

Gaussian process regression for ultrasound scanline interpolation

Alperen Degirmenci[Ⓢ],^a Robert D. Howe[Ⓢ],^a and Douglas P. Perrin[Ⓢ]^{b,c,*}

^aHarvard University, John A. Paulson School of Engineering and Applied Sciences, Cambridge, Massachusetts, United States

^bBoston Children's Hospital, Boston, Massachusetts, United States

^cHarvard Medical School, Boston, Massachusetts, United States

Abstract

Purpose: In ultrasound imaging, interpolation is a key step in converting scanline data to brightness-mode (B-mode) images. Conventional methods, such as bilinear interpolation, do not fully capture the spatial dependence between data points, which leads to deviations from the underlying probability distribution at the interpolation points.

Approach: We propose Gaussian process (\mathcal{GP}) regression as an improved method for ultrasound scanline interpolation. Using ultrasound scanlines acquired from two different ultrasound scanners during *in vivo* trials, we compare the scanline conversion accuracy of three standard interpolation methods with that of \mathcal{GP} regression, measuring the peak signal-to-noise ratio (PSNR) and mean absolute error (MAE) for each method.

Results: The PSNR and MAE scores show that \mathcal{GP} regression leads to more accurate scanline conversion compared to the nearest neighbor, bilinear, and cubic spline interpolation methods, for both datasets. Furthermore, limiting the interpolation window size of \mathcal{GP} regression to 15 reduces computation time with minimal to no reduction in PSNR.

Conclusions: \mathcal{GP} regression quantitatively leads to more accurate scanline conversion and provides uncertainty estimates at each of the interpolation points. Our windowing method reduces the computational cost of using \mathcal{GP} regression for scanline conversion.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JMI.9.3.037001](https://doi.org/10.1117/1.JMI.9.3.037001)]

Keywords: ultrasound; interpolation; Gaussian process; machine learning.

Paper 21297GR received Nov. 5, 2021; accepted for publication Apr. 26, 2022; published online May 17, 2022.

1 Introduction

Ultrasound imaging is one of the safest, cheapest, and fastest medical imaging modalities. It is widely used in many clinical applications, such as image-guided interventions, as a diagnostic as well as a therapeutic tool. Ultrasound image resolution and accuracy are important factors that can impact the outcome of such procedures.

The ultrasound transducer transmits and receives radio frequency data along scanlines. This scanline data are then processed, undergoing operations such as time gain compensation, low-pass filtering, envelope detection, and log compression.¹ Finally, the processed scanlines (A-mode) are interpolated such that a dense two-dimensional (2D) image [brightness-mode (B-mode)] can be displayed to the clinician.

In a probabilistic framework, the data along the scanlines can be considered as samples drawn from an underlying probability distribution, and then the goal of the interpolation step during scan conversion is to approximate from this underlying distribution the values of pixels for which we do not have a measurement. Commonly used interpolation methods, such as bilinear interpolation, only use local information from a few neighboring observations, which can

*Address all correspondence to Douglas P. Perrin, Douglas.Perrin@cardio.chboston.org

lead to inaccurate estimates. Methods that take into account more observations can better infer the underlying distribution, leading to more accurate interpolation.²

In this paper, we use Gaussian process (\mathcal{GP}) regression to improve the accuracy of ultrasound scanline conversion. In \mathcal{GP} regression, a Gaussian distribution function is fit to each observation point. A covariance function is used to compute the spatial correlation between observations. The parameters of this function can be tuned to capture different characteristic length scales that are specific to each dimension of the data. This adaptability is useful in ultrasound imaging, since observations along a scanline are dense, whereas the number of scanlines is small in comparison.

1.1 Previous Work on Improving Ultrasound Image Quality

Most work on ultrasound image enhancement has focused on speckle reduction^{3,4} and denoising.⁵ An anisotropic spatiotemporal smoothing filter was developed⁶ to improve scan conversion accuracy. Spline interpolation for three-dimensional (3D) ultrasound volume compounding was explored in Ref. 7. Kriging (another name for \mathcal{GP} regression in geostatistics) was used for 3D ultrasound reconstruction from multiple scan converted 2D slices in Ref. 8. A random field approach was used in Ref. 9 to construct confidence maps for ultrasound images, which are then used to improve 3D ultrasound interpolation accuracy. A comprehensive review of interpolation for 3D ultrasound compounding can be found in Ref. 10.

