Scattered Light in PFS

Robert Lupton

May 15, 2024

1 Introduction

Scattered light is a significant problem in PFS, coming both from bright targets in the field, and from the extremely bright OH Meinel series in the `r` and `n` arms. The estimated amplitudes of the scattering in the different cameras is given in Tab. 1. An example is given in Fig. 1.

We can learn more about the scattering by using data where some of the cobras are hidden behind the black spots. In practice not all the cobras are perfectly hidden (some are broken, and some don't move exactly as commanded) so we subtract a "dark roach" image with as many cobras hidden as possible from these black-spot-hidden ("roached") images. When appropriate, these dark roach images have also been scattered-light corrected.

Fig. 4 shows a horizontal slice through a "mod 32" roached image (*i.e.* one with every 32$^{\text{nd}}$ cobra illuminated) with a global background level removed; the scattering wings are obvious and the model fit (which uses the blue model profile and ignores the background) is oversubtracted. Fig. 5 shows a slice through a "mod 4"[1] image, and in this case it's clear that the scattering wings overlap strongly — with the global background removed the minima in the profiles are still well above the baseline.

We see that we have two problems to solve: the background level, and the overlapping profiles in the mod-4 data which are currently used to solve for the fibre profiles. We shall see that the first is by far the most important issue, and that the overlaps of the profiles of pixels separated by at least 4 fibre-spacings are are small once the scattering is handled properly.

[1] Or "shaka".

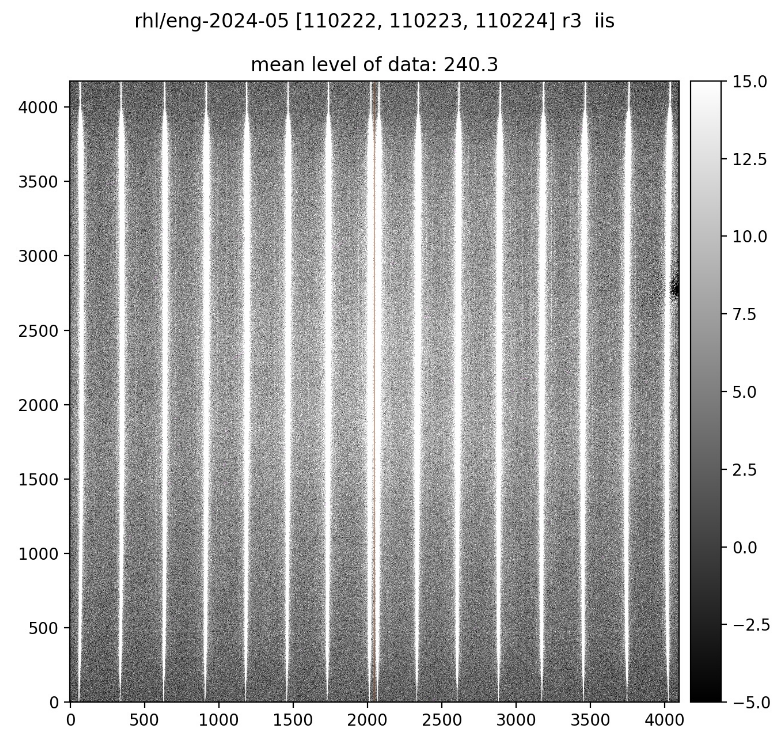

Figure 1: An IIS quartz-halogen image, the median-combined visits $110222 \cdots 110224$, stretched to show the scattered light. Note that only 16 fibres are illuminated, so the amplitude of the scattered light is only *c.* 2.5% of that in an exposure using all the fibres. The spectra have peak fluxes of *c.* 35000 ADU (see Fig. 2).

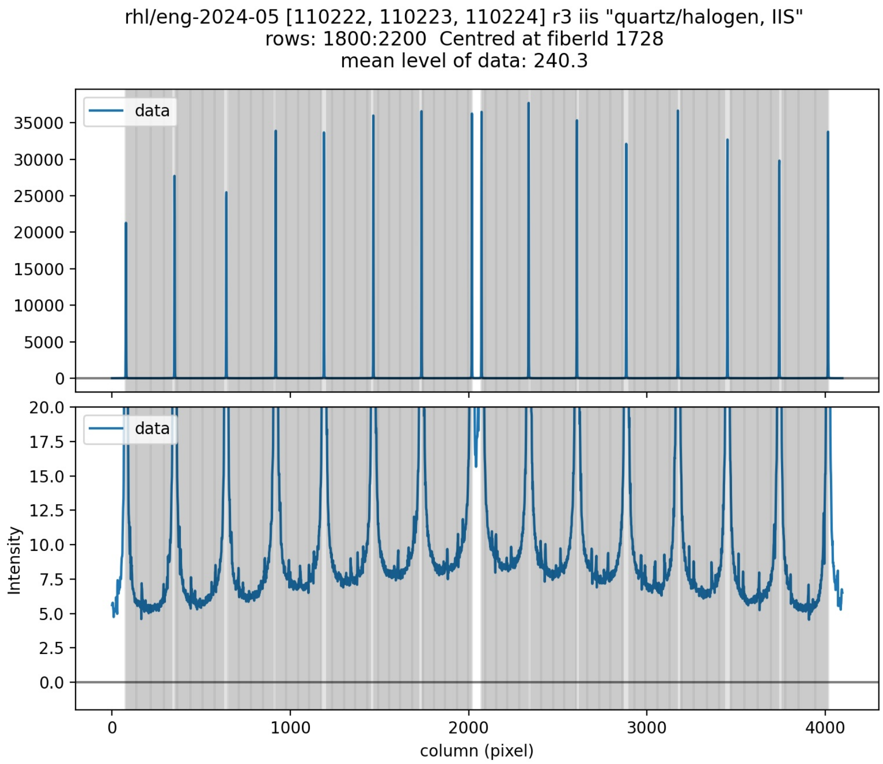

Figure 2: A horizontal cross-section (*i.e.* the spatial direction) through the centre of an IIS quartz-halogen image, *i.e.* one with only the 16 engineering fibres illuminated.

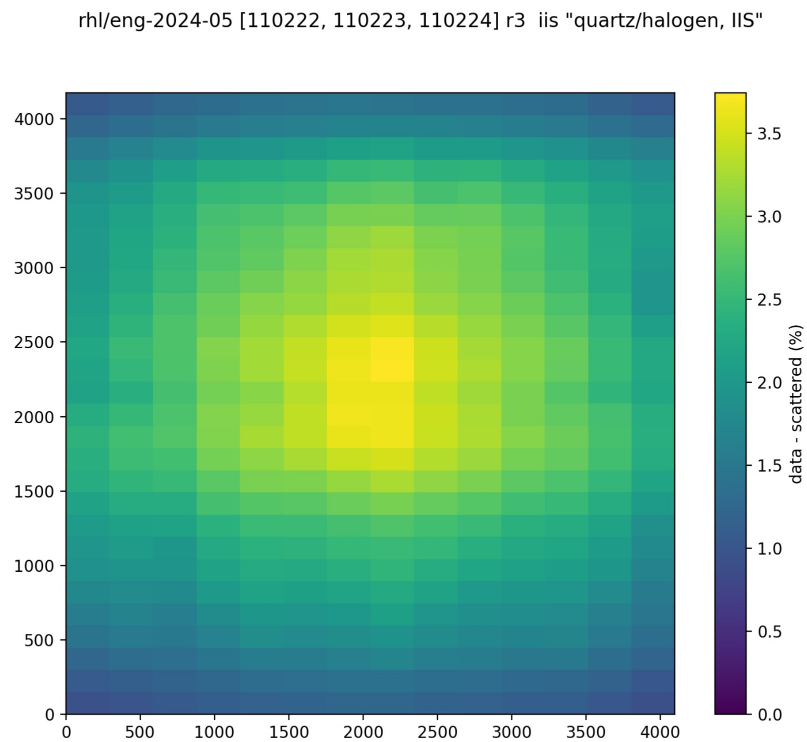

Figure 3: The scattering in an IIS image, as seen in Fig. 1, evaluated away from the illuminated pixels. The values are normalised by the *total* flux in the image. The horizontal band at row *c.* 2100 is unexplained, although it is near an inflection and dip in the direct spectrum.

