





Perceptually Guided Automatic Parameter Optimization for Interactive Visualization – Supplemental

D. Opitz^{1,2} , T. Zirr^{1,2} , C. Dachsbacher¹ , L. Tessari² 

¹Karlsruhe Institute of Technology
²Intel Corporation




Figure 1: We show the progression through an optimization of Phong material parameters (diffuse, specular coefficient and specular exponent), light intensity of one point light and softness of shadows. The iteration i increases from left to right and top to bottom. The large images show the image rendered from the parameter configuration that scored the best perception loss $L(I)$ up to the respective iteration i . This is also the image that would be presented to the user. The small images are the renderings of the parameter configurations that have been proposed by the optimization algorithm during the search but did not surpass the currently best perception loss.

Transfer Function Mapping

Transfer functions are an essential tool in many visualization techniques, however, choosing a suitable transfer function is often highly data-dependent and adjusting many parameters to data distributions by trial and error can be cumbersome. Frequently, data points are distributed unevenly, and manual windowing or outlier clipping are needed to infer suitable parameters in the data domain. An example is shown in Figure 5, depicting different renditions of a smaller subcube in the Illustris simulation, a cosmological simulation of galaxy formation. The dataset consists of particles representing the comoving mass density and is rendered with a particle splatting approach. In the selected crop, the majority of particles have a very small mass density, as opposed to a few where the mass density is comparatively large. This kind of distribution causes over 94% of the data points to only use 0.1% of the range of the transfer function when using a linear scaling, resulting in the leftmost image of Figure 5. We can compress the space taken up by outlier particles by a non-linear remapping of mass density values prior to classification by the transfer function:

$$\text{remap}(x) = \frac{\log(2^\alpha x + 1)}{\log(2^\alpha + 1)}. \quad (1)$$

The transformation in Equation 1, where α controls the degree of nonlinearity, is applied to values before the transfer function, mapping from $[0, 1]$ back to the same interval. A good choice for α maximizes the ranges used by particles in the transfer function, while retaining distinguishable high and low mass density areas. We use our optimization algorithm to automatically find $\alpha \geq 0$. The result is shown in the center of Figure 5: the optimization does not simply stretch the histogram towards equalization (see higher α s in the sequence), but also considers the spatial distribution of particles. Consequently, the optimal α -value would vary with the viewpoint (if the user chooses to continuously optimize the value during exploration). As an alternative to non-linear remapping, we can clamp the values to a predefined minimum and maximum before applying the transfer function. These parameters can be found by our optimizer, see the supplemental material (Fig. 5) for results and images.

