
THELIGUI Documentation

Release 2.7.0

Mischa Schirmer

July 09, 2013

CONTENTS

Version 2.7.0

Last update July 09, 2013

Latest changes ChangeLog

THELI GUI a is powerful and easy-to-use package for astronomical image reduction under Linux, offering e.g.

- Automated reduction of optical, NIR and MIR data
- 90 pre-configured instruments
- Parallelisation
- Flexible background correction
- Automatic distortion correction and creation of large mosaics
- Highly flexible coaddition: e.g. locking onto proper motion targets, various projections, etc
- Optimised weighting schemes and defect detection
- Crosstalk and non-linearity correction
- Atmospheric transparency correction
- Absolute flux calibration
- Extensive documentation and online help

HELP

- Looking for something, but you don't know what it is called? Browse the *genindex*.
- Or find it using the *Search page*.
- THELI is bugging you? Try the *THELI trouble-shooter and FAQ*.

DOCUMENTATION

These webpages are also available in pdf format, but some of the figures are missing due to the automatic conversion.

2.1 Download

THELI core components

- The pipeline core: `theli-1.6.1.tgz`
- The graphical user interface: `gui-2.7.0.tgz`

The THELI GUI is bound to a fixed version of the THELI pipeline core to ensure compatibility. You must use the pipeline core version linked above.

THELI GUI: Recent changes

External package: Astrometry and photometry

We highly recommend you download the latest [Scamp](#) (thanks Emmanuel Bertin!) version from the repository, and compile it. If you don't succeed, you may use *Scamp v1.4.6*, which is the last fully statically compiled version and we provide here for convenience.

- `scamp 32bit athlon`
- `scamp 32bit i686`
- `scamp 64bit x86_64`

SDSS ugriz photometric reference catalog

We provide a subset of the SDSS stripe 82 photometric calibration catalog (Ivezic 2007). The original published version of this catalog contains about 1000000 stars. We keep only sources with individual errors less than 0.05 mag in *all* filters, and remove very faint sources in *gri* filters which could lead to a systematic bias of 0.02 mag in the ZP. About 340000 photometric reference sources remain.

Stripe 82 is located between $20:30:00 < RA < 04:00:00$, and $-1.27 < DEC < 1.27$. By downloading the tarball below, you will obtain two versions of the catalog, one for the SDSS native *ugriz* system, and one for the USNO-1m *u'g'r'i'z'* system.

- Stripe 82 ugriz photometric catalog (40 MB)

In order to use these catalogs, you have to unpack the tarball in `gui/photstdcats/`.

2.2 THELI installation

2.2.1 Software pre-requisites

In order to use THELI GUI, your system needs to meet the following requirements:

Compilers

- gcc 3.4 (or higher)
- python 2.5 (or higher, including numpy, scipy and matplotlib; THELI is NOT yet compatible with python 3)

Libraries

- fftw3
- libtiff4
- GNU scientific library (GSL, including development version)
- xorg (development version)
- Qt3 (development version; **note that THELI GUI is NOT compatible with Qt4!**)

Note: All of these packages should be available as pre-compiled binaries from the online repository of your Linux distribution. This makes for the fastest and easiest installation.

Other requirements

- An internet connection (to download astrometric reference catalogues)
- FITS viewers such as [skycat](#) or [ds9](#) are very useful. If you have problems running the latest skycat version on your machine, try the oldest version 2.7.3 which has all necessary libraries statically linked.

Under **Debian/Ubuntu** the following commands update your system with all necessary packages (Ubuntu is very slim, hence the long list). Pre-compiled packages for other Linux flavours will likely have very similar names:

```
sudo apt-get update
sudo apt-get upgrade
sudo apt-get install csh g++ libqt3-mt-dev gawk wget tcl8.5-dev tk8.5-dev
sudo apt-get install netpbm imagemagick libfftw3-dev xorg-dev pgplot5
sudo apt-get install libtiff4-dev xorg-dev pgplot5 libpng3 binutils-dev
sudo apt-get install build-essential libncurses5-dev xterm gzip
sudo apt-get install python python-dev python-numpy libxxf86vm-dev
sudo apt-get install python-scipy python-matplotlib libgsl01db1 libgsl0-dev
```

2.2.2 Updating THELI

If you are updating, simply go to the corresponding section below and follow the instructions. Old installations do not need to be deleted, but you can do so without harm. A new GUI version can be installed anytime.

Note: If you install the THELI pipeline core, then the GUI must be (re-)installed.

2.2.3 Installing THELI

Installation takes place in an arbitrary directory, let's assume this is:

```
/home/user/THELI/
```

Move the downloaded files into this directory.

THELI core

Installing the pipeline core is really easy. Unpack and install it with

```
tar xvfz theli-<version>.tgz
cd theli-<version>/pipesetup
./install.sh -m "ALL"
```

Quite some of these programmes can also be used on the command line. To do that, you have to include

```
/home/user/THELI/theli-<version>/bin/Linux[_64]/
```

in your PATH environment variable.

THELI GUI

Warning: The THELI GUI does not (yet) compile under MacOS! If you want to use the GUI on your Mac, create a [virtual machine](#) running Linux and install THELI therein. At the moment I cannot predict when (or if) a native MacOS version can be made.

Unpack and install the GUI with:

```
tar xvfz gui-<version>.tgz
cd gui-<version>
./install.sh
```

Note: The installer assumes that the *libfftw3* and the *libtiff4* libraries are available on your machine (most likely located in */usr/lib/*).

Once the installation finished, you will have a *theli* executable

```
/home/user/THELI/gui-<version>/theli
```

and the following directory tree

```
/home/user/THELI/theli-<version>/
/home/user/THELI/theli-<version>/bin/Linux[_64]/
/home/user/THELI/gui-<version>/
```

including these three links

```
/home/user/THELI/gui -> /home/user/THELI/gui-<version>/
/home/user/THELI/theli -> /home/user/THELI/theli-<version>/
/home/user/THELI/theli/gui -> /home/user/THELI/gui-<version>/
```

The links are created during the installation of the GUI, and they point to the latest installed version. Like that you can install updates of the pipeline core and the GUI without interfering with older installations. Lastly, you have to copy the *scamp* executable to

```
/home/user/THELI/theli-<version>/bin/Linux[_64]/
```

and include that directory in your PATH environment variable. Collected therein are all binaries used by THELI. Many of those are also very useful command line tools, some of which are mentioned explicitly in this documentation.

2.2.4 Launching THELI

Warning: Do not start THELI by clicking on a Desktop icon linked with the executable. In this case a series of environment variables are not propagated correctly, and THELI will not function properly.

The correct way to start THELI is to call the *theli* executable directly from the command line. If you include

```
/home/user/THELI/gui
```

in your `PATH` variable, then the most recent version is always found, independent of your present working directory.

2.3 Getting started

2.3.1 What THELI does and what it does not

THELI is a generic data reduction package. Feed it a set of raw images together with suitable calibration files, and it will produce a fully calibrated, science grade stacked image.

At this point the responsibility of THELI ends. When it comes to extracting science from the final image, individual needs and requirements must be met. Their number is as large as the number of scientific questions that are out there to be answered. THELI will provide you with an excellent image to start with, but you have to push it through DAOPHOT, GALFIT or your specific measurement pipeline yourself.

A great deal of care has been taken in the development of THELI in order to make everything as instrument-independent as possible. For example, once you figured out how to reduce data from a single-chip camera mounted on one telescope, you are already familiar with the processing of multi-chip wide field camera data from a totally different observatory. In fact, there is no difference in processing these two data sets. The same holds for near-IR data. Processing is very similar to optical data, you just need to pick the right sky subtraction method.

2.3.2 How NOT to use THELI

While reducing data with THELI's graphical user interface is (or will be) probably one of the most convenient ways you have encountered so far, it does not relieve you from learning and understanding what you are doing and how the software works. It is common to see the brains of scientists (and others) shutting down when they come across a GUI with buttons to click and switches to toggle. If you blindly start somewhere in the middle, THELI will bite you (or at least gives you a sensible warning before it bites you).

THELI uses a strictly linear approach for data reduction, which makes for easy and transparent processing. And you have to stick to that. It is usually not possible to take data that has been processed half-way by other software and do the rest in THELI. FITS headers will not be understood, naming conventions not met, and data structures totally incompatible (think of multi-chip cameras).

2.3.3 Some good advice

It is of utmost importance that you know what you are doing. Blindly applying e.g. a defringing correction or a sky subtraction with default parameters can wreak havoc in your data if you do not know what is going on. For example, if you are interested in some faint extended features around some object, and you choose the wrong sky subtraction method, this feature may be removed entirely from the data. Or, because of the wrong defringing approach or near-IR sky background modelling, the noise in your final image is twice as high, and the faint supernova you would have caught otherwise and made you famous is gone. Fortunately, all default settings in THELI will lead to very reasonable

results in most cases. But you should know the implications of the various selections you make. These web pages will help you, as they explain the effects of all the whistles and bells THELI has.

Do not forget that a successful data reduction depends almost always critically on carefully planned and executed observations. If you fail to take calibrations, you will be lost. Don't settle for 4 flats and 5 biases, and don't skip the darks because "the camera is cooled anyway". Don't rely on data archives, maybe the last time flats were taken in your filter was half a year earlier, and that observer pointed right into cirrus clouds AND saturated the exposures. Over-calibrate your data, don't under-calibrate it, you never know what crops up during data reduction.

Lastly, the observing strategy is crucial as well. For example, if you observe extended targets and require accurate sky modelling, then even in the optical regime off-target blank fields can be essential. They cost extra time, but they save your data (infrared astronomers can tell you a thing or two about that). See my *guidelines for observers* for more things to do or not do do.

THELI isn't magic. It cannot save your data if the calibrations and/or the observations were messed up. Neither will any other software.

2.3.4 Processing overview

For a complete data reduction you have to work your way through the seven main sections of THELI. These are:

- Initialise (screen-shot)
- Preparation (screen-shot)
- Calibration (screen-shot)
- Background (screen-shot)
- Weighting (screen-shot)
- Astrom/Photom (screen-shot)
- Coaddition (screen-shot)

The last six sections deal with the data reduction itself. Therein you find a blue task area containing the various processing steps available. Those which are mandatory are highlighted with red text.

Almost all tasks can be configured individually. Make your parameter selection, then mark the task you want to execute. More than one task can be marked at the same time. In the *command window* the exact syntax of the script that is going to be executed is displayed. If you are familiar with THELI, you can edit them according to your needs, but that should rarely be the case.

Once you are happy with your settings, hit the *Start* button. In the yellow message window you get some feedback as of the task that is currently being executed. All tasks write an extensive log file to disk, which, upon the end of the execution, is scanned for characteristic error messages. If no errors are found, a success message ("Done") is printed. Otherwise the line in which the error is found is printed, and the logfile will be loaded and shown to you for inspection.

Each task that modifies the image as such inserts a unique string based on which you can tell what processing steps were applied already. The previous processing stage is stored in separate sub-folders. No data is ever deleted.

Initialise

Here you tell THELI where the data is located on your hard disks, and which instrument it was taken with. You also pick a meaningful name for this reduction task, under which THELI stores all the parameters used and tasks already executed. When launching THELI the next time, you will find it in exactly the same state as when you left it, and you can proceed with data reduction.

Preparation

There is only one mandatory task in this section: to transform the raw data and their FITS headers such that it conforms to THELI standards. Multi-extension FITS files are split into individual chips such that they are accessible for parallel processing.

Calibration

In this section the master calibration files (bias, dark, flat) are created and applied to the data. Should a superflat or sky background model be needed at this step for defringing or sky subtraction (near-IR data), then this can be configured and created here as well.

Superflattening

If your data still have instrumental signatures after the basic pre-processing steps, then you can attempt to remove those in this section. Available options are multiplicative superflattening, defringing, or sky subtraction (near- and mid-IR). You can also attempt to remove horizontal or vertical gradients which are often introduced by near-IR detectors (*reset anomaly*).

Weighting

THELI offers an advanced weighting scheme and defect detection mechanism. This ensures the cleanest coadded image possible with optimal S/N.

Astrom/Photom

This section basically deals with catalogs. Optionally, you can attempt an indirect absolute photometric calibration based on standard star observations, or a direct absolute calibration based upon comparison with SDSS or 2MASS magnitudes.

The remaining two tasks are mandatory. Object catalogs are produced for each image and compared to an astrometric reference catalogue downloaded from the web. A full astrometric solution is created including distortion correction and automatic mosaicing. Relative photometric zeropoints are determined as well to correct for varying atmospheric extinction. Results are written to separate FITS header files.

Coaddition

Here you can subtract individual sky models from each exposure. A large range of options is available to suit your particular needs for background modelling.

Lastly, the coaddition is performed. To this end a large number of configuration parameters is available, presenting you amongst other parameters various resampling kernels, the output pixel scale, or a possible proper motion vector for moving targets.

2.4 Instruments

2.4.1 Pre-configured instruments

The following **optical** cameras are available in THELI:

Instrument Name	# chips	Remarks
ACAM@WHT	1	
ALFOSC@NOT	1	
ALTAU16M@VYSOS06	1	
AltaU42_LOWRES@ASV	1	
AltaU42_HIGHRES@ASV	1	
CFH12K@CFHT99	12	
CFH12K@CFHT	12	
DECam@CTIO	62	
EFOSC2@ESO3.6m	1	
EFOSC_2x2@ESO3.6m	1	
EMMI_BIMG@NTT	1	
EMMI_RILD@NTT	1	
ENZIAN_CAS@HOLI_1M	1	
FORS1_1CCD@VLT	1	Old configuration
FORS1_2CCD@VLT	2	New configuration
FORS2_1CCD@VLT	1	Old configuration
FORS2_2CCD@VLT	2	New configuration
GMOS-N@GEMINI	3	2x2 binning
GMOS-S@GEMINI	3	2x2 binning
GOODMAN@SOAR	1	1x1 and 2x2 binning
GPC1@PS1	64	
GROND_OIMG@MPGESO	1	
IMACS_F2_OLD@LCO	8	
IMACS_F2_NEW@LCO	8	
IMACS_F4_OLD@LCO	8	
IMACS_F4_NEW@LCO	8	
LAICA_2x2@CAHA	4	
LAICA@CAHA	4	
LBC_BLUE@LBT	4	
LBC_RED@LBT	4	
LDSS3@LCO	2	
LORRI@NewHorizons	1	
MEGAPRIME@CFHT	36	
MEGAPRIME_ELIXIR@CFHT	36	ELIXIR pre-processing
MEROPE@MERCATOR	1	
MOSAIC-I_old@KPNO_0.9m	8	Old configuration (before Aug. 2010)
MOSAIC-I_old@KPNO_4.0m	8	Old configuration (before Aug. 2010)
MOSAIC-II_8@CTIO	8	8-channel mode
MOSAIC-II_15@CTIO	8	16-channel mode, 1 dead (after 2009-03)
MOSAIC-II_16@CTIO	8	16-channel mode (before 2009-03)
MOSCA_2x2@NOT	4	
OASIS4x4@WHT	1	
OASIS@WHT	1	
OMEGACAM@VST	32	
PFC@WHT	2	
SDSS	1	Download SDSS FITS images and process them
SOI@SOAR	2	
SuprimeCam_200101-200104@SUBARU	9	One chip dead
SuprimeCam_200105-200807@SUBARU	10	
SuprimeCam_200808@SUBARU	10	New red-sensitive detectors

Continued on next page

Table 2.1 – continued from previous page

Instrument Name	# chips	Remarks
SuprimeCam_200808_SDFRED@SUBARU	10	New red-sens. det., SDFRED pre-processing
SuSI2_2x2@NTT	2	
SuSI2old_2x2@NTT	2	
VIMOS@VLT	4	
WFC@INT	4	1x1 and 2x2 binning
WFI@AAT	8	
WFI@SSO_40inch	7	
WFI@MPGESO	8	
Y4KCam@CTIO	1	

The following **infrared** cameras are available in THELI:

Instrument Name	# chips	Remarks
Dolores@TNG	1	
FLAMINGOS2@GEMINI	1	
FourStar@LCO	4	
GROND_IRIM@MPGESO	1	
GSAOI@GEMINI	4	
HAWKI@VLT	4	
INGRID@WHT	1	
ISAAC@VLT	1	
LIRIS_POL@WHT	4	
LIRIS@WHT	1	
MMIRS@LCO	1	
MOIRCS@SUBARU	4	
NACOSDI@VLT	1	
NEWFIRM@CTIO	4	
NICS@TNG	1	
NICI@GEMINI	2	
NIRI@GEMINI	1	
NOTcam_highres@NOT	1	
NOTcam_lowres@NOT	1	
Omega2000@CAHA	1	
OSIRIS_F3@SOAR	1	
OSIRIS_F7@SOAR	1	
OSIRIS@GTC	1	
PISCES@LBT	1	
SOFI@NTT	1	
SPARTAN@SOAR	4	
VIRCAM@VISTA	16	
VISIR@VLT	1	mid-infrared
WIRCam@CFHT	4	

2.4.2 Setting up a new instrument

If your instrument is not contained in the *pre-defined lists of instruments*, you can include it in the *Initialise → Create instrument (user defined)* dialog. For a camera with 8 CCDs it would like this:

Instrument configuration | Splitting script

General properties

Instrument name: WFI@MPGESO

Number of CCDs: 8

Geographic latitude: -29.257 x center: 1000 x extent: 1000

Pixel scale (arcsec): 0.238 y center: 2000 y extent: 1000

Image statistics

CCD geometries

	# CCDs
Overscan xmin	10 10 10 10 2105 2105 10
Overscan xmax	40 40 40 40 2135 2135 40
Cut-on x	51 51 51 51 51 51 51
Cut-on y	31 31 31 31 2 2 2
x size	2032 2032 2032 2032 2032 2032 2032
y size	4092 4092 4092 4092 4092 4092 4092
CRPIX1	3867 1721 416 -2551 -2552 -412 1725 3864
CRPIX2	-233 -231 -224 -225 3893 3889 3885 3886

Fill in examples

ST-10 (1 CCD) WFI@MPGESO (8 CCDs) WFC@INT (4 CCDs) current instrument

Guess... Save Load... Clear All

Automatic creation

Configuring a new instrument for THELI is really easy. THELI can guess the chip geometries including overscan regions based on a well-exposed flat field (illumination level: 20-30 kADU). For this purpose an edge-detection algorithm is used. Just click on the blue **Guess...** button and select a suitable flatfield exposure. This can be a simple FITS image or, in case of a multi-chip camera, a multi-extension FITS file. In the latter case all extensions will be probed.

DSLR RAW format is also supported. In this case a simpler algorithm is used that works with any image.

Note: Fields that remain blank after this process must be filled in manually. You should check that the automatically determined values make sense, in particular overscan regions. For cameras with incomplete or ill-defined FITS headers the reference pixel should be checked as well.

General properties

Fields are self-explanatory.

Instrument name:

Stick to the *Instrument@Observatory* or *Instrument@Telescope* scheme. The name of the instrument chosen must be unambiguous. THELI will complain later on if it encounters two instruments with the same name.

CCD geometries

In this section the detector layout is defined. For multi-chip cameras you must enter a blank-separated list of values in each field, one number for each detector. At the right end you get some visual feedback: the counter turns green when you hit the same number of entries as specified in the *Number of CCDs* field above. For cameras with more than a half a dozen detectors this is quite useful. If you are uncertain how this should look like, try the *automatic approach*, or fill in some of the examples, or select one of the pre-configured instruments (the *Load...* button).

The fields are as follows:

1. **Overscan xmin:** The leftmost pixel of the overscan region. If the detector has no overscan, then enter a zero.
2. **Overscan xmax:** The rightmost pixel of the overscan region. If the detector has no overscan, then enter a zero.
3. **Cut-on x:** The first pixel in x-direction that receives light.
4. **Cut-on y:** The first pixel in y-direction that receives light.
5. **x size:** The number of pixels in x-direction receiving light.
6. **y size:** The number of pixels in y-direction receiving light.

7. **CRPIX1**: The x-coordinate of the astrometric reference pixel.
8. **CRPIX2**: The y-coordinate of the astrometric reference pixel.

Warning: The parameters entered here must reflect the chip status **after** the data run through the *splitting script*, as the latter will rotate or flip images if you entered a non-zero *transformation code*.

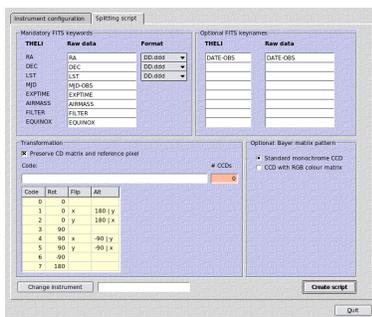
Note:

1. Do not confuse **xly size** with the coordinate of the last pixel receiving light on the xly axis.
2. Your choice of the first six parameters (**overscan**, **cut-on** and **size**) does not need to reflect the actual physical properties of the detector. For example, if your overscan region covers pixels 1 to 50, but you think that only pixels 20 to 45 are well-behaved, then you would use the latter values in the camera configuration. Same for the detector area that receives light.
3. The reference pixel does not need to be accurate. The astrometric solution will determine it precisely. In exceptional cases it can be helpful if you get it right within, say, 10 arcseconds, but usually errors as large as several arcminutes are recovered easily.
4. You must save the instrument configuration before proceeding to the creation of the *splitting script*.
5. If you are an amateur astronomer and use a chip with Bayer matrix, then **xly size** must consist of even numbers (see *Bayer matrix*)

Once the instrument has been created, it can be selected from the list of user-defined instruments in the *Initialise* section.

2.4.3 Creating the splitting script

In the previous section you told THELI about the geometries of your detectors. In addition, you have to provide information about the raw data format itself, such that it can be converted into something THELI can read and make sense of. This is done by an instrument-specific splitting script, which you should define immediately after finalising and saving the instrument configuration. Just switch to the *Splitting script* tab menu to invoke the corresponding dialog:



The splitting script will do several things to your data:

1. In case of a multi-chip camera data is mostly stored in a multi-extension FITS file. The script will break this FITS up and write individual 32bit floating point FITS files for each detector, such that full advantage can be taken from parallel processing. The chip number will be introduced into the file name just before the suffix. For example, if the file *image.fits* contains 4 extensions, the individual files will be named:

```
image_1.fits
image_2.fits
```

```
image_3.fits
image_4.fits
```

2. The script will write a very basic and compact *FITS-compliant header*, containing only the essential information about the data itself and a valid WCS solution.
3. If desired, all chips in the detector mosaic can be brought to the same orientation.
4. Amateur astronomers with colour-coded detectors using a Bayer matrix have to specify if their data is in DSLR RAW or FITS format. In the latter case one must indicate how the Bayer matrix is encoded. This can involve some trial and error. See also the section about the *Bayer matrix* below.

Mandatory FITS keywords

You must tell THELI the names of the FITS keywords in your data, such that your raw FITS header can be properly translated into a THELI-conform header. For some keys such as RA, DEC and LST you must define the format as well.

EXPTIME

For the exposure time keyword you can either specify a single FITS keyword, or in case of near-infrared cameras, two comma-separated keywords, e.g.:

```
DIT,NDIT
```

where DIT represents the individual detector integration time, and NDIT the number of such exposures averaged for the final image. The resulting image will be multiplied by NDIT in order to maintain a consistent background level (otherwise exposures with different NDITs cannot be processed/combined consistently).

FILTER

If two filter wheels are present in the instrument, you can specify two comma-separated filter keywords, e.g.:

```
FILTER1,FILTER2
```

The resulting FILTER keyword in the FITS header will then read:

```
FILTER1+FILTER2
```

Optional FITS keywords

You can propagate up to 8 other FITS keywords from your raw data into the THELI conform headers. You must provide an unambiguous key name for the THELI header, and the name of that key in the raw data. If the key name for the THELI header is longer than 8 characters, it will be truncated without warning.

Transformation

Preserve CD matrix and reference pixel

If your data come with a valid astrometric WCS header, you should leave this setting activated. If you unmark it, THELI will re-orient the CD matrix in the FITS headers such that North is up and East is left. In any case the precise (or correct) orientation will be determined during the astrometric part of the reduction.

Code

Often the chips in multi-chip cameras have different orientations. While this is no problem for THELI, it can be easier to work with the data if the orientation of all chips is identical, in particular if the data have to be eyeballed frequently. You must enter a blank separated list of numbers for the transformation code, one number for each detector. If no

transformation should be done, then enter zeros. If a rotation and flip are combined, then the rotation is performed first. Positive rotation angles are counted counter-clockwise. The coding is as follows:

Code	Rotation	Flip
0	0	
1	0	x
2	0	y
3	90	
4	90	x
5	90	y
6	-90	
7	180	

Warning: If non-zero transformation codes are entered, the *chip geometries defined previously* must reflect the status AFTER any flipping or rotation has been applied.

Bayer matrix

All professional instruments have monochrome CCDs. If you are a scientist, leave the default selection unchanged and ignore the rest of this section.

If you are an amateur astronomer, and you work with DSLR CMOS chips (data in RAW format) or CCDs with a Bayer colour matrix imprinted on the chip (data in FITS format), then select the second option. If your data is in FITS format, then you must also specify the colour coding of the Bayer matrix as it is usually not declared in the FITS header. If you do not know the encoding of the matrix this requires some trial and error.

Note: If you use a CCD with Bayer matrix and the data are in FITS format, then the Bayer matrix coding also depends on the settings for the *CCD geometries*, in particular *Cut-on xly*. In addition, *xly size* must be made of even numbers, as otherwise the de-bayering later on will not work correctly.

2.4.4 THELI FITS headers

This is a short primer for the THELI FITS header (and FITS headers in general). The FITS header contains some basic information about the data, encoded in keywords. After running the splitting script it will look like this:

```
SIMPLE = T / Fits format
BITPIX = -32 / bits per pixel
NAXIS = 2 / single image
NAXIS1 = 2048 / x size
NAXIS2 = 4092 / y size
BSCALE = 1.0 / pixel scale factor
BZERO = 0.0 / pixel value offset
CTYPE1 = 'RA---TAN' / WCS coordinate type
CTYPE2 = 'DEC--TAN' / WCS coordinate type
CRPIX1 = 1092.00 / WCS Coordinate reference pixel
CRPIX2 = 2136.00 / WCS Coordinate reference pixel
CD1_1 = -6.61111e-05 / Change in RA at CRPIX along 1st axis
CD1_2 = 0.0 / Change in RA at CRPIX along 2nd axis
CD2_2 = 6.61111e-05 / Change in DEC at CRPIX along 2nd axis
CD2_1 = 0.0 / Change in DEC at CRPIX along 1st axis
CRVAL1 = 246.41666667 / RA at CRPIX
CRVAL2 = -24.53870000 / DEC at CRPIX
RADECSYS= 'FK5' / Coordinate system for equinox (FK4/FK5/GAPPT)
```

```

FILTER = 'B/99' / Name of filter
OBJECT = 'IC4603' / observed target
AIRMASS = 1.357000 / (average) Airmass during observation
EXPTIME = 600.00 / effective exposure time
EQUINOX = 2000.00 / Equinox of coordinates
IMAGEID = 1 / Chip Number
GABODSID= 821 / Obs. date in days since 31/12/1998
EISID = 0 / no comment
ZP = -1.000000 / photometric zeropoint
COEFF = 1.000000 / photometric extinction coefficient
DATE-OBS= '2010-05-12T02:02:26.989'
DUMMY1 = 0 / Dummy for adding new FITS cards later
...
DUMMY50 = 0 / Dummy for adding new FITS cards later
HISTORY
...
END

```

Herein, the CD matrix (the CDi_j elements) encode the sky position angle, the pixel scale, and a possible flipping. CRVAL1 and CRVAL2 are the right ascension and declination, respectively. A list of 50 DUMMY keywords is introduced as well, they will be partially replaced by more meaningful values during data reduction.

Warning: You should never tamper with FITS headers if you do not know what you are doing. And even if you know what you are doing, it is a very smart idea to create a backup copy before you touch the headers.

2.5 File organisation

THELI places a number of files in different directories. Here is an overview of where you can find what, assuming that THELI has been installed in

```
/home/user/THELI/
```

2.5.1 Binaries

THELI's main binaries and other useful command-line tools can be found in

```
/home/user/THELI/theli/bin/Linux[_64]/
```

2.5.2 Data reduction scripts and configuration files

Scripts and configuration files are stored in separate places

```
/home/user/THELI/gui/scripts/
/home/user/THELI/gui/reduction/
```

You do not have to edit configuration files manually. This is done by choosing the desired parameters from within the GUI.