Arm	Spectrograph			
	1	2	3	4
r			5.8	
m			5.8+	

Table 1: Scattering amplitude, as a percentage of the total flux incident on the detector for each camera **XXX: "each" is a bit of an exaggeration**. The r and m cameras are treated separately as they use different gratings, and while the dichroic $r - n$ effectively restricts light to the imaging area in r, light extends beyond the CCD in the m arm, rendering the effective scattering amplitude (defined in terms of the *measured* light) larger.

While data with cobras hidden behind black spots allow us to explore the effects of illuminating any desired set of cobras, it proves convenient to use the engineering ("IIS") fibres to characterize the scattering. Figs. 2 and 2 show an example, where the spatially-varying structure of the scattered light is obvious. Not only does taking IIS data not involve the time-consuming step of hiding fibres behind spots, but the analysis doesn't involve the complexities involved in modelling and removing the effects of imperfectly hidden fibres. Because the IIS fibres differ from the science fibres (they have smaller cores), there is the potential for their illuminating the optics differently and thus producing different scattering, but there is no evidence for this (**XXX: not that we've looked**). The *cores* of the lines will of course be different.

Fig. 3 shows the amplitude of the scattered light in an IIS image. The peak amplitude is less than the corresponding value in Tab. 1 as even the central pixel is mostly affected by scattering of light which falls towards the centre of the detector, so to get an scattered amplitude of *c.* 3.5% requires nearer 4.8% scattered power.

2 Handling Scattered Light

An obvious approach is to estimate a spatially-varying model of the background from the light between the fibres; unfortunately, there are at least two problems with this approach. The first is that in reality there *aren't* any pixels "between the fibres", the fibre spacing is *c.* 6.5 pixels, and the

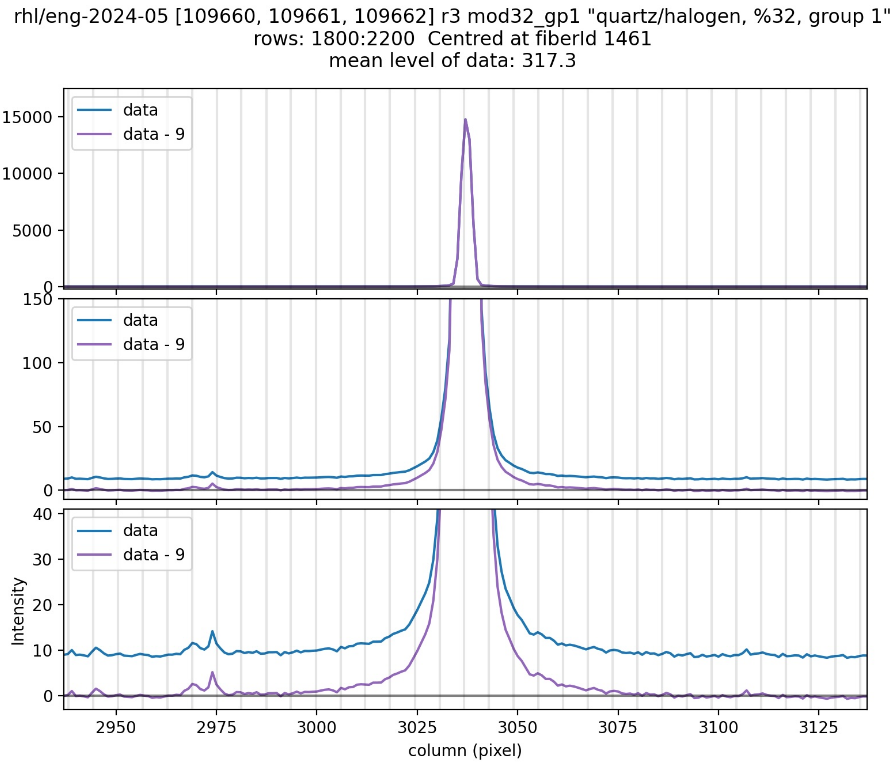

Figure 4: A horizontal cross-section through the centre of a "mod-32" quartz-halogen image (*i.e.* with every 32$^{\mathrm{nd}}$ fibre illuminated) image, with the "dark roach" image subtracted. The vertical gray lines show the positions of all the fibres, including unilluminated ones. The purple lines have had a background of 9 ADU removed.

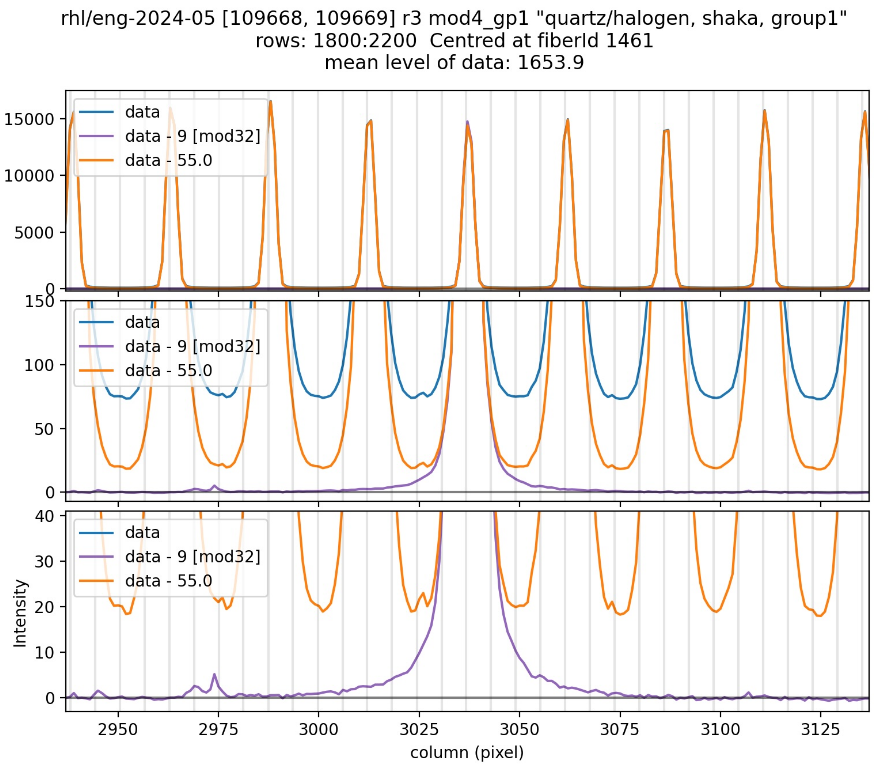

Figure 5: A horizontal cross-section through the centre of a "mod-4" quartz-halogen image (*i.e.* with every 4^{th} fibre illuminated) image, with the "dark roach" image subtracted. The vertical gray lines show the positions of all the fibres, including unilluminated ones. The orange lines had a background of 55 ADU removed; the purple lines are the same as those in Fig. 4. The orange lines are broadly inconsistent with a copy of the purple line, duplicated at the position of every illuminated fibre.

wings of the optical PSFs extend further than that **XXX: figure?**. A more subtle problem is that we cannot logically separate the cores of the PSF from the scattering wings, whereas an additive model handles the two processes quite differently. This would be unimportant if all fibres and spectra were identical, but in reality neither of these assumptions is correct.

Another way of expressing this concern is that the measured fibre profiles are the sum of the "optical" cores and the scattering. Modelling the scattering as a background means that it is independent of the brightness of the spectrum, and hence background subtraction would lead to an effective profile which was dependent on the brightness of the spectrum (which will depend on the properties of both fibre and the object being targetted). If these profiles are measured incorrectly, and given that the fibre profiles overlap, the extraction code will of necessity incorrectly assign flux in one fibre to its neighbours.

Accordingly, we shall treat the scattered light as the scattering *kernel* which it assuredly is. We can define a kernel κ which, when convolved with the image from the detector, produces an image of the scattered light. At first glance this is trivial, but it's worth thinking a little more deeply. Our mental model of the optical system should be that there is a converging beam of light focussed onto the detector, and it is *this* light which is scattered by various optical surfaces (the grating; the corrector, the Mangin mirror, ...). We can probably neglect the double scattering, but light which never hits the detector in a perfect system can still be scattered and detected. This is especially a concern in the m arm, as the grating directs light into the camera which is not imaged onto the detector.

