

High Dynamic Range Ultrasound Imaging

Alperen Degirmenci · Douglas P. Perrin · Robert D. Howe

Received: 26 January 2018

Abstract

Purpose High dynamic range (HDR) imaging is a popular computational photography technique that has found its way into every modern smartphone and camera. In HDR imaging, images acquired at different exposures are combined to increase the luminance range of the final image, thereby extending the limited dynamic range of the camera. Ultrasound imaging suffers from limited dynamic range as well; at higher power levels, the hyperechogenic tissue is overexposed, whereas at lower power levels, hypoechogenic tissue details are not visible. In this work, we apply HDR techniques to ultrasound imaging, where we combine ultrasound images acquired at different power levels to improve the level of detail visible in the final image. **Methods** Ultrasound images of *ex vivo* and *in vivo* tissue are acquired at different acoustic power levels and then combined to generate HDR ultrasound (HDR-US) images. The performance of five tone mapping operators is quantitatively evaluated using a similarity metric to determine the most suitable mapping for HDR-US imaging.

Results The *ex vivo* and *in vivo* results demonstrated that HDR-US imaging enables visualizing both hyper- and hypoechogenic tissue at once in a single image. The Durand tone mapping operator preserved the most amount of detail across the dynamic range.

Conclusions Our results strongly suggest that HDR-US imaging can improve the utility of ultrasound in image-based diagnosis and procedure guidance.

Keywords ultrasound imaging · high dynamic range · tone mapping · image enhancement · contrast enhancement

1 Introduction

Ultrasound imaging is an invaluable tool for medical diagnosis and intraoperative guidance. Images are acquired in real time at high frame rates, with no radiation exposure, and at much lower expense than CT and MRI. However, ultrasound imaging suffers from a limited dynamic range: at higher power levels, rigid surgical tools and hyperechogenic tissue in the imaging field get overexposed (Fig. 1c), while at lower power levels hypoechogenic tissue details are lost (Fig. 1b).

There is extensive literature on image processing methods for improving ultrasound image quality. Some popular techniques are speckle and noise reduction through spatial and frequency-domain filters [6], detail enhancement through spatial and frequency compounding [2, 3], and contrast enhancement using histogram equalization or tissue harmonic imaging. See [9] for a review of ultrasound image enhancement techniques. However, these methods do not address the issue of limited dynamic range.

Similar to ultrasound imaging, optical cameras have a limited dynamic range as well. Imaging shadows requires a long exposure, while imaging highlights requires a short one. High dynamic range (HDR) imaging is a method which extends the dynamic range of optical cameras by combining images taken at multiple exposures [1]. In the last few years, this method has become ubiquitous in smartphones and digital cameras.

This work was supported by the US National Institutes of Health under grant NIH 1R21EB018938, Toyota Motor North America Inc., and the NVIDIA Corporation Academic Hardware Grant Program.

A. Degirmenci, D.P. Perrin, and R.D. Howe
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA
E-mail: adegirmenci@seas.harvard.edu

D.P. Perrin
Department of Cardiovascular Surgery, Children's Hospital Boston, Boston, MA, USA




Fig. 1 (a) Experimental setup: A phased linear array ultrasound probe (Philips 15-6L) images a chicken thigh at 4 cm depth while an ablation catheter (ThermoCool SF, Biosense Webster, Diamond Bar, CA) is pressed against the tissue. (b) B-mode ultrasound image acquired at power level -18.0 dB. The catheter body and the skin surface exhibit strong reflection, showing a clear image of the contact between the tissue and the instrument, while the deeper tissue and the bone surface have poor detail. (c) Ultrasound image acquired at power level 0.0 dB. The bone surface and the hypodermis details are pronounced, whereas the instrument and skin surface details are lost due to overexposure as well as ringing from the instrument. Scale on the right shows the estimated drop in signal strength due to attenuation and reflection through the catheter and tissue layers

In this paper, we show how HDR imaging techniques can be applied to ultrasound images to improve their dynamic range, enabling the imaging of media over a wide range of reflection and attenuation characteristics. We also compare several tone mapping methods, which allow the HDR-US images to be displayed on regular (*i.e.*, low dynamic range) screens, while preserving the visibility of features from across the entire dynamic range. This leads to better visualization for clinicians and can improve the performance of subsequent image processing operations. Therefore the extended dynamic range of HDR ultrasound (HDR-US) can allow for better diagnosis and procedure guidance.