You do not have to edit scripts either. Should this be necessary, then THELI creates them on the fly and stores them in

```
/home/user/.theli/scripts/*.log
```

2.5.3 Instrument configuration files

Pre-defined and user-defined instruments are kept in different places

```
/home/user/THELI/gui/scripts/instruments_professional/  
/home/user/.theli/instruments_user/
```

2.5.4 Temporary files

Temporary files are mostly deleted when they are not needed anymore. However, a few of them are kept for trouble-shooting purposes. Files in the *tmp* directory can be deleted safely anytime does not perform any tasks

```
/home/user/.theli/tmp/
```

If you want THELI to keep temporary files in a different directory, then link to it from the nominal place

```
ln -s /home/user/.theli/tmp /preferred/path/to/tmp
```

2.5.5 Data reduction LOGs

The LOG files keep a memory of everything that happened during a task. They appear somewhat cryptic and can only be interpreted if you have some bash-scripting knowledge, but are invaluable for trouble-shooting

```
/home/user/.theli/scripts/*.log
```

Each data reduction section has its own LOG. They are overwritten when a new task is executed. You can also access them through *View* in THELI's menu bar

```
theli_pre_comm.sh.log (Preparation)  
theli_cal_comm.sh.log (Calibration)  
theli_bg_comm.sh.log (Superflattening)  
theli_w_comm.sh.log (Weighting)  
theli_ap_comm.sh.log (Astrom/Photom)  
theli_cc_comm.sh.log (Coaddition)
```

Other LOGs are

```
theli_color_convolve.sh.log  
theli_color_correction_avgwhite.sh.log  
theli_color_correction_photcat.sh.log  
theli_color_getseeing.sh.log  
theli_combima_comm.sh.log  
theli_create_chisquare_image.sh.log  
theli_create_privatestdcat.sh.log  
theli_delete_comm.sh.log  
theli_fits2tiff.sh.log  
theli_get_coadd_zp.log  
theli_getmaxoverlap.sh.log  
theli_gettiffpreview.sh.log  
theli_makemovie.sh.log  
theli_psf_plot.sh.log  
theli_psf_analysis.sh.log  
theli_statistics.log
```


- **NFRAMES**: This is the number of chips that can be processed simultaneously in memory without causing swapping. This number is automatically and conservatively estimated based on the physically available RAM, the number of CPUs selected, and the image size. Should your computer become sticky, reduce this parameter, or stop other memory-intensive applications.

You cannot select a value larger than the one determined automatically.

Parallelisation

THELI offers a convenient parallelisation scheme for multi-core CPUs: simply choose the number of CPUs you want. It can be changed any time while THELI is not reducing data. Each CPU will be assigned one or several chips. For example, if your camera has 16 detectors and you work with 4 cores, the detectors will be distributed by THELI's parallel manager in the following way:

#CPU	works on chips
1	1 5 9 13
2	2 6 10 14
3	3 7 11 15
4	4 8 12 16

Ideally, the number of detectors should be an integer multiple of the number of chips. In the previous example there would be no gain in speed if you used 5 CPUs instead of 4, as 1 out of the 5 CPUs would still have to process data from 4 chips. The other CPUs will finish sooner, but you have to wait until CPU #1 has finished too:

#CPU	works on chips
1	1 6 11 16
2	2 7 12
3	3 8 13
4	4 9 14
5	5 10 15

If you reduce single-chip data, you can still select more than one CPU. You will in general not gain speed as the additional CPUs will simply be idle. The only three exceptions are astrometry using *Scamp*, sky subtraction, and the final coaddition step with *Swrap*, as these are multi-threaded.

Note: If all your data of a multi-chip camera is on the same disk, it does not make much sense to use more than 4 CPUs at the same time. The I/O load on that disk becomes very large when all 4 processes read and write to it, thus the gain in speed is negligible (alternative: *linked file systems*)

Exception (1): You run a suitable RAID configuration that allows for high performance I/O, and/or have a fast SSD.

Exception (2): Processes that have low I/O rates or high CPU demand can efficiently make use of all CPUs available. These are:

- Catalog creation
- Astrometry (*scamp*)
- Sky subtraction
- Coaddition

2.6.3 Data directories

It is the user's responsibility to sort data into calibration files and target observations. THELI will not check if e.g. a flat field is actually a flat field.

THELI does not rely on the information given in the FITS headers of raw data. Even though FITS is supposed to be a standard, not everyone sticks to it. In many cases FITS headers are essentially incomplete, keywords contradict each other, or information in them is wrong. Therefore you have to manually sort your exposures into a directory tree which you then declare in THELI. If you KNOW that your headers are reliable, THELI offers you a little support in the *preparation section*.

Directory tree

This is where you tell THELI where the data are. THELI requires all data to lie in some **main directory**, being distributed in individual sub-directories for biases, darks, flats and the exposures of your target (hereafter: *science*). Standard star exposures and observations of a blank sky field must be copied to separate sub-directories as well.

Example 1: You observed NGC 1234 in a single filter, then your data directory tree should look like:

```
/thats/the/mainpath/
/thats/the/mainpath/BIAS
/thats/the/mainpath/FLAT
/thats/the/mainpath/SCIENCE
```

and you would enter in THELI:

```
Main path: /thats/the/mainpath
Bias: BIAS
Flat: FLAT
Science: SCIENCE
```

Example 2: You observed NGC 1234 and a standard star in filters B, V and R:

```
/thats/the/mainpath/
/thats/the/mainpath/BIAS
/thats/the/mainpath/FLAT_B
/thats/the/mainpath/FLAT_V
/thats/the/mainpath/FLAT_R
/thats/the/mainpath/NGC1234_B
/thats/the/mainpath/NGC1234_V
/thats/the/mainpath/NGC1234_R
/thats/the/mainpath/STD_B
/thats/the/mainpath/STD_V
/thats/the/mainpath/STD_R
```

In THELI you would start with

```
Main path: /thats/the/mainpath
Bias: BIAS
Flat: FLAT_B
Science: NGC1234_B
Standard: STD_B
```

Once finished, switch to the next filter. See the Section about multi-colour data processing for more details.

Warning: When clicking upon the *Restore ORIG* button, all data in the corresponding sub-directory will be deleted, and the original raw images will be restored.

Fully process SKY data

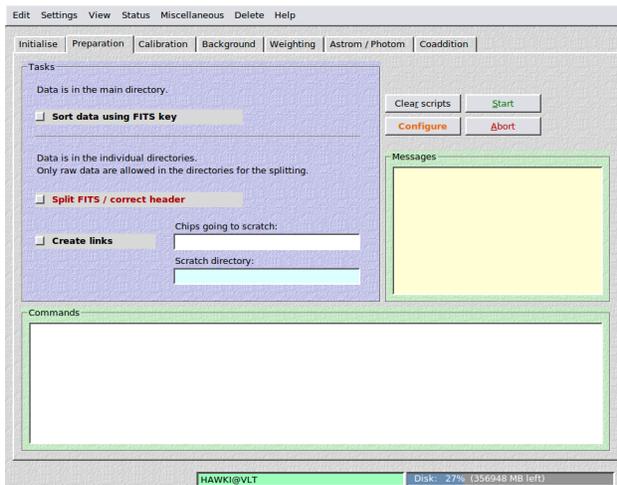
In case of extended targets one often needs observations of a blank sky field (hereafter: *SKY*) to perform proper background correction. THELI will process the *SKY* data automatically when present up to the point when it is no

longer needed, i.e. once the correction images calculated from it were applied.

However, in particular in the near-IR the total exposure time for the SKY field is often equal to that of the primary target. In this case it makes sense to have THELI continue processing the SKY up to and including the coaddition. The extra field may either double your field of view, provide some important calibration statistics, or just be useful for one of your colleagues. If you want this full scale processing, then activate this setting.

2.7 PREPARATION

This section is focused on preparing the raw data such that THELI can digest it. You can automatically sort the raw data, have to modify the fits headers, and in case of multi-chip cameras split the images into individual chips. Lastly, multi-chip data can be redistributed to different hard disks for faster processing.

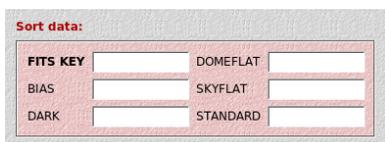


2.7.1 Sort data using FITS key

If the scientific data and their corresponding calibration files are all mixed up, then you can attempt to identify different image types based on the value of a certain FITS keyword, e.g. the *OBJECT* key.

All images belonging to one category will be collected in a correspondingly named sub-directory. All other images that remain unmatched are assumed to be target observations and will be put in a sub-directory called *SCIENCE*.

Parameters



- **FITS KEY:** The keyword that identifies the data type
- **BIAS:** The string for a bias
- **DARK:** The string for a dark
- **DOMEFLAT:** The string for a domeflat

- **SKYFLAT:** The string for a skyflat
- **STANDARD:** The string for a standard star field

Check the sorting

After the data have been sorted, you should go into the individual sub-directories and run the following command in a console to verify data consistency (possibly adjusting the FITS keyword names)

```
dfits *.fits | fitsort EXPTIME FILTER OBJECT DATE-OBS
```

In this way you can easily see if e.g. a flat-field slipped into your biases, or a bias into your *SCIENCE* directory.

Another very useful command to check for outliers or bad images in the BIAS, DARK and FLAT directories is to look at basic image statistics:

```
imstats *.fits
```

If you cannot find these binaries, then you did not include THELI's binary directory in your PATH variable (*details*).

Note: No sorting with respect to various filters is done. If you e.g. took flat fields in UBVRI filters, then all these flats will be put into the same FLAT directory. You must redistribute them manually.

Warning: This task relies on the integrity of the chosen FITS header keyword. If a flat is accidentally identified as a bias, then it will be put in the BIAS sub-directory etc. No further automatic consistency checks are performed by THELI. You still have to make sure that images end up in the correct sub-directories (see above for some helpful commands).

2.7.2 Split FITS / correct header

The main job of this task is to split multi-extension FITS files into the single chips, thus allowing for parallel processing. It also writes a new FITS header conform with the THELI pipeline. If single-chip images are provided, only the FITS header will be updated. Optionally, a crosstalk correction can be applied, and the image can be renamed.

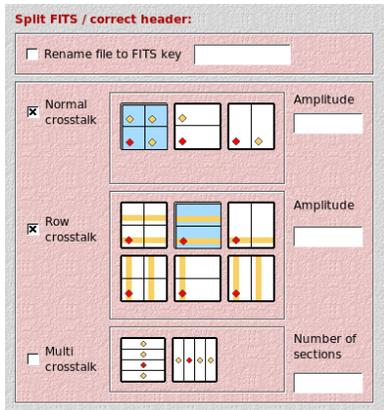
This task will be applied to all sub-directories that have been specified in *INITIALISE*. If a master calibration file (bias, flat or dark) is found, then the corresponding directory will be skipped. THELI will try to correct all other files present. If a FITS image is found that was run through this step before (judging from certain header information), then it will be skipped.

At this point THELI will also replace the following characters in the file names and in the FILTER header keyword with underscores ('_'):

```
$ , | " ' ; [ ] ! % / ( )
```

Warning: Only uncompressed FITS images, and no other files, should be present in the directory in which you attempt splitting. Otherwise an error will be returned, which may or may not have further consequences during processing.

Parameters



- **Rename file to FITS key:** The raw images can optionally be renamed to the string value of an arbitrary FITS keyword. This is useful if the software running at the telescope does not assign an unambiguous name to the file, but only to a certain header keyword.
- **Normal crosstalk:** Typical for CCDs and near-IR HAWAII detector arrays with several readout quadrants. Bright sources cause a ghost image in the other readout quadrants. You must specify the amplitude of the effect.
- **Row crosstalk:** Typical for CCDs and near-IR HAWAII detector arrays with one or several readout quadrants. The rows or columns going through a bright source have an enhanced background level (banding). The same rows in other readout quadrants can be affected, too. You must specify the amplitude of the effect.
- **Multi crosstalk:** Typical for HAWAII-2 detectors with their larger number of readout stripes. Bright sources leave ghost images in the other readout sections. You must specify the number of readouts. This task will calculate the median of all readouts and subtract it from each readout.

Crosstalk amplitudes can range from very small (0.0001) to very large (a few percent of the brightness of the source). THELI makes no assumptions about the crosstalk amplitude for the various instruments, which means you have to figure it out yourself through trial-and-error. When the crosstalk becomes invisible in the coadded image, then you got it right.

Note: The crosstalk amplitude must be a negative number if bright sources cause dark crosstalk features.

Note: The crosstalk correction assumes that the crosstalk amplitude is the same for all readout quadrants. This is commonly the case for HAWAII-1 detectors with 4 readout sections, but can be different for HAWAII-2 arrays with multiple readouts. In the latter case crosstalk correction will be incomplete with THELI.

Note: Crosstalk correction across different physical CCDs (e.g. a ghost image of one detector shows up in another) is not yet supported in THELI.

2.7.3 Create links

This section is **not** of interest for you if

- you reduce data from a camera with less than 2-4 CCDs
- you use 4 CPUs or less and have a reasonably fast hard disk
- your disks are in a RAID configuration appearing as one single disk
- you are a beginner with THELI or not familiar with how the Linux file-system works

How it works

If you work on a multi-core machine with several hard disks that you can address explicitly, then you can fully exploit THELI's *parallel mode*.

The working principle is that different chips are processed by different CPUs on different hard disks, minimising I/O traffic on a single hard disk. This task will distribute the images taken with different CCDs to separate hard disks. At the same time a link structure is established that makes all the data appear to reside within the normal THELI *directory tree*.

What you have to do

Chips going to scratch: In this field provide a blank separated list of the chips you want to move to one scratch directory. For example, you are working on a camera with 8 CCDs using a machine with 4 CPUs and 4 hard disks. Then you would enter in this field the string "1 5" (without the quotes).

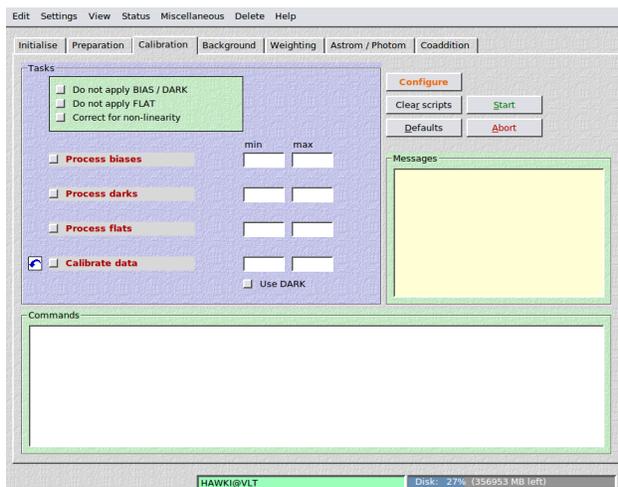
Scratch directory: This is the directory on the first hard disk where chips 1 and 5 go. You do not need to establish a directory tree in those scratch directories. All files end up in the same place, but are uniquely identified through the link system pointing to the various categories (bias, flat etc).

THELI will loop over all data sub-directories found in the *Initialise* section. Once done with chips 1 and 5, repeat the same for chips 2 and 6, and so on.

Note: The linked file system must be created after *splitting* the data.

2.8 CALIBRATION

In this section bias-, dark- and flatfield-correction take place.



2.8.1 Process BIASes / Process DARKs

Bias and dark images are overscan corrected and median combined without rescaling. At least 3 exposures are needed in order to reject spurious events such as cosmics. We recommend at least 10 exposures in order to keep the noise low.

The master bias and dark have the same name as the directory in which the individual frames reside, for example:

```
MYBIAS/MYBIAS_1.fits
MYDARK/MYDARK_1.fits
```

Options

Specify a lower and upper threshold in the **min** and **max** fields. Exposures with a statistical mode outside this interval are rejected from the processing and are moved to a *BADMODE* sub-directory.

Parameters

Rejecting nlow and nhigh pixels from a stack

	nlow	nhigh
Overscan	<input type="text" value="0"/>	<input type="text" value="1"/>
Bias	<input type="text" value="0"/>	<input type="text" value="1"/>
Dark	<input type="text" value="0"/>	<input type="text" value="1"/>
Flat	<input type="text" value="0"/>	<input type="text" value="1"/>

Rejecting nlow and nhigh pixels from a stack

	nlow	nhigh
Overscan	<input type="text" value="0"/>	<input type="text" value="1"/>
Bias	<input type="text" value="0"/>	<input type="text" value="1"/>
Dark	<input type="text" value="0"/>	<input type="text" value="1"/>
Flat	<input type="text" value="0"/>	<input type="text" value="1"/>

- **nlow**: The number of lowest pixels in the stack (overscan and bias) to be rejected. Default: 0
- **nhigh**: The number of highest pixels in the stack (overscan and bias) to be rejected. Default: 1

2.8.2 Process flats

Flat-field exposures are overscan corrected, have the master bias subtracted, are rescaled to the same statistical mode, and then median combined. Different gains in multi-chip cameras are preserved, thus after applying the master flat, gain differences in the target exposures are corrected for. At least 3 flat field exposures are needed, but we recommend to go for 5-10. The illumination level of individual exposures should stay clear from the on-set of any non-linearities.

THELI automatically subtracts a master bias from the master flat. If you do not have a master bias or do not want this correction to take place, then set the *Do not apply bias* switch at the top of the *Calibration* section.

The master flat has the same name as the directory in which the individual flat frames reside, for example:

```
MYFLAT/MYFLAT_1.fits
```

If the user also specified a DARK FLAT (*Flat off* in the *Initialise* section), i.e. a flat field with low illumination, then these exposures will be treated in the same way as the normal flat field. The resulting combined dark flat field will be subtracted from the bright flat, the name of the final flat field will be the same as the one without this correction. The original bright flat can still be found in the normal FLAT directory. This approach is commonly used for near-IR observations to remove thermal background from the flat field data. It depends on the instrument and the way the dark flats were taken whether this makes improves the data (more homogeneous photometric zeropoint) or not.

Options

Specify a lower and upper threshold in the **min** and **max** fields. Exposures with a statistical mode outside this interval are rejected from the processing and are moved to a *BADMODE* sub-directory. This is useful to reject under- or overexposed flats.

Parameters

Rejecting nlow and nhigh pixels from a stack

	nlow	nhigh
Overscan	<input type="text" value="0"/>	<input type="text" value="1"/>
Bias	<input type="text" value="0"/>	<input type="text" value="1"/>
Dark	<input type="text" value="0"/>	<input type="text" value="1"/>
Flat	<input type="text" value="0"/>	<input type="text" value="1"/>

- **nlow**: The number of lowest pixels in the stack (overscan and flat) to be rejected. Default: 0
- **nhigh**: The number of highest pixels in the stack (overscan and flat) to be rejected. Default: 1

2.8.3 Calibrate data

Images are overscan corrected, debiased (or dark-subtracted) and flat fielded.

Filename extension: After running through this step, images have the string *OFC* in their file name between the chip number and the suffix, e.g.:

NGC1234_1OFC.fits

Options

Specify a lower and upper threshold in the **min** and **max** fields. Exposures with a mode outside this interval are rejected from the processing and are moved to a *BADMODE* sub-directory. If left empty, no thresholding will be done.

Do not apply BIAS / DARK: If you work with near-IR data, you probably do not want to subtract a master bias as the so-called pre-read has already been subtracted. If you do not have a master bias or master dark, you also have to activate this switch.

Do not apply FLAT: Activate this switch if you do not want to calibrate your target exposures by a flat field. Mid-IR observers possibly want to do that, or e.g. if you are feeding THELI some pre-processed data from another instrument.

Use DARK: If you want to subtract a master DARK instead of the master BIAS, then indicate this here.

Dark rescaling

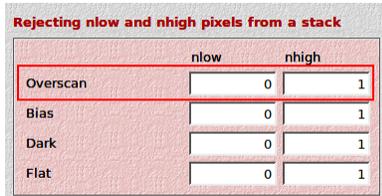
If you subtract a master dark from the data, then THELI looks up the dark's exposure time. If the latter deviates by more than 1% from the exposure time of the SCIENCE images, then the master dark is rescaled accordingly. Thus you can subtract a high S/N masterdark with longer exposure time, introducing less noise in the calibrated target frames.

If your camera does not have an overscan, then the bias offset is still present in the master dark. The offset would be rescaled, hence adding a constant to the sky background of the target images. While this as such is not critical because the background will be subtracted anyway, it would modify the sky background statistics (which is not used by THELI, but you might be interested in it). To avoid this, the mode of the master dark is subtracted before a rescaling takes

place. For cameras with overscan this correction has no effect as the mode is essentially zero after overscan correction took place.

Note: THELI does not perform a temperature rescaling of the master dark. If the dark current in your camera is temperature dependent, then make sure the darks are taken with identical settings. You can also attempt to *manually rescale the dark frames* in a trial and error manner.

Calibration parameters



	nlow	nhigh
Overscan	0	1
Bias	0	1
Dark	0	1
Flat	0	1

- **nlow:** The number of lowest pixels in the overscan stack to be rejected. Default: 0
- **nhigh:** The number of highest pixels in the overscan stack to be rejected. Default: 1

2.8.4 Non-linearity correction

Non-linearity usually arises when pixels approach their saturation limit and do not response in a linear manner anymore to the incoming flux. In some cases, e.g. near-IR detectors, non-linearities can also appear at very low exposure levels.

Implementation for new instruments

THELI can correct for the detector non-linearity provided that corresponding correction polynomials are known. Currently, this is only being implemented for WFC@INT and SPARTAN@SOAR. The polynomial is assumed to be of 3rd order and has the form

$$a_0 + a_1 x + a_2 x^2 + a_3 x^3$$

If you know the correction polynomial for an instrument called *SOME_INSTRUMENT* in THELI, then add the polynomial coefficients to a file *SOME_INSTRUMENT.nldat* and save it in the same place as *SOME_INSTRUMENT.ini*, i.e. in one of

```
/home/user/THELI/gui/scripts/instruments_professional/  
/home/user/.theli/instruments_user/
```

Each chip of the camera is represented in *SOME_INSTRUMENT.nldat* by a single line containing the four polynomial coefficients

$$a_0 \ a_1 \ a_2 \ a_3$$

Unknown coefficients must be set explicitly to zero. For example, if the correction polynomial is of 2nd order, then you must set $a_3 = 0.0$.

How THELI does it

In order to fully correct for non-linearity, one must :

1. Subtract the master bias from the flat exposures
2. Apply the non-linearity correction to the debiased flat fields
3. Create the master flat
4. Subtract the bias/dark from the science data
5. Apply the non-linearity correction to the debiased science data
6. Apply the corrected flat field

This processing scheme is significantly different from the standard data flow and requires the creation of intermediate data. THELI will do all of this automatically and invisibly for you. You just have to activate the *Non-linearity correction* switch and proceed as normal. The only difference is a significantly longer execution time.

2.9 About background modelling (superflattening, defringing, near-IR sky models)

This section provides some important information about background modelling. Its content must be understood in order to perform data reduction correctly. The THELI documentation makes use of the nomenclature and concepts presented here.

In particular in wide-field imaging, and/or observations at longer wavelengths, images will not appear flat after flat-fielding. Reasons are manifold:

- improper illumination or size of the domeflat screen
- scattered light in the dome while taking twilight flat fields
- moonlight
- differential airglow at higher airmass
- variable sky background (in particular in the near-infrared)
- sky concentration (reflection between CCD and corrector lenses)
- fringing
- pixel scale variations
- etc

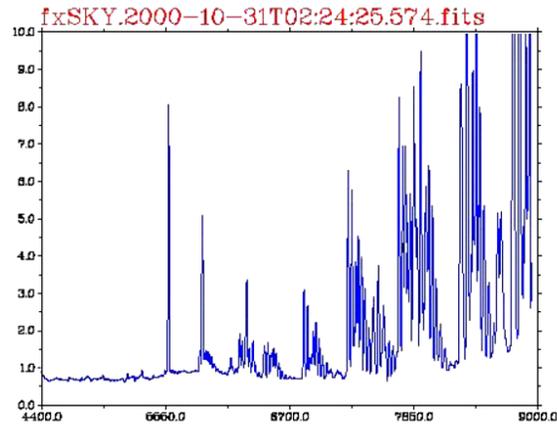
Apart from fringing, these effects show up as low frequency spatial variations of the background level. Amplitudes may reach or exceed 10% of the background level. A background model is required to correct for these effects.

2.9.1 Airglow

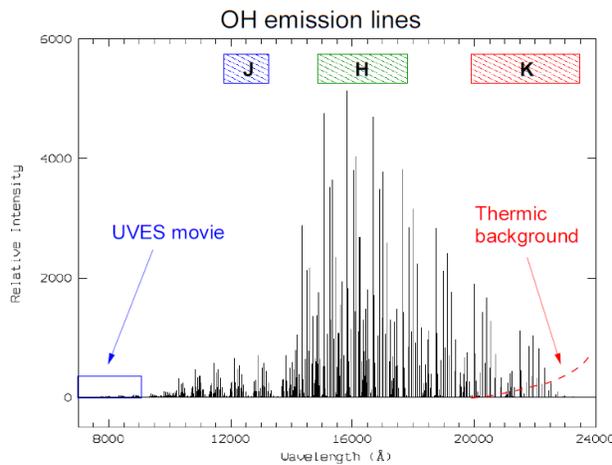
Temporal variations

The higher atmosphere contains OH^- molecules which are excited during daytime and slowly decay at night. The spectrum of this *airglow* is characterised by clusters of numerous emission lines, starting at around 600nm and beyond.

Movie of optical airglow spectra taken with UVES/VLT (credit: [Ferdinando Patat, ESO](#))



The airglow's intensity, temporal and spatial variability increase significantly with wavelength, dominating the near-IR sky:



The spectral range covered by the movie of the optical airglow is indicated by the small blue box to the lower left. Notice how much brighter the near-IR sky is.

Fringing

CCDs are made of silicon, which becomes transparent with increasing wavelength, starting at around 700nm. In back-illuminated, thinned CCDs airglow can cause internal reflections leading to a characteristic interference pattern in the background of an image, known as *fringing*. The pattern itself depends mainly on the thickness of the CCD and is thus essentially unchanged over the lifetime of the detector. However, the amplitude of the fringing pattern varies as the airglow spectrum at a given position on the sky changes with time. Fringing is also present in near-IR detectors.

In *R*-band this is usually not very critical and one fringe model will suffice for an entire night. In *I*-band things are more critical and in the worst case a new fringing model is needed every 30 minutes. In the near-IR the sky changes over time scales of minutes, and on angular scales of arcminutes, requiring new background models about every 5-30 minutes. In the mid-IR we talk about seconds. Fringing is less important in these cases, more problematic are the sky background variations themselves:

This is a movie of the near-IR H-band airglow over 1.5 hours. The field of view is 9 degrees, 1 second in the movie corresponds to 7 minutes real time. The mean background level has already been removed, shown are only the

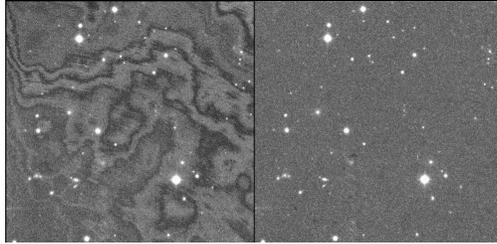


Figure 2.1: Fringing in *I*-band before and after correction.



variations around the mean. Hence reality is even worse than depicted here. (Credit: 2MASS wide field airglow experiment).

2.9.2 How to create a background model

A background model can be created from a set of dithered exposures of a non-crowded (or empty) field. The dither pattern has to be significantly larger than the largest object in the field. Extended objects must be absent or otherwise the correction image will be wrong.

THELI applies several steps in order to create a background model. First, all objects in an exposure are detected and masked, such their flux will not bias the result. The masked images are rescaled to the same mode and then median combined. If all goes well, the background model contains only background signal. This can be a pretty featureless image, or it can show significant large-scale variations and/or small-scale fringing.

The high-frequency fringing component can be extracted from the background model by smoothing the latter with a large (several hundred pixels wide) kernel, and subtract the smoothed image from the unsmoothed model. This is often used for optical data to separate a multiplicative illumination correction from the additive fringing component. In the near-IR, the original background model would be subtracted from the data.

Static background models

A *static background model* is created from a fixed set of exposures, and applied to all images in the data set. This will only work reliably well if the background is stable during the night or at least over longer stretches of time.

The sky becomes more unstable with increasing wavelength. In *R*-band the fringing is usually low and varies slowly, thus one static model (for fringe correction) usually suffices for the entire night. For the *I*-band this is also possible if conditions were stable, or if target observations did not extend over more than 30-60 minutes in bad conditions.

Dynamic background models

In the near-IR, static models can only be used if target observations do not exceed 5-30 minutes. Otherwise a *dynamic background model* is required, i.e. the correction image is calculated from a few exposures that were taken immediately before and/or after the image that is to be corrected. One would like to make this window as large as possible in order to improve the S/N in the correction image, but it has to be kept as small as possible in order to still properly sample the temporal sky variations. This can require some trial- and error in data reduction. Typical window sizes are 4-8, i.e. these many images closest in time to the exposure being corrected are used to create the model. The computational overhead for a dynamic model is significantly larger than for a static model.

2.9.3 Superflat or sky background model?

Background variations can have additive or multiplicate causes, such as scattered light (additive) or an incorrect flat field (multiplicative; which might be caused by an additive effect, such as scattered light or sky concentration).