Recent work has applied deep-learning methods to ultrasound imaging,¹¹⁻¹³ such as dynamic beam forming, adaptive spectral Doppler processing, improved noise suppression, and super-resolution. However, these techniques are yet far from clinical adoption, since acquiring high-quality ground truth data at massive scales required for model training and ensuring generalization to unseen anatomy remains a challenge in developing and validating such deep-learning-based solutions.

\mathcal{GP} regression was previously used in other imaging modalities: single-image super-resolution was explored¹⁴ for natural images¹⁵; explored magnetic resonance imaging (MRI)¹⁶ used the variance estimates from \mathcal{GP} regression to improve MRI image registration. To the best of our knowledge, \mathcal{GP} regression has previously not been applied to ultrasound scanline interpolation.

2 Methods

Let \mathbf{x} represent the spatial positions of our observations \mathbf{y} along the scanlines, and $f(\mathbf{x}) = \mathbf{y}$ be a nonlinear function that maps \mathbf{x} to \mathbf{y} . The core assumption in our work is that f can be approximated by a Gaussian process, allowing us to perform scanline interpolation by querying the process at the desired interpolation coordinates, \mathbf{x}_* .

2.1 Gaussian Process Regression

A \mathcal{GP} is a collection of random variables, any finite number of which have a joint Gaussian distribution.¹⁷ A \mathcal{GP} is completely specified by its mean function, $m(\mathbf{x})$, and covariance function, $k(\mathbf{x}, \mathbf{x}')$, represented as $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. In practice, a zero-mean function is commonly used when we have limited knowledge of f .¹⁸ Observations are considered as random variables that are drawn from a multivariate normal distribution

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K} + \sigma_n^2 \mathbf{I} & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_{**} \end{bmatrix}\right), \quad (1)$$

where \mathbf{y} is a vector of observed scanline intensities, \mathbf{y}_* is a vector of intensities we seek to estimate at the interpolation locations, and σ_n^2 is the observation noise. The covariance matrix is a positive semi-definite symmetric matrix, where $\mathbf{K} = k(\mathbf{x}, \mathbf{x})$ is the covariance between observations, $\mathbf{K}_* = k(\mathbf{x}, \mathbf{x}_*)$ is the covariance between observations and interpolation points, and $\mathbf{K}_{**} = k(\mathbf{x}_*, \mathbf{x}_*)$ is the covariance between interpolation points. The details of the covariance function will be discussed in Sec. 2.3.

The intensity vector \mathbf{y}_* is estimated at the interpolation locations \mathbf{x}_* as

$$\bar{\mathbf{y}}_* = \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}. \quad (2)$$

Concerns regarding the computational costs of evaluating Eq. (2), namely, inverting the covariance matrix, are addressed in Sec. 2.6.

The variance of the regressed values can also be computed, based on the distance between the observations and the interpolation locations, using

$$\mathbb{V}[\mathbf{y}_*] = \mathbf{K}_{**} - \mathbf{K}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{K}_*. \quad (3)$$

When computing Eqs. (2) and (3), we standardize the inputs to account for the zero-mean function used in Eq. (1), as well as to avoid numerical instabilities and rounding errors in floating-point arithmetic operations.

2.2 Scanline Conversion—Parallel and Diverging Scanlines

In ultrasound transducers, the scanlines can be either parallel to each other (e.g., from a linear-array probe), or diverging, spreading out radially from the source (e.g., from a phased-array or curvilinear probe). With parallel scanlines, the scanline data form a grid in Cartesian coordinates, and it is straightforward to interpolate and display them.

In contrast, with diverging scanlines, the scanline data form a grid in polar coordinates, but the clinician is often interested in the Cartesian space representation for ease of spatial reasoning. Figure 1 shows the arrangement of scanlines in polar and Cartesian coordinates. The desired




Fig. 1 Illustration of a simplified ultrasound image acquisition pipeline highlighting the differences in the scan conversion step for linear and diverging ultrasound scanlines. *In vivo* scanline data are shown to demonstrate the spatial effects of scan conversion. (Top) The posterior side of a human lower leg was scanned using an ultrasound probe with 256 parallel scanlines and 793 points along each scanline. The gastrocnemius and soleus muscles are clearly visible in the image. The scanlines and the interpolation points (colored) for generating the B-mode image lie on a grid. (Bottom) An ultrasound image of a human aortic valve was acquired using a 3D TEE probe with 64×48 diverging scanlines, with 412 points along each scanline. Here, we show a single 2D slice out of the 48. The scanlines lie on a grid in polar coordinates, however the interpolation points (colored) for generating the B-mode image lie on a grid in Cartesian coordinates. Mapping these interpolation points to polar coordinates results in a non-uniform spacing between interpolation points. Scan conversion is often performed in polar coordinates, then the interpolated points are mapped back into Cartesian coordinates.

