

Supplemental Material: Physically Meaningful Rendering using Tristimulus Colours

Johannes Meng* Florian Simon* Johannes Hanika Carsten Dachsbacher

Karlsruhe Institute of Technology

*joint first authors.

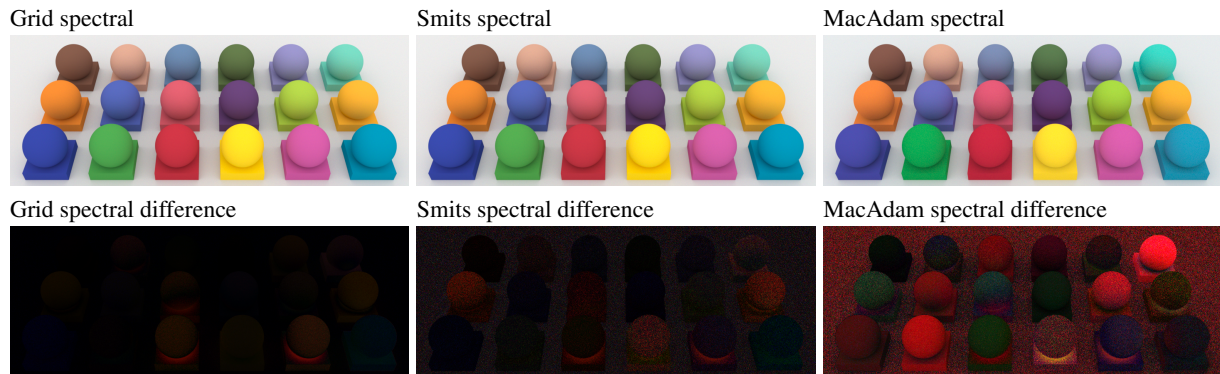


Figure 1: Renderings of measured spectral data (top) and difference to a spectral render using the original measured spectra (bottom, scaled $\times 8$ for display).

1. Comparing to Measured Spectra

Fig. 1 (top) shows spectral renderings of a scene with global illumination, and the difference to the ground-truth image rendered with measured spectra (bottom). The spectra for the top row were recomputed for the XYZ values of the measured spectra with Grid, Smits and MacAdam, respectively. The Grid method matches the original spectra the closest.

2. Surface Transport

The scene in Fig. 2 consists of four colourful chairs, with diffuse reflectance defined as $\text{rec709}(1, 0.01, 0.1), (1, 0.1, 1), (0.1, 0.1, 1)$ and $(0.1, 1, 0.1)$. The first three rows in Fig. 2 show spectral light transport using the indicated upsampling method in each row and the indicated solid mapping method in each column. Note that $\min \Delta E$ corresponds to a mapping into the solid of valid reflectances (MacAdam) for Smits and MacAdam, and to the space of smooth spectra for the grid-based method.

The bottom three rows show tristimulus transport in rec709 . The rows are labeled Smits, Grid, and MacAdam, to indicate the underlying set of spectra used for the respective solid mapping techniques, the three RGB renders without solid mapping are identical.

The RGB renders exhibit unnatural glow in the indirect

lighting, especially in the red and green tones. Upsampling these values to spectral reflectance without solid mapping results in a divergent render (see grid/uncorrected).

Note that RGB transport using solid mapping via Smits-upsampling and clamping (row 4, column 2) diverges as well. This is due to RGB coefficients greater than one, caused by the lack of round-trip precision of Smits' method; i.e. converting a RGB tuple $\in [0, 1]^3$ to spectral representation and back to RGB will not necessarily result in values within $[0, 1]^3$ again.

The solid mapped variants show a much improved tonality which gives more definition to the shape of the bunny, and more realistic appearance of the indirect light.




Figure 2: The chair scene with saturated colours. Rendering in `rec709` creates unrealistically bright colours (first column, bottom three rows). Creating spectra with Smits' or the Grid method requires additional solid mapping to fulfill energy conservation (first column, top three rows). We can solid map the RGB input values by clamping, scaling, or with $\min \Delta E$ (columns 3-4, respectively) such that the tristimulus transport images (rows 4-6) as well as the spectral renders (rows 1-3) produce more natural and converging results.

