

1 **An Astronomical Pattern-Matching Algorithm for Automated**
2 **Identification of Whale Sharks (*Rhincodon typus*)**

3
4
5
6 Z. ARZOUMANIAN

7 *Universities Space Research Association, NASA Goddard Space Flight Center, Greenbelt MD,*
8 *20770, USA*

9 zaven@milyway.gsfc.nasa.gov

10 (Corresponding author. Tel: 01-301-286-2547/Fax: 01-301-286-1684)

11
12
13 J. HOLMBERG

14 *3612 SE Stark St., Portland OR, 97214, USA*

15 holmbergius@yahoo.com

16
17
18 B. NORMAN

19 *ECOCEAN, c/o Centre for Fish and Fisheries Research, Murdoch University, South St.,*
20 *Murdoch WA 6150, Australia*

21 ecocean@ozemail.com.au

22
23
24
25
26
27
28
29
30
31
32
33 Suggested running title: *Automated Identification of Whale Sharks*

34
35 Word count: 8369

1
2 **Summary**

3 1. The largest shark species alive today, whale sharks (*Rhincodon typus*) are rare and
4 poorly studied. Directed fisheries, high value in international trade, a highly migratory
5 nature, and generally low abundance make this species vulnerable to exploitation. Mark-
6 and-recapture studies have provided our current understanding of whale shark
7 demographics and life history, but conventional tagging has met with limited success. To
8 aid in conservation and management efforts, and to further our knowledge of whale shark
9 biology, an identification technology that maximizes the scientific value of individual
10 sightings is needed.

11 2. We describe a novel technique for identifying individual whale sharks through
12 numerical pattern-matching of their natural surface 'spot' colorations. Together with
13 scarring and other visual markers, spot patterns captured in photographs of whale shark
14 flanks have, in the past, been used to make identifications by eye. We have automated this
15 process by adapting an algorithm originally developed in astronomy for the comparison of
16 star patterns in images of the night sky.

17 3. In tests using a set of previously identified shark images, our method correctly
18 matched pairs exhibiting the same pattern in more than 90% of cases. From a much larger
19 library of previously unidentified images, it has to date produced more than 100 new
20 matches. Our technique is robust in that the incidence of false positives is low, while
21 failure to match images of the same shark is predominantly attributable to projection
22 effects in photographs not ideally oriented with respect to the shark's flank.

23 4. We describe our implementation of the pattern-matching algorithm, estimates of its
24 efficacy, its incorporation into the new Web-based ECOCEAN Whale Shark Photo-
25 identification Library, and prospects for its further refinement. A subsequent paper
26 (Norman et al., in preparation) will discuss in greater detail the biological and
27 conservation implications of the capability to identify individual sharks across wide
28 geographical and temporal spans.

29 5. *Synthesis and applications.* An automated photo-identification technique has been
30 developed that allows for efficient 'virtual tagging' of spotted animals for population and
31 other studies through mark/recapture analyses. The pattern-matching software has been
32 implemented through a Web-based system designed for the management of encounter
33 photographs and derived data. The combined capabilities have demonstrated the
34 reliability of whale shark spot patterns for long-term identifications, and promise new
35 ecological insights. Applications to other species are anticipated.

36
37 *Key-words:* conservation, population studies, marine and fisheries management,
38 mark/recapture, whale shark

1 Introduction

2 The whale shark (*Rhincodon typus*) is one of approximately 370 species of shark alive today
3 (Last & Stevens 1994), and is a member of the order Orectolobiformes, which are predominantly
4 bottom-dwelling species, e.g. wobbegong and carpet sharks (Compagno 1988). Whale sharks
5 have a broad distribution in tropical and warm temperate seas, usually between latitudes 30°N
6 and 35°S (Last & Stevens 1994; Norman 1999). The species is regarded as rare, with as few as
7 320 sightings documented prior to the mid-1980s (Wolfson 1986). Many aspects of whale shark
8 biology remain poorly studied, particularly with regard to life history and demographics. The
9 World Conservation Union (IUCN) Red List of Threatened Species¹ lists the whale shark as
10 vulnerable to extinction (Norman 2000).

