

Weak Statistical Constraints for Variational Stereo Imaging of Oceanic Waves

Guillermo Gallego¹, Anthony Yezzi¹,
Francesco Fedele^{2,*}, and Alvisè Benetazzo³

¹ School of Electrical Engineering, Georgia Institute of Technology, Atlanta, USA

² School of Civil Engineering, Georgia Institute of Technology, Savannah, USA

³ CNR-ISMAR, Venice, Italy

Abstract. We develop an observational technique for the stereoscopic reconstruction of the wave form of oceanic sea states via a variational stereo method. In the context of active surfaces, the shape and radiance of the wave surface are obtained as minimizers of an energy functional that combines image observations and smoothness priors. To obey the quasi Gaussianity of oceanic waves observed in nature, a given statistical wave law is enforced in the stereo variational framework as a weak constraint. Multigrid methods are then used to solve the partial differential equations derived from the optimality conditions of the augmented energy functional. An application of the developed method to two sets of experimental stereo data is finally presented.

1 Introduction

In recent years there has been a growing interest in vision-based remote-sensing observational technology for the measurement of oceanic sea states [7,2,17,5]. This topic is a major concern in ocean engineering because it has a broad impact: the understanding of space-time dynamics of ocean waves enables better forecasting of extreme events, improved design of off-shore structures, validation of theoretical models, etc. Vision systems are non-intrusive, have economical advantages over traditional instrumentation (wave gauges and ultrasonic instruments or buoys) and provide spatio-temporal data whose statistical content is richer than that of previous monitoring methods, but they require more processing power to extract information from the observed video data. The application of vision tools, such as stereography, to oceanography dates back to the first experiments with stereo cameras mounted on a ship by Schumacher [12] in 1939. Stereography gained popularity in studying the dynamics of oceanographic phenomena during the 1980s due to advances in hardware. For example, Shemdin et al. [14] applied stereography for the directional measurement of short ocean waves. Recently, Benetazzo [2] successfully incorporated epipolar techniques in the Wave Acquisition Stereo System (WASS) and showed that the accuracy of

* Research supported by ONR grant BAA 09-012: “Ocean Wave Dissipation and energy Balance (WAVE-DB):toward reliable spectra and first breaking statistics”.

WASS is comparable to the accuracy obtained from traditional instrumentation. An alternative trinocular imaging system (ATSIS) for measuring the temporal evolution of 3-D surface waves was proposed in [17].

The three-dimensional reconstruction of an object's surface from multiple images is a classical problem in computer vision [10,6,13], and it is still an extremely active research area. There are many 3-D reconstruction algorithms available in the literature and they are designed under different assumptions that provide a variety of trade-offs between speed, accuracy and viability. Traditional *image-based* stereo methods typically consist of two steps: first, image points are detected and matched across images to establish local correspondences; then depth is inferred by back-projection of correspondences. This is the strategy used in recent observational systems [2,17], and it has the advantages of being simple and fast. However, it also has some major disadvantages that motivated the research on improved stereo reconstruction methods [4,18,8] based on variational theory. Firstly, correspondences rely on strong textures and image matching. They can be poorly estimated if the objects in the scene have a smooth radiance, and can also suffer from the presence of noise and local minima. Furthermore, each space point is reconstructed independently. Therefore, the recovered surface of an object is obtained as a collection of scattered 3-D points. Thus, the hypothesis of the continuity of the surface is not exploited in the reconstruction process. The breakdown of traditional stereo methods in these situations is evidenced by "holes" in the reconstructed surface, which correspond to unmatched image regions [10,2]. This phenomenon may be dominant in the case of the ocean surface, which, by nature, is generally continuous and contains little texture.

Modern *object-based* computer vision methods that rely on Calculus of Variations and Partial Differential Equations (PDE), are able to overcome the disadvantages of traditional stereo [4,18,1,8]. For instance, unmatched regions are avoided by building an explicit model of the smooth surface to be estimated rather than representing it as a collection of scattered 3-D points. Thus, variational methods provide dense and coherent surface reconstructions. Surface points are reconstructed by exploiting the continuity (coherence) hypothesis in the full two-dimensional domain of the surface.

