MOTION TRACKING OF TRANSIENT REFRACTIVE EFFECTS IN SAS IMAGERY USING OPTICAL FLOW

Daniel A. Cook — Georgia Tech Research Institute, Atlanta, USA
Roy E. Hansen — Norwegian Defence Research Establishment (FFI), Kjeller, Norway
Anthony P. Lyons — Applied Research Laboratory, Penn State University, USA
Anthony J. Yezzi — Georgia Institute of Technology, Atlanta, USA

1. INTRODUCTION

Traveling boluses induced by breaking internal waves[1] have recently been observed in synthetic aperture sonar imagery,[2,3] where their appearance resembles natural seabed ripple patterns. The boluses cause refraction that concentrates the acoustic energy in a manner that results in regions of increased intensity followed downrange by reduced intensity. They were discovered because the same area was surveyed multiple times, and the apparent sand ripples shifted in a manner that was inconsistent with known sediment transport behavior in the area. Example imagery showing the refractive effect of these boluses is given in Figure 1. Detecting them can be a time consuming process if repeat surveys are required. The motivation therefore exists to find detection methods that can be applied to single-pass imagery. Optical flow is a well-established means of estimating frame-to-frame motion of video image content, and this paper explores its potential for single-pass detection and motion tracking of internal waves.

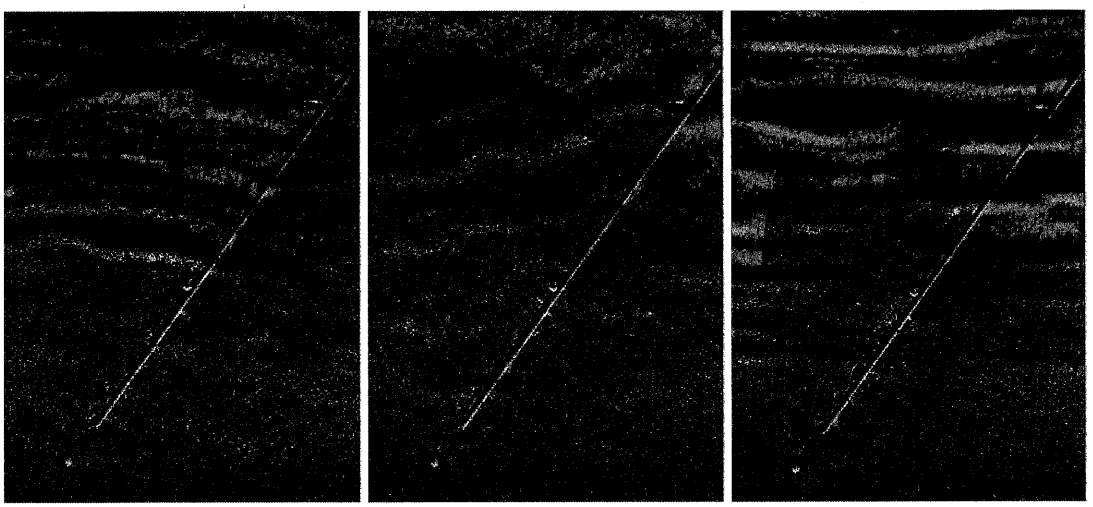

Figure 1: SAS images of the same scene were collected on three different passes on the same day, at times 07:19 (left), 08:18 (center), and 12:05 (right). The scene size is 80 m x 100 m and shows a pipeline with anchors on the seabed. The data were collected by FFI's HUGIN AUV carrying an interferometric SAS during the ARISE12 trials on board the NATO research vessel *Alliance*, outside Elba island, Italy.

The authors gratefully acknowledge the Norwegian Defence Research Establishment and the US Office of Naval Research for their support of this work.

