

Coherent space-time stereoscopic reconstruction of sea states

Guillermo Gallego*, Anthony Yezzi†, Francesco Fedele‡ and Alvisè Benetazzo§

Abstract

We present a variational method for the coherent space-time reconstruction of oceanic sea states from stereo video data. This remote sensing observational method yields accurate estimates of the spatio-temporal dynamics of ocean waves. The main idea is to pose the stereoscopic reconstruction problem in a variational optimization framework and design an energy functional whose minimizer is our desired temporal sequence of wave heights. Our functional combines photometric observations as well as spatial and *temporal* smoothness priors. A nested iterative scheme is devised to numerically solve via 3-D multigrid methods the system of partial differential equations resulting from the optimality condition of the energy functional. The output of our method is the coherent, simultaneous estimation of the wave surface height and radiance at multiple snapshots. We demonstrate our algorithm on real data collected off-shore. Statistical and spectral analysis are performed. Comparison with respect to an existing sequential method is analyzed.

1 Introduction

The study of the dynamics of oceanographic phenomena using vision systems has a long tradition that dates back to the first half of the previous century ([13, 14, 8], etc.). This topic has gained popularity in recent years due to the economical and practical benefits of these systems [12, 2, 16, 10, 4, 9].

In this paper, we build upon the graph variational method presented in [5, 6] for the stereoscopic reconstruction of oceanic sea states. The variational method for still images was extended to process stereo video on a sequential, snapshot-by-snapshot basis. Besides this, the variational framework also allows for more ways to incorporate temporal coherence on the reconstructed surface. The ultimate goal of such an observational technique is to include the spatial and temporal physics of the waves in the reconstruction step [7]. Because ocean waves are governed by the wave equation, it would be desirable to include such law in the estimation process. This is however, a challenging problem. Before considering this approach, a natural way to enforce temporal coherence of the reconstructed wave height, besides the purely sequential processing above mentioned, is to include a temporal regularizer in our energy model and solve the resulting variational problem. This is the object of this work and it implies a *simultaneous* estimation of the unknowns (wave height, radiance, etc.) for *all* snapshots in a sequence. The resulting reconstruction is called a *manifold* reconstruction (MR) because the developed method estimates a manifold of graphs (wave heights or elevation maps).

2 Theoretical model

In a multi-camera setup recording synchronized videos of the ocean surface, let the video signal acquired by the i -th camera be denoted by $I_i(\mathbf{x}_i, \tau)$, where $\mathbf{x}_i = (x, y)^\top$ and τ denote the (continuous) spatial and temporal variables, respectively. In reality, a digital video signal only specifies the intensity values at discrete locations of the space-time axes (pixels and time instants), however, it is

*Grupo de Tratamiento de Imgenes, Universidad Politécnica de Madrid, Madrid, Spain.

†School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, USA.

‡School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, USA.

§CNR-ISMAR, Venice, Italy.

more beneficial to model the problem in the continuum. Following the physical model of the scene presented in [5], consider the surface shape (wave height) and surface radiance as functions of space *and time*, i.e., as 3-D functions: $Z(u, v, \tau)$ and $f(u, v, \tau)$, with $\mathbf{u} = (u, v) \in U$ and $\tau \in [0, T]$. Recall that the surface radiance is a function that lives on top of the surface and represents the “color” of the wave that is seen from the viewpoints of the cameras. Let us denote the domain of Z and f by $U_T = U \times [0, T]$. Let the symbol ∇Z denote the gradient of Z with respect to all variables, not only the spatial ones (u, v) , i.e.

$$\nabla Z = (Z_u, Z_v, Z_\tau)^\top,$$

and similarly for ∇f . With this notation, the three variables u, v and τ are treated uniformly.

Next, we define an energy functional to measure the goodness of fit of candidate space-time functions Z and f to the observed data (i.e., the stereo videos):

$$E(Z, f) = E_{\text{data}}(Z, f) + \alpha E_{\text{geom}}(Z) + \beta E_{\text{rad}}(f), \quad (2.1)$$

with weights $\alpha, \beta \in \mathbb{R}^+$. Let the data fidelity term, which measures the photo-consistency throughout the video for a candidate wave height function, be

$$E_{\text{data}} = \sum_{i=1}^{N_c} E_i,$$

with

$$E_i \doteq \int_0^T \int_{\Omega_i} \phi_i \, d\Omega_i \, d\tau$$

and photometric criterion

$$\phi_i = \frac{1}{2} (I_i(\mathbf{x}_i, \tau) - f(\mathbf{x}_i, \tau))^2. \quad (2.2)$$

The definition of the data fidelity term as an integral over the image domain Ω_i (rather than over U) has two advantages: the data term is independent of the choice of domain for the graph and the resulting optimality conditions for the minimization of (2.1) lack image derivatives. This desirable property is inherited from the modeling and mathematical principles that we follow from [11]. The resulting method is less sensitive to image noise than other variational approaches for stereo 3-D reconstruction.