Variational stereo methods combine correspondence establishment and shape reconstruction into one single step and they are less sensitive to matching problems of local correspondences. The reconstructed surface is obtained by minimization of an energy functional designed for the stereo problem. The solution is obtained in the context of active surfaces by deforming an initial surface via a gradient descent PDE derived from the necessary optimality conditions of the energy functional, the so-called Euler-Lagrange (EL) equations. In parallel to the advances in vision tools, the oceanographic community has developed statistical and spectral models for the characterization of oceanic sea states [9,19,15,16] that clearly indicate that oceanic waves are quasi-Gaussian in nature.

Up to date, both traditional and variational stereo techniques do not include in the reconstruction process the prior information of Gaussianity of waves, which is usually verified a posteriori [2,5]. In this paper, we present a novel

variational framework in which a statistical distribution is enforced as a prior into the stereo reconstruction of water waves via a weak constraint. Motivated by the characteristics of the target object in the scene, i.e., the ocean surface, we first introduce the graph surface representation in the formulation of the reconstruction problem. Then, we cast the problem as a variational optimization problem and show how *a priori* knowledge of statistical wave height models can be weakly enforced in the variational framework to aid the recovery of the surface shape. The performance of the algorithm is validated on experimental data and the statistics of the reconstructed surface are also analyzed. Concluding remarks are finally presented.

2 The Variational Framework

2.1 Multi-Image Setup and Graph Surface Representation

Let S be a smooth surface in \mathbb{R}^3 with generic local coordinates $(u, v) \in \mathbb{R}^2$. Let $\{I_i\}_{i=1}^{N_c}$ be a set of images of a static scene acquired by cameras with known calibration parameters $\{P^i\}_{i=1}^{N_c}$. Space points are mapped into image points according to the pinhole camera model [6]. A surface point (or, in general a 3-D point) $\mathbf{X} = (X, Y, Z)^\top$ with homogeneous coordinates $\bar{\mathbf{X}} = (X, Y, Z, 1)^\top$ is mapped to point $\mathbf{x}_i = (x_i, y_i)^\top$ in the i -th image with homogeneous coordinates $\bar{\mathbf{x}}_i = (x_i, y_i, 1)^\top \sim P^i \bar{\mathbf{X}}$, where the symbol \sim means equality up to a nonzero scale factor and $P^i = K^i [R^i \mid \mathbf{t}^i]$ is the 3×4 projection matrix with the intrinsic (K^i) and extrinsic (R^i, \mathbf{t}^i) calibration parameters of the i -th camera. Point $\mathbf{C}_i = (C_i^1, C_i^2, C_i^3)^\top$ satisfying $P^i \bar{\mathbf{C}}_i = \mathbf{0}$ is the optical center of the i -th camera. Let $\pi_i : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ note the projection maps, $\mathbf{x}_i = \pi_i(\mathbf{X})$, and $I_i(\mathbf{x}_i)$ be the image intensity at \mathbf{x}_i .

In the variational context of active surfaces, we present a different approach to the reconstruction problem presented in [18,4] (level set approach) by exploiting the hypothesis that the surface of the water can be represented in the form of a graph or elevation map:

$$Z = Z(X, Y), \quad (1)$$

where Z is the height of the surface with respect to a domain plane that is parameterized by coordinates X and Y . Indeed, slow varying, non-breaking waves admit this simple representation with respect to a plane orthogonal to gravity direction. The graph representation of the water surface presents some clear advantages over the more general level set representation of [4,8,18,5]. Surface evolution is simpler to implement since the surface is not represented in terms of an auxiliary higher dimensional function (the level set function). The surface is evolved directly via the height function (1) discretized over a fixed 2-D grid defined on the $X - Y$ plane. The latter also implies that for the same amount of physical memory, higher spatial resolution (finer details) can be achieved in the graph representation than with the level set. The $X - Y$ plane becomes the natural common domain to parameterize the geometrical and photometric properties of surfaces. This simple identification does not exist in the level set

approach [18]. Finally, the graph representation allows for fast numerical solvers besides gradient descent, like Fast Poisson Solvers, Cyclic Reduction, Multigrid Methods, Finite-Element Methods (FEM), etc. In the level set framework, the range of solvers is not as diverse.

