

Learning the Manifold of Quality Ultrasound Acquisition

Noha El-Zehiry¹, Michelle Yan¹, Sara Good², Tong Fang¹, S. Kevin Zhou¹,
and Leo Grady^{1,*}

¹ Siemens Corporation, Corporate Technology

² Siemens Healthcare

Abstract. Ultrasound acquisition is a challenging task that requires simultaneous adjustment of several acquisition parameters (the depth, the focus, the frequency and its operation mode). If the acquisition parameters are not properly chosen, the resulting image will have a poor quality and will degrade the patient diagnosis and treatment workflow. Several hardware-based systems for autotuning the acquisition parameters have been previously proposed, but these solutions were largely abandoned because they failed to properly account for tissue inhomogeneity and other patient-specific characteristics. Consequently, in routine practice the clinician either uses population-based parameter presets or manually adjusts the acquisition parameters for each patient during the scan. In this paper, we revisit the problem of autotuning the acquisition parameters by taking a completely novel approach and producing a solution based on image analytics. Our solution is inspired by the autofocus capability of conventional digital cameras, but is significantly more challenging because the number of acquisition parameters is large and the determination of “good quality” images is more difficult to assess. Surprisingly, we show that the set of acquisition parameters which produce images that are favored by clinicians comprise a 1D manifold, allowing for a real-time optimization to maximize image quality. We demonstrate our method for acquisition parameter autotuning on several live patients, showing that our system can start with a poor initial set of parameters and automatically optimize the parameters to produce high quality images.

1 Introduction

Ultrasound imaging requires the adjustment of multiple parameters, e.g. the depth, focus, the frequency and the frequency operation mode (general or Tissue Harmonics Imaging (THI)). The correct choice of parameters has a great impact on the quality of the output image and, in practice, the default parameters recommended by the manufacturer do not always produce a good quality image. The acquisition of a good quality image is a very challenging task especially for difficult patients who have large body habitus. In particular, an abdominal scan involves multiple organs at different depths, is strongly affected by body habitus and requires significant manual tuning (20-45 minutes on average). Previous efforts to

* Leo Grady is currently the VP of research and development at HeartFlow Inc.

produce a good quality image have focused on the hardware aspect of the acquisition by designing probes that have the potential to provide better images [1,2,3]. For example, a curvilinear probe enables larger tissue penetration at the expense of the anatomic image resolution, while a linear array probe provides fine details but can only scan superficial structures. Other hardware solutions include introducing new materials to the sensors used in the transducer [4] and adaptive beamforming with its variations [5,6,7]. Image analysis approaches focus on postprocessing to enhance the image after the acquisition is complete (e.g. [8]). Postprocessing methods transform the acquired data rather than improve the acquisition.

We revisit the problem of autotuning the acquisition parameters by presenting a novel software-only approach based on image analytics and inspired by the autofocus system in a conventional digital camera. However, tuning the acquisition parameters of an ultrasound device is significantly more challenging than conventional camera autofocus due to the larger number of parameters and the challenge in measuring the quality of an ultrasound image. In contrast, autofocus in a digital camera optimizes a simple measure of image sharpness over just one parameter (focal length). The key contribution of our work is to learn a low-dimensional manifold on which lie all acquisition parameters that result in sonographer preferred images. We then train a machine learning system to model the image quality assessment given by experts and show how to efficiently optimize the image quality over the low-dimensional manifold of sonographer preferred acquisition parameters.

2 Methods

Our method for automatic tuning the ultrasound acquisition parameters is inspired by the autofocus in a digital camera. The autofocus in digital cameras works by optimizing the focal length to obtain the image with the best contrast. In ultrasound acquisition, we aim at developing an *ultrasound autotuning* in a similar manner, except that ultrasound acquisition is significantly more challenging since it requires optimization of several parameters instead of just focal length. The second major challenge is that the quality of the ultrasound image, unlike the optical image captured with the autofocus of a conventional digital camera, cannot be simply assessed by measuring the contrast. In this section we present the formulation of the autotuning problem for ultrasound and we will keep the parallel analogy to the autofocus in digital cameras to enhance the exposition of our solution.

Let the configuration of ultrasound parameters consisting of the depth, focus, frequency and operation mode (THI or GEN) be denoted as x and the image acquired with x be denoted as $I(x)$. Assume that the quality of the image can be represented by a function $\mathcal{Q}(I(x))$. The autotuning problem is described as

$$\max_x \mathcal{Q}(I(x)). \quad (1)$$

In the autofocus for digital cameras, x is simply the focal length and $\mathcal{Q}(I(x))$ is the image contrast. In ultrasound autotuning, there are two challenges that we

must address in the paper: First, the development of a quality measure $\mathcal{Q}(I(x))$ for the ultrasound image. Second, the solution of (1), i.e., finding the optimal parameter configuration x that provides an image $I(x)$ with the maximum quality.

