

1 WASS: an Open-Source Pipeline for 3D Stereo
2 Reconstruction of Ocean Waves

3 Filippo Bergamasco^{a,*}, Andrea Torsello^a, Mauro Sclavo^b, Francesco
4 Barbariol^b, Alvisè Benetazzo^b

5 ^a*DAIS, Università Ca'Foscari, Venice, Italy*

6 ^b*Institute of Marine Sciences, Italian National Research Council (ISMAR-CNR), Venice*
7 30122, Italy

8 **Abstract**

9 Stereo 3D reconstruction of ocean waves is gaining more and more popular-
10 ity in the oceanographic community and industry. Indeed, recent advances of
11 both computer vision algorithms and computer processing power now allow the
12 study of the spatio-temporal wave field with unprecedented accuracy, especially
13 at small scales. Even if simple in theory, multiple details are difficult to be
14 mastered for a practitioner, so that the implementation of a sea-waves 3D re-
15 construction pipeline is in general considered a complex task. For instance,
16 camera calibration, reliable stereo feature matching and mean sea-plane esti-
17 mation are all factors for which a well designed implementation can make the
18 difference to obtain valuable results. For this reason, we believe that the open
19 availability of a well tested software package that automates the reconstruction
20 process from stereo images to a 3D point cloud would be a valuable addition for
21 future researches in this area.

22 We present WASS (<http://www.dais.unive.it/wass>), an Open-Source
23 stereo processing pipeline for sea waves 3D reconstruction. Our tool completely
24 automates all the steps required to estimate dense point clouds from stereo im-
25 ages. Namely, it computes the extrinsic parameters of the stereo rig so that
26 no delicate calibration has to be performed on the field. It implements a fast
27 3D dense stereo reconstruction procedure based on the consolidated OpenCV
28 library and, lastly, it includes set of filtering techniques both on the disparity
29 map and the produced point cloud to remove the vast majority of erroneous
30 points that can naturally arise while analyzing the optically complex nature of
31 the water surface.

32 In this paper, we describe the architecture of WASS and the internal algo-
33 rithms involved. The pipeline workflow is shown step-by-step and demonstrated
34 on real datasets acquired at sea.

35 *Keywords:* Sea Surface Stereo Reconstruction, Stereo Reconstruction, 3D
36 wave field, Camera Calibration, Open-source software

*Corresponding author

Email address: filippo.bergamasco@unive.it (Filippo Bergamasco)

37 1. Introduction

38 The 3D geometry is a crucial feature of ocean waves and it is important
39 in many problems, covering both the scientific research, off-shore and coastal
40 engineering applications. Indeed, natural ocean waves evolve in time spanning a
41 defined sea surface region where the wave field is assumed statistically homoge-
42 nous. However, at short-time scale (comparable to characteristic wave periods)
43 the spatial wave field shows distinctive features (e.g. the short-crestedness) that
44 cannot be retrieved by traditional one-point observational systems (for example
45 wave probes, buoys, etc.) Instead, a spatio-temporal observational system is
46 required.

47 Spatio-temporal wave measurements are typically somewhat hard to obtain
48 in the open ocean, but significant progress has been made in recent years. Dif-
49 ferent techniques were used for this purpose, ranging from polarimetric imagery
50 (Zappa et al., 2008), X-band wave radar (Borge et al., 2013; Young et al., 1985)
51 to satellite-borne synthetic aperture radar (Hasselmann et al., 1985). A direct
52 observation of the 3D wave field is also provided by the stereo imaging tech-
53 nique (Jahne, 1993). To give a general taxonomy of the spatial scales captured
54 by each aforementioned technology, we can state that polarimetry captures short
55 (wavelength range: $0.001 - 1m$) gravity-capillary and gravity waves, stereo imag-
56 ing captures short to mid size wavelengths (wavelength range: $0.2 - 50m$), and
57 X-band radar captures mid-length to long waves (wavelength range: $10 - 300m$).

58 The pipeline described here concerns specifically the processing of stereo
59 images for 3D reconstruction of ocean waves. Indeed, stereo wave imaging has
60 been proving to be an accurate technique for direct observations of 3D wave
61 fields at scales important for many scientific and applicative problems, with the
62 only limitation being the image processing that has to be performed a efficient
63 and effective way. This is nowadays achieved given the advances in the field
64 of computer vision. History of stereo wave imaging traces back of about one
65 century, when (Schumacher, 1939) installed a stereo-photographic system on a
66 ocean going ship. Afterwards, a major attempt in stereo-photography for ocean
67 waves was the Stereo Wave Observation Project in 1954 (Cote et al., 1960).
68 In the eighties, advances on stereo processing include the studies of (Shemdin
69 et al., 1988; Shemdin and Tran, 1992; Banner et al., 1989). However, all early
70 analysis processed manually the stereo pairs, resulting in a limited amount of
71 image pairs to be processed. Only more recently has automatic computer-based
72 processing of stereo images been used to investigate the 3D wave geometry.

73 Scholars used stereo 3D reconstruction of ocean waves to explore a
74 variety of relevant scientific questions. For example, in the field of extreme
75 wave analysis, (Benetazzo et al., 2015) investigated the presence of rogue waves
76 within space-time sample of sea surface elevations, and (Fedele et al., 2013)
77 compared time and space-time statistics of high waves. (Banner et al., 1989)
78 examined the shape of the wavenumber wave spectrum, while (Leckler et al.,
79 2015) studied the shape of the wavenumber-frequency spectrum computed via
80 direct Fourier transform of the stereo data (Benetazzo et al., 2012; Gallego et al.,
81 2011). (Banner et al., 2014) verified, using laboratory data and field stereo ob-

82 servations, the slowdown of unsteady wave groups under wave crest maxima.
83 (Sutherland and Melville, 2013) used stereo-processed infrared imagery to es-
84 timate the dissipation rate of breaking waves. (Campbell et al., 2014) studied
85 the interaction of the wave field with different ice types. (Mironov et al., 2012)
86 proposed a methodology to extract short-scale statistical characteristics of the
87 sea surface topography. All studies so far used stereo system mounted over fixed
88 (non-moving) structures. This condition eases the stereo processing but highly
89 limits the conditions under which the 3D characteristic of the wave field are in-
90 vestigated. Hence, recently, studies have been conducted to verify limitations of
91 stereo systems mounted over moving structures (typically a ship), the difference
92 being the requirement of the motion parameters to correctly map the stereo
93 data into a reference system coherent with the gravity field. The ship motion
94 compensations for stereo data are obtained using external instruments (Brandt
95 et al., 2010; Benetazzo et al., 2014), the horizon line (Bergamasco et al., 2016),
96 or the 3D sea surface field itself (Benetazzo et al., 2016).