However, there are also some minor disadvantages. A world frame properly oriented with the gravity direction must be defined in advance to represent the surface as a graph with respect to this plane. This is not trivial *a priori* and might pose a problem if only the information from the stereo images is used [2]. Surface evolution is constrained to be in the form of a graph and this may differ from the evolution obtained for an unconstrained surface. As a result, more iterations may be required to evolve the active surface to reach convergence.

2.2 Proposed Energy Functional

Consider the 3-D reconstruction problem from a collection of $N_c \geq 2$ images (we will exemplify with $N_c = 2$). We investigate a generative model of the images that allows for the joint estimation of the shape of the surface S and the radiance function on the surface f as minimizers of an energy functional. Let the energy functional be the sum of a data fidelity term E_{data} and two regularizing terms: a geometry smoothing term E_{geom} and a radiance smoothing term E_{rad} ,

$$E(S, f) = E_{\text{data}}(S, f) + \alpha E_{\text{geom}}(S) + \beta E_{\text{rad}}(f), \quad (2)$$

where $\alpha, \beta \in \mathbb{R}^+$. The data fidelity term measures the photo-consistency of the model: the discrepancy between the observed images I_i and the radiance model f ,

$$E_{\text{data}} = \sum_{i=1}^{N_c} E_i, \quad E_i = \int_{\Omega_i} \phi_i \, d\mathbf{x}_i, \quad (3)$$

where a possible photometric matching criterion is

$$\phi_i = \frac{1}{2} (I_i(\mathbf{x}_i) - f(\mathbf{x}_i))^2. \quad (4)$$

The region of the image domain where the scene is projected is denoted by Ω_i . Assuming that the surface of the scene is represented as a graph $Z = Z(u, v)$, a point on the surface has coordinates

$$\mathbf{X}(u, v) = (u, v, Z(u, v))^{\top}. \quad (5)$$

The chain of operations to obtain the intensity $I_i(\mathbf{x}_i)$ given a surface point with world coordinates $\mathbf{X}(\mathbf{u}) \equiv S(\mathbf{u})$, $\mathbf{u} = (u, v)^{\top}$, is

$$\mathbf{X}(\mathbf{u}) \mapsto \tilde{\mathbf{X}}_i = \mathbf{M}^i \mathbf{X} + \mathbf{p}_4^i \mapsto \mathbf{x}_i \mapsto I_i(\mathbf{x}_i), \quad (6)$$

where $\tilde{\mathbf{X}}^i = (\tilde{X}_i, \tilde{Y}_i, \tilde{Z}_i)^{\top}$ are related to the coordinates of \mathbf{X} in the i -th camera frame, $\mathbf{x}_i = (\tilde{X}_i/\tilde{Z}_i, \tilde{Y}_i/\tilde{Z}_i)^{\top}$ is the projection of \mathbf{X} in the i -th image plane and $\mathbf{P}^i = [\mathbf{M}^i \mid \mathbf{p}_4^i]$, with $\mathbf{M}^i = \mathbf{K}^i \mathbf{R}^i \equiv (\mathbf{n}_1^i, \mathbf{n}_2^i, \mathbf{n}_3^i)^{\top}$ and $\mathbf{p}_4^i = \mathbf{K}^i \mathbf{t}^i$. Also, $|\mathbf{M}^i| = \det(\mathbf{M}^i)$.

The radiance model f is specified by a function \hat{f} defined on the surface S . Then, f in (4) is naturally defined by $f(\mathbf{x}_i) = \hat{f}(\pi_i^{-1}(\mathbf{X}))$, where π_i^{-1} denotes the back-projection operation from a point in the i -th image to the closest surface point with respect to the camera. By abusing notation, let us use f to denote the parameterized radiance $f(\mathbf{u})$, understanding that $f(\mathbf{x}_i)$ in (4) reads the back-projected value in $\hat{f}(\mathbf{X}(\mathbf{u})) = f(\mathbf{u})$.

