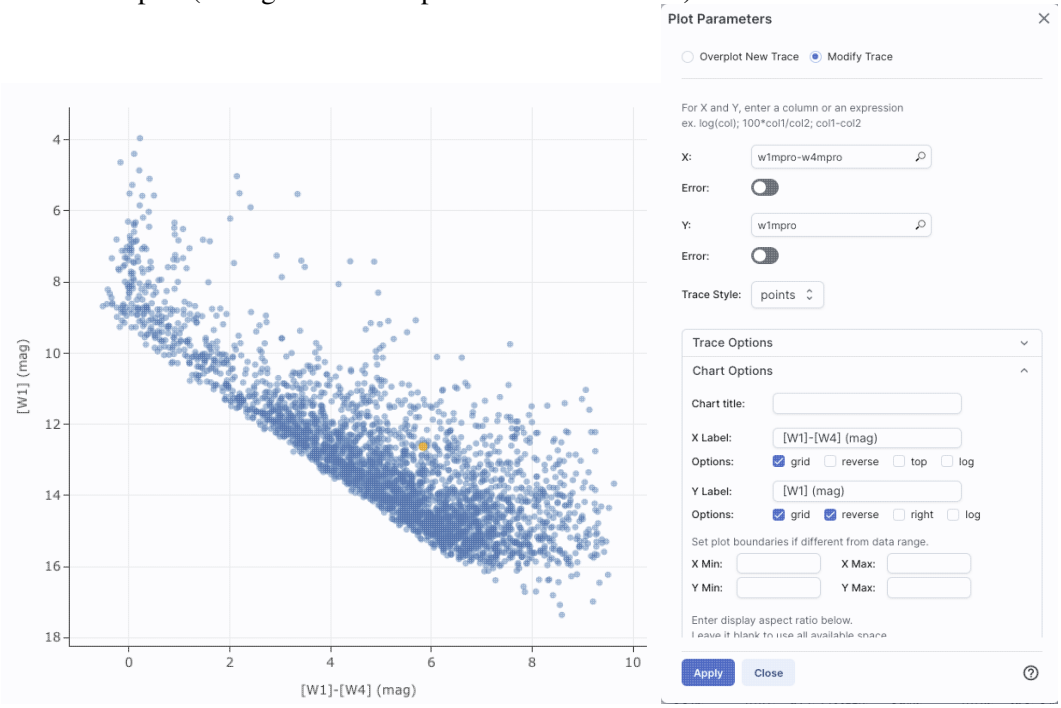

11. Ensure that the select icon is "active" in the plot ( ). Click and drag from corner given approximately by (2,6) to (7,11).
12. The icons in the upper right change after you do this, and we want to filter on this region. Click on the filter icon. This filters the table and the image overlays as well as the plot.
13. After this filtering, you should have about 100 points left in the catalog. The scatter plot after the filtering looks something like this:

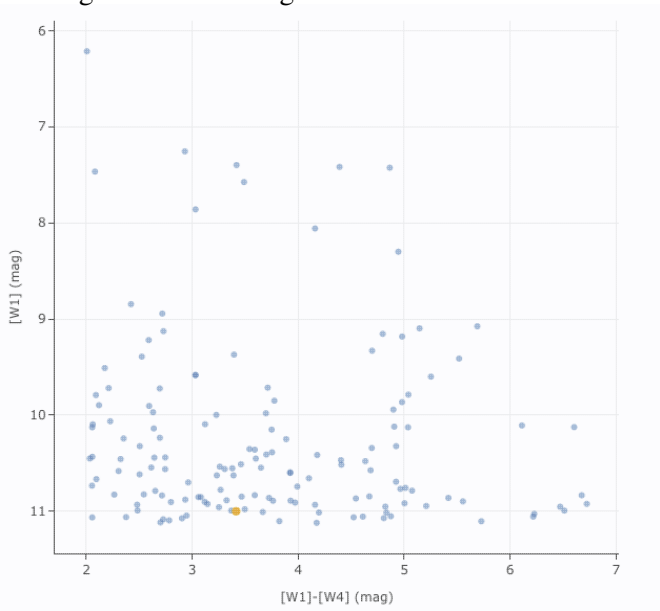

The points are shown as blue circles. Where they overlap, the blue appears darker.

14. We want to find an (astrophysically) bright and red object in the plot, catalog, and images. In order to find it easily in the images, first, zoom in relatively closely to any of the images. Select the lock drop-down in the image toolbar ( ) and select "Align and lock by WCS." Ensure "pan by table row" is selected from the recenter icon ( ).
15. Now, find the brightest source near [W1]-[W4]~4.5 in the plot. Click on that point. It is highlighted in the overlays on top of the image, the images have jumped to center it, and it is highlighted in the catalog (though you may need to scroll slightly in the catalog to find the highlighted line). (If you have been following closely, it should be J213808.44+572647.6.)

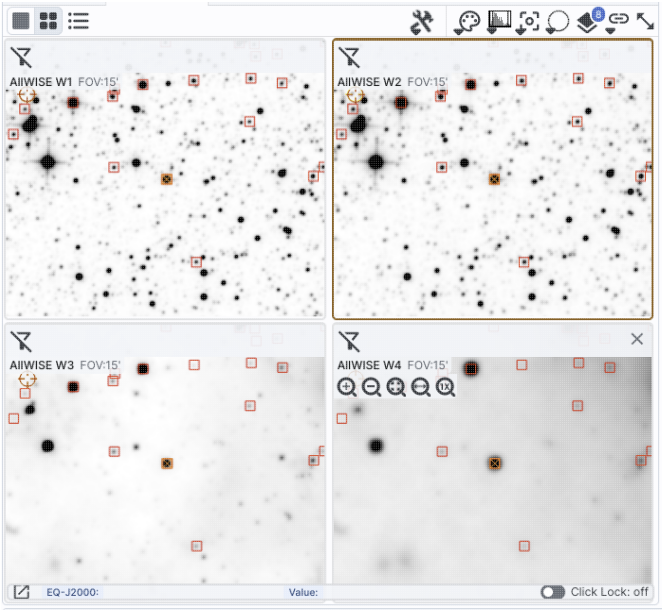

Abbreviated Example: Plotting [W1] vs. [W3]-[W4] with errors

Because you can do simple mathematical manipulations when specifying what to plot or use for errors, you can calculate these things on the fly. From the end of the prior example, cancel the filters, and impose new filters on the catalog by requiring w*snr>10 for all four bands. Intitiate a new plot by clicking on the plot's gears icon, and select "Add New Chart."

For the x-axis, ask it to plot w3mpro-w4mpro. Ask it to use symmetric errors and use sqrt(power(w3sigmpro,2) + power(w4sigmpro,2)). For the y-axis, ask it to plot w1mpro, and use w1sigmpro for the errors. Under "Chart Options", add the x label [W3]-[W4] and the y label [W1]. Reverse the y-axis to get bright objects at the top. Under "Trace Options", select red points just to be fancy. Obtain a plot something like this, which shows the calculated error bars for each point. Note that there are errors in both directions, though they are hard to see in the y direction.

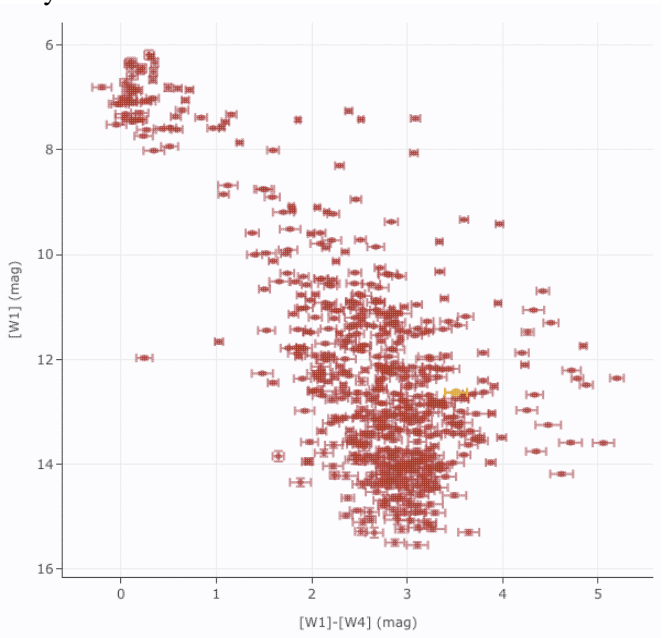

Abbreviated Example: Plotting a histogram of [W1]

Picking up from the end of the prior example. Click on the blue 'charts' tab near the top of the window. Select "histogram." Enter w1mpro for the column. If you leave everything else to the defaults, you can obtain a plot like the below. As for xy plots, mousing over a column tells you more information about the contents of the column.

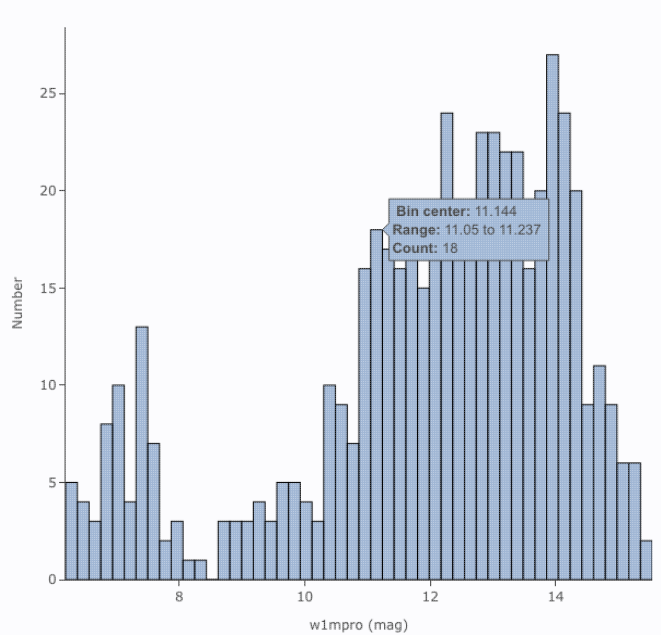

Still More Plots

Here are several more examples of plots made with IRSA tools.

Phase-folded light curve from K2 data:

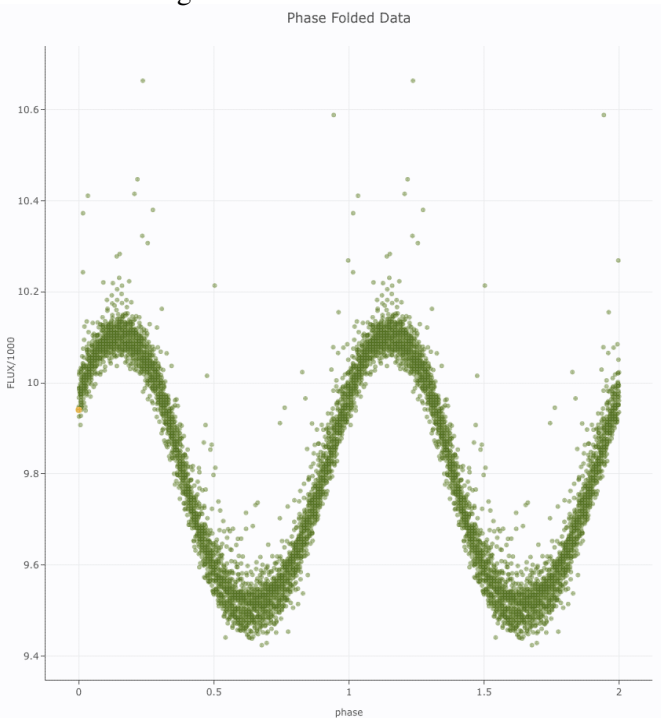

Plot on the sky of stars where the color of the point is scaled to brightness in WISE-4:

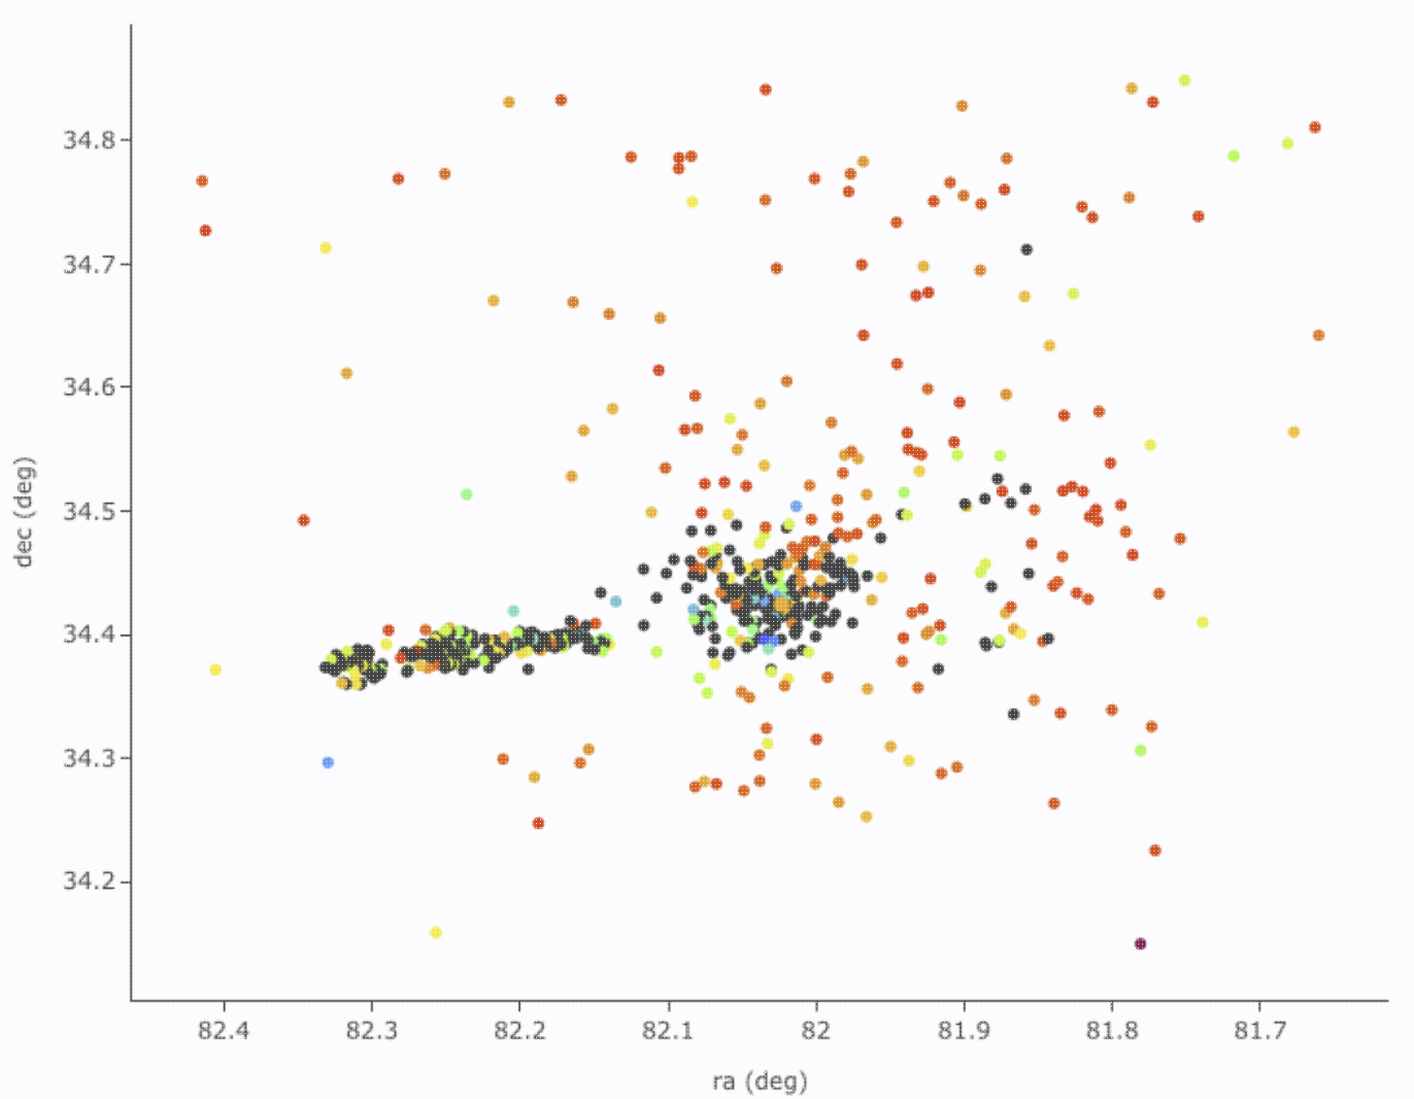

Gaia distance (in kpc, from Bailer-Jones et al. 2018), with asymmetric errors, as a function of Gaia G magnitude, with colors of the point scaled to brightness in WISE-4:

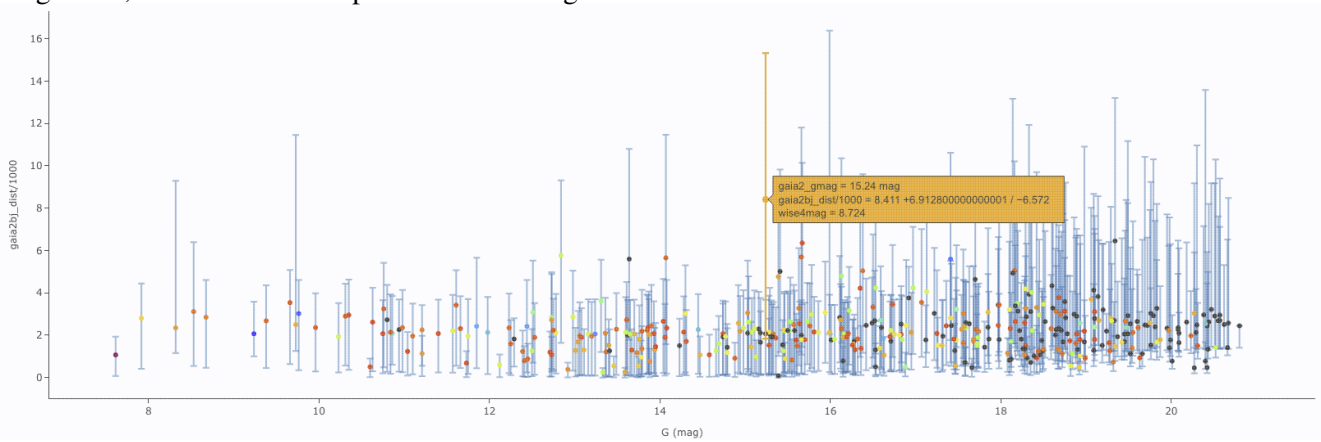

[W1] light curve of Neptune over several years, with colors of the point scaled to heliocentric distance:

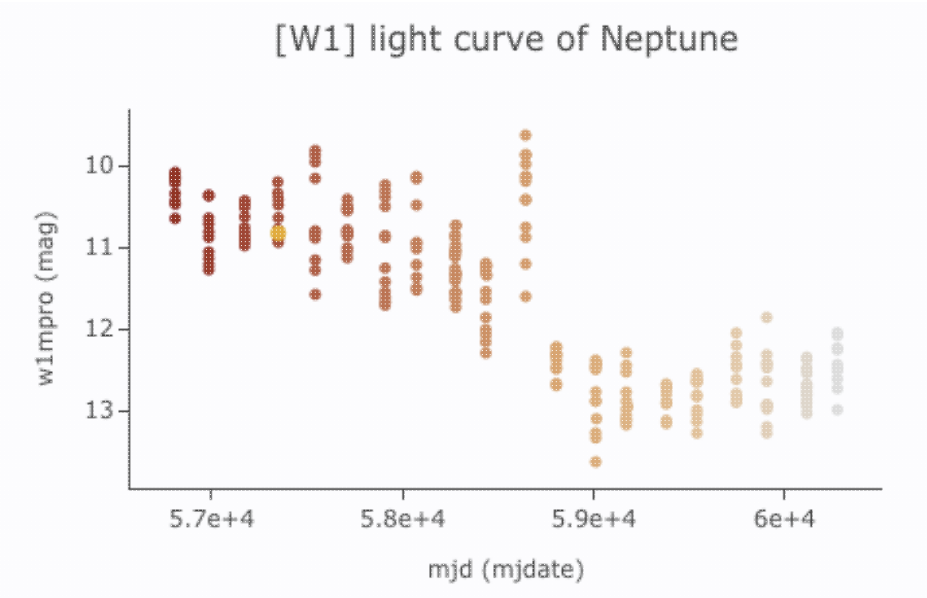

Absolute Gaia color-magnitude diagram of candidate members of a star-forming region (note some background giants still in the list), where point size is scaled by WISE-4 brightness:

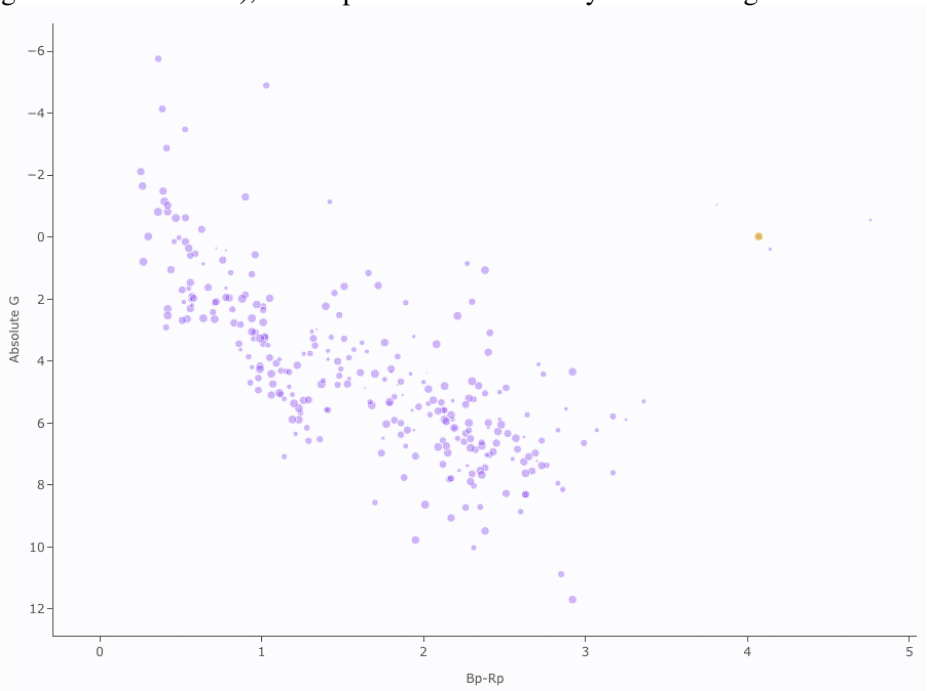

IRSA Viewer: Spectra

Visualization of spectra use capabilities of Tables, (image) Visualization, and Plots. Generic help on those capabilities can be found in those other sections; spectra are a special case of all of those, and this section attempts to build on that by collecting all the spectra-specific information in one place.

Contents of page/chapter:

Loading Spectra

The first step to using the spectral features in this tool is getting the tool to recognize that you have a spectrum.

When you upload your own files from disk, if you load a table, there is an option that appears at the bottom left of the screen:

If you tick this box, and your data are formatted consistently with (or at least close to) the IVOA SpectrumDM v1.1 data model, then the tool will recognize the data as a spectrum.

There is some flexibility in the interpreter, so if you think your file is close to the standard, give it a try. The computer might be able to figure it out.

Extracting Spectra

Alternatively, you can get spectra into the tool by extracting them from an image file. If you have a multi-HDU file where the planes are different wavelengths (such as can be found in data from SOFIA), use the drill extraction tool, and "pin" the extraction, then the tool will also recognize the extracted data as a spectrum. It will also (attempt to) propagate the errors correctly.

Plotting Spectra

Spectra are plotted automatically when loaded. By default, they are plotted as F-sub-nu vs. lambda in microns, with connected points. At least, this is what it attempts to do, if it understands the units of the file you have given it. If it understands the units, it will convert them appropriately. If it doesn't understand the units, it may serve them to you in the same units it received.

When spectra are plotted, changing what is plotted by clicking on the gears is similar to, but not quite the same as, the generic case. Now, because it knows it is plotting a spectrum, you can select the x- and y-axis columns and units from a pre-defined set of choices in the drop-down menus, where it will convert the units when necessary.

It is probably most efficient to demonstrate plotting of spectra via specific examples.

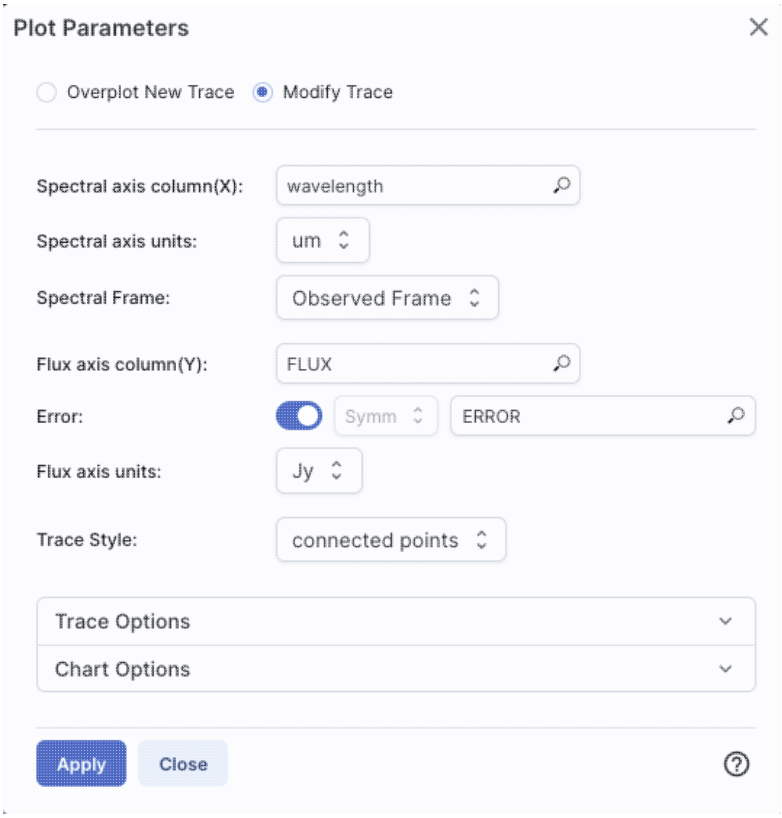

In this example, the tool has identified the spectral axis a 'wavelength', understood the units, converted them to mi is showing them in the observed reference frame. It has i the flux axis column as 'FLUX' and the corresponding er 'ERROR', and understood the units as Jy. From the drop- menus, you can choose to convert the wavelength to Ang nanometers, microns, millimeters, centimeters, or meters choose to convert the flux density to Janskys or Watts/meter^2/Hertz. It is plotting the spectrum as conne points, with error bars.

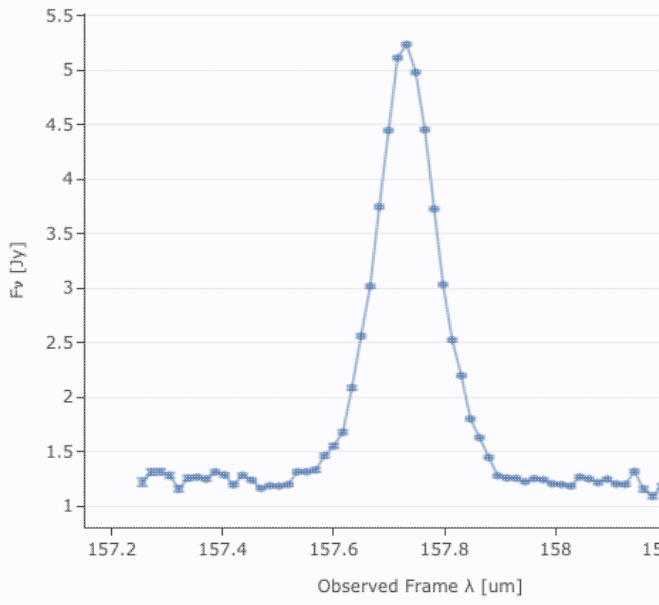

In this second example, the tool is struggling a bit more. It identified the spectral axis column as 'wavelength' and is sh it in the observed frame, but it is guessing about the units; t only appear in the plot. It has identified the flux axis colum "flux_density", it has identified the errors as "error", and it understood those units as Jy. It is showing the spectrum wit connected points. Note too that in this case, there are severa spectral orders; at the top, you can see that there is "choose -- it turns out that there are four orders in this spectrum, and overplotting all four on the same plot, with errors.

