

SHADOW BASED PHASE GRADIENT AUTOFOCUS FOR SYNTHETIC APERTURE SONAR

J Prater	Naval Surface Warfare Center, Panama City, Florida, USA
D Bryner	Naval Surface Warfare Center, Panama City, Florida, USA
S Synnes	FFI Norwegian Defence Research Establishment, Norway

1 INTRODUCTION

Synthetic Aperture Sonar (SAS) systems use long synthetic apertures to achieve long-range high-resolution imagery^{1,2}. These long apertures require precise localization of the sonar array in order to maintain array coherence¹⁻³. Motion estimation and compensation techniques are often applied to provide the precise localization, but residual errors that defocus the imagery are common in data when unresolved system motion is present. Autofocus techniques were developed to arbitrarily detect and remove phase errors in data to ensure focused imagery independent of system motion estimation⁴⁻⁶. These autofocus techniques often rely on the presence of bright point-like objects to extract the deviation from the ideal phase response across the aperture to determine the phase shift required to focus the imagery. While these algorithms have successfully been used to resolve any residual errors in SAS imagery, insufficient bright point-like objects in some environments may reduce the quality from point-driven autofocus algorithms.

Previous work where SAS images were constructed from normalized-intensity data⁷ demonstrated that the phase history of shadow areas in imagery did not consist of random phase values and exhibited a more destructive summation in beamforming than the mean background response. The normalized-intensity images were constructed similarly to typical SAS beamforming with the exception that the data is intensity normalized prior to beamforming. An example of normalized intensity imagery can be seen in Figure 1 which shows data prior to beamforming, a standard SAS image, and a normalized-intensity image. From this example, it is apparent that the normalized-intensity imagery contains highlights and shadows similar to the standard image. Under the assumption that the mean background is speckle dominated, the depth of shadows compared to the mean background indicate that summation of the data used to form shadows is closer to zero than would be expected if the phase of the data was random. This assertion is clearly demonstrated in Figure 2 where the same data used in Figure 1 was beamformed with random phase data inserted across multiple range bins. In Figure 2, the intensity of the random phase stripes matches the mean background, which is consistent with the assumption that the mean background in the image is speckle dominated and further underlines the assertion that the phase of data that forms shadows contains information and is not random.

The observation that phase history of SAS imagery shadows contain information opens up the possibility that shadow region phase history data can be used as a source for autofocus algorithms. In this paper, a novel autofocus algorithm is presented where phase history from SAS image shadow regions is used to focus the entire image. A description of the algorithm is presented along with a comparison of performance between similar point-driven and shadow-driven algorithms. Finally, samples of shadow-driven autofocused in situ SAS data are presented to illustrate the efficacy of this technique.




Figure 1: Unbeamformed SAS data (left), beamformed SAS data (center), and beamformed SAS data where the input data is normalized prior to beamforming.




Figure 2: A SAS image formed where the unbeamformed data was intensity normalized and the phase was randomized in 8 approximately 3 meter wide range bins prior to beamforming.

2 SHADOW PHASE GRADIENT ALGORITHM

2.1 Phase Gradient Algorithm

The phase gradient autofocus (PGA) algorithm uses the assumption that all data in a scene are blurred by the same blur function and measures the phase error responsible for the blur across multiple points in a scene⁴⁻⁵. This constant blur assumption requires that the blur is caused by a relatively uniform environmental effect or uncorrected system motion and not induced by object motion. Implementation of the PGA algorithm often uses subsets of an imaged area (snippet) for analysis to reduce the variability in the environment or uncorrected motion and relies on averaging phase error measurements across many points, typically the brightest point per range bin, in order to reduce the chance that the measured error is inconsistent with these requirements.