11 A number of outstanding questions in whale shark ecology may be addressed through
12 collection and subsequent collation of sighting data around the world (Norman 2004). Mark-and-
13 recapture studies can be used in any situation where animals can be ‘marked,’ or otherwise
14 identified, and ‘recaptured,’ or identified later by resighting (Lettink & Armstrong 2003).
15 Analysis of resultant data can be used to estimate abundance, survival, recruitment, and
16 population growth rates over time (Thompson et al. 1998). Importantly, this research can enable
17 an improved assessment of the global conservation status of this species.

18 Whale sharks are born with unique body patterning on their skin that is retained throughout
19 their lives (Norman 2004). Similar to a fingerprint in humans (Taylor 1994; Norman 1999), this
20 natural patterning of lines and spots—specifically behind the gill slits—shows no evidence of
21 significant change over years and can therefore be used to identify individual sharks. Through the
22 combination of photographed encounters and spot-pattern matching, a shark may be ‘tagged’
23 without any physical contact or interference with the animal. In an early effort, Norman (1999)
24 established a photo-identification library of whale sharks at Ningaloo Reef, Western Australia,
25 with photographs of individual sharks examined by eye for identifying characteristics including
26 spot patterns.

27 While it may be possible to manage small numbers of shark identification photographs and
28 distinguish individuals by eye, the process becomes inefficient and unreliable when collating data
29 from many individual animals sighted in a large number of regions throughout the world. The
30 availability of large quantities of data has rendered manual photo identification unfeasible,
31 motivating instead the development of a Web-based library with the capability of scanning an
32 entire database of encounter photographs in an automated way.

33 In this paper, we present a numerical method for identifying individual whale sharks by the
34 unique patterning of their surface spots. Our technique is adapted from an algorithm developed
35 within the astronomical community for stellar pattern recognition. It has been incorporated into
36 the ECOCEAN Whale Shark Photo-identification Library,² an online database facility that
37 archives digital images of whale sharks submitted by researchers and other interested parties.
38 With a sophisticated pattern-matching capability and a growing library of images, a scientifically
39 valuable number of individual sharks may be identified across wide geographic and temporal
40 spans, enabling significant new advances in the study of whale shark life histories, migration

¹ <http://www.redlist.org/>

² <http://photoid.whaleshark.org/>

1 patterns, and demographics.

2 **Materials and Methods**

3 THE SHEPHERD PROJECT AND THE PROVENANCE OF PHOTO-ID LIBRARY DATA

4 To support the collection and centralization of biological data by wildlife researchers, the
5 Shepherd Project was begun in 2002 with the goal of creating a reusable World Wide Web-based
6 catalog framework that allows for the management of mark/recapture data accumulated by a
7 global research community and interested third parties, such as ecotourists or government
8 management agencies. This framework, which combines an object-oriented database, image
9 management and protection functionality, an extensible programming interface, parameter search
10 and multi-format data export capabilities, was completed in 2004 and first used in the
11 ECOCEAN Whale Shark Photo-identification Library. The Library, built upon a J2EE software
12 platform (Sun Microsystems, Inc., California, USA), is a repository for whale shark spot-pattern
13 data and the photographs from which they are derived. Basic information required to accompany
14 whale shark identification photographs includes a) date and location of the sighting, b) sex and
15 size of the animal, and c) contact details of the submitter. The Library also served as the platform
16 upon which our automated pattern-matching algorithm was developed and tested (Fig. 1). Its
17 Web-based nature allows for worldwide access to data and the results of pattern-matching
18 database ‘scans,’ which seek out and rank similarities in whale shark markings in a manner
19 similar to an Internet search engine. Finally, the Library provides a data export facility to support
20 trending and population analyses using software such as Microsoft Excel or Program Mark,
21 allowing submitted data to become instantly useful to researchers and management agencies.

