

# What's new in TSG8, 28 January 2020

## TSG build 8.0.6.1

This round of updates is mostly 'backroom' work. None of the new arrivals really warrant a feature spot but I will attempt a couple. The first of these goes by the name 'Plot format tweaks' and may be of interest to TSG users who would like a little more control over plot formatting, e.g., for publications. The second is a collection of obscure developments that I'll label 'Production workflow enhancements'. These were motivated by multiple-drillhole HyLogger work but may be of broader interest. After that the document will finish off with other additions and fixes.

The usual log of changes is shown in

<https://research.csiro.au/static/tsg/tsglog1620.docx>.

## Plot format tweaks

Some TSG users have expressed frustration over TSG's limited plot-formatting options, especially when preparing a plot for publication. A few global options are offered in **File -> Settings** [Gfx1] and [Gfx2] but that's about all.

A new dialog is available for most of TSG's main screens<sup>1</sup>. It is brought up for the current screen (or sub-screen) via the main **View -> Plot format tweaks** menu. In some cases it relates to the whole TSG screen and in others (e.g., scatterplots) to the current sub-screen.

The dialog offers some control over four parts of the plot: the main title, X axis, Y axis and colourbar. It also offers sizing control for graphics saved to raster *files*<sup>2</sup>. The dialog has the expected buttons for saving or cancelling, but also has a **Try** button. It's a modal dialog, meaning that it will block you from interacting with TSG while it is up, but you can click **Try** to refresh the current plot and see how your changes work out.

<sup>1</sup> Floater support has not been done yet.

<sup>2</sup> Rasters saved to the *clipboard* are unchanged – they are not affected by this dialog's sizing controls.

## Checkboxes along the top

Each section is **enabled or disabled** via a checkbox at the top of the dialog. All are off by default, and for some screens there are checkboxes that are greyed out and blocked<sup>3</sup>. You must first turn on a checkbox (presuming it isn't greyed out) before you can adjust that part of the plot.

Active controls generally override relevant plot formatting settings in **File -> Settings [Gfx1]** and **[Gfx2]**.

## Titles

Each of the four plot sections has controls for adjusting a title. There's a **Title** field for entering text and a list for selecting the **Font size**. In some cases<sup>4</sup> the **Title** field may be greyed out, blocking customisation.

If you leave the **Title** field blank then you get TSG's automatically-generated title text.

You can adjust the text size<sup>5</sup> via the **Font size** list. The list's first choice is called '**Auto**' and gives you TSG's default. The other choices are font height in points.

## X and Y axes

There's an axis control set for each of X and Y. A set includes controls for the title and tick-marks. Some controls may be greyed out for some screens.

Titles have been discussed above, but it's worth noting that the axis title is the text like '**Wavelength (nm)**' or '**Depth (m)**'; not to be confused with the tick-mark labels. Until now, TSG has largely<sup>6</sup> made up its own mind about how many major tick-marks to draw and where to start them. This can result in unsightly tick intervals that aren't ideal for publication. Also, TSG has offered no control over the font size for the tick-mark labels.

If the plot allows it<sup>7</sup>, there is a **Ticks** list with four options:

1. **Auto**  
Leave it to TSG as in the past.
2. **Number of majors**  
You get a **list** for selecting the number of major tick-marks.
3. **Interval**  
You get a **field** where you can type in the interval between major tick-marks – in plot data units. Major tick marks are at whole multiples of the interval. The first tick-mark usually isn't at the origin.

---

<sup>3</sup> For example, the Log screen's dialog only has the Main title, X axis and Raster size checkboxes enabled.

<sup>4</sup> For example, customisation of title text for the Log screen is not allowed because the dialog deals with the whole screen and there's almost always more than one log column. Customisation of the title font size is allowed however. In case you're wondering *where* the log-screen titles are, for saved graphics and prints you get column titles to take the place of the 'header controls' that you see on the screen.

<sup>5</sup> The font family is still hard-coded in TSG.

<sup>6</sup> In File -> Settings [Gfx2] there are very basic, *global* 'axis rounding' controls that affect major tick intervals.

<sup>7</sup> Some screens / plots do not allow tick control, e.g. the notorious Log screen. Also, if a **class scalar** drives the axis then it does not make sense to offer tick control.

#### 4. Interval and 'nice' origin

As above, except TSG places the first tick at the origin. This will generally change the plot (automatically). The first tick is at the left (x) or bottom (y) edge of the plot, so the plot's range must generally be adjusted.