interpolation points (colored) lie on the image grid in Cartesian space. It is noteworthy that with diverging scanlines, the density of observations in Cartesian space gets sparser as the distance from the ultrasound transducer (r) increases. In practice, when using linear interpolation methods, scan conversion for diverging scanlines is performed in polar coordinates, where the observation points lie on a grid. The interpolated intensity values are then simply displayed in Cartesian coordinates. However, this method ignores the increasing distance between observations due to the mapping between polar and Cartesian spaces, which can be problematic especially when the beam spacing is sparse, or if the clinician wants to zoom in on regions that are further away from the probe. To achieve better accuracy, the interpolation method should account for the change in inter-beam distance, capturing the underlying process better. In Sec. 2.3, we discuss how this can be achieved by selecting a different kernel length per dimension in the \mathcal{GP} covariance function.

2.3 Covariance Functions

The covariance function (i.e., kernel) captures the spatial dependence between observations. There are various kernels that can be used to compute the covariance matrix, and a comprehensive list can be found in Ref. 17. In this work, we consider two covariance functions: the squared exponential (SE) and the Matérn.

Kernels usually have a parameter that controls the characteristic length scale of the dependence between observations. It is possible to set a different length scale for each dimension of the data, which is useful when dealing with a non-isometric sampling density, such as in the case of ultrasound imaging with diverging scanlines.

2.3.1 Squared exponential kernel

The SE covariance function, also known as the radial basis function kernel, is an infinitely differentiable function, which results in a smooth process. The SE covariance function with separate length scales for each dimension is expressed as

$$k(\mathbf{x}_i, \mathbf{x}_j | \theta) = \sigma^2 \exp \left[- \sum_{n=1}^2 \frac{(x_{i,n} - x_{j,n})^2}{2l_n^2} \right], \quad (4)$$

where θ is the function parameters (σ and l_n), σ^2 is the signal variance, l_n is the characteristic length scale for dimension n , and $x_{i,n}$ is the n 'th component of the i 'th observation.

In this work, we will denote the length scale along the scanlines as l_r , and normal to the scanlines as l_θ .

2.3.2 Matérn kernel

The Matérn autocovariance function has the form

$$k(x_i, x_j | \theta) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{l} r \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{l} r \right), \quad (5)$$

where Γ is the gamma function, K_ν is the modified Bessel function of the second kind, r is the distance between x_i and x_j , l is the characteristic length scale, and ν is the smoothness parameter. In D dimensions, the distance metric is expressed as

$$r(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{\sum_{n=0}^D \frac{(x_{i,n} - x_{j,n})^2}{l_n^2}}. \quad (6)$$

According to Ref. 19, the strong smoothness assumptions of the SE covariance function are unrealistic for modeling many physical processes, and the Matérn class of covariance functions

is recommended instead.¹⁷ The Matérn covariance function is $\lceil \nu \rceil - 1$ times differentiable. In general, ν is chosen to be a half-integer, which reduces Eq. (5) to an exponential multiplied by a polynomial, and the spectral density becomes rational.¹⁹ When $\nu = 1/2$, the Matérn covariance function reduces to the exponential covariance function, and as $\nu \rightarrow \infty$ the kernel converges to the SE covariance function. In the literature, ν is usually set to $3/2$ or $5/2$.¹⁷

We tested both the SE kernel and the Matérn covariance function with $\nu = 3/2$ and $5/2$ in our experiments. The Matérn covariance function with $\nu = 3/2$ yielded the most accurate interpolation, therefore all results reported in Sec. 3 are computed using this kernel.

2.4 Ground Truth Data and Measuring Interpolation Accuracy

Quantifying improvement in image quality is a challenging task in the absence of ground truth data. In this study, the original ultrasound scanline data are treated as the ground truth. The interpolation methods outlined are then only applied to a subset of the scanlines, and the resulting estimates are then compared with the excluded scanlines to measure the interpolation accuracy. We will refer to this as leave- N -out in Sec. 3. We compare the performance of commonly used interpolation methods (nearest neighbor, bilinear, and cubic spline interpolation) against the performance of \mathcal{GP} regression, where the mean absolute error (MAE) and the peak signal-to-noise ratio (PSNR) between the ground truth and the interpolated images are measured.