Also, let the spatio-temporal regularizers on the smoothness of the surface shape and radiance be:

$$E_{\text{geom}} = \int_{U_T} \frac{1}{2} \|\nabla Z\|^2 \, d\mathbf{u}_T, \quad (2.3)$$

$$E_{\text{rad}} = \int_{U_T} \frac{1}{2} \|\nabla f\|^2 \, d\mathbf{u}_T. \quad (2.4)$$

where $d\mathbf{u}_T = d\mathbf{u} \, d\tau$. The new terms with respect to the still-image case [5], Z_τ^2 and f_τ^2 , are the source of temporal coherence in the model.

The radiance function is an auxiliary variable with a twofold purpose: it simplifies pairwise image-to-image comparisons with fewer image-to-model comparisons and it also extends the influence of the point-wise photometric criterion by means of neighborhood effects due to the regularizer E_{rad} . The latter endows the energy functional with a global influence and improves the robustness of the energy functional against image noise.

The data fidelity term corresponding to the i -th camera can also be expressed as

$$E_i = \int_{U_T} \phi_i \mathbf{J}_i \, d\mathbf{u}_T,$$

where the Jacobian \mathbf{J}_i (the geometric conversion factor between integrals) has the same expression as in the still-image case despite the fact that now the Jacobian is also a function of time τ (since

J_i depends on the surface shape, which is a function of time). Therefore, the energy (2.1) can be expressed in the common and fixed domain of integration U_T as the integrand of the so-called Lagrangian $L(Z, \nabla Z, f, \nabla f, u, v, \tau)$:

$$E = \int_{U_T} L \, d\mathbf{u}_T.$$

To minimize the proposed energy, let us compute the necessary optimality condition, which follows from the first variation (directional derivative) of the energy with respect to the unknowns:

$$D_{(h,w)}E(Z, f) = \left. \frac{d}{d\epsilon} E(Z + \epsilon h, f + \epsilon w) \right|_{\epsilon=0}.$$

Setting to zero the directional derivative for all admissible perturbations (h, w) we arrive at a coupled system of partial differential equations (PDE), the so-called Euler-Lagrange (EL) equations:

$$\begin{aligned} L_Z - (L_{Z_u})_u - (L_{Z_v})_v - (L_{Z_\tau})_\tau &= 0 & \text{in } U_T, \\ L_{Z_u} \nu^u + L_{Z_v} \nu^v + L_{Z_\tau} \nu^\tau &= 0 & \text{on } \partial U_T, \\ L_f - (L_{f_u})_u - (L_{f_v})_v - (L_{f_\tau})_\tau &= 0 & \text{in } U_T, \\ L_{f_u} \nu^u + L_{f_v} \nu^v + L_{f_\tau} \nu^\tau &= 0 & \text{on } \partial U_T. \end{aligned}$$

After calculations, similar expressions to those of the still-image optimality conditions are derived. That is why this approach is considered a natural extension of the aforementioned still-image case. Now, however, variables Z, f , and consequently J_i and I_i depend on the new temporal variable τ :

$$g(Z, f) - \alpha \Delta Z = 0 \quad \text{in } U_T, \quad (2.5)$$

$$b(Z, f) + \alpha \frac{\partial Z}{\partial \boldsymbol{\nu}} = 0 \quad \text{on } \partial U_T, \quad (2.6)$$

$$-\sum_{i=1}^{N_c} (I_i - f) J_i(Z) - \beta \Delta f = 0 \quad \text{in } U_T, \quad (2.7)$$

$$\beta \frac{\partial f}{\partial \boldsymbol{\nu}} = 0 \quad \text{on } \partial U_T, \quad (2.8)$$

where the non-linear terms due to the data fidelity energy have the same expression as those of the still-image energy since the data fidelity energy does not depend on the new derivatives Z_τ, f_τ :

$$g(Z, f) = \nabla_{\mathbf{u}} f \cdot \sum_{i=1}^{N_c} |\mathbf{M}^i| \tilde{Z}_i^{-3} (I_i - f) (u - C_i^1, v - C_i^2), \quad (2.9)$$

$$b(Z, f) = \sum_{i=1}^{N_c} \phi_i |\mathbf{M}^i| \tilde{Z}_i^{-3} ((u - C_i^1) \nu^u + (v - C_i^2) \nu^v).$$