The PGA algorithm is performed on an image snippet by selecting high intensity points for analysis⁴⁻⁵. A typical implementation will include the point with the highest intensity per range bin for the snippet. These points are circularly shifted such that they will be in the center. A window is applied to the data to include the point blur is contained and exclude additional data that may confuse the analysis. A Fourier transform is performed in the cross range direction and a maximum likelihood estimator is used to determine the phase gradient of the error function associated with the blur in the snippet. The phase gradient is averaged across the range bins and applied to the data to correct the image. This process is often repeated until the desired level of focus is achieved.

An example of the PGA process is shown in Figure 3. For this example, blur is induced in SAS data by beamforming the data with an incorrect speed of sound (error of 20m/s). This error produces a quadratic phase error in the data prior to summation and results in image blur (Figure 3, far left). The brightest point in each range bin is selected and circularly shifted to the center of the frame (Figure 3, center left). An example window is shown in Figure 3 (right center) for illustrative purposes only. The actual window dimensions are determined by the expected width of the point response and will change over multiple iterations. The phase error was estimated, applied to the input data, and the resulting final focused image is shown on the right.




Figure 3: Unfocused 20m by 15m image data (left) for PGA correction, center shifted points for analysis (center left), windowed data prior to Fourier transform (center right), and PGA corrected imagery (right).

2.2 Sonar Simulation of Phase Errors

Simulated data was constructed in order to verify the efficacy of both highlight- and shadow-based PGA methods under evaluation. The first simulation method was to simply induce errors on existing SAS data prior to beamforming and to evaluate the PGA methods by comparing results to the original imagery. An example of this was discussed in the previous section where an incorrect speed of sound was used to induce a quadratic phase error in the data prior to beamforming and PGA corrected imagery was produced.

In order to verify that the methods are robust to arbitrary error functions, simulated data was also constructed with a variety of error functions as well as to ensure that the shadow response was not isolating some other characteristic of the data (underlying noise, etc.). This data was produced using an inverse imaging procedure where a random phase is assigned to an intensity image and inverse imaged to reconstruct the input data capable of producing the desired image⁸. This data was constructed with a background with a similar speckle intensity response when compared with real SAS imagery and included a single bright point in the center of the scene. The modeled SAS had a 100 kHz center frequency with 40 kHz bandwidth and a 50 kHz sample rate in a uniform environment with a speed of sound of 1500 m/s. Prior to beamforming, phase errors were applied to the data with a variety of forms including quadratic as well as non-quadratic asymmetric functions. An example of the simulated data is shown in Figure 4 which depicts the focused simulated image, a blurred image, and the PGA reconstructed images for both highlight- and shadow- driven methods. Examples of the error functions are shown in the next section.




Figure 4: An example of simulated data. The focused image on the left, blurred imagery (center left), highlight-driven PGA results (center right), and shadow-driven PGA results (right). Imagery is peak normalized.

2.3 Observed Phase Errors From Shadow Data

Development of a shadow-driven autofocus was initiated by simply observing the error estimated by the PGA algorithm when shadows rather than highlights were used for analysis. The resulting error functions were clearly inverted relative to the highlight results so the PGA algorithm was modified by multiplying the error function by -1 prior to image correction and a working shadow PGA (SPGA) algorithm was evaluated against a variety of simulated and real data to determine efficacy. The inverted response was uniform across all samples evaluated.

The relationship between highlight and shadow-driven results was further evaluated using simulated data with a quadratic phase error where the data used for the PGA algorithm was selected based on the intensity of the data relative to the image statistics of the snippet. The results of this analysis are shown in Figure 5 which shows the measured error function from PGA analysis for 9 equally populated datasets divided based on relative intensity. Each error estimate was produced as the mean from analysis of 10000 points out of the image snippets total of 90601 points. It can be seen that the error function is inverted when the intensity of samples used is below that of the mean intensity and that the samples near the mean intensity (~50th to 80th percentile) did not produce an error capable of focusing the imagery.