3. Differences Between Smits' Method and our Grid Method

Our Grid method can be seen as a generalisation of Smits' upsampling method [Smi99] that uses more spectral bins and allows for wider gamut colour spaces.

Our method differs from Smits' in several details, and we will discuss below the effects of this. We will show results obtained using an implementation of [Smi99] that is mostly a port of Smits' original code to Python. This implementation is contained in the submitted code, in the file `smits.py`.

Validation For validation, we compare the tabulated spectra shown in appendix A in [Smi99] to spectra obtained us-

ing our implementation. The results are almost identical, as can be seen in Fig. 3. Note that these spectra are defined on [380nm, 720nm].

Wavelength range The CIE recommends in [cie04] to use the wavelength range of [380nm, 780nm] for most applications, and we would like to follow this advice. As can be seen in Fig. 5 (left), this leads to larger values in the red and magenta spectra, however, which means that energy conserving spectra will overall be lower in intensity.

Discretisation In the original paper, Smits uses 10-bin and 20-bin discretisations. We use 80-bin discretisations, because the resulting spectra are a lot smoother (Note that we

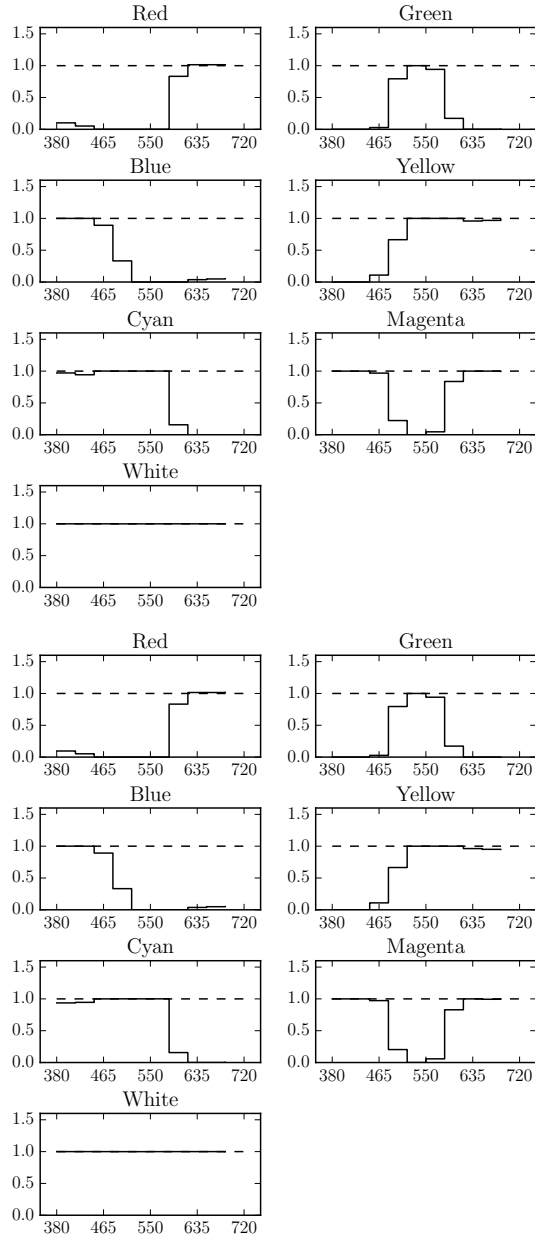


Figure 3: Basis spectra as shown in appendix A of [Smi99] (top) compared to spectra obtained using our implementation (bottom) for validation. The match is almost perfect. Slight differences can be seen, for example, at the long end of the yellow spectrum.

display the 80-bin spectra as piecewise linear functions, as recommended in [cie04]. This is in contrast to the 10-bin spectra, which we display as step functions to match Smits' paper). However, as can be seen in Fig. 5 (right), the number of bins actually influences the shape of the spectra obtained by the optimisation process. Specifically, the magenta spec-

trum has far higher values with 80 bins compared to the 10-bin variant, and the 80-bin red spectrum actually conserves energy.