The **Font size** list allows you to control the size of the tick labels.

### Example

Say you have a drillhole with depth ranging [101.8, 252.4] m and you're plotting a Summary-screen spatial histogram with depth on X. TSG gets creative with the X tick-marks. The interval between major ticks jumps around depending on the size of your TSG window. It could be 20, 30 or 40m, or even something else.

Now you can nail it down. Select the sub-screen and bring up the **Plot format tweaks** dialog. Enable the **X axis** checkbox at the top, then look for the **Ticks** list in the X axis control set lower down. Set it to **Interval** and type what you'd like (say 50) in the field that appears alongside the list. Click the **Try** button to see the effect. Not bad, but the first major tick-mark (150m) is quite far from the origin. Maybe you'd like a major at the origin. If so, set the list to **Interval and 'nice' origin** and click **Try** again. You should find that the plot has been adjusted automatically to have X start at 100 instead of 101.8, allowing an origin tick to be drawn at 100m.

### Colourbar

This control set has **Title** controls and a **Width** list. TSG has still not surrendered full control of the colourbar and making adjustments can be a bit of an adventure. Most bets are off if colour is driven by a class scalar (TSA Group for example), as the number of blocks is then determined by the number of classes. Numeric colourbars are more amenable to adjustment.

The **Font size** selection controls the height of the little title at the top *and* each colourbar block. The **Width** selection is in characters and controls the colourbar width. TSG controls the overall colourbar height, meaning that you don't have *direct* control over the number of colour blocks.

### Raster size

This set gives you control over the physical size of a raster (JPEG, Bitmap or PNG) saved for the plot, e.g., using the **View -> Save graphics to file** menu. It does not affect rasters saved to the clipboard.

You get to specify the raster's width (**X**) and height (**Y**) in millimetres, and its resolution (**DPI**) in

Raster size (saved graphics)

Size (mm), X:  Y:   Lock aspect    DPI:

dots per inch. Normally the generated raster has the same aspect ratio (width/height) as the screen plot, with the Y field being locked from editing and getting updated automatically when you change X. If you wish you can turn off the **Lock aspect** checkbox and change Y separately. Note that TSG uses a separate rendering pass to generate saved graphics, so if you change the aspect ratio then the raster will be rendered at the given aspect ratio. It isn't a stretched copy of the screen plot.

## Support in TSG screens

### Summary screen

The dialog is brought up separately for each sub-screen. All options are available.

### Log screen

The dialog deals with the whole screen – all log columns. The only options available are main title font size, X-tick-label size, and raster sizing.

### Spectrum screen

All options are available.

### Stack screen

Most options are available. Blocked options: Y-axis tick control and (for combo stacks) X-axis tick-interval control.

### Scatter screen

The dialog is brought up separately for each sub-screen. All options are available.

### Tray screen

The dialog deals with the whole screen. Y-axis and colourbar controls are blocked.

### Hole screen

The dialog deals with the whole screen. Colourbar controls are blocked, and for the X-axis only the tick-label font-size control is enabled.

### PLS screen

No support yet.

### Floaters

No support yet.

## Production workflow enhancements

These developments were largely motivated by a TSG power user who was faced with the challenge of doing systematic processing and downsampling on a set of several dozen HyLogging drillholes. The sequence was repeated from time to time to assess a change.

The new enhancements might seem petty and particular, hardly worth mentioning, but they can be significant time savers under the right circumstances. Even if you don't work with HyLogging datasets, I encourage you to look through this section if you are involved in large-scale or repetitive production-oriented work. There might

be something here for you, or you might have ideas for workflow-related feature requests yourself.

## Downsampler

### Absolute expanded

The previous ‘new since’ notes went over the Downsampler’s ‘absolute expanded’ option. A brief recap: When downsampling a class scalar or ‘consolidated scalar set’ (e.g., a TSA unmixing result set) to CSV, the ‘expanded’ option delivers a tabular output with one column per class item (e.g., one column per TSA mineral). Each row in the table is normalised to sum to 100% unless the ‘absolute’ option is also used, in which case the rows are not normalised.

There has been a subtle change in the handling of weight scalars.

- If you use a **weight scalar** (page 2 of the Downsampler) then weight contributions are usually normalised out for each row (bin) of downsampled outputs. Now with the absolute+expanded option, they are not normalised. There’s no row normalisation at all. For sensible results, a suitable weight scalar like **recovery rate**<sup>8</sup> should be used.