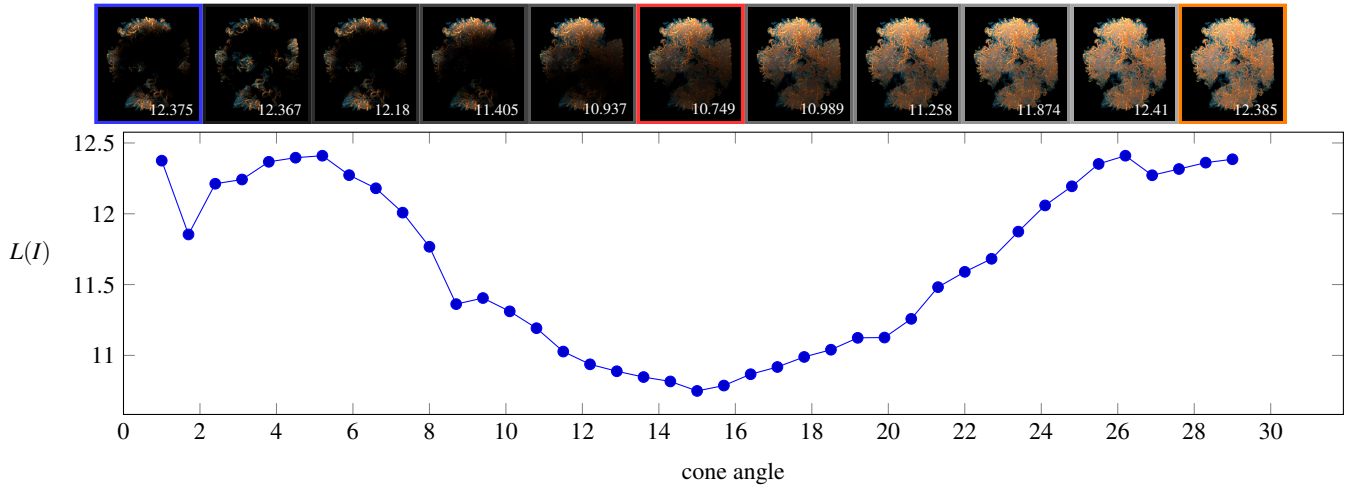


Figure 2: We show a series of tuning a single parameter—here the cone angle controlling the softness of LPFV shadows—for the vortex cascade dataset. The cone angle can assume floating values that are larger than zero, in our case we plotted the perceptual loss $L(I)$ for 40 equidistant cone angles in $[1, 29]$. The top images are representatives for the respective cone angle on the x-axis, their bottom-right number also showing the $L(I)$. The red framed image marks the minimum perceptual loss in the series. The blue and orange framed images mark the beginning and end of the series with cone angles 1 and 29, respectively.

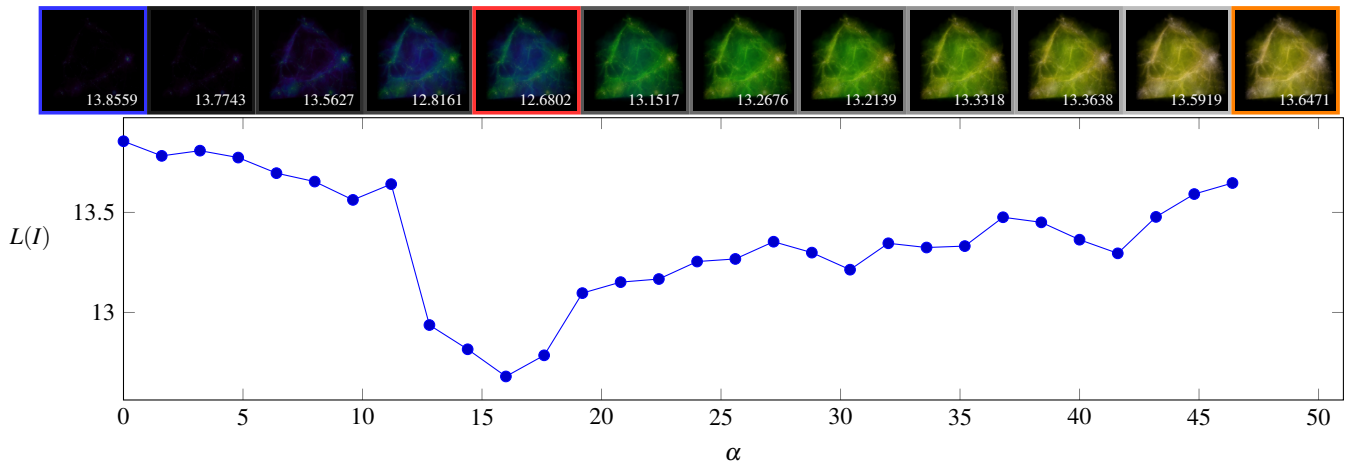


Figure 3: We show a series of tuning a single parameter for the illustris dataset. The images are rendered with point splatting, and the tuned parameter is a mapping factor α reweighing the transfer function, see Section 4.2.3 in the paper for more details. The mapping factor α can assume floating values that are larger than zero, in our case we plotted the perceptual loss $L(I)$ for 40 equidistant values in $[0, 47]$. The top images are representatives for the respective α on the x-axis, their bottom-right number also showing the $L(I)$. The red framed image marks the minimum perceptual loss in the series. The blue and orange framed images mark the beginning and end of the series with $\alpha = 1$ and $\alpha = 47$, respectively.




Figure 4: We show results of the parameter optimization from different experiments. The left two image pairs (top: resulting visualization, bottom: model output image) were obtained after optimizing the cone angle for softness of shadows and Phong material parameters (diffuse, specular coefficient and specular exponent) simultaneously. In the two right images additionally the intensity of the light source has been optimized. We can see that the parameter search finds different results in each experiment, but still converges to a viable solution in every case. Furthermore, parameters for light intensity converge to similar values on the right hand side as in the hand tuned version on the left hand side.




Figure 5: We show particle renderings of the Illustris simulation. The bottom sequence of images shows increasing settings for α , the parameter controlling the power of the logarithmic scale before particle values are classified with the transfer function. The lower image in this sequence is the output image of our perceptual model. The top left image uses a linear mapping of particle values to the transfer function. The top right uses an aggressive logarithmic mapping with a high value for α . The center image shows the result of the parameter search, which in this case is the global minimum. Particle values are stretched over a larger range in the transfer function.




