Chapter 23
Visual Contrast Preserving Representation of High Dynamic Range Mathematical Functions

Juha Jeronen

Abstract When Gaussian distributed inputs, representing model parameters with some measurement error, are mapped through certain mechanical vibration models, the corresponding output probability distribution exhibits an approximately logarithmic data value distribution (in the histogram sense) with a high dynamic range (HDR). We look at applying tone mapping techniques from HDR photography to produce a low dynamic range, visual contrast preserving representation of such high dynamic range mathematical functions—thus enabling HDR plotting. This makes it possible to visualize HDR functions, displaying their structure in a clear manner on standard low dynamic range media such as computer screens and print. The advantages over simple logarithmic scaling are the visual contrast preservation and data adaptivity. Comparing to histogram equalization, the present approach has the advantage of not exaggerating small contrasts. Three methods are suggested and demonstrated on two mechanical vibration problems: transverse waves in a classical vibrating string, and the dynamic out-of-plane behaviour of an axially travelling panel submerged in axial potential flow.

23.1 Introduction

What are high dynamic range mathematical functions, and where would one want to visualize them? The motivation for this study comes from physics and engineering problems where model input is never exact. To obtain reliable analysis results, it is desirable to find out how stable the predictions of a given model are with respect to small perturbations in model input, and how large the expected range of output is. When input uncertainties are present, instead of a single solution, one obtains a solution set corresponding to the admissible inputs.

Adopting a direct statistical approach to uncertainty analysis, it is possible, in the case of computationally lightweight models, to approximate the solution set

J. Jeronen (✉)
Department of Mathematical Information Technology, University of Jyväskylä, P.O. Box 35 (Agora), 40014 Jyväskylä, Finland
e-mail: juha.jeronen@jyu.fi

S. Repin et al. (eds.), *Numerical Methods for Differential Equations, Optimization, and Technological Problems*, Computational Methods in Applied Sciences 27,
DOI 10.1007/978-94-007-5288-7_23,

directly. This is done via sampling the multidimensional input probability distribution (corresponding to the set of admissible parameter combinations when the uncertainty is accounted for; Latin hypercube techniques can be used, see, e.g., [5, 23, 24, 33, 34, 46]), and mapping the samples through the model being analyzed. Different statistical quantities can then be computed based on the discrete output sample. See, e.g., [25] for a general review of statistical uncertainty and sensitivity analysis techniques.

In order to facilitate an intuitive understanding of the results, it is also possible to visualize the output probability density directly.[1] The density of the resulting output probability distribution is estimated from the discrete output sample (using, for example, kernel density estimation; see, e.g., [7, 8, 16, 43]), and the resulting probability density field is then plotted. See Figs. 23.7–23.9 below for examples.

Applying this methodology to mechanical vibration problems, it was found that the resulting probability field does not lend itself to traditional linear scaling for visualization. The reason is that the histogram of the data is distributed, approximately, in a logarithmic manner, and the dynamic range spanned by the data far exceeds the representable range of a computer display: it is a high dynamic range mathematical function.

Traditionally, logarithmically distributed data with a high dynamic range occurs in contexts such as audio and high dynamic range (HDR) photography. Thus, considering the task of visualizing a HDR function, it is natural to seek methods for representation of HDR data that have been developed in these fields.

The need for methods to represent HDR data in both of these classical fields is clear. The human eye has a range of five orders of magnitude in light intensity (e.g., [12]), while computer displays (in terms of photometric light intensity) are limited to two. Similarly, standard digital audio is sampled at 16 bits per sample, which gives a range of about four orders of magnitude (logarithmically, 90 dB).

In audio processing, the standard solution is to work on the decibel scale, which is logarithmic, allowing an at-a-glance representation of the HDR data. However, in HDR photography literature, it is well known that simple logarithmic scaling has a tendency to eliminate contrast (e.g., [28]).

For the purposes of HDR photography, special *tone mapping* algorithms have been developed for the purpose of representing, in a visually accurate manner, high dynamic range scenes on low dynamic range media such as regular computer displays and print. The basic idea behind tone mapping is that the human visual system is sensitive to differences in light intensity, but not to absolute intensities [28]. In terms of signal processing, tone mapping can be seen as data-adaptive dynamic range compression. As tone mapping algorithms are important for the focus of our study, we will review the related literature in the next section.

The application of tone mapping techniques for plotting HDR mathematical functions can be seen as a natural extension of the ideas of Park and Montag [35],

[1] SAVU, **S**ample-based **A**nalysis and **V**isualization of **U**ncertainty: https://yousource.it.jyu.fi/savu/codes/ Link cited 13 Jan 2012.

who investigated the use of tone mapping for representation of data from astronomical and medical imaging, captured at wavelengths other than those of visible light.

There is a large body of research in HDR signal processing, which is not limited to tone mapping only. Some examples follow. Display adaptation in different ambient lighting conditions is discussed in [30], and feature classification in areas obscured by shadows in [9]. Adapting HDR images to target devices with drastically different dynamic ranges is considered in [50].

The study [11] concentrates on the problem of obtaining HDR radiance maps by stitching together multiple LDR (i.e. the usual kind of) photographs taken with different exposure parameters; this is similar to stitching together a panorama, but along the light intensity axis. The same problem has been discussed in [29] earlier, and [1] is a recent technical report on the subject. Integration of computer-generated (3D rendered) objects into HDR photographs has been investigated in [10].

23.2 Tone Reproduction Operators and HDR Plotting

Considering visual contrast preserving representation of high dynamic range mathematical functions, it seems that the task has received very little attention. Some partly relevant studies exist; for example, [51] discusses volumetric visualization of HDR data. The study [36] talks about the potentially deceptive appearance of logarithmic scaling for data obtained from cytometry, and suggests a new scaling method for that particular application.

Most interestingly, in [35] a psychophysical study was carried out on applying different tone reproduction operators (TROs, see below) on HDR data from, e.g., medical, astronomical, and radar sources. The authors conclude that aside from some general trends, the appropriate tone reproduction operator ultimately depends on the kind of data, and on expert opinion in the specific field of application. The study concentrates mainly on qualitative aspects of the user experience, and does not consider extension into visualization of mathematical functions.

Tone reproduction operators (TROs) are used in the conversion of HDR images for display on standard dynamic range (low dynamic range; LDR) devices and media, which include regular computer displays and print. The conversion is known as *tone mapping*. TROs scale the data in an adaptive manner to maintain visibility of detail, and sometimes also simulate aspects of the human visual system.

The methods can be broadly divided into two categories. There are global methods (such as [14, 15, 28]; see also [12, 39, 42] for more references), which apply the same mapping function to each pixel in the image, and local methods (e.g., [2, 17, 18, 22]), which may vary the mapping across the image. Global methods operate on the histogram of the image, often on a logarithmic intensity scale, while local methods operate directly on the image data.

The problems with classical scalings, motivating the creation of TROs, are as follows. If a simplistic linear scaling and quantization procedure is used, often a large portion of the light intensity data falls into the first few bins (e.g., [35, 39] note this,

providing examples; see also Fig. 23.2 below). Scaling the logarithm of the intensity linearly and quantizing the result (i.e. plotting on a logarithmic scale), on the other hand, eliminates contrast [28]. Finally, the technique of histogram equalization makes the density histogram constant, which not only compresses large contrasts, but also exaggerates contrast in sparsely populated parts of the histogram ([28]; see also our example in Fig. 23.2 below). Histogram equalization may serve better in preprocessing input for pattern recognition in machine vision, as suggested in [19].

