

# Real-time Uncertainty Visualization for B-Mode Ultrasound

Christian Schulte zu Berge, Denis Declara, Christoph Hennersperger, Maximilian Baust, and Nassir Navab




Figure 1: **Uncertainty visualization for B-mode abdominal ultrasound.** Given the original B-mode ultrasound image (a), we first compute the corresponding Confidence Map (b), which maps regions of low attenuation to high confidence. Using different schemes, we then compute the proposed fused uncertainty visualizations. For educational purposes, we introduce an uncertainty color overlay (c). For clinical applications, we further propose mapping to chroma (d) and to fuzziness (e) as perceptual visualization schemes maintaining the original diagnostic value.

## ABSTRACT

B-mode ultrasound is a very well established imaging modality and is widely used in many of today's clinical routines. However, acquiring good images and interpreting them correctly is a challenging task due to the complex ultrasound image formation process depending on a large number of parameters. To facilitate ultrasound acquisitions, we introduce a novel framework for real-time uncertainty visualization in B-mode images. We compute real-time per-pixel ultrasound Confidence Maps, which we fuse with the original ultrasound image in order to provide the user with an interactive feedback on the quality and credibility of the image. In addition to a standard color overlay mode, primarily intended for educational purposes, we propose two perceptual visualization schemes to be used in clinical practice. Our mapping of uncertainty to chroma uses the perceptually uniform  $L^*a^*b^*$  color space to ensure that the perceived brightness of B-mode ultrasound remains the same. The alternative mapping of uncertainty to fuzziness keeps the B-mode image in its original grayscale domain and locally blurs or sharpens the image based on the uncertainty distribution. An elaborate evaluation of our system and user studies on both medical students and expert sonographers demonstrate the usefulness of our proposed technique. In particular for ultrasound novices, such as medical students, our technique yields powerful visual cues to evaluate the image quality and thereby learn the ultrasound image formation process. Furthermore, seeing the distribution of uncertainty

adjust to the transducer positioning in real-time, provides also expert clinicians with a strong visual feedback on their actions. This helps them to optimize the acoustic window and can improve the general clinical value of ultrasound.

**Keywords:** Ultrasound, Uncertainty Visualization, Confidence Maps, Real-time.

**Index Terms:** Computer Graphics [I.3.m]: Miscellaneous Computer Applications [J.3]: Life and Medical Sciences—Health.

## 1 INTRODUCTION

Ultrasound provides a very effective and low-cost real-time imaging modality for both diagnostic and intra-operative use and is thus part of the clinical routine in a variety of medical applications. However, acquiring a good image (e.g. in terms of high diagnostic value) is a non-trivial task due to the highly complex ultrasound image formation process, which is dependent on many explicit and implicit imaging parameters. Furthermore, some target anatomies can not be directly reached but need to be scanned by circumventing strong reflectors such as bones, which prevent the acquisition of images underneath. Thus, an optimal acoustic window has to be found in order to allow for a reliable diagnosis based on ultrasound imaging. For instance, to image the kidney, sonographers usually use the liver as an acoustic window since it has the best transmission properties of the surrounding anatomy. Hence, optimizing the imaging parameters in order to get an optimal view is a non-trivial task, which requires not only a lot of experience of the clinician, but also patience by both the clinician and the patient. In particular medical trainees and ultrasound novices have difficulties in getting the right image needed for their clinical objectives since traditional




Figure 2: **Kidney ultrasound using the liver as acoustic window** with applied uncertainty visualization using chroma as visual variables (cf. Section 3.3). A slight repositioning of the transducer results in a considerable increase of confidence in image (b) compared to (a).

ultrasound imaging does not provide a direct qualitative feedback on the image quality.

In this paper we present novel concepts for uncertainty visualization, specifically designed for ultrasound imaging, allowing sonographers to intuitively and interactively observe the quality of the ultrasound image for each anatomical target (cf. Figure 2). This simplifies their search for the best acoustic window while manipulating the probe and provides an entirely new approach to facilitate ultrasound acquisitions. We introduce a system, which is capable of computing per-pixel ultrasound Confidence Maps [8] in real-time, leveraging the highly dynamic nature of ultrasound imaging. We interpret these Confidence Maps as per-pixel uncertainty information, which we expose to the user. By integrating the uncertainty information directly into the original B-mode image, we can provide the sonographer with interactive feedback on the local image quality and credibility without requiring him to perform an additional mental mapping.

To this end, we introduce three different visualization schemes. A color overlay mode is primarily intended for educational purposes to support medical students understanding the complex ultrasound image formation process. For clinical usage, we suggest mapping of uncertainty to chroma using a perceptually uniform color model to augment the image with the uncertainty information while maintaining its original B-mode intensities and thus its original diagnostic value. As an alternative, which keeps the image in its original grayscale format, we propose mapping uncertainty to fuzziness.