Wide-gamut colour spaces We are interested in using wider colour spaces than sRGB. This is in theory possible with Smits' method, if one uses primaries closer to the spectral locus. Fig. 6 shows basis spectra for the Adobe RGB colour space, which has a more saturated green primary than sRGB, and thus a larger gamut.

It can be seen that with Adobe RGB primaries, more basis spectra have values greater than one, and so conservation of energy is a challenge. Also, the saturated green primary has a relatively narrow peak (cf. Fig. 5).

This narrow peak is visible in interpolated spectra, as well. Consider the Adobe RGB colour (0.5, 0.8, 0.5), which corresponds to sRGB (0.41, 0.8, 0.49). The resulting spectra are shown in Fig. 4.




Figure 4: Smits spectra for Adobe RGB (0.5, 0.8, 0.5) interpolated using sRGB and Adobe RGB primaries, and the equivalent spectrum computed using our Grid method.

The spectrum interpolated for sRGB primaries is smoother, even though both spectra represent the same colour. In the extreme case of monochromatic primaries (e.g. CIE RGB), the interpolated spectra would contain delta peaks.

Our Grid method offers two benefits: it is able to compute spectra for almost all visible colours, and the output spectrum does not depend on any specific input RGB colour space. Also, while colours close to the spectral locus necessarily contain narrow peaks, this effect is confined locally simply because the grid has more basis spectra which are distributed over the xy plane.

4. $1 - \epsilon$ Scaling Towards the White Point

After establishing grid points in x^*y^* colour space, as explained in Sec. 3.2. of the paper, we move all grid points towards the white point by some small distance. This removes some (highly saturated to monochromatic) colours from the set of colours that we can upsample to spectra.

The reason for this scaling is the spectrum discretisation. Discretised spectra cannot properly represent peaks narrower than the bin size. Spectra with such peaks occur very close to the spectral locus – monochromatic colours are simply delta peak spectra.

This makes the optimisation step unstable near the spectral locus. Moving grid points slightly toward the white point avoids this problem, while retaining the regular grid structure we exploit for fast interpolation.

References

- [cie04] *Colorimetry*. Tech. rep., Commission Internationale de l'Eclairage, 2004. [2](#), [3](#)
- [Smi99] SMITS B.: An RGB-to-spectrum conversion for reflectances. *Journal of Graphics Tools* 4, 4 (1999), 11–22. [2](#), [3](#)