In the paper [39], a large number of different local and global TROs are reviewed and tested, and a new one is proposed. The authors suggest that most TROs, regardless of whether they are global or local, can be approximated to a satisfying degree by simple, fast image processing operations; this observation is also made in [50]. In the paper [31], the observation is tested quantitatively by approximating a number of different operators via a simplified model and parameter fitting. The authors suggest their model as an approach for validation and comparison of TROs.

Regardless of whether a global or a local tone mapping method is used, several authors caution against gradient reversal, i.e. flips of the local gradient direction in the dynamics compressed LDR image when compared to the original HDR data. Care must be taken because gradient reversal may create dark halos around bright objects. It can be avoided by careful design of the TRO. See, e.g., [18, 28, 47].

The study [41] discusses both global and local methods for tone mapping, and notes the globally order-preserving property of some global TROs; mathematically, it is obvious that the required property is the monotonicity of the mapping function. Specifically for photographs, [41] notes that, due to the well-known perceptual illusion that may make the same intensity look different when surrounded by brighter or darker shades, local tone-mapping techniques are preferable.

For the purposes of HDR plotting, it is clear that a global method is more appropriate. This is because then (right until quantization into pixel values) there is a one-to-one mapping between the original data value and the dynamics compressed function value.[2] Hence, a global colour bar (or, e.g., a global vertical scale, if the y axis is compressed instead of pixel intensity) can be made; this would be impossible for a local method.

In the study [47], it is cautioned that even though by using global techniques one can easily avoid gradient reversals, they may cause reversals in gradient magnitudes. An originally small difference between two pixels falling into a densely populated part of the histogram may look larger in the resulting LDR image than an originally large difference between two pixels that fall into a sparsely populated part. This is because the TRO may attenuate more aggressively in the sparsely populated parts of the histogram, in order to make more room on the output intensity scale for the densely populated parts. This may happen even if the method prevents contrast expansion, since what matters here is the ratio of contrast attenuation factors

[2] Strictly speaking, in the case of data-adaptive histogram remappers, if some of the histogram bins are empty, there may be a flat region in the mapping function. In this case the mapping is not globally one-to-one. However, since such regions contain no samples in the data, for the existing data it is one-to-one.

in different parts of the histogram. Local methods can work around the problem by locally adapting the mapping; however, as was noted above, this is of little use for HDR plotting. In the present study, we have chosen to ignore this problem.

It should be noted that in HDR-to-LDR conversion of photographic images in particular, various kinds of perception models of the human visual system are commonly used. Perceptual models account for factors such as gradual loss of colour sensitivity at low intensities [27, 28], loss of visual acuity at low intensities [27, 28], veiling glare [27, 28, 32, 45], global adaptation of vision to the intensity at the foveal point (e.g., [47]), and time-dependent adaptation [14, 21, 37], [48, Refs. 15–21]. See also, e.g., [13, 20, 44]. In the present study, we will not consider perceptual modelling.

The review [12] presents a comprehensive and quickly readable overview and classification of different kinds of TROs up to the year 2002. In 2010, global TROs specifically were reviewed in [42]. A psychophysical study comparing user preferences for different TROs in the photographic context was carried out in [49]. The 2007 paper [40] reviews TROs, and raises an important point about the input and output domains of the operators. The authors argue that operators based on perceptual models (unlike ones based on engineering principles such as histogram remappers) require both forward (light intensity to luminance) and backward (luminance to light intensity) passes to produce valid results. In the present study, we will concentrate on histogram remappers only; hence a single pass is sufficient.

23.3 The Dynamic Range

For grayscale (scalar) data, the representable dynamic range of a given medium is defined as (e.g., [41])

$$B' = \frac{|d_{\max}|}{|d_{\min}|}, \tag{23.1}$$

where B' denotes the dynamic range, and $|d_{\max}|$ and $|d_{\min}|$ refer to the largest and smallest nonzero representable data values (in terms of absolute value), respectively.

It is convenient to use a logarithmic representation. On the decibel scale, we have the equivalent definition

$$B = 20 \cdot \log_{10} \frac{|d_{\max}|}{|d_{\min}|}, \tag{23.2}$$

where B is the dynamic range in decibels.

With the standard eight bits per colour channel, we have for a regular computer display $d_{\max} = 255$ and $d_{\min} = 1$, which leads to a dynamic range of $B' = 255$, or equivalently, $B \approx 48.1$ dB.

Below, when we speak of the dynamic range of a set of scalar data, we also refer to (23.2). In this case, we take $|d_{\max}|$ and $|d_{\min}|$ as the maximal and minimal nonzero data values (in terms of absolute value), respectively.

It is possible to extend the representable range by using a trick mentioned in, e.g., [35]. Because the colour channels are independent, one can define a colour palette utilizing the range of all three channels independently. However, this comes with the cost of nontrivial interpretation, and losing grayscale representability. Because grayscale print is still an important medium for scientific publishing, techniques which allow for easy grayscale conversion are preferable in the context of the present study.

Another classical engineering trick, which is not applicable to print media, is the use of pulse width modulation (PWM) to represent fractional pixel values [48]. The physical dynamic range of the display stays the same, but because fractional values are represented, less information is lost in quantization. The dynamic range extension comes from making the effective smallest representable pixel value $d_{\min}^{\text{eff}} \in (0, 1)$. When the refresh rate of the display device is high enough, the illusion can be convincing; otherwise the picture may flicker.

23.4 Tone Mapping Methods Used in the Present Study

In this study, we chose two tone reproduction operators from literature, and tested one of our own (specifically for HDR plotting). Minor changes (documented in this section) were made to the operators from literature, in order to adapt them into the HDR plotting context. For the full technical details, a GNU Octave compatible MATLAB implementation is available.[3]

All three methods are order-preserving; gradient reversal cannot occur. In all the methods, the data is first histogrammed, taking the logarithm and then binning linearly. This creates a logarithmic binning of the data. The methods then operate on the obtained (discrete) logarithmic density histogram.

Due to the logarithmic processing, positive-valued data, such as light intensity or probability density, is the easiest to handle. In practice, zeroes require some extra care. If the data to be displayed contains also negative values (such as PCM audio waveforms), there are two options. Taking the absolute value of the data, and handling it all at once, produces a scale that is symmetric with respect to the origin. Another option is to handle the positive and negative parts of the data separately, producing an independent scale for each. For the sake of simplicity, in this text we concentrate on positive data only.

All three methods are histogram remappers. The resulting remapped histogram is used the same way as in histogram equalization. For a quick review, let $p(x) : (0, +\infty) \mapsto [0, +\infty)$ be a density histogram (hence, piecewise constant) and $C(x) = \int_0^x p(\xi)\,\mathrm{d}\xi : (0, +\infty) \mapsto [0, 1]$ the corresponding cumulative histogram. The normalization is $\lim_{x\to+\infty} C(x) = 1$; in fact, this maximum will be reached at the end of the last histogram bin that contains a nonzero value.

[3] Files `hdr*.m` in https://yousource.it.jyu.fi/savu/codes/. Link cited 13 Jan 2012.

Given such a function $p(x)$, histogram equalization works by mapping $H(D) := C(D)$. In the mapped data, each element (pixel) belongs to the interval $[0, 1]$. This can then be scaled linearly to the range appropriate for the display device (and quantized to available pixel intensities). If the given function $p(x)$ is the actual density histogram of D, then the cumulative histogram of $H(D)$ will be a straight line from 0 to 1; the data will be histogram-equalized. The idea of the general histogram remapper is to first modify $p(x)$ in some appropriate way, before computing $H(D)$. See Figs. 23.4–23.5 below for an illustration of some mapping curves in a test example.