Recall that, for the i -th camera, $\mathbf{C}_i = (C_i^1, C_i^2, C_i^3)^\top$ is the optical center, $|\mathbf{M}^i|$ is the product of the focal lengths and \tilde{Z}_i is the depth of the surface point with respect to the camera. The 3-D Laplacians $\Delta Z = Z_{uu} + Z_{vv} + Z_{\tau\tau}$ and, similarly, Δf arise from the regularizing terms (2.3) and (2.4), respectively. $\partial * / \partial \boldsymbol{\nu} = (\nabla *) \cdot \boldsymbol{\nu}$ is the usual notation for the directional derivative along $\boldsymbol{\nu} = (\nu^u, \nu^v, \nu^\tau)^\top$, the normal to the integration domain U_T in the parameter space. $\nabla_{\mathbf{u}} f$ is the gradient of f with respect to the original spatial variables only. To simplify the model, we replace (2.6) by homogeneous Neumann boundary conditions, as in the still-image case.

Having introduced the manifold reconstruction model, let us make it more flexible by including an extra parameter $\rho^2 \geq 0$ to control the amount of temporal regularization of the solution relative to the amount of spatial regularization. The effect of this parameter is the substitution of the 3-D Laplacian $\Delta Z = Z_{uu} + Z_{vv} + Z_{\tau\tau}$ by the weighted sum $Z_{uu} + Z_{vv} + \rho^2 Z_{\tau\tau}$, and similarly for Δf . There are two possible interpretations to this anisotropic diffusion operator. The simplest one is that the operator arises by replacing the integrand in (2.3) by $\frac{1}{2}(Z_u^2 + Z_v^2 + \rho^2 Z_\tau^2)$, thus using a weighted

norm instead of the Euclidean norm (and similarly for (2.4)). The second interpretation is that the anisotropic diffusion operator arises by using the Euclidean norm (2.3) in a deformed space where variable τ is scaled by ρ with respect to variables (u, v) to yield the desired non-uniform scaling of the gradients and the Laplacians, $\nabla Z, \nabla f, \Delta Z$ and Δf .

3 Numerical solution

An iterative method is used to find the minimum of energy (2.1) via the solution of the coupled system of equations that arise from the the necessary optimality condition of the energy. Observe that, for a fixed height, equation (2.7) is a linear PDE in the radiance, which is easier to solve than the non-linear PDE in the height (2.5) for a fixed radiance. By exploiting this asymmetry one may devise a minimization strategy consisting of a nested iterative scheme: an outer loop performing a gradient descent in the height, and an inner loop implementing a direct optimization for the radiance. The scheme is initialized by an approximate solution, usually consisting of the zero height function and the corresponding optimal radiance without regularizer ($\beta = 0$). The linear PDE in the radiance is faster to solve using classical stationary iteration methods such as Jacobi or Gauss-Seidel rather than setting up a gradient descent equation and using time-stepping¹ solvers. The PDE in the height is more complicated and it is solved via time-stepping methods.

Numerical discretization of the PDEs is carried out using Finite Difference Methods (FDM). Therefore, the integration domain $U_{\mathcal{T}}$ is discretized by means of a 3-D grid. Forward differences in fictitious time and central differences in (u, v, τ) approximate the derivatives in the PDEs. The quantity $\tilde{\rho} = \rho h / \Delta\tau$ can be interpreted as the step ratio that states the relationship between the grid steps $h = \Delta u = \Delta v$ and $\Delta\tau$ in the anisotropic space previously mentioned.

Both updating schemes (stationary method for f and time-stepping method for Z) are used as relaxation procedures inside a 3-D multigrid method [3, 15] that approximately solves the EL equations. Multigrid methods are the most efficient numerical tools for solving elliptic boundary value problems. Each relaxation iteration implies the update of the values of the approximate solution (either Z or f) at *all* the grid points, that is, all physical time slices of the 3-D grid are updated in the same iteration. This simultaneous processing property is the main difference with respect to the sequential scheme [5].

Due to the extra dimension added to the problem, memory becomes a precious resource. Storage is required not only for the unknowns Z and f at the different levels of multigrid, but also for the coefficients of the linear/linearized PDEs and other temporary variables. In-place updates of the unknowns are preferred because they allow longer sequences of images to be processed for the same amount of memory.

Full Multigrid [3] with zero initial condition is used to initialize the 3-D multigrid solver. It is also possible to initialize it with the solution from a fast sequential 2-D multigrid, possibly computed using a coarser version of the problem.

Comparison to sequential 2-D multigrid

Although the sequential and manifold reconstruction methods (SR and MR, respectively) share a common base theoretical modeling, their major difference is that SR relies on 2-D multigrid (snapshot-wise), whereas MR is based on 3-D multigrid. A comparison of the speed performance between both methods is not the focus of this work, but we discuss this topic of practical relevance. In general, we may conclude in a first analysis that SR is simpler and faster but more sensitive to noise than MR.