## 2 RELATED WORK

Compared to other anatomical medical imaging modalities such as CT or MRI, medical ultrasound exhibits several benefits, such as being comparatively cheap, portable and real-time capable. It is therefore widely used in today's clinical practice, for instance for abdominal, pediatric, head and general vascular applications [13]. However, as previously motivated in the introduction, the ultrasound image formation process is influenced by various physical ultrasound parameters such as frequency, focus, and depth, as well as by external factors such as probe positioning, probe pressure, patient positioning and patient breathing cycle [1]. Furthermore, a wide range of image artifacts may occur, which need to be interpreted correctly [17]. Sonographers need to be aware of all these caveats of ultrasound imaging in order to correctly understand the image. Thus, training is an important aspect for medical students when learning ultrasound [5]. In particular in trauma applications, where time is critical, the surgeon has to determine possible frac-

tures and lesions as quickly as possible and has a minimal margin for error [10]. With this motivation in mind, we developed our work with the aim to support both medical students in learning sonography as well as expert clinicians for a more quick and intuitive interpretation of ultrasound images by providing an interactive feedback on the image quality and uncertainty distribution.

Uncertainty is an essential property of all kinds of information and is formed of many different facets, which can be categorized into nine different types: Accuracy/error, precision, completeness, consistency, lineage, currency, credibility, subjectivity and inter-relatedness [21]. Visualizing uncertainty in a helpful manner requires mapping these facets to visual variables of the image domain as intuitively as possible. MacEachren et al. identify nine different visual variables and empirically explore their intuitiveness to represent uncertainty [12].

In the last years, the visualization community has made great efforts to improve uncertainty visualization [15]. However, applying such techniques to medical data poses particular challenges since clinicians often need to evaluate very specific details of an image, whose preservation is crucial. Furthermore, the spatial domain is fixed by the pixel/voxel structure, the color domain is often constrained by the applied transfer function, and the temporal domain is also fixed as soon as it comes to real-time imaging. We provide a detailed in-depth discussion of the different approaches to uncertainty visualization in Section 3.3.1

Different works have introduced the concept of uncertainty to ultrasound with the objective of improving ultrasound registration, compounding or tissue classification [6, 23, 19]. Karamalis et al. present a method for estimating per-pixel confidence in ultrasound images and demonstrate the applicability of the method to different problems including ultrasound reconstruction and registration [8]. However, all this information was never exposed to clinicians and always remained an internal component of image processing algorithms. The way B-mode ultrasound is presented today is still the same as many years ago even though there is a lot more information available.

Advanced color spaces are often employed in both visualization and computer vision tasks. The HSL and HSV color spaces [20] can be used in order to measure or alter the presence or absence of single physiological color criteria such as hue, saturation or lightness. However, they are not perceptually uniform meaning that changing the saturation component may also lead to a change in the perceived brightness. As for all visualization tasks, it is crucial to take the human visual perception into account. When integrating the uncertainty information directly into B-mode ultrasound images using

color hue or saturation as visual variable, one has to make sure that the perceived intensity or brightness of a pixel remains the same, which renders HSL and HSV useless for our application. In contrast, the CIE  $L^*a^*b^*$  color model provides a perceptually uniform color space where perceptual similarity is measured by the Euclidean distance of two color points. For a direct mapping to physiological color criteria the  $L^*C^*h^*$  color space was defined as a polar coordinate derivative where the three components represent lightness, chromacity and hue [14].

### 3 METHODS

Our work for real-time ultrasound uncertainty visualization consists of two main parts: First, we need to generate per-pixel uncertainty information for B-mode ultrasound images in real-time in order to use them for subsequent visualizations. We therefore leverage the continuous nature of ultrasound and formulate an incremental computation of ultrasound Confidence Maps. In a second step we fuse the uncertainty information into the B-mode image using perceptual visualization schemes. It is important to mention that these two parts are technically independent from each other and each component is essentially a separate and replaceable black box.

#### 3.1 Generating Uncertainty Information Through Ultrasound Confidence Maps

In order to compute the needed uncertainty information, our implementation uses Confidence Maps, originally proposed by Karalis et al. [8]. Even though they model ultrasound physics only to a limited extent, and for instance do not detect reverberation artifacts, their per-pixel attenuation maps can be interpreted as an approximation of uncertainty information. In fact, due to the complex nature of sound-tissue interactions it may be impossible to compute exact uncertainty information from B-mode images alone without having detailed information about the processing pipeline of the ultrasound machine and all acquisition parameters. Since this is beyond the scope of this work, we want to emphasize that our method for estimating uncertainty is independent of the visualization and could later be replaced or extended.

Ultrasound Confidence Maps are defined as the solution to a random walks equilibrium problem modeled using ultrasound specific constraints. Building an 8-connected graph on the ultrasound B-mode image, one computes the probability for a random walker starting at the respective pixel to reach one of the virtual transducer nodes at the top before reaching one of the virtual sink nodes at the bottom. Since the underlying graph is undirected, this diffusion problem can effectively be described by a symmetric Laplacian matrix  $L$ , which yields a linear system of the form


$$Lx = b, \quad (1)$$

where  $x$  is the Confidence Map in vectorized form and  $b$  encodes the Dirichlet boundary conditions. Figure 3 shows an exemplary Confidence Map for liver ultrasound that clearly depicts shadowed regions behind strong reflectors.