Extending the dynamic range is of interest to other medical imaging modalities as well. A similar treatment was applied to MR imaging in [8], but there was no investigation of tone mapping in this study.

2 Methods

2.1 Ultrasound Image Formation

We begin by considering the process that generates the ultrasound image. The intensity of a pixel in the B-mode ultrasound image is proportional to the pressure of the reflected sound signal. The pressure of a sound wave that is scattered by a target away from the transducer can be expressed as

$$p_s(r, t) = A_0 \frac{R e^{-\mu_a(d+r)}}{r} F\left(t - \frac{d+r}{c}\right) \quad (1)$$

where r is the radius from the target, t is time, A_0 is the amplitude, R is the reflection coefficient, μ_a is the amplitude attenuation factor, d is the distance between the target and the transducer, $F(\cdot)$ is the acoustic envelope, and c is the speed of sound [11]. Increasing the amplitude increases the pressure. The only constant intrinsic to the ultrasound transducer in this equation is the amplitude; the rest of the terms are specific to the media being imaged.

The relative amplitude loss of the returning signal due to attenuation and reflection losses can be calculated as

$$\frac{A_z}{A_0} = \sqrt{R_I} e^{-\mu_a z} \quad (2)$$

where R_I is the intensity reflectivity, and z is the total distance traveled by the signal. Assuming a normal incidence at the boundary between two media, the reflected power loss in dB can be calculated as

$$20 \log_{10} \frac{A_z}{A_0} = 20 \log_{10} \left| \frac{Z_2 - Z_1}{Z_2 + Z_1} \right| e^{-\mu_a z} \quad (3)$$

where Z_1 and Z_2 are the acoustic impedance of media 1 and 2. The amplitude attenuation factor is a function of the signal frequency f and media-specific values a and b , written as $\mu_a = \frac{a f^b}{20 \log_{10} e}$. Fig. 1c shows the estimated drop in signal strength due to attenuation and reflection through the catheter and tissue layers. The acoustic impedance and the attenuation coefficients were taken from [11].

2.2 High Dynamic Range Formulation

In photography, the intensity Z of a pixel in a photograph is a nonlinear function of the exposure X , expressed as $Z = f(X)$. Malik and Debevec [4] derived a method to recover the response curve of a camera from multiple exposures. Exposure can be defined as $X = E\Delta t$, where E is the sensor irradiance, and Δt is the exposure time. Assuming f is monotonic, we can write $f^{-1}(Z) = E\Delta t$. Taking the logarithm of both sides, we have $g(Z) = \log f^{-1}(Z) = \log E + \log \Delta t$. The response curve g and the irradiances E_i can be then recovered by minimizing the loss function

$$\mathcal{L} = \sum_i^N \sum_j^P [g(Z_{ij}) - \log E_i - \log \Delta t_j]^2 + \lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} g''(z)^2, \quad (4)$$

where N is the number of pixels in each image, P is the number of photographs, and λ is a regularization parameter that controls the smoothness of g [4]. The response curve is then used to map the pixel intensities in each low dynamic range (LDR) image to a relative radiance value, by rearranging the expression for $g(Z)$ given above to $\log E_i = g(Z_{ij}) - \log \Delta t_j$ [4]. In practice, radiance is computed using a weighted average of all exposures from the input images, which increases robustness and reduces noise. Once the response curve is computed, it can be used with any image taken with the system and does not need to be recomputed, given other imaging settings remain constant.

A similar response curve is recoverable for ultrasound probes. In Eq. 1, the only term that depends on the ultrasound transducer is the signal amplitude, A_0 . Grouping all other terms in U , which describes the scattering and attenuation properties of the medium, Eq. 1 can be written as $p_s = UA_0$. Assuming that the mapping f from the pressure of the scattered wave to the pixel intensity Z in the ultrasound image is monotonic, we can write $f^{-1}(Z) = p_s = UA_0$. This has the same form as $f^{-1}(Z) = E\Delta t$.

The reflection equation (Eq. 2) is also linear in the amplitude, but it depends on the frequency as well. Assuming that the ultrasound probe is held stationary, kept at the same imaging depth and frequency, and the target did not move during imaging, a change in source signal amplitude should result in a linear change in the received signal amplitude. This is similar to keeping the camera stationary with the same aperture and focus while changing the exposure duration.