2. OPTICAL FLOW

The Lagrangian view of fluid mechanics may be used to motivate the derivation of optical flow methods. In an analogy with fluid flow, consider a small patch of fixed area whose image intensity is represented as $I(x(t), t)$, where the position vector $x(t) = [x_1(t),\ x_2(t)]^\mathsf{T}$. The patch is able to translate in space and its intensity may vary. Taking the total derivative with respect to time gives

$$\frac{\partial I}{\partial x_1}\frac{dx_1}{dt} + \frac{\partial I}{\partial x_2}\frac{dx_2}{dt} + \frac{\partial I}{\partial t} = \left(\frac{\partial}{\partial t} + v \cdot \nabla\right) I = \frac{DI}{Dt}, \tag{1}$$

where $v = [v_1,\ v_2]^\mathsf{T} = [dx_1/dt,\ dx_2/dt]^\mathsf{T}$, and the shorthand at the end of (1) is also known as the material, or substantial, derivative.[4] The material derivative gives the instantaneous change in I at any given position and time.

When the intensity of each image patch is constrained to be unchanging as it moves (a condition known as *brightness constancy* or the *optical flow condition*), the resulting PDE becomes $DI/Dt = 0$ or $\partial I/\partial t = -\nabla I \cdot v$. By itself, this equation cannot be solved for the velocity field. It only captures the component of the optical flow in the direction of ∇I, that is, the flow normal to contours of constant intensity. This ambiguity of motion is known as the *aperture problem*, and there are a number of approaches for resolving it.[5] A common approach is to impose a smoothness constraint on v.

The method of Horn and Schunck[6] was chosen for this initial investigation using optical flow for detecting the presence of transient refractive effects. Here, brightness constancy is combined with a smoothing term to obtain the energy functional to be minimized:

$$\begin{aligned} E &= (1-\lambda)\int_\Omega \left(v \cdot \nabla I + \frac{\partial I}{\partial t}\right)^2 dx + \lambda \sum_{j=1}^{2} \int_\Omega \|\nabla v_j\|^2\, dx \\ &= \int_\Omega \mathcal{L}(v, v', x, t)\, dx \end{aligned} \tag{2}$$

where Ω is the domain integration (i.e., the image), v' represents the possible derivatives of v with respect to its arguments (e.g., $\partial v/\partial x_1$ and $\partial v/\partial x_2$), and $\lambda \in [0, 1]$ is a user-chosen parameter whose purpose is to balance the relative contributions of the data fidelity and smoothness terms comprised by the Lagrangian $\mathcal{L}$. Horn and Schunk use a slightly different weighting α, which is equivalent to $\lambda/(1-\lambda)$. They suggest that the square of this term should be roughly equal to the expected noise level of the image gradient.

The calculus of variations is used to find the function $v(x, t)$ that minimizes the energy functional E. This is accomplished by solving the Euler-Lagrange equations[7,8] derived from $\mathcal{L}$ according to the following condition necessary for minimizing E:

$$\frac{\partial \mathcal{L}}{\partial f_i} - \sum_j \frac{\partial}{\partial \xi_j}\left(\frac{\partial \mathcal{L}}{\partial f'_{i,j}}\right) = 0, \tag{3}$$

where the subscript i indicates the components of the vector function f, ξ_j is the j^{th} argument of f, and $f'_{i,j} \equiv \partial f_i/\partial \xi_j$. For this problem, v is identified with f. Carrying out the calculations, a pair of coupled equations is obtained:

$$\begin{aligned} (1-\lambda)\left(v \cdot \nabla I + \frac{\partial I}{\partial t}\right)\frac{\partial I}{\partial x_1} - \lambda \nabla^2 v_1 &= 0 \\ (1-\lambda)\left(v \cdot \nabla I + \frac{\partial I}{\partial t}\right)\frac{\partial I}{\partial x_2} - \lambda \nabla^2 v_2 &= 0. \end{aligned} \tag{4}$$

There are a number of suitable numerical approaches for solving the Euler-Lagrange equations for v. A simple one is gradient descent, in which an artificial dependence on a time marching parameter τ is assigned to v. Re-deriving the Euler-Lagrange equations results in the same pair of equations above with $[\partial v_1/\partial\tau,\ \partial v_2/\partial\tau]^{\mathsf{T}}$ on the right-hand side instead of $[0,\ 0]^{\mathsf{T}}$. Starting from an initial estimate v_0 and the measured image pair, the solution is iterated until $[\partial v_1/\partial\tau,\ \partial v_2/\partial\tau]^{\mathsf{T}}$ is arbitrarily close to $[0,\ 0]^{\mathsf{T}}$.