There is no theoretical reason to separate the "scattering" and "optical" kernels, but in practice the distinction is important. We shall see that the scattering is on a *very* large scale (with a divergent (!) profile in 2-D of *c.* -1.5), so modelling the per-fibre profile directly would be inconvenient, leading to *c.* 4000^2 matrices. By separating the two scales (scattering: $O(4000)$ pixels; optical: $O(10)$ pixels) we make the problem computationally simpler, and more local. In theory we could iterate to make the two approaches formally identical (the iteration would be equivalent to an iterative matrix inversion), but in practice this hasn't proved necessary as the first-order scattering correction is no worse than *c.* 5×10^{-2} **XXX: I haven't actually checked**.

The kernel κ is probably a function of wavelength, *e.g.* it's reasonable to suppose that the effect of surface roughness is proportional to the wavelength

λ (as the roughness in units of wavelength is larger, and the diffraction angle scales as λ^{-1}), but for now we shall neglect these effects, and look for a single kernel for each arm. With this assumption, the kernel in r and m should be the same if scattering from the grating is similar, and we shall briefly examine this assumption in Sec. 4.2.

3 The Scattering Kernel

The simplest plausible assumption is that there is a single isotropic kernel κ for each arm, and that it is a power law. Experiment reveals that better results are obtained with a double power law, so we take

$$\kappa(r) = A\frac{c_A}{(r^2 + a^2)^{\alpha/2}} + B\frac{c_B}{(r^2 + b^2)^{\beta/2}} \tag{1}$$

where $c_{[AB]}$ are normalisation coefficients designed to normalize the (2-D) profiles to 1:

$$\int_0^\infty \frac{c_A}{(r^2 + a^2)^{\alpha/2}} 2\pi r \, dr = 1.$$

The second term is not theoretically necessary — it only modifies the shape of the profile which we measure for the fibres — but it's easier to interpret our models if we keep as clean as possible a separation between "scattering" and "optics". For the initial fits to the $r3$ profiles, we found $A = 0.048$, $b = 1$, $\alpha = 1.5$, $B = 0.01$, $b = 5$, and $\beta = 3$. The total scattering powers, $A + B$ (5.8% in this case), are given in Tab. 1. **XXX: Confess that RHL did this by hand for $r3$**

It's worth noting that profiles that fall this slowly to arbitrarily large radii are unphysical, as the total scattered flux diverges as $\sqrt{r}$. This means that the imputed amplitude of the scattering kernels depends on their size, or equivalently that light from distant parts of the image has a significant effect. We have taken the kernel size to be twice the size of the detector, to correctly allow for the influence of distant bright emission lines.

3.1 Modelling the Scattered Light

As mentioned in the introduction (Sec. 1) we shall start by assuming that all the light entering the camera is detected. In that case, for an image I the scattered light model S is simply

$$\mathfrak{S} = \kappa \otimes (I - S) \tag{2}$$

or, assuming that $|\kappa| \ll 1$,

$$\mathfrak{S} = \kappa \otimes I. \tag{3}$$

It is important that the image I be cleaned of signal that is not produced by light coming from the fibres (*e.g.* cosmic rays, dark current, persistence) before being used to estimate the scattered light. Our dark currents are too low to be more than a theoretical concern, but cosmic rays generate a significant signal. Another issue is saturation in the cores of emission lines lines. Providing that the charge in the bleed trails reflects the total (and if the CCDs bleed at levels that the A/D can handle; the H4RG ramps avoid the problem) saturation doesn't affect the far wings of the scattering. It does, however affect the nearby pixels once smeared by the kernel κ.

There is a problem with this approach, namely that the scattered light model $\mathfrak{S}$ is generally slightly smoother than the image I, so the corrected image $I - \mathfrak{S}$ will exhibit structure around sharp features (*e.g.* slight negative rings around bright emission lines). This is not a fundamental limitation to the approach as the real (unknown) scattering also modifies the PSF, but it makes finding a practical kernel harder**XXX: I didn't try very hard**.

An alternative is to model the image I in terms of the extracted spectra S_i, the known fibre profiles P_i, and detectorMap D: $I = \sum_i S_i D \otimes P_i$. This has the advantage of having lower noise than I itself, and potentially allows for better removal of image defects such as cosmic rays. The extracted spectra are also less sensitive to bleeding, as we can recover the correct spectrum from our knowledge of the fibre profiles. An advantage of the modelling approach is that we are not required to use the true profiles P_i in our model, but can instead choose to use narrower profiles, which helps with the deconvolution problem, at least in the spatial direction. Most of these results refer to continuum lamps, although Fig. 9 implies that this is not a major problem. Using a line model code such as `RAGNAR` would allow us to solve problem of the deconvolution in the wavelength direction by using a line list to construct the model spectrum rather than using an extracted one (we'd still use a smoothed version of the continuum). It is not clear how well this would work in practice.

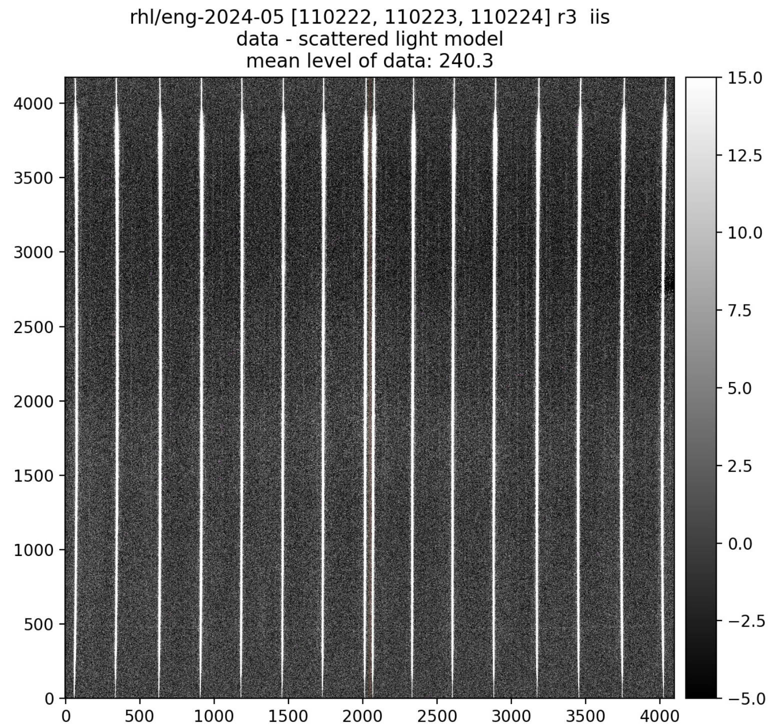

Figure 6: The same IIS image as Fig. 1, but with a the background removed using a simple isotropic spatially-invariant kernel.

The results presented in this note use the modelling approach, although it makes the processing slightly more complex, as we need the extracted spectra early in the processing and are thus forced to extract them twice.

4 Scattering-corrected Data

With this model in hand, I used the IIS data to determine amplitudes, core radii, and power-law indices by visual inspection of horizontal slices such as those shown in Fig. 8 and manual adjustment, for the SM3 r/m camera only. This could and should be automated, minimising the residuals well away from any lines. Fig. 6 shows the same IIS image as Fig. 1, but now