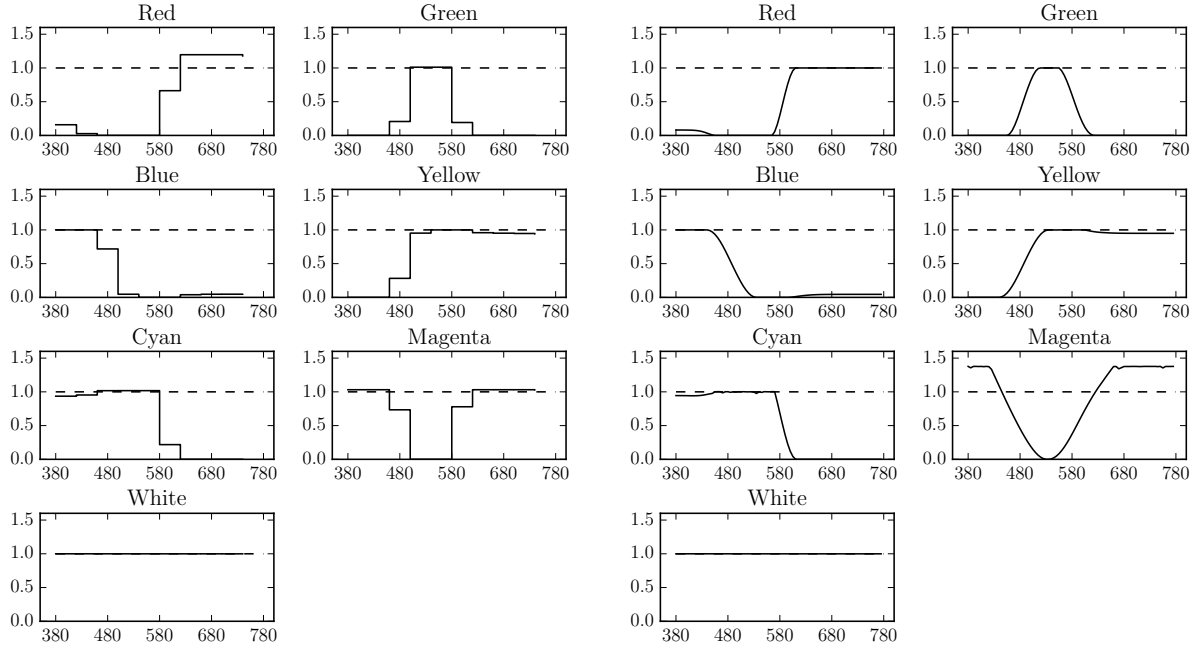


Figure 5: Smits basis spectra computed with our implementation for $\lambda \in [380\text{nm}, 780\text{nm}]$ with 10 (left) and 80 (right) bins. The RGB colour space is sRGB.

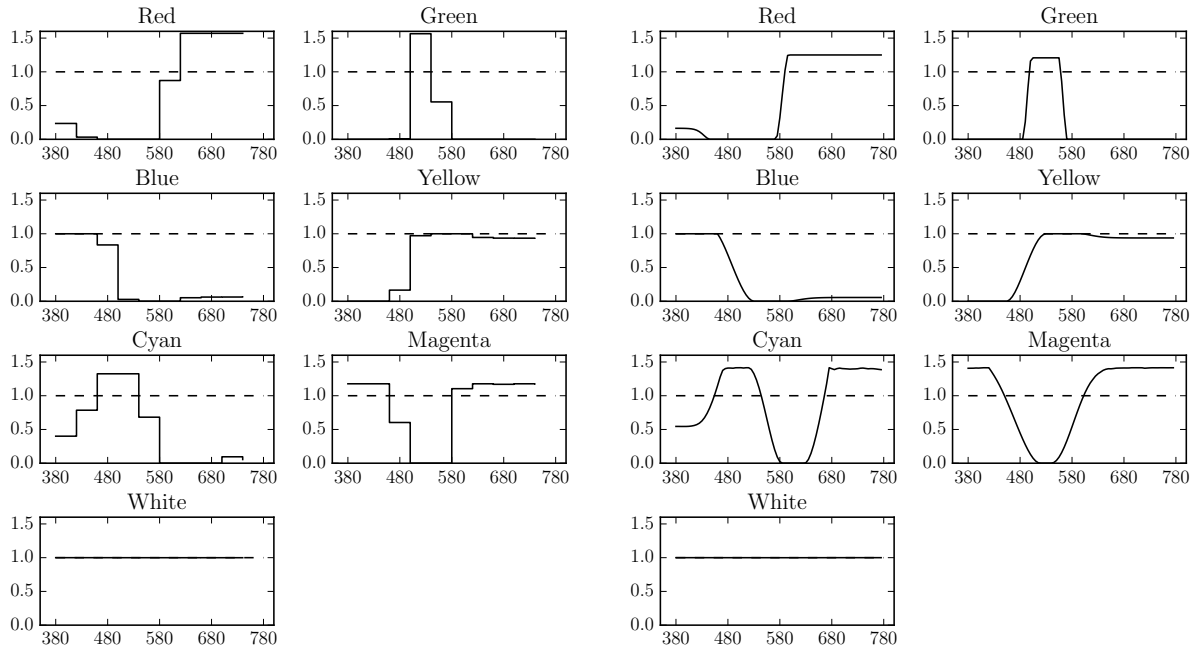


Figure 6: Smits basis spectra computed with our implementation for $\lambda \in [380\text{nm}, 780\text{nm}]$ with 10 (left) and 80 (right) bins. The RGB colour space is Adobe RGB.