We assume ultrasound Confidence Maps to be inversely related to the amount of uncertainty in the image, more precisely its facets of accuracy, precision and credibility. Thus, we obtain the per-pixel uncertainty information by applying a direct inverse linear mapping

$$U(x) := 1 - CM(x), \quad (2)$$

where  $CM(x) \in [0, 1]$  is the confidence map value at pixel  $x$  in the image domain. Due to the nature of the Confidence Map formulation where the whole value range of  $[0, 1]$  is always fully used for each image, the computed values exhibit relative information with respect to the current image. Thus, special attention has to be paid on this fact when comparing images showing different anatomy.



(a) Original liver ultrasound image. (b) Corresponding confidence map.

Figure 3: **Confidence maps on liver ultrasound:** (a) shows the original ultrasound image, (b) shows the corresponding confidence map, bright regions depicting regions of high confidence, dark regions depicting regions of low confidence.

#### 3.2 Real-time Confidence Maps Through Incremental Computation

Due to the size and complexity of  $L$ , it is not feasible to solve Equation 1 with direct methods (e.g. pivoted LU decomposition) in real time with today's hardware. In order to yield real-time Confidence Maps for our system, we leverage the dynamic nature of ultrasound acquisitions and the temporal coherency of its images based on the high acquisition rates of today's systems. Consecutive images in ultrasound sequences usually differ very little and therefore the corresponding Confidence Maps will likely be very similar. We exploit this fact by using an iterative Jacobi-preconditioned Conjugate Gradient solver instead of directly solving Equation 1 through matrix decompositions. This allows us to execute an incremental computation scheme, where we directly use the resulting Confidence Maps as initialization for the subsequent frame. This greatly reduces the number of needed iterations per frame as we have a better convergence to the true solution in the limited time budget available per frame.

As already discussed in Section 3.1, Confidence Maps exhibit only a relative measure of confidence. Their strength is not the exact confidence value at a single pixel, but rather the confidence distribution within the image. Thus, due to noise in the original ultrasound, the Confidence Maps of multiple consecutive frames may exhibit a flickering behavior when being watched in a sequence. We eliminate this effect by applying an additional alpha beta filter [3] to reduce such artifacts. Being a variant of the Kalman filter, it recursively operates on the stream of computed Confidence Maps and produces a smoothed version by averaging the current image with a prediction based on the previous images. We empirically selected a configuration of  $\alpha = 0.36$  and  $\beta = 0.005$  providing good with both damping of flickering and preservation of the original confidence distribution and temporal responsiveness of the estimates.

#### 3.3 Uncertainty Visualization Techniques for B-mode Ultrasound

The second part of our system is the actual visualization aspect on how to fuse the additional uncertainty information with the original B-mode image. Essentially, we target two different applications with our work: On the one hand, we are convinced that exposing uncertainty information to ultrasound novices and trainees helps them to better understand the complex ultrasound image formation process. At the same time, we also expect expert sonographers to benefit from uncertainty visualization as the additional information may improve the diagnostic value of the image. Regarding these two target applications, we yield the following requirements for ultrasound uncertainty visualization schemes.

- The uncertainty information should be fused directly into the original B-mode sequence. Thus, both the spatial and the temporal domain are fixed.

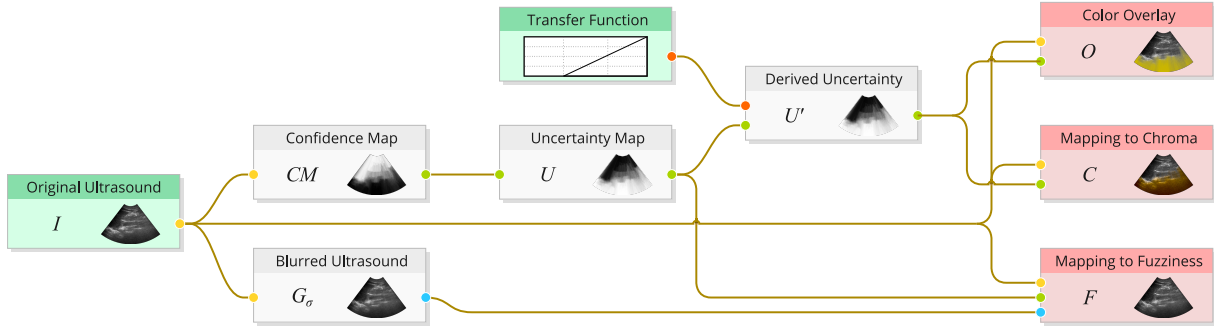


Figure 4: **Schematic diagram** of the different proposed uncertainty visualization schemes. Given the original B-mode ultrasound image  $I$ , we compute its Confidence Map  $CM$  and a Gaussian blurred version  $G_\sigma$ . The uncertainty map  $U$  is derived from  $CM$  by inverse linear mapping (cf. Equation 2). For the color overlay and the uncertainty mapping to chroma, we compute a derived uncertainty measure  $U'$  using a transfer function. Finally, the uncertainty information is fused into the original ultrasound image using one of the three visualization schemes.