The first chosen method (below Method A) was based on [28]. The perceptual modelling was not included; only the contrast expansion limiting histogram remapping algorithm was used. To this, no modifications were needed. In this method, the growth in display intensity is limited from above to at most the growth in world intensity; i.e., contrast in the image is never expanded.[4] Contrast expansion is prevented by capping and renormalizing the histogram. This acts as a slope limiter for the cumulative histogram. By [28], the maximum allowed data value in one bin is

$$\text{ceiling} = T \frac{\Delta B}{B_{\text{display}}}, \tag{23.3}$$

where T is the current sum of the histogram data (values from all bins summed together), ΔB is the width of one histogram bin in decibels, and B_{display} is the dynamic range of the display device, as per (23.2). The quantity ΔB can be computed as $\Delta B = B_{\text{data}}/N$, where N is the number of histogram bins used and B_{data} is computed by (23.2). The bins are looped over, and any bin which has a higher data value than (23.3) is clipped to the ceiling. The excess is summed to a total removed. The whole histogram is then iterated over several times. Once the total removed during one iteration of the outer loop falls below a prescribed tolerance (in the present study, we chose 0.5 % of the largest value in the original, unmodified histogram data), the algorithm terminates. The updated histogram is then normalized so that it becomes a density histogram. For details, see the original paper.

The second method (Method B) was based on [18]. The original method is a local one, working directly on the 2D image data. A corresponding global method, applicable for HDR plotting, was created by simply applying the method in one dimension to the logarithmically binned histogram. This method is a gradient attenuator. First, a Gaussian pyramid is constructed, and finite differences are used at each level of the pyramid for gradient approximation. Then the gradient field at each level is updated adaptively, compressing large values of the gradient while (optionally) amplifying small ones. Finally, an auxiliary Poisson problem is solved in order to reconstruct scalar potentials that match the updated gradient fields as well as possible, and the solutions are assembled to form the output. The authors use a full multigrid (FMG) solver to obtain the solution in linear time. Our case is simpler;

[4] Similar ideas were explored ten years earlier in a medical imaging context by [38]; however, we use [28] since it explicitly provides an algorithm.

in one dimension, the Poisson problem can be skipped. The function corresponding to a given derivative field can obviously be found directly by numerical integration. The constant of integration and final scaling are then fixed by requiring that the resulting function is a density histogram.

The third method (Method C), for comparison, was based on a simple observation. The aim of the other two methods is to remove the highest peaks in the histogram, while retaining the overall shape approximately. Thus, it should be possible to retain the overall location of the most massive peaks, while smoothing out the histogram, by applying linear diffusion to the logarithmic histogram. The one-dimensional time-dependent heat equation was set up, with zero right-hand side, zero Neumann boundary conditions at both ends, and the original histogram as the initial condition. The diffusion simulation was run until the highest peak in the histogram (i.e. the L_∞ norm of the solution) fell under the ceiling (23.3).

It should be noted that Methods A and C will fail if $B_{\text{data}} \leq B_{\text{display}}$ (i.e. if the data is LDR); this is because (23.3) will then produce a ceiling that cannot be satisfied. This case must be detected at the outset before applying the methods.

For LDR data, Method C simplifies to logarithmic scaling. To see this, let $t \to +\infty$ in the linear diffusion simulation; the end result is a constant function in logarithmic histogram space. Thus, if $B_{\text{data}} \leq B_{\text{display}}$, we can skip the tone mapping and use logarithmic scaling instead. When this occurs, the display device has enough dynamic range to display the data without losing contrasts.

23.5 Results

The results from the tone mapping algorithms are shown in Fig. 23.1. After preclipping out anything smaller than 10^{-3}, the data has a dynamic range of $B \approx 94.8$ dB ($B' \approx 54750$). Recall that a display with 8-bit colour channels, without tone mapping, can display $B \approx 48.1$ dB ($B' = 255$). How the displayed HDR data was produced will be explained in the next section; for now, the item of interest is the relative performance of the different methods. The colour scale runs from white at zero to the data maximum at black. Note that the scale is neither linear nor logarithmic; as was discussed earlier, tone mapping methods produce a data-adaptive scaling.

We have also provided, as a baseline for comparison, corresponding plots using six naïve scalings in Figs. 23.2–23.3. These methods are, respectively, linear scaling (full range), linear scaling with 0.1 % of the top end (in terms of histogram mass) clipped out, logarithmic scaling (full range), logarithmic scaling showing the highest 48 dB only, the 10th root of the data,[5] and classical histogram equalization.

The updated histograms and corresponding cumulative histograms are provided in Figs. 23.4–23.5. The histograms are given for the histogram remapping methods only, plotted on a log-log scale (with zeroes deleted from data). However, since

[5] According to [18], root-taking is a popular naïve approach.

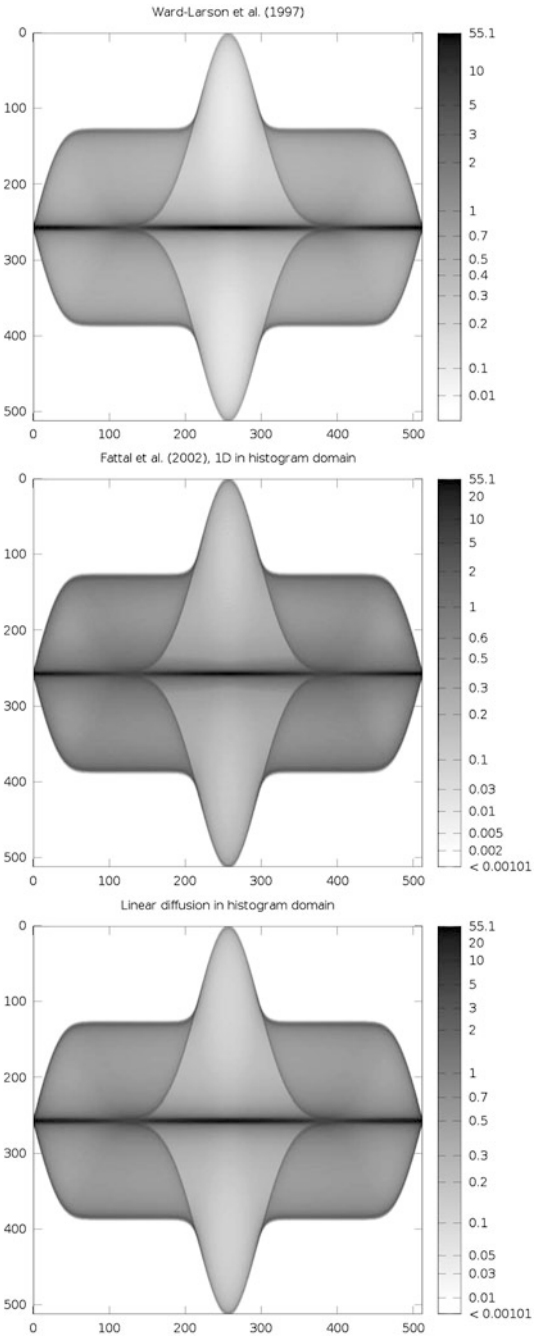

Fig. 23.1 HDR plotting, using tone mapping methods. Top to bottom: Method A [28], Method B [18], and Method C (linear diffusion). Note data-adaptive (non-linear, non-logarithmic) colour scale