97 Stereo matching has been one of the most active area in computer vision
98 in the recent past. A deluge of different methods have been proposed over the
99 years building on the same common sequence of operations, namely: matching
100 cost computation, cost aggregation, disparity computation and disparity refine-
101 ment. Given the intrinsic ambiguity of the problem, spatial smoothness of the
102 disparity map is often considered a good prior to obtain a physically reasonable
103 surface reconstruction. In particular, many different strategies were studied for
104 the first two steps depending whether the disparity computation is determined
105 by matching intensity values over a finite window (*local methods*) or by consid-
106 ering the whole image to perform a global matching cost minimization (*global*
107 *methods*). The first have the advantage to be faster in general but more prone
108 to outliers for the local nature of the matching cost computation. Conversely,
109 *global methods* usually result in a smoother and more accurate reconstruction
110 but require far more computational resources.

111 Almost all the seminal oceanographic studies (Benetazzo, 2006; Wanek and
112 Wu, 2006; Brandt et al., 2010; Bechle and Wu, 2011; de Vries et al., 2011;
113 Kosnik and Dulov, 2011; Benetazzo et al., 2012; Leckler et al., 2015) adopted
114 *local methods* to compute the disparity map and hence suffer the delicate trade-
115 off between the disparity window size (which influence the match localization
116 accuracy) and the required surface smoothness that is particularly important
117 while reconstructing a continuous injective surface like the sea. The state-of-
118 the-art global method applied to waves reconstruction was proposed by (Gallego
119 et al., 2008, 2011) that use a variational approach to directly estimate a con-
120 tinuous smooth surface by optimizing the photometric consistency between the
121 two views. Additionally, the formulation allows also to impose physical con-
122 straints in the recovered points elevations. Unfortunately, such approach is so
123 computationally intensive that is unlikely to be used in practice especially for
124 long sequences.

125 The open source pipeline proposed here is a refined implementation of the
126 WASS system which has been extensively used and tested in the last five years
127 (Benetazzo et al., 2012, 2014, 2015; Bergamasco et al., 2016). For stereo match-

128 ing, we make use of a *semi-global* approach that combines the good performances
129 of local methods (in terms of CPU time) with the accuracy and robustness of
130 global methods. Additionally, the proposed WASS pipeline incorporates all the
131 features needed to perform the extrinsic calibration on the field which has been
132 demonstrated to increase the calibration accuracy easing its installation at sea.

133 This manuscript is organized as follows. In Section 2 the proposed WASS
134 software is described in detail, following the general workflow that is performed
135 during the processing of the stereo frames. In Section 3 we give a theoretical
136 discussion of the expected reconstruction errors together with the internal and
137 external factors that influence the overall reconstruction accuracy. Finally, in
138 Section 4 we demonstrate our pipeline with data we acquired from an off-shore
139 oceanographic platform.

140 2. The WASS pipeline

141 Our proposed sea-waves reconstruction tool completely automates the cre-
142 ation of a sequence of dense 3D point clouds from stereo images providing three
143 important functionalities.

144 First, WASS can automatically recover the extrinsic parameters of the stereo
145 rig (up to scale) so that no delicate calibration has to be performed on the
146 field. Our past experience (Benetazzo et al., 2016) has demonstrated that if the
147 intrinsic calibration is easy to obtain with all the commodities of a dedicated
148 laboratory, extrinsic calibration is not. This is due to the fact that, for common
149 setups covering areas ranging from 20×20 to 70×70 meters, a distance of at
150 least a couple of meters is required between the two cameras. Consequently,
151 it’s rather difficult to calibrate the extrinsics on the field, forcing the usage
152 of a big calibration target in a possibly dangerous environment (ie. out of a
153 ship deck, on an offshore platform, etc). Moreover, the auto-calibration offers
154 the advantage of letting the rig geometry reconfigurable on-the-fly before each
155 acquisition to accommodate different requirements. For instance, it may be
156 reasonable to take the device closer to the sea when the waves are small, so
157 a small but highly resolved sea surface region can be acquired. On the other
158 hand, large waves demand a broader sea surface region, requiring the device to
159 be repositioned farther from the surface. Finally, the “calibrate once and for
160 all” strategy is not reliable since vibrations of the support and environmental
161 factors, as wind, can modify the relative angle between cameras and jeopardize
162 the reconstruction accuracy.

163 The second feature provided by WASS is a fast 3D dense stereo reconstruc-
164 tion procedure so that an accurate 3D point cloud can be computed from each
165 stereo pair. We rely on the well consolidated OpenCV library (Bradski and
166 Kaehler, 2008) both for the image stereo rectification and disparity map recov-
167 ery as described in the following section 2.4.

168 Third, a set of 2D and 3D filtering techniques both on the disparity map and
169 the produced point cloud are implemented to remove almost all the erroneous
170 points that can naturally arise while analyzing the optically complex nature of

171 the water surface (examples are sun-glares, large white-capped areas, fog and
172 water aerosol, etc).

173 As with the vast majority of dense stereo reconstruction methods, we make
174 the implicit assumption of the Lambertian response of the water surface. More
175 specifically, to search for corresponding matching features in the two stereo
176 images we expect that the optical appearance on the two cameras is mostly
177 preserved in terms of pixels intensities. In general, this is true only if the sur-
178 face intensity response depends only on the reciprocal position between the
179 illuminant and the surface normal, as in the Lambertian model. Any other
180 phenomena like specularities and refractions are function of the mixed contri-
181 bution between the illuminant, the surface normal and the viewpoint. Since the
182 appearance of sea water surface is essentially described by the Fresnel law, the
183 Lambertian assumption is quite limiting for many aspects. Nevertheless, under
184 good illuminant conditions (no direct sun, cameras almost parallel and angled
185 with respect to the sea plane, etc) classical dense stereo 3D reconstruction still
186 provides excellent results compared to other methods like mechanical buoys or
187 radars (Benetazzo et al., 2012).

188 *2.1. Software architecture*

189 In the present version, WASS is composed by 4 highly optimized executables
190 written in C++ whose parallel execution in a pipeline is orchestrated by a script
191 (*WASSjs*) running on the node.js runtime environment (Tilkov and Vinoski,
192 2010). The four executables are named respectively: *wass_prepare*, *wass_match*,
193 *wass_autocalibrate* and *wass_stereo*.

194 *wass_prepare* is responsible for initializing of each couple of stereo images
195 placing them into a proper “working directory” together with the necessary
196 calibration and configuration data. Images are also undistorted according to
197 the provided intrinsic calibration data.