A good background model will always result in a flat image, no matter how it is applied. However, if you accidentally correct a multiplicative effect by subtraction, then the photometric zeropoint will no longer be constant across the field. Likewise, an exposure with constant zeropoint may not have a constant sky background, which must be subtracted afterwards. If the goal is to achieve a flat background and photometry isn't important for your work, then you most likely do not care about all of this. However, if consistent photometry matters, you are in for a good deal of trouble, as multiplicative and additive components are hard to distinguish, in particular when they occur simultaneously.

- **Multiplicative character:** The background model falls in this category if the flatfield wasn't working correctly. This is very often the case with wide field imagers, leaving residuals on the level of 1-2 percent. Reasons can

be manifold, e.g. the flat screen is insufficiently illuminated or too small, sky concentration, or else. In this case the data must be divided.

- **Additive character:** Even if the flatfield itself is perfectly fine, it happens that the flatfielded data still exhibits significant background variations. In the optical this is usually the case due to scattered light finding its way onto the detector. In the near-infrared, the dominant component is the sky background itself (airglow). In these cases the background model has to be subtracted from the data.

If one erroneously subtracts a background model instead of dividing by it, or vice versa, the resulting background will still be perfectly flat, but the photometric zeropoint will not be constant anymore (variations of up to 10% are not uncommon).

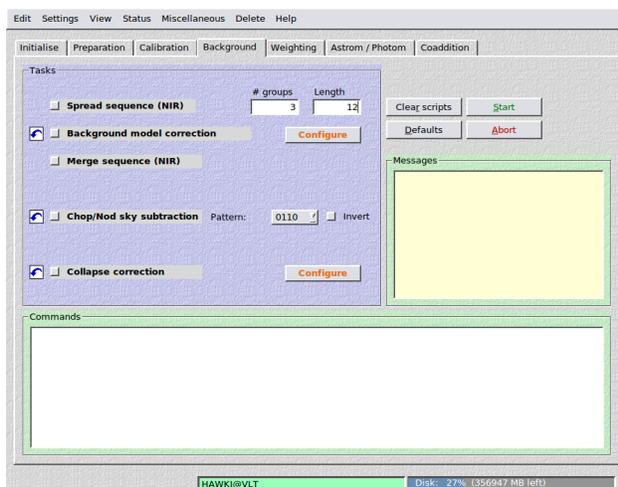
Interpreting the background structures sometimes helps to identify the correct approach, but it can be misleading. For example, flat field effects are often radially symmetric whereas scattered light can cause one side of the image to be brighter than the other if caused e.g. by moonlight. It may also be symmetric, e.g. when caused by light being reflected from the detector to a field corrector lens, and from there back again to the detector. In such cases one would have to correct the flat field for scattered light first, but for the observer this is usually not possible as it requires extensive testing at the telescope.

If both additive and multiplicative effects are mixed, there is only way to fully correct for it: one must compare the data with photometric reference data from the same field and in this way establish a map of the zeropoint variations, from which a (multiplicative) correction image can be calculated. This requires a comparatively high density of reference objects and a good photometric catalogue of your field in the same filter (if your field is covered e.g. by SDSS, then you have a good chance to correct for it). Alternatively, you could observe an external standard star field with excessive dithering and calculate a correction image which you apply to your data, assuming that the effect is stable. After this multiplicative correction, residual background variations can be modelled and subtracted individually.

2.10 BACKGROUND

This section deals with correction of background variations. This can either be superflating, defringing or subtraction of a background model. For mid-IR data a suitable chop-nod sky subtraction is offered. Lastly, in order to remove typical effects such as reset anomalies in near-IR cameras, a collapse correction can be performed.

If you are unfamiliar with background models, then please read this [background information](#) first.



2.10.1 Spread sequence (NIR)

LIRIS@WHT is the only near-IR detector I'm aware of that needs this treatment. So most of you can simply skip this section. Nevertheless, I'm going in some details in case it crops up in other instruments (likely using older HAWAII-1 arrays).

Resetting and image equilibrium

Near-IR detectors have very different properties than the CCDs used in optical instruments. While they are not exposed, they are continuously resetted and rapidly reach a *resetting equilibrium*. On the other hand, when a series of exposures is made, the detector reaches an *imaging equilibrium*. I have made up these two terms myself, so if you know a more precise or commonly used wording, then please let me know.

If you have dithered observations with LIRIS@WHT, and a series of e.g. 10 exposures was taken per dither point, then the detector will have reached the imaging equilibrium with the third exposure. All subsequent images will have the same (instrumental) background. While the telescope acquires the next dither position, the detector goes into the resetting equilibrium again and thus is in the same state as when you started the exposure series at the first dither position. Therefore, the image backgrounds are the same for the i -th exposures in the n -th sequence (apart from slower sky background variations). The consequence is that the background modelling has to be done separately for these exposures. This script sorts the data in according directories where it is then processed automatically in the right manner.

Warning: It is essential that you do not mix sequences with different lengths in the *SCIENCE* directory. THELI assumes that only complete sequences of the same length are present.

Example

You have 12 images per dither point. From subtracting one image from the next in the sequence, you found out that the detector settled into its imaging equilibrium starting with the third exposure. You would then enter 3 into the field **# groups**, and 12 into the one labelled **Length**.

The script will create three directories next to the *SCIENCE* directory, and redistribute the exposures in the following manner:

```
* SCIENCE_S1: exposures 1, 13, 25, ...
* SCIENCE_S2: exposures 2, 14, 26, ...
* SCIENCE_S3: exposures 3-12, 15-24, 27-36, ...
```

The background modelling must be done separately for each of these directories. THELI will do this automatically for all *SCIENCE_Si* directories and, if applicable, *SKY_Si* directories. You do not have to make any changes in the *directory tree* defined in the *Initialise* section.

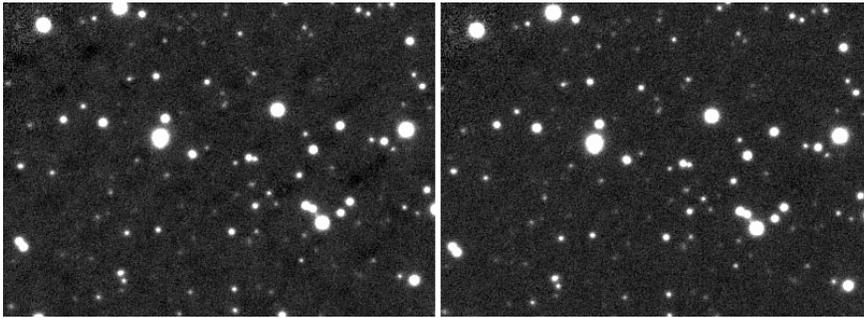
All other processing tasks do not (and do not have to) loop over the *SCIENCE_Si* directories. THELI will know that it has to do this loop if the *Spread sequence (NIR)* task has either been executed, or is activated.

After the background has been corrected, you merge the images again into the previous *SCIENCE* directory. This is done further below (see the *Merge sequence (NIR)* task). SKY fields are processed automatically after the *SCIENCE* data have been dealt with.

2.10.2 Background model correction

This task will create the background model, and apply it in a variety of ways, depending on whether you process optical or near-IR data. The resulting images may have suitably flat background after a first pass, but faint and

unmasked haloes will have contributed to the background model and then show up as darker patches, reflecting the dither pattern (see the left panel below):



To correct for this effect, simply run this task a second time with improved detection parameters (e.g. $DT=1.5$, $DMIN=10$), as is commonly done in the near-IR. It will create much better masks (because the objects are visible much more clearly). The improved result can be seen in the right panel above. A *collapse correction* may be run between the two passes to get rid of final residuals.

Filename extension: Images will have the character *B* appended to their filenames, e.g.

```
NGC1234_1OFCB.fits
```

THELI will recognise when the task is run a second time. In this case, the object detection will take place on the temporary *OFCB* images created in the first pass, and the corresponding pixels will be masked in the original, flat-fielded only, *OFC* images. The newly created mask images from these data are then also applied to the *OFC* images, resulting in a new batch of *OFCB* exposures. The previous ones will be kept in a *OFCB_IMAGES_IPASS* directory for comparison.

Parameters

Background modelling in THELI is highly configurable, to match the different demands and characteristics optical and near-IR data may have. Four main parameter groups are available to control this task, described in detail below.

Background model:

(1) Mask objects so that they cannot contribute to the background model (optional)

DT DMIN Median FILTER

(2) Reject pixels from the stack

nlow nhigh

0 3

(3) How to apply the background model

Subtract model

Smoothing scale for the background model:

Smoothing scale for the fringing model:

Rescale model

Apply model to SKY data, too (needed for 2-pass)

Adjust gains between chips

(4) Static or dynamic model

Window size: Model is: STATIC

Max gap size (hours) 1.0

1. Mask objects so that they cannot contribute to the background model (optional):

- **DT:** This is the *SExtractor* detection threshold per pixel, given in units of sigma of the sky background noise. If you leave it empty, no object masking will take place.

- **DMIN:** The minimum number of connected pixels above the detection threshold making up an object. The smaller *DT* and *DMIN*, the fainter the objects masked. If you leave *DMIN* empty, no object masking will take place.
- **Median/mean:** The method used for combining images. The median delivers a more stable result for a small number of stacked images, whereas the mean has lower noise when more images are stacked, but is susceptible to outliers. The default setting is median.
- **Filter:** If a lot of hot pixels are present in the data, then consider to switch off this option. Filtering increases the footprint of a pixel and thus the fraction of the masked area can become unacceptably large, e.g. for some particularly bad HAWAII-2 arrays.

The final mask images are stored in a separate *MASK_IMAGES* directory where they are available for inspection. If masking is switched off, then you will simply find links to the unmasked images from which the background model will be created. If masking takes place, then you should inspect the mask images to make sure that masking is not too aggressive. The latter would be the case if not only objects, but also sky features are masked.

2. Reject pixels from the stack:

- **nlow:** The number of lowest pixels in the stack to be rejected. Default: 0
- **nhigh:** The number of highest pixels in the stack to be rejected. Default: 1. Therefore, even if you decided to not do an object detection (leaving *DT* and *DMIN* empty), you will get a rather clean background model.

3. How to apply the background model:

This is the most important part of the background modelling process. Here you decide how the model is applied. Whatever suits your optical or near-IR data best. Note that individual images may still present some background residuals afterwards, which have to be removed using *sky subtraction*. The following options and settings are available:

- **Subtract model:** This is your choice for near-IR data. The background model will by default be rescaled (see below) to match the illumination of the individual images, and then subtracted. Possible fringes in the data are removed as well.
- **Divide model:** This corresponds to classical superflattening, often used for optical data. Note that background gradients in optical data are often caused by differential *airglow* or scattered light, and should therefore be subtracted, not divided.
- **Subtract fringes:** In this mode the background model will be smoothed with a large kernel (usually several 100 pixels, see below), resulting in an *illumination correction (superflat)*. The smoothed model is subtracted from the unsmoothed model, yielding the desired *fringing model* with the higher frequency variations. The fringing model will be rescaled to match the variations in the science image, and subtracted. This mode would be applied to optical data when fringe correction is desired, but no superflattening.
- **Divide smoothed model, subtract fringes:** Same as the previous, but in addition to defringing the science image gets divided by the normalised *illumination correction*, i.e. a superflat is applied.
- **Create masks only, do not apply them:** If this option is chosen, only the masks will be created. Change the object detection parameters *DT* and *DMIN* and rerun the task until you are satisfied, then choose one of the four other modes.

These modes can be fine-tuned by the following settings:

- **Smoothing scale for the background model:** This would be a several 100 pixel wide kernel to smooth the background model. Required if you want to subtract fringes, or divide by a smoothed superflat. Not recommended to be combined with near-IR data and the *Subtract model* mode, as it will leave systematic high frequency features common in HAWAII-2 detectors in the data.
- **Smoothing scale for the fringing model:** This is generally a very small kernel (1-3 pixels) for a median filter of the fringing model, resulting in less noise in the correction image. Comes in handy when defringing

optical data, such as *r*- or *i*-band images with strong fringing. A value of 1 (2,3...) means that pixels in a 1 (2,3...) pixel wide border (i.e. the 3x3 (5x5,7x7...) superpixel) are taken into account.

- **Rescale model:** This will rescale the background / fringing model, which is calculated from several dithered exposures, to the illumination level of the image to be corrected. Affects the 1st, 3rd and 4th mode listed above (where subtraction occurs). The assumption is that the background variations and fringing are caused by *airglow*, and that the amplitudes of the variations scale with the strength of the airglow. This should therefore always be switched on, unless this assumption does not hold anymore, for example when a series of exposures is affected by changing twilight, zodiacal or lunar contributions.
- **Apply model to SKY data, too:** If blank sky exposures are needed to create a background model, then a single application of the model to the science data may be sufficient. If, however, refined object masks are needed, then the SKY observations have to be corrected themselves before the objects therein can be identified and masked properly. If you need such a 2-pass correction, then you have to have this box switched on. No harm is done to leave that box always on. In case SKY data are present, the overall execution time will simply be increased.
- **Adjust gains between chips:** This is for multi-chip cameras only. Gain differences between chips are corrected for during flat-fielding. If, however, the flat fields are old, or the electronic of the camera instable, residual gain variations may occur. They can be corrected for using one of the two multiplicative modes (*Divide...*). Usually, this should not be necessary.

4. Static or dynamic model:

- **Window size:** The number of images closest in time to the exposure that is to be corrected. If left empty or set to zero, then a static model will be used, otherwise the model will be dynamic.
- **Max gap size:** The maximum amount of time, in hours, by which a series of exposures may be interrupted. If gaps longer than this parameter divide the sequence, then separate sky models (either dynamic or static) will be calculated for each block of exposures.

Note: THELI will automatically do the right thing when SKY data are present. Based on the modified Julian date (*MJD-OBS* keyword) the correct SKY images are selected also in dynamic mode. *MJD-OBS* header entries are created during splitting, using date and time of the observation if the *MJD* is not already present. This holds for all pre-defined cameras in THELI, but you may want to check that the MJDs for your data are correct.

2.10.3 Merge sequence (NIR)

At this point one can merge the images again, given one has run the *spread sequence* task before. It will read the *#groups* parameter from the *spread sequence* task. The fully calibrated exposures in the *SCIENCE_Si* directories are merged again in the original *SCIENCE* directory.

2.10.4 Chop/nod sky subtraction

This is for mid-infrared data, only. THELI assumes that all science observations, i.e. on-target AND off-target, resume in the same directory, and that their alphanumerical order is equivalent to their temporal order.

You can choose from four different chop-nod patterns, where “1” represents an image with the target, and “0” an image of a blank sky area. If your target is very small the chop-nod pattern will not move it off the detector area, in which case “0” can be considered as another target observation. Image “0” will be subtracted from image “1” in a pairwise manner. The available patterns are:

- **0110:** 2nd minus 1st, 3rd minus 4th
- **1001:** 1st minus 2nd, 4th minus 3rd
- **0101:** 2nd minus 1st, 4th minus 3rd

- **1010:** 1st minus 2nd, 3rd minus 4th

THELI assumes that this pattern is repeated, i.e. for the pattern **0110** the sequence of exposures is:

0110-0110-0110-0110-...

Invert: If this switch is selected, every second group is reversed, i.e. for the pattern **0110** the sequence of exposures becomes:

0110-1001-0110-1001-...

Filename extension: Images will have the character *H* appended to their filenames, e.g.

`NGC1234_10FCH.fits`

Note: Images belonging to the “0” chop-nod positions are not present afterwards anymore. If your source is point-like and also on the detector for the “0” positions, then they will form a negative image.

2.10.5 Collapse correction

If your data exhibits horizontal or vertical linear gradients, such as a residual reset anomaly in near-infrared detectors, then use this task to get rid of them. It calculates an average row (or column, or both) from all rows (or columns, or both) and subtracts it from the latter. Objects are (optionally) masked before the average rows/columns are calculated. Typical reset anomalies look like this:

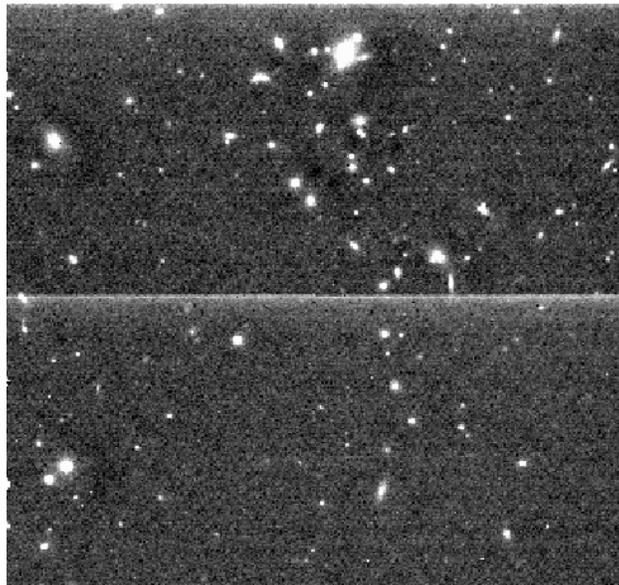


Figure 2.2: LIRIS@WHT reset anomaly

Parameters

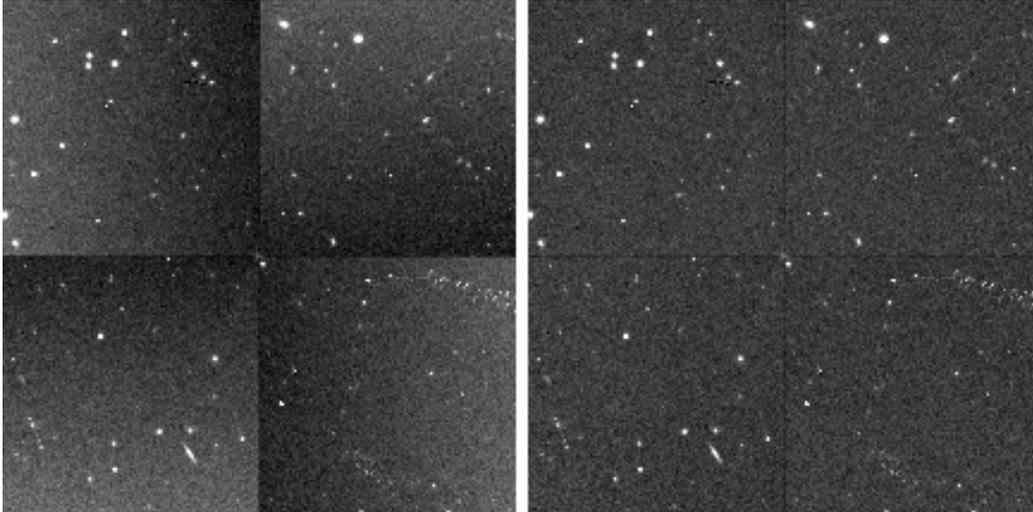


Figure 2.3: MOIRCS@SUBARU reset anomaly, before and after correction

Collapse correction

Mask objects so that they cannot contribute to the background model

DT | DMIN

Store object masks

Automatic revision of DT

Collapse parameters

Rejection threshold

Collapse direction 

Exclude this region

 xmin | xmax

ymin | ymax

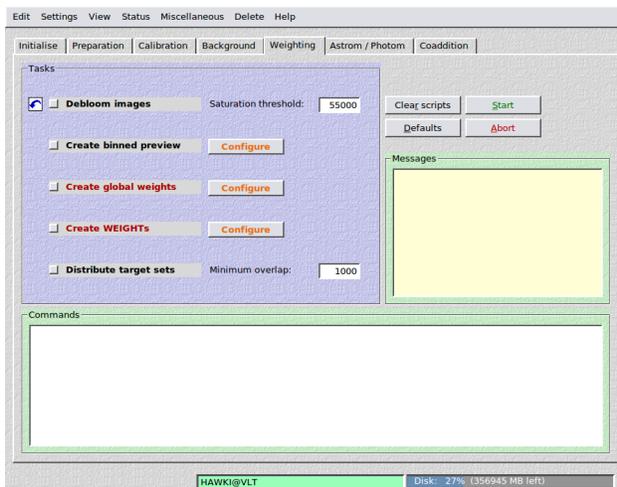
- **DT (optional):** The *SExtractor* detection threshold per pixel, given in units of sigma of the sky background noise. If left empty, no masking will take place.
- **DMIN (optional):** The minimum number of connected pixels above the detection threshold making up an object. The smaller *DT* and *DMIN*, the fainter the objects masked. If left empty, no masking will take place.
- **Store object masks:** On by default, so you can inspect the masks and fine-tune *DT* and *DMIN* if needed. Background features you want to correct must not be masked.
- **Automatic revision of DT:** Checks the value of *DT*. If a larger value is found to yield better results, then the user-supplied value is overridden. Can be tried for correcting an unstable reset anomaly, otherwise should be switched OFF.
- **Rejection threshold:** A kappa-sigma clipping is performed when calculating average rows/columns. This is the threshold in units of sigma. It should not be chosen too low (lower than 1.5 or 2), in particular if the distribution of background values is non-Gaussian with a bright tail. Check the resulting mask images in *MASK_IMAGES* to make sure the process is not too zealous.
- **Collapse direction:** If the brightened feature is horizontal (vertical), select **x** (**y**) as the collapse direction. THELI will calculate average columns (rows) in these cases. You can also subtract both horizontal and vertical lines in a single pass (**xy**). Some HAWAII-2 arrays feature 4 readout quadrants with readout directions rotated

by 90 degrees. In these cases you can use either **xyyx** or **yxyx**. A small graphical display will visualize the pattern chosen.

- **Exclude this region:** If an object with a faint extended halo is present in the data, then you must make sure that the halo does not contribute to the measurement area. The halo can be so faint that you do not see it in an individual exposure, and thus it will slightly bias the result. If you stack a large number of images (as is the case with near-IR data), then this can lead to a significant over-correction of the data, visible as a dark horizontal or vertical bar running through the extended source. To avoid this problem, you can define a region that is excluded from the calculation (you must define the left, right, lower and upper boundary in pixel coordinates).

2.11 WEIGHTING

This section focuses on the determination of individual weight images that will be used in the final coaddition. In addition, you can perform a cosmetic correction of blooming spikes, and obtain some binned preview images for quick visual inspection of your data set. Satellite trails and other bad areas can be masked manually.



2.11.1 Debloom images

This is a purely cosmetic task and should only be used if you plan to publish a pretty picture based on your data. You should not base a scientific analysis on images that underwent this task. It will replace all pixels with values above the saturation threshold with values estimated from the local neighbourhood, i.e. remove blooming spikes. The profiles of saturated stars are significantly corrupted afterwards.

Filename extension: Images will have the character *D* appended to their filenames, e.g.

```
NGC1234_10FCD.fits
```

Parameters

Saturation threshold: Pixels with values higher than this threshold will be interpolated.

2.11.2 Create binned preview

This task calculates binned images of the data. In case of multi-chip cameras all chips of one exposure are collected and arranged in the correct manner, such that you can see them all at once. Both FITS and TIFF versions are calculated,

the latter can easily be eyeballed in your favourite image viewer programme. The binned versions are stored in *BINNED_FITS* and *BINNED_TIFF* sub-directories.

Parameters

The screenshot shows a dialog box titled "Create binned preview". It has three input fields: "Binning" with a value of 2, "TIFF min" with a value of -50, and "TIFF max" with a value of 250. There is also an unchecked checkbox labeled "Outlier rejection".

- **Binning:** The binning factor.
- **Outlier rejection:** Optionally, an outlier rejection can be switched on for the data binning. The 10% highest and 10% lowest pixels in the binning stack will then be rejected. If your data has a lot of hot pixels this will result in a much cleaner binned image.
- **TIFFmin / TIFFmax:** The TIFF version of the binned FITS image has a dynamic range of only 8 bit. You must therefore specify the black- and the white-point, i.e. pixels with values below *TIFFmin* are represented black, and those above *TIFFmax* are white. Since the sky background varies from exposure to exposure, this would lead to a very inhomogeneous representation. Therefore the mean value of the binned exposure is subtracted before conversion to TIFF (this does not happen for the binned FITS). *TIFFmin* should therefore be negative, and *TIFFmax* positive. This can require some trial and error in order to obtain a pleasing result. As a guide line, *TIFFmax* should be about 5-10 times larger than the absolute value of *TIFFmin*.

User-defined multi-chip cameras

Users of single-chip cameras can ignore this section.

The idea behind the binned images is to offer the user a visual representation of the data that displays all chips in a multi-chip camera at once. Like that chips affected by a satellite trail can be readily identified, or the quality of the pre-reduction verified (e.g. defringing, background subtraction, gain correction etc...).

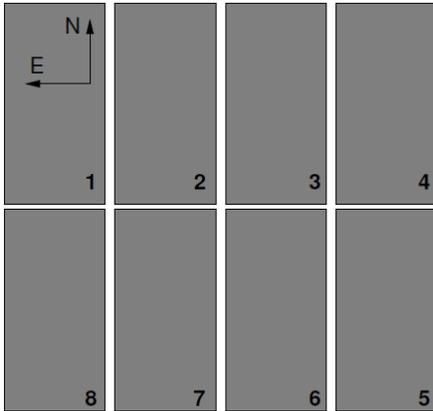
This requires that THELI knows how to arrange the individual chips of a multi-chip cameras. If you defined a new instrument and you want to use this task, then you must manually create the following two files:

Chip configuration

The spatial distribution of the chips in your instrument (let's assume it is called SOMEINSTRUMENT) must be stored in

```
~/the.li/scripts/album_SOMEINSTRUMENT.conf
```

This file must start with a commented header line, followed by one line for each chip defining the x and y position of its lower left edge at full resolution. In other words, think of your detector mosaic arranged in a global pixel grid. Then you would write down the x and y coordinates of the chips' lower left edges, starting with chip 1 and ending with chip n. For WFI@MPGESO, an camera with 8 CCDs,



album_SOMEINSTRUMENT.conf looks like

```
#
1      4150
2100  4150
4200  4150
6300  4150
6300  1
4200  1
2100  1
1      1
```

Binning script

The configuration file created above must be loaded by an instrument-specific binning script. You can use any of the pre-defined *make_album_xxx.sh* scripts located in

```
THELI/gui-<version>/scripts/
```

as a template. Copy and rename it to

```
~/theli/scripts/make_album_SOMEINSTRUMENT.sh
```

There is only one part in the script that needs to be updated, the call to a function `_${P_ALBUM}`. It needs the total unbinned size of the mosaic in pixels (no precise value are needed, just round it up to the next 100 or so), and each chip as an individual argument. In case of WFI@MPGESO the corresponding part of the script reads

```
_${P_ALBUM} -f ${DIR}/album_${INSTRUMENT}.conf ${ARGUMENT} -b ${V_WEIGHTBINSIZE} \
-p -32 8500 8300 \
${BASE}1$3.fits \
${BASE}2$3.fits \
${BASE}3$3.fits \
${BASE}4$3.fits \
${BASE}5$3.fits \
${BASE}6$3.fits \
${BASE}7$3.fits \
${BASE}8$3.fits \
> BINNED_FITS/${BASE}$3mosaic.fits
```

Herein, “8500 8300” would be the mosaic size in pixels (the `-p -32` must not be changed), and then follow eight lines for each of the eight chips. If your camera has only four chips, you’d simply remove the lines for chips 5-8. If your camera has more, just add more.

Note: The order of chips in *make_album_SOMEINSTRUMENT.sh* must be the same as in *album_SOMEINSTRUMENT.conf*.

2.11.3 Create global weights

Each image that enters the coaddition process has its own, individual, weight map. The basis of the latter is the normalised flat field telling how much information a particular pixel contains compared to all other pixels. Therefore pixel weights are not simply 1 for good pixels and 0 for bad pixels, but have a floating point representation.

Optionally, static pixel defects and other features that should be masked can be set to zero, forming the *global weight* that is the same for all images of a particular detector. This global weight is individually modified in the *Create WEIGHTS* task below.

The globalweight is stored as

```
/mainpath/WEIGHTS/globalweights_1.fits
```

Parameters:

- **Same weight for all pixels:** If you don't have a master flat for the basis of the global weight, or don't want this approach for other reasons, you can force all pixels in the global weight to have a value of "1.0". Bad pixel detection as described above based on a master dark can still take place.
- **Flat norm:** Pixels with values outside the [min:max] interval in the normalised flat field are set to zero in the global weight. In this manner significantly vignetted regions and other underexposed detector areas can be masked. If left empty, default values ($min=0.6$, $max=1.5$) will be used.
- **BIAS/DARK:** Optionally, the master dark can be used to identify bad pixels. Pixels with values outside the [min:max] interval will be set to zero.
- **Use BIAS instead of DARK:** If you don't have a master dark, the master bias can be used to identify bad pixels. However, this works less reliably than with a master dark.