3. IMAGE FRAME GENERATION

Optical flow computations are typically applied to sequential video frames. For SAS, repeat passes are the obvious way to obtain such a sequence, but this is highly undesirable. The time between passes is often long, and re-imaging an area severely reduces the area coverage rate of the survey. The goal of this work is to find a reliable approach for detecting transient refractive effects within a single image.

There are a number of ways in which a single SAS data collection might be divided into subsets that can be processed into sequential images. Some of these are discussed by Hansen et al.[3] The approach chosen for this paper begins with a complex-valued stripmap SAS image, as opposed to the raw sonar data it was derived from. The image is Fourier transformed into the two-dimensional (k_x, k_y) wavenumber domain, allowing the data to be easily filtered according to aspect, θ_{view}, and integration angle, θ_{int}, by applying the appropriate mask. The region of interest is defined by $\mathcal{M}(k_x, k_y) = 1$ for

$$\left.\begin{aligned} k_x &= 2k\sin\theta \\ k_y &= 2k\cos\theta \end{aligned}\right\} \qquad \theta \in \left[\theta_{\text{view}} - \frac{\theta_{\text{int}}}{2},\ \theta_{\text{view}} + \frac{\theta_{\text{int}}}{2}\right], \tag{5}$$

and zero otherwise. Additional masking and/or weighting to control the shape of the point scatterer response could be built into $\mathcal{M}$. Applying an inverse Fourier transform to the masked spectrum provides a single image frame.

The performance of the optical flow depends on θ_{view} and θ_{int}. Nonzero values of θ_{view} correspond to beams steered fore/aft, and there is a corresponding time delay between the image formed from these beams. Since the images are separated by angle, the delay between images depends on range according to $\Delta t = v^{-1}y(\tan\theta_{\text{fore}} - \tan\theta_{\text{aft}})$, where v is the forward speed of the vehicle, y is range, and θ_{fore} and θ_{aft} are the values of θ_{view} for the images comprising the pair. Consequently, the image-to-image displacement vector measured by optical flow can be converted to a velocity by dividing it by Δt.

There is a trade off associated with the selection of integration angles for the generation of frames. Larger integration angles are generally favored for imaging because they give improved cross-range resolution. When the object of interest is moving, smaller integration angles are often more suitable because the motion may be considered to be negligible or perhaps constant over the time required to collect the data. This fact has been used by the radar community for detecting moving objects in SAR imagery.

4. EXPERIMENTAL APPROACH

All optical flow schemes are affected by noise in the image, resolution, and the amount of change between consecutive images, which are typically fixed by the sensor and imaging conditions. The SAS application is somewhat unique in that it offers the flexibility to vary all of these. These parameters are coupled, but the user has tremendous latitude nevertheless.

A key property of SAS is that an image can be modeled as an underlying reflectivity function multiplied by the speckle caused by the coherent rough scattering from the sea floor. The speckle modulation varies between 0 and 1 and imparts a grainy texture to the image. For high-quality SAS imagery, the speckle is more pronounced than the noise. Assuming this to be the case, it seems reasonable to follow the advice of Horn and Schunck and set the weighting parameter α^2 equal to the average value of the squared magnitude of the noise (or speckle) gradient. To constrain the parameter space for this initial proof of concept study, this is how α^2 was set. The desired value can be easily computed from a speckle realization created separate from the image, but with the same wavenumber spectrum support and weighting. This is done by populating the wavenumber spectrum with complex phasors, whose phases are uniformly distributed, and then multiplying it by $\mathcal{M}$. Alternatively, the excellent book by Goodman[9] contains a closed-form analysis of speckle statistics that could be used to compute α directly.