198 *wass_match* and *wass_autocalibrate* are used in case that the extrinsic auto-
199 calibration is needed. The first matches corresponding features of a stereo image
200 pair with a robust state-of-the-art game-theoretic framework (Albarelli et al.,
201 2012). After the feature matching, the essential matrix is recovered and the
202 extrinsic parameters between the two cameras are factorized (up to translation
203 scale). Matched features are also filtered according to epipolar consistency.
204 *wass_autocalibrate* collects all the matches between multiple stereo frames and
205 uses Sparse Bundle Adjustment (SBA) to simultaneously optimize the extrinsic
206 camera parameters and the 3D points triangulated from all the matches. Final
207 extrinsic parameters are saved to all the workspaces.

208 *wass_stereo* performs stereo rectification and dense stereo 3D reconstruction
209 on a given stereo frame. After the reconstruction, multiple filters are applied to
210 the point cloud to remove erroneous points. Finally, the mean plane is robustly
211 fitted to the reconstructed data so that it can be further aligned to a common
212 sea reference frame.

213 Each WASS component runs as a single-threaded process performing a spe-
214 cific task on a working directory. To exploit the intrinsic parallelism of the stereo

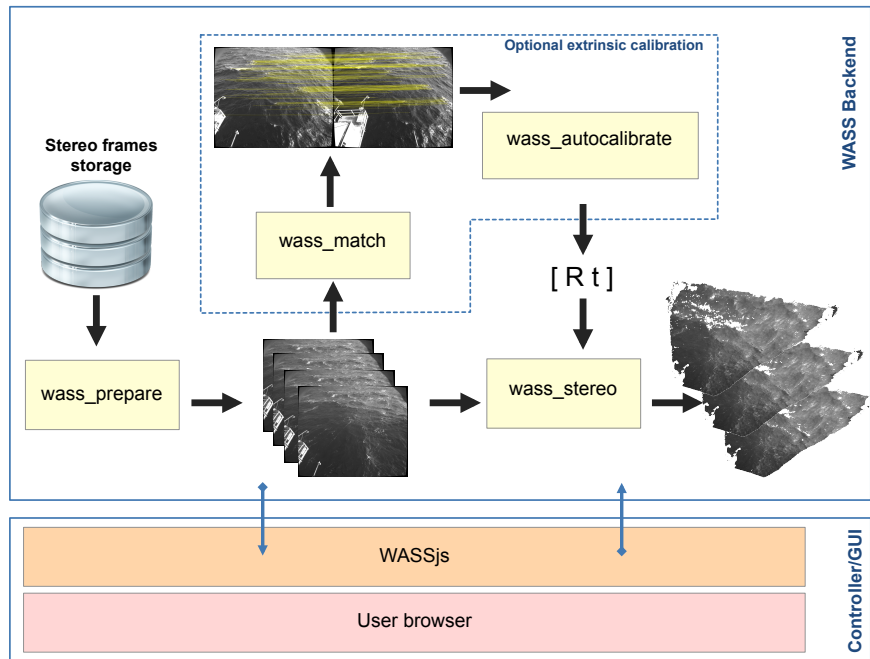


Figure 1: Components and workflow of the WASS reconstruction pipeline. User interacts with the pipeline via web browser communicating with the WASSjs component. *wass_prepare* and *wass_stereo* are responsible of the stereo frames undistortion and reconstruction respectively. The optional *wass_match* and *wass_autocalibrate* can automate the extrinsic parameters recovery of the stereo rig.

215 reconstruction task, the *WASSjs* controller manages multiple parallel instances
 216 of each component. As a consequence, WASS can scale well in distributed mem-
 217 ory machines (like virtualized environments) with the only requirement to have
 218 a consistent shared view of a common filesystem.

219 2.2. Image preparation

220 The first step performed by the pipeline is to organize the stereo images into
 221 an ordered sequence of working directories so that all the subsequent steps can
 222 be easily managed in parallel. Apart from the images, intrinsic calibration for
 223 each camera has to be provided by the user. The pipeline assumes a pinhole
 224 camera model for each camera with (optionally) non square pixels and zero
 225 skewness. We also support a polynomial (5-coefficients) radial distortion (Fryer
 226 and Brown, 1986) that may be provided or set to zero.

227 The user is free to use whichever method he prefers for intrinsic calibration.
 228 For our work, we used both the Bouguet’s Matlab calibration toolbox, imple-
 229 menting the method of (Heikkila and Silven, 1997), or the method proposed in
 230 (Albarelli et al., 2010) that is less sensible to the manufacturing accuracy of the

231 calibration target. Once the calibration is performed, data must be saved in an
 232 *xml* format compatible with the OpenCV serialization standard (see the official
 233 documentation for more informations).

234 2.3. Automatic stereo-rig calibration

235 The auto-calibration of the stereo rig is a complex task that can be option-
 236 ally executed if no extrinsic calibration is provided by the user. It's based on
 237 the assumption that the relative pose of the two cameras remains unchanged
 238 throughout the sequence so that it can be recovered up to a positive scale by
 239 relating multiple corresponding features that have to be reliably extracted be-
 240 tween the two cameras.

241 Supposing to have scene composed by a set of n 3D points projecting into
 242 the first and second camera as points $p_1 \dots p_n$ and $p'_1 \dots p'_n$ respectively, the
 243 following well-known epipolar constraint holds:

$$p_i = \mathbf{K}_1^{-T} \mathbf{E} \mathbf{K}_2^{-1} p'_i, \quad i = 1 \dots n \quad (1)$$

244 where \mathbf{E} is the essential matrix and \mathbf{K}_1 and \mathbf{K}_2 are the intrinsic camera
 245 matrices of the two cameras. This implies that, if a set of point-point corre-
 246 spondences can be matched from the two stereo images, the essential matrix can
 247 be recovered by solving equation (1). Then, the essential matrix can be further
 248 decomposed via Singular Value Decomposition (SVD) to recover the rotation
 249 \mathbf{R} and translation direction \vec{t} of the two views, as described in (Hartley and
 250 Zisserman, 2004).