Defect detection

Warning: Defect detection will NOT work with CMOS detectors, such as near-IR HAWAII arrays or DSLR cameras, NOR will it work with amateur astronomer cameras using a Bayer colour filter matrix.

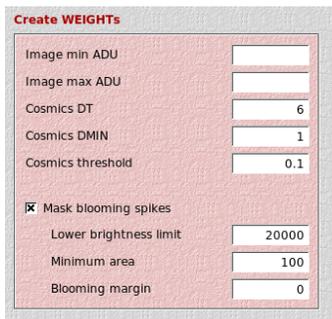
THELI can attempt to identify bad rows, columns, and clusters of bad pixels running the flatfield through a highly sensitive feature detection filter. The flatfield is divided by a smoothed version of itself, yielding a normalised detection map with all illumination patterns removed. For the detection of bad rows (columns) a mean row (column) is calculated in an iterative manner with internal outlier rejection, and individual rows (columns) are compared against it.

- **Smoothing size:** This is the size of the smoothing kernel. It should be chosen comparatively small, such as a few 10 pixels.
- **Cluster tolerance:** Pixels in the detection map are masked if they deviate by more than this fraction from their local neighbourhood.
- **Row tolerance:** If a row deviates by more than this fraction from the mean row, it will be masked. Useful thresholds are around 0.01-0.02.
- **Row tolerance:** If a column deviates by more than this fraction from the mean column, it will be masked. Useful thresholds are around 0.01-0.02.

2.11.4 Create WEIGHTS

In this step the *global weight map* is modified for each image individually and stored in the WEIGHTS directory.

Parameters



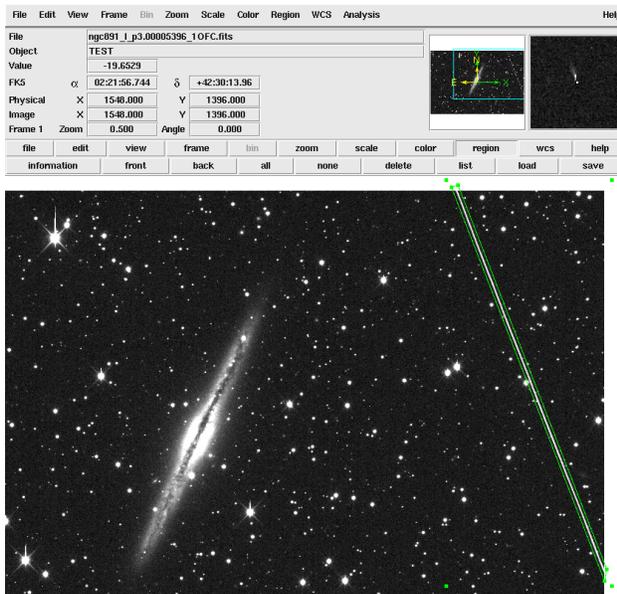
Parameter	Value
Image min ADU	
Image max ADU	
Cosmics DT	6
Cosmics DMIN	1
Cosmics threshold	0.1
<input checked="" type="checkbox"/> Mask blooming spikes	
Lower brightness limit	20000
Minimum area	100
Blooming margin	0

- **Image min (max) ADU:** Pixels in the image with values below (above) this threshold will be set to zero. In this way you can explicitly mask pixels that are e.g. close to saturation, nonlinear, or suffer from other problems.
- **Cosmics threshold:** If you set this parameter (default for well-sampled images: 0.1) THELI runs a neural network algorithm over the image that has been trained to detect cosmics. The latter appear significantly sharper than the PSF in well-sampled images and can thus be masked (without running a computationally expensive comparison between all images). This approach breaks down for undersampled data, in which case the cores of stars will be masked. If you observe this in your stacked image (or in the individual weight images) then switch off this feature (i.e. leave the field empty). You can use an outlier detection algorithm during the coaddition instead to get rid off spurious pixels.
- **Cosmics DT:** The minimum S/N per pixel for a cosmic. The higher, the less cosmics will be detected. The default setting of 6 should be fine for most cases.
- **Cosmics DMIN:** The minimum number of connected pixels with S/N above *Cosmics DT* that make up a cosmic. The default setting of 1 should be fine for most cases. If *DMIN* is larger than 1, then individual hot pixels will not be masked anymore.

Mask blooming spikes: If you want to explicitly set the weight of pixels affected by blooming spikes to zero, activate this switch. The algorithm uses some adaptive statistics in order to deal with detectors where the saturation limit is spatially varying. The idea is that in a histogram of pixel values the brighter pixels are less and less common. However, once saturation kicks in, pixels bleed over into neighbouring pixels at specific ADU levels, which is recognisable as a pile-up in the histogram at high values. This algorithm attempts to detect this turn-up by means of comparison with the number of pixels with lower ADU values. This turn-up can be different from chip to chip, and even vary within a CCD.

- **Lower brightness limit:** This is NOT the threshold above which pixels are considered to be part of a blooming spike. This value should be chosen such that it is well above the background level (in order to remove the latter from the pixel statistics), but still significantly lower than the lowest value of pixels being part of a blooming spike (see the description above). You could set it about twice as high as the typical background level, or about to 30-40% of the saturation level.
- **Minimum area:** This is the number of pixels a blooming spike should at least consist of in order to be masked.
- **Blooming margin:** If the values of pixels in blooming spikes vary depending on detector position, then you should set this parameter to a value other than zero. For example, if you observe variations of 5000 ADU between the top and the bottom of the CCD, you should set this parameter to 5000. The dynamic range within which pixels are masked is therefore increased.

Masking satellites



Satellites can be masked using `ds9` polygons:

1. Open the image containing a satellite track in `ds9`
2. From the main menu, select `Region -> Shape -> Polygon`
3. Click in the image and drag a rectangle with the mouse button pressed. When finished, click into the rectangle.
4. Point the cursor precisely over one of the vertices, then click on it and pull it towards the satellite trail. Repeat with the other vertices until the trail is fully enclosed in the polygon. If you need more vertices, just click onto a connecting line.
5. Click on `Region -> Save` to save the polygon mask. Save it in the same directory as your images, with the same name as the image but without the `status string`. For example, if the image file name is

```
NGC1234_30FC.fits
```

then the `ds9` mask should be called

```
NGC1234_3.reg
```

When saving the mask, `ds9` will ask you in which format to store it. Choose `REG`, then `ds9` and `physical`.

The masks created in this manner will be taken into account when the individual weight map are created. They will be slightly reformatted and moved to

```
SCIENCE/reg/*.reg
```

You can mask any image artifact in this way using arbitrarily shaped *ds9* polygons.

2.11.5 Distribute target sets

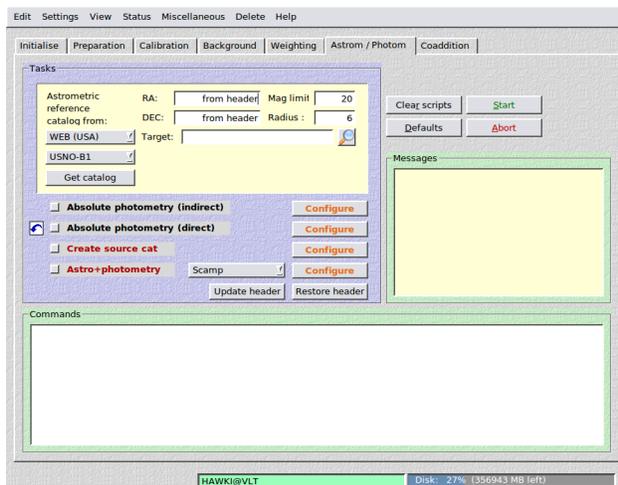
If you observed different targets in the same filter in a given night, you will probably have kept them in the same directory up until now (at least if you observed in the optical). The following steps however (astrometry etc.) must be performed on each pointing individually, and thus different targets must be moved to separate SCIENCE directories.

If images overlap by more than **minimum overlap** pixels (linear extent), then they are considered to belong to the same pointing. Otherwise they will be moved to newly created *set_i* directories (in the directory tree at the same level as the SCIENCE directory in which they lived before this task run). The *set_i* directory will be automatically inserted in the *Initialise* section.

The largest value accepted by *minimum overlap* is 1024.

2.12 ASTROMETRY / PHOTOMETRY

In this section astrometric and photometric solutions are obtained, i.e. images are registered to sky coordinates and are (optionally) flux calibrated. Note that the determination of an absolute photometric zeropoint is optional. The relative zeropoint of exposures, i.e. transparency variations etc are always determined. Results are written to separate FITS headers which are later-on read during the coaddition process.



2.12.1 Astrometric reference catalog

Making the right choice

This little section controls the download of the astrometric reference catalogue. While programmes such as *scamp* can do that internally and thus invisible to the user, there are numerous cases where such an approach is problematic. The principle behind all astrometric solvers is to match a catalog of reference sources to catalogs containing sources in your images. In order to identify the correct solution, a certain minimum contrast (in the one or other statistical sense, depending on the method) is required. This can fail in several situations:

1. You observed a field at low galactic latitude (i.e. in the milky way) with an extremely high source density. If you download too many (i.e. too faint) sources, then the solution space is so densely populated that the right solution is difficult to recognise (and the computational overhead is very large).
2. You observed a field at high galactic latitude with very low object density. Then you want your astrometric reference catalog to be as deep as possible in order to have sufficient overlap with your image catalogue.
3. You observed with a very large telescope and most stars in your reference catalogue are saturated in your images. Or the field of view is so small that none or only very few reference sources overlap with it. In this case you have to consider a deeper reference catalog, or you create one using secondary imaging of the area with a smaller telescope yielding deeper catalogs than what is available in online databases.
4. Wavelength mismatch: You observe a region with high galactic extinction (e.g. a dark cloud) in the optical, and the reference catalog does not contain any sources for that area. You might consider switching to a near-IR catalog.
5. Wavelength mismatch: You want to match a very blue optical catalog to a near-infrared reference catalog. They hardly have any targets in common and thus the astrometry fails.
6. One of the online servers is down and you have to switch to a different one.
7. Any combination of the above, or something else.

Astrometry in THELI usually works fine with the default settings, provided that the information in your FITS headers is halfway reliable. You can try and run it blindly, and if it fails then fine-tune some settings.

The enumeration above shows that in the end the success of the astrometric solution depends on the overlap between the reference catalog and the image catalogs. That is, you should not only think of the reference catalogue but also of the [source catalog creation](#) as the latter allows you to control how many sources are detected in your images.

You should aim for 100-1000 sources per chip, in the reference catalog as well as in the image catalogs. THELI will also work with a few 10s or with 10000+ per chip, but if you push it to the limits it can become difficult.

Parameters

The screenshot shows a web form titled 'Astrometric reference catalog from:'. It contains several input fields and buttons:

- RA:** A text input field with the value 'from header'.
- Mag limit:** A text input field with the value '22'.
- DEC:** A text input field with the value 'from header'.
- Radius :** A text input field with the value '40'.
- Source Selection:** Two dropdown menus. The first is set to 'WEB (France)' and the second is set to 'SDSS-DR7'.
- Action:** A button labeled 'Get catalog'.

- **RA / DEC:** This is only needed if your FITS header does not yet contain astrometric coordinates, or if these are significantly offset from the correct value. Otherwise the centre for the reference catalog will be determined from the FITS header.
- **Resolve target:** Enter the name of your target. Once you click the button with the magnifying glass, THELI sends a query to the Simbad, NED and VizieR databases to retrieve the target's coordinates (if not already contained in the FITS header).

Note: After entering coordinates, THELI will prompt you with a notification window when you hit the *Start* button. Therein, you have three options:

1. **Update RA/DEC:** Will overwrite the CRVAL1/2 keywords.
2. **Update RA/DEC, reset CD matrix:** Will overwrite the CRVAL1/2 keywords and reset the CD matrix (CDi_j keywords) to a North up and East left orientation.

3. **Leave header unchanged:** Select this option if you changed your mind and do not want any changes being made to the FITS header.

-
- **Mag limit:** Only sources brighter than this limit will be retrieved. This is your control over the density of the reference catalog.
 - **Radius (optional; in arcminutes):** Determined automatically, taking into account detector size and dither offsets. Objects within this radius of the nominal coordinates in the FITS header (or those entered manually) will be downloaded.
 - **Source:** Choose from 9 different download locations. If you have problems with one of them, switch to another. You can also choose *IMAGE*, in which case you will be presented with a file selection dialogue from which you can pick an astrometrically calibrated image that covers your field of view. In that case you have to provide the usual SExtractor detection thresholds *DT* and *DMIN* (defaulted to 10 each).
 - **Catalog:** Choose from 8 different reference catalogs

Once your selection is made, click on **Get catalog** to retrieve the reference catalog. The number of sources found will be displayed.

Available reference catalogs

Choose from one of these astrometric reference catalogs:

Catalog	Mag-limit	Remarks
SDSS-DR9	23.0	Covers 14000 sq.deg around the northern galactic cap
PP-MXL	21.0	Combines USNO-B1 with improved astrometry and 2MASS
USNO-B1	21.0	0.3"-0.4" rms uncertainty; can exhibit large systematics
2MASS	18.0	Near-infrared, low systematics
UCAC4	16.5	Good proper motions, all CCD based
GSC-2.3	19.0	
TYC	11.0	The TYCHO catalog
ALL-SKY	10.0	Based on GSC-2.3, filtered to mag<=10. Comes with THELI as downloading huge areas from CDS often resulted in incomplete sky coverage.

TYC and ALLSKY should only be used for calibration of extremely wide fields (obtained e.g. with photo lenses) with very large pixel scale (much larger than 1.0"/pixel).

2.12.2 Absolute photometry

There are two ways to perform absolute photometry: an indirect approach using external standard star fields, and a direct approach using sources with known brightness in the field of view. THELI supports both methods, returning the photometric zeropoint *ZP* (for an image normalised to an exposure time of 1s).

Before you start

Doing absolute photometry in THELI is very much simplified from a user's point of view as compared to what actually happens in the background. The reason is that the absolute zeropoints for the images must be known *BEFORE* object

catalogs are created and *BEFORE* they run through astrometry and relative photometry. Otherwise they cannot be propagated properly and an absolute zeropoint for the coadded image cannot be obtained.

THELI must do these steps to obtain a ZP:

Indirect approach: The standard star field has to be calibrated astrometrically in order to identify the standard stars unambiguously. This requires the download of a reference catalog, the creation of object catalogs, and a (simple) astrometric solution (no distortion fitting). Lastly, THELI will determine the ZP.

Direct approach: Astrometric and photometric reference catalogs must be downloaded for the target area, followed by catalog creation and an astrometric solution (such that targets in the field can be identified with those in the reference catalogs). Lastly, THELI will determine the ZP.

Note: As you can see, obtaining the ZP requires catalog creation and astrometric solutions. The configuration parameters for these tasks will be taken from the configurations of the *Create source cat* task and from the *Astro+photometry* task described further below. You must make sure that these settings are sensible, otherwise the absolute photometry will fail. It is a good idea to familiarise yourself with these two tasks first before you dive into absolute photometry.

Indirect approach: using external standard stars

Prerequisites

1. You must have observed at least one field with photometric standard stars in the same filter as your observations. Select the following standard star catalogs for observations in these filters (note that in particular in the near-IR the filter curves of these standard stars may not be the same as the ones used in your filter set. Check the pertinent publications for more details).

FILTER	Catalog
<i>UBVRI</i>	Landolt (Landolt 1992)
<i>UBVRI</i>	Stetson (Stetson 2000)
<i>ugriz</i>	Stripe82 (Smith 2007)
<i>u'g'r'i'z</i>	Stripe82 (Smith 2007)
<i>JHK</i>	MKO (Leggett 2006)
<i>JHK</i>	HUNT (Hunt 1998, aka ARNICA)
<i>JHKs</i>	PERSSON (Persson 1998)
<i>JHKs</i>	PERSSON_RED (Persson 1998)
<i>YJHKLM</i>	JAC (WFCAM+MKO+MKO_LM)
<i>LM</i>	MKO_LM (Leggett 2003)

2. Your observations must have been conducted in one of the filters listed in the table above
3. The night should have been photometrically stable, i.e. no change in transparency at a given airmass.
4. The standard star field(s) should have been visited several times at night at significantly different airmasses (for extinction calculation).
5. You may include standard star and target observations from many different nights. An independent solution will be calculated for each night.

Before you start reducing your target data, create a separate directory at the same level in the directory tree as the SCIENCE directory containing your main target. Include the name of your standard star directory in the *Initialise* section in the STANDARD field. When you reduce your target observations, the standard data will be processed too, including the same calibration files (master bias, master flat, if applicable also superflat and defringing) as the target observations.

Note: You can collect different standard star fields (taken in the same filter) in the same STANDARD directory.

THELI will automatically obtain a photometric solution for each night and present you with a dialog in which you can select if you prefer the 1-, 2- or 3-parameter fit, or if the solution is not acceptable.

Note that if your data set extends over several nights, it is sufficient for an absolute photometric calibration if only one night was calibrated successfully.

Parameters

- **REFERENCE CATALOG:** Choose one that matches your filters
- **FILTER keyword in observation:** The name of the filter in which you observed.
- **FILTER name in the standard catalog:** The name of the filter in the standard star catalog which matches your filter
- **COLOR INDEX in the standard catalog:** The colour term for the photometric solution is based on the colours in the reference catalog, i.e. you do not have to observe the standard star field in different filters. Select the colour index that matches your observations, respectively the corresponding filter in the standard catalogue.
- **EXTINCTION:** If your standard star observations do not cover a sufficiently large range in airmass, then the atmospheric extinction cannot be estimated reliably from the data. In this case the value entered here will be taken as the extinction. It must be negative. If no value is entered, the default of -0.1 will be used.
- **COLOR TERM:** If the colour term cannot be determined reliably from the data, then the value entered here will be used. If left empty, it is set to 0.1.
- **ZEROPOINT brighter/fainter acceptable limit:** THELI will reject zeropoints that lie outside this interval. This is just to avoid catastrophic outliers.
- **MINIMUM MAG for standard catalog:** Stars brighter than this magnitude will not be used for the fit.

Once the task is started, it calculates three different photometric solutions for each night:

- 3-parameter fit: extinction, colour-term and ZP are estimated
- 2-parameter fit: extinction and ZP are estimated, colour term is fixed
- 1-parameter fit: ZP is estimated, extinction and colour term are fixed

For each night, the three fits are graphically displayed as a line plotted over the corresponding data points. At the bottom of this dialogue, select which fit should be chosen, or if none of the fits is acceptable. Based on the plots you may fine-tune your manual selection of a fixed extinction and colour term coefficient, and then re-run the task. The zeropoints and coefficients will be written into FITS headers. If no ZP is found, the keyword is set to -1.

Note that if your data set extends over several nights, it is sufficient for an absolute photometric calibration if only one night was calibrated successfully.

Direct approach: using field sources with known brightness

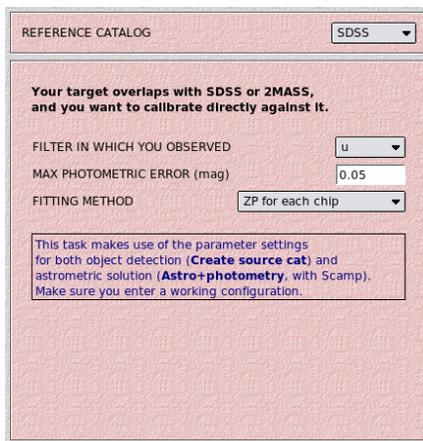
The other possibility for absolute photometric calibration is to determine the zeropoint (ZP) directly from the data. You do not need observations of an external standard star field.

Warning: Direct photometric calibration in THELI does not take into account colour terms arising from different total throughput curves between your telescope/filter/camera combination and the one used for SDSS or 2MASS. If you need absolute zeropoints better than about 0.1 mag then you should consider fine-tuning the ZPs obtained with this method using a comparison of instrumental stellar tracks against those in the Pickels library.

Prerequisites

1. Your field is covered either by SDSS (optical data) or 2MASS (for near-IR observations). SDSS has no all-sky coverage but a high density, whereas 2MASS covers the entire sky but is rather shallow.
2. Your observations must have been conducted in one of the *ugriz* filters (SDSS), or in *JHKs* (2MASS).

Parameters



- **REFERENCE CATALOG:** Select either SDSS or 2MASS
- **FILTER IN WHICH YOU OBSERVED:** Select one of *ugriz* or *JHKs*
- **MAX PHOTOMETRIC ERROR (mag):** The maximum photometric measurement error of sources in your data that go into the fit.
- **FITTING METHOD:** You can select between *ZP for each chip* or *ZP for entire mosaic*. The latter assumes that the pre-processing went fine, in which case no zeropoint variations are expected across the field of view. If you choose to calculate a ZP for each chip, then all chips will be rescaled to the highest ZP found, and this ZP will be written into the FITS headers (i.e. the ZPs in the headers of different chips are identical).

Filename extension: After doing direct photometric calibration with *FITTING METHOD = ZP for each chip*, images have the character *P* appended to their filename extension, e.g.

```
NGC1234_10FCP.fits
```

In all other cases no chip-specific treatment is necessary and thus no indication has to be made in the file names.

2.12.3 Create source cat

In order to detect objects in the images for astrometry and photometry purposes, THELI uses *SExtractor*. We encourage you to make yourself more familiar with this tool (see also the other links on the *SExtractor* website).

The task creates the following sub-directories:

- SCIENCE/cat : contains the binary object catalogs
- SCIENCE/cat/skycat : object catalogs for overlay in *skycat*
- SCIENCE/cat/ds9cat : object catalogs for overlay in *ds9*
- SCIENCE/cat/lownum : dodgy catalogs with too few sources end up here

Parameters

- **DETECT_THRESH:** The minimum S/N of a pixel above the background rms. If this value is decreased, fainter objects will be included in the catalogs. The lowest sensible value is about “1”, combined with *DETECT_MINAREA*=5. The default value is “5”. If your field is very crowded, consider rising it to 10 or higher.
- **DETECT_MINAREA:** The number of connected pixels with a S/N larger than *DETECT_THRESH* required for an object. The smaller this parameter, the fainter the objects included.
- **DEBLEND_MINCONT:** This is *SExtractor*'s internal deblending contrast. You can leave it at its default setting of 0.0005, unless you need to detect objects against a very crowded or structured background, such as a large spiral galaxy. You can e.g. make it as small as 0.00001. If a large number of detections e.g. in a globular cluster seems to shadow the right astrometric solution, increase it to e.g. 0.01, in which case less sources will be detected in the cluster.
- **Minimum FWHM:** Objects with a FWHM smaller than this value will be rejected from the catalogs. Normally you do not have to touch this setting and can leave it to its default setting of 1.5. THELI runs several filters over the object catalogs in order to remove spurious detections such as hot pixels. If your data is *significantly undersampled*, i.e. the FWHM is less than two pixels broad, then you should decrease this parameter to a sensible lower limit. Values less than about 0.7 or 0.8 do not make much sense.
- **Saturation level:** At this level saturation occurs. If nothing is entered, then THELI will assume a conservative value of 50000 ADU. For 12bit cameras such as DSLRs this value should be set much lower. For near-IR detectors it can also be much lower after *background subtraction*, but also much higher if on-chip coadding took place.
- **Maximum FLAG:** Per default, THELI keeps only objects that were not flagged (FLAG=0) as problematic by *SExtractor*. However, there are cases when you need to include flagged objects as well in order to arrive at an astrometric solution. For example, for mildly saturated stars one can still determine a good centroid, which can be very helpful if the faint end of the *astrometric reference catalog* is close to saturation in your images. Similar considerations apply to objects close to image borders, etc.

Quoting from the *SExtractor* manual: The FLAGS parameter contains all the extraction flags as a sum of powers of 2:

- **FLAGS=1:** The object has neighbours, bright and close enough to significantly bias the MAG_AUTO photometry, or bad pixels (more than 10% of the integrated area affected)
- **FLAGS=2:** The object was originally blended with another one

- **FLAGS=4:** At least one pixel of the object is saturated (or very close to)
- **FLAGS=8:** The object is truncated (too close to an image boundary)
- **FLAGS=16:** Object's aperture data are incomplete or corrupted
- **FLAGS=32:** Object's isophotal data are incomplete or corrupted
- **FLAGS=64:** A memory overflow occurred during deblending
- **FLAGS=128:** A memory overflow occurred during extraction

For example, an object close to an image border may have $FLAGS=16$, and perhaps $FLAGS=8+16+32=56$.

- **Min. # of objects:** If a catalog contains less objects than this number, then the catalog and the corresponding image are moved to *lownum* sub-directories and will not enter astrometry and coaddition.
- **Background level:** *SExtractor* usually gets the background level correctly. However, it may fail if a large number of NAN pixels is present, and then the object detection will also fail. In this case you can manually set the background level, assuming it is the same for all images (I was forced to use this setting only once, for background-subtracted Pan-STARRS images).
- **Data is well sampled, but plagued with hot pixels:** Activate this switch if your images have REALLY A LOT OF hot pixels and/or other strange small-scale features that might be mistaken for real objects. *A lot of* means *comparable to* or *more than* real objects. This probably makes only sense for HAWAII near-infrared arrays with very bad cosmetics. Optical astronomers can ignore this switch.

2.12.4 Astro+photometry

5 different methods are available for astrometry. The first two methods, *Scamp* and *Astrometry.net*, perform automatic mosaicing, whereas the two *shift* approaches do not base their solution on WCS coordinates. All but the last method (*Header*) perform relative photometry as well, i.e. measure transparency variations in a series of images.

1. **Scamp:** The preferred method, fast, creates meaningful check-plots. Can be difficult to find a match with the reference catalog if both the position as well as the position angle (and possibly a flip) are highly uncertain or unknown.
2. **Astrometry.net:** Matching is done differently as compared to *Scamp*. Relative zeropoints, and optionally distortion, is calculated afterwards with *Scamp* (automatically, using the preset *Scamp* parameters). Checkplots are provided by *Scamp*. Comparably fast as *Scamp*.
3. **Shift (float/int):** A very simple solver that calculates image shifts only. No rotations and no WCS parameters are involved. The shifts are computed in pixels with respect to the first image in the series. This is only useful for mid-IR data where only very few or only one object at all is visible. The astrometry of the coadded image will NOT be accurate as no comparison with reference catalogs is made. If you use the *integer* approach, only integer pixel shifts and no resampling will be performed during coaddition. This method cannot be used for mosaicing.
4. **Header:** This will simply copy the unchanged raw zero-order solution (CRVAL, CRPIX, CDi_j) present in the FITS headers, albeit without distortion terms. This can be chosen if the raw data already has a good astrometric solution in the headers (e.g. determined by other software) and you prefer to use this over the THELI solution.

Warning: No relative photometric calibration will be done when choosing this method.

All methods create separate FITS header files with the relevant astrometric (and, if applicable, photometric) information. The headers are stored in

SCIENCE/headers/

Update header: Writes the zero-order astrometric solution (CRPIX, CRVAL, CD-matrix) into the FITS headers. This is usually not necessary unless you need accurate sky coordinates for some reason. The only such reason in THELI is when you have to extract a background estimate from a small sky section and you want this area to be exactly the same despite a dither pattern. See the *sky subtraction* for details.

Restore header: Restores the previous FITS headers if a zero-order solution was written into them using *Update header*.

Scamp

We strongly recommend to have a look at the [scamp manual](#) for details as of how *Scamp* works.

The following parameters are available.

The screenshot shows the Scamp Astrometry.net interface with the following parameters and values:

- POSANGLE_MAXERR (deg): 2.0
- POSITION_MAXERR (arcmin): 2.0
- PIXSCALE_MAXERR (factor): 1.05
- DISTORT_DEGREES: 3
- FGROUP_RADIUS (deg): 1.0
- CROSSID_RADIUS (arcsec): 1.0
- ASTREF_WEIGHT: 1.0
- SN_THRESHOLDS (low | high): 5 | 20
- Add. DISTORT_GROUPS | DEGREES: (empty)
- Additional DISTORT_KEYS: (empty)
- ASTRINSTRU_KEY: FILTER
- PHOTINSTRU_KEY: FILTER
- STABILITY_TYPE: INSTRUMENT
- MOSAIC_TYPE: README | SAME_CRVAL
- FOCAL_PLANE: Use default FP
- Match flipped images

- **POSANGLE_MAXERR:** Maximum uncertainty in the position angle (up to 180 degrees)
- **POSITION_MAXERR:** Maximum uncertainty in the pointing in arcminutes
- **PIXSCALE_MAXERR:** Relative uncertainty of the pixel scale
- **DISTORT_DEGREES:** Distortion polynomial order. Distortion in the sense of *Scamp* is a variation of the pixel scale as a function of position. All other (linear) transformations are explicitly modelled as such and encoded in the CD matrix. The distortion degree entered in this field means:
 - *distort=1*: no distortion modelled; pixels have the same scale everywhere
 - *distort=2*: pixel scale may vary linearly across the chip (makes sense only for multi-chip cameras)
 - *distort=3*: pixel scale may vary quadratically
 - *distort=4*: cubic terms are included ...