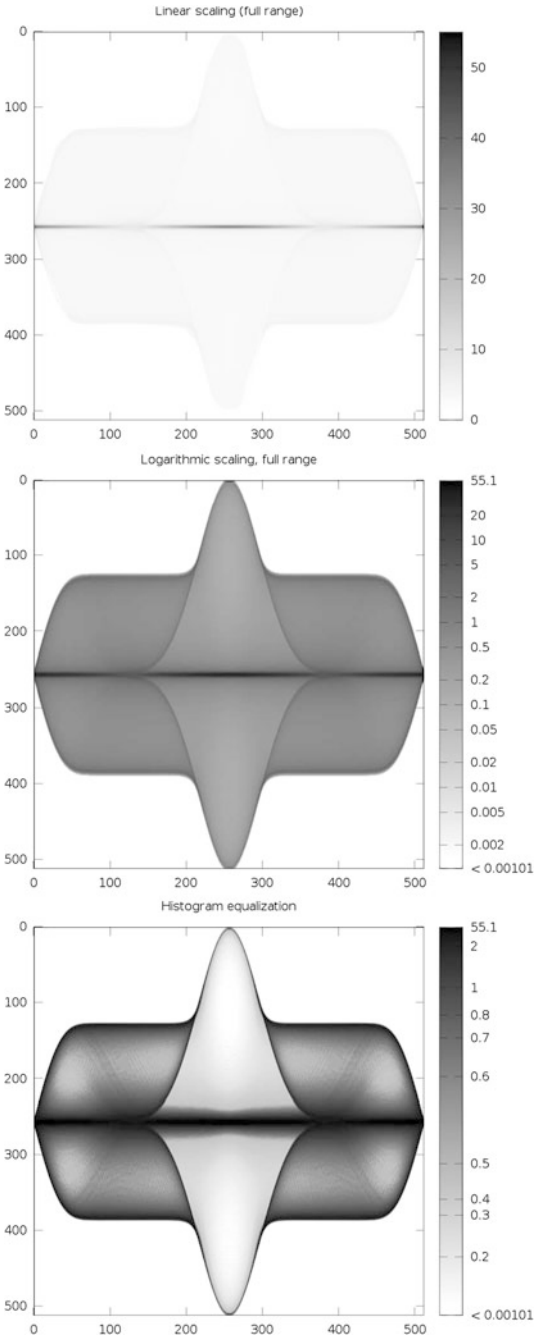

Fig. 23.2 HDR plotting, baseline (naïve) methods (set 1 of 2). Top to bottom: linear (*full scale*), logarithmic (*full scale*), histogram equalization

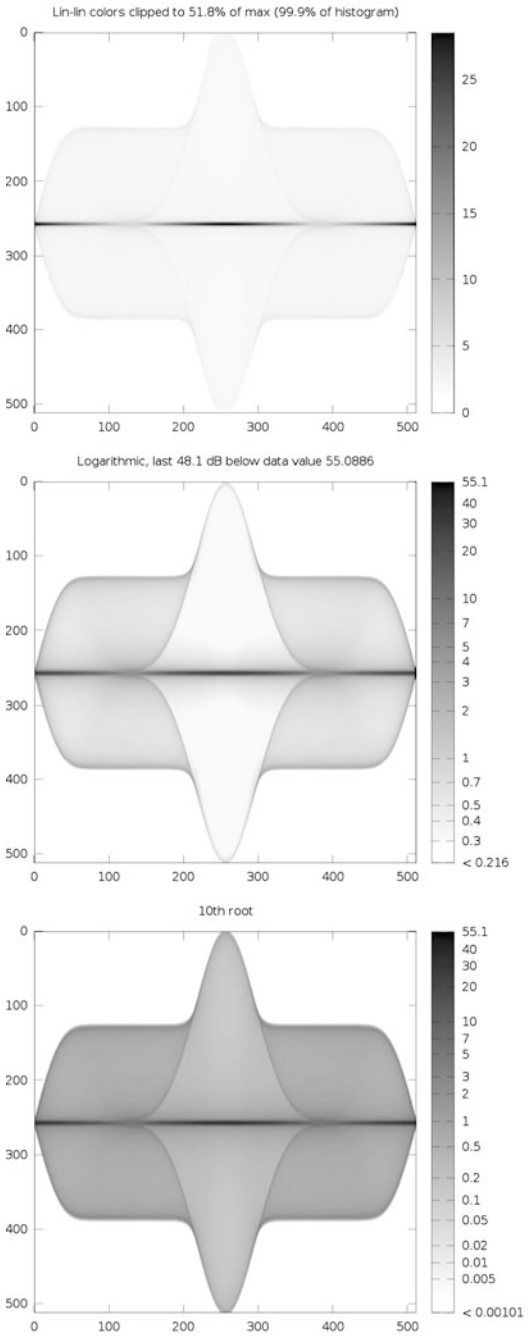

Fig. 23.3 HDR plotting, baseline (naïve) methods (set 2 of 2). Top to bottom: linear (showing only bottom 99.9 % of data), logarithmic (showing top 48 dB only), 10th root of data (showing all data)

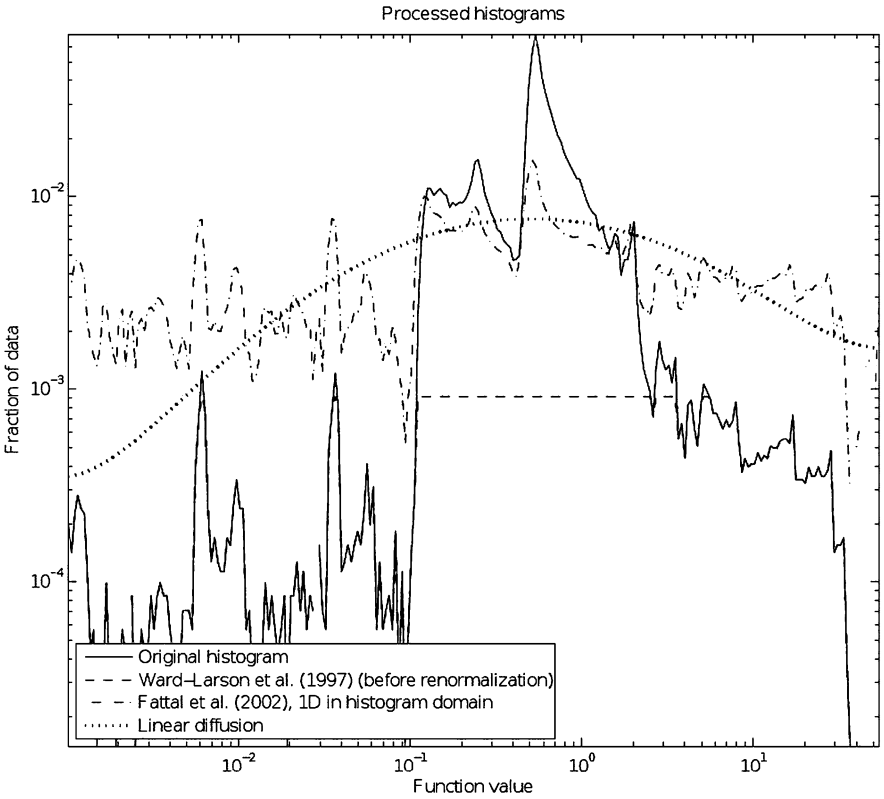

Fig. 23.4 Processed histograms of the methods used, corresponding to the images in Fig. 23.1. Note the log – log scale

the cumulative histogram is effectively the tone mapping function (as explained above), these are provided for the baseline methods, too. Here the scaling is log-lin; logarithmic in input data value (x axis) and linear in remapped output (y axis). Note that in the cumulative plot, the slope of the Ward-Larson mapping (Method A) never exceeds that of the logarithmic mapping of the last 48 dB; this is because the method prevents contrast expansion.