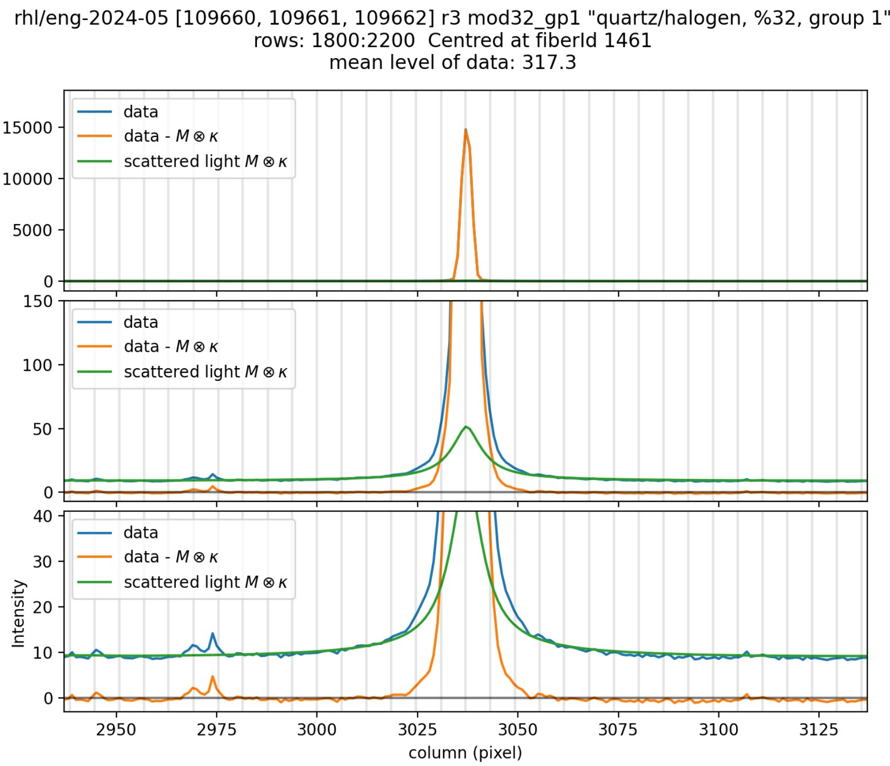

Figure 7: A horizontal cross-section (*i.e.* in the spatial direction) through the centre of an mod32 quartz-halogen image, but now with a model of the scattering removed (*cf.* Fig. 5).

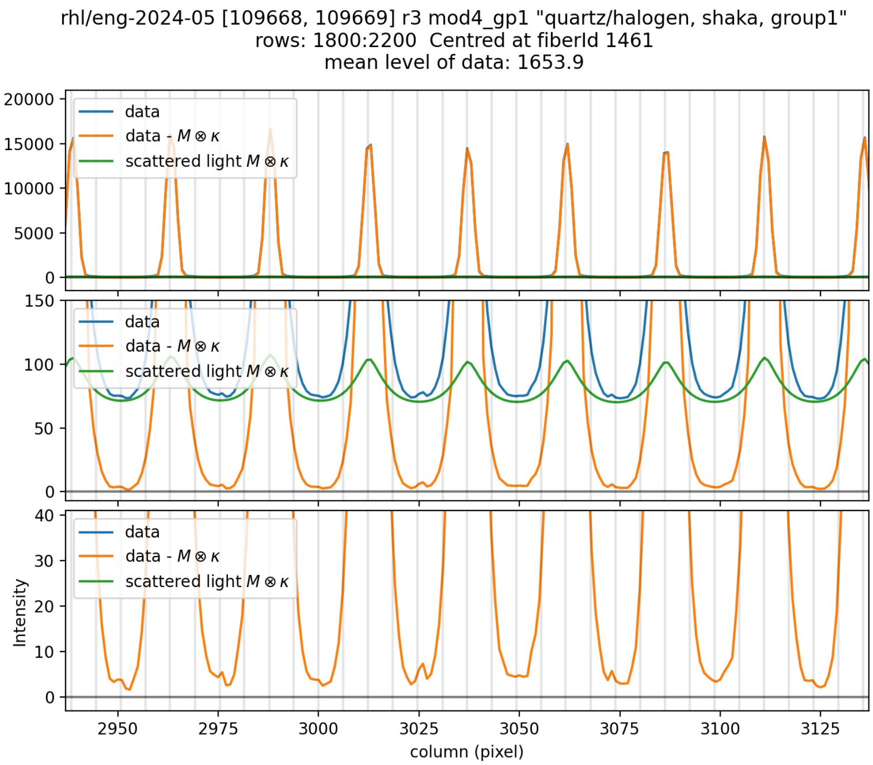

Figure 8: A horizontal cross-section (*i.e.* in the spatial direction) through the centre of an IIS quartz-halogen image, *i.e.* one with only the 16 engineering fibres illuminated, but now with a model of the scattering removed (*cf.* Fig. 2).

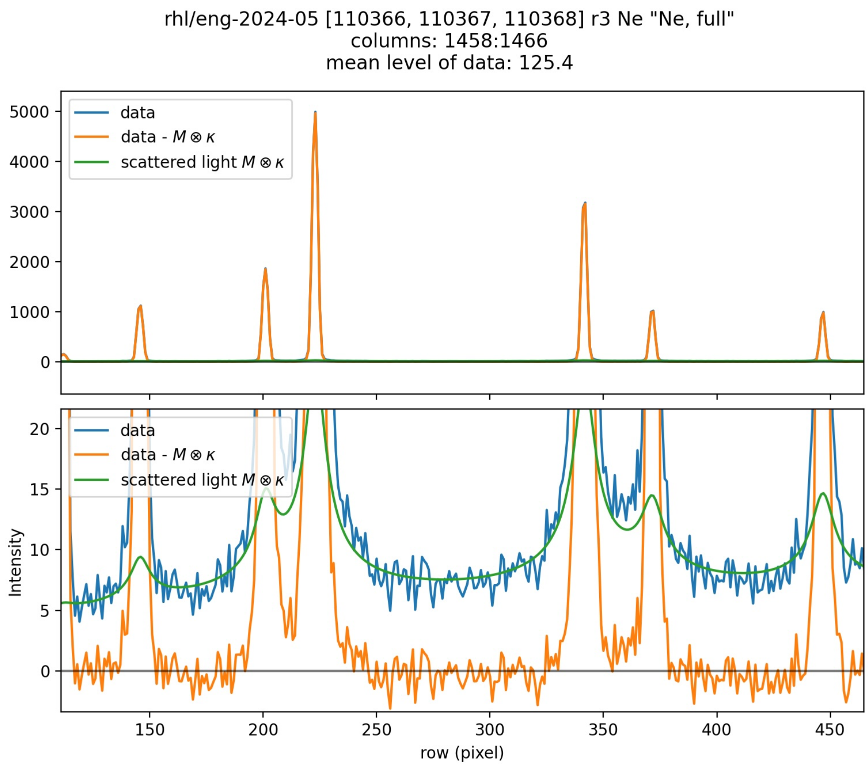

Figure 9: A vertical cross-section (*i.e.* parallel to the dispersion) through an illuminated fibre on a "full" Ne image (*i.e.* with all fibres illuminated). Note that there is no obvious oversubraction of the bright emission lines.

with an estimated background removed.

Having determined the scattering kernel κ, I applied it to black-spot data. The mod 32 data is simular, but requires dark-roach corrections; unsurprisingly the results are comparably good (Fig. 8; *cf.* Fig. 2). More ambitiously we can apply an identical kernel to the mod 4 (shaka) data; the results are shown in 8 which is to be compared with Fig. 2). The shaka residuals midway between the fibres are *c.* 3 ADU (*c.* 0.02% of the peak at the fibre centres).

4.1 Spatial Inhomogeneity of Scattering

Fig. 6 shows some structure, and a harder stretch brings it out even more clearly. We can use the mean spectrum in the blank regions between the IIS fibres to investigate this more quantitatively; see Fig. 10. It is not surprising to see a larger amplitude in the blue and efficient scattering back onto the detector for light falling just off the top edge of the CCD would explain the apparent excess at the very red end. The pixels adjacent to the serial register (at the blue end) are not used in PFS, and their absence would partially explain the lack of a similar upturn at the bottom; but this is all speculation.

4.2 Comparison of Scattering in r and m

We can use the $r3$ kernel to examine scattering in $m3$, which shares all components except the disperser. The original and scattering-corrected images are shown in Figs. 11 and 12. There are two obvious differences: there are bright bands at the top and bottom, and the overall amplitude is higher. Both of these can be explained if light passed by the dichroic but diffracted off the CCD is scattered back onto the detector. The second effect, at least, can plausibly be modelled by extrapolating the m spectra to the full r spectral range, and allowing that light to be scattered back onto the chip. Whether short r exposures to accompany all long m exposures remains to be explored.