251 WASS auto-calibration procedure is composed by the following steps:

- 252 1. A subset of all the available stereo frames are selected randomly as cali-
 253 bration images
- 254 2. A sparse set of *Speeded Up Robust Feature* (SURF) (Bay et al., 2008) is
 255 extracted from each image (Fig. 2 top row)
- 256 3. For each stereo pair, features are robustly matched with the approach of
 257 (Albarelli et al., 2012)
- 258 4. Erroneous matches are removed with a RANSAC-based epipolar filtering.
 259 (Fig. 2.c) As an outcome of this process, an initial estimate of the essential
 260 matrix is given for each stereo pair. The essential matrix associated to
 261 the pair with the highest number of feature matches is used as the initial
 262 coarse estimate for the next step.
- 263 5. All the matches between all the stereo frames of the calibration subset are
 264 used to simultaneously recover a 3D sparse point cloud of the matched
 265 features together with the camera poses with respect to a common un-
 266 known reference frame so that the global reprojection error is minimized.
 267 This effectively averages all the independent essential matrices estimations
 268 between each stereo pair used in the matching step. We perform this op-
 269 timization as a Sparse Bundle Adjustment implemented as described in
 270 (Lourakis and Argyros, 2009)

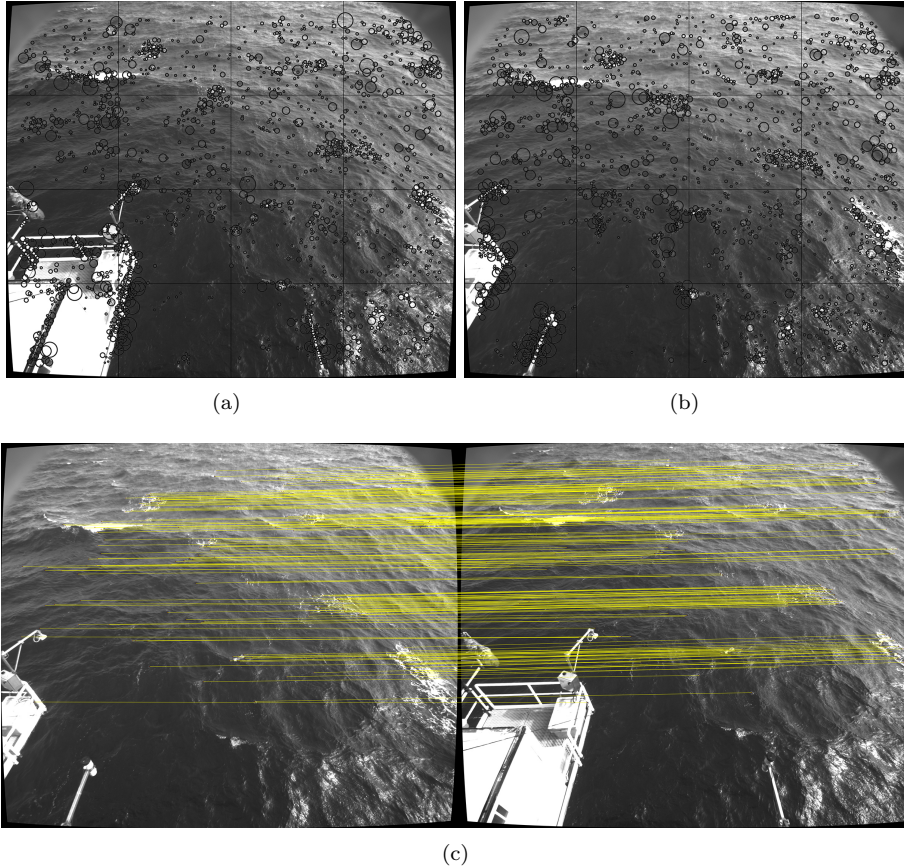


Figure 2: (a) Left and (b) right image captured by the stereo rig with the extracted SURF features superimposed. (c) Matched features (after epipolar filtering) computed by *wass_match* (best viewed in colors).

271 Steps 1 to 4 are performed by the *wass_match* executable that is run in
 272 parallel among all the randomly selected calibration images. Step 5 is performed
 273 by *wass_autocalibrate* as soon as the results of the previous matching process is
 274 available for all the calibration images. When *wass_autocalibrate* completes its
 275 job, the recovered rotation \mathbf{R} and normalized translation \vec{t} are saved on each
 276 workspace so that they can be used by the subsequent 3D reconstruction process.
 277 If the extrinsic parameters are provided by the user (either because they have
 278 been previously calibrated by the pipeline or with a calibration target) they can
 279 be automatically used to bypass the auto-calibration step described before.

280 2.4. 3D reconstruction and sea-plane estimation


281 The 3D reconstruction process that computes a dense point cloud from the
 282 stereo frames is implemented in the *wass_stereo* executable. Again, we exploit

283 the temporal independence of the frames to execute the operation in parallel on
284 all the available computation units so that the total computing time to recon-
285 struct a whole sequence can be heavily reduced even on a common consumer
286 workstation.


287 When *wass_stereo* is started, the two stereo frames are rectified using the
288 extrinsic and intrinsic calibration data provided. At the end of this operation,
289 all the corresponding epipolar lines will be perfectly aligned with the image rows
290 (Fig. 3.a) so that the search of corresponding stereo pixels on the two images
291 can be reduced along the same row during the dense stereo operation. The
292 rectified stereo pair is also saved on each workspace for inspection purposes. In
293 fact, the user can manually observe if the same corresponding visual features
294 on left and right images are aligned on the same image row. If this happens,
295 it gives a good feedback that both the intrinsic and extrinsic auto-calibration
296 were performed with a reasonable accuracy.

297 After the rectification, a dense stereo algorithm is executed to obtain the
298 disparity map between the two frames (Fig. 3.b). In our software we used
299 the implementation provided by the OpenCV library (Bradski and Kaehler,
300 2008) that represent a de-facto standard reference implementation for many
301 specialized computer vision algorithms. Specifically, we have chosen to use
302 the semi-global dense stereo method described in (Hirschmüller, 2008) as it
303 provides an excellent trade-off between the execution speed of a local method
304 (ie. methods for which the cost aggregation step is performed locally along the
305 disparity ranges) and the accuracy in low-textured areas of a global method (in
306 which the disparity is recovered by means of an optimization involving all the
307 image pixels). Since the obtained disparity map tends to exhibit errors at the
308 boundaries of the matching regions, we perform some preliminary morphological
309 filters directly on the disparity map to partially limit the outliers (Fig. 3.c).
310 Specifically, we dilate the disparity map by n steps (usually 1 or 2 is sufficient)
311 before eroding the results $m > n$ times. The initial dilation closes many of
312 the small holes that may happen inside the observed sea surface due to sun
313 glares or other artifacts. The subsequent erosion removes the disparity values
314 at the boundary of the disparity map where most of the errors are located (since
315 both the dense stereo cost computation and aggregation steps cannot be reliably
316 performed).