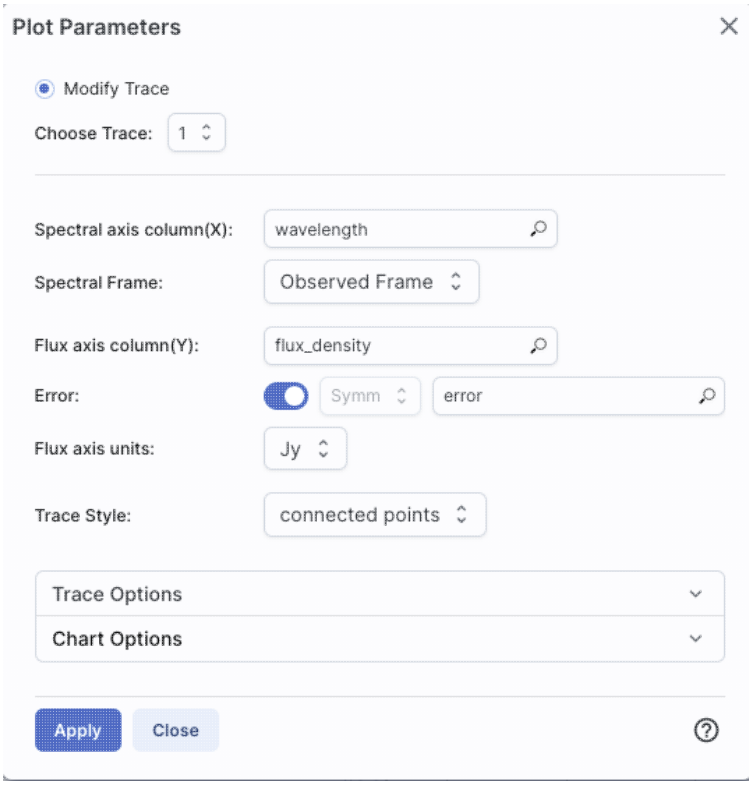

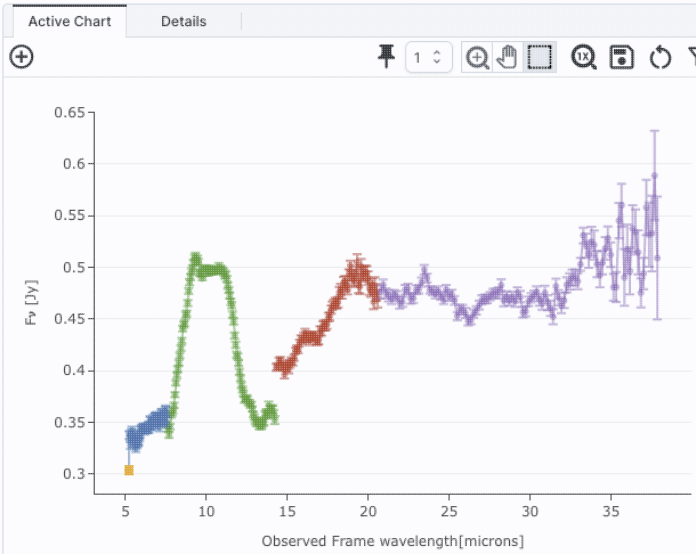

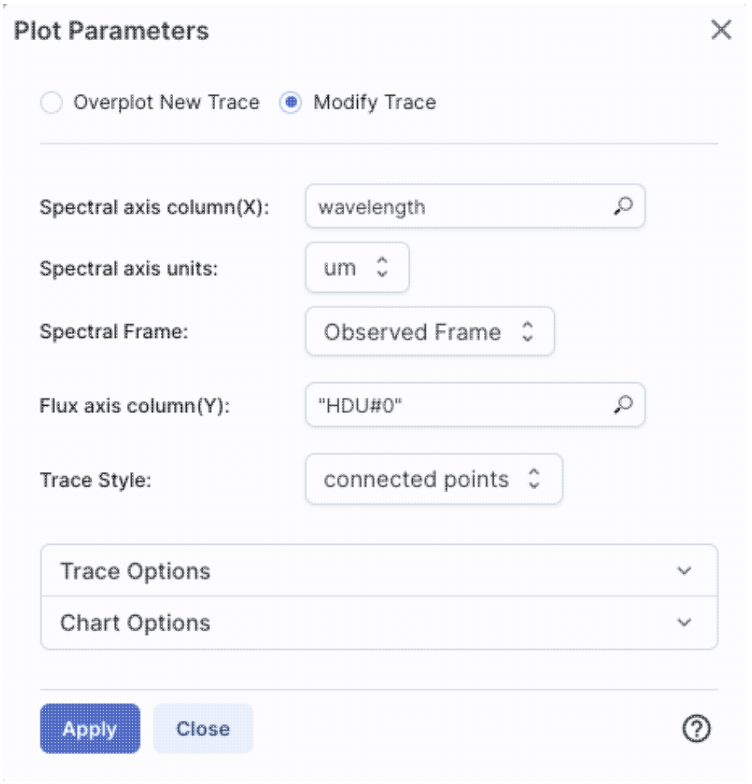

In this less well-behaved example, the tool has identified the spectral axis column as 'wavelength' and has figured out the of it, shown in in microns, in the observed frame. The flux a column in this case is just "HDU#0", so it is completely con It hasn't figured out what to do about errors at all. It is show spectrum with connected points.

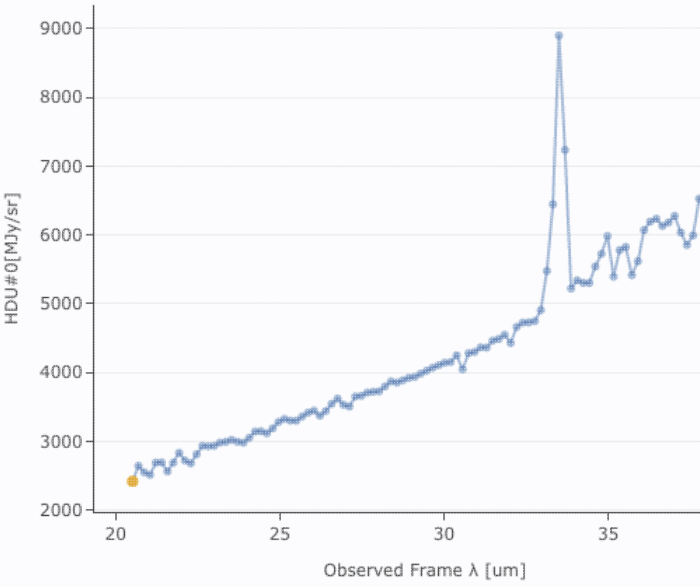

In all of these cases, you don't have as much flexibility in these plots as you do for plots in general, but the options you do have are highly customized to spectra, such as redshifts. See the next section!

Redshifting Spectra

When the tool recognizes the wavelength axis, it offers you a choice of redshifting the spectrum. By default, it assumes you want to plot the data as observed:

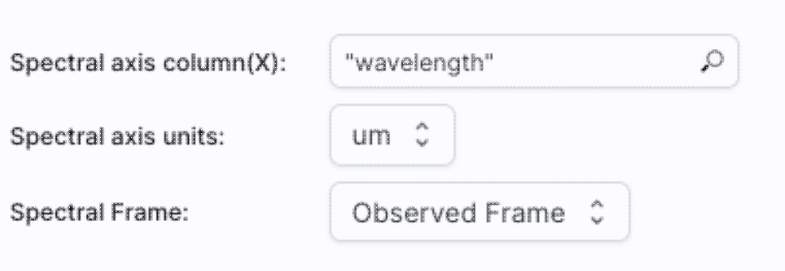

but if you pick "Rest Frame":

you can enter a redshift, and it shows you how it adjusts the wavelength axis accordingly:

Click 'Apply' to implement these changes in the plot. The axis labels on the plot correspondingly change.

To change back to the data as observed, simply pick "Observed Frame" from the drop-down menu.

Overplotting Spectra

The Plots section covers overplotting in general, and the example given there is a color-magnitude diagram. However, the reason why there is overplotting capability in the tool at all is to support overplotting spectral orders from a given spectrum on the same plot.

If you load in a spectrum with multiple orders and just try to plot it without convincing the tool that the data are a spectrum, you get, by default, the plot on the left. But, if you let the tool know that the file you have uploaded is, in fact, a spectrum, you get the plot on the right, where the four orders in this spectrum are overplotted in different colors.

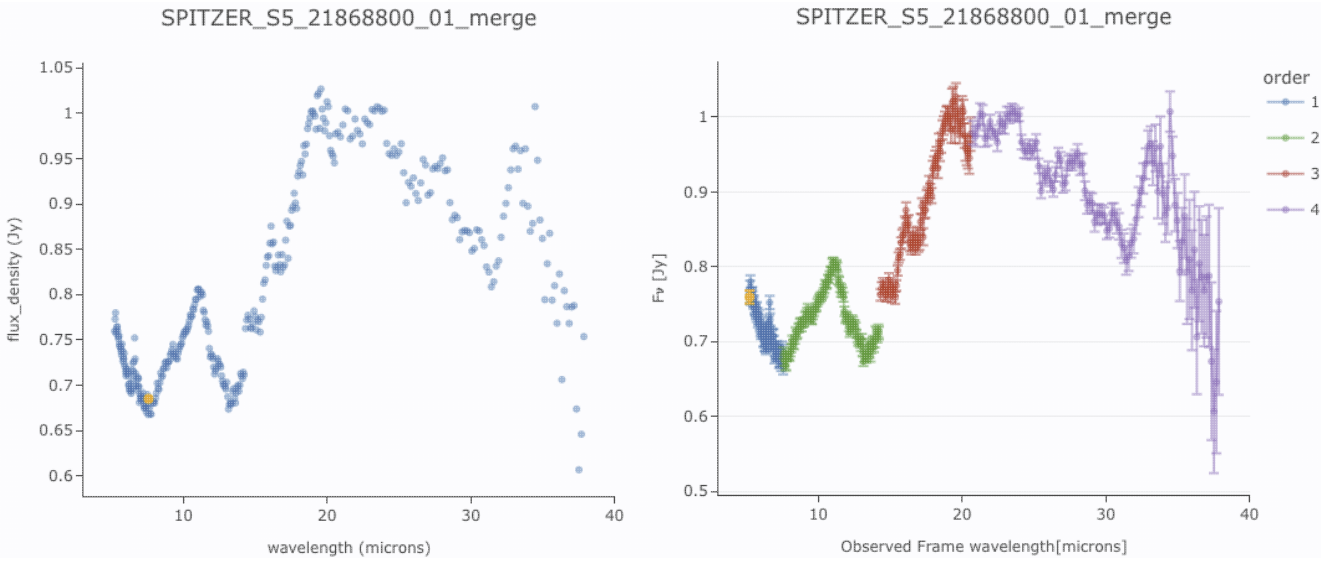

If you use the plot options (click on the gears icon), you can change the plot parameters. It has options similar to what you get for generic plotting, but limited as in the above, specifically for plotting spectra:

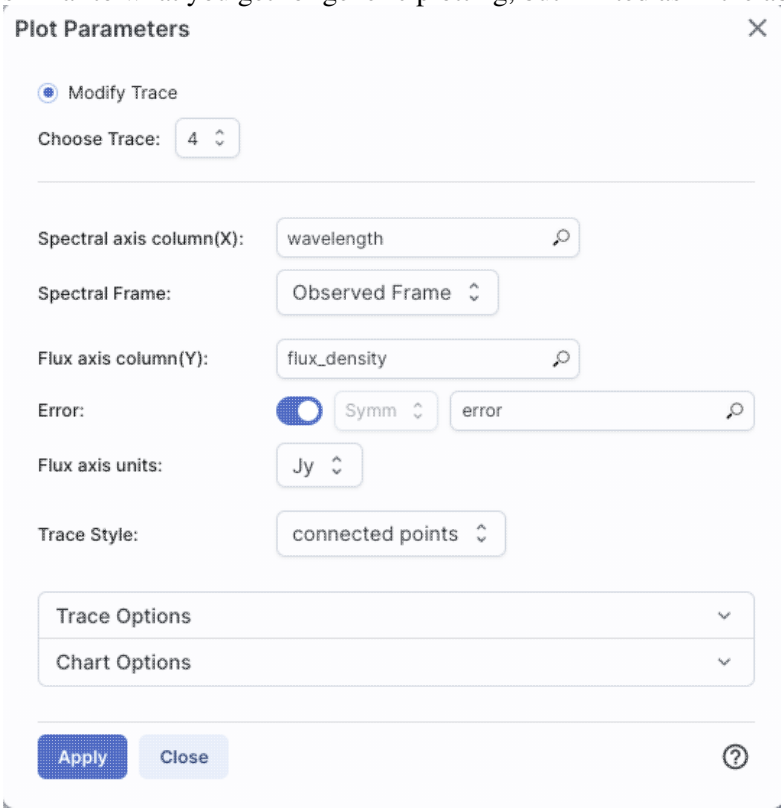

Combining Plots

Pinning plots is covered in the generic plots section, but in brief, pinning allows you to temporarily 'save' a given plot.

When you have more than one plot pinned, this icon may appear at the top of the plot pane: This means "Combine Chart".

This option only appears if you have pinned at least two plots, and it will only let you combine plots if it recognizes that you have spectra loaded.