To ensure that the inverted response was not isolated to quadratic phase error, imagery was simulated with random asymmetric error functions and the resulting highlight- and shadow- driven PGA error functions were compared. These error functions were generated as the sum of two sine functions with random intensities, initial phases, and frequencies. Some example error functions are shown in Figure 6 which show the induced error (black) and the resulting estimate of the error from the highlight driven PGA algorithm (blue) and the inverted response resulting in using data below the mean intensity for PGA analysis (red). The inverted response for error estimated with data below the mean intensity was consistent for all simulated data.

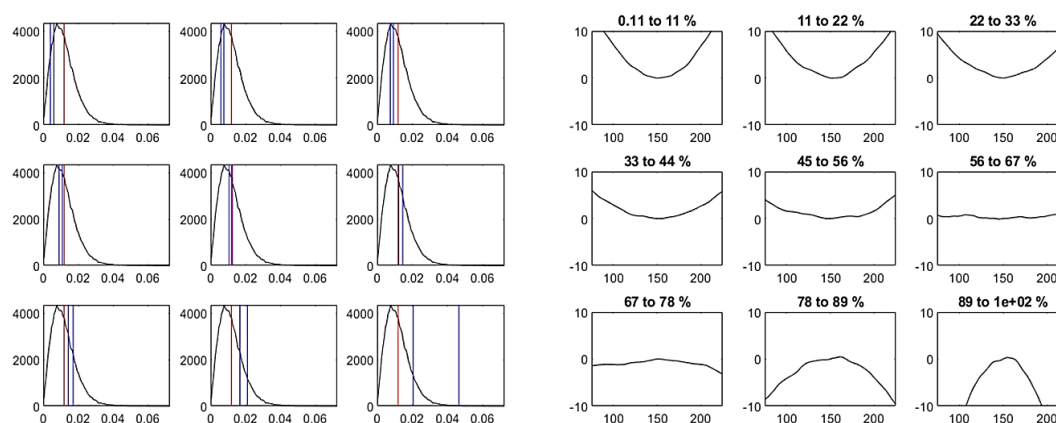


Figure 5: Histogram of data intensities used for PGA analysis (left) and the resulting measured error function (right). The data used for analysis is outlined in blue in the histogram of the snippet. The red line is the mean intensity of the snippet. The numbers on the error functions correspond to the ranked intensity percentile range of data used for analysis.




Figure 6: Asymmetric error functions used to blur imagery prior to PGA correction. Induced error from the simulation is shown in black and the PGA estimated error function is shown in blue for highlight-driven results and red for shadow-driven results.

2.4 Modification to PGA for Shadow Data

The standard highlight PGA approach was modified to enable shadow-driven error estimation for application to field collected SAS data described in Section 3. Primarily, points for analysis are selected to prefer the lowest intensity vice the highest and the estimated error is multiplied by -1 prior to image correction. These simple changes were sufficient to enable existing PGA algorithms to work on shadow data, but some additional modifications were able to provide faster convergence of the solution. Specifically, more data was used (i.e. not limited by a single value per range bin), larger windows were used on the initial pass, and more iterations were allowed if needed. The conditions applied to *in situ* data are described below where the algorithm was applied to multiple system's data.

3 RESULTS

3.1 Application to HISAS Data

The shadow-driven PGA algorithm was applied to data collected by a HISAS sensor. The original data was well focused, so a 20m/s error in the speed of sound was added prior to beamforming to induce quadratic phase error in the data and the associated image blur. The original focused image is shown in Figure 7 (upper left) as well as the blurred version (lower left) and both a highlight driven PGA resulting image (upper right) and a shadow-driven PGA resulting image (lower right). Both PGA methods produce resulting imagery that is sufficiently focused.

The method used to modify the existing PGA algorithm to enable shadow focusing was simply to select the pixel with the lowest intensity per range bin for analysis and to invert the sign on the resulting phase error estimate prior to image correction. An example of the intermediate step of PGA after point selection and circular shifting and prior to the Fourier transform is shown in Figure 8.