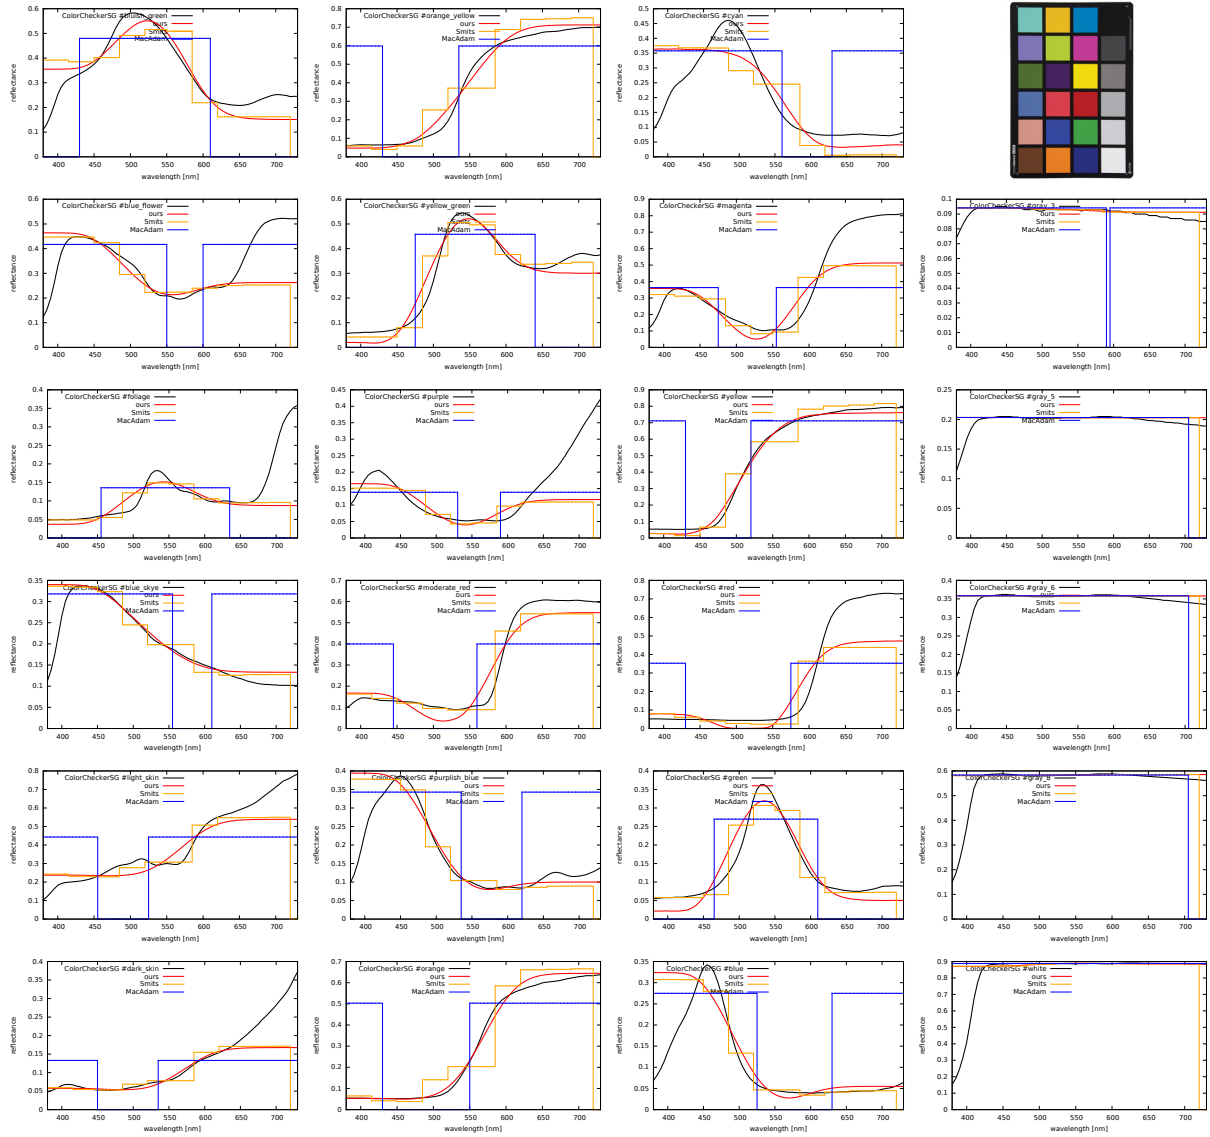


Figure 7: Comparison of Smits', MacAdam's and the Grid method for measured spectral data. Naturally occurring reflectances have a very smooth spectrum which is best fitted by the grid method.