Ultrasound machines generally express acoustic power in decibels (dB). This is not an absolute measure of power, but rather a logarithmic ratio between two amplitudes. We can convert from power to a relative amplitude A_r using

$$A_r = 10^{(Power/20dB)}. \quad (5)$$

This quantity can be used in the HDR calculations. Fig. 2 shows the response curve recovered using fifteen ultrasound images acquired at different power levels. The curve has a form that is very similar to that of digital cameras computed using natural photographs (see [4] for examples). The nonlinearity in this function is due to nonlinearities in the transducer behavior, as well as the log compression applied during B-mode image formation in the ultrasound machine.

The recovered response curve is then used to compute $\log U_i = g(Z_{ij}) - \log A_j$, which can be interpreted as mapping B-mode image pixel intensities to the relative scattering and attenuation characteristics of the tissue.

2.3 Tone Mapping – Display of HDR-US Images

Once the HDR ultrasound image is computed, there remains the problem of displaying the information to the user on a low dynamic range (LDR) screen. The human ocular system can only resolve up to 30 dB,




Fig. 2 Response function recovered from the fifteen ultrasound images acquired at different power settings

and conventional monitors are limited to two orders of magnitude of dynamic range. Mapping an HDR image to the LDR is called tone mapping, where the goal is to preserve details and color appearance from the HDR image, which would be lost in a simple linear compression of pixel intensities. There is extensive work on tone mapping for HDR images; a detailed treatment is presented in [12]. Note that this step is only necessary if the HDR-US images will be shown to the clinician for diagnosis and procedure guidance. If the images are instead going to be used in an image processing pipeline, the rich HDR information can be retained to improve performance.

Tone mapping operators (TMOs) are often distinguished into *global* and *local* operators. Global operators find a single compression function that changes the luminance of all pixels in the image [12]. In contrast, local operators compress the luminance of each pixel using a mapping that varies spatially based on the neighboring pixels [5]. Therefore, local TMOs preserve detail better, whereas global TMOs can lead to a loss of contrast. However, local TMOs can cause halo artifacts around sharp features due to the Gaussian filter that is commonly used to compute the local parameters. The use of edge-preserving smoothing operators, such as the bilateral filter, can help prevent halos.

The simplest tone mapping method is normalization, where the HDR pixel intensities are linearly scaled down to the LDR range. However, this operation leads to a loss of detail, especially in the darker regions. In this work, we focus on five common tone mapping operators: (1) contrast-limited adaptive histogram equalization (CLAHE) [10], (2) a global luminance compression TMO [13], (3) a local luminance compression TMO that computes the local parameters based on a difference of Gaussians (DoG) approach [13], (4) a local luminance compression TMO based on a bilateral filter [12], and lastly (5) a two-scale decomposition method using bilateral filtering [5].

In CLAHE, the image is divided into regions, and the histograms of these regions are redistributed to have a wider, more uniform distribution, effectively increasing the contrast.

The luminance compression approach suggested by Reinhard in [13] has several implementations. The global implementation computes the geometric average of each pixel's luminance to steer the compression [12]. The DoG approach computes a difference of Gaussians to determine the appropriate filter size to be used in determining the local compression parameter. The bilateral filter approach replaces the DoG with an edge-preserving smoothing operator, which is better suited to preserving sharp contrasts and can be faster to compute than the DoG.

The decomposition approach by Durand in [5] preserves sharp gradients by separating the HDR image into a base layer (low frequency components) and a detail layer (high frequency components), and only reducing the contrast of the base layer, which preserves details in the final LDR image.

2.4 Quality Metrics

Quantifying the improvement (or degradation) in image quality of a tone-mapped HDR image over the original images is not straightforward since there is no ground truth. Quality metrics, such as mean-squared error (MSE), apply a direct comparison between pixel intensities. However, comparing the tone-mapped HDR image to any one of the input images necessarily leads to a mismatch between intensities, resulting in lower scores.

Instead, the analysis technique we adopt in our work is to use a sliding window to compute the similarity between the tone-mapped HDR image, and each of the original LDR ultrasound images that are used to generate the HDR image. Using only a small patch instead of the entire image makes it possible to determine local correspondences between the HDR and LDR images. This analysis shows how much of the dynamic range detail is preserved when the HDR-US image is tone mapped for display on an LDR medium: if the tone-mapped HDR-US image only has high agreement with a small subset of the original LDR images (*i.e.*, uses a narrow power level range), this means that the tone mapping failed to preserve the HDR information. Through this analysis, we aimed to determine the TMO that maintains the most detail.