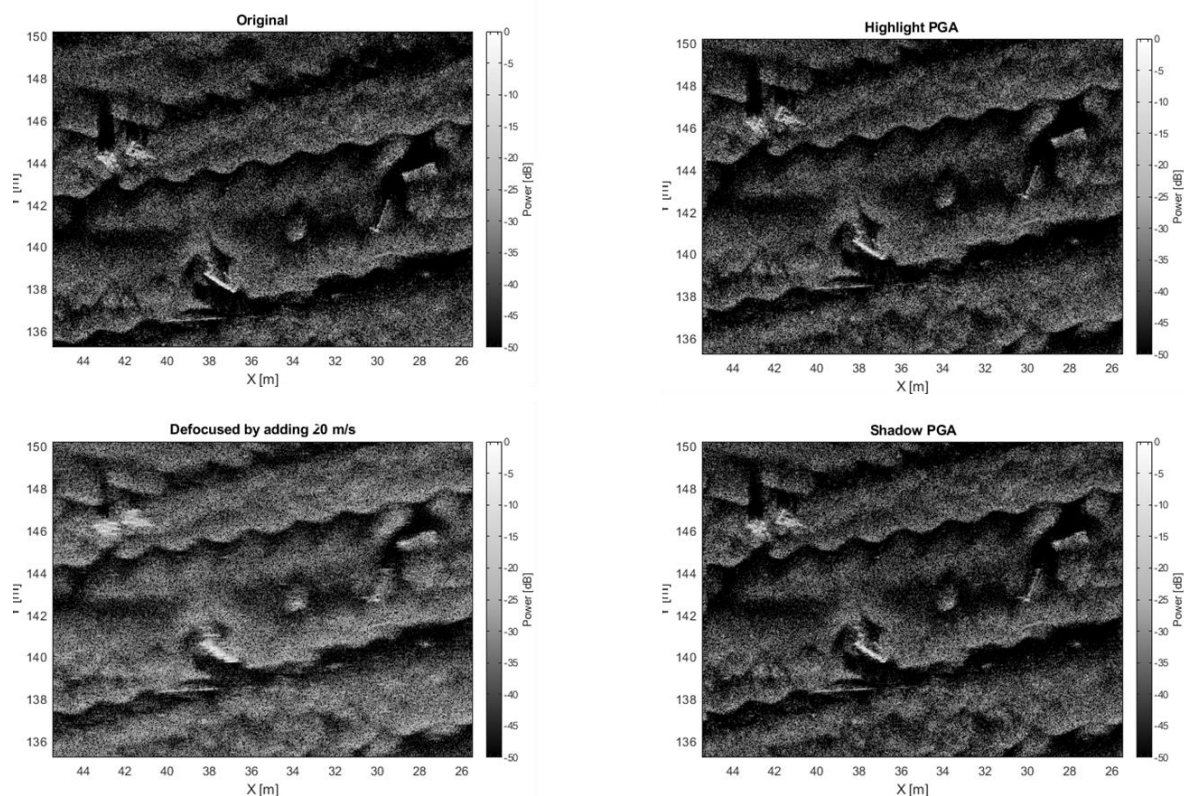


Figure 7: Focused HISAS imagery (upper left), blurred imagery from an incorrect speed of sound (lower left), and the resulting PGA focused images for both highlight (upper right) and shadow (lower right) algorithms.




Figure 8: The center shifted snippet pixels for PGA analysis for the highlight-driven (left) and the shadow-driven (right) PGA algorithms.

3.2 Application to SSAM2 Data

The shadow-driven PGA algorithm was applied to data from the Small Synthetic Minehunter 2nd generation (SSAM2) and compared with results from the standard highlight-driven approach. The data used for analysis are small snippets of randomly selected seafloor imagery from a variety of ranges with obvious image blur that could be resulting from uncorrected motion error or environmental phenomena. The median point spread width⁹ was used to determine blur in the original imagery and to quantify the improvement in PGA focused imagery. The dataset was randomly selected from samples with measurable blur that also showed improvement for standard PGA analysis.