Scamp can also make use of higher distortion polynomials, but they get increasingly more unstable and require a large number of sources.

- **FGROUP_RADIUS:** The maximum angular separation between to pointings. If two images are further apart, then independent astrometric solutions will be calculated for them.

- **CROSSID_RADIUS:** The search radius (in arcsec) for cross-identifications. This value can remain at its default setting of 2, unless the pixel becomes larger than 0.7 arcsec/pixel. In the latter case, THELI sets

$$\text{CROSSID_RADIUS} = 2.5 * \text{PIXELSCALE}$$

- **ASTREF_WEIGHT:** In *Scamp* the object catalogs and the reference catalogs have equal weights in the solution process. Some cases may require to put more weight on the reference catalog. If your reference catalog is very precise and accurate to within e.g. 1/5th of a pixel of your camera, then you may choose a very large weight such as 10 which would make the reference catalog control your solution. Usually, the resolution of the camera is equal to or better than the uncertainty of the reference catalog. The default value for this parameter is 1.
- **SN_THRESHOLDS (low|high):** Minimum S/N of a source to be kept in *Scamp* (low), and the minimum S/N for some high-S/N statistic calculations (high).
- **Additional DISTORT_GROUPS | DEGREES | DISTORT_KEYS:** In almost all cases one can ignore these three settings (i.e. leave them empty). One would enter here a comma-separated list of additional distortion groups, degrees and keywords. *DISTORT_KEYS* would also accept FITS keywords, which must be preceded with a ":".

Example: A set of images with arbitrarily distributed position angles is being stacked, and it turns out the distortion pattern is a linear function of the position angle. One would then define a new distortion group, *DISTORT_GROUPS=2*, with *DISTORT_DEGREES=2* and *DISTORT_KEYS=:POSANGLE* to reflect the dependence of the distortion of the position angle. For more details see the **Scamp* manual.

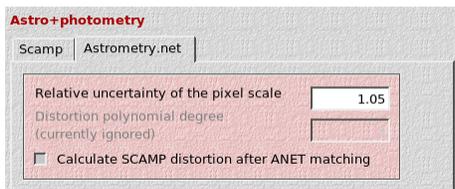
- **ASTRINSTRU_KEY:** The FITS header keyword used to identify exposures for which one astrometric solution should be calculated. The default value is *FILTER*, meaning that images with the same *FILTER* header keyword will have one solution calculated. Images with different values in those keywords will get their own solution. You can add several comma-separated keywords in order to break down the astrometric solution into smaller groups. If you want to calculate one solution for images taken in different filters, then just enter a string that does not exist as a keyword in the FITS header, such as "none".
- **PHOTINSTRU_KEY:** The FITS header keyword used to identify exposures for which one photometric solution should be calculated. The default value is *FILTER*, meaning that images with the same *FILTER* keyword will have one solution calculated. If you want to calculate one solution for images taken in different filters (for flux-normalisation purposes for example, e.g. difference imaging), then just enter a string that does not exist as a keyword in the FITS header, such as "none".
- **STABILITY_TYPE:** *INSTRUMENT* can be used if one is sure that the distortion pattern does not change during a series of exposures. *EXPOSURE* should be used if all parameters shall be determined individually per exposure (and chip).
- **MOSAIC_TYPE:** Quoting from the *Scamp* user manual: *Scamp* can manipulate mosaics in a number of ways to perform the matching of sources on the sky, and the astrometric calibration itself. For single-chip cameras only the *UNCHANGED* mode makes sense.
 - **UNCHANGED:** The relative positioning of detectors on the focal plane, as recorded in the WCS keywords of the FITS headers, is assumed to be correct and constant from exposure to exposure. Matching with the reference catalogue will be done for all the detectors at once.
 - **SHARE_PROJAXIS:** The relative positioning of detectors is assumed to be constant and correct, but the different extensions within the same catalogue file do not share the same projection axis (the CRVAL FITS WCS keywords are different): although this does not prevent *Scamp* to derive an accurate solution, this is generally not an efficient astrometric description of the focal plane. This option brings all extensions to the same centred projection axis while compensating with other WCS parameters.
 - **SAME_CRVAL:** Like *SHARE_PROJAXIS* above, brings all extensions to the same centred projection axis (CRVAL parameters), but does not compensate by changing other WCS parameters. This option is useful when the CRPIX and CD WCS parameters are overridden by some focal plane model stored in a global .ahead file.

- **FIX_FOCALPLANE:** Applies first a *SHARE_PROJAXIS* correction to the headers and then attempts to derive a common, median relative positioning of detectors within the focal plane. This mode is useful to fix the positions of detectors when these have been derived independently at each exposure in an earlier not-so-robust calibration. A minimum of 5 exposures per astrometric instrument is recommended.
- **LOOSE:** Makes all detector positions to be considered as independent between exposures. Contrary to other modes, matching with the reference catalogue will be conducted separately for each extension. The *LOOSE* mode is generally used for totally uncalibrated mosaics in a first *Scamp* pass before doing a *FIX_FOCALPLANE*.
- **Which focal plane:** What *Scamp* should assume about the detector mosaic:
 - **Use default FP:** Use the focal plane model implemented in THELI
 - **Create new FP:** Create a new FP model based on your data (will not replace the default model)
 - **Do not use FP info:** Do not assume anything about the focal plane
- **Match flipped images:** If set, then *Scamp* will attempt to match flipped images, i.e. the image is flipped with respect to what its CD matrix says.

Astrometry.net

For more information about this tool visit www.astrometry.net.

The following parameters are available.



- **PIXSCALE_MAXERR:** Relative uncertainty of the pixel scale
- **Distortion polynomial degree:** Currently deactivated, as *SWarp* does not understand SIP polynomials
- **Calculate SCAMP distortion after ANET matching:** The distortion output of ANET is not understood by *SWarp* (image coaddition). To still get distortion correction, you can run *Scamp* after ANET has done the matching. In this case the matching function of *Scamp* is deactivated. In addition, *Scamp* will also calculate the **relative photometric zeropoints** as usual, which are not provided by ANET. Hence this checkbox should always be activated. THELI will use the distortion polynomial degree (and all other *Scamp* parameters) as entered in the corresponding configuration dialog.

2.12.5 Trouble-shooting *Scamp*

It happens that *Scamp* returns a wrong astrometric solution without recognising it. This usually happens when the images were not correctly matched to the sky. You do not need to go through a possibly lengthy coaddition to find out if the solution was good or not. In most cases it is sufficient to have a look at the various check-plots created by *Scamp*. You can find them in

SCIENCE/plots/

In the same directory is also a *scamp.xml* file with further valuable information that can help you identify trouble-making images and other issues. In order to look at the XML file with *firefox*, you must first enter in the address field

```
about:config
```

and then locate the following parameter and set it to *false*:

```
security.fileuri.strict_origin_policy = false
```

Otherwise the XML file will not be displayed correctly.

If you can't get it done with *Scamp*, you may equally well try *Astrometry.net* which uses a different matching algorithm that is superior if WCS parameters such as pointing and position angle and a possible flip are unknown.

Interpreting *Scamp* check-plots

Image distribution on the sky: fgroups

This should be the first plot you should look at. It displays the dither pattern together with unmatched (red) and matched (green) reference sources. You should recognise your dither pattern here.

Good fgroups plots

- Example 1: Megacam@CFHT
- Example 2: A mosaic of ACS/HST images. The overlap with the ground-based reference catalog is not very compelling, but sufficient.
- Example 3: WFC@INT (wide-field camera at the 2.5m Isaac Newton Telescope)
- Example 4: A large set of rotated GPC1@Pan-STARRS exposures
- Example 5: LBC_RED@LBT (the red camera at the Large Binocular Telescope, with a cubic distortion polynomial). Compare to example 4 below.
- Example 6: A mosaic of mosaics with WFC@INT

Bad fgroups plots

- Example 1: A wide field view obtained with a photo-lens. Only part of the image (green) is matched with the reference catalogue, due to distortions. Increase the *CROSSID_RADIUS*.
- Example 2: Total failure. Reason: very small field of view, insufficient sources in the image catalogs. Try lowering the detection thresholds and set *distort=1* because not enough sources are available to determine many distortion coefficients.
- Example 3: Inaccurate CD matrix (wrong position angles and/or pixel scale). Try increasing *POSANGLE_MAXERR* and *PIXSCALE_MAXERR*. RA and DEC in the FITS header could also be way off. Try increasing *POSITION_MAXERR*.
- Example 4: LBC_RED@LBT (the red camera at the Large Binocular Telescope, with a quadratic distortion polynomial). Note that only parts of the detectors are matched properly to the sky. Compare to example 5 above.
- Example 5: A wide field of view of a milky way area. While the depth of the reference catalog was ok (about 200 sources returned) the object catalogs contained 10000+ sources per exposure. Increase the object detection threshold DT from e.g. 5 to 50 would bring the source density down to a reasonable level.
- Example 6: That's a tricky one, based on MOSAIC-II@CTIO. Note that almost all frames are registered nicely. However, the rightmost exposure shows no green matches (only visible at the rightmost edge; the coadded image would be clearly corrupted, though). Reason here was that the WCS information in the FITS header of this single exposure was 7' offset, whereas the others were all within about 1'. Increasing *POSITION_MAXERR* to 8' solved the problem instantly, after many days of trouble shooting on the user side. Mostly problems with scamp have very easy solutions,

directly in front of your nose. Lesson: Even modern telescopes mess with you every now and then (not the first time I have seen this; WFC@INT sometimes has had its headers off by 20').

Image distortion: `distort`

Distortion can be displayed as a variation of the pixel scale across the detector area. Make sure these plots look as symmetric as possible.

Good `distort` examples

- Example 1: Megacam@CFHT. Note the slight irregularity in the upper half of the image centre. One can do better than that, however this distortion model did not lead to any visible ill effects in the coadded image.
- Example 2: Pan-STARRS. Due to the very short exposure time (30s) and the high galactic latitude only very few sources were visible in this field. Hence only a linear distortion model could be fit to the data. A quadratic fit looked smoother, but showed artifacts with similarly large deviations from reality.
- Example 3: WFC@INT. Note the slight mismatch between the two detectors at the lower left.
- Example 4: WFC@INT. Same as example 3, but using more objects in the reference and the source catalogs.
- Example 5: WFI@MPGESO. The wide-field imager at the 2.2m MPG/ESO telescope is one of the very few instruments where the largest pixel scale does not coincide with the optical axis. You can model it with `distort=3`, but a higher order model (`distort=4..5`) (shown: `distort=5`) yields better results as the curvature of the “ring” is better reproduced.
- Example 6: SuprimeCam@SUBARU
- Example 7: LBC_RED@LBT

Bad `distort` examples

- Example 1: A wide-field exposure through a photo-lens. Try increasing the number densities in the catalogs or `CROSSID_RADIUS` to obtain a more concentric pattern.
- Example 2: Pan-STARRS. Note how the chip at the upper right is completely off. It contained only very few sources due to excessive masking. The corresponding images from this chip were removed by hand, then astrometry was repeated.
- Example 3: Megacam@CFHT. Try e.g. to increase the overlap between object catalogs and the reference catalog.
- Example 4: HAWKI@VLT. This is based on a sparse field. Adding another, neighbouring, sparse field fixed the problem. Wide dither pattern patterns or independent data fields help.

External astrometric residuals: `astr_referror1d`

Astrometric residuals with respect to the reference catalogue after a solution was found. You'd expect a tight and featureless horizontal distribution whose width is solely determined by the precision of the reference catalog.

Good `astr_referror1d` examples

- Example 1: Megacam@CFHT, USNO-B1
- Example 2: Pan-STARRS. Note the slight deviations at the outermost edges. Those chips are largely masked and were difficult to match.

Bad `astr_referror1d` examples

- Example 1: Megacam@CFHT. Adjust the reference and object catalog settings, or try a different reference catalogue.

- Example 2: Pan-STARRS. The only trick that helped here was putting a lot more weight onto the reference catalogue, $ASTREF_WEIGHT=10$. Normally this should be avoided as the images have better resolution and higher object density than the reference catalogue. In this case the reference catalogue was SDSS and significantly deeper than these 30s exposures, hence increasing the weight for the reference catalog is justified.

Internal astrometric residuals: `astr_interror1d`

Astrometric residuals between the individual exposures after a solution was found. You'd expect a tight and featureless horizontal distribution whose width should be on the order of 1/5th-1/20th of a pixel.

Good `astr_interror1d` examples

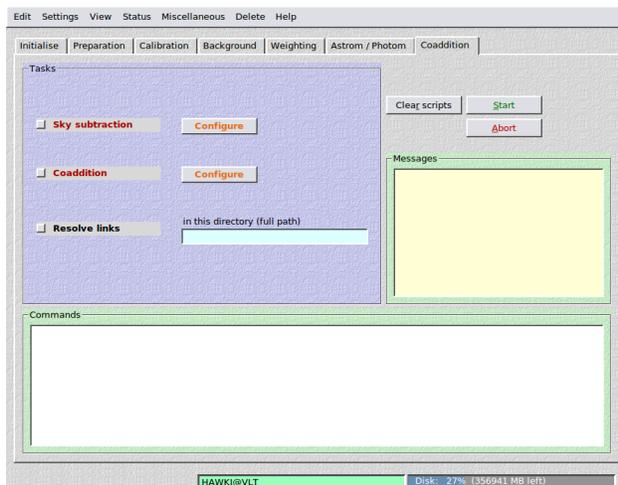
- Example 1: Pan-STARRS
- Example 2: Good residuals for a one-chip camera (pixel scale 0.85 arcsec/pixel)

Bad `astr_interror1d` examples

- Example 1: Megacam@CFHT. Note the small bifurcation in the upper line. The rest is ok.
- Example 2: WFI@MPGESO. The magnitude limit for the reference reference catalog was not deep enough in this case, resulting in badly matched chips.
- Example 3: Note the huge range (10 arcsec) of the y-axis. Normally the range should be about one pixel. Reason here was that the reference catalog downloaded had no overlap at all with the data.
- Example 4: FORS2@VLT observations of a globular cluster. The match is generally good (very tight horizontal line), but there are obvious mis-identifications within the very crowded globular cluster. Repeat the object catalogs with e.g. $DEBLEND_MINCONT = 0.01$, and try 2MASS instead of USNO-B1. 2MASS has significantly better resolution in crowded objects than USNO-B1 which features many blended or otherwise inaccurate sources in such cases.

2.13 COADDITION

This section comprises the final sky subtraction and the coaddition.



2.13.1 Sky subtraction

This step is generally considered to be mandatory, at least for multi-chip cameras, as otherwise the gaps between the chips show up as discrete jumps in the background (if the latter was variable). If you prefer not to do subtract the sky, then simply omit this step.

Sky modelling has to be done with care if very large objects are present in the field. THELI offers different and very flexible modes for sky subtraction.

Filename extension: After running through this step, images have the string `.sub` appended to their filename extension, e.g.

```
NGC1234_10FC.sub.fits
```

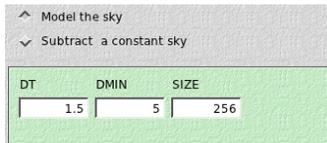
Model the sky

Automatic sky modelling is the default method and useful for all exposures where the largest object is significantly smaller than the field of view of the detector.

Warning: Do not use automatic modelling if you are interested in very faint and extended structures that are not or just barely visible in a single exposure. They will not be masked, therefore enter the background model and as a result will be removed from the data. Subtract a constant sky in this case.

In a first pass, *SExtractor* is run to remove all objects from the image. The resulting gaps in the field are filled in with a dynamically variable estimate, after which a convolution with a Gaussian kernel is performed. You must provide the usual detection threshold, minimum number of connected pixels, and the FWHM (pixels) of the Gaussian kernel. Optionally, the created sky model can be saved.

Parameters



DT	DMIN	SIZE
1.5	5	256

- **DT:** This is the *SExtractor* detection threshold per pixel, given in units of sigma of the sky background noise.
- **DMIN:** The minimum number of connected pixels above the detection threshold making up an object. The smaller *DT* and *DMIN*, the fainter the objects masked.
- **FWHM:** The FWHM of the Gaussian convolution kernel for the sky background. If you want to remove background features with a linear extent n (in pixels), then you should set this parameter the same or a smaller value. Values less than about 60 do not make much sense.

Subtract a constant sky

A constant sky value should be subtracted if very extended targets are present that prevent the sky from being modelled at the target's position.

If desired, you can manually fine-tune the estimates for each image. This approach is useful if the target is in size comparable or even larger than the field of view, or if extended very low surface brightness objects are observed which might be mistaken as a background feature and removed by sky modelling.

Parameters

The configuration is split in three parts, not all of them have to be filled in:

1. Choose statistics

Define how the sky background estimate is determined. Choose between

- **Mode:** the most common value in the smoothed distribution of pixel values
- **Median:** the *median* of the distribution, more stable against outliers than the *mean*
- **Mean:** the *mean* of the distribution
- **Lower quartile:** the *lower quartile* of the distribution. Useful if you are not sure about the contamination of some faint extended object.

Mode or *median* are the best choices in most cases. The other two are unlikely to yield better results. **From each chip:** If you are using a multi-chip camera, measure the sky independently in all CCDs. **From chip i:** Alternatively, estimate the background from one chip only where you think the sky is representative for the entire mosaic, and subtract this value from all other chips. In this case you can make use of windowing (see below).

Note: This task DOES run *SExtractor* over the images, using the detection thresholds provided under *Model the sky*. The statistical estimates are obtained from the object-subtracted images. If you do not provide thresholds, default values will be used ($DT=1.5$, $DMIN=5$).

2. Region with empty sky (optional)

You can specify a small window that does not contain any flux from extended objects and obtain the sky estimate from it. There are two options:

- **Window fixed on the sky:** By entering min and max RA and DEC values, define a window on the sky from which the background is determined. For example, this could be a small dark cloud superimposed on brighter nebulosity. This task also works if the window falls partially into a gap between chips of a multi-chip camera as a consequence of dithering. If no sky could be obtained because the window falls entirely into a gap, then an error message is returned. Click on the *update header* button in the astrometry section before running this task. In this way you make sure that no differential WCS offsets are present in the header and the window will be centred on the same spot on the sky in all exposures.
- **Window fixed on the chip:** By entering Cartesian $xmin$ $xmax$ $ymin$ $ymax$ chip coordinates, the window from which the sky is estimated is always taken from exactly the same detector area, independent of any dither pattern. This task is only meaningful for multi-chip cameras if the *sky is estimated from a single chip* only.

If you leave one or more of the fields empty, the entire detector area will be used.

3. Manual override (optional)

Use this feature only if you want to manually fine-tune individual sky background values.

Note: This task DOES NOT run *SExtractor* to remove objects from the images, the thresholds given under *Model the sky* are not applied. The statistical estimates are obtained from the images as they are.

The following settings can be made:

- **Get estimates:** Determines from each image (no objects subtracted) the mode, median, mean and lower quartile. If a *window* was defined, the sky will be estimated from it. This has to be done only once, as the output is stored internally. The result will be shown in the table, one value for each image. The numeric values are editable. This task will always run on all chips, no matter what selection you make elsewhere.
- **Load/Save:** The statistics obtained by *Get estimates* is saved in:

```
SCIENCE/skybackground_xxx.stat.
```

You can *Load* such a statistics if it already exists, or *Save* it under a particular filename. You do not need to save the table explicitly. THELI will use the values displayed in the table when you leave this dialog (the *what you see is what you get* principle).

- **Fill:** Replaces any value in the table with the numeric value entered in the field to its right. This does not erase the values obtained through *Get estimates*. If you select e.g. *Mode* again for the display, the previous values will be restored. This button can be clicked instead of *Get estimates* at the very beginning, for example if you know that you only want to subtract a value of e.g. 587.3 from all images.
- **Show mode:** Can be used to display either the mode, the median, the mean or the lower quartile in the table. You can cycle through the various measures without losing data.
- **Use all chips:** If you have data from a multi-chip camera, select whether the sky value displayed is taken from each chip individually, or only from a particular chip whose background you consider representative for all other chips. For example, all 8 CCDs of the camera cover a large object, and only chip no. 5 contains a small region with empty sky. Then you would define the *window*, click on *Get estimates* and select *Use chip 5*. You can cycle through the various chips without losing data.
- **Show mosaics only:** Allows you to display exposures only, i.e. the table will contain only one line for a group of n images, where n is the number of chips in your camera. The name of the image of the first chip will be displayed, representative for all other images of the same exposure. The other option is to show all images of all chips.

WYSIWYG principle for constant sky configuration

When you close the dialog, sky subtraction will be performed according to your settings. If you manually fine-tuned the sky, the content of the table will be saved to

```
SCIENCE/skybackground.use
```

together with the associated image names. If a file with the same name is present already, it will be overwritten without warning. This file will be read once you launch the sky subtraction. If you selected *Show mosaics only*, then all images of one exposure have the same sky value assigned. That is, *what you see is what you get*.

2.13.2 Coaddition

This last step is internally broken down into several other steps. First, images are resampled to a common output grid taking into account the *astrometric solution* obtained previously. This step is parallelised. In an optional second step the resampled pixels contributing to one output pixel are compared and outliers can be rejected (their weight is set to zero). This task is not (yet) parallelised and can be lengthy for larger mosaics. Afterwards, the resampled images are combined into the coadded image, and the header of the latter is updated with various parameters and measurements. Resampling and stacking are done with *Swarp*.

The coadded image can be found under

SCIENCE/coadd_XXXX/coadd.fits

SCIENCE/coadd_XXXX/coadd.weight.fits

where *XXXX* stands for an unambiguous, user-provided identifier. Thus different coadditions can be created without overwriting older stacks.

Parameters

- **Reference RA / DEC:** The reference coordinates for the projection on sky. You can leave them empty, in which case they will be determined automatically. However, if you have data of the same target taken in different filters, you should use identical values for all coadditions. You **MUST** use identical values if you plan to create a colour picture or to do difference imaging. The reference coordinates should not be too close the border of the image as otherwise the coadded image might be truncated.
- **Identification string:** This string will be appended to the name of the coadd directory. Therefore different coadditions can be made without overwriting each other. If left empty, it is set to *default*.
- **Coadd these chips:** For a multi-chip camera you can choose to coadd only certain chips instead of all of them. If left empty, all chips will be stacked.
- **Coadd this FILTER:** If you run the astrometry over images taken in different filters, then these will be mixed in the SCIENCE directory. Tell THELI here which filter it should stack. You do not need to move images taken in other filters to parking directories.
- **Proper motion:** If you observed moving targets, you can enter here the proper motion vector. THELI will look up the DATE-OBS keyword in the FITS header and apply a linear offset to the WCS solution contained in the astrometric header. Like that the moving target will fall on top of itself in the coadded image instead of leaving a trail:



Good if you want to detect e.g. faint TNOs not visible in single exposures.

- **Coadd image size (x,y):** For very wide fields of view and all-sky projections, the size of the output image is difficult to predict automatically. In many such cases it will be too small, and you can provide the output image size in pixels manually. You can estimate the size in pixels as follows, for both axes: field of view in degrees divided by pixel scale in degrees, and add some margin for good measure.
- **Pixel scale:** The pixel scale of the output image. You can decrease or increase the resolution. The latter only makes sense if your data is undersampled and if you have a sufficiently large number of dithered exposures, or if you want to match your image to that taken with another telescope or camera. Upsampling by more than 50% is rarely useful with ground-based data.

Note: If left empty, the native pixel scale given in the instrument configuration will be used. For instruments with exchangeable foreoptics (e.g. SOI@SOAR, FORS1+2@VLT) this is the most frequently used pixel scale.

- **Sky position angle:** If left empty then THELI will orient the coadded image such that North is up and East is left. If you want a different sky position angle, then enter it here (counted positive from North over East). If your data was taken under a (unknown) non-zero position angle, you can estimate it by clicking on the *From image* button. This will not rotate your images into standard orientation. However, note that a flipping of axes will always be corrected.
- **Edge smoothing length:** Mosaic data of very extended targets often cannot be fully sky-subtracted as there is no way to estimate the real background level without blank field exposures. In this case one usually subtracts a constant value, which can lead to small discontinuities in the overlap areas. By smoothly reducing the weight map at its borders to zero these transition areas can be blended into each other, eliminating the discrete jumps while preserving photometry. The value entered is the width of the softened edge. As a rule of thumb, it should approximately be the same as or smaller than the extent of the overlap in pixels (for example 50 to 100).
- **Maximum image seeing:** If you want to exclude images with a seeing above a certain threshold, then enter this value here.
- **Minimum relative zeropoint:** If you want to exclude images with a relative zeropoint worse than this threshold, enter it here. The relative zeropoints are normalised such that their mean is zero. Images with worse transparency have negative values, images with better transparency have positive values.
- **Outlier rejection threshold:** The clipping threshold in units of sigma. If a (resampled) pixel deviates by more than this from the mean value of resampled pixels at this location, then it will be removed from the stack (i.e. it has its weight set to zero). If left empty, no rejection will be performed.

A value smaller than 4 should only be used if you have 15-20 exposures or more, ideally with a *cluster size* larger than 1 (see below). Values less than 3 hardly make sense. If you feel you need such aggressive rejection, probably something else went wrong in your data reduction process, such as mismatched darks, bad sky subtraction (near-IR only), or you forgot to dither your telescope. A warning message will be shown.

- **Outlier rejection cluster size:** Only mask pixel groups consisting of that many bad pixels above the threshold. If you find that masking is too zealous, then increase this to values larger than 1. Values larger than 9 are not

permitted and will be reset to 9. Requires the rejection threshold to be non-empty. For the Lanczos2/3 kernels values of 3-6 work appear plausible.

- **Outlier rejection border width:** If a pixel is found to be bad, then also mask a ring of pixels around it. This parameter specifies the width of this ring in pixels. Requires the rejection threshold to be non-empty.
- **Resampling kernel:** Choose from
 - **Nearest neighbour (NN):** No resampling (flux distribution between output pixels) is done. Leads to significant astrometrical artifacts.
 - **BILINEAR:** Use only for extremely undersampled images
 - **LANCZOS2:** Use for significantly undersampled images
 - **LANCZOS3:** Optimal for well-sampled images. If you notice dark edges in your star profiles, then you are likely undersampled and should go with *LANCZOS2* instead.
 - **LANCZOS4:** Best from a sampling point of view, but leaves large artifacts around hot pixels and other steep image features. Slow.

Animation of images resampled with *LANCZOS3* (top) and *NN* (bottom) kernels.

Note the erratic wiggling of stellar positions using the NN kernel:

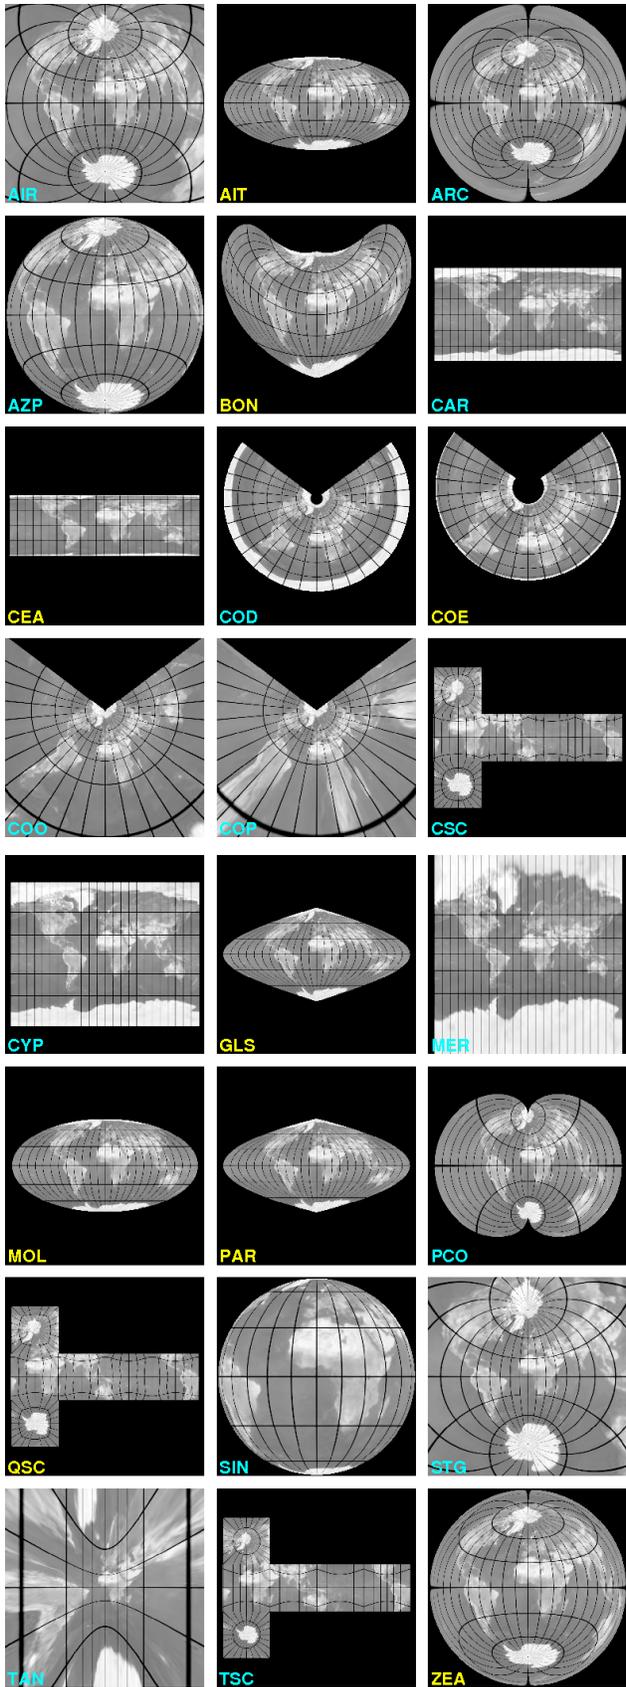
The *NN* kernel can also lead to significant distortions within the PSF, depending where on the pixel grid a source is projected. Look at the bright star immediately above the diffuse large galaxy. In one of the images it appears significantly elongated, whereas this is absent with the *LANCZOS3* kernel.