Figure 6: We show optimization results for different mappings of density values and transfer functions. In the left the density is mapped linearly to the transfer function. In the right column density is mapped with a non-linearity as defined in Equation 5. In the middle column density is mapped linearly but is clamped to a minimum t_{min} and maximum t_{max} value. To ensure that the optimization algorithm selects parameters with consistent $t_{min} < t_{max}$, we are optimizing t_{min} and $\delta_t = t_{max} - t_{min}$ instead. The result of the parameter search always finds that $t_{min} = 0$ and $t_{max} \in [1e^{-7}, 4e^{-7}]$




Figure 7: We demonstrate that the optimization can tailor to specific display devices with different tone reproduction curves. We chose an easy gain-offset-gamma tone reproduction curve (trc) of a monochromatic CRT display, because differences in the resulting images are not subtle but quite obvious and better recognizable when viewed in pdf or printout. We show the optimized result and then apply our trc operator on the left hand side, while on the right hand side we apply the trc every time before the model evaluation. The $L(I)$ scores have been computed for the image after our chosen trc has been applied. The optimization result of the left hand side image is the same as in Fig. 1, converged at $L(I) = 10.135$ and did not use a special trc prior to the model evaluation.

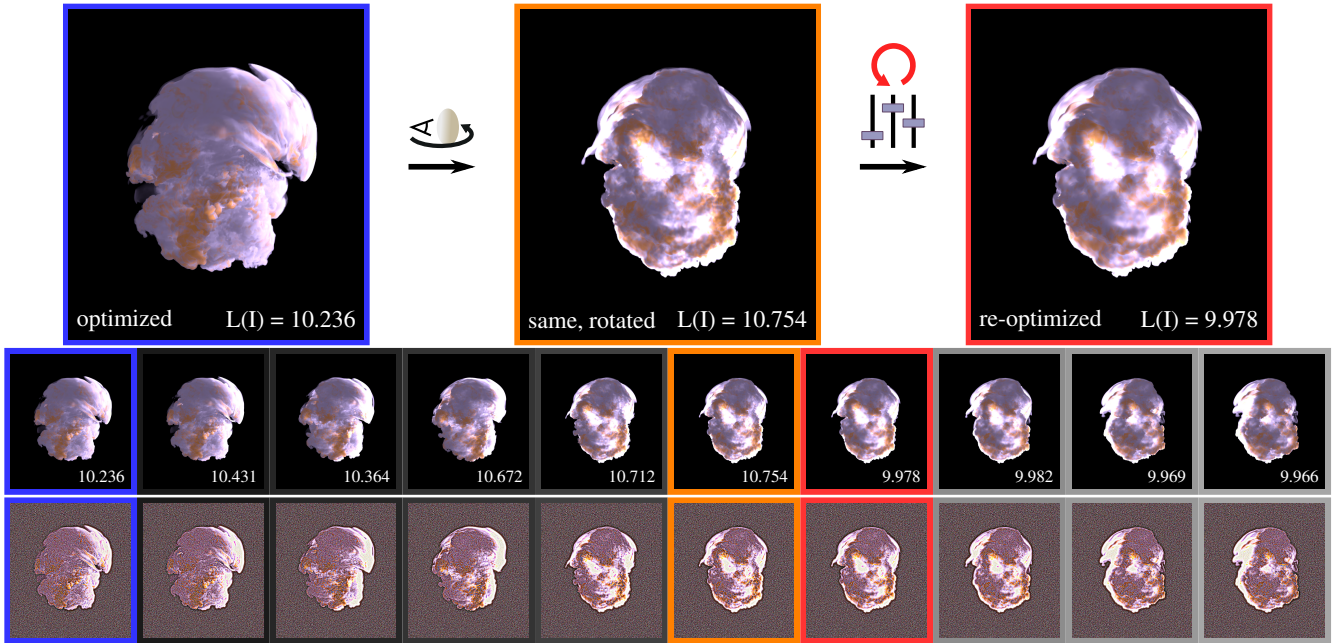


Figure 8: We show the supernova dataset in an animated sequence, where it is rotated counter-clockwise while keeping the parameter optimization enabled in the background. The bottom sequence shows different time steps, the lower image is the perceptual model's output image. The top left image shows the beginning of the rotation where parameter optimization has already been finished. Optimized parameters are the filter size for the softness of shadows and diffuse material coefficient. In the top center image the dataset has been rotated between viewer and light source which triggers the parameter search to find a better configuration because of high frequency features. The top right image shows the new found optimum.