The existing PGA algorithm was modified for shadow-based measurements by multiplying the phase error by -1 prior to image correction. Additionally, all points with an intensity less than the mean intensity value of the snippet were used for analysis for shadow PGA whereas the highlight based method used all points greater than 6 standard deviations from the mean intensity or the top 0.5% of intensities, whichever is greater. Finally, the shadow PGA maximum window size was increased from 32 to 128 samples and the algorithm was allowed 6 iterations vice the 4 that were allowed for the highlight case.

A sample of the seafloor snippets used for this analysis can be seen in Figure 9 which shows the original seafloor textures that exhibit obvious signs of blur with unknown origins. The samples are all capable of being focused by the existing highlight-driven PGA algorithm (Figure 10). The results of a shadow-driven approach are shown in Figure 11 which shows similar levels of focus indistinguishable from the highlight-based method. The image resolution metric⁹ results for the 9 samples does show some bias towards the highlight method. The original images have a mean along-track resolution of 3.9 pixels, the highlight PGA images have a mean resolution of 1.3 pixels, and the shadow PGA images have a mean resolution of 1.5 pixels. However it is important to point out that the point spread resolution metric used was designed to identify blurred samples and may not accurately discern between well focused samples.




Figure 9: Seafloor texture samples with obvious signs of blur.




Figure 10: Highlight PGA focused samples from Figure 9.




Figure 11: Shadow PGA focused samples from Figure 9:

3.3 Comparison of Methods

A larger set of snippets with existing blur were evaluated for a comparison between highlight- and shadow-driven PGA algorithms. A set of 577 samples used to train a machine learning algorithm to discriminate between image quality¹⁰ was used to test highlight and shadow PGA approaches. The results are shown in Figure 12 which shows the resulting resolution⁹ for both methods vs. the initial (blurred) resolution and a 1:1 line for reference. The highlight data is in blue and the shadow results are in red. Both methods appear to routinely improve focus. The shadow results are also plotted against the highlight results and reinforce the result seen in Section 3.2 where on average the shadow results appear to have slightly worse resolution compared with the highlight PGA results.

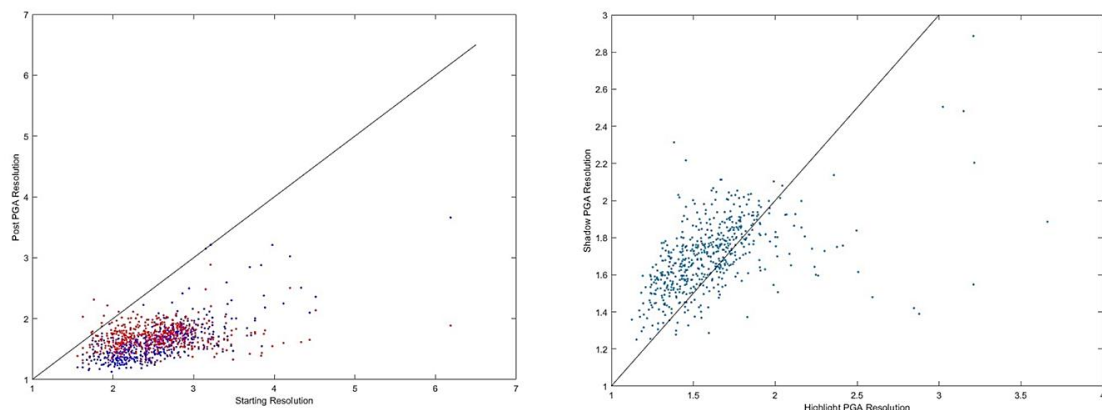


Figure 12: PGA resulting resolution vs starting resolution for blurred snippets (left) and shadow vs. highlight PGA results (right).