Fine-resolution SAS imagery has resolution on the order of 2-4 cm, while the refractive distortion in the image is typically on the scale of meters. The SAS image therefore has excess resolution that can be traded off to reduce speckle via multilooking.[10] Multilooking carries the added benefit of reducing the pixel count in the image, greatly accelerating the optical flow computation. This reduction is advantageous because SAS imagery often contains on the order of 10-100 million pixels, and the optical flow can require thousands of iterations resulting in impractically-long computation times. Since the image frames are generated by excising a subset of the available beamwidth, the multilooking is applied only in the k_y dimension. Traditional multilooking tiles the wavenumber spectrum in both k_x and k_y.

In addition to the integration angle (resolution), the angle between frames (time separation) must be decided. These angles can be such that there may or may not be overlap in the spectral support between two or more consecutive frames. It is possible to construct a sequence of images with arbitrarily small time separation. There are drawbacks to making the time separation too large or too small. If the inter-frame time separation is too small, then slow-moving features in the image may not change enough to be reliably detected. If it is too large, then the scene content may change drastically resulting in a meaningless optical flow field.

Regardless of the integration angle, a key consideration is the amount of spectral overlap between consecutive frames. If there is no overlap, then the speckle realizations will be completely uncorrelated. Otherwise, the speckle will exhibit some degree of correlation. It is important to note that the correlation of the speckle is in addition to any correlation due to the underlying seafloor reflectivity, for example, from specular or point-like reflectors. These facts make it difficult to predict the correlation between frames, but it is clear that as the overlap of the frames' spectral regions of support decreases, the optical flow field can be expected to exhibit some amount of fine scale randomness.[11] This randomness can be masked or altered by the smoothing term in Equation (2).

5. EXPERIMENTAL RESULTS

The optical flow estimation was applied to data collected by the HUGIN autonomous underwater vehicle carrying a HISAS wideband interferometric SAS. This system is operated by the Norwegian Defence Research Establishment (FFI). Two images were chosen for comparison. Both are scenes of a smooth sea floor, where one image is free of refractive anomalies while the other is known to contain them. The presence of the anomalies was confirmed using imagery from repeat passes. (Recall Figure 1, and see the article by Hansen, et al.[3]).

The same 40 m by 40 m range and cross-range interval was excised from the original imagery. The refraction-free image contains some man-made objects large enough to cast shadows, but otherwise shows a uniform sea floor. The image containing the refractive effect was taken from