The two main components of any multigrid method are the integrid transfer operators and the smoothers to update the approximate solution. Let us compare both qualitatively for SR and MR. On the one hand, 3-D restriction is faster than sequential 2-D restriction because each coarser level of 3-D multigrid has roughly half the number of snapshots of the previous level, whereas in sequential 2-D multigrid the number of snapshots remains constant and maximal for all levels. On the other

¹artificial time, not the physical time τ .

hand, assuming full weighting restriction, each point in the coarse grid requires a weighted sum of 9 points in case of 2-D restriction and 27 points in case of 3-D restriction. Thus, each point-wise 3-D restriction involves more neighbors and calculations than each 2-D restriction. Similar comments apply to the prolongation operators: trilinear interpolation is applied to a smaller number of grid points than sequential bilinear interpolation, but the former is more expensive point-wise. For details on 2-D and 3-D inter-grid transfer operators, see [15].

The update of each grid point in MR is slower than the update of each point in SR because the former depends on more neighbors (temporal derivatives and 3-D Laplacian). In addition, SR only performs full multigrid for the reconstruction of the first snapshot (and forward propagation of the solution to initialize the remaining snapshots), whereas MR carries out full multigrid for all snapshots. Moreover, the artificial time step to evolve the non-linear gradient descent PDE in the height is smaller in MR than in SR because the former takes the minimum step over all snapshot-wise steps of the latter. Hence, MR requires more iterations than SR to reach the same final, artificial time. On the other hand, MR is more robust than SR because it is less prone to be trapped in a local minimum caused by the poor reconstruction of a snapshot. Errors are distributed better among all snapshots in MR than in SR.

4 Experiments

Our variational method has been tested on stereo video acquired at an off-shore platform near the southern seashore of the Crimean peninsula, in the Black Sea. Two cameras mounted 12 m above the mean sea level and 2.5 m apart provide images of 1624×1236 pixels, acquired at a frame rate of 10 Hz.

Current implementation of the MR method requires all relevant information (input images, processing grids, etc.) to be stored in the computer's main memory (RAM). Thus, the limit in the amount of available main memory in the computer performing the reconstruction imposes a constraint on the size of the achievable reconstruction. For a fixed amount of memory, there is a trade-off between the spatial and temporal sizes of the grid within the MR method: one can either reconstruct a few snapshots with high spatial resolution or a larger number of snapshots at a coarser spatial resolution.

For example, in a computer with 2 GBytes of RAM, a 6-level multigrid solver can roughly handle a reconstruction on a spatial grid of 513×513 points (resolution $h = 2.5$ cm) on 65 snapshots of the original dataset or a reconstruction on a grid of 129×129 points ($h = 10$ cm) on 1025 snapshots of a $\times 4$ coarser version of the images. Fig. 4.1 shows an example of the reconstructions obtained at both spatial resolutions previously mentioned. This figure shows that the reconstruction at the low resolution captures the gist (i.e., low spatial frequency components) of the wave heights, thus resembling the reconstruction at the high spatial resolution.

The MR algorithm acts on the entire sequence of images by reconstructing pieces (sub-sequences) of consecutive snapshots. Other coarse-to-fine strategies are also possible to ensure a smooth transition near the temporal boundaries of the sub-sequences.

Snapshots can also be decimated in time. The linear dispersion relation that relates spatial and temporal frequencies of waves in deep water, $k = \omega^2/g$ (where $k = 2\pi/\lambda$ is the wave number, $\omega = 2\pi f$ is the angular frequency and g is the gravity acceleration), serves as a physical criterion to choose a reasonable frame rate for a given spatial resolution, and vice versa. For instance, in the example with grid resolution $h = 10$ cm, assuming the minimum spatial wavenumber that the algorithm reconstructs reliably is $\lambda = 4h$, the corresponding frequency of the wave is

$$f = \frac{\omega}{2\pi} = \frac{1}{2\pi} \sqrt{g \frac{2\pi}{\lambda}} = \sqrt{\frac{g}{2\pi 4h}} \approx 2 \text{ Hz}.$$

If this is the maximum temporal frequency that the algorithm reconstructs reliably, the corresponding Nyquist rate (minimum sampling rate required to avoid aliasing) is twice as much, 4 Hz, meaning that snapshots should be at most $\Delta\tau \leq 1/4 = 0.25$ seconds apart. In the experiments, we used a smaller resolution $\Delta\tau = 0.1$ s since the acquisition rate (10 Hz) allowed so, but we could have