- **Projection type:** Use your preferred projection type. Will only make a difference for fields of view larger than several degrees. For a complete list and more explanations see *below*.
- **Celestial type:** The coordinate system into which you want to project your images:
 - **NATIVE:** use whatever is defined in the FITS header. In case of THELI this will be *EQUATORIAL*.
 - **EQUATORIAL:** the classic RA and DEC sky coordinates
 - **GALACTIC:** galactic coordinates
 - **SUPERGALACTIC:** supergalactic coordinates
 - **ECLIPTIC:** useful for solar system stuff
- **Combine type:** How the resampled images are combined:
 - **WEIGHTED:** yields the best S/N
 - **AVERAGE:** better S/N than *MEDIAN* but sensitive to outliers
 - **MEDIAN:** worse S/N than *AVERAGE* but stable against outliers
 - **MIN:** use the smallest pixel value in the stack
 - **MAX:** use the highest pixel value in the stack
 - **CHI2:** uses chi-square for the stack (all pixel values will be non-negative)

Sky projections

The *projections* of the proposed *FITS WCS standard* are available in THELI (copied from the *SWarp* manual, thanks Emmanuel!). Nine of them are equal-area projections, i.e. they preserve surface brightness (indicated with yellow labels).

Which of the projections is chosen does not really matter as long as the field of view is about less than 10 degrees. For larger fields projection effects become significant.



Zenithal projections

- **AZP:** Zenithal perspective
- **TAN:** Distorted tangential
- **STG:** Stereographic
- **SIN:** Slant orthographic
- **ARC:** Zenithal equidistant
- **ZPN:** Zenithal polynomial
- **ZEA:** Zenithal equal-area
- **AIR:** Airy

Cylindrical projections

- **CYP:** Cylindrical perspective
- **CEA:** Cylindrical equal-area
- **CAR:** Plate carree
- **MER:** Mercator

Conic projections

- **COP:** Conic perspective
- **COE:** Conic equal-area
- **COD:** Conic equidistant
- **COO:** Conic orthomorphic

Pseudoconic and polyconic projections

- **BON:** Bonne's equal-area
- **PCO:** Polyconic

Pseudocylindrical projections

- **GLS:** Global sinusoidal (Sanson-Flamsteed)
- **PAR:** Parabolic
- **MOL:** Mollweide
- **AIT:** Hammer-Aitoff

Quad-cube projections

- **TSC:** Tangential spherical cube
- **CSC:** COBE quadrilateralized spherical cube
- **QSC:** Quadrilateralized spherical cube

Note: Some projections require additional PVi_j parameters in the header, i.e. you must run *Scamp* with *distort>1*:
CYP, CEA, COD, COE, COO, COP, BON

2.13.3 Resolve links

If you spread the images of a multi-chip camera over several hard disks under *Create links*, then you can recollect them in their main directory (the one where the links are). Just enter the name of the directory where you want to resolve the link structure.

2.14 Miscellaneous THELI modules

THELI offers a few additional tasks you can run outside the framework of the main reduction thread.

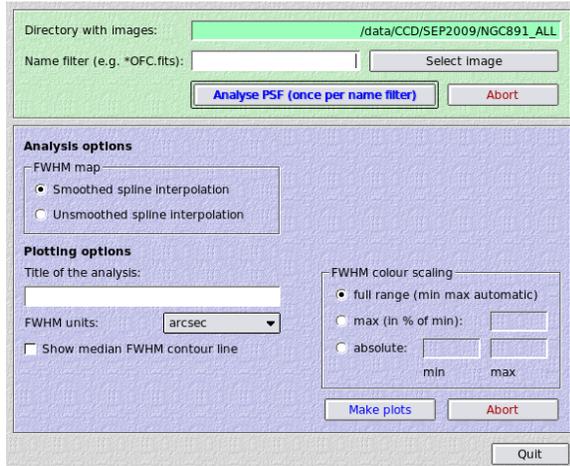
2.14.1 Combine images

Combine a set of arbitrary images in a given directory (without taking into account image shifts etc). Images are stacked as they are.

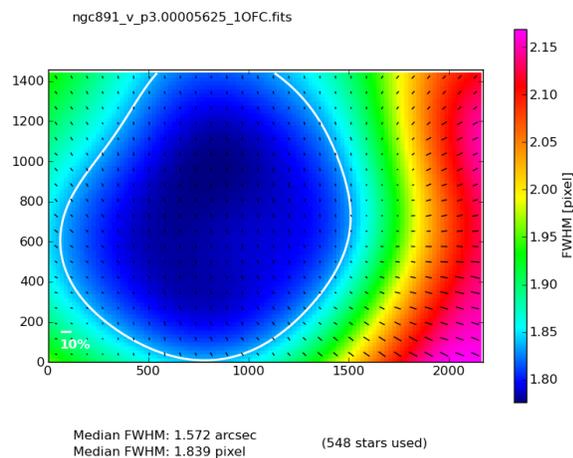
Parameters

- **Name filter:** This should be the *status string* between the chip number and the file suffix, for example *OFC*. If left empty, all images in the directory will be combined. Separate stacks will be created for multi-chip camera data.
- **Name of the combined image:** If left empty, the combined image will have the name of the directory. For example, if the images are in a directory called *NGC1234*, the stacks will be called *NGC1234_i.fits* where *i* is the chip number.
- **DT/DMIN:** Objects will be removed from the images using *SExtractor* before combination takes place. These are the detection threshold per pixel above the sky noise and the minimum number of connected pixels above this threshold forming an object. If these fields are left empty, then objects are left unmasked.
- **Number of pixels to reject:** The number of lowest and highest pixels in the stack to be rejected.
- **Combination method:** Choose between *median* and *mean*.
- **Rescale images:** Optionally, images can be rescaled to the same background level before combination. That could be useful e.g. for the creation of a superflat.

2.14.2 Ialyzer (image analysis)



If your individual exposures (per chip) contain about 100 stars or more, a more detailed image analysis using THELI's *Ialyzer* can be performed. It automatically extracts stellar sources and performs a very sensitive weak-lensing style PSF analysis. FWHM variations across the detector areas are modelled using a 2-dimensional spline fit. The PSF anisotropy (ellipticity) is displayed as a stick plot superimposed over the FWHM map.



The anisotropy is calculated from the second brightness moments of the object's filtered flux distribution. In the idealised case of a star with an elliptical Gaussian brightness distribution, the ellipticity e is reduced to

$$e = \left| \frac{1 - a/b}{1 + a/b} \right|$$

where a and b are the major and minor axes. This should simply serve for illustration purposes, as in general stellar profiles are very different. In the moment description, a perfectly round source has an ellipticity of 0, and a highly elliptical one will be close to 1.

Parameters

- **Name filter:** The Ialyzer will process all images in a given directory, unless a name filter such as *ngc1234_night1_**OFC.fits is provided. You can also enter here the name of one individual image.
- **Title of the analysis:** You can specify an unambiguous string which will serve as a title for the analysis. In this way you can run Ialyzer with different configuration settings without overwriting previous results.

- **Smoothed spline interpolation:** Applies a strong smoothing to the FWHM field during spline interpolation. The result is more idealised and representative for what the optics delivers.
- **Unsmoothed spline interpolation:** Only a minimum of smoothing is applied to the data field before spline interpolation, in order to reject outliers.
- **Show median FWHM contour line:** Overplots a thick white line indicating the median level of the FWHM map.
- **FWHM units:**
 - **arcsec:** The FWHM map will be displayed in units of arcseconds
 - **pixel:** The FWHM map will be displayed in units of pixels
 - **minFWHM:** The FWHM map will be normalised to the smallest FWHM, i.e. this is a relative scaling as compared to the previous two.
- **FWHM colour scaling:**
 - **Full range:** For each image the min and max FWHM are determined, and the full colour range will be used for display. Images are treated independently of each other.
 - **max (in % of min):** If set to a numeric value of e.g. 50, the upper FWHM limit shown will be 50% larger than the smallest one. Images are treated independently of each other.
 - **absolute:** The lowest and highest FWHM shown. This is the same for all images. Units are either in arcsec or pixels, depending on which ones you chose for the display.

How to run Imalyzer

Imalyzer requires that you run the *Create source cat task* first.

Select the directory and optionally a subset of images, then click on *Analyse PSF*. This has only be done once for that particular set of images. It will create a sub-directory called

```
SCIENCE/cat/ksb/
```

in which the PSF analysis results are stored per image. Once done, set your plotting options and click on *Make plots*. This will do the FWHM spline fit and also read the output from the PSF ellipticity analysis for the stick plot.

Results are collected in an interactive html file (example) which is automatically displayed provided you have ‘firefox’ installed. If you use a different browser, point it to

```
SCIENCE/imagequality/<yourtitle>.html
```

Interpreting Imalyzer results

With several hundred stars the Imalyzer results should be fairly robust. Given that no polynomials are used in the FWHM fit, it represents the actual FWHM variations measured, albeit smoothed on a scale on which FWHM variations appear plausible. The ellipticity indicated is a blend of individual data and an iterative fit.

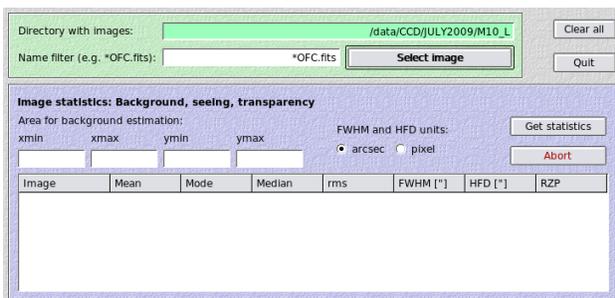
- With well-aligned optics the area of lowest FWHM should appear close to the centre of the image (or the optical axis), which might be outside the currently investigated chip. If this is not the case, then probably something is misaligned, or the FWHM variations are so small that the variations indicated are negligible. If the plot is clearly skewed, then chances are that there is something misaligned. This could either be the optics, or a tilt of the CCD with respect to the focal plane, or else.

- Ellipticity and FWHM can be correlated: If something is misaligned in the optics, then then the PSF e.g. becomes more elliptical in the corners of the field (and therefore the FWHM measured rises, too). It can also work the other way round: if for example the seeing increases, the PSF gets rounder as optical imperfections are blurred. Likewise, if the seeing gets better, the PSF reveals more imperfections. With undersampled data the PSF can become quite elliptical, but since it is so compact this might very well be of no importance for your scientific purposes.
- With significantly undersampled data the accurate determination of FWHM and ellipticity can be unstable. Imagine a perfectly round PSF with a FWHM comparable to one pixel. If the star falls right on the boundary between two pixels, it will essentially appear twice as long as broad and the FWHM will be artificially increased, compared to the situation where the star falls onto the centre of a pixel.

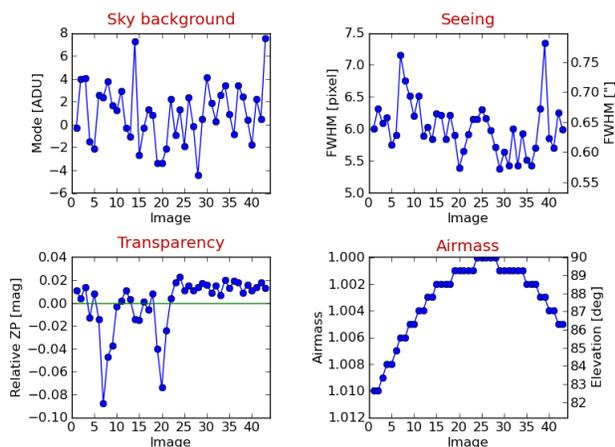
Therefore, when interpreting IAnalyzer results, you should always take both FWHM and ellipticity into account. Don't forget to look at the colour scale bar indicating the FWHM range displayed, in particular if one of the two non-absolute scalings was chosen. Knowing the observing conditions and the telescope also helps.

For example, a slow change in optical alignment with increasing zenith distance will result in a slowly drifting FWHM pattern, whereas optical realignment between exposures shows up as a sudden jump. In images which short exposure times the FWHM map may be misleading as seeing variations can cause noticeable distortions. Or in a closed tube assembly rising warm air on one side of the tube can lead to increased FWHM in some part of the image, mimicking a tilt of the CCD.

2.14.3 Image statistics



This module tabulates basic image properties and displays some of them graphically:



Parameters

- **Name filter:** The statistics module will process all images in a given directory, unless a name filter such as *ngc1234_night1_*.OFC.fits* is provided. You can also enter here the name of one individual image.
- **Area for background estimation:** Use these 4 fields to define an area of empty sky should the rest of the detector be covered by an extended source.
- **FWHM and HFD units:** Whether you want the FWHM and the HFD (half flux diameter) displayed in pixels or arcseconds.

2.14.4 Absolute photometric ZP

Results :		
Description	Value	FITS KEY
Photometric ZP		ZP_D
Zeropoint error		ZP_D_ERR
RMS of (MagOBJ - MagREF)		ZP_D_RMS
Number of sources used		ZP_D_NUM

This module can be used to obtain a crude absolute photometric zeropoint for an image, provided that you observed in similar passbands (either *ugriz* or *JHKs*). The image must have a valid WCS solution in the FITS header. The method used is essentially the same as the *direct photometric calibration* in the main reduction framework.

THELI will return:

- the zeropoint for the exposure normalised to an integration time of 1s, written to a *ZP_D* keyword
- the zeropoint error (*ZP_D_ERR*)
- the rms scattering between objects in the image and those in the reference catalogue (*ZP_D_RMS*)
- the number of sources that went into the fit (*ZP_D_NUM*)

Parameters

- **Photometric reference catalogue:** SDSS or MASS
- **Download server:** Switch between different servers should one of them not respond
- **Filter in which you observed:** Either one of *ugriz* or one of *JHKs*
- **Maximum photometric error:** Only sources with a *SExtractor* magnitude error equal to or less than this value will enter the fit.

Warning: This approach does not take into account colour terms arising from different total throughput curves between your telescope/filter/camera combination and the one used for SDSS or 2MASS. If you need absolute zero-points better than about 0.1 mag then you should consider fine-tuning the ZPs using a comparison of instrumental stellar tracks against those in the Pickels library, or a classical approach based on standard star observations.

2.14.5 Animate

The resampled images created during coaddition have accurate WCS information in their headers and can therefore be animated to check for moving or variable objects. THELI will create animated GIFs from the data and stores them as

SCIENCE/coadd_XXX/movie/anim_i-j.gif

Examples, with and without *proper motion vector* during coaddition (left, respectively right):

Parameters

- **Select coadditions:** After entering a SCIENCE directory, you will be presented with a list of the coaddition directories found. You can select either one or several.
- **Number of fields along xly:** One usually does not want to produce a full-scale animated GIF, but concentrate on smaller sections. These two parameters control into how many sections the image is split along x- and y- directions. Counting starts at the lower left. For example, if you set $x=3$ and $y=2$, then the field labelled 2-3 would correspond to the third image from the left, in the second row from the bottom.
- **Field overlap:** This is the overlap in pixels between fields.
- **Select field(s):** Choose here for which field(s) you want to create or display an animation.
- **Select chip:** For multi-chip cameras, you must specify from which chip the resampled images should be taken. This option is not visible for single-chip cameras.
- **Dynamic range:** The animated GIF has a dynamic range of 8 bit, only. Therefore you must select which range should be displayed. Two options are available:
 - **Manual:** Enter *min* and *max* levels (black- and white-point) for the GIFs. Since images are sky-subtracted, *min* should be negative and *max* positive, the latter with an absolute value about 5 times larger than *min*. For example, $min=-100$ and $max=500$.

- **Automatic:** Will apply a ds9-style automatic z-scale range. You can adjust the contrast value from -9 to +9.
- **Mask pixels with weight zero:** Pixels which have a value of zero will be masked (set to zero) in the resampled image. If your data has a lot of hot pixels, you can suppress them in this manner in the animated GIF.
- **Frame delay:** The delay between images in the animated GIF in units of 1/10th of a second. Note that this delay is ignored by the *Animate field* button which simply displays the images as rapidly as possible (it does not load the animated GIF but the individual frames)
- **Create animation(s):** Creates animated GIFs for all selected fields
- **Animate field:** Shows the animation for the currently selected field. Note that the frame delay is ignored by this function.

2.15 Near-infrared data

While near-infrared detectors are fundamentally different than CCDs, the general data reduction steps required are essentially the same. Therefore, once you understood how to reduce optical data with THELI, you already have most of the knowledge required for the near-IR regime. However, instrumental and atmospheric effects make reduction of near-IR data tricky. The sections below summarise how you deal with this in THELI.

2.15.1 Crosstalk

Crosstalk can in general be corrected well, provided that it is spatially stable. The latter is not always the case for near-IR detector arrays. In particular recent HAWAII2 sensors with multiple parallel readout sections can show crosstalk in form of compact positive and negative ghost images whose amplitude varies between readout sections. THELI assumes that the amplitude is the same, therefore the correction will only partially remove the effect (if at all). If you know in advance that this will be a problem for your science case, then consider choosing different camera rotator angles for your observations.

2.15.2 Reset anomaly and imaging equilibrium

Some detectors will exhibit a *horizontal or vertical brightening* towards one edge of the readout quadrant. It is unstable and in general depends on temperature, exposure time, the number of previously executed exposures, and can also have an erratic component. You can correct for it using the *Collapse correction*.

In very rare cases, such as for LIRIS@WHT (the only one the author is currently aware of), another problem occurs. If you take n exposures at a certain dither point and then move on to the next dither point where you repeat the same n exposures, piece-wise background correction can become necessary. In other words, even if the physical sky background does not change, the instrumental background pattern can change in a repeated fashion: it is the same for all k -th exposures in the series, but different from the pattern inherent to all $(k+1)$ -th exposures (and so on). To apply such a group-wise background correction, use the *spread sequence* task and then proceed as normal. THELI will take care of the rest.

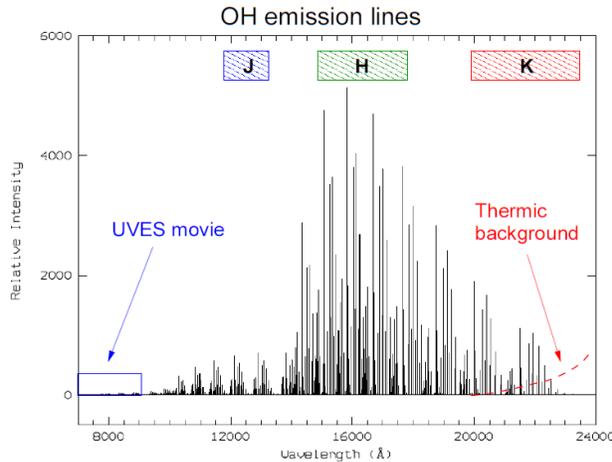
2.15.3 Sky background

Characteristics

Optical and near-infrared imaging differ most when it comes to atmospheric sky background. In the near-IR the background level changes on time scales of a few minutes, and on angular scales of a few arcminutes. Some nights are very stable and one can calculate a single static background model that will suit all exposures taken within a e.g. 20

minute window. Other nights will be highly unstable and it is necessary to create individual background models for each exposure (dynamic modelling). More about dynamic and static models can be found [here](#).

In general, *H*-band is affected most by this highly variable *airglow*, as shown by this night sky spectrum:



Exposure times

Exposure times in near-infrared imaging are rather short, for two reasons:

- **The background is variable:** Frequent dithering is necessary to achieve good background subtraction.
- **The background is high:** Exposure times must be short to avoid non-linearity or object saturation.

The following table contains typical night sky surface brightnesses in mag/sq. arcsec, and characteristic exposure times for detectors with 0.1-0.2 arcsec/pixel scale. Integration times for *Z* are for near-IR detectors, for optical detectors they are usually longer (optical and near-IR *Z*-band filter curves are not necessarily identical):

FILTER	<i>B</i>	<i>V</i>	<i>R</i>	<i>I</i>	<i>Z</i>	<i>J</i>	<i>H</i>	<i>Ks</i>
Surface brightness	22.7	21.9	21.0	20.0	19.0	16.0	14.5	13.5
Exposure time [s]	600	600	300	300	60	30	20	10

Narrow-band filters

Exposure times with near-IR narrow-band filters should be kept as short as possible, and made as long as necessary in order to reach background limited images. There is a trade-off depending on which narrow-band filter you are using. Exposure times up to 300s are possible for filters where less airglow is present.

Co-averaging, or co-addition

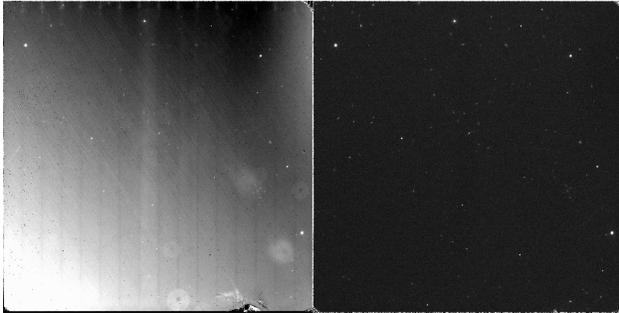
One problem due to short exposure times in the near-infrared is that for high-galactic latitude fields the number density of discernible sources can become very low. This can be critical from an astrometric and photometric point of view, as internal data calibration is difficult. Most observatories offer the possibility to average a sequence of *k* exposures on-chip before the image is read out. In this manner the S/N is improved without running into problems with non-linearity, saturation or the 16bit dynamic range. This process is sometimes called *co-averaging* (or *co-adding*, depending on the implementation), the individual integration time is called *DIT*, and the number of co-averaged or co-added exposures *NDIT*.

2.15.4 Background modelling

When you reduce near-IR data the first time, you'll probably be surprised that after flat-fielding the image appears everything else than flat. The amplitudes of the residual background variations from sky and instrumental contributions

are usually much larger than the fluxes from astrophysical objects. Frequently, you won't even be able to discern any objects at all in the data after flat-fielding.

The following example displays a 6x10s exposure ($DIT=10$, $NDIT=6$) taken with one of the four 2kx2k detectors of HAWKI@VLT in K_s -band. Shown in the left panel is the appearance after flat-fielding, and the right panel displays the same image after subtraction of a (dynamic) background model:



1-pass modelling

Which background model is most suitable for the data does not only depend on the atmosphere, but also on the target. If the field is sparse and contains point sources only, then very often a simple (static or dynamic) background model with basic *min-max rejection* and median combination is sufficient. Leave the DT , $DMIN$ and $SIZE$ thresholds empty.

If the field is crowded and you have many exposures, switching on additional object masking can avoid the background model to be biased by object fluxes. Use some explicit settings for the *SExtractor* detection thresholds, e.g. $DT=10$ and $DMIN=10$. Thresholds must be chosen high enough such that no background features are masked. The mask images are called **OFC.mask.fits* and are collected under

```
SCIENCE/MASK_IMAGES
```

once processing is finished. You should eyeball them and see if sky features were detected. In addition, you can use min-max rejection as outlined above.

Iterative or 2-pass modelling

The main problem with the 1-pass approach just outlined is that good object masks can only be created from a flat image. This is particularly important if the field contains extended sources, parts of which are too faint to be visible in the flat-fielded exposures. If this hidden flux is not masked, the background model will be biased, leading to significant over-subtraction. Very often this will only be visible in the coadded image, showing up as dark areas or patches scattered around brighter sources. In such cases an iterative approach is advisable:

1. Create the static or dynamic model as normal. You can safely switch off object masking by leaving the DT and $DMIN$ *parameter fields* empty.
2. If you chose a static (dynamic) background model, set the *window size* to a zero (non-zero) value. In both cases *rescale* the model.

Once THELI has finished, your images will have the *status string OFCB* written into their file names and should look very close to flat (if no iterative modelling is required you can leave here and jump to the *weighting* process). The following sub-directories are now in the *SCIENCE* directory:

```
SCIENCE/SPLIT_IMAGES
SCIENCE/OFC_IMAGES
SCIENCE/MASK_IMAGES
```

The first contains the raw data after splitting. The second contains the flat-fielded (and possibly dark-subtracted exposures), and the last one contains the mask images.

- Repeat step 1, but this time set *DT* and *DMIN* to their default values to trigger object masking. THELI will recognise that the task is run a second time (due to the presence of *OFCB* images). In this case the object detection will take place on the temporary *OFCB* images with good background, and the corresponding pixels will be masked in the original, flat-fielded only *OFC* images. The newly created mask images from these data are then also applied to the *OFC* images, resulting in a new batch of *OFCB* exposures. The previous ones will be kept in a *OFCB_IMAGES_IPASS* directory for comparison.

The process is straight forward and should provide you with fairly flat images.

Using separate blank SKY exposures

If your targets are really extended and/or very faint, then you should seriously consider observing blank SKY fields. They will be automatically used in the correct manner once collected in a separate directory and defined in the *Initialise* section. No additional settings have to be made. You should check though that the images have correct *MJD-OBS* header keys after *splitting*.

2.16 Advanced usage

This section provides some more background information and contains a collection of useful tricks and techniques. Some of them are meant for troubleshooting, others will help you reduce data more efficiently.

2.16.1 Change processing status

THELI indicates the processing status of the data in the file names, between the chip number and the suffix. A typical file name will look like

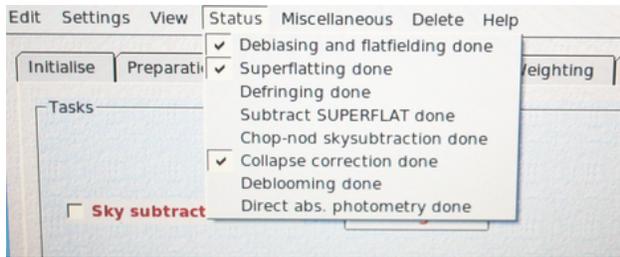
```
filename_<chipnumber><status>.fits
```

The *status* string encodes the various steps you have done and which resulted in a modification of pixel values. With advancing data reduction, more identifiers are appended to *status*. The following identifiers are used:

ID	Meaning
OFC	Data has been overscan corrected, flatfielded, and the overscan regions have been cut away. This string will always be appended by <i>Calibrate data</i> , even if you did not overscan correct or flat-field the data.
B	<i>Background modelling</i>
H	<i>Chop-nod sky correction</i>
C	<i>Collapse correction</i>
D	<i>Debloating done</i>
P	<i>Direct absolute photometry per chip</i>
.sub	<i>Sky subtraction</i>

THELI stores this status internally and provides it as an essential argument to many scripts. It can happen that *status* does not reflect anymore the actual status files have on disk. Reasons are manifold, for example the user started a new data reduction without resetting the logfile, or data was modified/deleted manually outside THELI and the processing is resumed with data at an earlier processing status.

Before a script is executed, THELI checks if images with *status* in their file name do exist, and if none are found the user will be notified. In this case one can tell THELI the current processing status by setting or removing the corresponding tickmarks in the *Status* menu:



Afterwards, simply re-select the task you wanted to execute and continue.

2.16.2 Processing only a particular chip from a multi-chip camera

Assume you have a multi-chip camera, and one of the following situations occurs:

- One of the chips needs additional treatment (for example *collapse correction*) that you do not want to apply to the other chips
- You want to process one of the chips with slightly different parameter settings
- You want to quickly test some parameters without running through all chips

To understand how to proceed, let's have a look at how THELI's *parallelisation* works.