4.3 Scattering Corrections for Bright Emission Lines

In Sec. 3.1 I fretted about the scattering correction oversubtracting light near the cores of lines; Fig. 9 suggests that this is not a major problem, but much

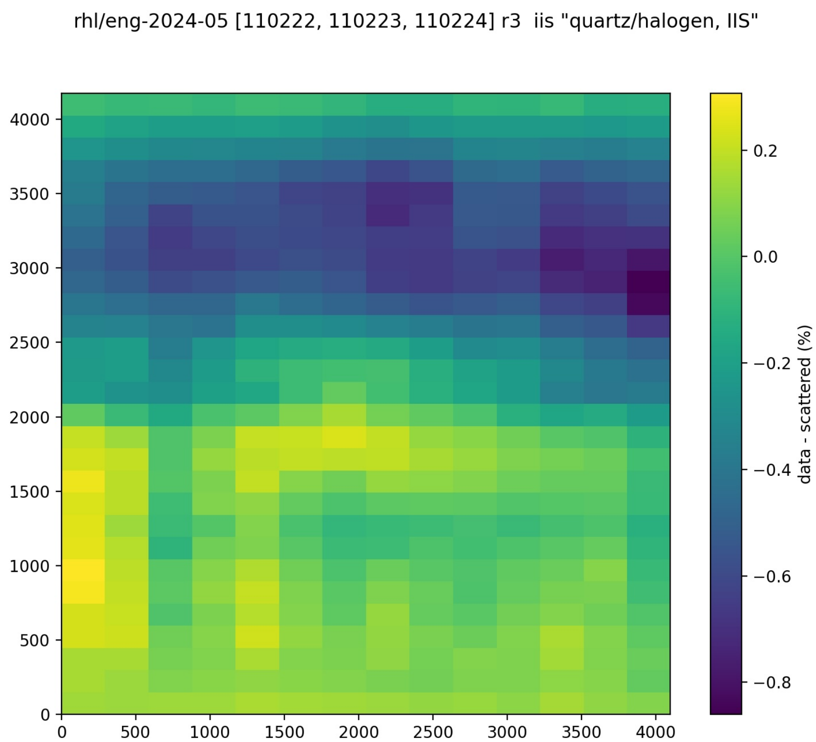

Figure 10: The residuals from the scattering for an IIS image, as seen in Fig. 6, evaluated away from the illuminated pixels. The values are normalised by the *total* flux in the image.

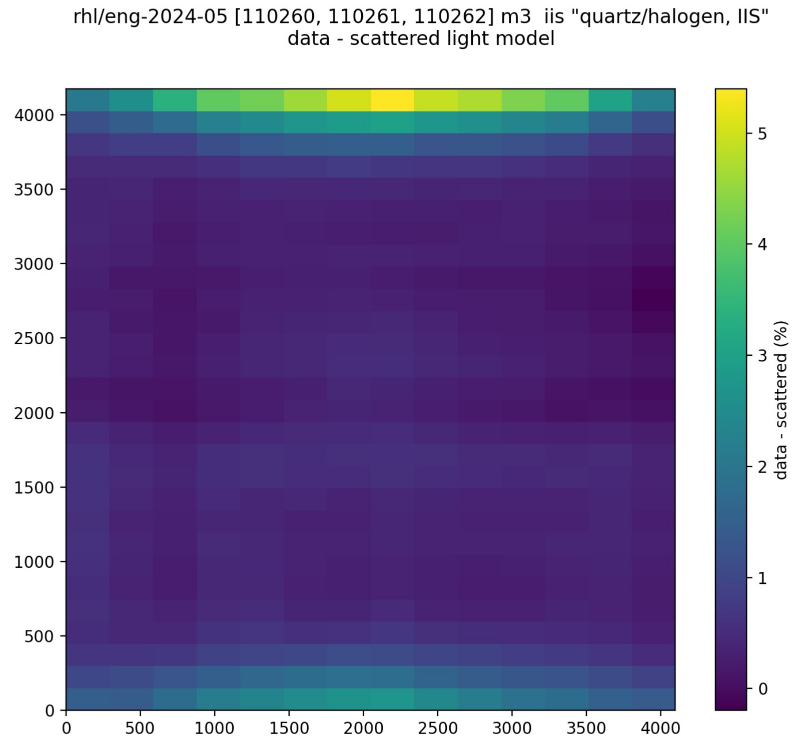

Figure 11: The scattering in an IIS image, as seen in Fig. 1 but now with the m grating; evaluated away from the illuminated pixels. The values are normalised by the *total* flux in the image. The horizontal band at row *c.* 2100 is still seen, but the dominant features are the excess light at top and bottom of the CCD.

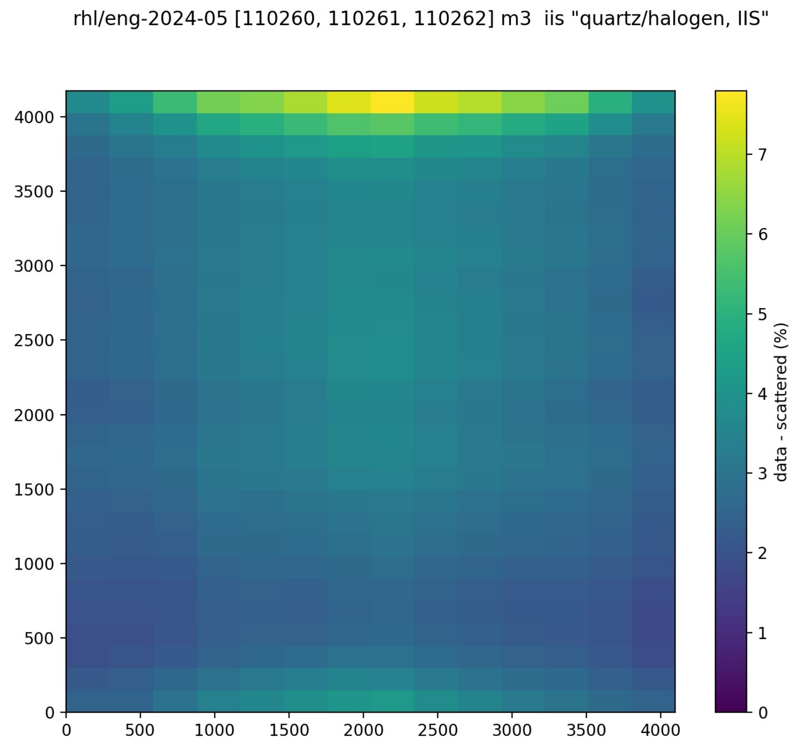

Figure 12: The scattering residuals in an IIS image using the *r* scattering kernel but now with the *m* grating; as before, evaluated away from the illuminated pixels. The values are normalised by the *total* flux in the image. The horizontal band at row *c.* 2100 is still seen, but the dominant features are the excess light at top and bottom of the CCD.

more work is required, especially on the `OH` Meinel lines which dominate the measured spectra in the red and near-IR.

5 Does this Work?

Using this approach to processing the PFS data, we can proceed to construct fiber profiles and extract spectra. The former is a pre-requisite to the latter as profiles are constructed using shaka data, with quite strongly overlapping scattering wings. A more complete analysis, using all available data and a model for the scattering kernel would be better in the long run.

5.1 Profiles

I rebuilt the profiles using the following visits, in four "groups" with different slit dithers:

What	Group 1	Group 2	Group 3	Group 4
quartz	109668..109676	109730..109738	109783..109791	109831..109831 109835..109841
dark roach	109652..109658	109722..109728	109770..109776	109823..109829

including a scattered light correction; these are the same visits and groups as were used previously to build the "old" (*i.e.* no scattered light subtraction) profiles.

A subset of the resulting profiles are shown in Figs. 13 and 14. The removal of the profile's wings in the new version is quite obvious.

5.2 Residuals

Fig. 15 shows an image of the residuals produced by subtracting a model of the data (the external product of the extracted spectra and the known fibre profiles, allowing for the updated detectorMap). It's smoothed with a $N(0, 1.5^2)$ Gaussian to suppress aliasing; Figs. 16 and 17 show unsmoothed zoomed-in regions. The residuals due to the astrometric "epoxy blob" errors in the detectorMap are prominent, along with various other interesting features. The overall structure in the scattering residuals seen in Fig. 10 is also present.