2.1 Image Quality Assessment

The assessment of ultrasound image quality is a perceptual characteristic that is difficult to model with an explicit formula, since it depends on several factors such as brightness, sharpness, contrast, resolution, and whether the organ of interest is in focus or not. In the absence of an explicit formula for $\mathcal{Q}(I(x))$, we propose to sample a range of images $I(x)$ and learn the $\mathcal{Q}(I(x))$ mapping for perceptual quality. We train a Support Vector Machine (SVM) regressor based on a set of biologically-inspired features [9,10]. The feature extraction scheme uses a hierarchical approach that consists of four layers of computational units, building an increasingly complex and invariant feature representation by alternating between simple S layers and complex C layers. We have chosen this hierarchical model as it emulates the object recognition in the human visual cortex. The training images were collected from 9 different subjects and we tested on 4 different subjects that were never scanned in the training phase. The data set consists of abdominal scans of seven different organs for each subject: aorta, liver, right kidney, left kidney, pancreas, spleen and gall bladder. A total of 192 images were used in the training. A sonographer provided a grade for each image. The convention used for the grading as suggested by the expert clinician is as follows: Grades 1-6 are given to a poor quality image that cannot be used for diagnosis and treatment. Grades 7-8 are given to minimally acceptable images. Grades 9-10 are given to images for which no further improvement is possible. The variability between 9-10 reflects the variability in an expert's preference. For each input image, we used the feature extraction in [9] and calculated a total of 4075 features. We have performed a Sequential Minimum Optimization (SMO) regression with a normalized polynomial kernel and a reduced feature set that has 10 features.

2.2 Optimization

The second challenge in designing the *ultrasound autotuning* is the optimization of the image quality or choosing the parameter configuration that produces the best quality image. For digital cameras, a gradient search is sufficient to solve the problem because it is a 1D search for the optimal focal length and hence can be done efficiently. However, ultrasound autotuning is more challenging than auto-focusing a digital camera because ultrasound autotuning requires optimization over several parameters. A naive solution would be to do a grid search for the parameter configuration that optimizes the image quality. However, this is very computationally expensive and cannot be performed in real-time acquisition systems. A key insight of our paper is that the known relationship of the acquisition parameters can be exploited to perform a search over a lower-dimensional space of virtual parameters. As an example of this relationship between acquisition parameters, the physics of ultrasound dictates that a deeper focal depth should

require a lower frequency. To perform a dimensionality reduction on our space of acquisition parameters, we employ manifold learning. Specifically, we applied manifold learning to determine if a lower dimensional manifold contains all configurations of acquisition parameters that produce large $\mathcal{Q}(I(X))$. Training data were collected from 9 different subjects. For each subject, 7 different organs were scanned. We obtained a total of 32 “good” configurations that produce images with grade 9 or 10 as judged by an expert. We applied *diffusion maps* manifold learning [11] on the 32 configurations to learn the intrinsic dimensionality of the acquisition parameters.

Although we suspected that the relationship between parameters would lead to a lower-dimensional manifold, we were surprised to find that **the manifold of configuration parameters leading to a good image is one-dimensional**, the manifold is depicted in Figure 1(a). This one-dimensional manifold of good parameters means that any good quality ultrasound image can be determined by optimizing a single virtual parameter. Consequently, we can perform the parameter optimization very quickly by projecting the input parameter configuration to the manifold of good parameter configurations (using k -nearest neighbors) and then optimizing the configuration parameter by a simple gradient ascent (relative to $\mathcal{Q}(I(X))$) along the manifold surface. The algorithm is shown in Algorithm 1.

3 Experimental Results

The objective of our experiments is two fold: First, to test the quality score produced by our system against the quality score assigned by an expert. Second,

Input: Default acquisition parameters x and the learned manifold pairs (x, y) . y is the representation of x on the learned 1D manifold