22 While most of the raw data available for the development of our pattern-matching technique
23 was collected through the research of one of us (see Norman 1999), data submissions to the
24 ECOCEAN Library from ecotourists, researchers, tour operators, managers, etc., have been
25 made, to date, from participants in 19 countries where whale sharks have been sighted (see
26 Acknowledgments). Of particular importance has been the availability of sighting data spanning
27 a 12-year period, 1992–2004, especially to confirm the efficacy of spot patterning as a reliable
28 long-term identification tool.

29 For the identification of individual sharks, we select an area (the ‘measurement region’)
30 located directly behind the gill slits on both the right and left sides of the shark. The region is
31 bounded as follows: a) anteriorly by the fifth gill slit, b) ventrally by the insertion plane of the
32 pectoral fin, c) posteriorly by a line drawn vertically from the insertion point of the trailing edge
33 of the pectoral fin, and d) dorsally by the most ventral of the three longitudinal ridges (Fig. 2).
34 This is an area that can be easily photographed by a diver or snorkeller while swimming
35 alongside the shark. To ensure accurate photo-identification using the tools described in this
36 paper, underwater photographers are encouraged to position their cameras as nearly as possible
37 over the center of the measurement region, with the field of view including both the vertebral
38 column above and the pectoral fin below. Photographs of any secondary identification features
39 are also encouraged, e.g. scarring on fins or body that can be used to further confirm the identity
40 of an individual shark. In addition to digital photographs and video, hard copy images can be
41 scanned and subsequently submitted to the Library for analysis.

1 SPOT EXTRACTION METHODOLOGY

2 A reliable computer-driven pattern matching system must be capable of clearly discerning
3 features of interest in the measurement region of an image. The contrast of white whale shark
4 spots on colored skin is well suited to a common image manipulation process called 'blob
5 extraction,' where, in this case, the blobs are the white spots to be distinguished from the darker
6 background. The spatial relationships between the spots, as represented by a set of derived (x,y)
7 coordinates, form the basis for the unique spot data pattern for each shark.

8 The variability of the underwater environment can compound the problem of computer-based
9 feature recognition by introducing limiting factors such as low visibility or bright surface sunlight
10 that may wash out spot patterns. Moreover, a whale shark may be photographed in configurations
11 such that its anterior-posterior line forms an angle to the image horizontal. A variety of image
12 processing techniques are available that can compensate for these undesirable environmental
13 conditions, such as image rotation and sharpening of contrast by balancing color levels.

14 **Rotation Correction.** To standardize the shark's orientation before spot extraction, a straight,
15 thin horizontal reference line is drawn, using a graphics package (*Fireworks MX 2004*,
16 Macromedia, San Francisco, USA), over each image. The underlying image is then rotated
17 until the segment of the shark's curved vertebral column directly above the measurement
18 region is made parallel to the reference line. Finally, the source image is cropped down to the
19 boundaries of the measurement region.

20 **Contrast Enhancement for Noise Reduction.** Because blob extraction algorithms rely on a
21 single color and, in some cases, allow for varying shades thereof, to differentiate a blob from
22 its local background, care must be taken to ensure that 'noise pixels' of that color do not
23 appear elsewhere in the image and cause false blobs, or in this case false white spots, to be
24 counted and measured. This is especially problematic in images of whale sharks because they
25 are most often photographed during daylight and near the surface. Surface sunlight washing
26 over the shark's body can cause white pixels to appear in the source image that may be
27 mistaken for spots by the blob extraction software. Artifacts of digital compression can also
28 cause spurious white spots. False spots can interfere with pattern matching and increase the
29 computation time for processing by forcing additional calculations.

30 To reduce white pixel noise, *Fireworks* is first used to paint pure white spots on top of the
31 natural shark spots, covering each with a best-fit circle. The underlying contrast and
32 brightness of the source image, but not of the painted white circles, are then reduced, thereby
33 increasing the overall contrast between the artificially superimposed white spots and any noisy
34 white pixels (see Figs 1 and 7). The likelihood of extracting spurious spots is thus essentially
35 eliminated.