- The primary information in the uncertainty maps are not the exact per-pixel uncertainty values but the distribution of the uncertainty with respect to the anatomy (cf. Section 3.1).
- For educational applications, the uncertainty distribution in the image should be easily and intuitively perceivable. Even small changes in the uncertainty distribution should be clearly observable when repositioning the ultrasound probe in order to maximize the learning effect.
- For clinical applications, the diagnostic information in the B-mode image must not be impaired. Thus, the original image intensities should be preserved as good as possible and no image regions should be occluded.

### 3.3.1 Selection of Visual Variables

Given the above requirements for our uncertainty visualization, the number of viable visual variables for depicting the uncertainty information is limited.

One traditional technique is the usage of glyphs for depicting uncertainty, with error bars in 1D visualization being the classic example. Glyphs have the advantage of offering a large number of visual variables that can be used to alter their appearance. For instance, MacEachren et al. evaluate 11 different mappings of uncertainty to point glyphs [12]. Glyphs excel in depicting uncertainty when used in sparse layouts allowing the observer to individually focus on single glyphs in order to read their information. However, this does not work well for our application where the goal is to visualize the distribution of a dense 2D scalar field. Although dense glyph fields have also been successfully used for depicting global information [9, 2], we do not consider them for our work since early experiments did not show promising results. The work of Sanyal et al. supports this fact as their experiments showed that in many uncertainty related tasks on 2D data sets the different glyph mappings perform significantly worse than surface color mapping [16]. Furthermore, adding glyphs to the B-mode image would occlude the original ultrasound image, which is undesirable for clinical applications.

One intriguing approach is to extend the 2D data to the third dimension and map uncertainty to the  $Z$  axis in a 3D rendering [4, 7]. While this may be a valid method for applications such as geospatial visualization, it can not be applied to our use case since, due to the 2D projection of a 3D scene, it requires the user to interact with the camera to get the full information. Furthermore, as mentioned above, the spatial domain of our visualization is fixed as clinicians expect a 2D image when performing 2D B-mode ultrasound.

Another approach is to exploit the spatial domain for depicting uncertainty, for instance through animations, animated jittering or

probabilistic animation [4, 11]. However, since we are working with real-time ultrasound sequences, the temporal domain is fixed and such approaches are not applicable.

Given these considerations and the particular requirements of our intended uncertainty visualization, we selected the visual variables of color and texture. They both do not affect the spatial and temporal image domain and are very powerful and intuitive for expressing general uncertainty [12]. In total, we propose three different uncertainty mapping techniques (cf. Figure 1) for the two applications, which we will discuss in detail in the following sections. Since both the B-mode image and the uncertainty map are in the same image domain, no coordinate transformation is necessary and the mapping techniques are focused on the optical properties.

### 3.3.2 Uncertainty as Color Overlay

For educational applications, the focus of the visualization should be on the uncertainty information and even small changes in the distribution should be clearly distinguishable by the observer. The corresponding ultrasound B-mode image should be shown as anatomical reference in order to allow for an understanding of the connection between image features and their effects on the uncertainty. Therefore, we combine the visual variables of hue and value in our proposed color overlay scheme (cf. Figure 4). Compared to the other presented mapping schemes, the combination of the two visual variables makes even subtle changes of uncertainty visible to the observer. We deem this an important feature to teach ultrasound novices the caveats of the ultrasound image formation process.

As previously introduced in Section 3.1, for each pixel  $x$ , let  $U(x)$  be its uncertainty value and  $I(x)$  be its original B-mode intensity. Since the color overlay is a very obtrusive mapping scheme, we use a transfer function to apply a thresholding to the uncertainty measure and define the derived uncertainty as  $U'(x) := \max(0, 2U(x) - 1)$ . Using this derived measure instead of the original  $U(x)$  avoids overlaying regions of negligible uncertainty. We first generate the color overlay in HSV color space as

$$C(x)_{HSV} := (H, U'(x), V), \quad (3)$$

with constant hue  $H \in [0, 1]$  and constant value  $V \in [0, 1]$ . We chose a bright orange color with  $H = 0.15$  and  $V = 0.8$  to avoid lowering the contrast to Doppler ultrasound utilizing the colors blue and red. In a second step, we linearly mix the color overlay transformed to RGB color space with the original B-mode to yield the final pixel color

$$O(x) := U'(x) \cdot C_{RGB}(x) + (1 - U'(x)) \cdot I(x). \quad (4)$$

Figure 1 (c) shows the uncertainty color overlay applied to the original abdominal ultrasound image from Figure 1 (a).




Figure 5: **Processing pipeline** of our reference implementation. We receive the original ultrasound B-mode image through OpenIGTLink and perform an Gaussian blur and a resampling in order to improve the PCG convergence performance. The Confidence Map solver then incrementally computes the B-mode image’s Confidence Map by using the previous image’s Confidence Map as initialization. Finally, we perform the uncertainty mapping using one of the three presented schemes (cf. Section 3.3) and transform the image from polar to cartesian coordinates.