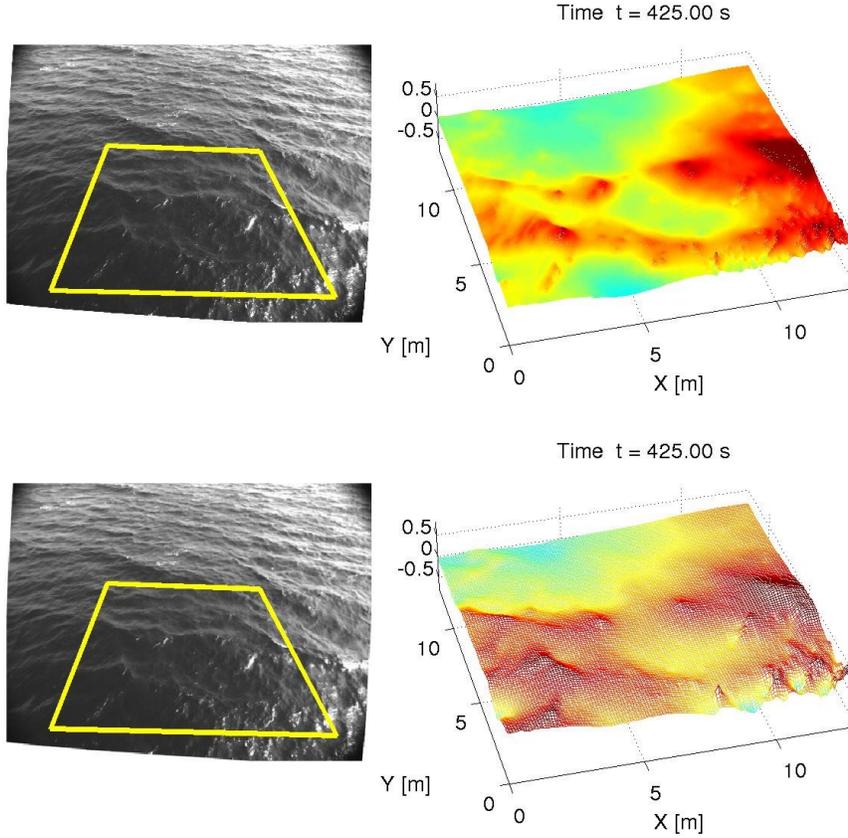


Figure 4.1: Right: reconstruction of a snapshot at spatial resolutions $h = 2.5$ cm (top) and $h = 10$ cm (bottom). Wave height is pseudo-colored, from blue (low) to red (high). Left: one of the two input images at the corresponding resolution (top: 1624×1236 pixels, bottom: 406×309 pixels), with highlighted reconstructed region of interest.

used $\Delta\tau = 0.2$ s (i.e., temporal snapshot decimation by a factor of 2) to achieve results with similar validity and interpretation, albeit expanding twice the physical time interval.

4.1 Comparison to the sequential reconstruction method

Experiments were carried out using MR for a sequence of 4100 consecutive snapshots, split in subsequences of 1025 snapshots. A 6-level full multigrid method with 1000 iterations per level and 1500 iterations at the finest level was performed, with 2 V-cycles per iteration and one pre- and post-relaxation sweeps per level. The weights $\alpha = 0.1$ and $\beta = 0.025$ were empirically determined. The initial manifold surface consisted of the zero height surface (for all snapshots). Different values of the parameter ρ were tested: $\tilde{\rho} = \rho h / \Delta\tau = \{0.1, 0.2, 0.5, 1\}$. Observe that the temporal coherence of the MR decreases as $\rho \rightarrow 0$. In the limit, $\rho = 0$ is equivalent to the reconstruction of each snapshot independently, using full multigrid on each of them, but with a common artificial time step.

Figs. 4.2 through 4.4 show slices of the surface height and radiance functions obtained by the MR method. The computational grid has $129 \times 129 \times 1025$ points, with spatial and temporal resolutions $h = 10$ and $\Delta\tau = 0.1$ s, respectively. In particular, the results correspond to the case $\tilde{\rho} = 0.1$. Observe the oscillating patterns of the ocean waves in both the height and radiance functions of the u and v slices (Figs. 4.2 and 4.3). The radiance function also captures the location and motion of sea foam caused by breaking waves. Fig 4.4 gives the reconstructed surface height and radiance for a particular snapshot of the sequence.



Figure 4.2: A slice at constant $u = u_0$. Top: surface height $Z(u_0, v, \tau)$ (grayscale encoded, from dark (low) to white (high)). Bottom: surface radiance $f(u_0, v, \tau)$. Horizontal axis is time τ .