To start this process, click to select the first chart you want to combine, then click on the "combine chart" icon. You get a pop-up like this:

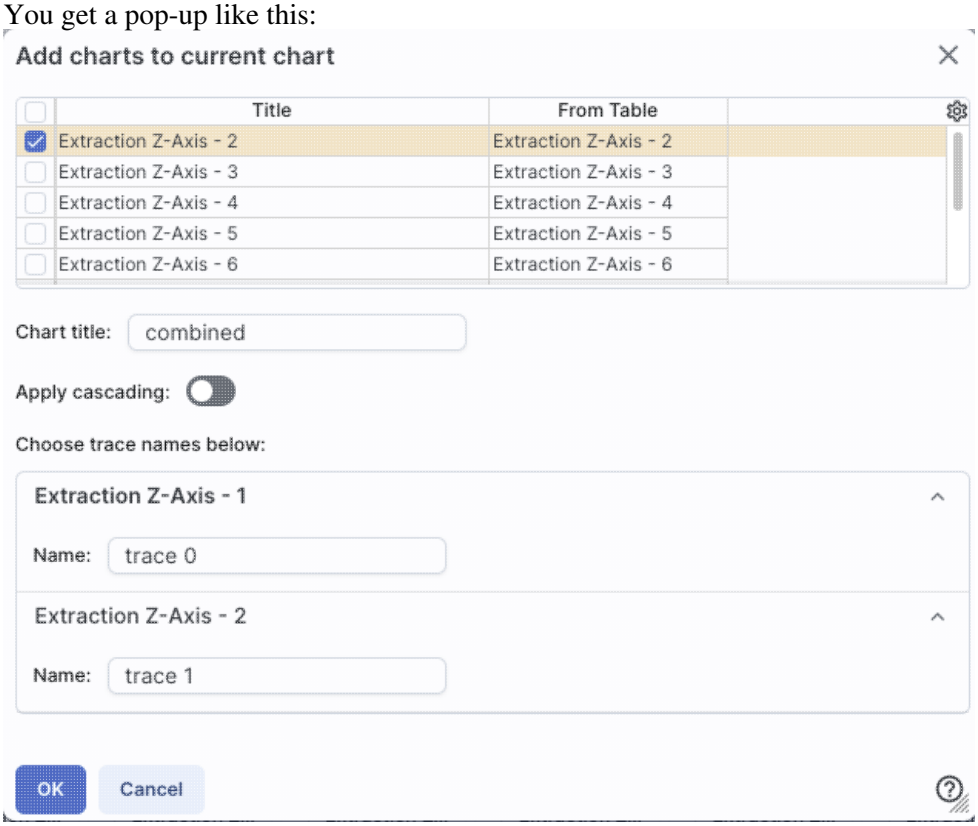

All of the remaining pinned charts that can be combined appear as a list at the top. Once you select them via the tickboxes on the far left of the list (the first one in the list is selected here), they appear as options on the bottom of the pop-up window. For this example, I extracted the observed spectra from a SOFIA FIFI-LS data cube at several sky positions. The extractions are the default "Extraction Z-Axis - *n*".

Continuing through this pop-up, you can choose to set the title of the new plot you are about to create -- the default is "combined".

The next choice is "Apply cascading." Here is what this is and why it matters. If you are combining spectra that are nearly all the same brightness, the spectra will be plotted on top of each other. Sometimes that is what you want, and sometimes that is not. If you click on the "Apply cascading" option, you have an additional choice:

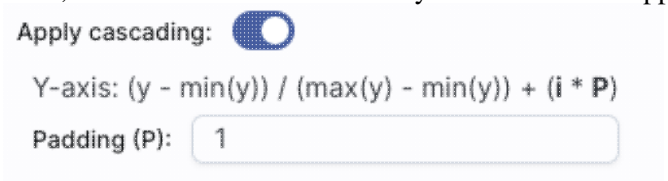

This is telling you how it is going to stack the spectra on the final plot. See below for examples with and without cascading. You can adjust the amplitude of the cascade by changing the size of the padding, as shown.

Finally, you can change the name of the trace as displayed on the plot (and in the pull-down menus in the tool) for each of the spectra you are combining.

Click "OK" to actually make the new plot.

Here are two examples of combined spectra that were extracted from a SOFIA FIFI-LS data cube, one without cascading, and one with cascading. Both are useful, but in different contexts.

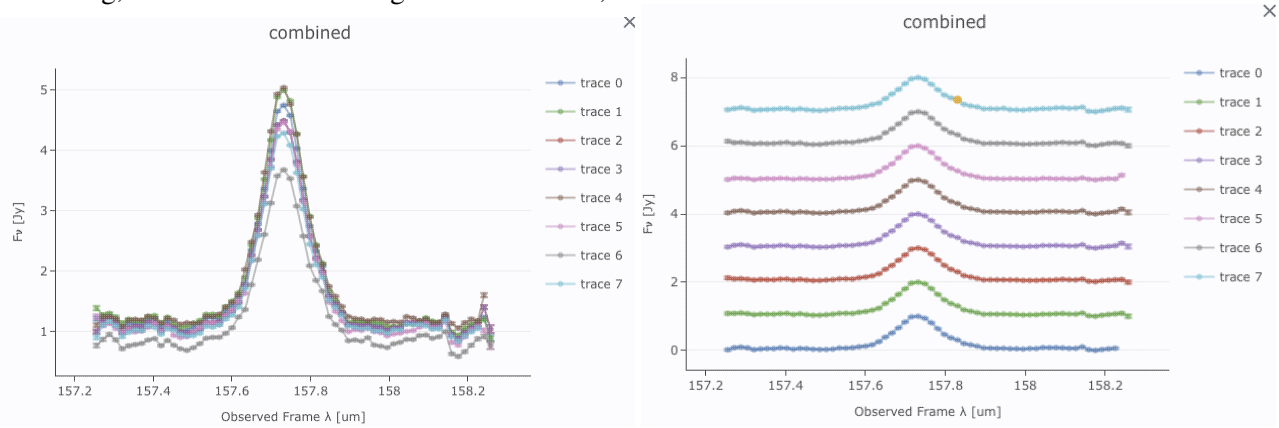

Note that after you combine a plot, there is a new drop-down at the top of the plot that controls which trace is in the 'foreground' for changing plot parameters or selecting points, but you can also simply click on points in the plot to bring that trace to the foreground.

Tips and Troubleshooting

- The plot you have selected when you click on "Combine Chart" is implicitly part of the combining process, so it's not available to select with a tickbox in the pop-up.
- If you have loaded spectral files with multiple orders, each order appears as a separate color in the plot. You can change the order labels before or after combining.
- At this time, you can't combine charts that the tool doesn't recognize as spectra, even if they are in the same parameter space.
- To delete a combined plot that you have created, just click on the 'x' in the upper right of the plot.

IRSA Viewer: Download

Contents of page/chapter:

Introduction

You can download images and tables and plots from IRSA Viewer. The Background Monitor helps keep track of downloads placed in the background.

Images

If you want to save an image, click on the image you want to save so that it has the brown highlighted box around it. The diskette icon () in the image toolbar is within the tools drop-down (), and that will allow you to save that currently selected image as a FITS file (or other formats, including saving just the overlays). Note that **you** control where the file is saved on your disk through your browser; your browser may be configured to store all downloads in a particular location on your disk.

Tables

On the upper right of any table shown in the tool, there is a diskette option:  . Click on this icon to save the table; you have a variety of choices of file format.

Plots

To save a plot, click on the diskette icon (). The plot will be saved to your disk as a png file.

Prepare Download

Some windows will have a button that reads "Prepare Download." If that button appears, then you just have to select the rows corresponding to the data (or data product) you wish to download -- tick the tickbox on the far left of the corresponding row, or tick the box at the top of the column to "select all" -- and then click "Prepare Download." It may ask you for additional information, like how much ancillary information you want, and what the filename you want for the zip files. If you ask for a lot of data, the packaging will likely need to spin off to the Background Monitor.

File Location

The files that you are saving can go to your own disk, or to the IRSA Workspace.

Tips and Troubleshooting

- **You** control where the file is saved on your disk through your browser; your browser may be configured to store all downloads in a particular location on your disk. Try looking in a "Downloads" directory, or for "recently modified files."
- You need to be logged in to use the IRSA Workspace.

Background Monitor

The Background Monitor appears as blue text 

in the upper right of your window to keep track of the data downloads you have requested. It actively changes to reflect what it is doing (e.g., this is preparing a download or catalog: 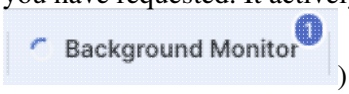

) Larger catalog requests can be sent to the Background Monitor.

A pop-up window can be called up at any time by clicking the "Background Monitor" tab. You can watch your data being retrieved. It will update that window when the data are available for download and/or overlay on your image, providing a link (or links) for obtaining the data. Remove them from the list by clicking on the 'x'.

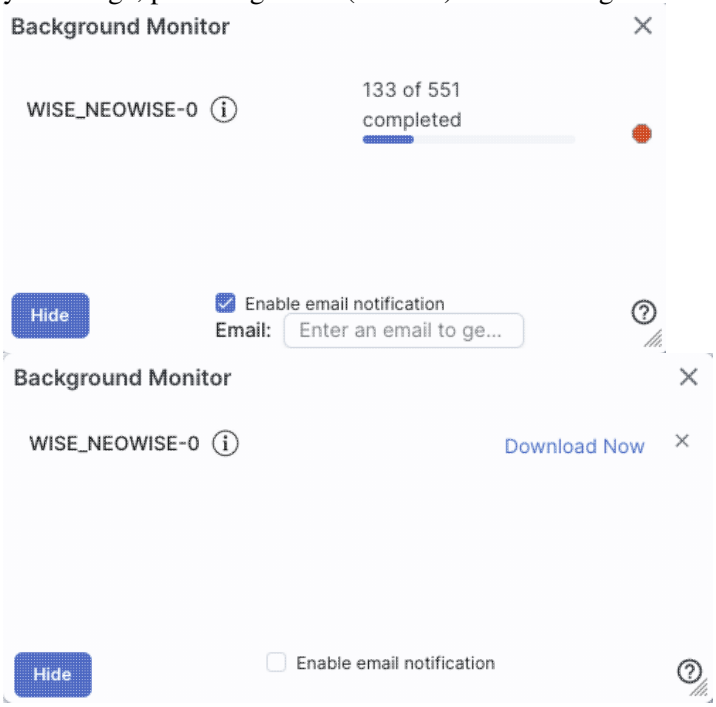

If you have made, say, a large catalog request and don't want the pending catalog request to occupy screen space while it loads, you can click on the button marked "send to background" to reclaim your screen space, seen here:

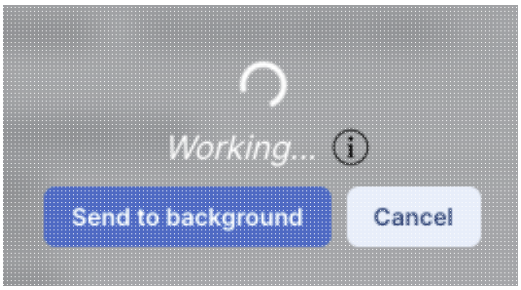

When the background monitor finishes, however, you will have to actively tell it to display results; it doesn't do it automatically if you have sent it to the background. You can ask it to send you email when it is ready by clicking on the "enable email notification" and giving it an email address.
To stop any query mid-way through, click on the little red octagon ("stop sign") that appears next to the query in the Background Monitor pop-up.

If you forgot to put in your email at the beginning, or if the packaging is taking longer than you expected, you can click on "Enable email notification" from the bottom of the Background Monitor popup and add your email during the packaging process.

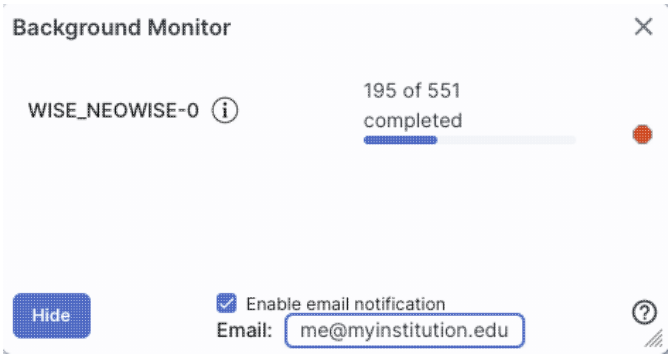

If you are having technical difficulties, click on the circle with an "i" in it to get additional information, like this:

In this case, the job is executing, at the time as given, with a job id as shown.

Acknowledgments

Note that you can use the \facilities call in AASTeX to acknowledge IRSA as a facility you used in your journal article.

The standard IRSA acknowledgment is:

> This research has made use of the NASA/ IPAC Infrared Science Archive, which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

Please consult the documentation that accompanies the data that you use for whatever corresponding acknowledgement and/or citation they want you to use. You should use the DOI corresponding to IRSA's copy of the data, if the data are available in more than one place.

Time Series Tool: Overview

Contents of page/chapter:

Purpose of Tool

The Time Series Tool allows exploration and analysis of time series observations.

For WISE/NEOWISE and PTF/ZTF, you can view measurements as a function of time, simultaneously visualize the single-epoch images, and optionally find the period of variability. For other data sets, visualization of the single-epoch images is not currently available, but period analysis is available.

Terminology

The words in blue rectangles at the top of the screen are "**tabs**" : 

Your browser window is divided into "**panes**", like "window panes." The contents of the panes depends on what you are doing with the tool, but could include a pane that includes a set of images, a pane that has the catalog used for the time series, and/or a plot pane. You can expand any of the window panes by clicking on the expand icon: 

Each of the 'panes' in the display has its own toolbox in its upper right corner which operates on things in that pane, and the basic functionality for each of these panes is covered elsewhere in this document:

- The Tables section covers (among other things) the tables toolbox 
- The Plots section covers (among other things) the plots toolbox 
- The Visualization section covers (among other things) the images toolbox 

After loading a light curve, you are in the "**Viewer**." From the Viewer, you can launch the "**Period Finder**" to set a period. After you set a period, you return to the "**Viewer**."

Data Sets Accepted

For WISE/NEOWISE and PTF/ZTF, you can view measurements as a function of time, simultaneously visualize the single-epoch images, and optionally find the period of variability.

Note that WISE/NEOWISE light curves obtained from IRSA have two times of interest:

- hjd: Heliocentric Julian Date of the midpoint of the exposure (using the mean ra and dec of the input catalog)
- mjd: Modified Julian Date of the start of the exposure

For other data sets, visualization of the single-epoch images is not currently available, but period analysis is available.

Documentation Contents

This documentation covers things in the order in which you are likely to encounter them in the Time Series Tool:

Basic information such as how to interact with images in general, tables in general, and plots in general is covered elsewhere in this document.

Time Series Tool: Loading in a Time Series

When you launch the Time Series Tool, the first screen you could encounter is the screen in which you can upload a time series.