Since the original B-mode intensities are altered in this mapping scheme, we propose to use it only for educational purposes to give ultrasound novices a better understanding of the image formation process, but do not consider it for clinical usage. Furthermore, the color overlay also partially occludes the original ultrasound intensities. Although image regions of low confidence may not be reliable enough for diagnosis, they still contain structural information, which can help the sonographer in navigating towards the correct anatomy and optimizing the acoustic window. Thus, hiding these parts completely is disadvantageous for clinical usage, which was later also confirmed by some candidates during our evaluation (cf. Section 5). Therefore, we additionally propose two alternative mapping schemes maintaining structural information also in highly unreliable regions.

### 3.3.3 Uncertainty Mapping to Chroma

For clinical applications, we propose additional uncertainty visualization schemes that maintain the structural information in the ultrasound B-mode image also in unreliable regions. Similar to the color overlay, uncertainty mapping to chroma also uses color to depict uncertainty but uses chroma as visual variable. In order to preserve the original ultrasound image’s diagnostic value, we need to ensure that the perceived intensity remains the same when augmenting the image with uncertainty information. Therefore, we perform the chroma modification in the perceptually uniform CIE  $L^*C^*h^*$  color space, the polar coordinate derivative of the CIE  $L^*a^*b^*$  color model [14].

As illustrated in Figure 4, we compute the derived uncertainty from the original Confidence Map as  $U'(x) := \max\left(0, \frac{3U(x)}{2} - \frac{1}{2}\right)$  to again avoid coloring regions with negligible uncertainty. The final pixel color in  $L^*C^*h^*$  space is given by

$$C(x)_{L^*C^*h^*} := (I(x)_{L^*}, U'(x), H), \quad (5)$$

where  $I(x)_{L^*}$  is B-Mode intensity transformed to  $L^*$  space and  $H = 0.23$  is a bright orange. Again, we chose bright orange as hue for depicting uncertainty to avoid lowering the contrast to Doppler ultrasound.

### 3.3.4 Uncertainty Mapping to Fuzziness

Many clinicians prefer to look at grayscale ultrasound images as they have been trained to do so. Since texture is a very effective visual variable that keeps the spatial, temporal and color domain of the original image, we selected it as third visualization scheme for real-time B-mode ultrasound. More precisely, we use the visual cue of fuzziness to map the uncertainty information. Therefore, we fuse the ultrasound image with its uncertainty map such that regions of low uncertainty appear sharp and regions of high uncertainty appear fuzzy. As a matter of fact, according to MacEachren et al., this visual variable is also the most intuitive to represent uncertainty [12].

Our proposed uncertainty mapping to fuzziness combines a slight Gaussian blur of the original ultrasound image with its un-

sharp mask (subtraction of the blurred image from the original image). As illustrated in Figure 4 we compute the confidence map of the original image and apply Equation 2 to obtain the uncertainty value  $U(x)$  for each pixel  $x$ . Since this is a diverging mapping scheme, where regions with high uncertainty will be blurred, and regions of low uncertainty will be sharpened, we can directly use  $U(x)$  and do not need to compute a derived uncertainty measure. We compute a Gaussian filtered version of the original image and combine the two to yield the final pixel value  $F(x)$  as

$$F(x) = U(x)G_\sigma(x) + ((1 - U(x)) \cdot (2I(x) - G_\sigma(x))), \quad (6)$$

where  $I(x)$  is the original ultrasound intensity and  $G_\sigma(x)$  the corresponding intensity in the Gaussian with parameter  $\sigma$ . We choose  $\sigma = 2.5$  to limit the blurring to a tolerable amount but still yield the effect of perceived differences in fuzziness. An exemplary result can be seen in Figure 1 (e).


It should be noted that this mapping scheme certainly alters the original B-mode image in a way that may reduce the amount of original information in regions of high uncertainty. However, discussions with clinicians showed that they nevertheless like uncertainty mapping to fuzziness and appreciate the intuitiveness of the visual variable. Our evaluation results in Section 5 underline this fact.

## 4 SYSTEM


Our reference implementation consists of an Ultrasonix RP (Analogic Corporation, Peabody, MA, USA) device equipped with a curvilinear ultrasound transducer (model C5-2/60) for acquisition, as well as a workstation where all the computation is performed. We use OpenIGTLink [22] for streaming the ultrasound frames and imaging parameters to our computation framework, where we perform the necessary processing steps before eventually routing the fused image back to the ultrasound device’s display. The processing pipeline is implemented in the CAMPVis framework [18] and entirely executed on the GPU using both CUDA (solving of Equation 1) and OpenGL/GLSL (all other processing steps) in order to achieve optimal performance.

Our image processing pipeline consists of several steps and is illustrated in Figure 5. After acquiring the B-mode image, we perform a Gaussian filtering as well as a resampling. Gaussian filtering is required in order to remove high-frequency noise as well as to allow for our uncertainty mapping to fuzziness. We perform the downsampling to speed up the computation of the Confidence Maps and achieve better convergence (cf. Section 5.1). For our experiment setup we used a 0.5 scaling factor and a smoothing factor  $\sigma$  of 2.5.

The next step in our processing pipeline is the iterative solver that computes the Confidence Maps. To ensure a smooth real-time experience with constant frame rate, we restrict our iterative solver to perform as many iterations as possible within a fixed time budget, inversely proportional to the frame rate of the ultrasound device. We use a CUDA implementation of the Conjugate Gradient method with a Jacobi preconditioner provided by the CUSP library to solve



(a) Iterations/resample scale configurations to achieve 30ms execution time on average.