Motivated by the common parameterizing domain of the shape and radiance of the surface and to obtain the simplest diffusive terms in the necessary optimality conditions of the energy (2), let the regularizers be

$$E_{\text{geom}} = \int_U \frac{1}{2} \|\nabla Z(\mathbf{u})\|^2 d\mathbf{u}, \quad E_{\text{rad}} = \int_U \frac{1}{2} \|\nabla f(\mathbf{u})\|^2 d\mathbf{u}, \quad (7)$$

where $\nabla Z(\mathbf{u}) = (Z_u, Z_v)^\top$, $\nabla f(\mathbf{u}) = (f_u, f_v)^\top$ and subscripts indicate the derivative with respect to that variable.

The definition of the data fidelity term as an integral over the image domain (rather than over the parameter space U) has two advantages: (i) the data term is independent of the choice of domain for the graph, and (ii) the resulting optimality conditions for the minimization of (2) lack image derivatives, which are transferred to the radiance model and can be controlled by the regularizer E_{rad} . This desirable property is inherited from the modeling and mathematical principles that we follow from [18]. The resulting algorithm is less sensitive to image noise than other variational approaches for stereo 3-D reconstruction.

Once all terms in (2) have been specified, they are expressed over a common domain: the parameter space. The Jacobian of the change of variables between integration domains for the data term is, by applying the chain rule to (6),

$$\mathbf{J}_i = \left| \frac{d\mathbf{x}_i}{d\mathbf{u}} \right| = -|M^i| \tilde{Z}_i^{-3} (\mathbf{X} - \mathbf{C}_i) \cdot (\mathbf{X}_u \times \mathbf{X}_v), \quad (8)$$

where $\mathbf{X}_u \times \mathbf{X}_v$ is proportional to the outward unit normal \mathbf{N} to the surface at $\mathbf{X}(u, v)$, and $\tilde{Z}_i = \mathbf{n}_3^i \cdot (\mathbf{X} - \mathbf{C}_i) > 0$ is the depth of the point \mathbf{X} with respect to the i -th camera (located at \mathbf{C}_i). With this change, energy (3) becomes

$$E_i = \int_{\Omega_i} \phi_i d\mathbf{x}_i = \int_U \phi_i \mathbf{J}_i d\mathbf{u}, \quad (9)$$

where the last integral is over U : the part of the parameter space whose surface projects on Ω_i in the i -th image. Observe that the Jacobian weights the photometric error ϕ_i proportionally to the cosine of the angle between the unit normal to the surface at \mathbf{X} and the *projection ray* (the ray joining the optical center of the camera and \mathbf{X}): $(\mathbf{X} - \mathbf{C}_i) \cdot (\mathbf{X}_u \times \mathbf{X}_v)$. After collecting terms (7) and (9), and noting that the shape \mathbf{X} of the surface solely depends on the height (Eqn. (5)), energy (2) becomes

$$E(Z, f) = \int_U L(Z, \nabla Z, f, \nabla f, u, v) d\mathbf{u}. \quad (10)$$

where subscripts indicate the derivative with respect to that variable, and the integrand is the so-called *Lagrangian* L .

2.3 Energy Minimization. Optimality Condition

The energy (10) depends on two functions: the shape Z and the radiance f of the surface. To find a minimizer of such a functional, we derive the necessary optimality condition by setting to zero the first variation of the functional, yielding a coupled system of PDEs (EL equations) along with boundary conditions:

$$g(Z, f) - \alpha \Delta Z = 0 \quad \text{in } U, \quad (11)$$

$$b(Z, f) + \alpha \frac{\partial Z}{\partial \boldsymbol{\nu}} = 0 \quad \text{on } \partial U, \quad (12)$$

$$-\sum_{i=1}^{N_c} (I_i - f) \mathbf{J}_i(Z) - \beta \Delta f = 0 \quad \text{in } U, \quad (13)$$

$$\beta \frac{\partial f}{\partial \boldsymbol{\nu}} = 0 \quad \text{on } \partial U, \quad (14)$$

where the non-linear terms due to the data fidelity energy are

$$g(Z, f) = \nabla f \cdot \sum_{i=1}^{N_c} |\mathbf{M}^i| \tilde{Z}_i^{-3} (I_i - f) (u - C_i^1, v - C_i^2), \quad (15)$$

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

The Laplacians ΔZ and Δf arise from the regularizing terms (7), and $\partial^* / \partial \boldsymbol{\nu}$ is the directional derivative along $\boldsymbol{\nu} = (\nu^u, \nu^v)^\top$, the normal to the integration domain U in the parameter space. A simple classification of the PDEs can be done as follows. For a fixed shape, (13) and (14) form a linear elliptic PDE (of the inhomogeneous Helmholtz type) with Neumann boundary conditions. On the other hand, for a fixed radiance, (11) and (12) lead to a nonlinear elliptic equation in the height Z with nonstandard boundary conditions.