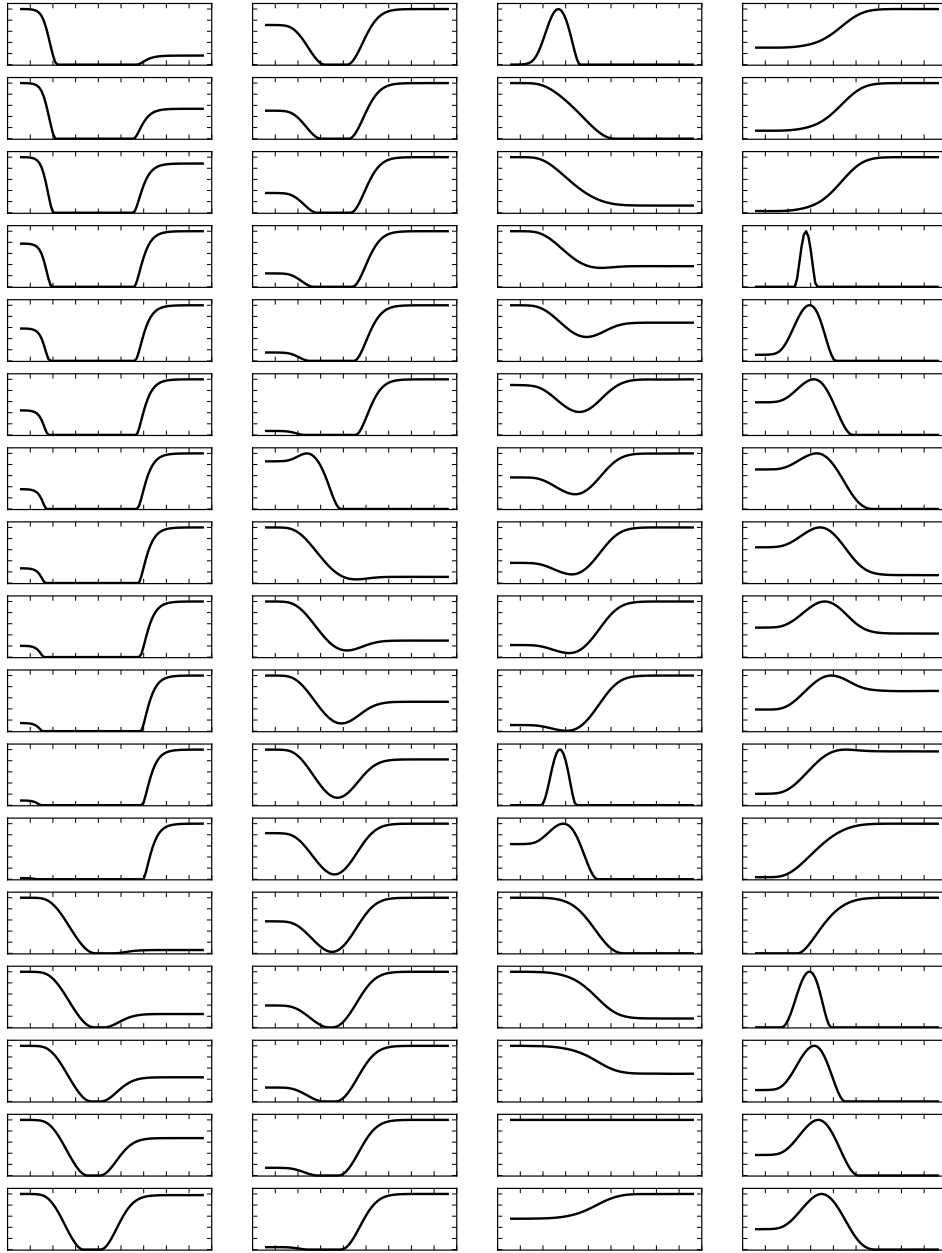


Figure 8: Complete set of the smooth spectra optimised for our grid points, first page.