(b) Error (residual norm) for the different iterations/resample scale configurations when maintaining 30ms execution time.

Figure 6: Graphs showing the **dependency of iteration count and resample scale with respect to the runtime and error** (residual norm). When maintaining a target runtime of 30ms, a lower resample scale yields a lower error.

the underlying Dirichlet problem. The equation system matrix  $L$  (cf. Equation 1) is explicitly constructed on each frame and stored on the GPU. Since it has only 9 diagonals with non-zero entries, it can be efficiently stored using the sparse DIA matrix storage format. After solving the system, the resulting confidence map is used as initialization for the upcoming image since it provides a good initial approximation, due to the temporal coherency of ultrasound.

Finally, one of the discussed uncertainty visualization schemes is applied and the rectilinear B-mode image in polar coordinates is scan converted to Cartesian coordinates using the known probe geometry and dynamically queried imaging parameters.

## 5 EVALUATION AND RESULTS

We structured the evaluation of our proposed system into three parts: To quantitatively analyze the system performance, we executed several simulated runs on pre-recorded clinical ultrasound scans to assess error and speed with respect to various system parameters. To evaluate the educational purpose of our system, we ran a large user study with medical students and ultrasound trainees. Finally, we presented our visualizations to ultrasound experts in order to evaluate the potential of our system in clinical practice.

### 5.1 System Evaluation

For the performance evaluation, we ran our system on a Linux notebook with an Intel i7 processor and a nVidia GTX 750M GPU. Instead of streaming live images from the ultrasound machine, we streamed a pre-recorded sequence of patient abdominal ultrasound containing 392 images of 512x512 pixels resolution from the hard disk via OpenIGTLink link. To examine the influence of the system performance with respect to the different parameters, we ran our real-time solver and visualization in different configurations. The analytical solving of Equation 1 using standard Cholesky decomposition required an average of more than 2.2 seconds per frame and is thus far from real-time. As shown in Figure 6, computing the Confidence Maps on a lower resolution is beneficial in terms of error. In 30ms execution time, our system ran 30 PCG iterations on full resolution, 50 iterations with a resample scale of 0.75, and 100 iterations with a resample scale of 0.5.

Table 1, as well as Figure 7, show how the residual norm, expressed as  $\|\mathbf{r}\|_2$ , with  $\mathbf{r} = L\mathbf{x} - \mathbf{b}$ , evolves over time, with respect to the iteration count. When performing 110 PCG iterations, the residual norm never exceeds 0.005. For comparison, the dashed line shows the error progression when the solver is reinitialized at each iteration, which yields worse results than 20 incremental iterations. This demonstrates the importance of leveraging the dynamic




Figure 7: **Error (residual norm) over time** for different PCG iteration counts. The resample scale was fixed to 0.5.

nature of ultrasound and using the solution of the previous frame as initialization.

As it can be observed in Figure 6a, a significant jitter is present in the processing time when using a fixed number of iterations. Therefore, in order to ensure real-time visualization, our system does not use a fixed amount of iterations as stopping criterion, but rather a fixed time budget wherein we perform as many iterations as possible.

Table 1: Aggregated **performance results in terms of runtime and error** with different configurations for the number of iterations. The resample scale was fixed to 0.5. Data set was a patient kidney ultrasound sequence of 392 frames (first 10 frames omitted).

# Iterations	Avg Runtime (ms)	Avg. Error (Residual Norm)	Peak Error (Residual Norm)
Direct solver	2214.68	0.0	0.0
20 (incremental)	10.84	0.014	0.033
50 (incremental)	18.07	0.006	0.014
80 (incremental)	25.37	0.003	0.008
110 (incremental)	32.46	0.002	0.005
110 (discarding)	32.59	0.020	0.026

### 5.2 User Study with Ultrasound Novices

In a first user study we evaluated the educational value of our proposed system. We equipped our ultrasound machine with an Ultrasonix C5-2/60 convex abdominal transducer and asked ultrasound novices to locate different structurally deep anatomies in a CIRS




Figure 8: Evaluation results on user study with ultrasound novices. While there is no clear favorite visualization scheme, almost all test subjects appreciate the uncertainty visualization.

abdominal phantom while using our uncertainty visualization techniques. In total we interviewed 13 medical students, which all had very limited experience with ultrasound (average of 5.6 performed ultrasound examinations, minimum 0, maximum 20).

For a quantitative evaluation, we recorded the acquisitions and measured the time the users required to optimize the view on anatomical targets. After giving the participants some time to familiarize themselves with the phantom anatomy, we asked them to find an optimal view onto the vessel targets in the left and right kidney, as well as onto the portal vein. We then measured the required time to optimize the view for each visualization scheme by counting the number of frames between the first frame where the target anatomy was in the field of view until the frame where the student defined the view as optimal in his personal opinion. To avoid biasing the results, we shuffled the order of the visualizations for each user. While the results in Table 2 show no significant differences between color overlay and chroma mapping, the time needed with fuzziness mapping is consistently lower (in average 0.79 seconds) than with the one of other two mapping schemes. Furthermore, the students performed significantly worse (in average 1.86 seconds longer) with only the original B-mode image compared to all of our proposed visualization schemes, which supports our idea that visualizing uncertainty helps the user in interactively assessing the quality of the acquired ultrasound image.