Output: Parameter configuration that produces the best quality image

INITIALIZE x_i to x and calculate $\mathcal{Q}(I(x_i))$

while $\mathcal{Q}(I(x_{i+1})) > \mathcal{Q}(I(x_i))$ **do**

1. Project the set of parameters x_i to the lower dimensional manifold using an interpolation of the k NN with $k = 5$, to obtain lower dimensional configuration y .
2. Find y_m the closest point to y on the manifold.
3. Take a small step t along the manifold to obtain the new low-d parameters y_{i+1} .
4. From the database of pairings (x, y) , obtain the back projection x_{i+1} that corresponds to the adjusted low-d parameters y_{i+1} . x_{i+1} is the new set of parameters in the original parameter space.
5. Acquire a new image $I(x_{i+1})$ and calculate $\mathcal{Q}(I(x_{i+1}))$

end

if $\mathcal{Q}(I(x_{i+1})) < \mathcal{Q}(I(x_i))$ *and direction of movement has never been changed* **then**
 | Change the movement direction and GOTO 1.

else

| Terminate

end

Algorithm 1. Steps for automated tuning of ultrasound acquisition




Fig. 1. Left: Manifold of parameters producing good images. Different colors indicate different organs. Right: Correlation between the expert’s grade and our system.

to test, on live patients, whether our ultrasound autotuning system succeeds to provide a parameter configuration that produces a good quality image. Our experiments are performed using a Siemens S2000 Ultrasound scanner.

3.1 Quality Assessment Evaluation

For the testing of the quality assessment, we have used 280 images that were not used in training. Figure 1(b) depicts a scatter plot of the grades given by the expert clinician and the grades obtained by our system for the training data set. The correlation coefficient on the ratings was 0.82 and the average error on the training set was 1.3 (in units of 1-10 scale).

3.2 Live Ultrasound Acquisition with Our Autotuning System

We tested our autotuning system with a live ultrasound acquisition where the expert clinician starts with default preset parameters that usually yields bad/moderate image quality and our system generates the new set of parameters to acquire a second image. The clinician is asked to rate the output image and evaluate whether or not this image can be used in the clinic. If the expert clinician rates the image acquired with the new parameters as a bad image, we re-adjust the parameters based on a small step along the manifold shown in Figure 1 (a) and we repeat this process until a good quality image is reached or no further improvement can be made. We chose to run the testing on abdomen scans because each of the seven organs is scanned with a different set of parameters. Manually changing the parameters when moving from one organ to another is very tedious and requires 20–45 minutes per scan. Replacing the manual acquisition tuning with an automatic tuning would provide a great benefit to the workflow.

Figure 2 shows the detailed steps of one of our experiments. In this experiment, we aim at scanning the aorta. We started from the abdomen default set of parameters. The default parameters generated a poor quality image as judged by the expert (grade = 2). Our system applied Algorithm 1 to autotune the parameters until termination. In Figure 2 the top left corner of each image has

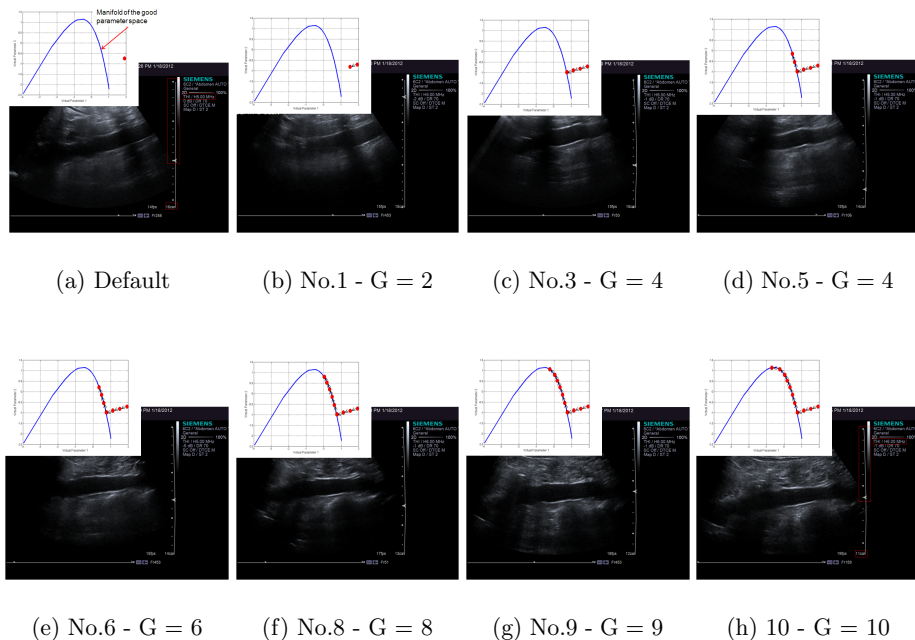


Fig. 2. Autotuning of the acquisition parameters for an aorta scan. The first image is acquired with the Siemens abdomen default preset (Frequency = THI/H 5MHz, depth = 16cm and focus = 10cm). The last image is acquired using the auto tuned parameters (Frequency = THI/ H 6MHz, depth = 11cm and focus = 6cm).