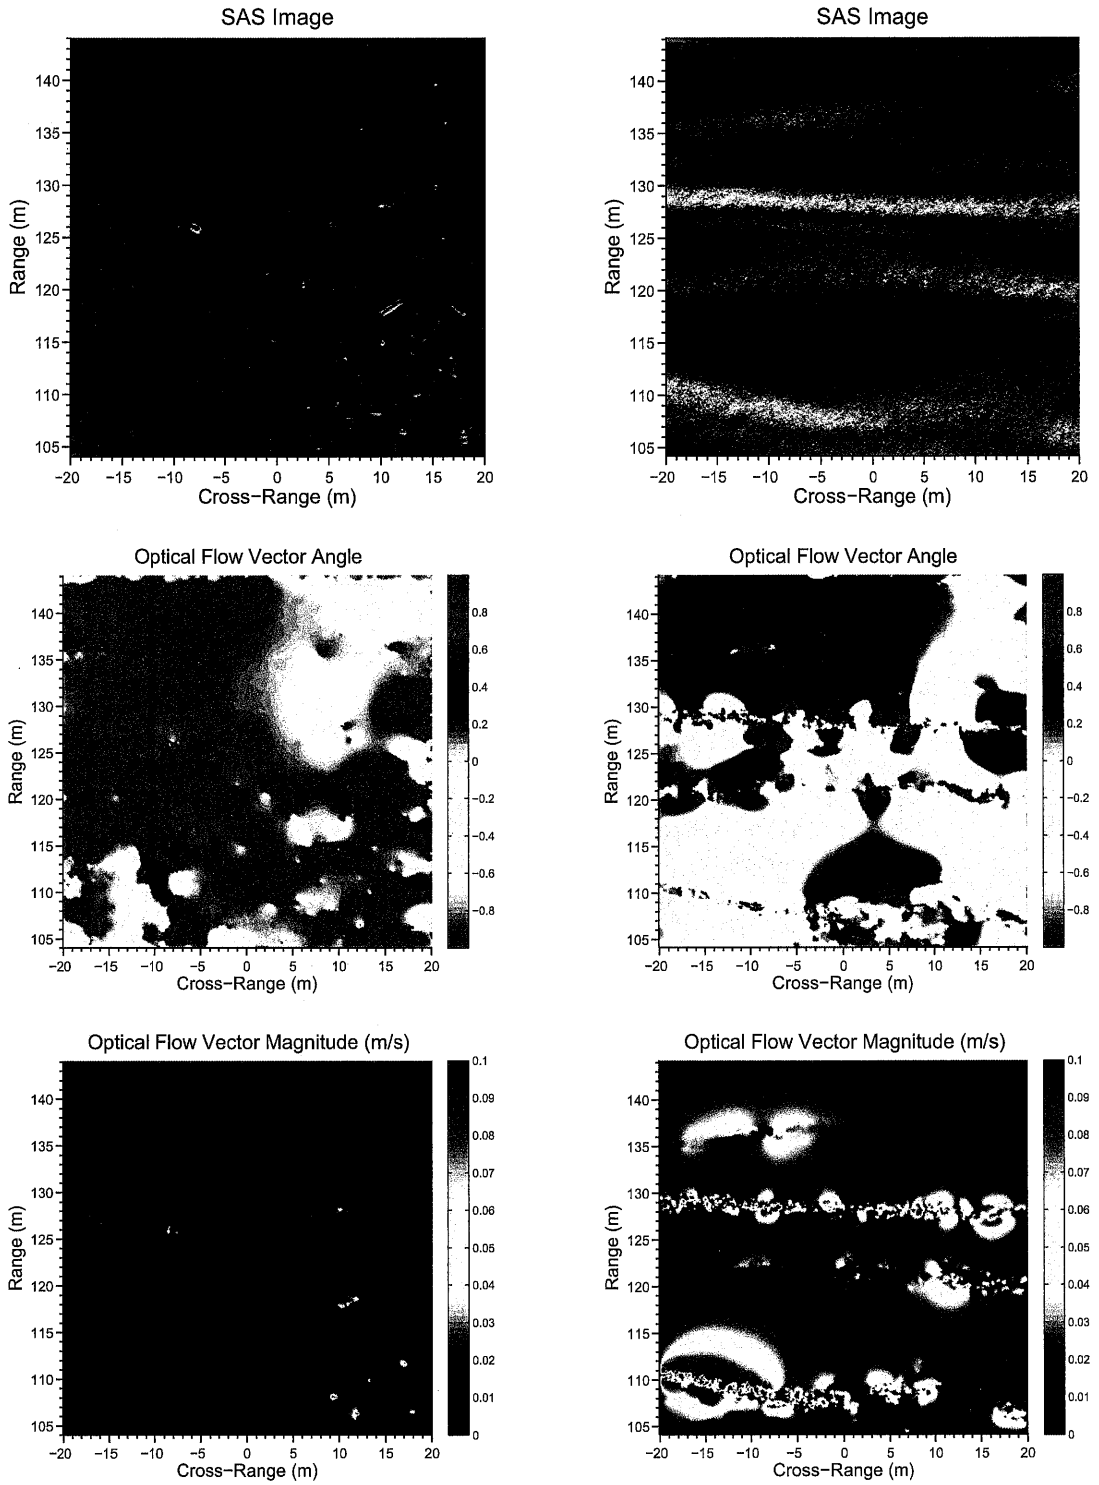

Figure 2: Comparison of optical flow for images without (left column) and with (right column) refractive effects. The top row shows the underlying SAS imagery, while the middle and bottom rows show the angle and magnitude of the optical flow vector field.