Difficult EL equations, such as (11)-(14), are commonly solved by the steady-state of gradient descent PDEs that evolve the unknown functions in artificial time t . This is the context of the so-called active surfaces. Due to the asymmetry in the complexity of the PDEs, a minimization strategy consisting of a nested iterative scheme is proposed: an outer loop performs a gradient descent in the height, and an inner loop implements a direct optimization for the radiance. Starting from an initial approximate solution, there are two phases within each iteration: (1) compute the optimal radiance for a fixed shape, and (2) evolve the shape, leaving the radiance fixed. To simplify the equations, we approximate the boundary condition (12) by a simpler, homogeneous Neumann boundary condition. This can be interpreted as if the data fidelity term vanished close to the boundary and it is a reasonable assumption since the major contribution to the energy is given by the terms in U , not at the boundary.

Numerical Solution. The optimality PDEs are discretized on a rectangular 2-D grid in the parameter space and then solved numerically using finite-difference methods. Direct optimization of the radiance is achieved by using stationary iterative methods (Jacobi or Gauss-Seidel). Forward differences in time and central differences in space approximate the derivatives in the gradient descent PDE for the height, yielding an explicit updating scheme. The von Neumann stability

analysis of the linearized PDE yields a time step $\Delta t \leq 1/(\frac{4\alpha}{h^2} + \frac{1}{2} \max |\dot{g}(Z)|)$, where $\dot{g}(Z)$ is the derivative of (15) and the maximum is taken over the 2-D discretized grid at current time t . The time step may change at every iteration.

Both updating schemes (stationary methods for f and the time-stepping method for Z) are used as relaxation procedures inside a multigrid method [3] that approximately solves the EL equations. Multigrid methods are the most efficient numerical tools for solving elliptic boundary value problems.

3 Weak Enforcement of Wave Height Distributions

The flexibility of the variational framework allows us to incorporate properties of the *physics* of the waves in the model that would be otherwise difficult to take into account in image-based stereo methods. For example, we may include global statistical properties in the form of a weak constraint by considering an extra energy term that penalizes the deviation of the statistics of the reconstructed surface with respect to some target statistics derived from a physical model. In particular, we may penalize the deviation of the height distribution of the water surface with respect to a physically-justified Gaussian model and drive the surface evolution toward (weakly) satisfying such a global property.

If $Z(u, v) = Z(\mathbf{u})$ is the height of the surface (wave) and it is interpreted as a random variable, then its cumulative distribution function (CDF) is

$$\text{cdf}^Z(Z_0) = P(Z \leq Z_0) = \frac{1}{A} \int_U H(Z_0 - Z(\mathbf{u})) d\mathbf{u},$$

where $H(\cdot)$ is the Heaviside function and $A = \int_U d\mathbf{u}$ is the area of the (fixed) domain of integration. Suppose (2) is augmented with an extra energy term $\gamma E_{\text{cdf}}(S)$, $\gamma > 0$, that measures the discrepancy between a target height CDF that we wish to enforce, $G(Z)$, and the experimental CDF of the height:

$$E_{\text{cdf}}(Z) = \int_{-\infty}^{\infty} \frac{1}{2} (G(\hat{z}) - \text{cdf}^Z(\hat{z}))^2 d\hat{z}. \quad (16)$$