After the experiment, the test subjects additionally answered a short questionnaire on educational value, helpfulness and intuitiveness of the different visualizations. The results in Figure 8 show that all test subjects appreciate the uncertainty visualization. Interestingly, there is no clear favorite visualization scheme. While some students preferred the color overlay, others preferred the mapping of uncertainty to chroma. However, fuzziness mapping yielded an overall worse response, as many students considered this scheme as not helpful for diagnosis nor intuitive to read, which is interesting as it contradicts the quantitative results of Table 2.

Table 2: Average time (seconds) ultrasound novices required to optimize the view on target (aggregated results from 9 of the 13 users, since not all acquisitions were complete).

	Left Kidney	Right Kidney	Portal Vein
Original B-Mode	6.73 ± 4.3	4.96 ± 1.6	4.78 ± 1.2
Color Overlay	4.74 ± 3.8	4.26 ± 1.5	3.30 ± 1.7
Chroma Mapping	4.19 ± 2.3	3.42 ± 1.8	3.44 ± 1.5
Fuzziness Mapping	3.49 ± 3.9	3.15 ± 0.7	2.67 ± 1.1

### 5.3 User Study with Ultrasound Experts

In a second user study, we presented our system to expert sonographers in order to evaluate the clinical significance of real-time ultrasound uncertainty visualization. Therefore we applied our visualization schemes to clinical data of abdominal ultrasound and presented the results to 7 experts (5 clinicians, 2 senior researchers). We presented them two different sequences with patient data of three different anatomies (liver, kidney, spleen) and asked them about perception, clinical value as well as whether our techniques assist in finding the optimal acoustic window. Since, for clinical usage, we do not want to hide information, we showed them only mapping to chroma and mapping to fuzziness.




Figure 9: Questionnaire results on expert sonographers' uncertainty perception. Each question was answered independently for the two different sequences. Therefore, there are 14 answers in total.

Interestingly, the expert sonographers clearly prefer mapping to fuzziness over mapping to chroma (Figure 9). This discrepancy compared to ultrasound novices is probably due to the fact that experts got used to the monochrome appearance of B-mode ultrasound over the years. Thus, they prefer a visualization scheme that keeps the image in its original grayscale domain.

Regarding clinical significance, the study shows significant improvements compared to the default visualization in today's ultrasound devices (Table 3). The clinicians reported that seeing the amount of uncertainty dynamically adapting to the ultrasound view provides them with a very strong feedback on the image quality and its credibility. Here, 6 out of 7 stated that our uncertainty visualization helps with the correct interpretation of the images and also 6 out of 7 participants confirmed that the proposed real-time visualization schemes assists in optimizing the acoustic window. One clinician found that the computed Confidence Map was unexpected in one of the sequences and therefore confused him with the correct

interpretation of the image. Nevertheless, he stated that the technique was helpful for optimizing the acoustic window. One other clinician confirmed the uncertainty visualization being helpful for the correct interpretation but had doubts that it helps with optimizing the acoustic window.

Table 3: Results on expert sonographers evaluating the clinical value.

	Yes	No
Helps with Correct Interpretation	86%	14%
Helps Optimizing Acoustic Window	86%	14%

## 6 CONCLUSION

With this work we introduce novel visualization techniques for B-mode ultrasound imaging to provide the user with interactive feedback on the level of uncertainty present in the current image. Our evaluation demonstrates that these techniques support the clinician in finding the optimal acoustic window, which is a highly important though non-trivial task, especially for less experienced users struggling to cope with the complex ultrasound image formation process. We propose three different visualization schemes for depicting the per-pixel uncertainty present in ultrasound images. Both schemes show the level of uncertainty (i.e. accuracy, precision, credibility) directly integrated in the ultrasound image and take the human visual perception into account to preserve the diagnostic value. Discussions with clinicians showed that many clinicians are not fully aware of the amount of uncertainty in ultrasound and highly appreciate the additional feedback on the credibility of the image our techniques provide.

To our knowledge, we propose a unique technique to facilitate the acquisition and enable interactive assessment of the current quality in everyday ultrasound. The evaluation results of both perception and clinical value suggest its great potential to become an impactful tool enriching existing ultrasound imaging protocols. Clinicians can use the information to optimize the acoustic window or acquisition parameters and thereby the overall image quality and diagnostic value. We are currently in the process of testing our system with other anatomies, for instance with the very challenging transcranial ultrasound. Additionally, we would like to focus our future work on integrating algorithms that capture further aspects of uncertainty, for instance reverberation artifacts.

## ACKNOWLEDGEMENTS

This research was performed in the course of the *CAMPVis* project funded within the *SoftwareCampus* program of the German Federal Ministry of Education and Research (*BMBF, Förderkennzeichen 01IS12057*). We thank Dr. Reza Ghotbi from Klinikum Pasing, Dr. Florian Muschalla from Klinikum Rechts der Isar München, as well as Dr. Anna Maria von der Heide and Dr. Simon Weidert from LMU Klinikum München for their valuable input and support.