The similarity metric we used is the peak signal-to-noise ratio (PSNR), defined as

$$\text{PSNR}_j = 10 \log_{10} \left(255^2 / \text{MSE}_j \right), \quad (6)$$

in which

$$\text{MSE}_j = \frac{1}{N} \sum_i^N \left(Z_i^{\text{HDR}} - Z_i^{\text{LDR}_j} \right)^2, \quad (7)$$

where N is the number of pixels in the sliding window, j is the index of the reference image, Z_i^{HDR} is the intensity of pixel i in the tone-mapped HDR image, and $Z_i^{\text{LDR}_j}$ is the intensity of pixel i in the reference image j . Computing the $\arg \max$ of PSNR_j yields the index (and therefore the power setting) of the reference image that has the best correspondence with the tone-mapped HDR image. Computing this




Fig. 3 (a–d) Four of the fifteen ultrasound images acquired at various power levels, and their histograms. In low power settings, the catheter body and the skin surface have good contrast and detail; however, the deeper structures are missing. At the higher power settings, the bone surface and the deeper tissue have good detail, but the catheter and the skin are overexposed and have lost fine detail. **(e)** Tone-mapped high dynamic range ultrasound image. The deep tissue structures, as well as the catheter and skin details, are easily visible




Fig. 4 (a) Ultrasound image acquired at 0.0 dB. **(b)** Tone mapped HDR-US image. **(c)** Regions discussed in the text. Highlighted in green is the ablation catheter body, purple is the ringing from the catheter, red is the fat layer, cyan is the bone, and yellow is the deep tissue structure. **(d)** HDR-US image intensities before tone mapping, shown in false color to convey the entire dynamic range. This image contains over three orders of magnitude of dynamic range, from 0.09 to 140.

index at each sliding window location yields a contour map (Fig. 6), which shows a power level landscape of the best corresponding reference image. The TMO that results in an HDR image with the widest set of reference images will be chosen as the TMO that is most suitable for ultrasound imaging.

3 Experiments and Results

We conducted an *ex vivo* and an *in vivo* experiment to test our methods. In both experiments, the ultrasound images were acquired using a Philips SONOS 7500 ultrasound machine (Philips Healthcare, Andover, MA, USA) with a 128-element linear phased array probe (Philips 15-6L), which has an operating frequency range of 6–15 MHz. The TGC controls were set to their lowest value. The grayscale B-mode ultrasound images were transferred to a PC using an S-video frame grabber (USB3HDCAP, StarTech, London, Ontario, Canada).

3.1 *Ex Vivo*

In our *ex vivo* experiments, we placed a chicken thigh into a water tank and pressed an ablation catheter (ThermoCool SF, Biosense Webster, Diamond Bar, CA) against the tissue surface. The bottom of the tank was covered with an ultrasound absorbing mat to prevent ringing due to reflections from the tank surface. See Fig. 1a for a picture of the experimental setup.

Fifteen B-mode ultrasound images were acquired at 4 cm depth and 6 MHz frequency, with power levels linearly spaced between -26.7 to 0.0 dB (mechanical index of 0.1 to 1.0, there was no return below -26.7 dB). A subset of these images can be seen in Fig. 3. The power settings were dialed in manually using the ‘power’ control knob on the ultrasound machine. The response function for these images, calculated using Debevec and Malik’s method [4], can be seen in Fig. 2. The HDR-US image computed from the original fifteen LDR images is shown in Fig. 4d, which contains over three orders of magnitude of dynamic




Fig. 5 Output of five different tone mapping algorithms: (a) CLAHE, (b) global, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand's bilateral separation.




Fig. 6 Similarity contours for (a) CLAHE, (b) global, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand's bilateral separation. The blue rectangle in the top left corner indicates the size of the sliding window (40 × 40 pixels) used for computing the local PSNR.

range. The tone-mapped HDR image is shown in Fig. 4b and demonstrates the utility of HDR in combining the low- and high-power details in a single image. Fig. 4c shows regions of interest.

The HDR Toolbox [1] was used to perform the HDR and tone mapping operations in MATLAB (MathWorks, Natick, MA, USA). Fig. 5 shows the output of five different tone mapping algorithms: (a) contrast-limited adaptive histogram equalization (CLAHE), (b) global Reinhard, (c) local Reinhard with DoG, (d) local Reinhard with bilateral filter, and (e) Durand's bilateral separation. For all TMOs, we used the default parameters. The Durand and Reinhard TMOs automatically adjust their parameters based on the input images. A gamma correction value of 2.2 was also applied, except for the CLAHE-based TMO.