The MAE is defined as

$$\text{MAE} = \frac{1}{I} \sum_i^I |Z_i^* - \hat{Z}_i|, \quad (7)$$

where I is the total number of pixels in the ultrasound image, Z_i^* is the pixel intensities in the ground truth image, and \hat{Z}_i is the pixel intensities in the interpolated image.

Assuming that pixel intensities are normalized to $\mathbb{R}: [0,1]$, the PSNR is defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{1}{\text{MSE}} \right), \quad (8)$$

where

$$\text{MSE} = \frac{1}{I} \sum_i^I (Z_i^* - \hat{Z}_i)^2. \quad (9)$$

The best interpolation method should ideally have the highest PSNR and lowest MAE and MSE among the interpolation methods tested.

2.5 Optimizing the Kernel Length Scale

In general, the kernel parameters for a \mathcal{GP} are determined by minimizing the log marginal likelihood,

$$\log p(y|X) = -\frac{1}{2} \mathbf{y}^T (K + \sigma_n^2 I)^{-1} \mathbf{y} - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{N}{2} \log 2\pi, \quad (10)$$

where N is the number of observations. This is equivalent to minimizing the mean squared error (MSE). We optimize the kernel parameters through a constrained optimization problem with the MSE metric as the cost function, described as

$$\min_{l_n} \frac{1}{I} \sum_i^I (Z_i^* - \hat{Z}_i^{\mathcal{GP}})^2 \quad \text{s.t. } l_n > 0, \quad n = \alpha, r, \quad (11)$$




Fig. 2 Incorrect choice of the kernel length can lead to reconstruction errors, such as missing data (shorter than optimal) or blurring (longer than optimal). Ground truth shown on the right uses all 64 scanlines, whereas the reconstructions on the left use 32 (top row) and 8 (bottom row) scanlines. When the kernel length is too short, any regression point that lies too far from an observation quickly approaches the mean of the \mathcal{GP} .

where Z_i^* is the ground truth pixel intensity and $\hat{Z}_i^{\mathcal{GP}}$ is the \mathcal{GP} estimate. In our optimization studies, we noticed that the landscape is convex, however, the gradients near the optimum are small.

Figure 2 shows how choosing a larger than optimal length scale results in blurring, while shorter length scales lead to striated mean-valued regions between the scanlines. Using a sub-optimal length scale results in a decrease in the PSNR.

2.6 Reducing Computational Cost—Patched \mathcal{GP} Regression

\mathcal{GP} regression has $O(N^3)$ memory and $O(N^2)$ time complexity, where N is the number of observations,¹⁷ which is perhaps the main reason why it has not been widely adopted in clinical imaging. Recent work²⁰ has pushed exact \mathcal{GP} training to over one million data points using multi-GPU parallelization and methods such as linear conjugate gradients. However, medical images can contain tens of millions of pixels or voxels and medical applications generally require real-time performance.

The main bottleneck in \mathcal{GP} regression is in inverting the covariance matrix in Eq. (2). The covariance function gives a measure of the spatial dependence between observations. It is noteworthy that with certain covariance functions (such as the ones used in this work), the covariance function quickly decays as the distance between observation pairs increases. Therefore, outside of a local region, the contributions from other observations can practically be ignored without a loss in regression accuracy. We exploit this behavior to reduce the computational demand of \mathcal{GP} regression. The scanline data are subdivided into overlapping patches and \mathcal{GP} regression is evaluated separately for each patch.

Experimenting with a range of window sizes, we determined that a window size of 15 observations is optimal. Figure 3 shows the percent increase in computation time and PSNR as the window size is increased from 5 to 29. Bold black line indicates the 15 observation window size we selected. The colored lines represent different levels of decimation of the scanlines. This plot is for the phased-array dataset. We can see that increasing the window size beyond 15 yields negligible or no increase in PSNR at the cost of increased computation time. It is noteworthy that the windowed \mathcal{GP} regression can be parallelized on a GPU or field programmable gate arrays (FPGA) to process the entire image in a single pass.