36 Once the digital source image has been reduced to a cropped, rotation-corrected, and contrast-
37 enhanced grayscale image, the spot pattern can be identified by blob extraction software and
38 stored in a database. In this case, a custom application was written using a commercial library
39 (*eVision EasyObject*, Euresys, Angleur, Belgium) to measure the coordinates of the center of
40 gravity of the spots in the processed image and to transmit these via the Internet to the
41 ECOCEAN Whale Shark Photo-identification Library for storage as a matchable digital
42 identifier. The entire extraction process requires approximately 10 minutes for an experienced
43 operator.

1 AN ASTRONOMICAL PATTERN COMPARISON ALGORITHM

2 Astronomers are frequently confronted with the task of identifying (and precisely locating within
3 a standard coordinate system) stars, galaxies, and other celestial objects that appear in images of
4 the night sky. A typical technique involves comparing newly acquired images, which may be
5 arbitrarily magnified, rotated, or inverted, with cataloged images of the same region of the sky—
6 the positions of objects common to both images are used to derive the geometric relationship
7 between the coordinate axes that underlie each image. With the advent of digital imaging and
8 large machine-readable datasets, automated methods were needed to carry out these tasks, and in
9 particular the difficult first step of identifying common objects. A useful approach might, for
10 example, identify an imaged region of the sky by its characteristic pattern of stars.

11 Groth (1986) developed a pattern-matching algorithm based on the comparison of two lists of
12 coordinates, i.e. the (x,y) positions of stars in astronomical images, that effectively identifies
13 individual points from one list with their likely counterparts in the other. The algorithm achieves
14 the desired insensitivity to image magnification, rotation, and inversion by forming triangles from
15 selected triplets of coordinate points (Fig. 3). Geometrically similar pairs of triangles, one from
16 each list, are then identified, and a ‘voting’ process provisionally flags points that appear in
17 multiple triangle pairs as being common to both lists. A second iteration of the technique using
18 the provisionally identified points as input confirms or refutes their identification. The method,
19 which has been implemented as part of several astronomical data reduction packages, e.g. the
20 Hubble Space Telescope’s STSDAS,³ and is cited in the literature (e.g. Schmidt et al. 1998), has
21 been demonstrated to be reliable even when the two lists of coordinates have as few as 25% of
22 their points in common.

23 We have adapted this pattern-matching algorithm to the problem of identifying whale sharks,
24 replacing stellar positions with the (x,y) coordinates of prominent spots in photographs of shark
25 flanks. Here, we summarize the algorithm in its original form, following much of Groth’s
26 notation and describing the algorithm’s basic functioning. In the next section, we describe
27 changes that we have made to optimize the method for use with whale shark spot data.

28 Groth’s triangle-based algorithm comprises the following steps. Hereafter, **A** refers to data
29 derived from, e.g. a newly acquired image, and **B** refers to cataloged data. We assume, for the
30 purposes of this description, that coordinate lists **A** and **B** contain the same number (n) of points,
31 but the algorithm does not require lists of equal length.

32 **Filtering of coordinate lists.** The coordinates of stars or spots (generically, ‘points’) in each list
33 are renormalized from their natural units, i.e. pixels, to the unitless interval $[0,1]$ while
34 preserving the aspect ratios of the original images. A user-adjustable tolerance parameter,
35 denoted ϵ , is defined to quantify the typical uncertainty of coordinate measurements; Groth
36 suggests the unitless value $\epsilon = 0.001$. To avoid confusion in pattern-matching, the coordinates
37 in each list are then inspected to flag pairs of points that are too close together: any separation
38 less than a fixed multiple of the uncertainty (e.g. 3ϵ) is deemed too small and one of the points
39 in the pair is purged from the list.