A qualitative assessment shows that the adaptive histogram equalization amplifies the noise in the deeper regions and the ringing around the catheter body. The DoG TMO results in saturated regions along the catheter and skin (Fig. 5c) that lead to a loss of detail. Of the remaining, the Reinhard bilateral TMO maintains finer texture along the catheter and the skin compared to the global TMO. However the Durand TMO has the best ringing suppression at the catheter tip and sharpest detail along the skin overall.

In our quantitative analysis, we calculated the PSNR for the output of each TMO, which is a measure of the similarity between the reference image and the tone-mapped HDR-US image, as described in Section 2.4. We iterated over each of the fifteen ultrasound images as the reference input in the PSNR calculations. A sliding window size of 40 × 40 pixels was chosen heuristically. For each patch, we calculated its $\arg \max_j \text{PSNR}_j$ to find the reference image power level that corresponded the most to the tone-mapped HDR-US image. We display this information in the form of a contour plot in Fig. 6. The colors in the contour plot correspond to the power levels of the reference images.

We expect the area around the catheter–tissue interface to have the highest PSNR score at lower power levels, since these areas have better contrast and detail at the lower power levels. Furthermore, when the region of interest focuses on the deep tissue structures and bone surface, the PSNR should peak near or at the highest power reference image, since these structures are only visible in the high power images.

Based on Fig. 6, the PSNR contours at the catheter–tissue interface helps us rule out the adaptive histogram and Reinhard DoG TMOs as viable candidates. For these two TMOs, the catheter–tissue interface shows a higher correspondence with the high-power images; however in reality, these structures are overexposed and have lost detail at the higher power levels. Therefore, these two TMOs fail to preserve the detail and contrast in these regions. Of the remaining three, the Reinhard global and the Reinhard bilateral TMOs have almost identical performance, suggesting that the high frequency texture in ultrasound images results in a large kernel size, which reduces the compressive power of the bilateral approach to that of the