Furthermore, the computation time and costs can be significantly reduced by pre-computing the matrix $\mathbf{K}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}$ in Eq. (2) and \mathbf{K}_{**} in Eq. (3). For a given probe and imaging depth, the observation and interpolation locations are not going to change, therefore, these matrices can be pre-computed to avoid performing the costly Cholesky decomposition with each regression. This reduces the computational complexity to just $O(NM)$ for the matrix-vector multiplication in Eq. (2), where $M < N$ is the number of interpolation locations.




Fig. 3 Results of the window size study on the phased-array scanline data, showing the percent increase in computation time and PSNR as the window size is increased from 5 to 29. Bold black line shows the 15 observations we selected. The colored lines represent different levels of decimation of the scanlines. The computation times reported here are the mean of 10 measurements per window size.

3 Experiments and Results

In vivo scanline data from two different ultrasound imaging systems were acquired. A linear-array ultrasound probe with parallel scanlines (LV8-4L65S-3, TELEMED Ltd., Vilnius, Lithuania) was used to acquire 2D images of the lower leg of a healthy male subject at a 61-MHz imaging depth and 6-MHz frequency. The 256 scanlines with 793 observations along each scanline are shown in Fig. 1(top). In another experiment, 3D scanline data of a human heart was acquired at 63-mm imaging depth and 7-MHz frequency using a phased-array 3D transesophageal echocardiography (TEE) probe (X7-2t transducer connected to an iE33 ultrasound imaging system, Philips Healthcare, Andover, Massachusetts, United States). A 2D slice (out of 48) from the 3D dataset is shown in Fig. 1(bottom), containing 64 scanlines and 412 observations along each scanline. Both studies were approved by the Harvard University Institutional Review Board (IRB) and the Boston Children’s Hospital IRB.

The observation noise σ_n^2 for \mathcal{GP} regression was set to 2×10^{-3} for the linear probe, and 8×10^{-3} for the phased-array probe through hyperparameter optimization. Computations used the image processing, parallel computing, and the statistics and machine learning toolboxes in MATLAB (The MathWorks, Inc., Natick, Massachusetts, United States).

3.1 Parallel Scanlines

Figure 4 shows the interpolated B-mode images for three different leave- N -out studies. Leftmost column shows the ground truth data. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (3), where darker regions indicate smaller variance (i.e., higher confidence). The normalization coefficient was determined across the entire image set.

It is worthy to note that the amount of uncertainty increases as the number of scanlines is reduced. Figure 4(a) shows the absolute PSNR scores, Fig. 4(b) shows the relative PSNR scores (computed by subtracting the lowest scores from the rest), and Fig. 4(c) shows the MAE scores. Nearest neighbor interpolation has the largest error as expected, and \mathcal{GP} regression has the lowest error and highest PSNR.

As the number of observations (scanlines) is reduced, we see a wider separation in performance between \mathcal{GP} regression and the other interpolation methods. Looking at the relative PSNR scores, we see that the scores for the bilinear and spline methods have a negative slope that tend




Fig. 4 (Top) Scan converted B-mode images of the parallel scanline dataset. Leftmost column shows the ground truth data with 256 scanlines. N indicates the number of scanlines used for interpolation, and the next four columns show the resulting B-mode images using four interpolation methods (nearest, bilinear, spline, and \mathcal{GP}). Rightmost column shows the normalized variance matrix generated during \mathcal{GP} regression, where darker regions indicate higher confidence in the estimates. PSNR scores are overlaid on each image. (Bottom) Interpolator performance for scan conversion of parallel scanlines: (a) absolute PSNR scores (higher is better), (b) relative PSNR scores (computed by subtracting the lowest scores from the rest, higher is better), and (c) MAE scores (lower is better).

toward the performance of the nearest neighbor interpolation. However, \mathcal{GP} regression maintains a constant separation in metrics from the baseline method (nearest neighbor).

3.2 Diverging Scanlines

Figure 5 shows the interpolated B-mode images for three different leave- N -out studies in polar coordinates. Leftmost column shows the ground truth data. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (3). The variance maps shown in Fig. 5 indicate that the amount of uncertainty increases as the number of scanlines is reduced. The variance also increases along the scanlines due to the $1/r$ scaling of the l_α parameter. Figure 5(a) shows the absolute PSNR scores, Fig. 5(b) shows the relative PSNR scores (computed by subtracting the lowest scores from the rest), and Fig. 5(c) shows the MAE scores. Nearest neighbor interpolation has the largest error as expected, and \mathcal{GP} regression has the lowest error and highest PSNR. \mathcal{GP} regression exhibits a larger increase in performance for diverging scanlines compared to the linear scanlines.