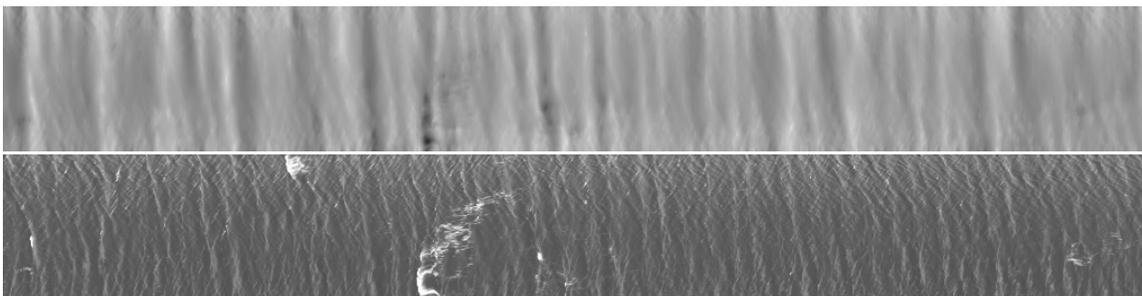


Figure 4.3: A slice at constant $v = v_0$. Top: surface height $Z(u, v_0, \tau)$ (grayscale encoded, from dark (low) to white (high)). Bottom: surface radiance $f(u, v_0, \tau)$. Horizontal axis is time τ .

By visual inspection of the reconstruction (see Fig. 4.5), one concludes that the values $\tilde{\rho} = \{0.5, 1\}$ are too large: temporal derivatives are penalized too much with respect to spatial derivatives, yielding a reconstructed surface shape that is very smooth in time and does not capture the wave patterns present in the stereo video data. Drawing an analogy with linear signal processing, the anisotropic diffusion carried out by the (weighted) Laplacian operator has a low-pass filtering effect: it limits the temporal bandwidth of the output signal, thus reducing noise but also destroying the desired wave signal.

Using the following formulas to measure the photometric quality of the modeled image given by the reconstructed surface at a single snapshot,

$$\tilde{E}_{\text{data}} = \frac{E_{\text{data}}}{N} \approx \frac{h^2}{N} \sum_{i=1}^{N_c} \sum_{j=1}^N \phi_i(\mathbf{u}_j) \mathbf{J}_i(\mathbf{u}_j), \quad (4.1)$$

$$\text{RMS} \left[\frac{E_i}{\text{Area}(\Omega_i)} \right] \approx \left(\frac{\tilde{E}_{\text{data}}}{2} \frac{N}{\text{Area}(\Omega_i)} \right)^{\frac{1}{2}} = \left(\frac{N}{2r \text{Area}(I_i)} \right)^{\frac{1}{2}} \sqrt{\tilde{E}_{\text{data}}}, \quad (4.2)$$

it is possible to measure the photometric quality of the modeled images given by the reconstructed space-time surface, i.e. for all snapshots. Table 1 compares the photometric error of the reconstructed sequence of wave heights obtained by means of the SR and MR methods with different values of the parameter ρ . Observe that the magnitude of the photometric error is small and very similar for all values of ρ , however the differences in reconstructed surface shape are more noticeable by means of other methods such as the statistics of the wave heights.

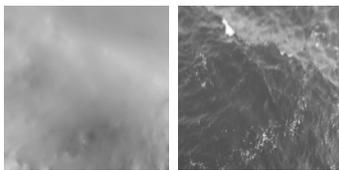


Figure 4.4: A slice at constant $\tau = \tau_0$. Size: 129×129 grid points. Left: surface height $Z(u, v, \tau_0)$ (grayscale encoded, from dark (low) to white (high)). Right: surface radiance $f(u, v, \tau_0)$.

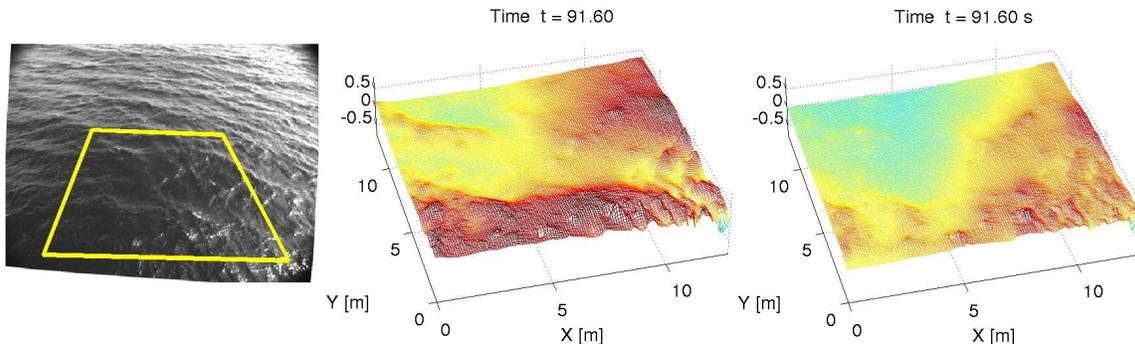


Figure 4.5: Manifold reconstruction at one of the snapshots. Left: a coarse input image of the stereo pair. Center: reconstructed wave height with $\tilde{\rho} = 0.1$. Right: reconstructed wave height with $\tilde{\rho} = 1.0$.