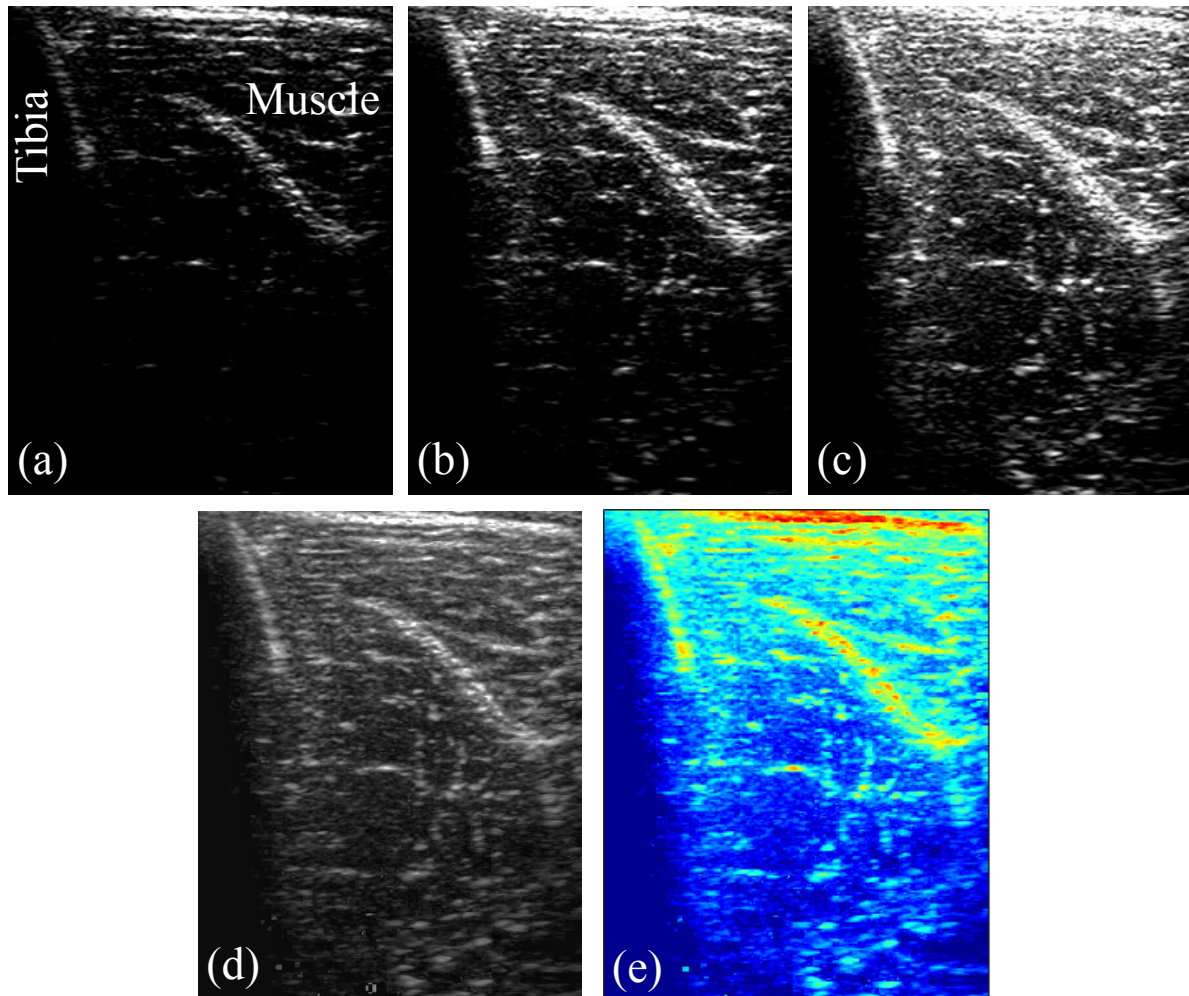


Fig. 7 HDR-US image of a human lower leg acquired in the transverse plane. The tibialis anterior muscle and the surface of the tibia can be seen. Images at fifteen power levels were acquired, three of which are shown in (a) -15.9 dB, (b) -6.3 dB and (c) 0 dB. (d) HDR image tone-mapped using the Durand TMO. (e) HDR-US image intensities prior to tone mapping, shown in false-color with values ranging from 0.1 to 49 , almost three orders of magnitude of dynamic range.

global. The Durand TMO has the widest range of power levels, indicating that this TMO preserves the widest dynamic range among all, and is therefore best suited for use in our application. This is expected, since the Durand TMO only compresses the low frequency components of the image and preserves the high frequency details in the image.

3.2 *In Vivo*

In a second experiment, we imaged the lower leg of a healthy human male subject. Fifteen B-mode ultrasound images were acquired in the transverse plane at 6 MHz using the same probe, imaging depth, and power settings as before. The study was approved by the Harvard University Institutional Review Board (IRB). Fig. 7 shows ultrasound images taken at three power levels, as well as the false color HDR-US image displaying the full dynamic range, and its tone-mapped LDR version. The tibialis anterior muscle and the reflection from the surface of the tibia are hard to discern in the low-power image (Fig. 7a), whereas the muscle fascicles and the skin surface are saturated in the high-power images (Fig. 7b-c). Note that the ultrasound image in Fig. 7b is already saturated near the skin surface, while the deep tissue structures are still not visible. This demonstrates that capturing all details of the anatomy is not possible at a single power level.

The HDR-US image in Fig. 7d is tone-mapped using the Durand TMO. This image shows a clear improvement in the contrast and visibility of the deep tissue structures and the muscle fascicles while preventing saturation in the skin and bone surface. Fig. 7e shows the HDR-US image prior to tone-mapping, which contains almost three orders of magnitude of dynamic range.

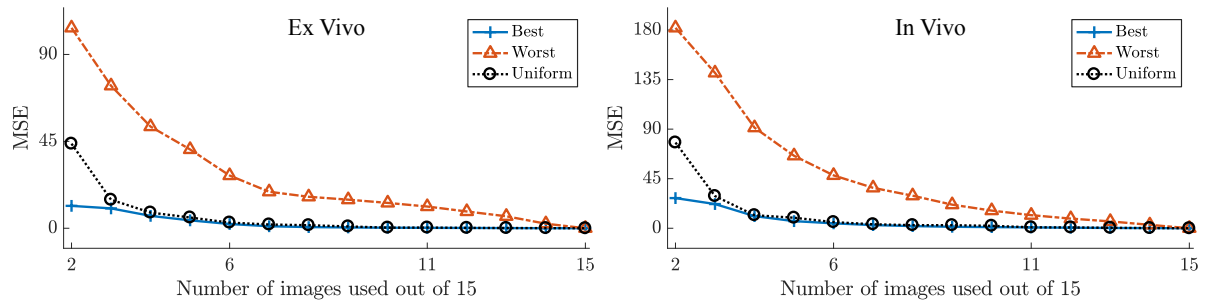


Fig. 8 Leave- p -out cross-validation results for *ex vivo* and *in vivo* studies. The best and worst mean squared error (MSE) scores provide lower and upper bounds on the deviation from the “ground truth” HDR image constructed using fifteen power levels. The uniform sampling strategy quickly converges to the best (*i.e.*, lowest) MSE score