Method A gives excellent results; the detail in the HDR function can be seen very clearly. The results from Method B are similar, but consistently darker than those from Method A. It seems that Method B allocates a relatively larger portion of the available bins to the low end (small data values), leaving less bins for the high end (large data values). With the chosen colour scale, this makes the image darker. Also Method C produces acceptable results. Surprisingly, the differences in the results between the three different methods are relatively minor.

Since Method A gives, subjectively, slightly better results than the other two, and is also very simple to implement, our recommendation is to use Method A.

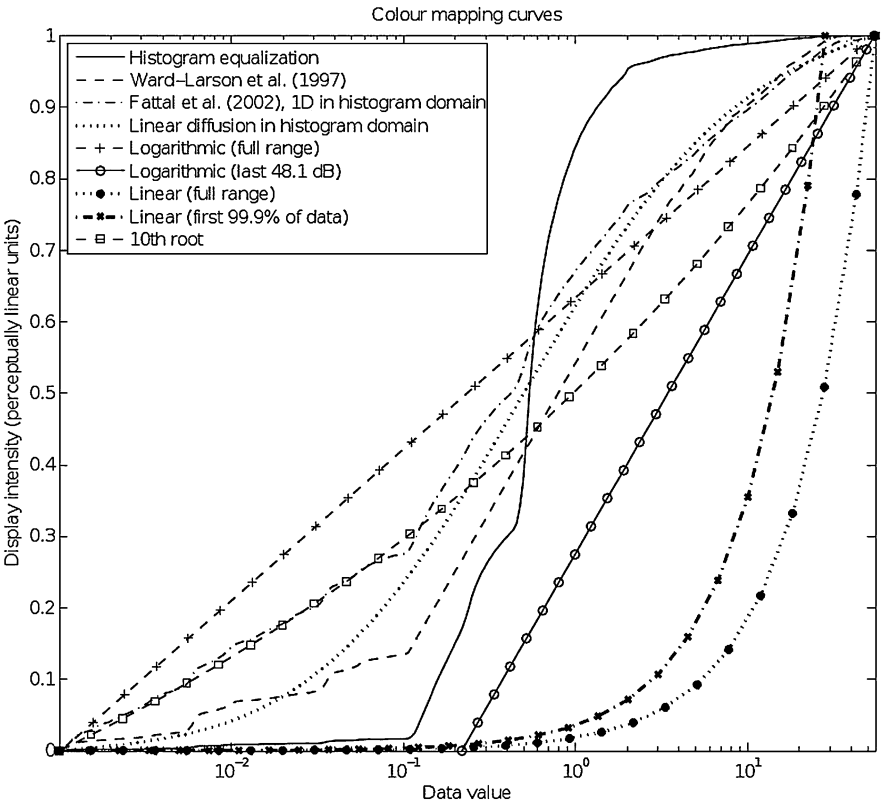

Fig. 23.5 Cumulative histograms (tone mapping functions) of the methods used, corresponding to the images in Figs. 23.1–23.3. Note the log–lin scale

It should be noted that all three methods can easily compress a dynamic range of over $B = 100$ dB ($B' = 10^5$) onto a regular display. This is not surprising, since the first two methods have been designed for use with HDR photography, where this is a requirement.

The methods were tested up to approximately $B = 300$ dB ($B' = 10^{15}$). While the end result in that case does not look as nice, it still gives some idea of the overall structure, see Fig. 23.6. (This is the same data as above, this time with the pre-clipping disabled. The extremely small numbers are ringing artifacts of the Fourier-based density estimator that produced the data.)

23.6 Application to Probability Densities in Mechanical Problems

As illustration of the use of HDR plotting in mechanics, we consider two vibration problems with uncertain data: transverse waves in a classical vibrating string, and

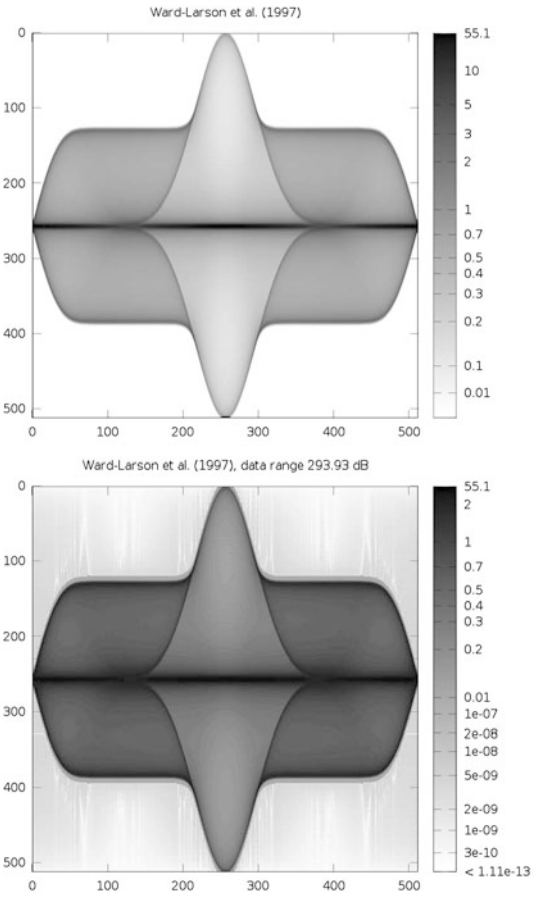

Fig. 23.6 A high dynamic range function tone mapped using Method A. *Top*: Pre-cutoff at 10^{-3}; dynamic range $B \approx 94.8$ dB. *Bottom*: No pre-cutoff; dynamic range $B \approx 294.0$ dB

the dynamic out-of-plane behaviour of an axially travelling panel submerged in axial potential flow.

For these problems, the behaviour of the displacement itself is of course not very interesting, because for a non-damped oscillator, any small disturbance in the frequency will cause two initially close trajectories to eventually diverge. Basically all that is needed for analysis are the eigenfrequencies, which behave in a stable manner. However, for the purposes of illustrating HDR plotting, we will consider displacements only, because this generates HDR data that is ideal for demonstration of the methods. From the modelling viewpoint, in the plotted time evolutions, we can visually observe the divergence of the solution set.

First, consider travelling transverse waves in a finite ideal string, with the endpoints fixed to zero level. As is well known, this situation is described by the one-dimensional wave equation $w_{tt} + cw_{xx} = 0$ for the displacement $w \equiv w(x,t)$. The boundary conditions are zero Dirichlet, $w(-1,t) = w(1,t) = 0$. As initial conditions, we choose an initial shape $w(x,0) \equiv w_0(x)$ (fulfilling the boundary conditions), and zero transverse velocity, $w_t(x,0) \equiv 0$. An analytical solution for this

case can be easily constructed by repeated reflection of the free-space d'Alembert solution; hence, we only need to plot a known function.

Now, let the wave propagation speed c be uncertain. We take $c = c_0 + X$, where X is a random variable with a zero-mean Gaussian distribution truncated at $\pm 3\sigma$. The parameter $c_0 = 0.05$ and σ is taken as 1 % of c_0. We then use SAVU to compute and plot the resulting probability density. See Figs. 23.7–23.8 for a time evolution simulation. (The last frame of this simulation is the example used in the previous section.)

The second problem comes from our research. We consider an axially travelling panel submerged in ideal fluid (potential flow), with an optional axial free-stream component. Details can be found in [3, 4, 26]. *Panel* is understood as a plate in the limit of cylindrical deformation (the *flat panel* of aeroelasticity; see, e.g., [6]).