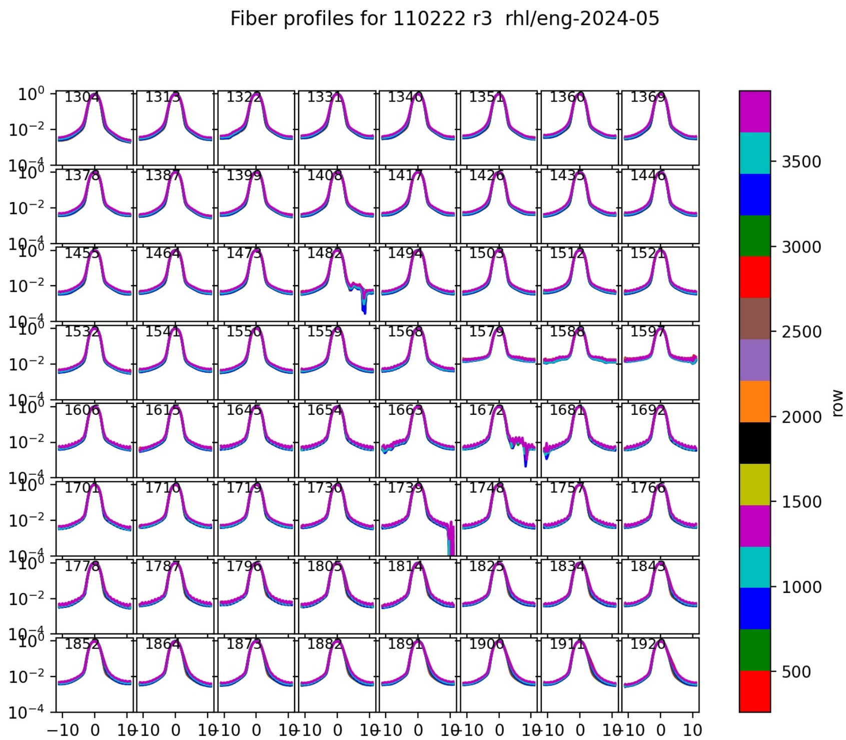

Figure 13: Fibre profiles estimated using `reduceProfiles.py` and `mergeFiberProfiles.py`, but without scattered light corrections.

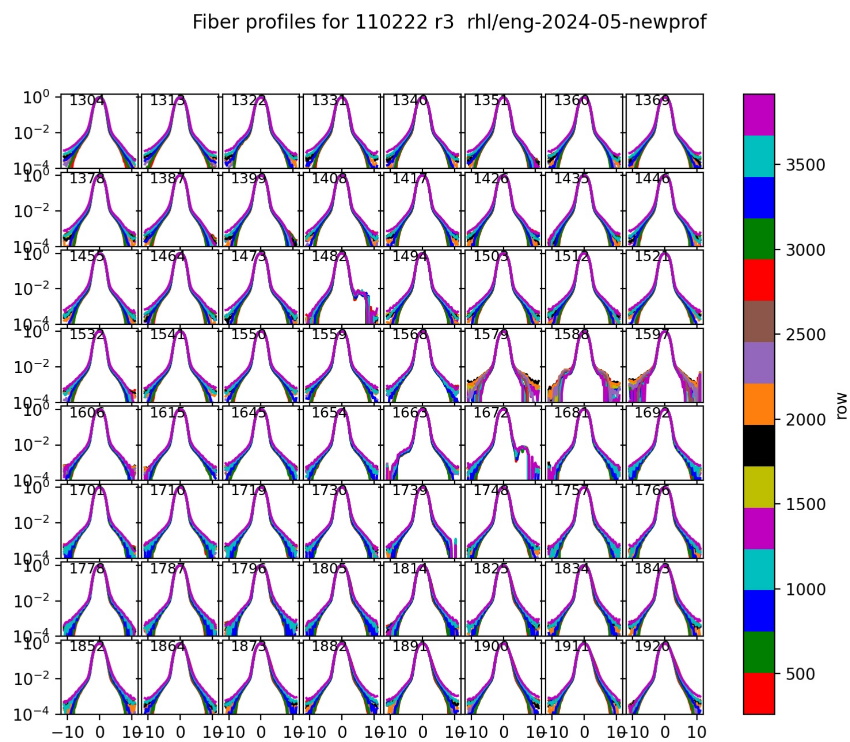

Figure 14: Fibre profiles estimated using `reduceProfiles.py` and `mergeFiberProfiles.py`, using the model-based scattered light corrections described in this note.

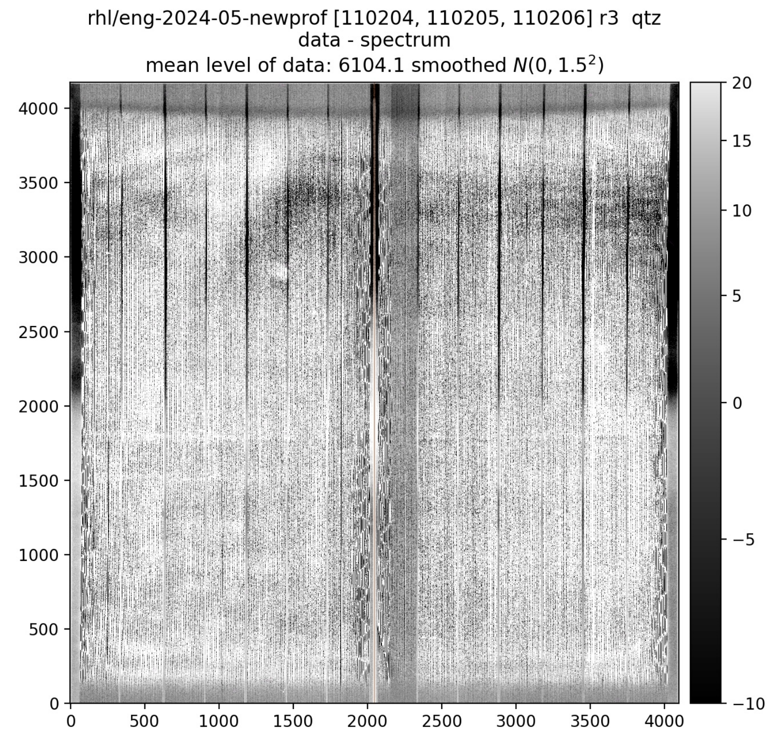

Figure 15: An all-fibre quartz-halogen image, after subtracting a model of the data (the external product of the known fibre profiles and the extracted spectra). The smoothing is only required to suppress aliasing artefacts at this zoom level. The spectra have peak fluxes of *c.* 35000 ADU (see Fig. 20).

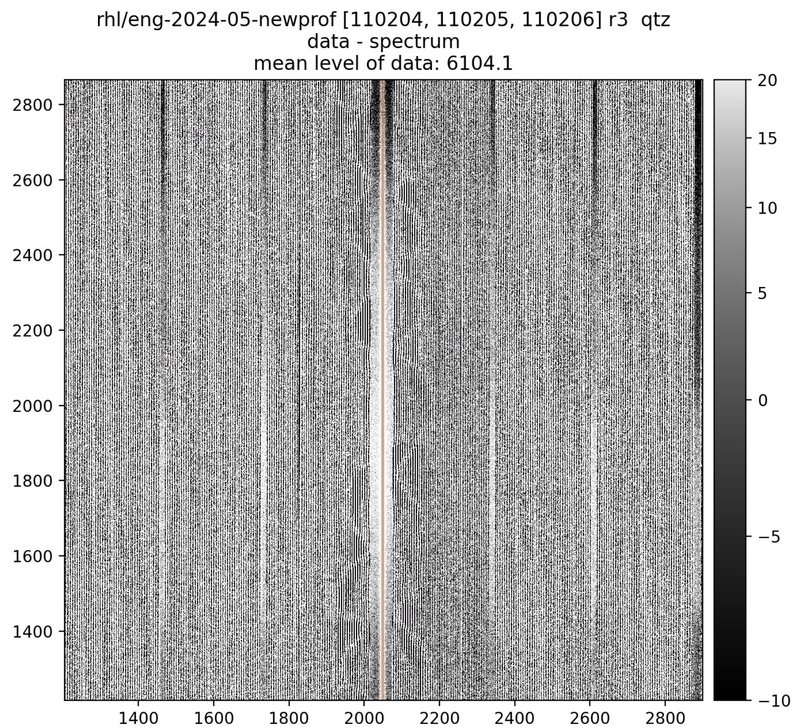

Figure 16: The central section of Fig. 15 at full resolution.