To compute the first variation of (16), we can directly use the definition of the Gâteaux derivative or augment Z with an artificial time variable, $Z = Z(\mathbf{u}, t)$, so that the energy depends on t , differentiate with respect to this variable and exploit the relationship between both derivatives. Carrying out operations in the distributional sense,

$$\begin{aligned} \frac{d}{dt} E_{\text{cdf}} &= \frac{d}{dt} \int_{-\infty}^{\infty} \frac{1}{2} (G(\hat{z}) - \frac{1}{A} \int_U H(\hat{z} - Z(\mathbf{u})) d\mathbf{u})^2 d\hat{z} \\ &= \int_{-\infty}^{\infty} (G(\hat{z}) - \text{cdf}^Z(\hat{z})) \left(\frac{1}{A} \int_U \delta(\hat{z} - Z(\mathbf{u})) Z_t d\mathbf{u} \right) d\hat{z} \\ &= \int_U \frac{1}{A} \int_{-\infty}^{\infty} (G(\hat{z}) - \text{cdf}^Z(\hat{z})) \delta(\hat{z} - Z(\mathbf{u})) d\hat{z} Z_t d\mathbf{u} \\ &= \int_U \nabla_Z E_{\text{cdf}} Z_t d\mathbf{u}, \end{aligned}$$

where δ is the Dirac delta function and the gradient of (16) with respect to Z is

$$\nabla_Z E_{\text{cdf}}(Z(\mathbf{u})) = \frac{1}{A} (G(Z(\mathbf{u})) - \text{cdf}^Z(Z(\mathbf{u}))). \quad (17)$$

As a result of the statistical penalty, a new non-linear term of the form (17) appears in the EL equation (11), while the boundary condition remains unchanged. It is as if the nonlinear term (15) in the PDEs (11) was replaced by $g(Z) \leftarrow g(Z) + \gamma \nabla_Z E_{\text{cdf}}(Z)$. Multigrid methods are still suitable to efficiently solve the new non-linear PDE. However, the time-stepping smoother requires an additional constraint on the time step: the maximum height increment must be of the order of the bin size used to estimate the experimental CDF so that each iteration does not drastically change the CDF of the surface height.

Another reasonable energy to measure the statistical discrepancy between the empirical distribution of the wave field and the one dictated by the physical model is the L^2 difference between probability density functions (PDFs):

$$E_{\text{pdf}}(Z) = \int_{-\infty}^{\infty} \frac{1}{2} (\dot{G}(\hat{z}) - \text{pdf}^Z(\hat{z}))^2 d\hat{z}, \quad (18)$$

where $\dot{G}(Z)$ is the target PDF that we wish to enforce. Following similar steps as before, the EL equation (11) would have instead an extra term of the form

$$\nabla_Z E_{\text{pdf}}(Z(\mathbf{u})) = -\frac{1}{A} \frac{d}{dZ} (\dot{G}(Z(\mathbf{u})) - \text{pdf}^Z(Z(\mathbf{u}))). \quad (19)$$

Enforcing the statistical constraint via the L^2 difference of characteristic functions (i.e. the Fourier transform of the PDFs) is, by Parseval's theorem, equivalent to the above PDF approach.

Theoretical probabilistic models that can be used as target physical wave height distributions are presented in [15,16]. These models are quasi-Gaussian distributions that capture the asymmetry present in real life water waves, which have steep crests and shallow troughs.

4 Applications

Experiment 1. Images of "Canale della Giudecca" in Venice (Italy). Figs. 1 and 2 show an example of a reconstructed water surface from images of the Venice Canal. Cropped images in Fig. 1 are of size 600×450 pixels and show the region of interest to be reconstructed. Fig. 1 also displays the modeled images created by the generative model within our variational method. The data fidelity term compares the intensities of the original and modeled images in the highlighted region. As observed, the modeled image is a good match of the original image. Fig. 2 shows the converged values of the unknowns of the problem: the height and the radiance of the surface, as well as the 3-D representation of the reconstructed surface obtained by combining both 2-D functions. In this experiment, the values of the weights of the regularizers were empirically determined: $\alpha = 0.035$ and




Fig. 1. Left: projection on image 1 of the boundary of the estimated graph, which has been discretized by a grid of 129×513 points. Center: modeled image (computed from surface height and radiance) superimposed on original image 1. Right: modeled image 2 superimposed on original image 2.




Fig. 2. From left to right: (1) estimated height function $Z(u, v)$ (shape of the water surface) in pseudo-color; (2) height represented by greyscale intensities, from dark (low) to white (high); (3) estimated radiance function $f(u, v)$ (texture on the surface); (4) perspective, three-dimensional wire-frame representation of the estimated surface shape (height) according to grid points; (5) texture-mapped surface obtained by incorporating the radiance function in the wire-frame model. In (4) and (5) the vertical axis has been magnified by a factor of 5 with respect to the horizontal axes for visualization purpose.