We take as the uncertain parameters the axial panel velocity V_0, applied axial tension T, the Young modulus of the panel E, and the mass per unit area of the panel m. A Gaussian distribution truncated at $\pm 3\sigma$ is used for each parameter, representing a typical measurement error. The input is thus a four-dimensional hypercube. We choose σ (arbitrarily) as 1 % of the reference value for each parameter. The governing equation and reference values for the parameters can be found in [26, p. 100 and 155]. The initial conditions for w and w_t are taken as zero, and an external disturbance (force) is applied for a finite time at the beginning.

In Fig. 23.9 we have snapshots of the time behaviour of the displacement, with the four simultaneously uncertain parameters.

23.7 Conclusion

In this study, we presented and tested three methods for data-adaptive dynamic range compression of high dynamic range (HDR) mathematical functions, achieving a visual contrast preserving representation on low dynamic range media such as regular computer displays and print. This produced a scaling that is neither linear nor logarithmic, but instead data-adaptive. It is also global across the picture, allowing a colour bar to be created in the usual manner. One of the methods was seen to perform slightly better than the other two. As it was also the simplest to implement, it was recommended.

High dynamic range functions occur, for example, as probability densities in some uncertain data problems in mechanics. When such data is encountered, in our opinion it is natural to look for a visualization that shows the structure clearly. Compared to logarithmic scaling, the described methods have the advantage of visual contrast preservation, making the structure of the function more clearly visible.

Finally, it should be especially emphasized that what the presented methods do is dynamic range compression, and dynamic range compression only. For data that already fits into the dynamic range of the display (LDR data), classical methods are sufficient. In that case, these methods will either do nothing or, in the worst case, possibly harm the visualization. But if the data to be visualized is HDR, then these methods are very useful and can significantly improve the visualization.

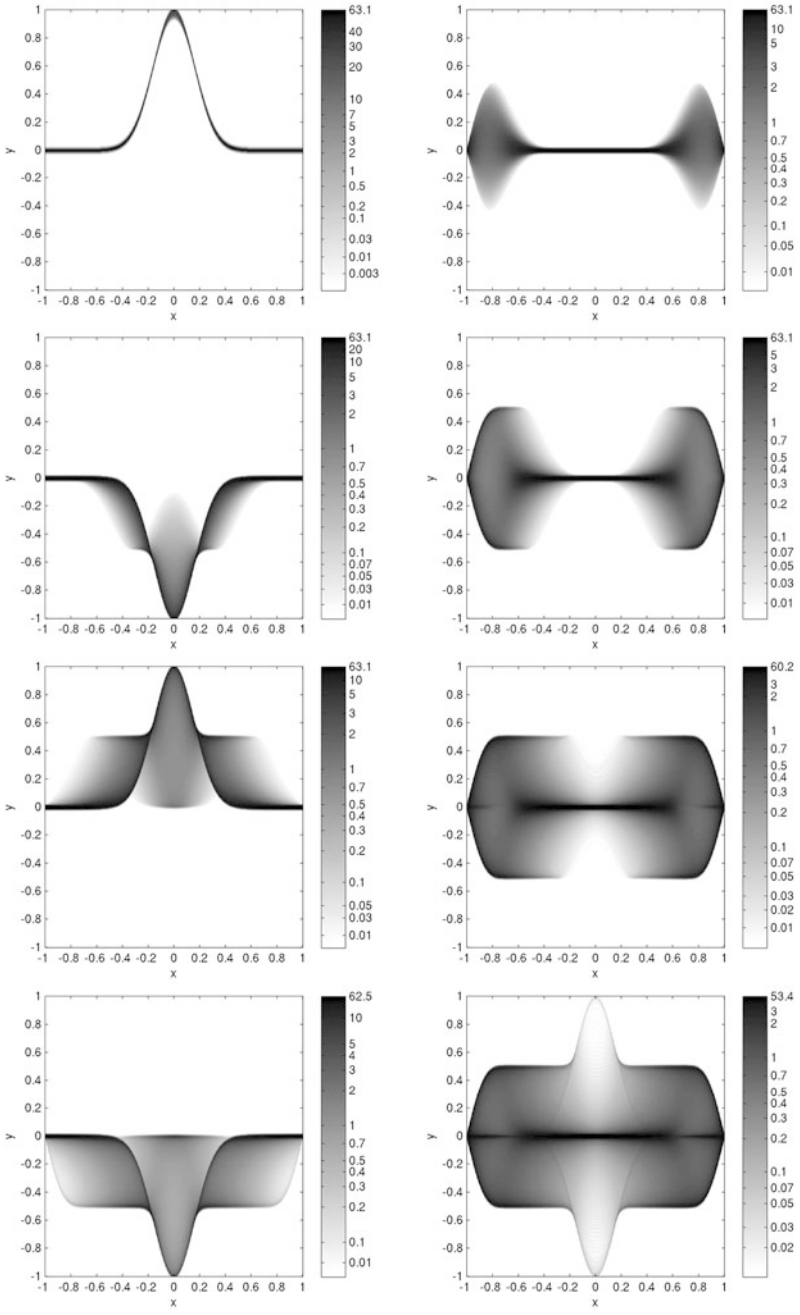

Fig. 23.7 Probability density of travelling transverse waves in a string with uncertain wave propagation speed. Snapshots taken at regular intervals; time increases as top left, top right, middle left and so on. *Top left*: initial pulse

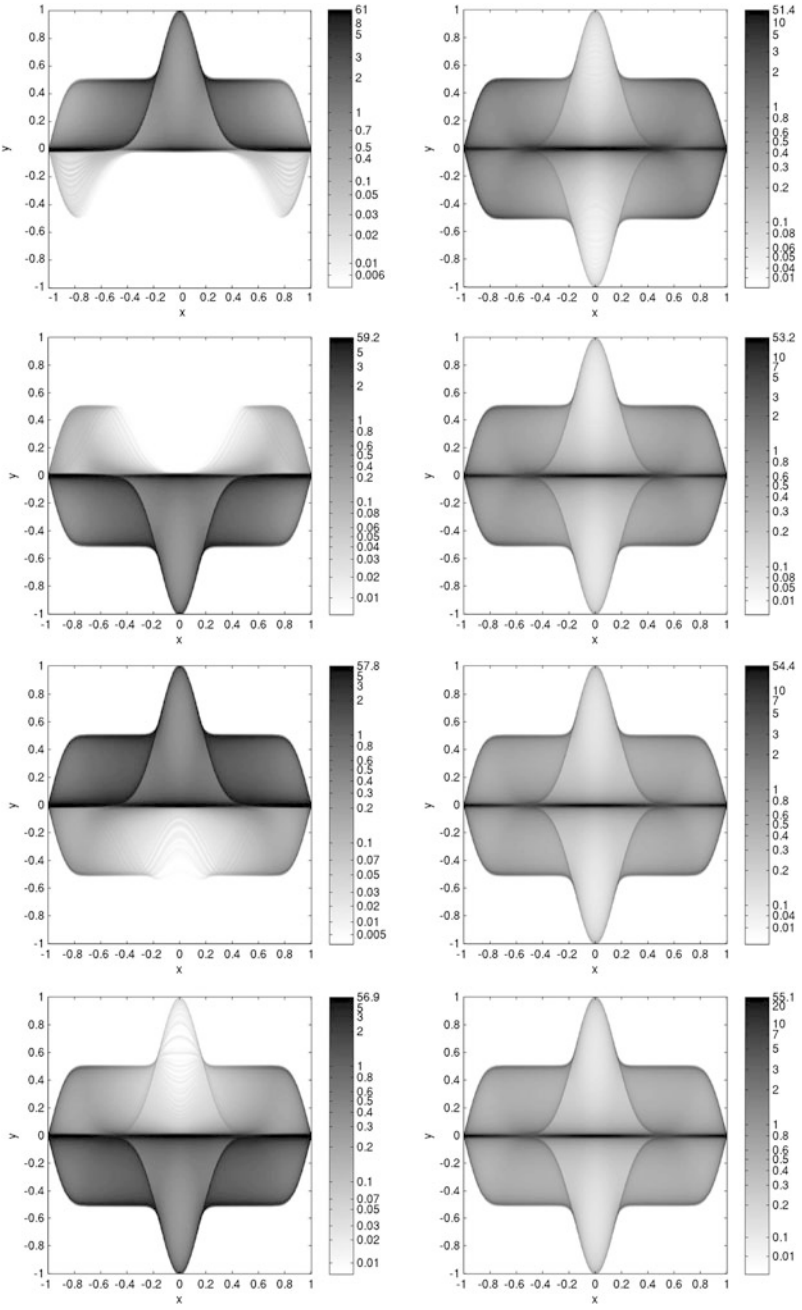

Fig. 23.8 Probability density of travelling transverse waves in a string with uncertain wave propagation speed. Snapshots taken at regular intervals; continued from Fig. 23.7