THELI's parallel manager

In all relevant cases, you deal with parallelised scripts, easily recognisable by their syntax displayed in the command window, e.g.

```
./parallel_manager.sh <script_name.sh> [arguments]
```

What *parallel_manager.sh* does is to call *script_name.sh* as many times as *CPUs have been specified*. It also appends additional arguments (the chip numbers to be processed by each CPU), which are required by *script_name.sh* but not shown to the user. If, for example, we work with a camera having 8 chips, and we want to use 4 CPUs, *parallel_manager.sh* would launch the following four instances:

```
./<script_name.sh> [arguments] "1 5"  
./<script_name.sh> [arguments] "2 6"  
./<script_name.sh> [arguments] "3 7"  
./<script_name.sh> [arguments] "4 8"
```

Therefore, if you want to just process e.g. chip number 2 with different settings, you would manually remove *parallel_manager.sh* and write instead:

```
./<script_name.sh> [arguments] 2
```

If you want to reprocess chips 1, 3 and 4, you would modify the syntax to:

```
./<script_name.sh> [arguments] "1 3 4"
```

Note: Since you manually remove *parallel_manager.sh* from the command, no parallelisation takes place anymore.

Different data reduction statuses for various chips

Using the scheme outlined above, you can not only change parameters for a particular chip, but you can also run it through additional data reduction steps. In this case, however, extra care needs to be taken, because the *processing*

status of that chip will change. Hence the status string in the file names is not the same anymore for all chips, and thus THELI will not be able to identify all chips in the next reduction steps. This can be corrected as outlined in the following.

Example: Data were debiased and flat-fielded, and the status string *OFC* appears in the file names. You then want to run chip 3 through the *collapse correction*. The following steps must be done:

1. Activate the collapse correction, and change the syntax to:

```
./process_collapsecorr_para.sh MAINPATH SCIENCEPATH OFC 3
```

2. Data from chip 3 will now show the status string *OFCC*, whereas the others are still called *OFC*. The previous *OFC*-images of chip 3 are located in a new sub-directory called *OFC_IMAGES*. You must now manually rename all other *OFC* images to *OFCC*, using:

```
cd /MAINPATH/SCIENEPATH

rename OFC OFCC *OFC.fits
```

Alternatively, you could also rename all *OFCC* images to *OFC*:

```
cd /MAINPATH/SCIENEPATH

rename OFCC OFC *OFCC.fits
```

In the latter case you must also change the *processing status* in THELI, as THELI assumes that images have been collapse corrected and therefore carry *OFCC* in their file names (which they don't, because you renamed everything to *OFC*). Simply remove the corresponding tickmark next to *Collapse correction done* in the *Status* menu.

2.16.3 Processing multi-colour data sets

Imagine you observed a target in several filters, *B*, *V* and *R*, and you arranged the raw data in the following way:

```
/mainpath/BIAS
/mainpath/FLAT_B
/mainpath/FLAT_V
/mainpath/FLAT_R
/mainpath/TARGET_B
/mainpath/TARGET_V
/mainpath/TARGET_R
```

If you follow the standard reduction approach in THELI then you would reduce one filter at a time: you declare one flat and target directory in the *Initialise section*, e.g. *FLAT_B* and *TARGET_B*.

This is a bit tedious with several filters, and can be done more efficiently, leaving the directory entries unchanged.

Preparation

When you tell THELI to *split* the raw data, it will issue the following commands in the large white *command window*:

```
./check_rawfilenames.sh /mainpath/ BIAS
./process_split_<instrument>.sh /mainpath/ BIAS
./check_rawfilenames.sh /mainpath/ FLAT_B
./process_split_<instrument>.sh /mainpath/ FLAT_B
./check_rawfilenames.sh /mainpath/ TARGET_B
./process_split_<instrument>.sh /mainpath/ TARGET_B
```

The entries listed in the command window can be edited. By means of copy & paste you can manually rearrange them a bit, including the other filters:

```
./check_rawfilenames.sh /mainpath/ BIAS
./check_rawfilenames.sh /mainpath/ FLAT_B
./check_rawfilenames.sh /mainpath/ FLAT_V
./check_rawfilenames.sh /mainpath/ FLAT_R
./check_rawfilenames.sh /mainpath/ TARGET_B
./check_rawfilenames.sh /mainpath/ TARGET_V
./check_rawfilenames.sh /mainpath/ TARGET_R
./process_split_<instrument>.sh /mainpath/ BIAS
./process_split_<instrument>.sh /mainpath/ FLAT_B
./process_split_<instrument>.sh /mainpath/ FLAT_V
./process_split_<instrument>.sh /mainpath/ FLAT_R
./process_split_<instrument>.sh /mainpath/ TARGET_B
./process_split_<instrument>.sh /mainpath/ TARGET_V
./process_split_<instrument>.sh /mainpath/ TARGET_R
```

The *check_rawfilenames.sh* scripts have to be executed before the *process_split* scripts, hence they can be grouped together for better visibility.

Calibration

Process the bias images normally.

When processing the flats, expand the commands listed as follows:

```
./parallel_manager.sh process_flat_para.sh /mainpath/ BIAS FLAT_B
./parallel_manager.sh process_flat_para.sh /mainpath/ BIAS FLAT_V
./parallel_manager.sh process_flat_para.sh /mainpath/ BIAS FLAT_R
./parallel_manager.sh create_norm_para.sh /mainpath/ FLAT_B
./parallel_manager.sh create_norm_para.sh /mainpath/ FLAT_V
./parallel_manager.sh create_norm_para.sh /mainpath/ FLAT_R
```

Same with the target exposures. Do not forget to edit the *FLAT_X* string (!):

```
./parallel_manager.sh process_science_para.sh /mainpath/ BIAS FLAT_B TARGET_B
./parallel_manager.sh process_science_para.sh /mainpath/ BIAS FLAT_V TARGET_V
./parallel_manager.sh process_science_para.sh /mainpath/ BIAS FLAT_R TARGET_R
```

Weighting

The globalweight depends on the filter. Therefore, **make sure you do create the weights for each filter immediately after you create the globalweight**, as the latter will be overwritten:

```
./par... create_global_weights_para.sh /mainpath/ FLAT_B_norm TARGET_B
./par... create_weights_para.sh /mainpath/ TARGET_B OFC
./par... create_global_weights_para.sh /mainpath/ FLAT_V_norm TARGET_V
./par... create_weights_para.sh /mainpath/ TARGET_V OFC
./par... create_global_weights_para.sh /mainpath/ FLAT_R_norm TARGET_R
./par... create_weights_para.sh /mainpath/ TARGET_R OFC
```

If you masked satellites by hand, you must include the *transform_ds9_reg.sh* scripts before *create_weights_para.sh*.

Astrom/Photom

Now that the pre-processing is done you can combine all images in a common directory, e.g.

/mainpath/TARGET_ALL

and treat them as if they belong to the same filter. Enter *TARGET_ALL* as the new SCIENCE directory in the *Initialise* section. However, there are three things you have to keep in mind:

Note:

1. If you want to do absolute photometry you must create parking directories for filters *V* and *R* while dealing with filter *B*, and so on.
2. When creating source catalogs, remember that the object density can vary significantly from filter to filter, e.g. you could have much less sources in *B* than in *R*. Make sure that the detection settings return a sufficient number of sources in all filters. If this is not possible, treat filters independently (i.e. use parking directories).
3. During astrometry with *Scamp*, make sure that you set
 - *ASTRINSTRU_KEY* = *NONE*
 - *PHOTINSTRU_KEY* = *FILTER*
 - *STABILITY_TYPE* = *INSTRUMENT*

In this way *Scamp* will calculate a **common astrometric solution** for all filters, while doing independent photometry for each of them.

A **common astrometric solution** can be very useful if e.g. one of the filters features very few sources (such as observations in *U*-band), making distortion correction with that data alone difficult. By setting *ASTRINSTRUKEY* to a non-existing keyword, e.g. *NONE*, all exposures in all filters will be used to create a common distortion model. This is also useful for difference imaging or other applications where you want the same source to end up exactly in the same pixel.

Coaddition

Sky-subtraction can be performed normally, pretending all exposures in the *TARGET_ALL* directory belong to the same filter.

The coaddition itself has to be done per filter. In the *configuration menu* for the coaddition you can select from a small pull-down menu which filter you want to stack. There is no need to make images in the other filters invisible by moving them to parking directories.

Identical image sizes

Note: For a consistent astrometric projection of the coadditions it is essential that you choose *identical reference coordinates* in the configuration. Do not leave these fields empty with multi-colour data sets! If you do not do that, images cannot be automatically registered to identical sizes with objects in identical positions (should this be important for your research).

Once you created all coadditions with *identical reference coordinates*, you will realise that they still have different image geometries. This is due to variations in the dither pattern etc. Fortunately, this can be fixed very easily due to a consistent astrometric projection on sky and without further resampling:

Open the *Prepare colour picture* dialog in the *Miscellaneous* section of the GUI's main menu bar. A list of all coadd_XXX directories will appear, in our case

```
coadd_B
coadd_V
coadd_R
```

(provided that the FILTER keyword in the headers is *B*, *V* and *R*). All you have to do next is to click on *Get coadd images*. When this step is done, you will find the registered coadded images and their weights as

```
/mainpath/color_theli/B.[weight].fits  
/mainpath/color_theli/V.[weight].fits  
/mainpath/color_theli/R.[weight].fits
```

Different processing steps required

What, for example, if the *R*-band images need to be defringed, whereas this step is not necessary for the other filters? This will lead to a problem with the combined processing approach outlined above, as the *processing status* of some of the images will have changed. Therefore THELI will not identify the other images correctly anymore.

The work-around is straight forward once you are done with *background modelling*, or *collapse correction*, etc):

1. A backup copy of the data just before the defringing step already exists (in that case it would be called *OFC_IMAGES*).
2. The defringed images will have the `status` string *OFCB* in their file names. Rename the images manually. THELI provides you with a useful command line tool called *rename*, which you can find in

```
THELI/theli/bin/Linux[_64]/
```

The syntax is easy, in this case

```
cd /mainpath/TARGET_R  
rename OFCB OFC *OFCB.fits
```

This will replace the string *OFCB* by *OFC* in all files **OFCB.fits*.

Note: Some Linux distributions come with a (perl) executable that is also called *rename*. It does something similar, albeit with a different syntax. If you included the THELI binary directory in your PATH variable, it might shadow the perl command.

3. Lastly, you need to make THELI forget that it defringed some data: simply remove the defringing entry from the *Status* menu bar.

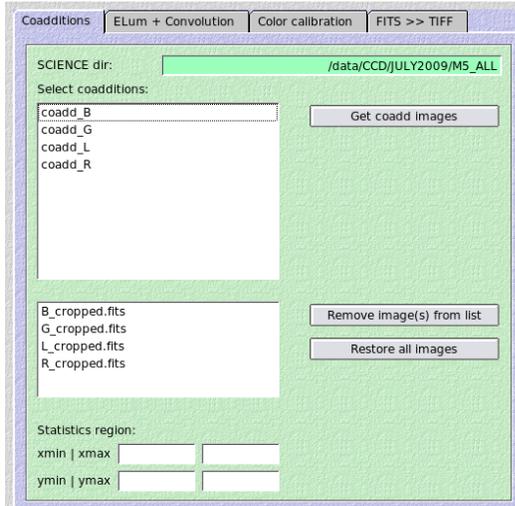
2.17 Prepare colour picture

If you ever tried to render a nice-looking three-colour picture out of CCD images, you will have noticed that this is far from trivial. The main difficulties are:

- Exact alignment of the coadded images in all filters (hereafter: *channels*)
- Find a good colour balance
- Truncate the dynamic range in order to create 16-bit TIFF images

THELI does not go all the way to a final image. However, it can prepare suitable scaled images that can be readily processed in third-party software without much further processing.

2.17.1 Coadditions



The first step that has to be done is to select those stacked images that should be combined for a colour picture. THELI assumes that all necessary *coadd_*xxx directories can be found in the same SCIENCE directory. If you followed the *guidelines for processing multi-colour data sets*, then this will be the case.

All processing takes place in a new sub-directory

```
SCIENCE/color_theli/
```

If this directory already exists, it will be unambiguously renamed into

```
SCIENCE/color_theli_backup_<timestamp>
```

Preparations

1. **SCIENCE dir:** Upon opening this dialog, THELI will fill in the currently processed SCIENCE directory. You can change to any other directory you want. If coadditions are present, then these will be shown in the *Select coadditions* list.
2. **Select coadditions:** Mark the coadditions you want to use for the colour image. Several entries can be highlighted at the same time.
3. **Get coadded images:** Upon clicking here, THELI will link the coadded images selected into

```
SCIENCE/color_theli/
```

If a coaddition was labelled *coadd_*xxx, the corresponding image and its weight will appear as

```
SCIENCE/color_theli/xxx.fits
SCIENCE/color_theli/xxx.weight.fits
```

Once all images are linked, an automatic cropping process kicks in that trims all images such that they have precisely the same size, and such that a particular object will have identical pixel coordinates in all channels. Once this is done, the images will be called

```
SCIENCE/color_theli/xxx_cropped.fits
SCIENCE/color_theli/xxx_cropped.weight.fits
```

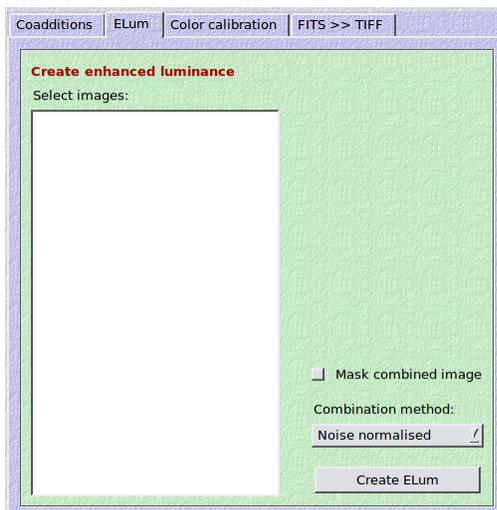
Note: This works only if the coadditions were created with *identical reference coordinates*! This step does not

involve any resampling anymore as the images already have identical astrometric projections (if the coadditions were done with identical reference coordinates).

Optionally, you can remove some of the cropped images from the list (they will not be deleted from the hard disk, just ignored), or the full list can be restored. The latter is useful if you want to resume work in an already existing *color_theli* sub-directory.

4. **Statistics region:** Define an area which is (ideally) free of any stars and nebulosity. This can be as small as 5x5 pixels if necessary and should correspond to the darkest region in the images. In the later colour picture, this area will appear with a neutral grey background. Though the coadded images are already sky-subtracted, small residuals could be left, which will be removed in this way.

2.17.2 Elum



Create enhanced luminance

If you took red, green and blue exposures, you will obtain a significantly better (deeper) colour image if you combine the three exposures into a synthetic luminance image.

If you have an unfiltered (clear filter, luminance filter) coadded image, you can add the filtered exposures to it to get even greater depth. However, this will only yield an improvement if

- the combined transmission curves of the filtered exposures cover a similarly large spectral range as the luminance filter
- if the total exposure time of each of the filtered images is comparable to or larger than the one of the unfiltered/clear/luminance exposure.

Mask combined image

Pixels which are not covered by all exposures (as judged from the individual weight maps) will be set to zero in the combined image.

Combination method

The following two options are available:

- **Noise normalised:** The background noise of the individual images will be measured in the *statistics region* specified previously. Images are normalised with respect to their noise and then added (this is similar to the chi-square sense). In this way images with lower background noise contribute more than those with higher noise. The result is called

```
SCIENCE/color_theli/elum_chisquare.fits
```

You need to be careful, though. For example, at a dark sky location the background increases significantly from blue to red to unfiltered exposures. If you noise-normalise such data, it can happen that the blue exposure which might carry comparably little signal starts dominating the other exposures, in particular the deep luminance channel.

- **Mean:** The average of the images is calculated. The result is called

```
SCIENCE/color_theli/elum_mean.fits
```

If you want to combine plain red-, green- and blue-filtered images, then this is the best approach as it mimics what a luminance filter would have seen.

Note that if you set `PHOTINSTRU_KEY = FILTER` during *astrometry*, objects in the three channels will have the same average brightness. If one of them was taken under bad transparency conditions, it will be up-scaled which increases its noise, and can therefore lead to inferior results. Same if you use an image which very low signal-to-noise ratio as compared to the others.

2.17.3 Color calibration

The idea behind colour calibration is to have a solar-type G2V star appearing white in the colour picture. If this only depended on the total telescope throughput (mirror reflectivity, corrector lenses, filter transmission and quantum efficiency), fixed calibration factors could be used. However, the atmosphere changes its transparency continuously, either because of changing airmass, dust, cirrus, humidity etc. Atmospheric extinction has a significant impact even in very clear nights when targets are rising or setting, affecting the calibration factors by 10% or more. If exposures are taken in different nights, the balance can be off even more. Variations on the level of a few percent are easily seen by the eye in a colour picture.

The best way to obtain the colour calibration is to identify solar-type stars directly in the images. Like that

- no external standard star observations are necessary
- it works under non-photometric conditions
- no extrapolation from external observations is necessary.

However,

- G2 stars are rare (only 6% of all stars are have type G2)
- known G2 stars in the image could be saturated as only bright stars have spectra taken
- the reference star could be variable
- no reference stars are present in the field (or their spectral types are unknown).

Photometric colour calibration (PCC method)

To work around these problems, one can drop the requirement of selecting known G2 stars as standards, and take G2-alike stars instead. These can be selected photometrically from an external reference catalogue by means of their colour indices. In the classic Johnson-Cousins *UBVRI* system, G2 stars have on average $B-V = 0.65$ and $V-R = 0.5$. In SDSS *ugriz* filters we have $u-g = 1.43$ and $g-r = 0.44$. Once a sample of suitable field stars is known, they can be identified in the images and calibration factors are readily determined.

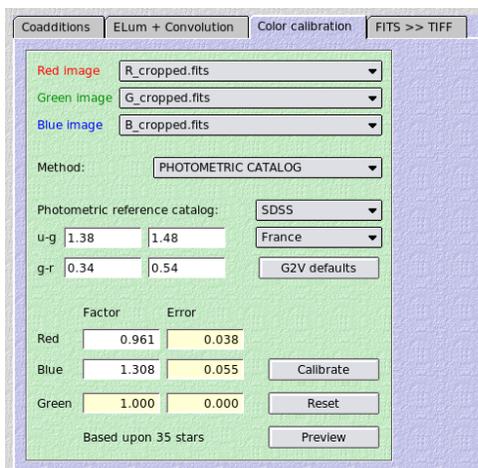
The average white approach

What if no G2-alike stars are found, i.e. in under-dense regions or in images with a very small field of view? In this situation a good white point can still be obtained by making *all stars on average white*. Compared to the PCC method the blue and green channels appear enhanced. The reason is that the average stellar population is somewhat reddish (80% of all main sequence stars in the solar neighbourhood are red dwarfs).

The calibration factors will also depend on the line of sight: stars in the galactic halo have different characteristics than those in the disk, and towards the galactic bulge things change once more. Thus, one will obtain slightly different calibration factors depending whether one looks outside the disk, into the disk, or towards the galactic centre. Galactic extinction (interstellar reddening) will be partially removed by this method, which might be a desirable effect (or not).

Calibration in THELI

THELI offers both the PCC method as well as the average white approach.

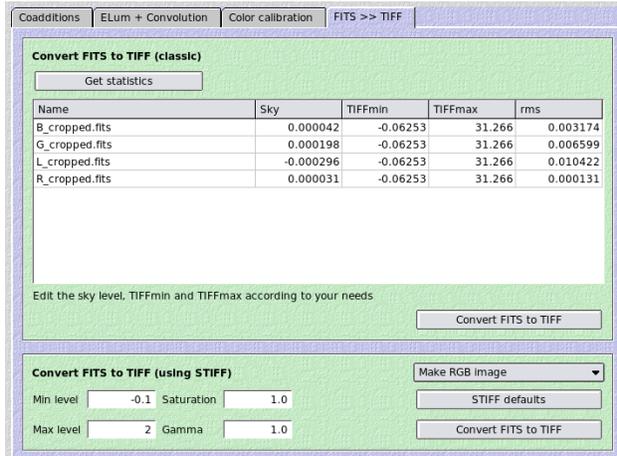


Calibration is fairly easy:

1. Identify the red, green and blue channels
2. **Method:** Select *Photometric calibration* or *Average white*
3. If you go with the PCC method, several other options are available:
 - **Photometric reference catalog:** Either *SDSS* or *NOMAD*. SDSS covers about 11600 square degrees around the northern galactic cap, with a few extensions down into the galactic plane. It is based on CCD photometry which is accurate to within about 1%. NOMAD on the other hand covers all the sky, however it is a conglomerate of different surveys mostly based on photographic plates. Photometric accuracy is significantly worse than SDSS, variations of 10% or more are not uncommon and also depend on sky position. As soon as SDSS-like photometry becomes available for all the sky (*Pan-STARRS*, *Skymapper*), NOMAD will be retired from THELI.
 - Choose from various download servers should one of them be unavailable
 - Optionally, edit the B-V and V-R (u-g and g-r) G2V star colour index ranges (i.e. widen them should no sources be found). This compromises the idea of the PCC approach though as more non-G2 stars will be included.
4. Click on *Calibrate*. Objects will be detected in the images and non-saturated stars extracted. In case of the PCC method only stars matching the G2-alike stellar sample will be used, whereas with the average white approach all non-saturated stars will be kept. The red and blue images are then scaled such that these stars have on average the same brightness as those in the green image.

- Preview the result. This is not a well-balanced image in terms of contrast and black- and white-points, but the colours are accurate. You can modify the calibration factors manually and see their impact. Sometimes, NOMAD calibrations should have their blue values reduced by 10%, and the red values increased by 5% to match those of SDSS. These are noisy though and depend on the line of sight.

2.17.4 FITS >> TIFF



The last step is to convert the FITS images to 16-bit TIFF format. The latter can be exported into e.g. *Photoshop* or other programmes for further processing. Note that for *Photoshop* *FITS Liberator* is available, which can directly import the FITS images. The next release of *FITS Liberator* will be a stand-alone tool that does not require *Photoshop* anymore.

Before you can proceed in *Photoshop* or with *FITS Liberator*, you must fine-tune the sky background and apply the colour calibration factors:

Click on *Get statistics*, which will measure the mean background level and its rms in the *statistics section* defined previously.

Then choose between two methods to create the 16-bit TIFF images:

Convert FITS to TIFF (classic)

This will convert each FITS image to a stand-alone TIFF image, which you have to combine in an external programme to render the final colour image. You can manually override THELI's suggestions:

- **Sky:** This is the average background level measured in the *statistics section*. This value can usually be left unchanged.
- **TIFFmin:** This will be the black-level of the TIFF image, i.e. all pixels in the FITS image with values below this threshold will be represented as black in the 16-bit TIFF. THELI identifies the image with the largest background noise and multiplies this value with -6. The resulting value, *TIFFmin*, will be the same for all exposures. In this way the sky background is not clipped. Usually you can leave this value unchanged.
- **TIFFmax:** This is the white-level of the TIFF image, i.e. all pixels in the FITS image with values above this threshold will be represented as white in the 16-bit TIFF. *TIFFmax* is set arbitrarily 500 times higher than *TIFFmin*.

If you want to display even brighter image parts without saturation, increase this value, otherwise leave it (or decrease it). For example, you observed a galaxy with a bright core and don't want this core to be saturated.

Measure its brightness in the three colour channels (e.g. red=25, green=15, blue=12) and set *TIFFmax* to a value that is somewhat higher than the maximum of the three (e.g. to 30).

Warning: *TIFFmin* and *TIFFmax* must be identical for the red, green and blue channel, otherwise the colour calibration performed previously will be distorted. You may choose different values for luminance channels.

Using the *Convert FITS to TIFF* button will convert all images listed in the statistics table into TIFF images:

```
SCIENCE/color_theli/*.tif
```

If you want to use the corresponding FITS images (colour calibrated and sky subtracted) for e.g. *FITS Liberator*, use

```
SCIENCE/color_theli/*_2tiff.fits
```

Convert FITS to TIFF (using STIFF)

A slightly different approach uses the *STIFF* package. *STIFF* creates individual 16-bit TIFF images as well, but can also combine the three colour channels into a colour TIFF, such that you don't have to do this in external software anymore. In addition, *STIFF* offers to increase the colour saturation and to apply a non-linear gamma correction. However, these things can also be done in any classic image processing software).

STIFF parameters

Note: *STIFF* applies the sky background values listed in the statistics table, the same as used by the classic FITS to TIFF conversion.

- **Min | max level:** Same as *TIFFmin* and *TIFFmax* described *above*.
- **Saturation:** Only meaningful when creating a RGB colour TIFF
- **Gamma:** Non-linear gamma correction to the dynamic range.
- **Make RGB image:** Select here which images should be converted. If you choose *Make RGB image*, the red, green and blue images defined in the *colour calibration section* will be converted into a 3-colour RGB TIFF image, saved as

```
SCIENCE/color_theli/rgb_stiff.tif
```

If you want to convert a specific channel into a grey-level TIFF, then just select it from this pull-down menu. The resulting TIFF image will be displayed automatically and is stored as

```
SCIENCE/color_theli/*_stiff.tif
```

Only one TIFF at a time will be created when clicking on *Convert FITS to TIFF*.

2.18 Guidelines for visiting observers

Some advice from somebody who has a long record of observing, of supporting visiting astronomers, of doing service observations for others, and of seeing programmes fail and time wasted because people didn't listen or didn't care.

I have seen it all! At least I hope so. Well, thinking about it twice, I think I haven't. Most of this stuff is really avoidable, but it happens again and again. If only one programme is saved because you as an observer read this, then it was worth it! All observatories will thank you! And I will thank you too, because I have to answer less questions about data reduction ;-)

2.18.1 Planning

Do plan your observations well in advance! You will save time at the telescope, be more efficient, and come home with data that can actually be calibrated! You'll be happy! In addition, you will make your support astronomer and/or the staff executing your observations happy, and you will save the telescope operator lots of nerves.

- Use the exposure time calculators. Even though some are infamous, they'll put you in the right ballpark, and some of them actually work pretty well!
- Make visibility plots of your targets such that you can observe them efficiently at night when they are high on the sky, and not when they are 10 degrees above the horizon making the telescope, the telescope operator and the support astronomer ache. Use Chris Benn's excellent [Staralt tool](#).
- Think about how large your *dither pattern* has to be. Don't think about NOT dithering; there are hardly any science cases for that. Hmm, thinking about it twice, I think there aren't any at all!
- If you need standard stars, select them in advance. Keep in mind that for spectroscopy the spectral type of your standard star can be essential for telluric correction. If you have to dig through long lists at 3 a.m. in the morning to locate a suitable standard next to your target you're not doing anybody a favour, at least not your science case. Remember to observe standard star fields at significantly different airmasses, otherwise no extinction correction can be performed.
- Make sure the filters/grisms/prisms/slits you need are actually present at the telescope or in the instrument. If the observatory staff needs longer lead times to set up the instrument, then tell them well in advance. Actually, make sure the instrument you want to use still exists at the telescope! I'm not kidding, I had one observer who wanted several hours of the night to be executed with another instrument, but never mentioned this at any point anywhere before: O: "*Could we please switch to XX in Nasmyth focus for the rest of the night? Then I could also do this and that...*" SA: "*Errrh, XX was decommissioned two years ago and replaced by YY, which is only operated at Cassegrain focus and back on sky in three months. May I suggest a service proposal?*". Another one applied for an instrument that was actually mounted at a telescope in Australia instead of in Spain; this was flagged in due time though before entering the time allocation process.
- Check when your observations are taking place, and plan your travel in advance. Arrive early! Observatories are usually in remote locations and in case of bad weather they can be difficult to access.
- Make sure you know the actual overheads of observations!
- Have a backup plan if transparency or seeing are not as good as you may want.
- Read the instrument manual and the observatory's guidelines for observers in advance. If possible, familiarise yourself with data reduction as this can reveal better observing strategies.
- Check the moon distance, and the lunar illumination. You can't do any deep observations in blue filters in bright moonlight.

2.18.2 Substitute observers

If you cannot observe yourself and you have to send somebody else, make sure this person knows how to observe, and arrange with staff at the telescope to ensure proper support.

Ideally, make sure that your observer has a scientific interest that the data taken is being useful. If you send unmotivated observers, they might go to bed at 4 a.m. wasting the several good hours before sunrise (which the observatory staff will then likely use for their own purposes). Or they observe an entire night with a totally defocused telescope, which HAS happened on a 2.5m telescope where the observers are by themselves, unsupervised by a telescope operator or support astronomer.

2.18.3 Dithering

When planning your observations, make sure that you dither your exposures (apply small telescope offsets after each image). If you observe with a multi-chip camera, make sure that the dither pattern is wide enough to yield a sufficient overlap between neighbouring CCDs, and does not just barely cover the gaps between the CCDs. This helps astrometry and photometry enormously. THELI was designed to handle data with large dither patterns, so make use of it. Inhomogeneous effective exposure times in the coadded image can properly be taken into account using weight maps.

Remember: No dithering,

- no defringing
- no superflating
- no background modelling
- no correction for chip defects
- no removal of certain instrumental signatures
- no good signal-to-noise
- no good astrometric distortion correction
- no filled inter-chip gaps
- no good science. At least not as good as it could be.