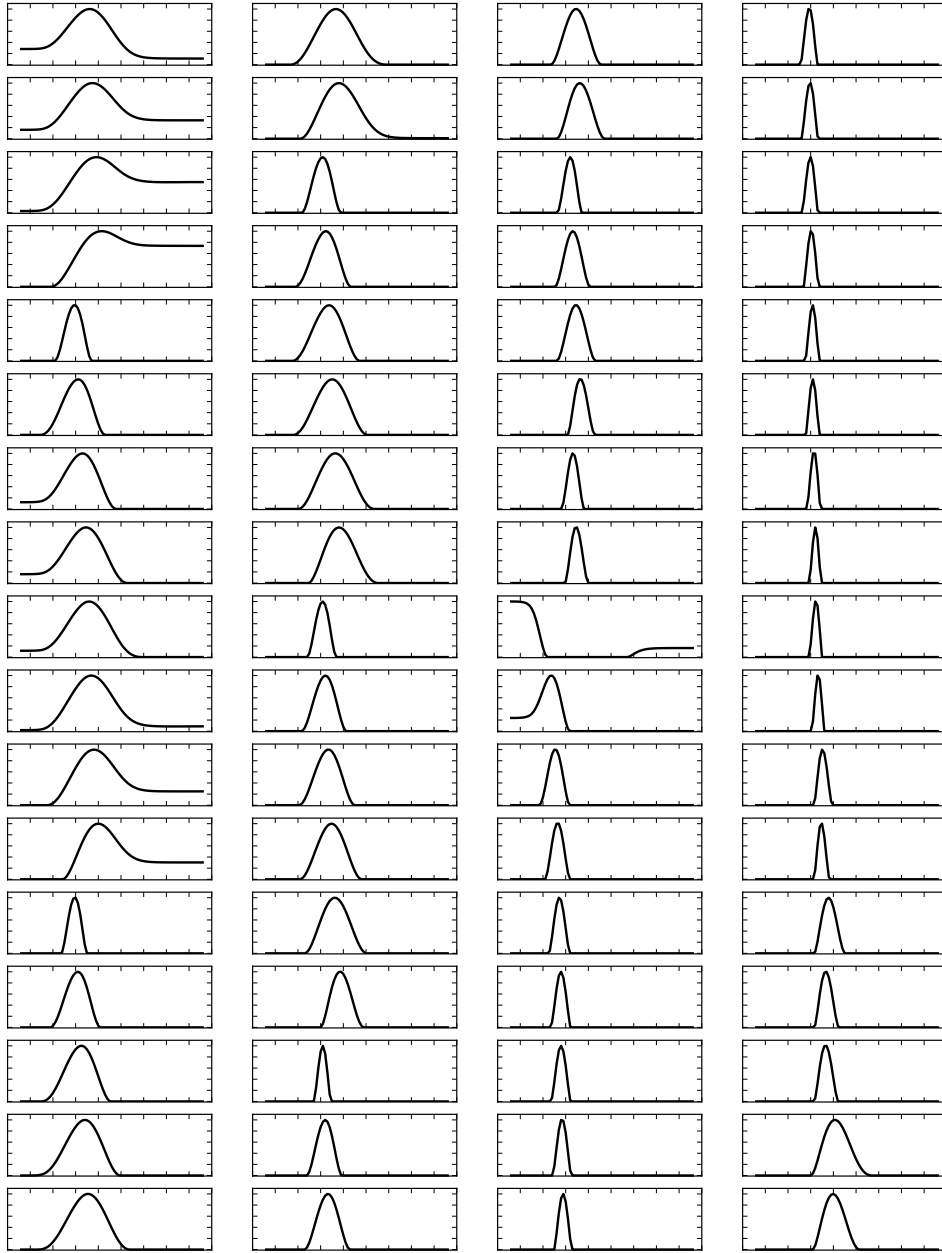


Figure 9: Complete set of the smooth spectra optimised for our grid points, second page.