We observe an even wider separation (compared to parallel scanline results) in performance between \mathcal{GP} regression and the other interpolation methods as the number of observations (scanlines) is reduced. Looking at the relative PSNR scores, we see that the scores for the bilinear and spline methods again have a negative slope that tend toward the performance of the nearest neighbor interpolation. Differently from the parallel case, the relative PSNR score for \mathcal{GP} regression has a peak around 13 scanlines, then rapidly decreases near the performance of bilinear interpolation. However, we should notice that this occurs when the number of scanlines is <15 , which is the window size we determined previously when analyzing the \mathcal{GP} kernel range




Fig. 5 (Top) Scan converted B-mode images of the diverging scanline dataset, presented in polar coordinates. In the top three rows, leftmost column shows the ground truth data with 64 scanlines. N indicates the number of scanlines used for interpolation, and the next four columns show the resulting B-mode images using four interpolation methods (nearest, bilinear, spline, and \mathcal{GP}). PSNR scores are overlaid on each image. Rightmost column shows the normalized variance matrix generated during \mathcal{GP} regression. Bottom three rows show a detailed view of the red and blue regions of interest. (Bottom) Interpolator performance for scan conversion of diverging scanlines: (a) absolute PSNR scores (higher is better), (b) relative PSNR scores (computed by subtracting the lowest scores from the rest, higher is better), and (c) MAE scores (lower is better).

of influence. This suggests that we simply do not have enough data points to accurately infer the underlying distribution of the data, thus \mathcal{GP} regression performance drops rapidly.

Figure 6 shows the interpolated B-mode images for three different leave- N -out studies in Cartesian coordinates. Leftmost column shows the baseline data, which was generated by converting the scanlines from polar coordinates to Cartesian coordinates using bilinear interpolation, and all 64 of the available scanlines were used in the scan conversion. We also tested cubic interpolation and \mathcal{GP} regression to generate the baseline, which resulted in similar PSNR and MAE metrics. The variance map on the right shows the normalized values of the variance reported by \mathcal{GP} regression from Eq. (3). We can see that the amount of uncertainty increases as the number of scanlines is reduced. The variance also increases along the scanlines due to the distance-dependent ($1/r$) scaling of the l_α parameter. At higher decimation rates with $N = 13$ and $N = 8$, the B-mode images generated by \mathcal{GP} regression are visibly clearer compared to the other interpolation methods. Bilinear interpolation results exhibit the laterally elongated features that is characteristic to ultrasound images seen in commercial ultrasound systems.

3.3 Relationship between Scanline Properties and Optimal Kernel Lengths

Optimal kernel lengths were calculated by minimizing the problem defined in Eq. (11). This function had to be minimized for each different number of scanlines used, which is a




Fig. 6 Scan converted B-mode images of the diverging scanline dataset, presented in Cartesian coordinates. Leftmost column shows the baseline image with 64 scanlines. PSNR scores are overlaid on each image. Rightmost column shows the normalized variance map generated during \mathcal{GP} regression.

computationally expensive process. However, in this particular application, since observations lie on a fixed grid, it should be possible to determine a relationship between the grid parameters and the optimal kernel lengths. This relationship can be then used to directly calculate the kernel length for a given grid, without having to run an expensive optimization cycle.

In the case of parallel scanlines, the distance between observations along the scanlines, Δr , and the distance between the scanlines, $|r\Delta\alpha|$, are constants that are determined by the ultrasound probe and imaging settings.

In the case of diverging scanlines, the regression variance along the rays r should be constant, since the distance between observations is constant. However, as the scanlines diverge, since there are less observations per unit area, the variance across scanlines (along the α axis) should increase. This can be achieved by scaling l_α by $1/r$, which reduces the characteristic length scale in the radial axis, thus, increasing the uncertainty of the pixels in between the scanlines.

Inspecting of the optimized kernel lengths, we noticed the following relationships between the kernel lengths and the physical parameters of the ultrasound scanlines:

1. For parallel scanlines,

$$l_\alpha = \frac{1}{\sqrt{3}} \Delta\alpha, \quad (12)$$

$$l_r = k_N \Delta r, \quad (13)$$

where $\Delta\alpha$ and Δr are the (standardized) distances between observations (i.e., voxel size), and k_N is the reduction factor for leave- N -out studies. For example, if only $N/4$ scanlines are used in the interpolation, then $k_N = 4$.