317 With the cleaned disparity map, the initial unfiltered 3D point cloud is
318 generated by triangulating all the corresponding pixels. In this step we filter all
319 the points with a depth smaller than 1.0 (since \vec{t} has unitary norm this essentially
320 filters all the points nearest than the cameras baseline) and higher than z_{max}
321 (ie. z_{max} times the camera baseline). Additionally, we filter all the points
322 for which the angle between the two intersecting camera rays is lower than 20
323 degrees. Note that, since the corresponding points to be triangulated are always
324 lying on the respective epipolar lines (the dense stereo algorithm operates along
325 the rows of the rectified images) there is no need to use a computationally-
326 expensive triangulation approach that minimizes the reprojection instead of the
327 algebraic error. This greatly speeds up the 3D reconstruction that usually takes
328 few seconds to process our 5 Mpixel images on a common Intel i7 CPU running



(a)



(b)

(c)

Figure 3: (a) Rectified stereo-pair used as input for the dense stereo algorithm. Note that the corresponding visual features of the images are aligned on the same image row between the two frames. (b) Unfiltered disparity map obtained by the dense stereo algorithm. (c) Disparity map after the dilation and erosion filters. Note how the small holes at the bottom of the disparity map have been closed and the boundary of the disparity area is less noisy

329 at 2.5Ghz.

330 The obtained point cloud is in general already quite clean but a limited
 331 amount of spurious points can still be present. In fact, the dense stereo algo-
 332 rithm is designed to be general purpose so it does not make any assumptions on
 333 the reconstructed scene. Conversely, we can take advantage of the fact that the
 334 sea surface is spatially continuous and smooth everywhere. To greatly reduce
 335 the number of point outliers we use an additional filtering step that exploits
 336 the surface smoothness along the camera optical z-axis. We start by building a
 337 graph with the vertices being the reconstructed 3D points and edges connecting
 338 each vertex with its 4-neighbors considering the adjacency relation of points in-
 339 duced by the image lattice topology. In other words, we consider two 3D points
 340 connected if they have been triangulated from two adjacent pixels in the right
 341 image. We compute a weight associated to each edge as the absolute differ-

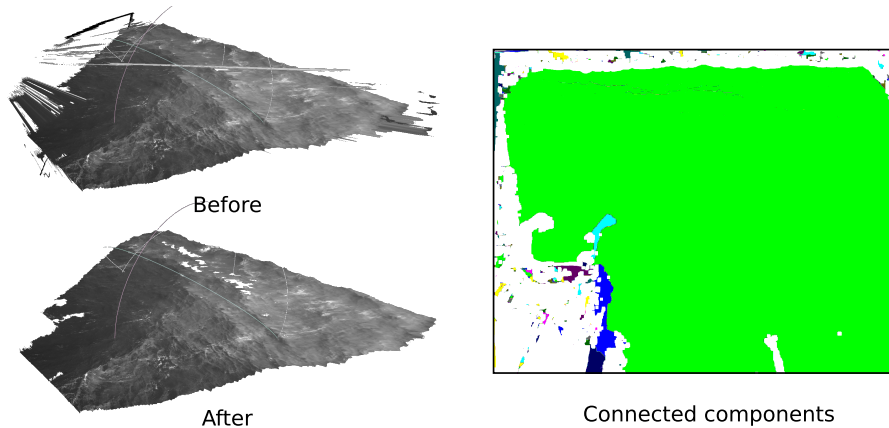


Figure 4: Effect of the graph-based filtering on the reconstructed point cloud. Left: Point cloud before and after the filtering. Right: Connected components extracted from the constructed graph. For the sake of visualization, a random color is assigned to each connected component except for the green which marks the biggest one.

342 ence along the camera z -axis of their respective points coordinates. Then, each
 343 edge whose weight is greater than the 98th percentile of the weight distribution
 344 is pruned by the graph. The idea is to disconnect all the vertices exhibiting
 345 abrupt changes along the z -axis with respect to the neighbors. Finally, we filter
 346 out all the vertices and edges not belonging to the biggest connected component
 347 of the graph. Since the erroneous points are likely to exhibit abrupt changes
 348 with respect to the neighbours, they are efficiently removed from the principal
 349 observed surface. In Figure 4 we show an example of all the connected components
 350 extracted from the graph and the filtering result on the reconstructed
 351 point cloud. We can clearly observe that most of the reconstruction artifacts
 352 that happen at the boundary of the sea surface are effectively removed since
 353 they exhibit abrupt changes along the camera z -axis.

354 The final operation performed by *wass_stereo* is the estimation of the mean
 355 sea plane from the reconstructed point cloud. This allows a subsequent rotation
 356 of the 3D data so that the new Z -axis is oriented upward with respect to the
 357 earth surface. To limit the effect of possible point outliers that may have passed
 358 the previous filters, we perform an initial robust RANSAC (Fischler and Bolles,
 359 1981) estimation of the mean plane before the least-squares fitting. Specifically,
 360 for a selected number of iterations we extract 3 random points from the point
 361 cloud and compute a candidate plane parameters. Then, we count the number
 362 of reconstructed points lying closer than a fixed threshold from the candidate
 363 plane. At the end of the RANSAC iterations, we keep the plane parameters
 364 corresponding to the candidate plane having the higher number of close points
 365 (inliers). This initial plane estimation is then refined with a least-squares fitting
 366 of all the inlier points by weighting each sample with a value proportional to the
 367 distance of the point from the camera. As previously assessed in (Bergamasco
 368 et al., 2016), this approach can improve the mean sea-plane recovery for wavy




Figure 5: The implemented point cloud filtering approach is able to discard image portions clearly not belonging to the reconstructed sea surface area. Left: Example of a clutter object (a seagull in this case) in the field-of-view of the two cameras. Note how the reconstructed points corresponding to the seagull have been properly removed. The point removal is evident as it causes a seagull-shaped hole in the surface due to the shadowing produced by the seagull itself with the other camera field-of-view. Right: Magnification of the area highlighted in the red rectangle on the left. (Best viewed in colors)

369 seas especially if no other external aids are available (like horizon-line position
 370 or IMU data).