### Template handling

Headless (batch / unsupervised) downsampling is important in a production workflow, and in TSG it involves a downsampler configuration file<sup>9</sup>. The idea is that you set the Downsampler up for one of your datasets and save a configuration file at the end. Afterwards this configuration can be applied to other datasets<sup>10</sup>.

Standardisation is important and some things can trip you up when applying a downsampling configuration to other datasets. A couple of areas have received some attention.

### Range handling for the main downsampling scalar (commonly ‘depth’)

Depth ranges can be a headache when working with a set of drillholes. It’s unlikely that each drillhole will have the same depth span, so how do you set up a Downsampling configuration file without compromising and working with a common subset? What if you would like ‘nice’ boundaries (like whole metres) for the outputs? The Downsampler has subtle ways of dealing with such issues. When you prepare a configuration file, the Downsampler knows the template dataset’s *real* basis (depth) range and compares it to the ‘**Range of interest**’ that you type in. Some assumptions are made and here’s the current state of play. It’s best described by example:

---

<sup>8</sup> A recovery-rate scalar is often used in core-logging work. Its values give the ratio [‘depth’ span in ground] / [span in tray]. A metre in the core tray doesn’t usually equate to a metre under the ground, and a recovery-rate scalar helps us produce plots and output that relate to the span in the ground rather than the tray. In a sense the recovery-rate scalar is ‘*ready to use*’. Its values are not arbitrary (or at an arbitrary scale) but have physical meaning.

<sup>9</sup> A Downsampler configuration file can also be used interactively – not just in headless mode. You can load the config file in page 1 of the Downsampler wizard. This sets up the Downsampler according to the saved settings.

<sup>10</sup> In the last page of the Downsampler wizard you will find an option to save a configuration file.

Say you're downsampling to a 5m bin and the template dataset has a real depth range of [31.7, 862.1] m. You have set up an active depth range of [35, 860].

- TSG sees that 35 is **within a bin's span** of 31.7, so when applying the template to dataset X it will take X's depth minimum instead of going with 35. It also sees that 35 is 'nice' (a whole bin multiple) so it will 'niceify' X's minimum.
- TSG sees that 860 is within a bin's span of 862.1, so when applying the template to dataset X it will take X's depth maximum instead of 860. It also sees that 860 is 'nice' so it will 'niceify' X's maximum.

Say a dataset X comes along with depth ranging [52.5, 1001.7] m. The Downsamplers should go with an output depth range of [55, 1000] m.

And a twist in the story to take us home:

In the example above, say you set the start of the range to 30 instead of 35.

- TSG sees that 30 is within a bin's span from the template's real 31.7m minimum and that it is 'nice' (a whole bin multiple). However it is outside the template's [31.7m, 862.1m] range so the rounding goes 'the other way'. For dataset X and its actual depth start of 52.5, the output start will be 50 (rather than 55).
  - The first output bin will have fewer contributors than the other bins. The first bin covers 2.5 meters while the others cover 5. Normally the Downsamplers compensate by output row normalisation but *it will not compensate at all for absolute+expanded*.

## CSV columns in scalar downsampling

In a production workflow that involves scalar creation and downsampling, it is important for each dataset to have the same scalars with consistent names – or at least the subset of scalars involved in the workflow. Scalar creation can often be addressed well by 'Copy processing', which can be run headless but won't be discussed here.

A remaining issue is that the scalar *order* in the datasets is not necessarily consistent. When datasets are downsampled to CSV, there might be variation in the CSV column order.

That's been dealt with now. The configuration file now includes information about the number of output CSV columns and their order. Output CSV files are created with these columns and each dataset is downsampled to match. You get consistent CSV files that can be merged (after dealing with CSV header rows). If a dataset has extra columns (not included in the configuration) then they are ignored, and if it is missing columns then these columns have space (no data before their comma) in the output.

This mechanism is automatic but will only work with configuration files that were created by this release of TSG (or a newer one).

## General

### Copying scalars between paired datasets

Users of paired Vis-SWIR + TIR datasets come across this issue from time to time. Currently, each dataset in the pair has its own scalars and there are some methods<sup>11</sup> in TSG that can only access the primary dataset.