3.3 Number of Images Required for HDR-US Computations

An important point to consider in HDR computations is the number of images that is sufficient for capturing the dynamic range in the imaging field. For a given dynamic range, minimizing the number of images used in HDR reconstruction will allow for faster acquisition, resulting in higher frame rates and improved robustness to tissue and transducer motion. On the other hand, increasing the number of images can reduce noise and increase image quality.

In order to understand how the number of images included in the HDR computations effects HDR quality, we conducted an exhaustive leave- p -out cross-validation. We calculated the mean squared error (MSE) between the “ground truth” (*i.e.*, the HDR image computed using all available images) and the HDR image computed using p fewer images.

In addition to the number of images, the spacing between power levels is also crucial to recovering the widest dynamic range available. Images that are close to each other in terms of their dynamic range will not contribute much toward extending the dynamic range. In our cross-validation studies, we calculated the MSE over all possible image combinations (32,767) to obtain bounds on best and worst MSE possible for each value of p . Fig. 8 shows these values both for the *ex vivo* and *in vivo* datasets, which exhibit almost identical trends. Using the popular structural similarity (SSIM) index [15] instead of the MSE score leads to similar trends; interested readers can run this analysis using our source code in the supplementary material.

We hypothesized that uniform sampling is a good strategy for choosing which images to include in the HDR reconstruction for a given number of images. For example, if three LDR images are to be used, we would pick -26.7 dB, -12 dB, and 0 dB. Our study results validate this hypothesis; the MSE scores for the uniform sampling strategy (shown in black in Fig. 8) quickly converge to the optimal MSE score when four out of fifteen images are used.

4 Discussion

Our results demonstrate that it is possible to enhance the amount of detail that is visible in an ultrasound image by combining images that were acquired at different power levels. This enables capturing a wide dynamic range, from hypoechogenic tissue details to highly reflective instruments, and presenting them to the user in a single image. We tested the performance of five TMOs through qualitative and quantitative analysis. The peak signal-to-noise ratio studies clearly show that the Durand TMO preserves the most detail during tone mapping. The leave- p -out cross-validation results indicate that four uniformly-sampled images are sufficient for capturing the full dynamic range of the specific imaging targets used in our studies. However, the minimum number of images required to construct an HDR-US image depends on the dynamic range of the imaging target. See [4] for a discussion on the number of images required to recover the response curve and the radiance map.

We did not conduct an exhaustive study of TMOs, and the parameters for the TMOs we tested could be optimized to further improve the level of detail in the tone-mapped LDR images. Additionally, the TMO can be a user-selected post-processing option. Just as clinicians can currently control the amount of log compression applied to the image using a knob or slider on the ultrasound machine, once HDR-US is implemented clinically, the users can switch between different TMOs.

One limitation of our study was in terms of acquisition time. We currently do not have a way of automatically adjusting the power settings on the ultrasound machine in real-time. Thus, acquiring an HDR stack of muscle activity, such as the heart, is not possible (although ECG-gating can be used to gate

cardiac imaging). However, ultrasound machine manufacturers, such as BK Ultrasound (Peabody, MA, USA), have already integrated real-time frame-by-frame power level adjustments into their system APIs. Implementation of HDR-US imaging in such systems would enable a set of HDR images to be acquired in just four consecutive frames, which would result in an acquisition time of 50 ms for a probe with an 80 Hz acquisition rate. Images at different power levels can be acquired fast enough such that tissue or probe motion causes minimal misalignment. These images can then be aligned using standard algorithms, such as SIFT or Ward's method [16]. The HDR computations, which are already optimized for mobile computing platforms, can easily be integrated into the ultrasound data processing pipeline, including 3DUS imaging.

In our studies, we were limited to using the S-video output for image capture, which introduces noise in the image. Once the HDR computations are integrated into the image formation pipeline, using the clean digital signals will lead to higher-quality HDR images.