371 3. Expected Errors

372 Five principal uncertainties can be expected in any stereo method: the un-
 373 certainty in the internal parameter calibration (internal calibration error), the
 374 uncertainty in the external parameter calibration (external calibration error),
 375 the uncertainty in the determination of the corresponding pixels (matching er-
 376 ror), the uncertainty in the recovery of 3D coordinates (quantization error),
 377 the uncertainty in the determination of the transformation between the cam-
 378 era reference system and the water reference system with the Z-axis vertical
 379 and pointing upward (camera orientation error). Following standard calibra-
 380 tion procedures (Bouguet, 2004) we have verified that uncertainty in the in-
 381 ternal calibration produces is of fractional of pixels, and we can consider them
 382 negligible for all applications. Yet this is not always the case for the external
 383 calibration. In previous WASS deployments (Benetazzo, 2006; Benetazzo et al.,
 384 2012), external calibration parameters were estimated by exposing an ad-hoc
 385 calibration target to both cameras, and by relating the known 3-D geometry of
 386 the target with its re-projection onto the image planes. However, even if this is
 387 the standard de-facto way to calibrate a stereo rig in laboratory conditions, this
 388 approach manifests several drawbacks when applied to stereo systems with large
 389 baseline, typical of field applications. In the latter the pipeline incorporates the
 390 auto-calibration procedure described in (Benetazzo et al., 2016), which has been
 391 proven to be as accurate as standard procedures. To determine the orientation and
 392 displacement between the camera reference system and the still sea water

393 surface, we have adopted the strategy proposed by (Benetazzo, 2006) that has
394 been proved to be accurate on field (Benetazzo et al., 2012; Gallego et al., 2008)
395 and synthetic (Benetazzo et al., 2016) data.

396 Concerning the matching and quantization errors, when designing a stereo-
397 camera system one must compromise to meet mainly the conflicting require-
398 ments of accurate 3D estimation and accurate image feature matching. In fact,
399 as observed by (Mironov et al., 2012) the major difference between stereo wave
400 imaging and the classical problem of stereo reconstruction is in the fact that the
401 water surface reflectance is not Lambertian. However, (Benetazzo, 2006) shown
402 that the matching error is small for highly-textured water surfaces, and (Jahne,
403 1993) indicates that the matching error is small when the wave slope is much
404 larger than the inclination of the stereo cameras optical axis. The geometry
405 of the stereo rig must be tuned so that the interplay between the light vector,
406 the point of view and the surface normal is similar for the two cameras. Thus,
407 we assume that the disparity of each corresponding pixel is dominated by the
408 spatial position of the 3D surface point and not by the rather complex water
409 surface bi-directional reflectance distribution function (BRDF). To provide a
410 good 3D reconstruction, care has to be taken to choose different parameters,
411 the principal being the camera cell size and pixel numbers, the focal length, the
412 baseline, the camera reciprocal orientation, and the distance from the stereo-
413 camera system to the scene of interest. Under general conditions, in order to
414 keep small the range error due to quantization (hence to provide high accuracy
415 in the 3D reconstruction) the baseline-to-distance ratio must be large; however,
416 accurate feature points matching require that this ratio be small (Benetazzo,
417 2006; Jahne, 1993; Rodriguez and Aggarwal, 1990).

418 Based on our experience on stereo wave imaging at the sea, the photometric
419 consistency between the two views can be rather good if the following general
420 principles are taken with care. First, the stereo camera optical axes should be
421 placed (almost) parallel in order to reduce the angle between them. Second, the
422 baseline-to-distance should be kept small, around 0.10, which is much smaller
423 than historical set-up used for field experiment like in (Jahne, 1993). Third, the
424 camera should be oriented so that no severe sun glitters are visible in the field of
425 view. Finally, the quantization error is alleviated by using subpixel (fractional
426 pixel) correspondence as performed by our pipeline. Sub-pixeling reduces errors
427 depending on the weight function adopted to estimate the cross-correlation map
428 (Nobach and Honkanen, 2005), and has a large impact on the accuracy of the
429 smaller wavelengths (see for instance Fig.7 of (Benetazzo et al., 2012)). To
430 give an idea of the reconstruction error expected with the proposed pipeline,
431 in Figure 6 we show a map of the quantization error for the same WASS setup
432 described in (Benetazzo et al., 2015).

433 4. Example

434 We use a field data example to demonstrate the capability of the open-source
435 pipeline for WASS. The pipeline has already been used and its performance as-
436 sessed in different studies using stereo cameras mounted on oceanographic tow-




Figure 6: Map of the quantization error for sea surface elevations (Z-axis) within the horizontal field of view of the stereo cameras. WASS setup as in (Benetazzo et al., 2015)

437 ers (Benetazzo et al., 2015) and moving vessels (Benetazzo et al., 2016; Berga-
 438 masco et al., 2016; Falcieri et al., 2016). Here, the pipeline is tested on images
 439 acquired by a stereo wave imaging system ad-hoc installed on the oceanographic
 440 research platform Acqua Alta (North Adriatic Sea, Italy; see Figure 2). Image
 441 pairs were acquired with a 5 mega-pixel digital cameras (with 2456 columns
 442 by 2048 rows array of $3.45\mu\text{m}$ square active elements) and mounting 5 mm
 443 distortionless lenses, placed 2.5 m apart and on the roof of the platform.

444 The 3D wave fields were observed during a mature sea state generated by
 445 a north-easterly storm with average wind speed of 11 m/s, resulting in a local
 446 significant wave height of about 1.3 m. The WASS sequence comprises 26971
 447 stereo-image pairs grabbed at 15 Hz. For internal parameter calibration of both
 448 cameras, we used the Bouguet’s Matlab calibration toolbox (Bouguet, 2004).

449 Using the WASS pipeline, the automatic stereo rig calibration is performed
 450 selecting a subset containing 51 stereo frames (uniformly spaced along the image
 451 sequence). A total of about 20000 features are matched, and 89% of them re-
 452 tained after application of the RANSAC-based epipolar filtering. Then the rota-

453 tion matrix $R = \begin{pmatrix} 0.998568 & 0.0297004 & -0.0444994 \\ -0.0311139 & 0.999022 & -0.0314147 \\ 0.0435229 & 0.0327543 & 0.998515 \end{pmatrix}$ and the normalized




Figure 7: 3D point cloud result of the stereo triangulation superimposed to the left camera image of the stereo imaging system.