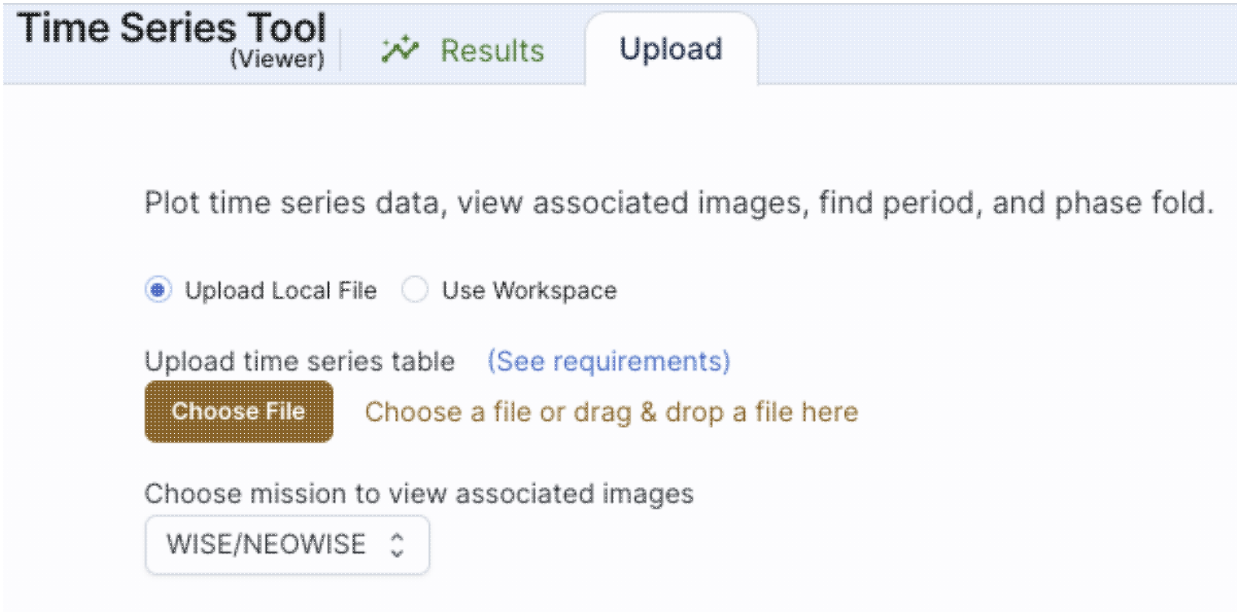

Contents of page/chapter:

Main Options

You can use a file on your disk or in the IRSA Workspace. In either case, click on the "Choose File" button to find the file on your disk or in the workspace to upload into the tool, or click-and-drag a file from your computer onto this space.

For "mission", you can choose "WISE/NEOWISE", "PTF", "ZTF", or "Other".

Data Sets Accepted

For WISE/NEOWISE and PTF/ZTF, you can view measurements as a function of time, simultaneously visualize the single-epoch images, and optionally find the period of variability.

For other data sets, visualization of the single-epoch images is not currently available, but period analysis is available.

Table Requirements

WISE/NEOWISE

To load the single-epoch images, the WISE/NEOWISE table you upload must have the columns frame_id and source_id, or both scan_id and frame_num. (It needs those columns to be able to grab the corresponding exposure image.) Catalogs obtained from IRSA's Catalog Search Tool for the AllWISE multiepoch photometry catalog will work transparently; catalogs from Catalog Search Tool for other single exposures **need the long-form results**, which include these columns, and then these catalogs will also work transparently. (Here is an example input table from AllWISE.)

PTF and ZTF

To load the single-epoch images, the PTF or ZTF table you upload must have the column "pid" included in your table. Catalogs obtained from IRSA's Catalog Search Tool for the PTF or ZTF light curve catalog will work transparently. (Here is an example input table from PTF.)

Table Formatting

If you upload your own catalog, your catalog needs to be in IPAC table format, which is a varietal of plain text. IRSA has a table validator which may be helpful, or you can download just about any catalog you find through IRSA, and mimic that format.

Default Columns

If you have a data table with an independent and dependent variable (time and flux, time and temperature, position and color, whatever), you should choose "Other" and you will still be able to work with your data; it just won't be able to pull individual images for each data point. If it finds a column called MJD, mjd, BJD, bjd, or really anything followed by "jd", it will assume that is your time column. It doesn't know how to assume your dependent variable column. If there is only one other column, it grabs that as the dependent variable. If you have many columns, you may need to help it realize what column should be the dependent variable.

New Tables

You can retrieve this upload window at any time by clicking on the blue "upload" button in the upper left. From here you can upload a new table and start again.

Cancelling

Click "cancel" to cancel the upload.

Go on to Main Screen, Viewer (Initial View).

Time Series Tool: Main Screen, Viewer (Initial View)

After you load in a time series, you get a screen where you are presented with several different panes: settings, table, plot, and images, highlighted in the screen shot below. The table, plot, and images are all connected in that clicking in one changes things in the others.

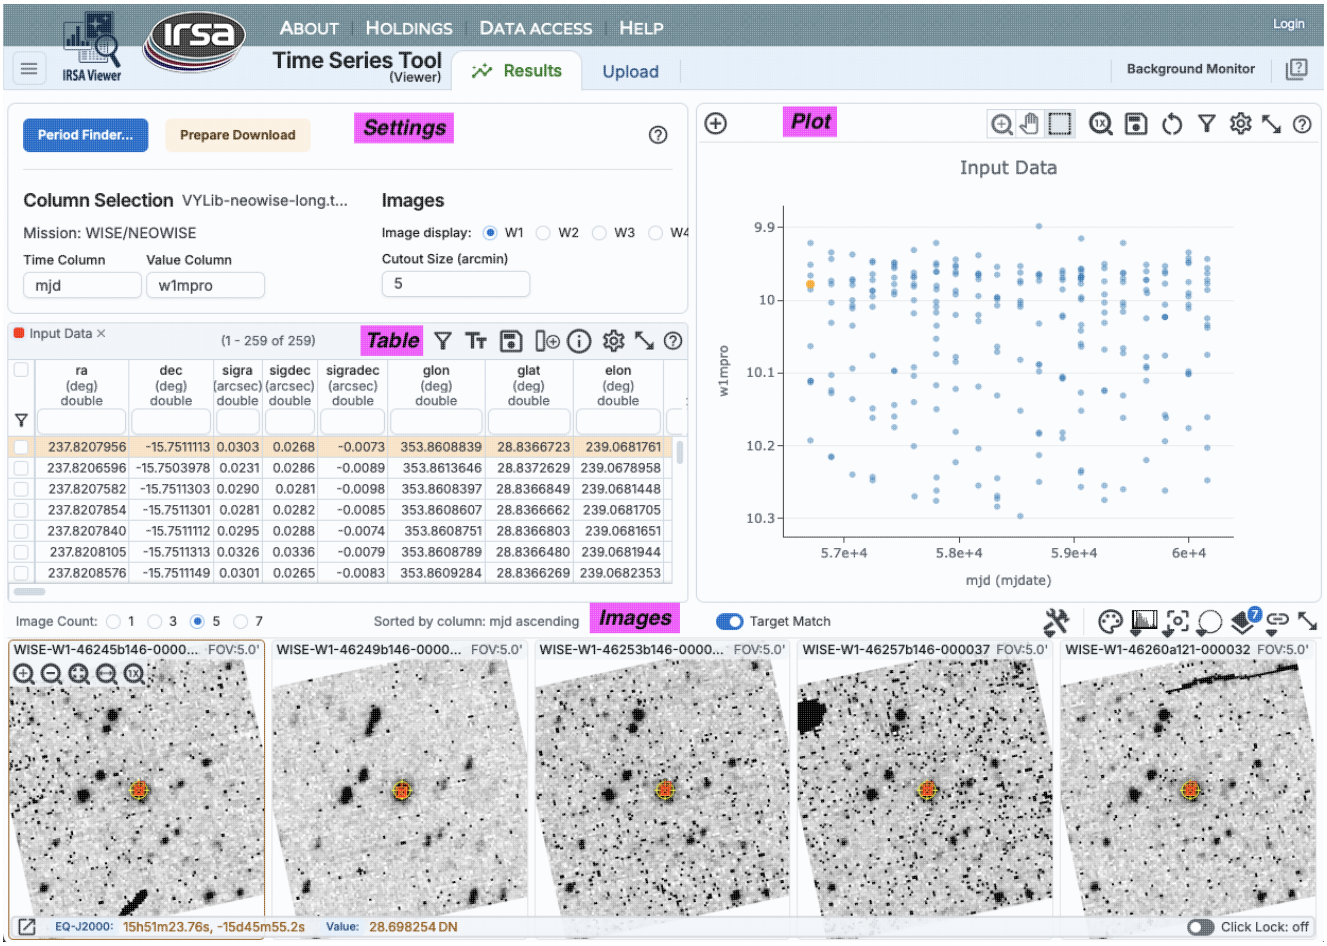

Contents of page/chapter:

Settings

This box in the upper left controls the time series with which you are working. For WISE/NEOWISE, you can view measurements as a function of time, simultaneously visualize the single-epoch images, and optionally find the period of variability. For other data sets, visualization of the single-epoch images is not currently available, but period analysis is available.

The upper left starts with the file you have uploaded, followed by the mission you have selected. Then it has a place to enter the time and time-dependent column.

For WISE/NEOWISE time series, it assumes that the time column you want is "mjd", and the time-dependent variable column is "w1mpro_ep". You can change this to any of the other 3 WISE channels (w2mpro_ep, w3mpro_ep, w4mpro_ep, assuming there are data for your selected channel). You can change both 'Time Column' and 'Value Column' to be any pther column in the uploaded catalog.

For PTF time series, it assumes that the time column you want is "obsmjd", and the time-dependent variable column is "mag_autocorr". You can change both 'Time Column' and 'Value Column' to be any pther column in the uploaded catalog.

If you have uploaded your own time series, you can control which columns are plotted (and eventually sent to the period-finding function) from this settings pane. If it finds a column called MJD, mjd, BJD, bjd, or really anything followed by "jd", it will assume that is your time column. It doesn't know how to assume your dependent variable column. If there is only one other column, it grabs that as the dependent variable. If you have many columns, you may need to help it realize what column should be the dependent variable. (If you're not careful, you may end up doing a periodic analysis of ra [since it is likely the first numeric column in your file] as a function of time, which could be interesting but is likely not your primary science goal.)

The top left of this box has a button, "Period Finder", which launches period finding; see Period Finding. The top left may also have a button, "Prepare Download", which gives you an option to download the lightcurve and images from WISE or ZTF/PTF, if applicable.

The upper right of this box controls the images displayed, if the images are displayed at all - which band, and what size of images.

By default, for WISE/NEOWISE time series, it displayes the W1 images at the bottom of the window (see images, below). To have different WISE images displayed (for any column in the uploaded table, not just another WISE magnitude), click on the radio buttons.

The "cutout size" controls the size of these image cutouts.

Table

Just below the settings pane is a pane with the data table displayed. This is the time series you have uploaded into the tool.

The table, if very long, is divided into 'pages' you can scroll through. You can filter and sort the table. You can control which columns are shown. For a more complete introduction to this tool's table manipulation functions, please see the Tables section.

Quick start: Sorting. Click on any column heading to sort the table by that column in ascending order; click it again to make it descending order, or click it a 3rd time to return the table to the unsorted state.

The table comes up, by default, sorted by the time column you have selected in the settings pane.

Quick start: Filtering. Depending on what you have loaded into the tool, you may have, for example, a table from ZTF that mixes g, r, and i photometry. The boxes near the top of each column are filters. Find the "filtercode" column, and filter the data table to limit it to data from just one filter. (See here for more on

filtering.)

After you run Period Finding, there will be another table loaded in to this pane. See the section on main screen, after period finding.

Plot

On the upper right, there is a plot of the input light curve you have loaded into the tool. It takes the time and dependent variable you have selected in the settings pane. If you change the settings pane, this plot changes. If you have > 50,000 points in your light curve, the plot will be binned, e.g., show the light curve as a density plot, where each bin can represent multiple points.

You can change what is plotted and how it is plotted. You can add error bars. For a more complete introduction to this tool's plotting functions, please see the Plots section.

Quick start: Adding errors. To add error bars to the plot of the light curve, click on the gears ( ) at the top of the plot to bring up options for the plot:

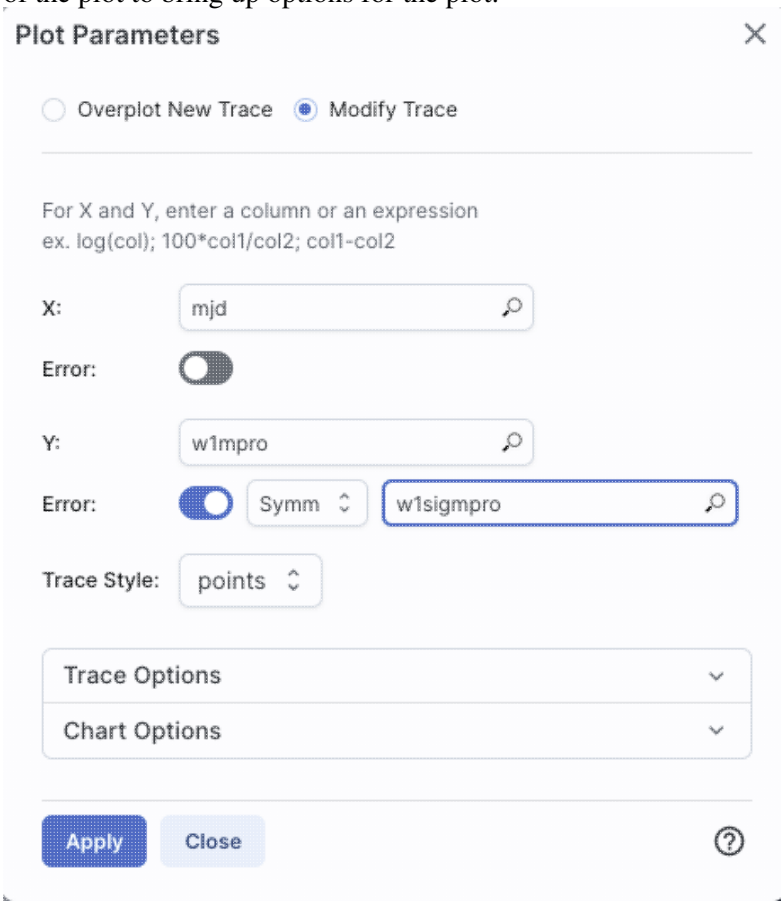

Toggle the switch for 'errors' and tell it which column of your catalog to use for the error bars. Note that the errors you add to this plot are for display purposes only and are not used in subsequent period finding.

