Co-adding spectra - in progress

Story

A user wants to co-add grism spectra from several different exposures within a visit (same PA) or from different visits (different PAs).

Inputs

Outputs

Computations

Drawbacks

Users need to keep in mind the limitations of co-adding 2D and 1D grism images! In the general case, this procedure cannot be done correctly because the spacial AND the spectral coordinate cannot be simultaneously interpolated on the same grid for each spectrum. You can do just the spacial axis (see emission line maps) or the spectral (and the co-added spectrum will smooth over the spacial variations) but not both. There are corner cases when a reasonable co-addition can be done:

- Spectra at the same PA (from images within the same visit) can be coadded somewhat reasonably if the dithers are not very large. Since the dispersion is pretty slowly varying as a function of position on the detector, cutouts from dithered FLTs within the same visit can be coadded. Also, when the spectra are at the same PA, the spacial axis is the same and there will be little smoothing of the morphology of emission lines (due to the distortion correction).
    - A possible approach to co-addition of spectra at the same PA without resampling is the interlacing method used by 3D-HST (Momcheva et al., 2017). The benefit is that there is no resampling of the flux or errors of the pixels and therefore bad pixels/missing data/etc. can be properly accounted for rather than creating spurious features in the spectra. Since this is a rare corner case, I don't think this method even merits effort because it cannot be generalized:

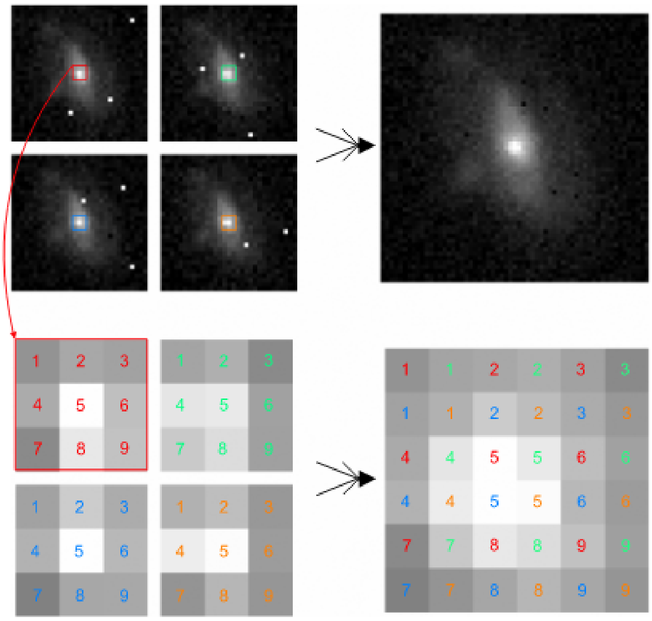

Figure 5. Illustration of interlacing with a small section of a JH_{140} direct image. The same process is also used to interlace the grism exposures. The top left shows the four individual exposures. These are combined to produce the interlaced image in the top right. In the bottom row, we show the same procedure for a 3×3 pixel part of the core of the galaxy to demonstrate how the pixels from the individual images are arranged in the final interlaced grid.

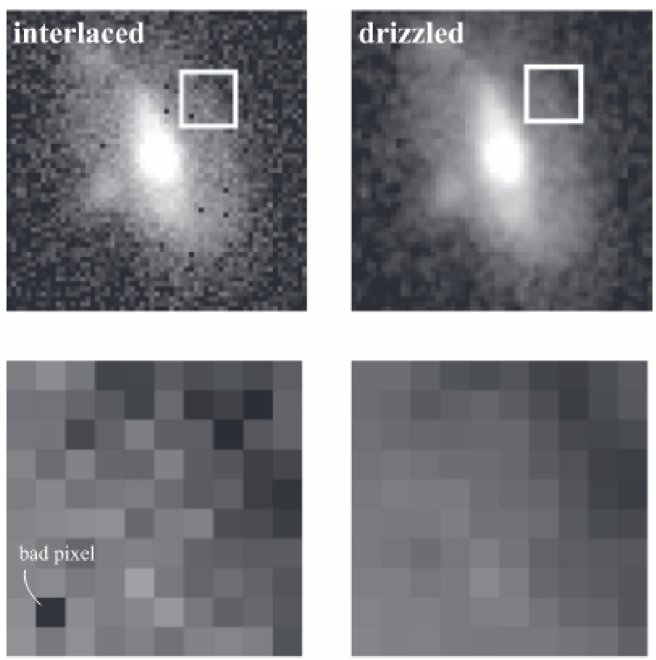

Figure 6. Comparison between interlacing and drizzling. Compared to the drizzled image, the interlaced image has higher resolution, as the pixels were not interpolated. Flagged pixels (due to cosmic rays and chip defects) are retained as single pixels in the interlaced image, whereas they are interpolated over in the drizzled image.

- Continuum spectra can be coadded across different PAs if there is no useful information in the spacial axis. It is probably better that 1D spectra are extracted form each FLT and then co-added so that the errors can be propagated properly.