454 translation vector $\vec{t} = (0.999686 \ 0.0101738 \ 0.0229187)^T$ between the two
455 views are recovered. After the calibration and rectification of images, all stereo
456 pairs are processed for 3D stereo reconstruction. The dense stereo algorithm
457 execution find on average 3.3 million corresponding points, which are reduced
458 of about 5% with the optimization and cleaning of the disparity map, including
459 the graph-based filtering (see for example Figure 7). Along the whole sequence,
460 about 80×10^9 3D sea-surface elevation points resulted from the stereo process.
461 Then, as standard procedure, for each point cloud, elevations are transformed
462 to a common earth reference frame (Benetazzo, 2006; Gallego et al., 2008). To
463 this end, mean sea plane parameters are derived from each reconstructed point
464 cloud, and then parameters averaged over the 26971 outputs. Finally, each re-
465 constructed point cloud is scaled by a factor of 2.5 corresponding to the distance
466 between the two cameras that have been empirically measured in the field. Af-
467 ter this procedure, cameras resulted to be placed 12.5 m above the mean sea
468 level, with an elevation angle of 50° . Finally, to make stereo data useful for
469 oceanographic applications, a patchwise planar surface is constructed by means
470 of 2D Delaunay triangulation of each point cloud, resampled over a regular
471 grid at uniform resolution of 0.2 m along the horizontal X- and Y-axis (in the
472 world reference frame), spanning the horizontal region $x \in [-42.5m, 42.5m]$ and
473 $y \in [-70.0m, -5.0m]$ (Fig.8).




Figure 8: Representation of the 3D sea surface elevation fields measured by WASS. Top-Left: Slices on the XY-plane of the space-time volume. Top-right: Time series of sea elevations at different spatial positions. (bottom panel) Hovmoller diagram along the transect $x = 0$

474 5. Conclusions

475 We presented WASS, an open-source implementation of our sea waves 3D
 476 stereo reconstruction pipeline. WASS is designed to be fast and reliable when
 477 processing thousands of stereo frames and offers an intuitive interface that assist
 478 the user in all the common operations from dense stereo matching to stereo
 479 camera calibration. Tested and used extensively in our past works, we decided
 480 to release it to the whole oceanographic community in the hope that it can help
 481 improve the research in the study of sea waves.

482 We described the details of WASS inner working and demonstrated the per-
 483 formance and the expected errors of a typical WASS installation on a fixed
 484 oceanographic platform. We remand to the official website <http://www.dais.unive.it/wass>
 485 for additional technical details on installing and running our
 486 software on various platforms.

487 6. Acknowledgments

488 Authors gratefully acknowledge the funding from the Flagship Project RIT-
489 MARE - The Italian Research for the Sea - coordinated by the Italian National
490 Research Council and funded by the Italian Ministry of Education, University
491 and Research within the National Research Program 2011–15

492 References

- 493 Albarelli, A., Rodolà, E., and Torsello, A. (2010). Robust camera calibration
494 using inaccurate targets. *British Machine Vision Conference, BMVC 2010 -*
495 *Proceedings*.
- 496 Albarelli, A., Rodolà, E., and Torsello, A. (2012). Imposing semi-local geometric
497 constraints for accurate correspondences selection in structure from motion:
498 A game-theoretic perspective. *International Journal of Computer Vision*,
499 97:36–53.
- 500 Banner, M. L., Barthelemy, X., Fedele, F., Allis, M., Benetazzo, A., Dias,
501 F., and Peirson, W. L. (2014). Linking Reduced Breaking Crest Speeds to
502 Unsteady Nonlinear Water Wave Group Behavior. *Physical Review Letters*,
503 112:114502, 1–5.
- 504 Banner, M. L., Jones, I. S. F., and Trinder, J. C. (1989). Wavenumber spectra
505 of short gravity waves. *Journal of Fluid Mechanics*, 198:321–344.
- 506 Bay, H., Ess, A., Tuytelaars, T., and Van Gool, L. (2008). Speeded-up robust
507 features (surf). *Computer Vision and Image Understanding*, 110(3):346–359.
- 508 Bechle, A. J. and Wu, C. H. (2011). Virtual wave gauges based upon stereo
509 imaging for measuring surface wave characteristics. *Coastal Engineering*,
510 58(4):305–316.
- 511 Benetazzo, A. (2006). Measurements of short water waves using stereo matched
512 image sequences. *Coastal Engineering*, 53(12):1013–1032.
- 513 Benetazzo, A., Barbariol, F., Bergamasco, F., Torsello, A., Carniel, S., and
514 Sclavo, M. (2015). Observation of extreme sea waves in a space-time ensemble.
515 *Journal of Physical Oceanography*, 45(9):2261–2275.
- 516 Benetazzo, A., Barbariol, F., Bergamasco, F., Torsello, A., Carniel, S., and
517 Sclavo, M. (2016). Stereo wave imaging from moving vessels: Practical use
518 and applications. *Coastal Engineering*, 109:114–127.
- 519 Benetazzo, A., Bergamasco, F., Barbariol, F., Torsello, A., Carniel, S., and
520 Sclavo, M. (2014). Towards an operational stereo system for directional wave
521 measurements from moving platforms. *Proceedings of the International Con-*
522 *ference on Offshore Mechanics and Arctic Engineering - OMAE*, 8B: Ocean
523 Engineering.

- 524 Benetazzo, A., Fedele, F., Gallego, G., Shih, P.-C., and Yezzi, A. (2012). Off-
525 shore stereo measurements of gravity waves. *Coastal Engineering*, 64(0):127
526 – 138.
- 527 Bergamasco, F., Benetazzo, A., Barbariol, F., Carniel, S., and Sclavo, M. (2016).
528 Multi-view horizon-driven sea plane estimation for stereo wave imaging on
529 moving vessels. *Computers and Geosciences*, 95:105–117.
- 530 Borge, J. N., Reichert, K., and Hessner, K. (2013). Detection of spatio-temporal
531 wave grouping properties by using temporal sequences of x-band radar images
532 of the sea surface. *Ocean Modelling*, 61:21 – 37.
- 533 Bouguet, J. Y. (2004). Camera calibration toolbox for Matlab. Technical Report,
534 Calif. Inst. of Technol. Available at: [www.vision.caltech.edu/bouguetj/calib](http://www.vision.caltech.edu/bouguetj/calibdoc/index.html)
535 [doc/index.html](http://www.vision.caltech.edu/bouguetj/calibdoc/index.html). Technical report, Calif. Inst. of Technol.
- 536 Bradski, G. and Kaehler, A. (2008). *Learning OpenCV: Computer Vision with*
537 *the OpenCV Library*, volume 1. Inc, OReilly Media.
- 538 Brandt, A., Mann, J. L., Rennie, S. E., Herzog, A. P., and Criss, T. B. (2010).
539 Three-Dimensional Imaging of the High Sea-State Wave Field Encompassing
540 Ship Slamming Events. *Journal of Atmospheric and Oceanic Technology*,
541 27:737–752.
- 542 Campbell, A. J., Beckle, A., and Wu, C. H. (2014). Observations of surface waves
543 interacting with ice using stereo imaging. *Journal of Geophysical Research:*
544 *Oceans*, 119(-):3266–3284.
- 545 Cote, L. F., Davis, J. ., Marks, W., McGough, R. F., Mehr, E., Pierson, W.,
546 Ropek, J., Stephenson, G., and Vetter, R. (1960). The directional spectrum
547 of a wind generated sea as determined from data obtained by the Stereo Wave
548 Observation Project,. *Meteorological Papers*, pages 1–88.
- 549 de Vries, S., Hill, D., de Schipper, M., and Stive, M. (2011). Remote sensing of
550 surf zone waves using stereo imaging. *Coastal Engineering*, 58(3):239–250.
- 551 Falcieri, F. M., Kantha, L., Benetazzo, A., Bergamasco, A., Bonaldo, D., Bar-
552 bariol, F., Malač?, V., Sclavo, M., and Carniel, S. (2016). Turbulence ob-
553 servations in the Gulf of Trieste under moderate wind forcing and different
554 water column stratification. *Ocean Science*, 12(2):433–449.
- 555 Fedele, F., Benetazzo, A., Gallego, G., Shih, P.-C., Yezzi, A., Barbariol, F., and
556 Ardhuin, F. (2013). Spacetime measurements of oceanic sea states. *Ocean*
557 *Modelling*, 70:103–115.
- 558 Fischler, M. and Bolles, R. (1981). Random sample consensus: A paradigm for
559 model fitting with applications to image analysis and automated cartography.
560 *Communications of the ACM*, 24(6):381–395.