2.18.4 Background models

If you observe empty fields and plan to superflat them or subtract a background model, then make sure that the dither pattern you apply is larger than the largest object in the field of view. Should this be unfeasible because you observed an extended target, then calculate the superflat from different pointings obtained in the same filter in the same night. If you know that you have to superflat your data or to create suitable background models, then there is **no way** around blank field observations.

If you observed very extended targets (comparable to or larger than your dither pattern), then you cannot calculate a superflat or sky background model from these data. In this case you must observe a neighbouring blank SKY field. Take 5-10 well-dithered images, exposure times can be shorter than for the main target as long as the background is recorded sufficiently well. Near-IR observers (hopefully) know all of this, but optical observers often don't. The overhead for these blank fields can be significant, but it's better to have somewhat shallower but well-calibrated data than a worthless pile of pixels.

THELI has full built-in support for SKY fields which are reduced automatically with the main stream of data and readily applied.

2.18.5 Extended low surface brightness objects

If you are observing extended very low surface brightness objects, visible or invisible in single exposures, then you must take great care in all background modelling and sky-subtraction steps. Choose a very wide dither pattern that is at least twice as large as an optimistic estimate of the extent of your target. If this is unfeasible, use a sufficiently large number of blank field exposures (at least 5-10).

Otherwise you cannot calculate a superflat or background model from your data that does not have a low-level imprint of your target. The same holds for the sky subtraction. If you choose to go the standard THELI way and model the sky background for each image, then choose very low detection thresholds and a smoothing scale that is much larger than the extent of your object. If you feel that this is too dangerous, then obtain a constant sky estimate using the methods offered in the *sky-subtraction task*.

2.18.6 DO THIS

- If you know in advance that there aren't any targets during some hours at night, then let the observatory staff know about it. They can slot in some service proposals and thus reduce the pressure on their queues. Or they'll use it for their own programmes and who knows, maybe you become co-author and famous in return for sharing your time. Same if observing conditions get unsuitable for your programme. Most likely there are several programmes in the service queue that can make use of bad seeing and/or bad transparency.
- Show up in time for the introduction at the telescope, even if you are an experienced observer. Things might have changed compared to last time when you visited, and there is almost certainly something which you forgot. A good support astronomer will also spend some time to discuss and refine your observing strategy.
- Take the time to fill out the online feedback forms or fault reports. It may seem tedious to you, but any detail you can think of, positive or negative, could be of great help to improve services. If nobody complains about anything, operations will assume that things are alright, and nothing will change. A lot of issues can be resolved the very next day and you have a better working telescope/instrument the following night.
- Try to sleep! Not at night, but during the day. If you pass out at 2 a.m. then you are not good for anything, not to mention for your science. Not all observatories offer full-time presence of a support astronomer who checks if whatever the tired observer is doing still makes sense (and support astronomers often have other duties at night as well). If you can't sleep because the shutters in your room are bad or because the air is too dry, ask for a different room or a humidifier. If you suffer from something, then most likely the observatory staff does as well, and thus there are solutions for it (different food, medication, a boxing sack for stress management or physical work-out).
- If you know you are getting car sick, bring medication. Access to some telescopes is only possible over steep or bad roads with many turns. Same for altitude sickness. Most observatories are between 2000m and 3000m altitude, which is high, but not so high that you wouldn't get adjusted within a day. Show up a day earlier to allow your body to adjust itself. Observatories at higher mountains will have specific regulations that are enforced. You don't do anybody a favour if you are sick in bed. Sun-blocker, lipsticks, mosquito repellent, etc are good ideas, and drinking a lot of water, too.
- Stick to the safety regulations. A broken leg because you didn't watch your tired steps at night will cause you surprisingly many painful hours before you reach a hospital, and it might not be a hospital you want to be treated in.
- Telescopes are complex systems and have a small fraction of technical downtime. If that happens to you at night, then relax, because there is nothing you can do (unless it is your visiting instrument). In almost all cases staff has seen that particular failure before and knows how to fix it. It may take some time. Meanwhile, re-evaluate your observing strategy so that you can resume observations, possibly of a new target, when the telescope becomes available again.
- Over-calibrate your data, don't under-calibrate it. Take plenty of biases and flats and darks, you never know what they might be good for until you start reducing the data. If you observed extended targets and there is some dead time until the next object is high enough, observe some blank fields or standards.

2.18.7 DO NOT DO THIS

Any of this will most likely make you end up on some black list. Not that I know of any, but memories are long. The following is no nonsense I'm making up but some of it actually happens quite regularly, and it is by far not a complete list.

- Do not call the support astronomer four hours before (past) sunset asking (complaining) when (why) your service data will be (hasn't been) taken, if you were supposed to actually show up yourself and observe yourself, instead of sitting 5000 kms away at home in your comfy warm office/beach/pool. If your support astronomer is really kind he'll do the observations for you, but most likely your night will be used for something else.

- Don't tell observatory staff at dinner time that instead of what you told them earlier in the afternoon you actually do need a lengthy reconfiguration of the instrument. Chances are that staff qualified for doing instrument modifications has left the mountain already. Read the *planning* section above.
- Don't freak out while taking sky flatfields in 4 four different polarisation angles at the same time. The twilight sky is highly polarised, and the fact that the polarisation channels have very different counts (or none at all) does NOT mean that the instrument is broken! Actually, it shows that it is working perfectly fine!
- If your boss sent you to do some observations for a long-term project or survey, and almost all scripts are automated, don't just lean back with a beer and watch movies. Do check the data the telescope returns. And do not write a lengthy fault report the next day complaining that the image quality is crap, because what happened in reality was that you never bothered to focus the telescope. Actually, your images were so badly defocussed that the black shadow of the secondary showed up, and you didn't recognise it all night long. That makes other people who could have used that dark time with 0.5 arcsec seeing angry. I mean really, really angry. And your boss too, by the way. Staff works hard to make everything work perfectly right for you. Don't flush their efforts down the drain pipe. A night at a medium-class telescope is easily worth EUR 10,000, and with 8m class telescopes we speak EUR 100,000.
- Do not show up two hours after sunset and then complain why no flat fields were taken for you if this was actually your responsibility. The support astronomer might have taken some for you, but he or she does not necessarily know all the filters you need or the illumination level you want.
- Do not show up at the telescope at night without a proper plan what to observe first. It's generally a very bad idea to have the support astronomer figure out from your proposal which targets are visible at the moment and should be observed first/next/etc. It's your time that is ticking away, and people might have other important things to do, for example health-checking that other instrument you are going to use later at night.
- Don't complain about bad reflections in your data when your target is just 10 degrees away from the full moon. Certainly not all time allocation committees will give you the nights you want, but please try and make the best out of it. Maybe that target can be replaced by another one?
- Don't attempt to observe a target at -50 degrees declination at an observatory at +30 degrees latitude. The Earth is spherical and not everything is visible from everywhere.

Please bear with me, but I just had to write this down. In this sense... happy observing everybody! :-)

2.19 THELI GUI trouble-shooter and frequently asked questions

This is the trouble-shooter for the THELI graphical user interface.

2.19.1 Before everything else

1. **Make sure you have the latest version.** The THELI GUI develops fast and is used by many people. If you find a bug, it is likely that someone else found it before you and that it has been fixed already.
2. **Think.** A data reduction package based entirely on scripts forces you to think through everything before you embark on launching the scripts. You have to come up with the right syntax, make sure the data are in the right place, and so on. The purpose of a graphical user interface is to remove this technical overhead, greatly facilitating data reduction. However, there is a downside. Together with the removal of the technical overhead some of the user's higher brain functions are shut down as well. Clicking a button merely involves a few neurons to move the index finger, whereas coming up with the right script syntax requires some serious thinking. Therefore, with GUI reduction, always try to stay in the picture and ask yourself this question:
 - *What am I currently doing to my data?*

3. **THELI has a steep and long learning curve.** 95% of all problems are encountered by beginners and are simply due to wrong handling or misconceptions. THELI makes a lot of effort to anticipate and eliminate such situations, yet they do happen. When they occur:
 - Ask yourself this question: *What am I doing right now?*
 - Try to understand what is happening to your data at the particular reduction step.
 - Think whether your parameter choice makes sense.
 - Look it up in the *documentation*, everything is explained there.
 - Check the trouble-shooter below.
4. **The problem is still around:** Please ask your more experienced colleague in the office next door.
5. **THELI still doesn't do what it's supposed to do:** Contact me (mischa[at]gemini.edu) and do not forget to send along the LOG file. Without the LOG file I can't help you. If you reduced data from a multi-chip camera using more than one CPU, please re-run the task with a single CPU before sending the LOG file, as otherwise it is scrambled by the multiple simultaneous outputs. The LOG files can be found *here*.

2.19.2 Use the integrated help layer

THELI comes with a substantial on-the-fly help layer. It assists you through

- **Tool-tips:** Little markers with short explanations that appear when you hover over any GUI element with the mouse cursor.
- **What's this?:** In the GUI, click on *Help*, select *What's this* and then click on a specific reduction task to obtain more extensive help. If that doesn't help, read the online documentation (i.e. these webpages).
- **Internal consistency checks:** THELI does not allow you to e.g. enter characters into numeric fields. It will also check that the data expected or necessary for a certain reduction task is actually present before a task is executed. THELI will attempt to solve a problem invisibly to you in the background. Should this be impossible or ambiguous, then you will be notified by visible warning or STOP messages. These are either self-explanatory or tell you how to proceed or fix a problem. Don't click them away. They are meant for you to read them, not to scare you.
- **Error scanner:** THELI scans the output of all reduction scripts for characteristic error messages. The line number in the LOG file in which the error occurred is then displayed, together with LOG file itself. With some bash scripting skills you should be able to identify quite some of the problems. Note that the error often occurs before the actual error THELI stumbles upon, so scroll upwards and see where it starts breaking down. If you contact me because of an error message or unexpected behaviour, please ALWAYS include the LOG file. They can be found *here*.
- **Disk space warning:** You can set two alarms that will notify you when the available disk space in your home partition and/or the main data partition drops below a certain threshold.
- **Ancillary files:** THELI produces quite a number of secondary files that help you trouble-shooting. For example, you can overlay the objects detected by THELI over the corresponding images, using *skycat* or *ds9*. Likewise, astrometry produces a number of check-plots, and so on.

2.19.3 General questions about THELI

How does THELI compare to other data reduction tools? We are confident that it provides you with close-to-perfect results. THELI has been extensively tested and compared with other large-scale data reduction packages, and uses some of the most advanced third party algorithms (in particular various *Terapix* tools). In terms of ease-of-use, flexibility and completeness our users say that it is one of the best tools available. Give it a try, you'll probably like it very much :-)

Can I make use of parallelisation? Just *select the number of CPUs* in the *INITIALISE* section.

Can I make mosaics with THELI? Yes, THELI will automatically assemble your images as they are located on sky.

My instrument is not pre-configured in THELI. Can I still reduce my data with it? If it is an imaging camera, yes! THELI will *assist you creating the corresponding configuration files and splitting scripts*.

How about difference imaging with narrow-band filters? THELI's responsibility ends with the production of coadded images. There is no module for sub-sequent difference imaging, or any other post-reduction analysis tool. However, THELI allows you to create precisely registered images in different filters which can be subtracted right away from an astrometrical point of view (see the *processing of multi-colour data sets*). There is no problem with processing narrow-band images in THELI.

2.19.4 Trouble-shooting

The following sections deal with the most frequently encountered errors NOT caught by THELI in advance, and describes some other problems. All command line tools mentioned require that you have the *THELI binary directory* in your *PATH* environment variable.

Preparation

The splitting script runs really fast, does not produce any output, or creates garbage. Most likely this is because you work with self-defined instruments, and the *splitting script* has not been set up correctly. When THELI creates the splitting script, it has to make some assumptions about the format of the input raw data. Likely, your raw data deviates significantly from what is usually the case. This can be either certain FITS header keywords or the way the FITS image is structured. Check the output of the LOG file for errors, and manually edit the splitting script such that it suits your data. The splitting script can be found in

```
~/theli/scripts/process_split_<YOURINSTRUMENT>.sh
```

If the problem occurs with a pre-defined splitting script, then please contact me and provide a sample raw data file. Likely, some format has changed or an invalid keyword is present.

error: cannot open/read fits: See the paragraph above. Maybe your image is compressed? Some systems store compressed images in files having a *.fits* suffix, from which alone the compression is not visible.

mmap: Invalid argument: You are (most likely) running THELI in a Virtual Machine, and keep the data in an exchange folder that can be read by the guest Linux OS as well as the host OS. This leads to problems with file access and permissions, depending on the settings for this shared folder. You must keep all images on the virtual Linux partition. This error occurs with Windows, I do not know if similar issues exist with MacOS.

Calibration

Many of the problems described here are caused by wrongly defined instrument configurations, i.e. they will appear mostly with self-defined instruments, and usually only once (i.e. until you fixed your camera config).

Cannot extract such zone: This error message appears when the geometry of your images is smaller than what has been specified in the *instrument configuration*. THELI then tries to extract pixels that do not exist. Likely, you got confused when entering chip geometries manually. If you work with self-defined instruments, you should always use the *Guess ...* function to determine the chip geometries based on an example image. Another explanation could be that you work with binned data, whereas the instrument configuration was made for unbinned data. In this case you would need to create a separate instrument configuration for the binned data. The corresponding splitting script can be created by simply copying the script for the unbinned data, as the data structure itself does not change.

mmap: Invalid argument: You are (most likely) running THELI in a Virtual Machine, and keep the data in an exchange folder that can be read by the guest Linux OS as well as the host OS. This leads to problems with file access

and permissions, depending on the settings for this shared folder. You must keep all images on the virtual Linux partition. This error occurs with Windows, I do not know if similar issues exist with MacOS.

Incompatible plane sizes in list: This error message appears when images of the same chip but with different sizes are present. These cannot be combined into a master bias/dark/flat. Either some binned data is present, or a corrupted file, or a calibration image that was taken e.g. during day-time by observatory staff in a different configuration. Check the file sizes to identify the offending exposure, or execute the following command in the corresponding directory:

```
dfits *.fits | fitsort NAXIS1 NAXIS2
```

Bias (dark, flat) and cube have incompatible sizes: Occurs when calibration image and the images to which it is applied (after overscan correction and trimming) have different sizes. This does not happen if you reduce the data with the same instrument configuration file, but for example if you changed some setting in the instrument configuration after creating the master bias/dark/flat and before applying them to the images. It can also happen if you import some calibration image reduced with a different software package and want to apply it in THELI; that is something you should NEVER do.

Flat-fielding does not at all work as expected: If your data does not appear as it normally does after flat-fielding, or shows significant over-correction (under-correction) of vignetting, and if it is not the usual *background* residuals, then check the following possible causes:

- One or more of the flat field exposures is significantly over- or underexposed, or saturated. Identify such exposures using basic image statistics with the the command

```
imstats *.fits
```

Use the *min-max* parameters right next to the *Process flats* task to reject bad exposures.

- You did not subtract a bias/overscan from the flat-field, or the overscan region is badly defined. In this way a large pedestal can be left in the flat, distorting the intensity ratios.
- You defined an overscan in your camera configuration even though the data has no overscan, or you subtracted a bias even though your data has no bias (e.g. because it has been subtracted internally as is the case with near-infrared data). Again, this distorts the intensity ratios in the flat. If you defined an overscan region but your camera has none, then *enter zeros in the corresponding fields* of the instrument configuration.
- The flat fields were taken in a different filter.
- Significant scattered light was present when the flats were taken. Try a different set of flats or obtain some archival flat fields.

Dark subtraction does not work as expected: Your camera has significant dark current, e.g. visible as individual hot pixels, and they are still present after you run through the *Calibrate data* task. There are three likely reasons for this:

- You must explicitly mark the **Use DARK** switch next to the *Calibrate data* task. Otherwise a bias is subtracted.
- The detector temperatures for your dark frames and the target exposures were different (*different exposure times are compensated for*). Solution: Use different darks. You can also try to manually rescale them by a few percent. For example, to increase the dark by a factor of 1.03, you would do

```
cd /mainpath/DARK
mv DARK_1.fits DARK_1_backup.fits
ic '%1 1.03 *' DARK_1_backup.fits > DARK_1.fits
```

and then re-run the calibration of the data: Click on the blue *Redo* arrow next to the *Calibrate data* task which restores the previous data status, and then repeat the calibration.

- If you are an amateur astronomer and took data with a DSLR camera, there is a chance that your camera electronics botched it up in different possible ways. In some cases results can be better without dark subtraction (use the *Do not apply BIAS/DARK* option at the top of the *Calibration* section).

Background

Defringing does not work: Fringes are caused by highly variable line emission in the upper atmosphere (*airglow*). The longer the wavelength, the worse it gets. If defringing does not work as expected, then the most likely cause is that the fringing pattern itself changed too much over the time window from which you created the correction image. Consider using a *dynamic instead of a static background model*, or choose a smaller *window size* for the dynamic model. Switching off the *fringe rescaling* can help if the background is not dominated by airglow but by scattered light (twilight, zodiacal light, moon).

A key to successful defringing are short exposure times and frequent dithering. In other words, the wrong observing strategy can render your data useless.

Weighting

Large parts of the coadded image are masked out: You were probably too zealous when applying thresholds to the global weight. Re-run the global weights with less restrictive settings (*FLAT_norm min/max*) and check *WEIGHTS/globalweight_i.fits* afterwards. Only those regions which you want to be masked should be masked in these images.

How can I get rid of those satellite tracks? Satellites can be masked *like this*. **Stars in the coadded image have masked cores:** If you are working with significantly undersampled data, stars will be mistaken for cosmics by the automatic cosmic detection. In the configuration of the *Create WEIGHTS* task, leave the *Cosmics threshold* field empty. To get rid of cosmics, you can use the *outlier rejection* in the coaddition.

Astrometry / Photometry

This is where your brains are required. Making use of it greatly helps.

Any problem with astrometry: *PLEASE READ THIS FIRST*.

Understanding Scamp check-plots: *PLEASE READ THIS FIRST*.

Sure you read these two links? Then continue.

In almost all cases problems with astrometry boil down to one common issue: insufficient overlap between object catalog and reference catalog. The number of reasons for this is almost infinitely large, but essentially it always comes down to this point. With increasing experience you'll become more proficient. **Solving an astrometry problem requires you to look into four places.** If you do not know where to start, then I strongly suggest that you check all of them:

1. Reference catalogue:

- **Catalog type:** Does the *reference catalog* chosen match your data, mostly in terms of depth, but sometimes also in terms of wavelength coverage?
- **Radius:** Too small (large) for your field of view?
- **Mag limit:** No sense in having this deeper than the depth of your individual exposures. It shouldn't be too shallow either. To check, overlay the reference catalog over your images in *skycat* or in *ds9*. These are stored under:

```
~/theli/scripts/theli_mystd.skycat
~/theli/scripts/theli_mystd.reg
```

Note that if the FITS header coordinates are significantly off, it can be difficult to visually match the catalog entries with the sources in the image.

2. Object catalogs (create source cat):

Overlay the object catalog over your image in *skycat* or *ds9* and see what you have detected! Only the brightest sources? Lots of spurious faint stuff?? Hot pixels??? Overlay the reference catalog at the same time (see above)! Do they have a significant amount of objects in common?

In *skycat* there is a default limit of 1000 objects displayed. If your catalog contains more, you must increase this limit. Object catalogs in suitable format can be found here:

```
SCIENCE/cat/skycat
SCIENCE/cat/ds9cat
```

In *skycat* use *Data servers* -> *Local catalogs* -> *Load from file* In *ds9* use *Region* -> *load*.

- **Empty fields:** In sparsely populated fields, in particular in the near-infrared, you may want to push the detection thresholds down to $DT=1.5$ and $DMIN=5$.
 - **Dense fields:** It is sufficient to get the brighter stars only, $DT=10...100$, $DMIN=5..20$. A few hundred sources per chip are sufficient!
 - **Maximum FLAG:** There are situations in which most reference sources you have appear saturated in the images. In this case, increase *FLAG* from 0 to 8. If your image quality is dodgy, and the number density is low, and other stuff happens, crank it up to 60 (seriously, I hope your data are better than that).
 - **Minimum FWHM:** Lots of hot pixels, but well sampled data? Set this to anything between 1.5 and 2.5, with strongly oversampled data even higher. This will remove the spurious detections (THELI should get rid of most of them automatically). Is your data undersampled? Make it 1.0, maybe even 0.8 with strongly undersampled images. Less than about 0.8 makes no sense.
3. **Scamp configuration:** Even if the overlap between reference catalog and object catalogs is sufficient, astrometry can still fail. The reason is that the solution space is in general very large and needs to be limited. This is controlled by the various *Scamp parameters*. If the true solution lies outside the search range specified, e.g. the offset between nominal and sky coordinates is 3 arcminutes but you set $POSITION_MAXERR=2$, you can vary all other parameters as much as you want, the correct solution will never be found. Same holds for the pixel scale or the sky position angle.

Users with cameras where the raw FITS data comes without WCS information should carefully check if the coordinates entered manually are correct and close to the reference pixel (usually, the middle of the image). $POSITION_MAXERR$ should be chosen accordingly. $POSANGLE_MAXERR$ can be set to 180, and *match flipped images* can be activated, too.

4. **FITS headers:** Don't trust them. If the astrometry just doesn't want to work out, maybe the position angle in the CD matrix is wrong? Or the RA/DEC coordinates are off by several (or even many) arcminutes? This even happens with modern telescopes: the telescope operators accidentally loaded the wrong pointing model, or forgot to update the TCS with the correct instrument after an instrument change (and the new camera has its main axis looking in a 90 degree different direction), or else. It has happened before, it just happened (which is why I am writing this), and it will happen again. Just don't think you are immune. To check, overplot the reference catalogue over the image. You should be able to find the pattern quickly. See also *item 2 above*.

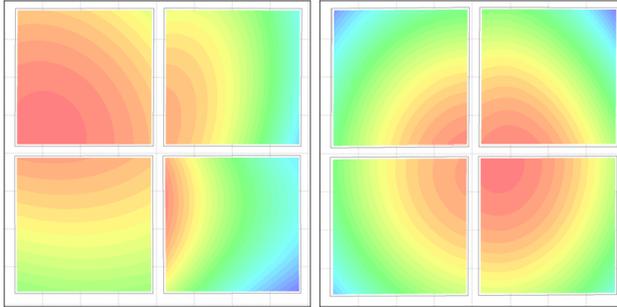
Inaccuracies likely to occur: Your astrometric solution is wrong. See above, and try again.

Not enough matched detections: Either the depth of your reference (or object) catalog is too low, in which case you should go deeper with the reference catalogue (or detect more objects in the images). Or the overlap between the reference catalogue and the object catalogue is insufficient for other reasons, e.g. manually entered coordinates are wrong, etc...

No match with reference catalog: Self-explanatory. See above for possible solutions.

It just doesn't want to work: *Scamp* check-plots look garbage, no matter what you did? Again: did you overlay the object and the reference catalog over the images? Do they agree reasonably well? What objects are detected?

A really dirty hack for sparse fields with ill-defined distortions: One problem with near-infrared observations is the very low number density of sources in the images, and often in the reference catalog as well. If that happens it can be really hard to get the distortion right. In such cases you can attempt a self-calibration: Try to achieve a half-way reasonable distortion model in first place, possibly one with *distort=2* or even *distort=1*. Create a coadded image, and *use it to extract the reference catalogue*. Then repeat the astrometry. Chances are good that the distortion model found will be significantly better, because the reference catalogue matches the image almost perfectly well. Proof of concept:



The left panel shows the HAWKI@VLT distortion measured from a sparse field. This was the best calibration possible, using the USNO-B1. The data was stacked and from the coadded image a new reference catalog was created, which was fed into *Scamp*. Using the same parameters, the much improved distortion model shown in the right panel emerged.

I have a multi-colour data set. What is the best way to do astrometry on it? See the [section about multi-colour data processing](#).

Coaddition

If there is a problem with coaddition, it's mostly the astrometry: The coaddition only does what is written in the astrometric FITS headers created by *Scamp*. If the solution in the headers is wrong, the coaddition will be wrong. A wrong astrometric solution is not necessarily recognised by the astrometric solver, and sometimes only recognised by eyeballing the coadded image. If you have a problem with coaddition, go back to the astrometry. This holds in particular for the following examples:

Significant inaccuracy likely to occur in projection: This is an astrometry problem and means that something is wrong with the astrometric distortion calculated by *Scamp*. It can well be the case that you don't notice anything in the coadded image, or that only one corner is affected and that you don't care. However, try to go back to the astrometry and get an improved solution with different settings. Try different reference catalogues, detection thresholds, depths, lower distortion degrees, and so on. It is sufficient to run the first coaddition task, *prepare_coadd_swarp.sh*, to check if the error occurs again, no need to go through the full coaddition. Simply erase the other commands from the command window. As for the reason why this error is detected by *SWarp* during the coaddition but not by *Scamp*, let me use Emmanuel Bertin's wording:

The reason is that *SWarp* needs to invert the astrometric solution, which otherwise is given in the FITS header as $(\alpha, \delta) = f(X, Y)$. The inversion is required because, contrary to *Scamp*, *SWarp* works predominantly on distorted astrometric solutions using a reverse mapping approach. In the presence of distortion terms, the inversion process involves fitting a polynomial on a 2D-grid, and checking that the direct mapping followed by the obtained reverse mapping gives back the same coordinates within a small tolerance all over the grid. The message *Significant inaccuracy likely to occur in projection* occurs whenever this condition is not met, which generally means that the solution is not invertible, and therefore is unlikely to be correct.

The coadded image looks distorted, warped, strange, badly registered etc: There is nothing wrong with the coaddition process. What is wrong is the astrometric solution. Fix it, and try again.

The coadded image wasn't created at all, is tiny, or huge, or empty (all pixels null): There are several explanations: Either you specified reference coordinates far outside the actual field of view, or the astrometry is wrong, or your

weights are all zero (wrong thresholding of flats during creation of the global weight).

Part of the coadded image is masked: You were probably too zealous when applying thresholds to the global weight. Re-run the global weights with less restrictive settings (*FLAT_norm min/max*) and check *WEIGHTS/globalweight_i.fits* afterwards. Only those regions which you want to be masked should be masked in these images.

The coadded image is too small and does not contain all pointings: This usually happens with very wide fields of view and all-sky projections. *Provide a manually chosen size for the coadded image.*

Null or negative global weighting factor: Either the globalweight is indeed null (check!), or the reference coordinates given for the coadded image are far off the actual target coordinates.

Resampling takes forever, one of the images is much larger than the others: Again, there is no problem with the coaddition. What's wrong is the astrometry.

The coadded image appears strongly (de-)magnified: Did you accidentally switch to a different instrument? In rare occasions it can happen that THELI cannot retrieve all information from a LOG file correctly, and selects the wrong instrument (with a totally different pixel scale) upon launch.

Stars in the coadded image have masked cores: See [here](#).

Wrong number of elements in key P_xxx: You deleted catalogs of some images and then re-run the coaddition. Re-run the catalog creation step to fix that.

Truncation in the coadded image: Two solutions:

- You probably specified reference coordinates which are significantly off the centre of the field or even outside the area on sky that is covered by your images. Give coordinates closer to the centre.
- If you are talking about extremely wide field mosaics and all-sky projections, it can be that you have to edit the header of the output image manually. To this end, run the *prepare_coadd_swarp.sh* step only, i.e. erase all other commands from the command window. Then, edit the NAXIS1 and NAXIS2 keywords in the following FITS header according to your needs:

```
~/the.li/scripts/coadd.head
```

Make sure that the “=” sign stays in the 9th column, the “/” sign in column 31, and that each line has precisely 80 characters in length (filled up with blanks if necessary). Otherwise the FITS header will be corrupted. Increasing NAXIS1 and NAXIS2 will provide more space to the right and to the top of the image. If you need more space to the left and at the bottom, increase CRPIX1 and CRPIX2 by the necessary amount (and possibly NAXIS1 and NAXIS2). Once done, click on the *Coaddition* task again. This time, remove the *prepare_coadd_swarp.sh* task and keep all the others.

Sorry, this is cumbersome, but predicting the necessary size with such extremely wide field images and all-sky projections is not trivial and does not always work well using an automated approach.

Cannot open file for reading / writing (setting ulimits): This is a classic with *SWarp*, and will happen if more files than allowed are opened at the same time for coaddition. Usually, the maximum number of files to be opened at any given time by a user is set to 1024, as you can verify with:

```
ulimit -a
```

To change this permanently (at least under Ubuntu) to a higher value, you have to do the following:

In */etc/security/limits.conf*, add/set the following two lines:

```
* soft nofile 100000
* hard nofile 100000
```

In */etc/pam.d/su* and */etc/pam.d/common-session*, add/uncomment the following line:

```
session required pam_limits.so
```

Afterwards, logout and login again, and verify with the *ulimit* command that your *open files* parameter is set to 100000.