2. For diverging scanlines,

$$l_{\alpha|\bar{r}_i} = \frac{1}{\sqrt{3}} \frac{|\bar{r}|}{2\bar{r}_i} \Delta\alpha, \quad (14)$$

$$l_r = k_N \Delta r, \quad (15)$$

where $|\bar{r}|$ is the overall (standardized) scanline length, and \bar{r}_i is the depth of the i 'th observation along the scanline.

It is worthy to note that we have not yet tested if these kernel length scale relationships hold true across different anatomy, scan depths, and ultrasound probes.

4 Discussion

In this paper, we investigated the use of \mathcal{GP} regression for ultrasound scanline interpolation. We analyzed *in vivo* ultrasound data acquired using both linear and phased-array probes to validate the performance of \mathcal{GP} regression. Using leave- N -out studies, we qualitatively and quantitatively showed that \mathcal{GP} regression leads to better B-mode conversion than other interpolation methods, indicated by the higher PSNR and lower MAE scores. We greatly reduced the computation time through patched computation of the inverse of the covariance function. In our implementation, the average time to estimate the value of one data point was 0.3 ms. It should be possible to enable real-time execution of \mathcal{GP} regression through an optimized implementation. The covariance matrix can also be precomputed, significantly reducing the computational complexity of the problem. Computational time can be further reduced by relaxing the requirements on exact regression, and taking advantage of the Kronecker structure of gridded data.²¹

We demonstrated that \mathcal{GP} regression leads to better scanline conversion using clinical data. We collected ultrasound scanline data from two distinct parts of the anatomy using a linear and phased-array ultrasound probe. The leg dataset exhibits high-frequency texture and has a dense set of scanlines, whereas the cardiac dataset comprises uniform patches of dark (blood) and bright (muscle) regions and has a sparse set of scanlines. The PSNR and MAE scores for \mathcal{GP} regression are better than those for nearest, bilinear, and spline interpolation across both datasets.

Using \mathcal{GP} regression, it may be possible to use fewer scanlines while maintaining the same image quality. This can reduce system complexity or help increase the frame rate by reducing number of transmitted beams.

Beyond improving scan conversion accuracy, our method can be further beneficial in downstream tasks, such as image registration. The uncertainty estimate of \mathcal{GP} regression can be used to improve ultrasound image registration accuracy. This technique was successfully demonstrated for MRI registration.¹⁶

One open question in our study is whether the kernel length relationships that we identified hold for different anatomy and ultrasound probes. We plan to address this in our future work.

Choosing a covariance function prescribes a prior on the spatial variation of the data. Even though the Matérn 3/2 function resulted in lower errors compared to the SE kernel, use of other covariance functions should be investigated as well.

In this work, we applied \mathcal{GP} regression to scanline data that was acquired after log compression to an 8-bit range. However, this compresses and distorts the dynamic range of the ultrasound transducer. We expect performing regression prior to these post-processing steps will lead to better scanline conversion. In combination with tone mapping techniques and high-dynamic range ultrasound imaging methods,²² the ultrasound images presented to the clinician can be further improved.

Disclosures

The authors declare that they have no conflict of interest to disclose.

Acknowledgments

This work was supported by the Harvard University John A. Paulson School of Engineering and Applied Sciences, the National Institutes of Health (Grant No. 1R21EB018938), Toyota North America Inc., and an NVIDIA Academic Hardware Grant.

References

1. M. Ali, D. Magee, and U. Dasgupta, "Signal processing overview of ultrasound systems for medical imaging," Tech. Rep. SPRAB12 - November 2008, Texas Instruments, Dallas, Texas (2008).