Note: We are only showing 7 iterations out of 10 due to space limitation.

a schematic diagram that illustrates the idea of the parameter adjustment along the manifold, each red circle represents the projection of the acquisition parameters to the 1D manifold and the arrows represent the movement towards/along the manifold of the good parameter space depicted by the blue curve in the figure. The caption of each subfigure shows the iteration number and the grade (G) given by the expert.

Figure 3 shows a sample of our results for different subjects and different organs. The examples of this figure reveal different aspects of strength for our autotuning algorithm. These images are acquired from three different subjects with varying body mass indices (the subject in the first row has the lowest body mass index and the subject in the third row has the highest body mass index) which reflects that our algorithm works equally well for these challenging image acquisition scenarios. The first example in the figure tests whether the algorithm is capable of providing a good set of parameters if the starting set deviates from the manufacturer recommendation. The starting set of parameters were recommended by the clinician. The clinician’s recommendation is generally based on the gender, race and body mass index of the patient. We have performed 5 acquisitions with a clinician’s recommended parameter initializations and in all




Fig. 3. Sample results for the autotuning system. Row 1: Left kidney. Row 2: Liver.




Fig. 4. Quantitative Assessment. Results for 16 different scans spanning four different patients and seven different organs.

cases our system autotunes the parameters to produce a good quality image as judged by the expert. The kidney image obtained by the autotuned parameters was graded 10 by the expert clinician. The example in the second row shows the scan of a fatty liver which is very challenging as it requires very deep penetration. The figure depicts that our algorithm was capable of producing the correct frequency, depth and focus for imaging.

Figure 4 shows a summary of the quantitative assessment of our algorithm. The testing scan were performed on four different patients (denoted as P1, P2, P3 and P4) that were not scanned in the training phase. Due to limited scanner time, we could not scan the seven organs for each subject, however, we managed to test 16 different scenarios. In all of our experiments, the algorithm provided a set of parameters that yielded a good quality image rated 9 or 10 by an expert. The summary is given by the bar chart depicted in Figure 4.

4 Conclusion and Future Work

In contrast to previous work on the topic of ultrasound acquisition autotuning which employed hardware solutions, we proposed a novel software-only solution that learns the manifold of the set of ultrasound acquisition parameters which produce high-quality images. Our experiments show the surprising fact that this set lies on a 1D manifold. We demonstrated the excellent performance of our system and its capability to produce a good quality image (100% accuracy) in live ultrasound acquisition experiments. Comparison to hardware-based solutions was not presented due to the lack of access to such hardware systems. The future work will focus on the improvement of the quality assessment performance to reach higher correlation between the expert clinician's grade and the grade obtained by our system. We also plan to run a large scale study that contains more data to validate our manifold.

References

1. Smith, S., Goldberg, R.: Multilayer piezoelectric ceramics for two-dimensional transducers. *IEEE Trans. on Ultrasonics, Ferroelectrics and Frequency Control* (1994)
2. Mooney, M., Wilson, M.: Linear array transducers with improved image quality for vascular ultrasonic imaging. *Hewlett-Packard Journal* (1994)
3. Dhanaliwala, A.H., Hossack, J., Mauldin, W.: Assessing and improving acoustic radiation force image quality using a 1.5-d transducer design. *IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control* (2012)
4. Zhu, J.: A holey-structured metamaterial for acoustic deep-subwavelength imaging. *Nature Physics* (2010)
5. Asl, B.M.: Minimum variance beamforming combined with adaptive coherence weighting applied to medical ultrasound imaging. *IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control* (2009)
6. Xu, M., Chen, Y., Ding, M., Ming, Y.: Adaptive minimum variance beamforming combined with phase coherence imaging for ultrasound imaging. *SPIE Medical Imaging* (2012)

7. Nilsen, C.I., Hafizovic, I.: Beamspace adaptive beamforming for ultrasound imaging. *IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control* (2009)
8. Su, Y., Wang, H., Wang, Y., et al.: Speckle reduction approach for breast ultrasound image and its application to breast cancer diagnosis. *European J. of Radiology* (2010)
9. Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., Poggio, T.: Robust object recognition with cortex-like mechanisms. *IEEE PAMI* 29, 411–426 (2007)
10. Mutch, J., Lowe, D.G.: Object class recognition and localization using sparse features with limited receptive fields. *IJCV* 80(1), 45–57 (2008)
11. Coifman, R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 5–30 (2006)