Figure 9: We show the progression through an optimization of Phong material parameters (diffuse, specular coefficient and specular exponent), light intensity of two point lights and softness of shadows. The model was rendered by extracting an iso-surface from a volume. The iteration i increases from left to right and top to bottom. The large images show the image rendered from the parameter configuration that scored the best perception loss $L(I)$ up to the respective iteration i . This is also the image that would be presented to the user. The small images are the renderings of the parameter configurations that have been proposed by the optimization algorithm during the search but did not surpass the currently best perception loss. This should serve as a proof of concept, that optimization with the perception loss may be used in surface shaded visualization techniques as well.

Automatic Viewpoint Selection

We also used our perceptual loss metric for the automated selection of engaging viewpoints on a dataset and show results in Figure 10. We optimize three camera parameters: The distance from the origin and two spherical coordinates. The polar angle is divided into N equisized areas, and our optimizer finds an optimal viewpoint for each. This division is needed as it restricts the search space and enforces a well-distributed set of different viewpoints.




Figure 10: We find good view positions of the Bonsai dataset with our parameter optimization. We configure the view position in spherical coordinates for 4 separate intervals in the upper hemisphere with a azimuthal angle $\phi \in [0^\circ, 360^\circ]$, polar angles $\theta_i \in [i30^\circ, (i+1)30^\circ]$, and a radius with only a lower bound such that the view position is not inside the dataset. The large images in the top are the converged results for every interval. The matrix of images below shows representative viewpoints out of the first three intervals with fixed radius (we omit the fourth row since the tree is mostly hidden when viewed from below). The approximate position of the converged images is marked with colored circles. Interestingly, the optimization does not choose to zoom in to the minimal pre-specified radius. Also, it hides the flat, uni-colored cutoff of the plant pot.

1. JND Color Noise in $L^*a^*b^*$

The LAB to XYZ transformation is given by:

$$X = X_n f^{-1} \left(\frac{L^* + 16}{116} + \frac{a^*}{500} \right), \quad (2)$$

$$Y = Y_n f^{-1} \left(\frac{L^* + 16}{116} \right), \quad (3)$$

$$Z = Z_n f^{-1} \left(\frac{L^* + 16}{116} - \frac{b^*}{200} \right), \quad (4)$$

$$f^{-1}(t) = \begin{cases} t^3 & \text{if } t > \delta \\ 3\delta^2 \left(t - \frac{4}{29} \right) & \text{otherwise} \end{cases} \quad (5)$$

$$\delta = \frac{6}{29}. \quad (6)$$

The first-order Taylor expansion around a given color coordinate $(L^*, a^*, b^*)^T$ yields the transform:

$$T = \begin{bmatrix} T_X/116 & T_X/500 & 0 \\ T_Y/116 & 0 & 0 \\ T_Z/116 & 0 & -T_Z/200 \end{bmatrix},$$

where

$$T_X = X_n f^{-1'} \left(\frac{L^* + 16}{116} + \frac{a^*}{500} \right),$$

$$T_Y = Y_n f^{-1'} \left(\frac{L^* + 16}{116} \right),$$

$$T_Z = Z_n f^{-1'} \left(\frac{L^* + 16}{116} - \frac{b^*}{200} \right),$$

and

$$f^{-1'}(t) = \begin{cases} 3t^2, & \text{if } t > \delta \\ 3\delta^2, & \text{otherwise} \end{cases}.$$

Let σ_{JND} be the just noticeable luminance difference in the Y -coordinate of XYZ, then the radius of the isotropic spherical distribution of colored JND noise in LAB is:

$$\sigma_{LAB} = \frac{116}{T_Y} \sigma_{JND} = \frac{116}{Y_n} \left[f^{-1'} \left(\frac{L^* + 16}{116} \right) \right]^{-1} \sigma_{JND}.$$

We obtain the anisotropic colored JND distribution in XYZ by applying the transformation matrix T . For a random sample r_1 within the unit sphere, the resulting JND noise sample in XYZ space, r_{XYZ} , is:

$$r_{XYZ} = T \sigma_{LAB} r_1.$$