4 DISCUSSION

Results presented here clearly indicate that there is information in snippet pixels below mean intensity that can be used to focus imagery. However, these data require slight modification to PGA algorithms and may require more computational power to perform on par with highlight data. The lower signal to noise ratio of the low intensity data can be offset by including more points for analysis. While a typical highlight PGA analysis uses $\sim <1\%$ of the data (one point per range bin), a shadow-driven PGA can use $\sim 50\%$ of the data (algorithm used in Section 3.2 acted on all data less than mean intensity).

While this work was intended to determine if information was present in low intensity data and algorithms were developed for highlights and shadows independently, a future PGA implementation to exploit this data would likely apply both low and high intensity data simultaneously to estimate error functions. By including a larger data set for analysis, the probability of inducing an error associated with selecting inappropriate points (like moving objects) is reduced. Furthermore, by incorporating all data available, smaller snippets can be used for analysis. By acting on smaller areas of the original image, the underlying assumption for PGA that residual phase error is constant over the scene is also less likely to be violated.

Finally, the observed phenomena where information is obtained from low-intensity data warrants further investigation. It is clear that low intensity data contains a signal that can be exploited for additional measurements. This signal should be better understood in order to appropriately model shadows in acoustic simulations. Modeling a bright point is relatively straight forward. A bright response is modeled with the correct phase associated with the distance to the point for the aperture samples. However, shadows are rarely observable in unbeamformed data and the shadow is formed during beamforming due to the phase of the data actively canceling the return. Models need appropriate signals to add to simulations in order to produce realistic shadows. The inversion of the phase for shadow-based PGA measurements is an interesting result and may provide insight on why this method works and potentially improve the understanding of how shadows work in wide beam acoustic data.

5 ACKNOWLEDGEMENT

The authors would like to express gratitude to the Office of Naval Research and to the FFI Norwegian Defence Research Establishment for support in the development of SAS signal processing algorithms, access to SAS data, and support of research in autofocus algorithm development and sonar simulations.

6 REFERENCES

- [1] D. W. Hawkins: "Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar", PhD thesis, Department of Electrical Engineering, University of Canterbury, Christchurch, New Zealand, 1996.
- [2] M. P. Hayes, P. T. Gough, "Synthetic Aperture Sonar: A Review of Current Status", *IEEE Journal of Oceanic Engineering*, vol. 34, no. 3, 2009.
- [3] D. A. Cook, "Synthetic Aperture Sonar Motion Estimation and Compensation", Thesis, Georgia Institute of Technology, May 2007.
- [4] D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and C. V. Jakowatz, Jr., "Phase gradient autofocus – a robust tool for high resolution SAR phase corrections", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 30, no. 3, pp. 827-835, July 1994.
- [5] C. V. Jakowatz, Jr., D. E. Wahl, P. H. Eichel, D. C. Ghiglia, P.A. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*, Boston, Kluwer Academic Publishers, 1996.
- [6] H. J. Callow, M. P. Hayes, and P. T. Gough. "Autofocus of stripmap SAS data using the range-variant SPGA algorithm". In OCEANS 2003. Proceedings, volume 5, pages 2422-2426 Vol.5, 22-26 Sept. 2003.
- [7] J. L. Prater, H. Schmaljohan, "Premiminary Results From Attempts to Determine SAS Array Coherence From Image Metrics" International Conference on Synthetic Aperture Sonar and Synthetic Aperture Radar, Proceedings of the Institute of Acoustics, Lercici, Italy, 5-7 September, 2018.
- [8] J. L. Prater, H. Schmaljohan, "Side-looking sonar modelling using inverse imaging techniques", in the MTS/IEEE Oceans '17 Anchorage Proceedings, AK, September 2017.
- [9] J. L. Prater, J. L. King, D. C. Brown, "Determination of Image Resolution from SAS Image Statistics", in the MTS/IEEE Oceans '15 Washington D.C. Proceedings, October 2015.
- [10] J. Dale, M. Emigh, J. Prater, "DAFI: A deep learning-based autofocus improvement metric for synthetic aperture sonar", in the Synthetic Apertures in Sonar and Radar Proceedings of the Institute of Acoustics, Lercici, Italy, September 2023.