- 561 Fryer, J. G. and Brown, D. C. (1986). Lens distortion for close-range pho-
562 togrammetry. *Photogrammetric Engineering and Remote Sensing*, 52:51–58.
- 563 Gallego, G., Benetazzo, A., Yezzi, A., and Fedele, F. (2008). Wave Statistics and
564 Spectra via a Variational Wave Acquisition Stereo System. In *ASME 2008*
565 *27th International Conference on Offshore Mechanics and Arctic Engineering*,
566 pages 801–808.
- 567 Gallego, G., Yezzi, A., Fedele, F., and Benetazzo, A. (2011). A variational stereo
568 method for the three-dimensional reconstruction of ocean waves. *Geoscience*
569 *and Remote Sensing, IEEE Transactions on*, 49(11):4445–4457.
- 570 Hartley, R. and Zisserman, A. (2004). *Multiple View Geometry in Computer*
571 *Vision*. Cambridge University Press.
- 572 Hasselmann, K., Raney, R. K., Plant, W. J., Alpers, W., Shuchman, R. A.,
573 Lyzenga, D. R., Rufenach, C. L., and Tucker, M. J. (1985). Theory of syn-
574 thetic aperture radar ocean imaging: A marsen view. *Journal of Geophysical*
575 *Research: Oceans*, 90(C3):4659–4686.
- 576 Heikkila, J. and Silven, O. (1997). A four-step camera calibration procedure
577 with implicit image correction. In *Proceedings of the 1997 Conference on*
578 *Computer Vision and Pattern Recognition (CVPR '97)*, CVPR '97, pages
579 1106–, Washington, DC, USA. IEEE Computer Society.
- 580 Hirschmüller, H. (2008). Stereo processing by semiglobal matching and mutual
581 information. *IEEE Transactions on Pattern Analysis and Machine Intelli-*
582 *gence*, 30(2):328–341.
- 583 Jahne, B. (1993). *Spatio-Temporal Image Processing: Theory and Scientific*
584 *Applications*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- 585 Kosnik, M. V. and Dulov, V. a. (2011). Extraction of short wind wave spectra
586 from stereo images of the sea surface. *Measurement Science and Technology*,
587 22(1):015504, 1–9.
- 588 Leckler, F., Ardhuin, F., Peureux, C., Benetazzo, A., Bergamasco, F., and
589 Dulov, V. (2015). Analysis and interpretation of frequency-wavenumber spec-
590 tra of young wind waves. *Journal of Physical Oceanography*, 45(10):2484–
591 2496.
- 592 Lourakis, M. A. and Argyros, A. (2009). SBA: A Software Package for Generic
593 Sparse Bundle Adjustment. *ACM Trans. Math. Software*, 36(1):1–30.
- 594 Mironov, A., Yurovskaya, M., Dulov, V., Hauser, D., and Guérin, C. (2012).
595 Statistical characterization of short wind waves from stereo images of the sea
596 surface. *Journal of Geophysical Research*, 117(October):C00J35.

- 597 Nobach, H. and Honkanen, M. (2005). Two-dimensional Gaussian regression for
598 sub-pixel displacement estimation in particle image velocimetry or particle
599 position estimation in particle tracking velocimetry. *Experiments in Fluids*,
600 38(4):511–515.
- 601 Rodriguez, J. and Aggarwal, J. K. (1990). Stochastic Analysis of Stereo Quan-
602 tization Error. *IEEE Transactions on Pattern Analysis and Machine Intelli-
603 gence*, 12:467–470.
- 604 Schumacher, A. (1939). Stereophotogrammetrische Wellenaufnahmen. *Wiss.*
605 *Ergeb. Dtsch. Atlant. Exped. Forschungs Vermessung. Meteor.* Technical re-
606 port.
- 607 Shemdin, H. and Tran, H. M. (1992). Measuring short surface waves with
608 stereophotography. *Photogram. Eng. Remote Sens.*, 93:311–136.
- 609 Shemdin, O. H., Tran, H., and Wu, S. (1988). Directional Measurement of Short
610 Ocean Waves With Stereophotography. *Journal of Geophysical Research*,
611 93(C11):13891–13901.
- 612 Sutherland, P. and Melville, W. K. (2013). Field measurements and scal-
613 ing of ocean surface wave-breaking statistics. *Geophysical Research Letters*,
614 40(12):3074–3079.
- 615 Tilkov, S. and Vinoski, S. (2010). Node.js: Using javascript to build high-
616 performance network programs. *IEEE Internet Computing*, 14(6):80–83.
- 617 Wanek, J. M. and Wu, C. H. (2006). Automated trinocular stereo imaging
618 system for three-dimensional surface wave measurements. *Ocean Engineering*,
619 33(5-6):723–747.
- 620 Young, I. R., Rosenthal, W., and Ziemer, F. (1985). A three-dimensional analy-
621 sis of marine radar images for the determination of ocean wave directionality
622 and surface currents. *Journal of Geophysical Research: Oceans*, 90(C1):1049–
623 1059.
- 624 Zappa, C. J., Banner, M. L., Schultz, H., Corrada-Emmanuel, A., Wolff, L. B.,
625 and Yalcin, J. (2008). Retrieval of short ocean wave slope using polarimetric
626 imaging. *Measurement Science and Technology*, 19(5):055503.