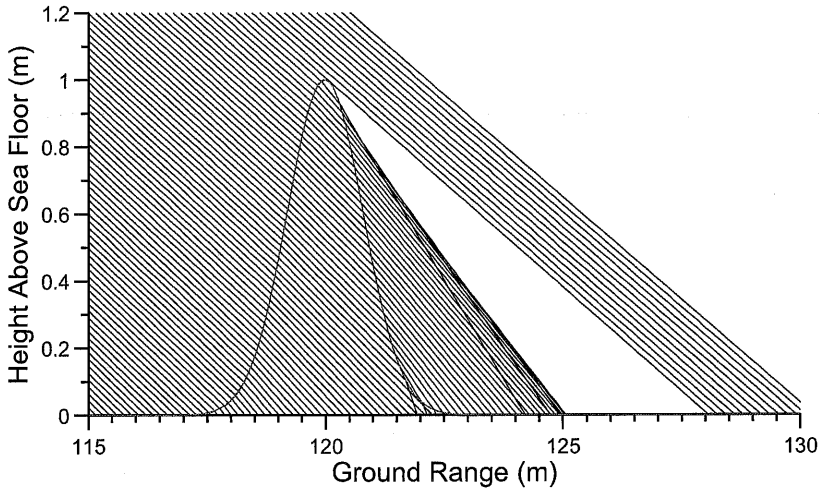

Figure 3: Rays were traced through a hypothetical internal wave bolus using a geometry similar to the collection used for this study. Note the region of concentrated acoustic energy followed downrange by a region that would appear as a shadow in an image.

the upper-left corner of the right-hand image in Figure 1. No objects are visible in the excised region, but it can be inferred from Figure 1 that the image is well-focused.

Example optical flow results are shown in Figure 2. The top row of the figure shows one frame from the image sequence used to compute the optical flow. The middle row shows the angle of the optical flow vector field, and the bottom row shows its magnitude. The effective integration angle is 1° and the frames correspond to ±2°. It is important to stress that the optical flow field is *not* a direct measurement of the motion of the transient refractive anomaly. For reasons mentioned in the previous section, even a totally quiescent environment is expected to exhibit some amount of optical flow. Furthermore, the multiaspect image frames will cause shadows to be cast in different directions.[12]

The refractive boluses shown in the figure are roughly parallel to the cross-range axis and are known to be moving quite slowly in the range direction. The induced optical flow would be expected to have significant components near $+\pi/2$, provided the temporal separation between frames is such that the movement of the boluses can overcome the quiescent optical flow field. This is indeed the case shown in Figure 2. It is clear that the flow is parallel to the range axis, but it is not uniform. Rather, it is characterized by distinct regions flowing both uprange and downrange. The flow direction was observed to switch between $\pm\pi/2$ from frame to frame, and the reason for this is not yet understood. Another visible effect are the unstructured regions appearing as nearly horizontal stripes at ranges of 110 m, 120 m, and 128 m. These locations coincide with the highlights in the SAS image.

This optical flow behavior is impossible to explain completely without knowledge of the spatial variation of the sound speed. However, a clue might be obtained by considering the refractive behavior through a single internal wave bolus. The rays crossing through the bolus can concentrate and dilute the intensity within the image, as illustrated in Figure 3. It is reasonable to expect that the expansion and contraction of the rays observed as the bolus travels can cause apparent localized motion within the image. There are also regions where rays can cross, and perhaps these are responsible for the distorted optical flow located along the image highlights.