One common issue with imaging rigid instruments with ultrasound is the ringing and shadowing artifacts from the tool. Here, we did not explicitly address these artifacts; however, the tone-mapping step will automatically prevent the instruments from getting saturated, and the HDR computations will reduce the prominence of these artifacts. See [7, 14] for in-depth analyses of ultrasound artifacts arising from metallic instruments, and methods to reduce such artifacts.

Readers familiar with ultrasound imaging may question whether time gain compensation (TGC) can compensate for signal loss due to attenuation, thereby obviating the need for HDR-US. It should be noted that the utility of TGC is limited to cases when the attenuation media is approximately homogeneous in the lateral direction (*i.e.*, at a given depth, the dynamic range of the media is narrow). Otherwise, adjusting the TGC will not prevent structures from getting overexposed or lost.

5 Conclusions

In this work, we have shown that standard HDR imaging techniques can be used to improve the dynamic range of ultrasound images. Based on acoustic equations, we argued that the power settings can be used as an analog to exposure time in HDR calculations. We evaluated five tone mapping operators based on the peak signal-to-noise ratio. The results based on images acquired from *ex vivo* and *in vivo* tissue showed that the Durand TMO produces LDR images with the widest perceivable dynamic range. In future work, we will investigate whether extending the dynamic range of ultrasound imaging improves the performance of further image processing steps. We hypothesize that the benefit of HDR-US will be the greatest in procedure guidance, where we face the broadest dynamic range, due to the presence of highly reflective surgical tools and deep tissue structures.

Acknowledgements The authors would like to thank Neil Tenenholtz, Ph.D. for insightful discussions, and Yashraj Narang, Richard Nuckols, Ph.D. and Mohsen Moradi Dalvand, Ph.D. for their help with data collection.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. Banterle, F., Artusi, A., Debattista, K., Chalmers, A.: Advanced High Dynamic Range Imaging: Theory and Practice. AK Peters (CRC Press), Natick, MA, USA (2011)
2. Cincotti, G., Loi, G., Pappalardo, M.: Frequency decomposition and compounding of ultrasound medical images with wavelet packets. *IEEE Transactions on Medical Imaging* **20**(8), 764–771 (2001)
3. Coupé, P., Hellier, P., Kervrann, C., Barillot, C.: Nonlocal Means-Based Speckle Filtering for Ultrasound Images. *IEEE Transactions on Image Processing* **18**(10), 2221–2229 (2009)
4. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97, pp. 369–378. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA (1997)
5. Durand, F., Dorsey, J.: 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, pp. 257–266. ACM, New York, NY, USA (2002)
6. Erez, Y., Schechner, Y.Y., Adam, D.: Ultrasound image denoising by spatially varying frequency compounding. In: Pattern Recognition: 28th DAGM Symposium, Berlin, Germany, September 12–14, 2006, pp. 1–10. Springer Berlin Heidelberg (2006)

7. Huang, J., Triedman, J.K., Vasilyev, N.V., Suematsu, Y., Cleveland, R.O., Dupont, P.E.: Imaging artifacts of medical instruments in ultrasound-guided interventions. *Journal of Ultrasound in Medicine* **26**(10), 1303–1322 (2007)
8. Hung, A.H., Liang, T., Sukerkar, P.A., Meade, T.J.: High dynamic range processing for magnetic resonance imaging. *PLOS ONE* **8**(11), 1–11 (2013)
9. Perperidis, A.: Postprocessing Approaches for the Improvement of Cardiac Ultrasound B-Mode Images: A Review. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **63**(3), 470–485 (2016)
10. Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., Romeny, B.T.H., Zimmerman, J.B.: Adaptive histogram equalization and its variations. *Comput. Vision Graph. Image Process.* **39**(3), 355–368 (1987)
11. Prince, J.L., Links, J.: *Medical Imaging Signals and Systems* (2nd Edition), 2 edn. Pearson (2014)
12. Reinhard, E., Kunkel, T., Marion, Y., Brouillat, J., Cozot, R., Bouatouch, K.: Image display algorithms for high- and low-dynamic-range display devices. *Journal of the Society for Information Display* **15**(12), 997–1014 (2007)
13. Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.: Photographic tone reproduction for digital images. In: *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '02*, pp. 267–276. ACM, New York, NY, USA (2002)
14. Ren, H., Anuraj, B., Dupont, P.E.: Varying ultrasound power level to distinguish surgical instruments and tissue. *Medical & Biological Engineering & Computing* (2017)
15. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004)
16. Ward, G.: Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures. *Journal of Graphics Tools* **8**(2), 17–30 (2003)