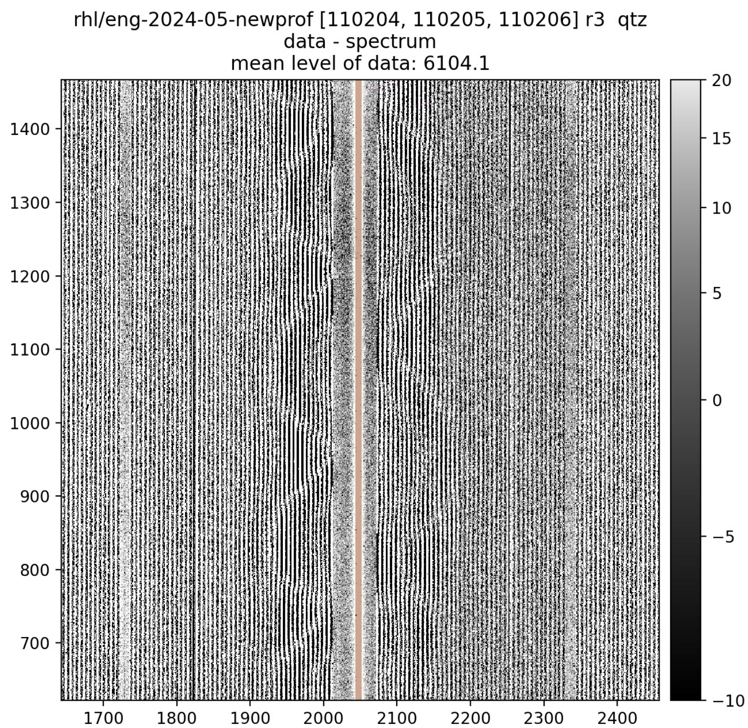

Figure 17: The central section of Fig. 16.

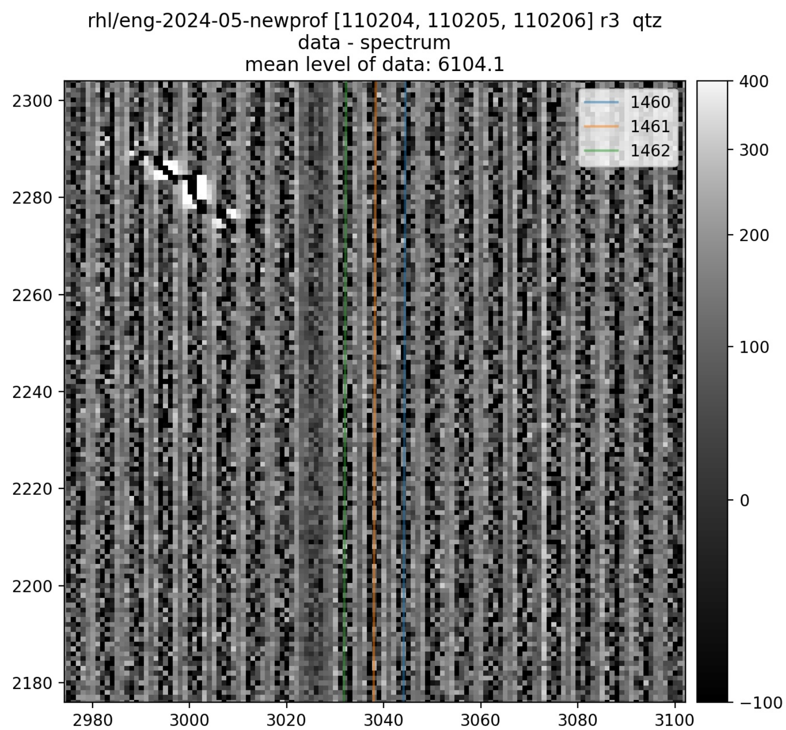

Figure 18: The section of Fig. 16 showing fiberId 1461 (*cf.* Fig. 20).

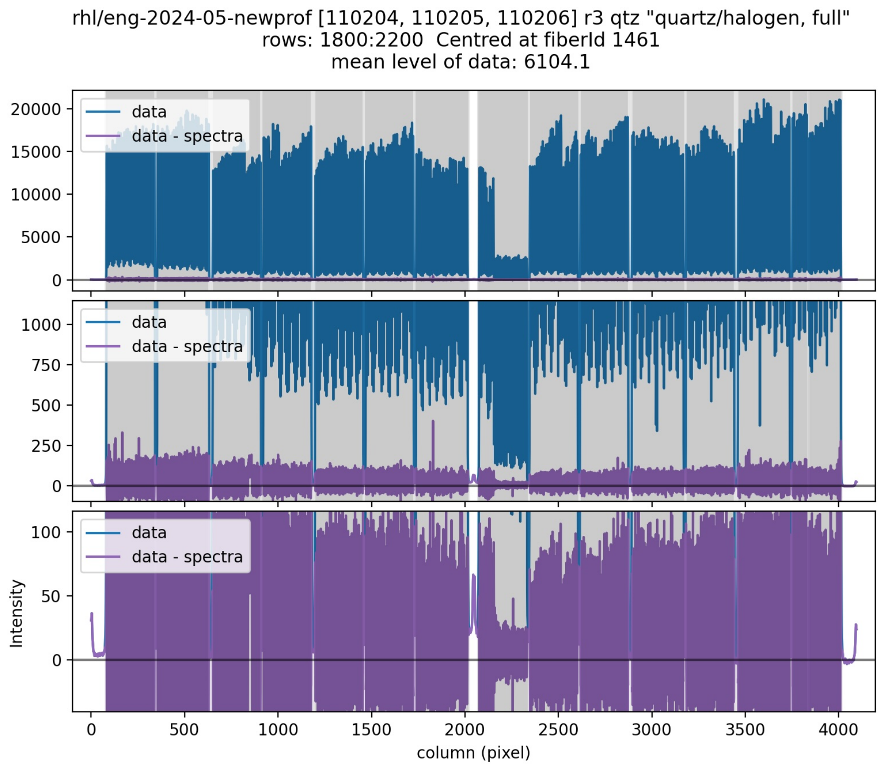

Figure 19: An all-fibre quartz-halogen image, after subtracting a model of the data (the external product of the known fibre profiles and the extracted spectra).

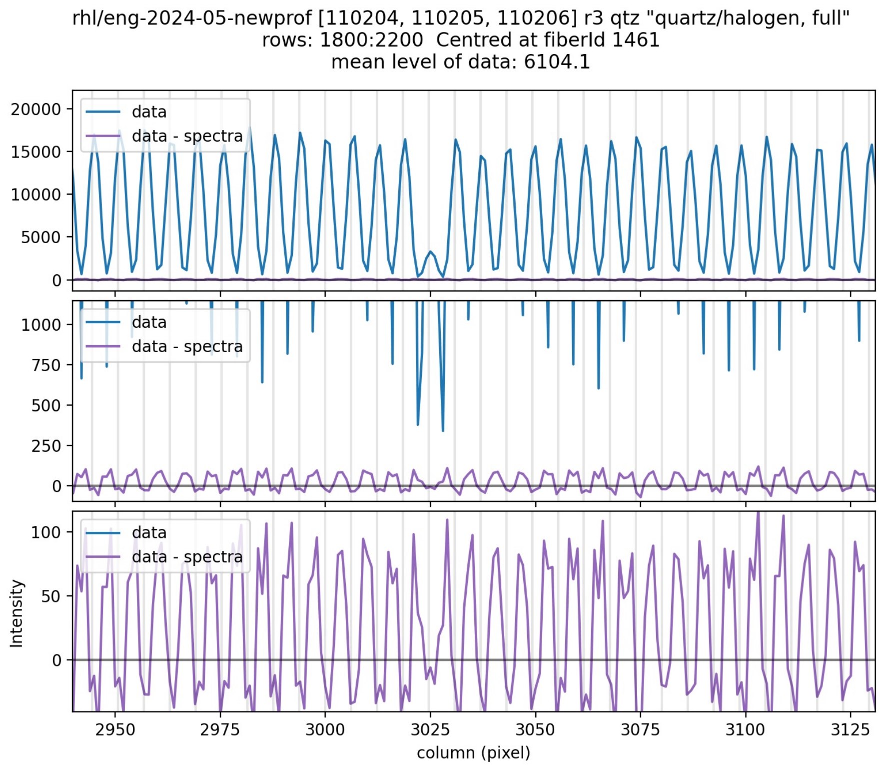

Figure 20: An all-fibre quartz-halogen image, after subtracting a model of the data (the external product of the known fibre profiles and the extracted spectra); the same data as Fig. 19.