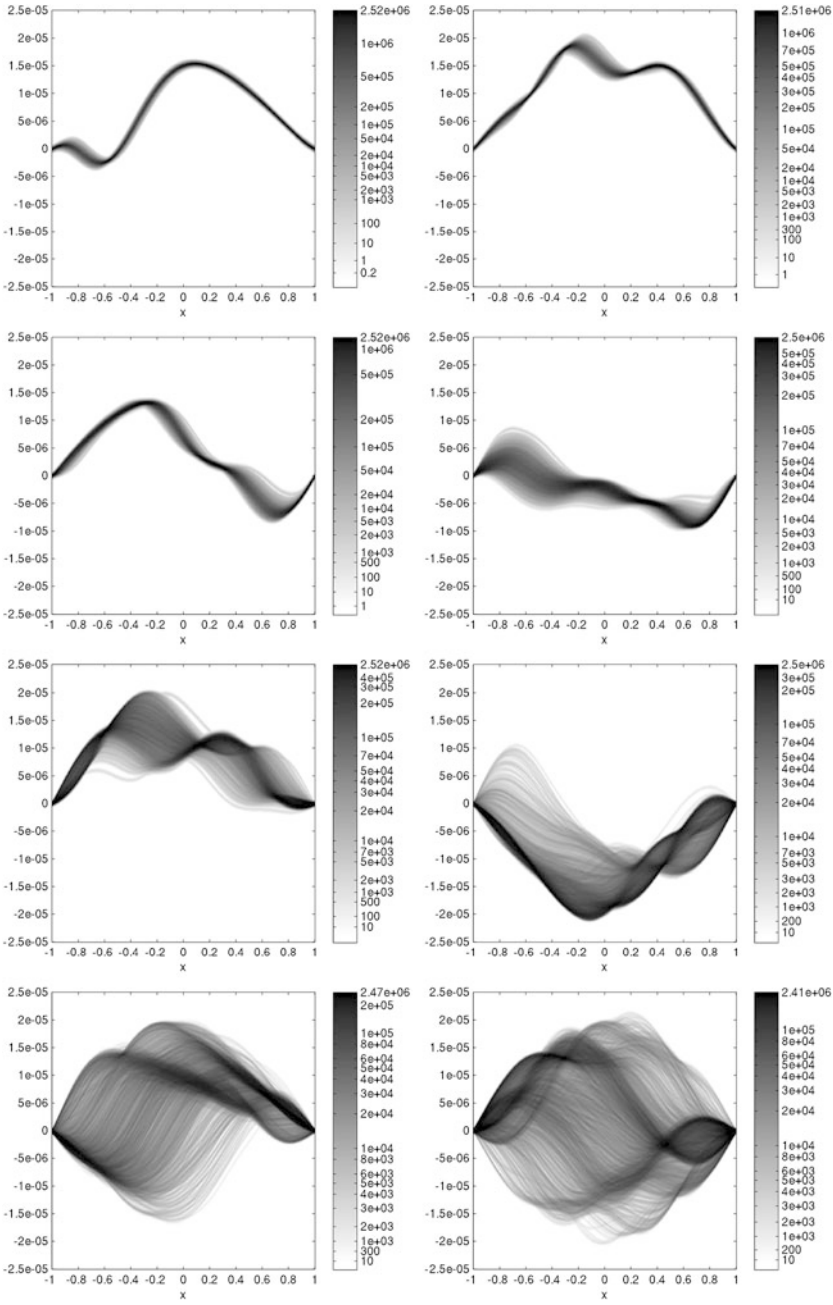

Fig. 23.9 Probability density of transverse displacement of an axially travelling panel submerged in ideal fluid, with four simultaneously uncertain parameters. After cutting the low end at 10^{-1}, the dynamic range is approximately $B \approx 147.6$ dB ($B' \approx 2.4 \cdot 10^7$). Snapshots taken at different times. Left to right by row: $t = 5,10$; 15,20; 50,100; 150 and 200 s

Hence, the presented methods are not intended to replace classical scalings, but to provide an extension where needed. By utilizing dynamic range compression, it is possible to extend the range of mathematical functions that can be represented on regular computer displays and print in a visual contrast preserving manner.

References

1. Abanoz B, Wang M (2008) A review of high dynamic range imaging on static scenes. Technical report 2008-04, Boston University
2. Ashikhmin M (2002) A tone mapping algorithm for high contrast images. In: Proceedings of the 13th eurographics workshop on rendering, pp 145–156. The Eurographics Association
3. Banichuk N, Jeronen J, Neittaanmäki P, Tuovinen T (2010) Static instability analysis for travelling membranes and plates interacting with axially moving ideal fluid. J Fluids Struct 26(2):274–291. doi:10.1016/j.jfluidstructs.2009.09.006
4. Banichuk N, Jeronen J, Neittaanmäki P, Tuovinen T (2011) Dynamic behaviour of an axially moving plate undergoing small cylindrical deformation submerged in axially flowing ideal fluid. J Fluids Struct 27(7):986–1005. doi:10.1016/j.jfluidstructs.2011.07.004
5. Beachkofski BK, Grandhi R (2002) Improved distributed hypercube sampling. In: Proceedings of the 43rd conference AIAA/ASME/ASCE/AHS/ASC on structures, dynamics and materials. Paper AIAA-2002-1274
6. Bisplinghoff RL, Ashley H (1975) Principles of aeroelasticity, 2nd edn. Dover, New York
7. Botev ZI, Grotowski JF, Kroese DP (2010) Kernel density estimation via diffusion. Ann Stat 38(5):2916–2957
8. Chacón JE, Duong T (2010) Multivariate plug-in bandwidth selection with unconstrained pilot bandwidth matrices. Test 19(2):375–398
9. Cox SE, Booth DT (2009) Shadow attenuation with high dynamic range images. Environ Monit Assess 158(1–4):231–241
10. Debevec P (1998) Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: Proceedings of the 25th annual conference on computer graphics and interactive techniques (SIGGRAPH'98). ACM, New York, pp 189–198
11. Debevec P, Malik J (1997) Recovering high dynamic range radiance maps from photographs. In: ACM SIGGRAPH 2008 classes (SIGGRAPH'08). ACM, New York. Article No. 31,
12. Devlin K (2002) A review of tone reproduction techniques. Technical report CSTR-02-005, University of Bristol
13. Drago F, Martens W, Myszkowski K, Seidel H-P (2002) Perceptual evaluation of tone mapping operators with regard to similarity and preference. Research report MPI-I-2002-4-002, Max-Planck-Institut für Informatik
14. Drago F, Myszkowski K, Annen T, Chiba N (2003) Adaptive logarithmic mapping for displaying high contrast scenes. Comput Graph Forum 22(3):419–426. doi:10.1111/1467-8659.00689
15. Duan J, Bressan M, Dance C, Qiu G (2010) Tone-mapping high dynamic range images by novel histogram adjustment. Pattern Recognit 43(5):1847–1862. doi:10.1016/j.patcog.2009.12.006
16. Duong T, Hazelton ML (2003) Plug-in bandwidth matrices for bivariate kernel density estimation. J Nonparametr Stat 15(1):17–30
17. Durand F, Dorsey J (2002) Fast bilateral filtering for the display of high-dynamic-range images. In: Proceedings of the 29th annual conference on computer graphics and interactive techniques (SIGGRAPH'02). ACM, New York, pp 257–266