## REFERENCES

- [1] J. E. Aldrich. Basic physics of ultrasound imaging. *Critical care medicine*, 35(5 Suppl):S1317, May 2007.
- [2] R. Borgo, J. Kehrler, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In M. Sbert and L. Szirmay-Kalos, editors, *Eurographics 2013 - State of the Art Reports*. The Eurographics Association, 2012.
- [3] E. Brookner. *Tracking and Kalman filtering made easy*. A Wiley-Interscience publication. Wiley, 1998.
- [4] R. Brown. Animated visual vibrations as an uncertainty visualisation technique. In *Proceedings of the 2Nd International Conference on Computer Graphics and Interactive Techniques in Australasia and*

- South East Asia*, GRAPHITE '04, pages 84–89, New York, NY, USA, 2004. ACM.
- [5] J. Butter, T. H. Grant, M. Egan, M. Kaye, D. B. Wayne, V. Carrin-Carire, and W. C. McGaghie. Does ultrasound training boost year 1 medical student competence and confidence when learning abdominal examination? *Medical Education*, 41(9):843–848, 2007.
- [6] P. Hellier, P. Coup, X. Morandi, and D. L. Collins. An automatic geometrical and statistical method to detect acoustic shadows in intraoperative ultrasound brain images. *Medical Image Analysis*, 14(2):195–204, 2010.
- [7] D. Kao, J. Dungan, and A. Pang. Visualizing 2D probability distributions from EOS satellite image-derived data sets: a case study. In *Visualization, 2001. VIS '01. Proceedings*, pages 457–589, Oct 2001.
- [8] A. Karamalis, W. Wein, T. Klein, and N. Navab. Ultrasound confidence maps using random walks. *Medical Image Analysis*, 16(6):1101–1112, 2012.
- [9] G. Kindlmann and C.-F. Westin. Diffusion tensor visualization with glyph packing. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1329–1336, Sept 2006.
- [10] J. L. Knudtson, J. M. Dort, S. D. Helmer, and R. S. Smith. Surgeon-performed ultrasound for pneumothorax in the trauma suite. *The Journal of trauma*, 56(3):527530, March 2004.
- [11] C. Lundstrom, P. Ljung, A. Persson, and A. Ynnerman. Uncertainty visualization in medical volume rendering using probabilistic animation. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1648–1655, Nov 2007.
- [12] A. MacEachren, R. Roth, J. O'Brien, B. Li, D. Swingley, and M. Gahegan. Visual semiotics & uncertainty visualization: An empirical study. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12):2496–2505, Dec 2012.
- [13] J. A. Noble, N. Navab, and H. Becher. Ultrasonic image analysis and image-guided interventions. *Interface Focus*, 1(4):673–685, 2011.
- [14] K. N. Plataniotis and A. N. Venetianopoulos. *Color image processing and applications*. Springer, 2000.
- [15] B. Preim and C. Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. The Morgan Kaufmann Series in Computer Graphics. Elsevier Science, 2013.
- [16] J. Sanyal, S. Zhang, G. Bhattacharya, P. Amburn, and R. Moorhead. A user study to compare four uncertainty visualization methods for 1D and 2D datasets. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1209–1218, Nov 2009.
- [17] K. Scanlan. Sonographic artifacts and their origins. *AJR. American journal of roentgenology*, 156(6):12671272, June 1991.
- [18] C. Schulte zu Berge, A. Grunau, H. Mahmud, and N. Navab. *CAMPVis – A game engine-inspired research framework for medical imaging and visualization*. Technical report, Technische Universität München, 2014.
- [19] C. Schulte zu Berge, A. Kapoor, and N. Navab. Orientation-driven ultrasound compounding using uncertainty information. In D. Stoyanov, D. Collins, I. Sakuma, P. Abolmaesumi, and P. Jannin, editors, *Information Processing in Computer-Assisted Interventions*, volume 8498 of *Lecture Notes in Computer Science*, pages 236–245. Springer International Publishing, 2014.
- [20] A. R. Smith. Color gamut transform pairs. *SIGGRAPH Comput. Graph.*, 12(3):12–19, Aug. 1978.
- [21] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. A typology for visualizing uncertainty. In R. F. Erbacher, J. C. Roberts, M. T. Gröhn, and K. Börner, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 5669 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 146–157, Mar. 2005.
- [22] J. Tokuda, G. S. Fischer, X. Papademetris, Z. Yaniv, L. Ibanez, P. Cheng, H. Liu, J. Blevins, J. Arata, A. J. Golby, T. Kapur, S. Pieper, E. C. Burdette, G. Fichtinger, C. M. Tempny, and N. Hata. OpenIGTLink: an open network protocol for image-guided therapy environment. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 5(4):423–434, 2009.
- [23] Y. Yu and J. Wang. Backscatter-contour-attenuation joint estimation model for attenuation compensation in ultrasound imagery. *Image Processing, IEEE Transactions on*, 19(10):2725–2736, oct. 2010.