A corner-cutting tool has been introduced, allowing one to copy a scalar from the associated dataset to the primary (or vice versa). The copy is an ‘import’ type scalar with the same name as the original, but it will not automatically stay in sync with the original. A *manual* sync is possible by copying it again – this will update the copy that’s already present instead of making a second copy.

### Interactive

Use the **Edit -> Copy a scalar Primary <-> Assoc** menu. You’ll get a scalar selection dialog. By default it opens for selecting a scalar from the primary dataset but you can select its special first entry to switch datasets. Select a scalar and it’ll get copied to the other dataset.

### Headless

There’s a new headless task called **COPYSLR**, and the ‘TSG headless reference’ document has been updated for it.

### Copy-processing – user TSA

A little-known fact about TSG, perhaps, is that Copy-processing can copy user TSA settings, causing a user TSA recalculation in the target dataset.

For interactive use, the master switch for this functionality is carefully hidden amongst the countless checkboxes in **File -> Settings [System]**. Turning this checkbox on adds TSA copying to the mix. You can then do just TSA copying on its own by using **File -> Copy processing -> Scalars only** and not selecting any scalars in the dialog that follows.

For non-interactive use, the **COPYPROC** headless task’s **CPROC\_OPTIONS** instruction has a **TSA** option. Well it took *quite a while* for some to tell me: ‘Alone, it is vanity.’ I had a technical reason for not hooking it up for solo use but overcame it – it works on its own now. You can now have COPYPROC just update user TSA without doing any scalar or layout copying.

## Other new functionality

---

<sup>11</sup> E.g., Downsampler and scalar construction wizard.



## TSA

### CLS screening

A CLS screening stage has been implemented. This means that the relatively fast CLS algorithm is used to thin out the list of available library classes before TSA gets started on unmixing a spectrum. For example screening might thin TIR unmixing's typical list of about 80 active library classes down to about 20. This makes TSA run much, *much* faster<sup>12</sup> at high mixture levels.

The screening pass is done at the start of a spectrum's unmixing calculation and includes 'subgroup preservation'. For example, say we come in with 3 of the 6 quartz classes active and screening leaves only 1 active. As there was a survivor in the quartz subgroup the other two are reinstated.

CLS screening is **automatically**<sup>13</sup> **enabled for mixture levels of 4 or higher**. It is mostly for TIR unmixing, where higher-level unmixing is often desired. TIR TSA at mix6 used to be impossibly slow but is now quite acceptable.

CLS screening *can* change results. A fair amount of testing has been done and currently it is felt that the changes are not detrimental in balance. In any case it has been uncommon for TSA to be run at mix4+ due to the long runtime.

### TSA internals

Those who dare to run TSA at high mixture levels might feel that it delivers a sparsity of high-mixture results. Screening tests revealed<sup>14</sup> a weakness in TSA's two-stage approach where (particularly at high mixture levels) some promising candidate results might not be retained from the first stage.

Some internal adjustments have been made to mitigate the issue. A future development may resolve it completely.

### Unmixing with external libraries

Support has been added for 'blacklisting' some of the external library's spectra so that they are not active in unmixing *by default*.

Also, there is support for the external library to blacklist some of the spectra in the system library that it is extending. This was done in support of extension libraries that contain system library *replacements* – it makes such libraries easier to use. You will need to contact me (Peter Mason) if you would like to use this functionality as it is obscure.

---

<sup>12</sup> If N is the number of active library classes and K is the current mixture level then stage 1 of TSA must evaluate 'N choose K' candidate results. Say we are working at mixture level K=6. This works out to 300500200 for N=80 and 38760 for N=20. So this reduction in N should make stage 1 run roughly seven and a half thousand times faster.

<sup>13</sup> It isn't used for lower mixture levels or if Domain Restricted Mineral Sets are in use.

<sup>14</sup> CLS screening sometimes yielded a good result that didn't make it past stage 1 of an unscreened TSA calculation.

## Licence

The licence dialog has been tidied up a bit and the tooltips improved. Mostly, there is a new button called '**Switch**'. It deactivates the current licence and prompts to activate another licence.

For example, say you have your own Base (viewer) licence and sometimes access a Premium licence that your organisation shares amongst TSG users. **Switch** will enable you to switch back and forth between the two more easily. TSG now remembers *two* sets of licence IDs and passwords, so switching is easy after you've done it once (no IDs and passwords to remember).