Some qualitative expectations can be proposed: The angular-temporal nature of the multiaspect filtering used to create the image frames should produce two kinds of effects. Any changes in

the speckle pattern should result in a randomization of the optical flow field. It might be true that the magnitudes of these random vectors is small. However, this is difficult to assert since speckle violates the brightness constancy assumption, behaving as a continuum of sources and sinks of intensity such that $DI/Dt \neq 0$. When objects are present in the scene, their highlights and shadows will shift primarily in the cross-range direction according to the change in θ_{view}. The corresponding optical flow vectors would point in the cross-range direction as well (angles near 0 and π). Even in the absence of shadows, the scene should experience a bulk cross-range shift, as is observed in the middle plot of the left-hand column of Figure 2.

6. CONCLUSION

Optical flow has been experimentally demonstrated to show promise for at least detecting the presence of transient refractive effects using a single SAS image. The next logical step would be to automate the detection since the visual cues hinting at the presence of refractive anomalies can be subtle, often mimicking the appearance of natural effects. Such a capability would be useful for identifying portions of a survey area that may suffer from reduced image quality. A far more ambitious goal is to measure the character of the refracting disturbance, for example, the shape of its boundary and the sound speed profile within it. Such measurements could have a significant impact on the science of oceanography.

Achieving any of these goals requires further consideration of the properties of synthetic aperture sonar and how they can be exploited to refine the optical flow estimation. One avenue of investigation is to develop optical flow models and constraints that are better matched to SAS than the classical Horn and Schunck approach. Another is to find ways of automating the process of setting parameters that are normally chosen empirically by the user. A first attempt at establishing some of these connections was made when α^2 was set according to the mean of $||\nabla I_s||^2$, where I_s is a speckle realization designed to have the same statistics as the speckle contained in the SAS imagery. The current body of knowledge regarding SAS imaging should provide a solid foundation for further links between the theory of imaging and the determination of optical flow using PDEs.

References

1. K. R. Helfrich, "Internal solitary wave breaking and run-up on a uniform slope," *Journal of Fluid Mechanics*, pp. 133–154, 1992.

2. A. P. Lyons, R. E. Hansen, T. O. Saebø, and H. J. Callow, "Refractive effects of internal waves on synthetic aperture sonar images," in *Proceedings of the 1st Underwater Acoustics conference, Corfu, Greece*, 2014.

3. R. E. Hansen, A. P. Lyons, T. Sæbø, H. J. Callow, and D. A. Cook, "The effect of internal wave-related features on synthetic aperture sonar," *IEEE Journal of Oceanic Engineering (In Press)*, 2014.

4. B. R. Munson, D. F. Young, and T. H. Okiishi, *Fundamentals of Fluid Mechanics*, 2nd ed. John Wiley & Sons, Inc., 1994.

5. G. Aubert and P. Kornprobst, *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*. Springer-Verlag, 2002.

6. B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

7. H. Goldstein, *Classical Mechanics*. Addison-Wesley Publishing Company, Inc., 1953.

8. L. C. Evans, *Partial Differential Equations*. American Mathematical Society, 2002.

9. J. W. Goodman, *Speckle Phenomena in Optics*. Roberts and Company Publishers, 2010.

10. C. V. Jakowatz, D. E. Wahl, P. H. Eichel, D. Ghiglia, and P. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Boston, MA: Kluwer Academic Publishers, 1996.

11. W. F. Walker and G. E. Trahey, "The application of k-space in pulse echo ultrasound," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 45, no. 3, 1998.

12. J. Groen, R. E. Hansen, H. J. Callow, J. C. Sabel, and T. O. Sæbø, "Shadow enhancement in synthetic aperture sonar using fixed focusing," *IEEE Journal of Oceanic Engineering*, vol. 34, no. 3, pp. 269–284, July 2009.