"Apply" to implement your changes, "close" to exit without further changes. You can also click on the 'x' in the upper right to remove this pop-up.

Images

Purpose

The purpose of having each image from each epoch displayed is to look for any irregularities (instrumental or otherwise) on or near your target in each individual exposure.

Default view

If you have loaded in a WISE/NEOWISE or ZTF/PTF time series, on the bottom of the screen there is a set of 5 images. Each of these images is centered on your target, and is of the cutout size specified in the settings pane.

Tips and Troubleshooting: Note that if you change the time series to be, say, "w2mpro_ep", the images will NOT automatically change to WISE-2; you must control what images are shown by clicking the radio buttons near the words "image display" near the top of the window.

Dynamic Linking

The images are dynamically linked to the table and the plot. Click on a row of the table, and the images (and plot) change; the image in the center of the row of 5 corresponds to the table row you have selected. Click on a point in the plot, and images (and table) change; the image in the center of the row of 5 corresponds to the plot point you have selected.

If you change the sort order of the table to, say, sort by brightness, then the order of the images also changes. (This will become more important after you run Period Finding.)

Image Displays

The text immediately above the images reminds you how you have sorted the table, and whether it is ascending or descending.

You can control whether 1, 3, 5, or 7 images are displayed in the sequence.

You can control the size of the image cutouts in the settings box on the top left.

The upper left corner of each image reminds you which survey, channel, exposure, size, and zoom you are using.

By default, the tool loads with each image centered on your target. Some targets are on the edge of WISE exposures, so the image may not fill the entire image box for each epoch. To turn this feature off, unclick the "Target Match" box near the upper right of the images pane.

Image Interactions

When you interact with these images, you are interacting with the FITS files, so you can change the color table and stretch, measure distances, etc. For a more comprehensive introduction to this tool's image manipulations functions, please see the Visualization section.

Quick start: Changing the color table or stretch.

To change the color table, go up to the image toolbar above and to the right of the images. Click on the color table icon () and select a new color table from the options presented. Similarly, for the color stretch, click on the color stretch icon () and select a new color stretch from the options presented.

Go back to loading a time series or go on to Period Finder.

Time Series Tool: Period Finder

Contents of page/chapter:

Initial View

When you click on the "Period Finder" button, this is the screen you get:

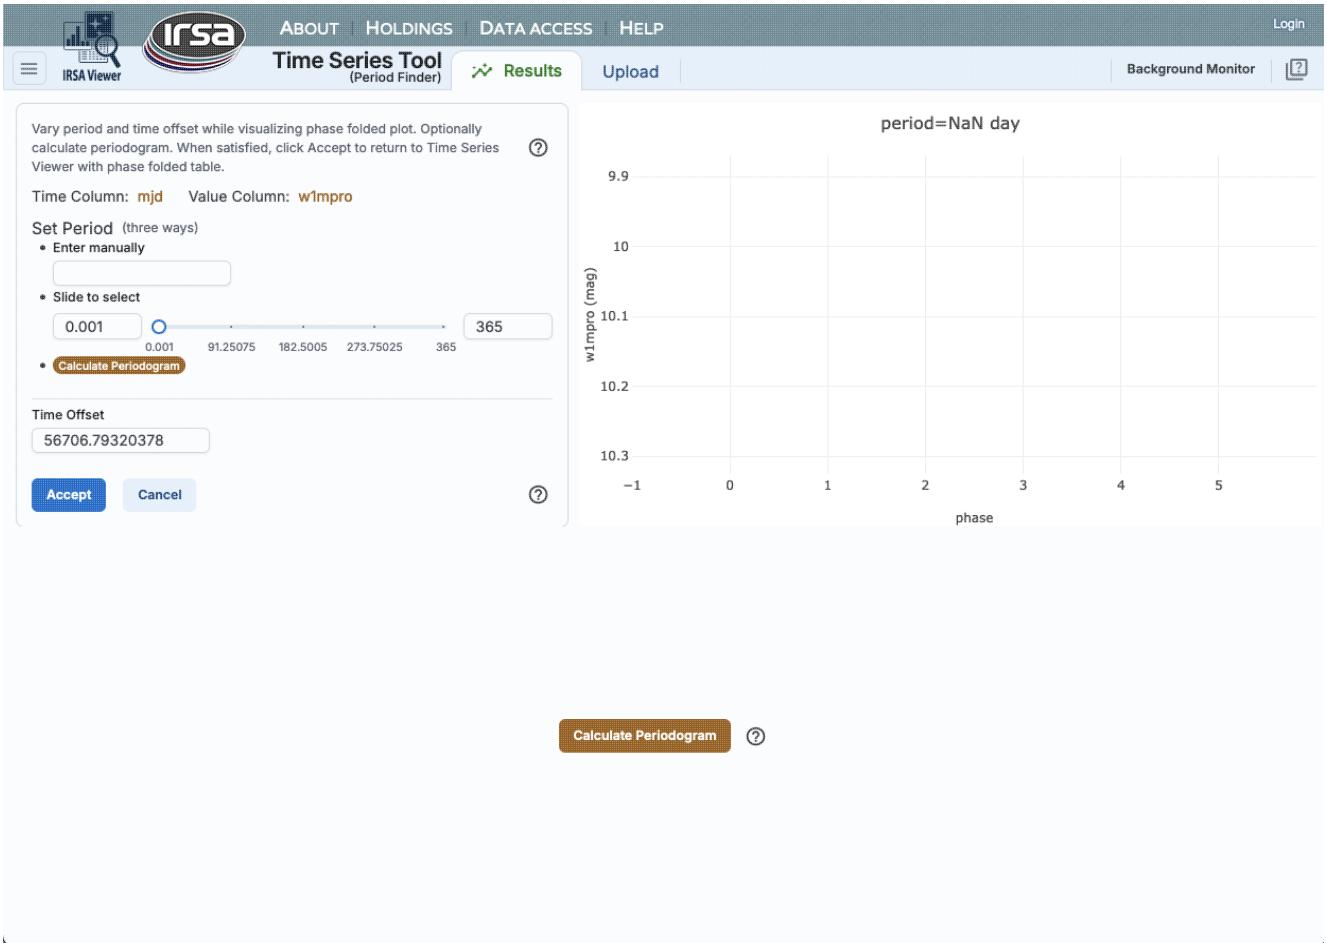

The upper left is another pane, and the upper right will have (eventually) a plot of the light curve. Initially, there is nothing here because no period has been determined yet.

Now we describe these portions of the screen, as well as what to do next.

Settings and Adjusting Periods

The settings pane carries with it the time and dependent variable columns you set in the Viewer (prior screen). These columns cannot be changed from here; you must return to the main screen to change them.

You can set the period from this screen in three different ways:

1. **Enter manually**. Just type the period you want in the box.
2. **Slide to select**. The boxes on the far left and right of the slider control the range over which the slider applies. By default, it comes up with 0.001 d and 365 d. If your target has, say, a period near 0.7 d, but you are not sure exactly what it is (perhaps due to changes in the source itself), you can reset the slider limits to be 0.5 and 1 d, then click and drag the slider slowly through the various periods. The plot of the phased light curve in the upper right changes dynamically as you move the slider.
3. **Calculate periodogram and click to select period**. To use this option, click on the "Calculate periodogram" button either here or in the lower portion of the screen (see below for more information).

It has also extracted the time of the first epoch in your data file and listed it in the box labeled "Time Offset." You can change this if you want.

Once you have a period you like, click "accept" to accept means to return to the main screen, and plot the phase folded light curve using the selected period on that screen. Cancel to return to the main screen without phase folding the light curve.

Plot

The plot in the upper right is a plot of the uploaded light curve, phased to the currently selected period. It comes up without an initial period. If you enter a different period in the settings pane, then this plot dynamically updates.

Tips and Troubleshooting : This plot is unlike any other plot within this tool in that you cannot control what is plotted or how. Because this tool was built to work primarily with ZTF/PTF and WISE/NEOWISE data, both of which are in magnitudes, at least for now, it's going to plot small numbers at the top of the y-axis in this screen. Apologies to everyone working in flux or temperature or anything else where big numbers should be at the top.

Calculate Periodogram

When you first load this portion of the tool, the bottom has only a "Calculate Periodogram" button. Clicking on it starts the process of the periodogram calculation. The code that is actually run by the tool is that from NExScI periodogram service.

Clicking on the button brings up this configuration pop-up:

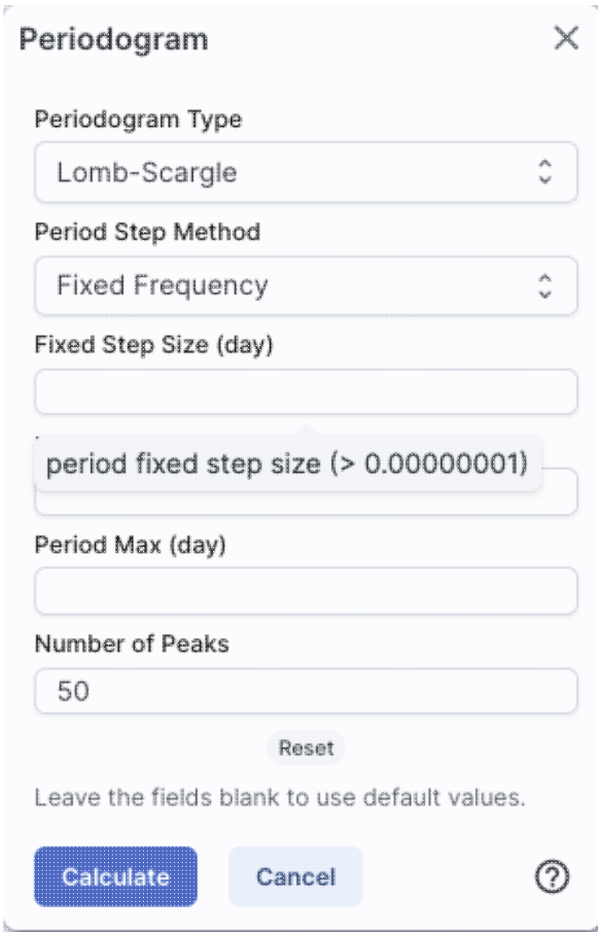

From here, you can select from drop-down menus:

- Periodogram type. Lomb-Scargle is the only choice. (Lomb-Scargle is an approximation of the Fourier transform for unevenly spaced time sampling. It identifies periodic signals that are simple combinations of sinsinusoids.)
- Period step method. You can select fixed frequency or fixed period.

The next set of choices are left blank by default: Fixed step size, period min, and period max. If you leave these variables blank, the routine will calculate values for these parameters as follows:

- fixed step size - if using fixed frequency, this is the delta f; if using fixed period, this is the delta P. The default value is (max-min)/(number of periods to sample), where number of periods to sample is 10 times the number of data points.
- period min - minimum period to calculate. The default value is twice the median time step in the data (which is about the Nyquist limit).
- period max - maximum period to calculate. The default value is half the time baseline.
- number of peaks - number of peaks to identify in a separate table returned by the calculation. The default is 50.

If you want to override any of these selections, you can do so in the boxes in the pop-up.

Click on 'Reset' to reset the parameters to the defaults.

Click on 'Cancel' to close the pop-up without doing anything.

Click on 'Calculate' to initiate the calculation.

After a periodogram calculation has been completed, a "Recalculate Periodogram" button is located in the settings pane.

Periodogram Results

The results of the periodogram calculation appear in two tables (left) and a plot (right), both on the bottom of the screen.

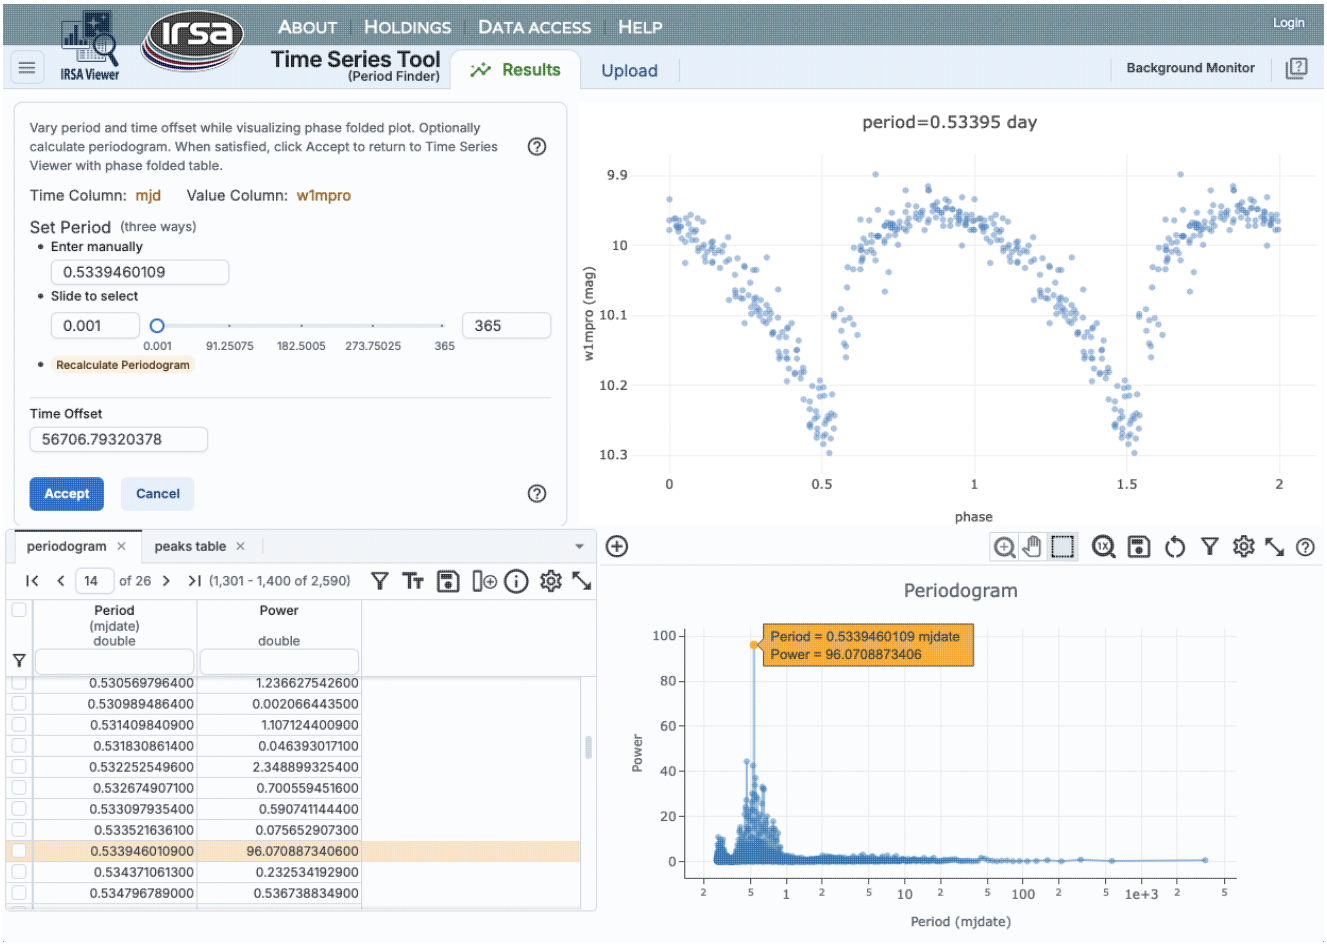

On the left, the two tables are (1) the periodogram itself (which is in the foreground by default) and (2) a table listing just the top *N* peaks, where *N* is a value you specify in the periodogram finding pop-up (the default is 50).