Method	mean	standard deviation
Sequential	3.613	0.398
Manifold, $\tilde{\rho} = 0.1$	3.615	0.393
Manifold, $\tilde{\rho} = 0.2$	3.621	0.391
Manifold, $\tilde{\rho} = 0.5$	3.657	0.392

Table 1: Comparison of photometric error (4.2) for several methods used to reconstruct the same coarse stereo image sequence: the variational graph sequential method (SR) [5] and the variational graph manifold method (MR).

The spectral and statistical analysis carried out in [5] can be repeated for the manifold method to validate the resulting reconstruction (see Figure 4.6). No significant differences with respect to the SR method are reported except for the fact that the MR method shows an improvement in the estimation of the crest-trough asymmetry of wave heights (see Fig. 4.7).

5 Conclusion and future work

A variational graph manifold method for the space-time coherent reconstruction of ocean waves has been discussed and developed. Due to the convenient representation of the ocean surface as a height function, the incorporation of a temporal dimension is straightforward. This benefit of the graph representation makes the reconstruction problem significantly simpler than what it would be if the level set framework was used instead. The manifold method has been tested on stereo video data from real ocean waves at an off-shore platform in the Black Sea. The qualitative and photometric performance of the method has been demonstrated on a coarser version of the dataset due to memory limitations of the implementation with respect to the size of the original dataset. By design, the manifold reconstruction method is more robust than the sequential method, at the expense of speed. In addition, it captures better some physical properties of wave heights than the sequential method,

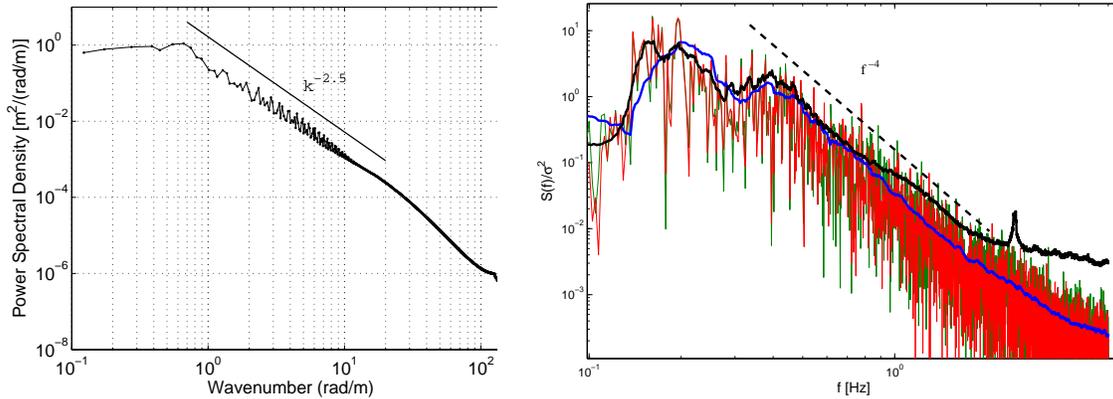


Figure 4.6: Left: Omni-directional wave number spectrum. Right: Normalized frequency spectrum (σ^2 is the variance of the wave surface) averaged over several virtual probes in time (blue line) and compared to the spectrum by the epipolar method [2].

such as the crest-trough asymmetry of wave heights.

Evidence shows that the incorporation of the physics and coherence in our variational method produces tangible improvements and encourages us to continue with this line of research to achieve better results. These insights justify the research on the incorporation of the wave equation in the reconstruction process. Finally, the manifold view developed in this paper can also be applied to the variational disparity method in [1]. This topic may be further investigated in the future.