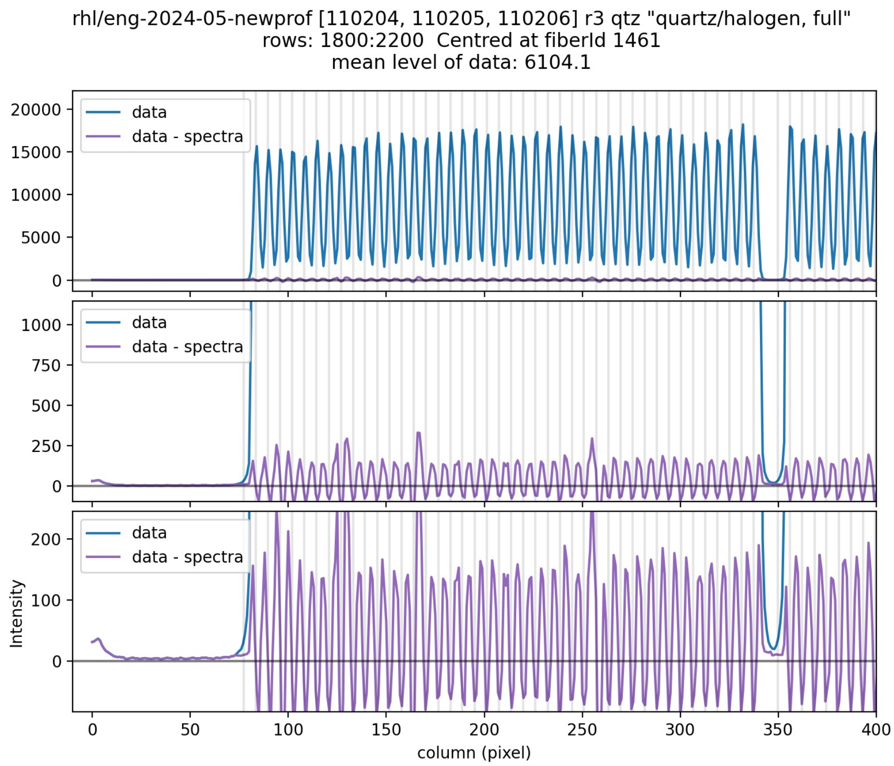

Figure 21: An all-fibre quartz-halogen image, after subtracting a model of the data (the external product of the known fibre profiles and the extracted spectra), showing the worst section of the residuals; the same data as Fig. 20. Note that the bottom panel's scale differs from that of Fig. 21.

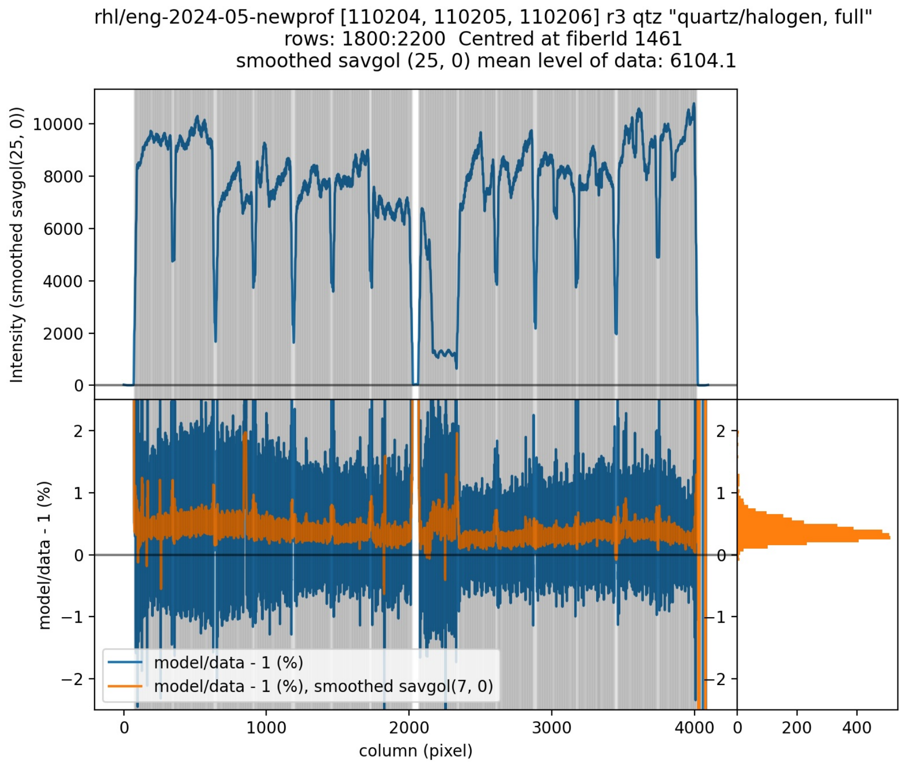

Figure 22: An all-fibre quartz-halogen image, showing the data smoothed with a 25-point boxcar filter, and the ratio of the residuals to that smoothed version of the data (corresponding to the signal level at the fibre). Additionally, the result of applying a 7-point boxcar to the residuals is shown, corresponding to smoothing over a fibre. The histogram reflects these smoothed residuals.

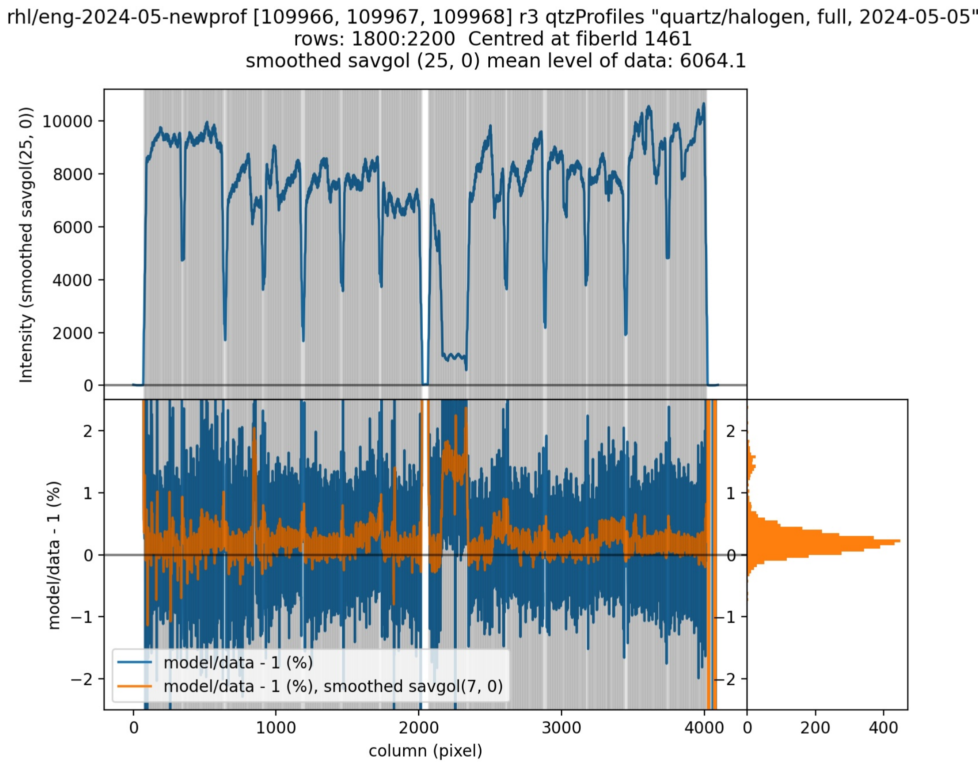

Figure 23: The same figure as Fig. 22, except for some data taken nearer to the time of the profile determination; unfortunately these data were taken at a different instrument angle, so the fibre fluxes are not identical. These changed fluxes appear to have little effect on the residuals, except perhaps for the faint fibres near $x = 2200$.

We can make this more quantitive by looking at the horizontal cross-section in Fig. 19; the errors seem to be less than *c.* 1%, and predominantly due to remaining issues in the fibre profiles. (Figs. 20; 21)

We can smooth these residuals to get an idea of the photometric accuracy. Fig. 22 shows the result; some data taken closer to the time we measured the profiles (Fig. 23) show smaller offsets, possibly due to evolution of the profiles. Both datasets were taken before we deliberately defocussed the instrument.