18. Fattal R, Lischinski D, Werman M (2002) Gradient domain high dynamic range compression. In: Proceedings of the 29th annual conference on computer graphics and interactive techniques (SIGGRAPH'02). ACM, New York, pp 249–256
19. Finlayson G, Hordley S, Schaefer G, Tian GY (2005) Illuminant and device invariant colour using histogram equalisation. Pattern Recognit 38(2):179–190. doi:10.1016/j.patcog.2004.04.010
20. Gatta C, Rizzi A, Marini D (2007) Perceptually inspired HDR images tone mapping with color correction. Int J Imaging Syst Technol 17(5):285–294
21. Goodnight N, Wang R, Woolley C, Humphreys G (2003) Interactive time-dependent tone mapping using programmable graphics hardware. In: Proceedings of the 14th eurographics symposium on rendering (EGSR'03), pp 26–37 Eurographics Association
22. Goshtasby AA High dynamic range reduction via maximization of image information, 2003. http://www.cs.wright.edu/people/faculty/agoshtas/goshtasby_hdr.pdf. CiteSeerX: doi:10.1.1.121.421
23. Helton JC, Davis FJ (2002) Illustration of sampling-based methods for uncertainty and sensitivity analysis. Risk Anal 22(3):591–622
24. Helton JC, Davis FJ (2002) Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. Sandia report SAND2001-0417, Sandia National Laboratories
25. Helton JC, Johnson JD, Sallaberry CJ, Storlie CB (2006) Survey of sampling-based methods for uncertainty and sensitivity analysis. Sandia report SAND2006-2901, Sandia National Laboratories
26. Jeronen J (2011) On the mechanical stability and out-of-plane dynamics of a travelling panel submerged in axially flowing ideal fluid: a study into paper production in mathematical terms. PhD thesis, University of Jyväskylä, Jyväskylä. http://julkaisut.jyu.fi/?id=978-951-39-4596-1
27. Krawczyk G, Myszkowski K, Seidel H-P (2005) Perceptual effects in real-time tone mapping. In: Proceedings of the 21st spring conference on computer graphics (SCCG'05). ACM, New York, pp 195–202
28. Larson GW, Rushmeier H, Piatko C (1997) Visibility matching tone reproduction operator for high dynamic range scenes. IEEE Trans Vis Comput Graph 3(4):291–306
29. Mann S, Picard RW (1995) On being 'undigital' with digital cameras: extending dynamic range by combining differently exposed pictures. In: Proceedings the 48th annual conference of IS&T, pp 442–448
30. Mantiuk R, Daly S, Kerofsky L (2008) Display adaptive tone mapping. ACM Trans Graph (TOG) – Proc ACM SIGGRAPH 2008, 27(3). doi:10.1145/1360612.1360667
31. Mantiuk R, Seidel H-P (2008) Modeling a generic tone-mapping operator. Comput Graph Forum 27(2):699–708. doi:10.1111/j.1467-8659.2008.01168.x
32. McCann JJ, Rizzi A (2007) Veiling glare: the dynamic range limit of HDR images. In: Human vision and electronic imaging XII. Proceedings of electronic imaging science and technology, vol 6492. IS&T and SPIE
33. McKay MD, Beckman RJ, Conover WJ (1979) A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 21(2):239–245
34. Owen Art B (1992) Orthogonal arrays for computer experiments, integration and visualization. Stat Sin 2(2):439–452
35. Park SH, Montag ED (2007) Evaluating tone mapping algorithms for rendering non-pictorial (scientific) high-dynamic-range images. J Vis Commun Image Represent 18(5):415–428
36. Parks DR, Roederer M, Moore WA (2006) A new "logicle" display method avoids deceptive effects of logarithmic scaling for low signals and compensated data. Cytometry A 69(6):541–551
37. Pattanaik SN, Tumblin J, Yee H, Greenberg DP (2000) Time-dependent visual adaptation for fast realistic image display. In: Proceedings of the 27th annual conference on computer graphics and interactive techniques (SIGGRAPH'00). ACM, New York, pp 47–54
38. Pizer SM, Amburn EP, Austin JD, Cromartie R, Geselowitz A, Greer T, Romeny BTH, Zimmerman JB, Zuiderveld K (1987) Adaptive histogram equalization and its variations. Comput

Vis Graph Image Process 39(3):355–368
39. Reinhard E, Devlin K (2005) Dynamic range reduction inspired by photoreceptor physiology. IEEE Trans Vis Comput Graph 11(1):13–24
40. Reinhard E, Kunkel T, Marion Y, Brouillat J, Cozot R, Bouatouch K (2007) Image display algorithms for high and low dynamic range display devices. J Soc Inf Disp 15(12):997–1014
41. Schlick C (1994) Quantization techniques for visualization of high dynamic range pictures. In: Photorealistic rendering techniques. Proceedings of the 5th eurographics workshop on rendering. Springer, Berlin, pp 7–20
42. Shan Q, Jia J, Brown MS (2010) Globally optimized linear windowed tone mapping. IEEE Trans Vis Comput Graph 16(4):663–675
43. Sheather SJ, Jones MC (1991) A reliable data-based bandwidth selection method for kernel density estimation. J R Stat Soc B 53(3):683–690
44. Smith K, Krawczyk G, Myszkowski K, Seidel H-P (2006) Beyond tone mapping: enhanced depiction of tone mapped HDR images. Comput Graph Forum 25(3):427–438
45. Talvala E-V, Adams A, Horowitz M, Levoy M (2007) Veiling glare in high dynamic range imaging. ACM Trans Graph (TOG) – Proc ACM SIGGRAPH 2008, 26(3)
46. Tang B (1993) Orthogonal array-based Latin hypercubes. J Am Stat Assoc 88(424):1392–1397
47. Tumblin J, Hodgins JK, Guenter BK (1999) Two methods for display of high contrast images. ACM Trans Graph 18(1):56–94
48. Wu Y, Qiu B (2010) Perceptually fractural pixel values in rendering high dynamic range images. In: Proc SPIE, vol 7744. doi:10.1117/12.863019
49. Yoshida A, Blanz V, Myszkowski K, Seidel H-P (2005) Perceptual evaluation of tone mapping operators with real-world scenes. In: Human vision and electronic imaging X. Proceedings of the SPIE, vol 5666. SPIE, New York, pp 192–203
50. Yoshida A, Mantiuk R, Myszkowski K, Seidel H-P (2006) Analysis of reproducing real-world appearance on displays of varying dynamic range. Comput Graph Forum 25(3):415–426
51. Yuan X, Nguyen MX, Chen B, Porter DH (2005) High dynamic range volume visualization. In: Proceedings of the conference on visualization 2005. IEEE Comput Sci, Los Alamitos, pp 327–334