References

- [1] L. Alvarez, R. Deriche, J. Sánchez, and J. Weickert. Dense disparity map estimation respecting image discontinuities : A pde and scale-space based approach. *Journal of Visual Communication and Image Representation*, 13:3–21, 2002.
- [2] Alvise Benetazzo. Measurements of short water waves using stereo matched image sequences. *Coastal Engineering*, 53(12):1013–1032, December 2006.
- [3] William L. Briggs, Van Emden Henson, and Steve F. McCormick. *A Multigrid Tutorial, Second Edition*. SIAM, 2000.
- [4] F. Fedele, G. Gallego, A. Yezzi, and A. Benetazzo. Understanding extreme waves via a variational wave acquisition stereo system. In *Workshop ROGUE WAVES*, 2008. Oct 13-15, Brest, France.
- [5] G. Gallego, A. Yezzi, F. Fedele, and A. Benetazzo. A variational stereo method for the three-dimensional reconstruction of ocean waves. *Geoscience and Remote Sensing, IEEE Transactions on*, PP(99):1–13, 2011.
- [6] G. Gallego, A. Yezzi, F. Fedele, and A. Benetazzo. A variational wave acquisition stereo system for the 3-d reconstruction of oceanic sea states. In *ASME 30th International Conference on Offshore Mechanics and Arctic Engineering*, 2011. Rotterdam, The Netherlands.
- [7] Guillermo Gallego. *Variational Image Processing Algorithms for the Stereoscopic Space-Time Reconstruction of Water Waves*. PhD thesis, Georgia Institute of Technology, Atlanta, GA, USA, 2011. Directors: Yezzi, Anthony and Fedele, Francesco.
- [8] K.T. Holland, R.A. Holman, T.C. Lippmann, J. Stanley, and N. Plant. Practical use of video imagery in nearshore oceanographic field studies. *Oceanic Engineering, IEEE Journal of*, 22(1):81–92, jan. 1997.

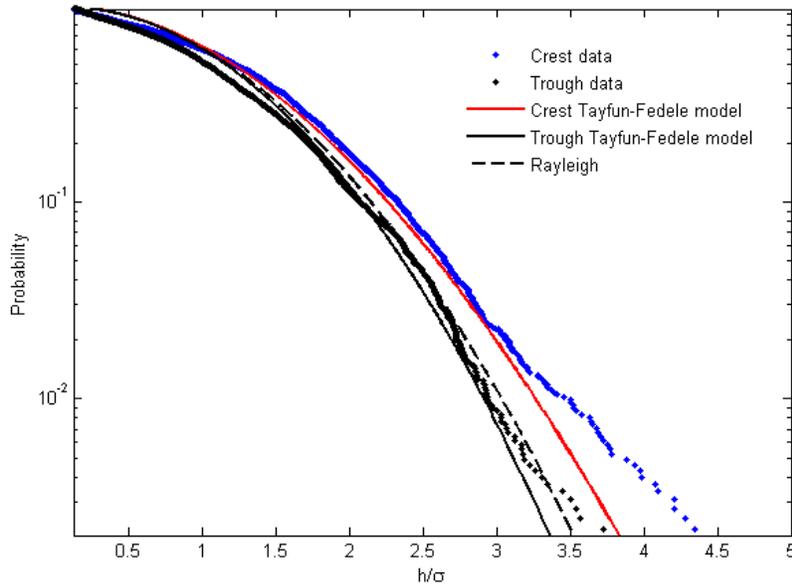


Figure 4.7: Empirical exceedance probabilities of wave crests and troughs from time series extracted at virtual probes of the manifold reconstruction with $\tilde{\rho} = 0.1$. Rayleigh and Tayfun-Fedele theoretical models are also shown for comparison.

- [9] Y.-H. Hsiao and M.-C. Huang. Application of active contour model in tracking sequential nearshore waves. *China Ocean Engineering*, 23(12):251–266, 2009.
- [10] K.R. MacHutchon and P.C. Liu. Measurement and analysis of ocean wave fields in four dimensions. In *ASME 26th International Conference on Offshore Mechanics and Arctic Engineering*, volume 1, pages 923–927, 2007.
- [11] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and variational problems. *Comm. Pure Appl. Math*, XLII(5):577–685, 1988.
- [12] F. Santel, W. Linder, and C. Heipke. Stereoscopic 3d-image sequence analysis of sea surfaces. *Proceedings of the ISPRS Commission V Symposium*, 35(5):708–712, 2004.
- [13] A. Schumacher. Stereophotogrammetrische wellenaufnahmen. wiss ergeb. dtsh. atlant. exped. forschung vermessung, meteor 1925-1927. *Ozeanographische Sonderuntersuchungen, Erste Lieferung*, 1939.
- [14] O.H. Shemdin, H.M. Tran, and S.C. Wu. Directional measurements of short ocean waves with stereography. *J. Geophys. Res.*, 93:13891–13901, 1988.
- [15] Ulrich Trottenberg. *Multigrid*. Academic Press, 2000.
- [16] Justin M. Wanek and Chin H. Wu. Automated trinocular stereo imaging system for three-dimensional surface wave measurements. *Ocean Engineering*, 33(5-6):723–747, April 2006.