On the right, the periodogram is plotted. This is a plot with all the same basic functionality provided by this tool in that you can change what and how it is plotted -- for example, if you want to remove the lines connecting the points, click on the gears ( ) to bring up options for the table and choose a different plot option from the pop-up.

The plots and tables are dynamically linked. If you click a point in the periodogram, the row is highlighted in the table, and the phased light curve changes. In the settings pane, the period next to "enter manually" is updated to reflect the period you selected.

If you select the other table, the table of peaks, on the left, all the same functionalities apply, except the plot on the lower right is now power versus peak number.

(This power spectrum and period are what you get if you use this example input table from NEOWISE.)

Troubleshooting 1: power spectrum does not have a single peak

If your light curve is sampled with large gaps (as a lot of WISE/NEOWISE and ZTF/PTF light curves are), use caution when interpreting the power spectrum. For example, here is the NEOWISE light curve for VY Lib, a variable star. If you plot that light curve, you will see big gaps between clumps of observations.

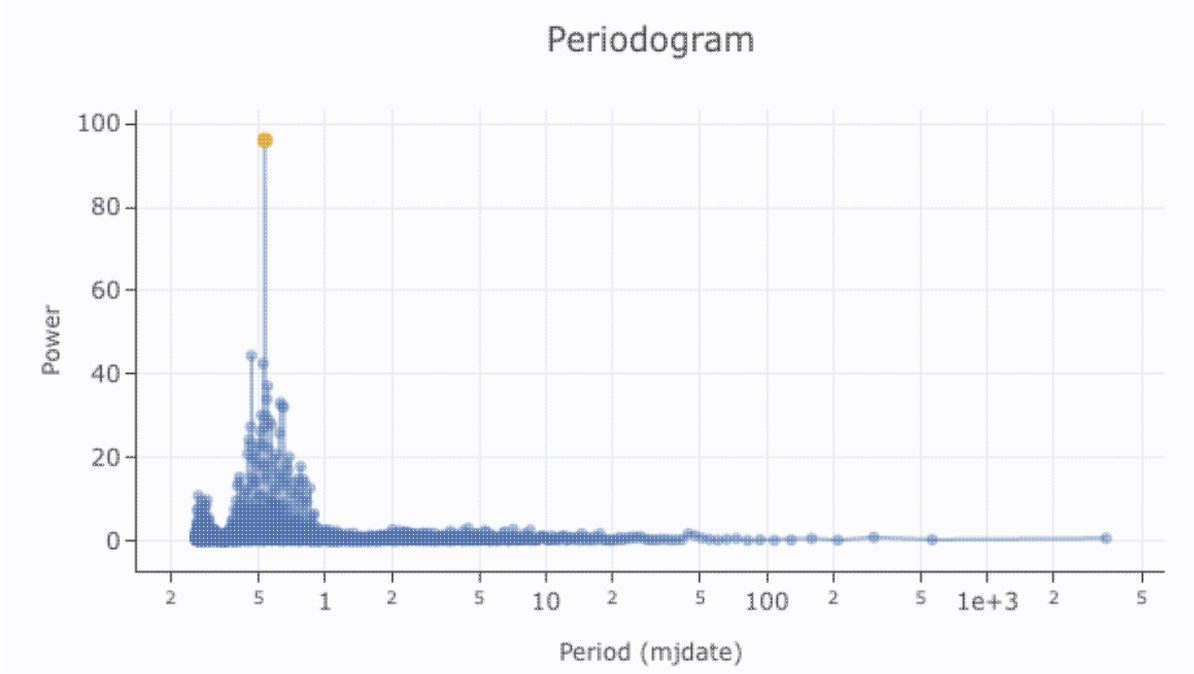

This is what you get if you let the period searching routine pick all its own defaults. The period is 0.5339 d. There is a peak at that period that sticks up above the noise.

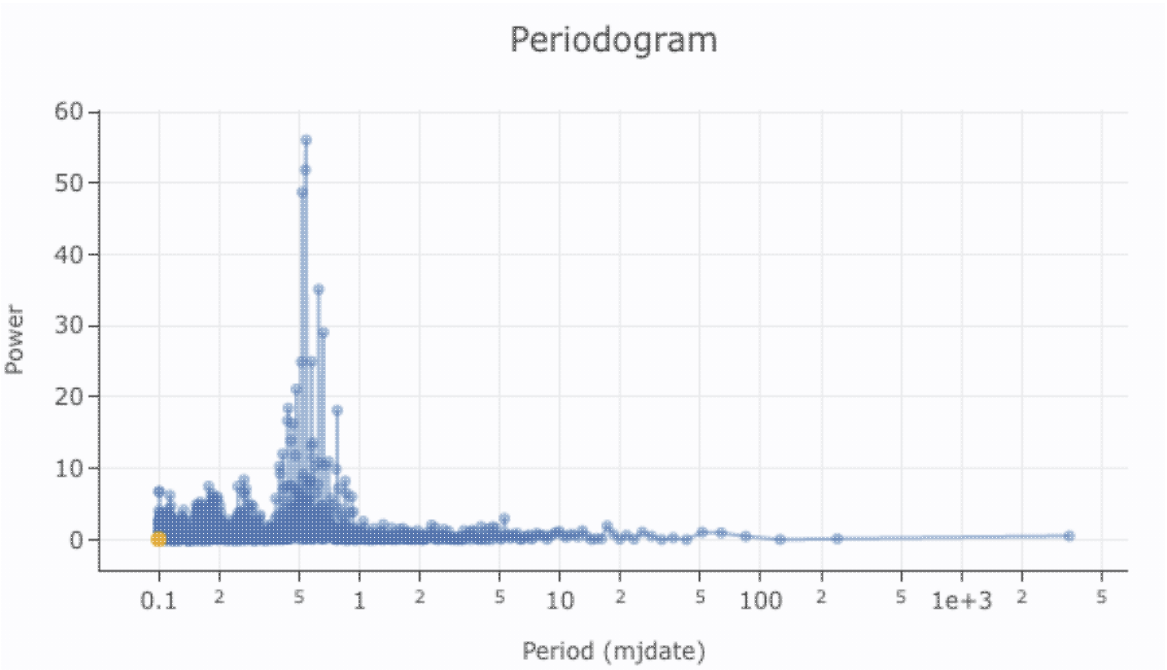

This is what you get if you force the minimum period to be 0.1 d, which changes the influence of the gaps. The 0.5339 peak no longer stands out from the other peaks, and with just this result, one could justifiably

conclude there is no significant period here. Except there really is one, and it's 0.5339 d.

This is what you get if you use only three of the campaigns (three of the groups of points). There are fewer points overall, but there are also fewer big gaps in the light curve. Now the distribution of peaks in the power spectrum has changed again.

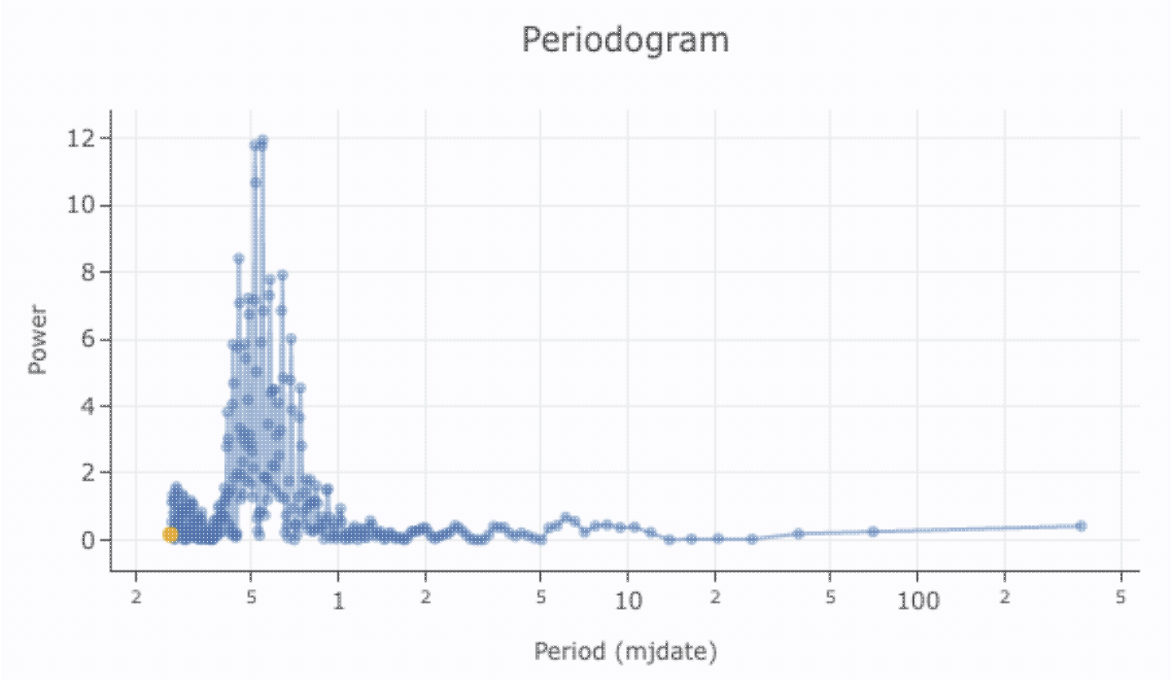

Caution is warranted when interpreting power spectra, especially if the time series has large gaps. Simulations are helpful to know what to believe.

Troubleshooting 2: The version of Plot.ly used by IRSA Viewer sometimes makes ...unusal label choices for logarithmic plots. The two minor ticks between major ticks are 2 and 5, so if they occur between 10 and 100, it means 20 and 50.

Accepting a Period

In practice, you may want to work with all of the available tools to find a period -- periodograms calculated in different ways, different peaks selected, sliders used to align particularly ill-behaved waveforms, etc. After you have a period that you like, put that value in the box next to the "enter manually" text in the settings pane.

The buttons in the lower left of that pane now read "accept" and "cancel." 'Cancel' closes the period-finding window without doing anything further. 'Accept' takes that period as the best possible period, and returns you to the main screen, but now that screen is a little different, because it is after Period Finding.

Go back to Main Screen, Viewer (Initial View) *or* *go on to Main Screen, Viewer (After Period Finder)*.

Time Series Tool: Main Screen, Viewer (After Period Finding)

Now you are back to the main screen, with settings, tables (now more than one), plot, and images.

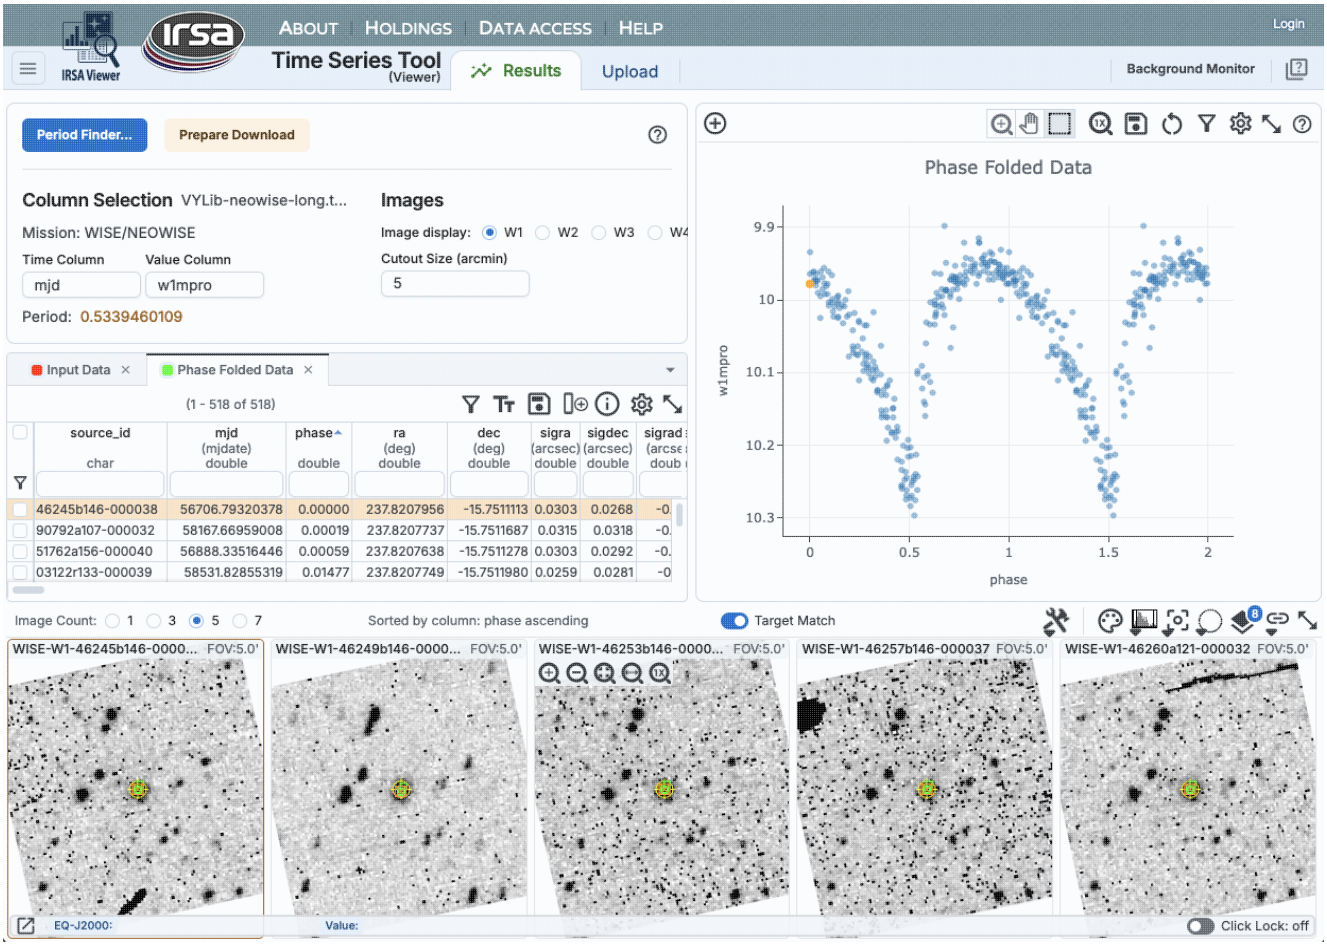

Contents of page/chapter:

Settings

The settings shown in this pane reflect the earlier choices you made, in the Viewer (initial screen). Make changes here if you want to work with another band, and then you can go to the Period Finder again from here.