The screenshot shows a 'TSG Licence' dialog box with the following fields and buttons:

- Name: PETER MASON
- Company: CSIRO
- Licence: 61840549
- Base: Standard
- Premium time: None!
- Affiliate: Virtual Ed's Efficacy Emporium
- Buttons: Refresh, Deactivate, **Switch** (highlighted with a red box), Purchase time, Purchase new, Activate a licence, Close.
- Customer License Portal button.

## Multiple-screen issues

It's becoming common for people to have more than one screen connected to their computer, and some such TSG users have experienced problems with TSG's interface. On-screen lists don't appear to respond, or they 'pop up' in unexpected places (like way off to the side on a different screen for example). One of the Windows controls that TSG uses was found to be problematic and has been replaced by a custom control.

## Import

Work has done to support the upcoming HyLogger-4 and its MIR (mid infrared) spectrometer. You will see a change in the Import wizard's '**Wavelength info and resampling**' page. There's a row of '**Component**' radio buttons along the top. (TSG can now support four incoming spectral components.)

Component radio buttons are only enabled for import types that have more than one spectral component.

Users of field-spectrometer data will experience no change here.

Users of HyLogger-3 and TIR-extended HyLogger-1 data will find things organised a little differently now. The dialog shows only one set of wavelength resampling controls (at a time), not two. Click the **Assoc 1** component radio button to access the TIR controls.

## Special dataset components

This dialog is brought up via the **File -> Special -> Special dataset components** menu. If any unique scalars are assigned in the dialog then they are synced to the associated dataset (if there is one) when the dialog is closed.

The reason this little change gets its own spot in these ‘what’s new’ notes is because I would like to tell you...

## How to edit a dataset’s HoleID scalar

This is probably relevant only to users of HyLogger data.

A HyLogger import delivers a dataset that includes a HoleID scalar that is locked from editing, but sometimes it happens<sup>15</sup> that HoleID really needs to be edited. Here’s what to do. In the case of a dataset pair you just need to work with one dataset.

1. Bring up the **Special Dataset Components dialog**. Under **Unique system scalars**, turn off the **HoleID** checkbox. This will free up the Hole ID scalar for editing. Close the dialog.
2. Plot the **Hole ID** scalar in a log column and bring up the class editor for it. Make your changes.
3. Bring up the **Special Dataset Components** dialog again. Turn on the **HoleID** checkbox and then locate the Hole ID scalar in the accompanying list. Close the dialog. The Hole ID scalar will automatically be synchronised<sup>16</sup> to the associated dataset (if there is one).

## Miscellaneous

- Imported class scalars are not forced to lowercase anymore.
- The Google ‘Turbo’ colour LUT is available for colouring numeric scalars and spectra. Reminder: Spectral colouring is set via **File -> Settings [Gfx2]** and scalar colouring via **Edit -> Scalar names and groups**.
- A clipboard-copying glitch has been fixed. TSG normally copies data to the Windows clipboard in several formats. Sometimes<sup>17</sup> the metafile is too big for the clipboard to accept. This used to spoil the whole copy, but now the other formats (like JPEG) will still be copied.
- Tray screen, tray and mosaic pic generation: The output resolution cap is effectively gone.
- NVCL Database:
  - The scalar upload could fail unexpectedly, or even crash. (Fixed.)
  - For a dataset that has already been uploaded before, the linescan upload will happen again automatically if the local raster file is newer than the database’s version.
- Sorting:
  - The original sample order is now preserved over intervals where the sort scalar has the same value.
  - Domained unmixing wasn’t working properly on sorted datasets. (Fixed.)
  - Tray screen navigation and pic generation were glitched on sorted datasets. (Fixed.)

---

<sup>15</sup> E.g., An incorrect ID was entered during scanning and the imported HoleID is wrong, or a drillhole was imported in two separate parts and the merged dataset has two drillhole IDs.

<sup>16</sup> This last bit is new, and in the past there was no way to sync the Hole ID scalar.

<sup>17</sup> For large datasets, plots like a coloured Summary Overview can yield very large metafiles.

- PLS screen: There was a bug in the session-file save / load that could cause a crash. (Fixed.)
- The interactive depth-logging tool:
  - Depth calculations weren't working properly on single-section trays. (Fixed.)
  - There was a glitch in recovery-rate calculation that was evident in a chips dataset - where the pixel size in mm is relatively big. (Fixed.)
- Edits to the Rockmarks scalar were not being synced to the associated dataset.