$\beta = 0.01$. At the finest of the 5-level multigrid [3] algorithm, the gradient descent PDEs are discretized on a 2-D grid with 129×513 points. The distance between grid points is $h = 5$ cm. Therefore, the grid covers an area of $6.45 \times 25.65 \text{ m}^2$. An example of a surface discretized at the finest grid level is also shown in Fig. 2. The high density of the surface representation is typical of variational methods. The step size h must be chosen so that it approximately matches the resolution in the images: a displacement of 1 pixel is observable at the finest grid level in the multigrid framework and it corresponds to a physical displacement of at least h . Due to perspective, the maximum value of h is determined by the grid points closest to the cameras.

Experiment 2. We apply our variational method, with and without statistical regularizer, to a pair of stereo images acquired at an off-shore platform in the Black Sea. Two cameras mounted 12 m above the mean sea level and with a




Fig. 3. Original image (left), modeled image superimposed on original image (center), error image (right).




Fig. 4. Pseudo-colored $Z(u, v)$ without statistical regularizer (left). $Z(u, v)$ (center) and $f(u, v)$ (right) with statistical regularizer.

baseline of 2.5 meters acquire images of size 1624×1236 pixels. Fig. 3 (left) shows a sample image from one of the cameras. A grid with 513×513 points and resolution $h = 2.5$ cm, covering an area of $13 \times 13 m^2$, is used to discretize the graph of the surface. Roughly, 1 image pixel corresponds to a physical displacement of 1.06 cm (1.88 cm) for grid points near (resp. far from) the cameras. Both displacements are of the same order as h . A 6-level full multigrid method [3] with 400 iterations per level, 2 V-cycles per iteration, and 1 pre- and post-relaxation sweeps per cycle, is performed on the linearized optimality PDEs to reach a local solution. The weights of the regularizers used are: $\alpha = 0.1$ and $\beta = 0.025$. Fig. 4 (left) shows the converged height function of the reconstructed surface without imposing a weak statistical constraint, i.e., $\gamma = 0$. Fig. 5 shows the corresponding observed PDF using normalized height $\xi = (Z - \mu_Z)/\sigma_Z$ (zero mean and unit variance). Note the deviations from Gaussianity with large kurtosis. Further, the associated omni-directional spectrum $S(k)$ is also shown in Fig. 5 (dashed line). In a polar-reference frame, $S(k)$ is computed from the two-dimensional power spectrum Ψ of the wave surface Z as $S(k) = \int_0^{2\pi} \Psi(k, \theta) k d\theta$, where k is the wavenumber and θ is the angle. According to the wave turbulence theory of Zakharov [19], the spectrum tail initially decays as $k^{-2.5}$ as a result of an energy cascade from large to small scales up to ~ 10 rad/m and then switch to a k^{-3} equilibrium range [11].

Next, 200 V-cycles of multigrid are carried out using the energy augmented by (16), γE_{cdf} with $\gamma/A = 10^{-2}$, to drive the surface toward the target




Fig. 5. Left: Observed PDF of the reconstructed wave surface Z with (solid line) and without (dash line) statistical regularization. The Gaussian distribution is plotted for comparison (dotted line). Right: Observed omni-directional spectrum $S(k)$ of the reconstructed surface Z with (solid line) and without (dash line) statistical regularizer.

distribution: Gaussian for simplicity, although other distributions could have been used [15,16]. The converged height and radiance functions are shown in Fig. 4. Both, Z and f generate the modeled image in Fig. 3. The absolute error image with respect to the input image is also displayed. There are subtle differences between height functions with and without the statistical constraint. Both solutions correctly capture the (almost breaking) wave front moving toward the camera. Now, two non-linear terms (photometric fidelity and statistical constraint) compete to evolve the surface. The regions that change the most due to the statistical regularizer are those with smooth texture, corresponding to small photometric error. The statistical regularizer leaves the photometric error and omni-directional spectrum (Fig. 5, right) almost unchanged while significantly modifying the PDF of the height map. The new reconstructed surface is quasi-Gaussian as clearly shown in Fig. 5.