Note that if you have done calculations on, say, WISE-1, and change the dependent column to WISE-2, the plot on the right will phase fold the WISE-2 data to the period you have selected from the WISE-1 analysis.

Tables

Now the pane with the data table has two tables displayed. The time series you have uploaded into the tool is one tab (as in the Viewer (initial screen)), and the other tab is the phase folded time series (based on the P you selected in period finding).

The phase-folded table comes up, by default, in the foreground, and sorted by phase. The images at the bottom are also sorted by phase.

Plot

On the upper right, instead of the input light curve you have loaded into the tool (dependent variable against time) as you had originally, now it is the phased time series.

Note that if you have done calculations on, say, WISE-1, and change the dependent column to WISE-2 in the settings pane on the left, the plot will phase fold the WISE-2 data to the period you have selected from the WISE-1 analysis.

Tips and Troubleshooting Now you have the plots toolbox in the upper right of the plot, so you can change what is plotted. If you have been working with flux (or something else) rather than magnitudes, this means you can fix it so that large numbers are at the top of the plot. See the plots section for general help on plots.

Images

If you have been working with a WISE/NEOWISE or ZTF/PTF time series, just as you had originally, on the bottom of the screen there is a set of 5 images. Because the table in the foreground is the phase-folded time series, with the table sorted in order of phase, the images are now also sorted in order of phase.

Everything is still dynamically linked. Click on a row of the table, and the images (and plot) change; click on a point in the plot, and images (and table) change.

Go back to Period Finder or *go on to Saving Results*.

Time Series Tool: Saving Results

There are several different ways to save results of analysis done in the Time Series Tool.

Contents of page/chapter:

Tables

On the upper right of any table shown in the tool, there is a diskette option:  . Click on this icon to save the table; you have a variety of choices of file format.

Plots

To save a plot, click on the diskette icon (). The plot will be saved to your disk as a png file.

Single Images

If you want to save a single image, click on the image you want to save so that it has the brown highlighted box around it. The diskette icon () in the image toolbar is within the tools drop-down (), and that will allow you to save that currently selected image as a FITS file (or other formats). Note that **you** control where the file is saved on your disk through your browser; your browser may be configured to store all downloads in a particular location on your disk.

Multiple Images

To save multiple images, identify the row(s) in the table you wish to save. Click on the tickbox on the far left of that row. Select as many as you want; click on the tickbox at the top of the column to select all.

Now, go up to the upper left and click "Prepare Download." It will give you a pop-up:

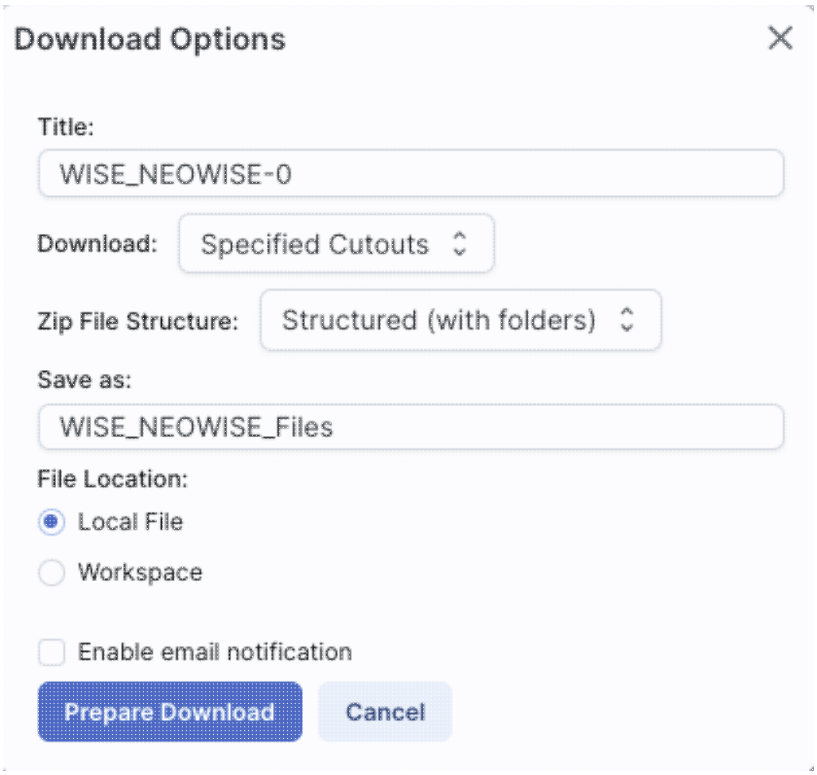

From here, you can specify:

- the title of the download (used in the background monitor to keep track of jobs);
- whether you want the original images or just the cutouts that are displayed in the tool;
- whether you want the file structure to be binned such that each epoch is in a different subdirectory (structured) or with all the files at the same level in the directory hierarchy (flattened);
- the filename of the saved zip file;
- whether that file should be saved on your local disk or in the IRSA Workspace;
- whether you want it to send you email when it is done (and if you tick the box, you can enter your email address).

If it is a small download, it might ask you right away where to put the file on your disk.

If it is a big download, the packaging process will spin off to the "Background monitor" in the upper right. Click on 

to watch it package, and enter an email if you want it to email you when it is done.

Note that if you are working with WISE data, even if you were working with just one WISE channel, the images from all available channels will be included in the download.

Go back to Main Screen, Viewer (After Period Finding).

IRSA Viewer: User Registration for the IRSA Archives

There is one user registration for all IRSA applications, but data access, preferences, search history, data tagging, etc., are all unique to each archive.

While you certainly **do not have to register** to search the archive, download data, analyze data, and write a paper, you may wish to register.

The IRSA Viewer does not access proprietary data, and does not contain any global preferences, so there is no strong reason to log in at this time.

On the other hand, ***if you are trying to gain access to your proprietary data in a particular archive (other than from the IRSA Viewer), you will need to log in*** so that the system grants you access to your data. Not all archives served by IRSA have proprietary data. If you got email from an archive's Help Desk about account information, you may already have an account; else you can set up an account (see below). Once you have set up an account, please send us an email at the IRSA Help Desk and we'll tie your new account to your proprietary data.

Login

The "Login" link is in the far top right of the page.

Select this option to log in if you know your password, or to create a new account.

Create New Account

Find this option by selecting "Sign in" in the far top right of the browser window, and then "Create an account" appears as an option in the lower left of the pop-up window.

Select this option to create a new account.

Forgot your Username or Password

If you do not remember your username or password, select this option to retrieve this lost information.

Find this option by selecting "Sign in" in the far top right of the browser page, and then "forgot your password?" appears as an option below the Login button.

Edit Profile

Find this option by logging in, and then clicking on your account name in the top right of the browser window. Then, select "Edit Profile" to change your password on an existing account.

Change Password

Find this option by logging in, and then clicking on your account name in the top right of the browser window. Then, select "Change Password" to change your password on an existing account.

Update Email

Find this option by logging in, and then clicking on your account name in the top right of the browser window. Then, select "Update Email" to change your email on an existing account.

IRSA Viewer and Catalog Search Tool: FAQs

Do you have any tutorial videos?

Yes. There are also IRSA Viewer Tool and Catalog Search Tool video tutorials, including quick start and longer AAS-demo style overviews, available at the IRSA YouTube channel. Also see the IRSA Viewer playlist (which collects all the IRSA Viewer tutorials together), or the Catalog Search Tool playlist (similarly, combines all the relevant vidoes together), as well as the playlist of tutorials relevant for more than one archive.

How do I get more help?

The "Help" blue tab leads you into this online help, as does "IRSA Viewer Help" from under the "Help" in the IRSA menu on the top of the IRSA Viewer page. You can also download a PDF version of this manual; look at the top of the help window.

You can submit questions to the IRSA Help Desk.

Found a bug? The known bugs and issues in this version of IRSA Viewer are listed here. If you think you have found a bug, before reporting it, please check this list, and read this online IRSA Viewer help. It may be a "feature" we already know about. If you have found a new, real bug then please do contact us via the IRSA Help Desk. Please include your operating system version and your browser software and version. If you can, please also include any specific error message you may have gotten. (NB: In our testing, copying shortcuts worked on Windows and Linux; the command-C did not work on Macs, but selecting and clicking the right mouse button did.)

IRSA Viewer and Catalog Search Tool: Notice to Users -- Privacy Notice

The data contained in this archive are managed by the NASA/IPAC Infrared Science Archive (IRSA), which includes an archive of images, catalogs, and spectra from multiple telescopes and missions, managed by the Jet Propulsion Laboratory. This website is maintained by the Infrared Processing and Analysis Center (IPAC), located on the campus of the California Institute of Technology (Caltech).

The information you provide on a Caltech website will be used only for its intended purpose. We will protect your information consistent with the principles of the Privacy Act, the e-Government act of 2002, the Federal Records Act and, as applicable, the Freedom of Information Act. This notice is posted pursuant to the California Online Privacy Protection Act of 2003 (Cal Bus and Prof Code Sections 22575-22579).

Submitting information is strictly voluntary. By doing so, you are giving Caltech your permission to use the information for the intended purpose. In addition, Caltech may also furnish this information to NASA at NASA's request. If you do not want to give Caltech permission to use your information, simply do not provide it. However, not providing certain information may result in Caltech's inability to provide you with the information or services you desire.

Caltech never collects information for commercial marketing. We will only share your information with a government agency if it relates to that agency, or as otherwise required by law. Caltech/JPL never creates individual profiles or gives your information to any private organization.

We collect no personal information about you when you visit this Web site, unless otherwise stated or unless you choose to provide this information to us. However, we collect and store certain information automatically for use in site management and security purposes. What we collect and store automatically in terms of site statistics is:

- The Internet Protocol (IP) address for the domain from which you access the Internet (e.g., 123.456.789.012) whether the domain is yours individually or is provided as a proxy by your Internet Service Provider (ISP)
- The date and time you access our site
- The pages you access (recorded by the text and graphics files that compose that page)
- The Internet address of the website from which you linked directly to our site.

We use the summary statistics to help us make our site more useful to visitors, such as assessing what information is of most and least interest to visitors, and for other purposes such as determining the site's technical design specifications and identifying system performance or problem areas.

The website also collects and stores information about your search options, such as

- Name resolver choice (NED/Simbad)
- Page size (number of rows)
- Which search results (tabs) should be displayed
- Email address, if provided, for email notifications
- Search parameters so that you can resubmit your search via your search history
- Data tags, if you create one
- Plus, additional preferences that may be developed in the future, such as those tied to the visualization options.

If you register as a user, these options will be kept in our database (along with your login ID and password via MD5 hash) and used for your session the next time you log in. If you do not register as a user, these options are set via cookies kept on your computer; if you clear your cookies and start a new session, these preferences are lost.

At no time is your private information, whether stored in persistent cookies or elsewhere, shared with third parties who have no right to that information. If you do not wish to have session or persistent cookies stored on your machine, you can turn them off in your browser. However, this may affect the functioning of the website on your computer.

IPAC will protect all such information consistent with applicable law.

Comments Sent by E-mail

You may choose to provide us with personal information, as in an e-mail containing your comments or questions. We use this information to improve our service to you or to respond to your request. There may be times when your message is forwarded, as e-mail, to other IPAC employees who may be better able to help you. We normally do not share our e-mail with any other outside organizations, unless determined necessary for security purposes or when required by law. Remember that email isn't necessarily secure. You should never send sensitive or personal information like your Social Security number in an email. Use postal mail or secure websites instead.

Security Notice IPAC is part of the Division of Physics, Mathematics and Astronomy at the California Institute of Technology ("Caltech"), and operates this website as part of a federally funded computer system used to accomplish Federal functions. Unauthorized attempts to defeat or circumvent security features, to use the system for other than its intended purposes, to deny service to authorized users, to access, obtain, alter, damage, or destroy information, or otherwise to interfere with the system or its operation is prohibited. Evidence of such acts may be disclosed to law enforcement authorities and may result in criminal prosecution under the Computer Fraud and Abuse Act of 1986 and the National Information Infrastructure Protection Act of 1996, codified at section 1030 of Title 18 of the United States Code, or other applicable criminal laws.

IPAC uses software programs to monitor this website for security purposes to ensure it remains available to all users and to protect information in the system. Any and all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized Caltech, JPL, NASA, law enforcement personnel, as well as authorized officials of other agencies. By accessing this website, you are expressly consenting to such interception, monitoring, recording, copying, auditing, inspection and disclosure at the discretion of Caltech or NASA. Users have no explicit or implicit expectation of privacy.

Disclaimers

Disclaimer of Liability

With respect to documents available from this server, neither Caltech, nor the United States Government, nor any of their employees, makes any warranty, express or implied, including the warranties of merchantability and fitness for a particular purpose, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Disclaimer of Endorsement

Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Caltech or the United States Government. The views and opinions of authors expressed herein do not necessarily state or reflect those of Caltech or the United States Government, and shall not be used for advertising or product endorsement purposes.

Copyright Status

For information on possible copyright infringement, please visit Caltech's "Copyright Infringement" page.