2. D. Robinson and P. Knight, "Interpolation scan conversion in pulse-echo ultrasound," *Ultrasonic Imaging* **4**(4), 297–310 (1982).
3. Y. Yu and S. Acton, "Speckle reducing anisotropic diffusion," *IEEE Trans. Image Process.* **11**(11), 1260–1270 (2002).
4. P. Coupé et al., "Nonlocal means-based speckle filtering for ultrasound images," *IEEE Trans. Image Process.* **18**(10), 2221–2229 (2009).
5. Y. Erez, Y. Y. Schechner, and D. Adam, "Ultrasound image denoising by spatially varying frequency compounding," *Lect. Notes Comput. Sci.* **4174**, 1–10 (2006).
6. I. Herlin and N. Ayache, "A new methodology to analyze time sequences of ultrasound images," (RR-1390) (1991).
7. R. Rohling et al., "Radial basis function interpolation for freehand 3D ultrasound," *Lect. Notes Comput. Sci.* **1613**, 478–483 (1999).
8. M. R. Stytz and R. W. Parrott, "Using kriging for 3D medical imaging," *Comput. Med. Imaging Graphics* **17**, 421–442 (1993).
9. A. Karamalis et al., "Ultrasound confidence maps using random walks," *Med. Image Anal.* **16**(6), 1101–1112 (2012).
10. O. V. Solberg et al., "Freehand 3D ultrasound reconstruction algorithms—a review," *Ultrasound Med. Biol.* **33**, 991–1009 (2007).
11. A. C. Luchies and B. C. Byram, "Deep neural networks for ultrasound beamforming," *IEEE Trans. Med. Imaging* **37**(9), 2010–2021 (2018).
12. R. J. G. van Sloun, R. Cohen, and Y. C. Eldar, "Deep learning in ultrasound imaging," *Proc. IEEE* **108**(1), 11–29 (2020).
13. A. Wiacek, E. González, and M. A. L. Bell, "Coherent: a deep learning architecture for ultrasound spatial correlation estimation and coherence-based beamforming," *IEEE Trans. Ultrasonics, Ferroelectr., Freq. Control* **67**(12), 2574–2583 (2020).
14. H. He and W.-C. Siu, "Single image super-resolution using Gaussian process regression," in *CVPR'11: Proc. 2011 IEEE Conf. Comput. Vision and Pattern Recognit.* (2011).
15. H. D. V. Cardona et al., "Gaussian processes for slice-based super-resolution MR images," in *ISVC*, pp. 692–701 (2015).
16. C. Wachinger et al., "Gaussian process interpolation for uncertainty estimation in image registration," *Lect. Notes Comput. Sci.* **8673**, 267–274 (2014).
17. C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, MIT Press, Cambridge, Massachusetts (2006).
18. I. Delbridge, D. Bindel, and A. G. Wilson, "Randomly projected additive Gaussian processes for regression," in *Int. Conf. Mach. Learn.*, PMLR, pp. 2453–2463 (2020).
19. M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer Science & Business Media (2012).
20. K. Wang et al., "Exact Gaussian processes on a million data points," in *Adv. Neural Information Processing Systems 32*, H. Wallach et al., Eds., pp. 14648–14659, Curran Associates, Inc. (2019).
21. A. Wilson and H. Nickisch, "Kernel interpolation for scalable structured Gaussian processes (KISS-GP)," in *Proc. 32nd Int. Conf. Mach. Learn.*, D. Blei and F. Bach, Eds., pp. 1775–1784 (2015).
22. A. Degirmenci, D. P. Perrin, and R. D. Howe, "High dynamic range ultrasound imaging," *Int. J. Comput. Assist. Radiol. Surg.* **13**(5), 721–729 (2018).

Alperen Degirmenci received his PhD in electrical engineering from Harvard University in 2018 with a secondary field in computational science and engineering. He was a graduate student at the Harvard Biorobotics Lab at the time of completion of this work, where his research focused on real-time, high-performance algorithm development for medical ultrasound image processing, and robotic procedure guidance in catheter-based cardiac interventions.

Robert D. Howe is the Abbott and James Lawrence Professor of Engineering at the Harvard School of Engineering and Applied Sciences. He received his doctoral degree in mechanical engineering from Stanford University, then joined the faculty at Harvard in 1990. He directs the Harvard BioRobotics Laboratory. His research interests focus on manipulation, the sense of touch, haptic interfaces, and robot-assisted and image-guided surgery.

Douglas P. Perrin is a senior scientist and principal investigator at the Boston Children's Hospital, directing a research group within the Department of Cardiovascular Surgery, and an instructor of surgery at the Harvard Medical School. He received his PhD in computer science from the University of Minnesota in 2002 and was a postdoctoral researcher at the Harvard School of Engineering and Applied Sciences, where he remains an affiliate. His current research interests include image-guided heart surgery, ultrasound imaging enhancement, and patient-specific models for surgical planning.