5 Conclusion

Variational stereo is more powerful, flexible, and rigorous, albeit computationally expensive, than earlier traditional, image-based stereo methods. Therefore, we follow this research path by developing a variational stereo method for the case of smooth surfaces representable in the form of a graph supporting a smooth radiance function. Moreover, we show how global properties of ocean waves, such as statistical distributions, can be incorporated in the variational stereo reconstruction framework via a weak constraint. We successfully apply this method in two experiments to reconstruct a small region of the surface of the ocean. The variational stereo method developed can be naturally extended in several ways to process sequences of stereo images to generate a coherent space-time reconstruction of ocean waves. In future research we plan to investigate new energy terms to incorporate more global and/or local properties of the dynamics of ocean waves such as the wave equation, etc.

References

1. Alvarez, L., Deriche, R., Sánchez, J., Weickert, J.: 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. Benetazzo, A.: Measurements of short water waves using stereo matched image sequences. *Coastal Engineering* 53(12), 1013–1032 (2006)
3. Briggs, W.L., Henson, V.E., McCormick, S.F.: *A Multigrid Tutorial*, 2nd edn. SIAM, Philadelphia (2000)
4. Faugeras, O.D., Keriven, R.: Variational principles, surface evolution, pdes, level set methods, and the stereo problem. *IEEE Trans. Image Proc.* 7(3), 336–344 (1998)
5. Gallego, G., Benetazzo, A., Yezzi, A., Fedele, F.: Wave statistics and spectra via a variational wave acquisition stereo system. In: *OMAE* (2008)
6. Hartley, R.I., Zisserman, A.: *Multiple View Geometry in Computer Vision*, 2nd edn. Cambridge University Press, Cambridge (2004)
7. Holland, K., Holman, R., Lippmann, T., Stanley, J., Plant, N.: Practical use of video imagery in nearshore oceanographic field studies. *IEEE Journal of Oceanic Engineering* 22(1), 81–92 (1997)
8. Jin, H.: Variational methods for shape reconstruction in computer vision. Ph.D. thesis, Washington University, St. Louis, MO, USA (2003), director: Soatto, S
9. Longuet-Higgins, M.S.: The effect of non-linearities on statistical distributions in the theory of sea waves. *Journal of Fluid Mechanics* 17, 459–480 (1963)
10. Ma, Y., Soatto, S., Kosecka, J., Sastry, S.: *An Invitation to 3D Vision: From Images to Geometric Models*. Springer, Heidelberg (2003)
11. Phillips, O.: The equilibrium range in the spectrum of wind-generated waves. *Journal of Fluid Mechanics* 4(4), 426–434 (1958)
12. Schumacher, A.: *Stereophotogrammetrische wellenaufnahmen. wiss ergeb. dtsh. atlant. exped. forschung vermessung. meteor 1925-1927. Ozeanographische Sonderuntersuchungen, Erste Lieferung* (1939)
13. Seitz, S.M., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *CVPR*, vol. 1, pp. 519–528 (2006)
14. Shemdin, O., Tran, H.: Measuring short surface waves with stereography. *Photogrammetric Engineering and Remote Sensing* 58(311-316) (1992)
15. Socquet-Juglard, H., Dysthe, K., Trulsen, K., Krogstad, H.E., Liu, J.: Probability distributions of surface gravity waves during spectral changes. *Journal of Fluid Mechanics* 542, 195–216 (2005)
16. Tayfun, A., Fedele, F.: Wave height distributions and nonlinear effects. *Ocean Engineering* 34(11-12), 1631–1649 (2007)
17. Wanek, J.M., Wu, C.H.: Automated trinocular stereo imaging system for three-dimensional surface wave measurements. *Ocean Engineering* 33(5-6), 723–747 (2006)
18. Yezzi, A., Soatto, S.: Stereoscopic segmentation. *International Journal of Computer Vision* 53(1), 31–43 (2003)
19. Zakharov, V.E.: Statistical theory of gravity and capillary waves on the surface of a finite-depth fluid. *Eur. J. Mech. B - Fluids* 18(3), 327–344 (1999)