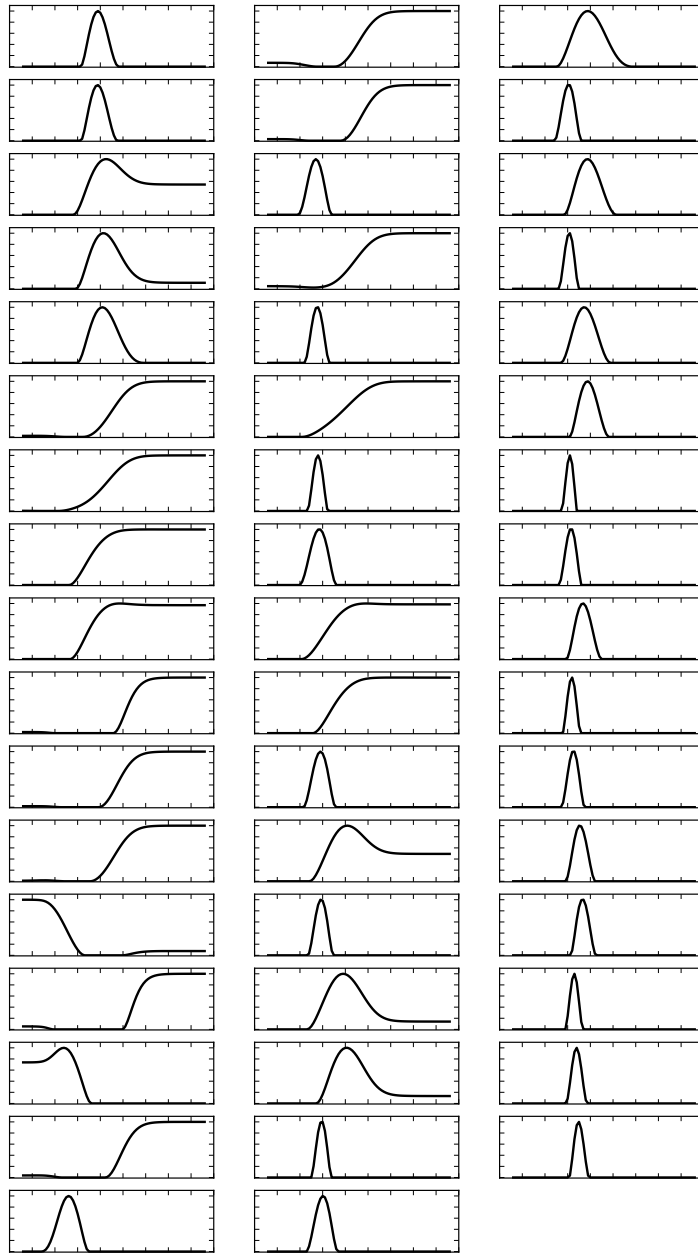


Figure 10: Complete set of the smooth spectra optimised for our grid points, last page.