40 **Formation of triangles.** Every combination of three points within each coordinate list describes
41 a triangle, with a point at each vertex. For **A** and **B** separately, all possible triangles are
42 formed and their vertices indexed according to each triangle’s shape: the shortest side is
43 defined to lie between vertices 1 and 2, the intermediate side between vertices 2 and 3, and the

³ http://www.stsci.edu/resources/software_hardware/stsdas/

1 longest side between vertices 1 and 3. The following geometric properties of all triangles are
 2 then computed, where (x_1, y_1) , (x_2, y_2) , and (x_3, y_3) are assumed to be the coordinates of
 3 vertices 1, 2, and 3, respectively:

- 4 • The ratio of the longest (r_3) to the shortest (r_2) sides, $R = r_3 / r_2$, where

$$5 \quad r_2 = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} \quad (1)$$

$$6 \quad r_3 = \sqrt{(x_3 - x_1)^2 + (y_3 - y_1)^2}. \quad (2)$$

- 7 • The cosine of the angle at vertex 1,

$$8 \quad C = \frac{1}{r_3 r_2} [(x_3 - x_1)(x_2 - x_1) + (y_3 - y_1)(y_2 - y_1)]. \quad (3)$$

- 9 • Tolerances in R (t_R) and C (t_C), assuming the coordinate uncertainty ϵ to be
 10 independent in x and y and propagating the measurement uncertainty through the
 11 expressions above,

$$12 \quad t_R^2 = 2R^2 F \quad (4)$$

$$13 \quad t_C^2 = 2S^2 F + 3C^2 F^2, \quad (5)$$

14 where

$$15 \quad F = \epsilon^2 \left(\frac{1}{r_3^2} - \frac{C}{r_3 r_2} + \frac{1}{r_2^2} \right)$$

16 is convenient shorthand and S is the sine of the angle at vertex 1.

- 17 • The logarithm of the triangle's perimeter, $\log p$.
- 18 • The orientation, i.e. whether the vertices 1, 2, and 3 are traversed in a clockwise or
 19 counterclockwise sense.

20 **Filtering of triangles.** For a coordinate list of length n , the number of triangles generated will be
 21 $n_t = n(n-1)(n-2)/6$. The results of the computations above are cumulated into new lists that
 22 record the properties of all n_t triangles for **A** and **B** separately. Figure 4 shows the
 23 distributions of the R and C values of triangles derived from the whale shark encounter image
 24 in the left-hand panel of Fig. 1.

25 Not all triangles are well-suited to pattern matching and some filtering is therefore necessary.
 26 Triangles with large length ratios (Groth recommends $R > 10$, we use $R > 8$ in Fig. 4) are
 27 discarded from both lists. Such elongated triangles produce large tolerances through eqn 4; as
 28 a result, they can be falsely matched (according to the criterion, eqn 6, below) with many
 29 triangles, weakening the algorithm's ability to discriminate between different patterns.
 30 Although not required in Groth's original formulation, Fig. 4 shows a cutoff in the value of C
 31 beyond which we remove triangles from further analysis; we describe in the next section the
 32 motivation for this additional filtering criterion.

33 **Matching of triangles across lists.** A given length ratio R and internal-angle cosine C together
 34 describe a unique class of geometrically similar triangles, within which triangles differ only in
 35 their relative size, i.e. by a magnification factor, and their orientations. At the heart of the
 36 pattern-matching algorithm, each **A** triangle's R and C values are compared with those for
 37 triangles from **B** according to matching criteria that depend on the tolerances t_R and t_C ,

$$38 \quad (R_A - R_B)^2 < t_{R_A}^2 + t_{R_B}^2 \quad (6)$$

$$39 \quad (C_A - C_B)^2 < t_{C_A}^2 + t_{C_B}^2, \quad (7)$$

1 where both inequalities must be satisfied to declare a pair of triangles successfully matched.
 2 This computationally intensive search for similar triangles across the two lists can be
 3 optimized by searching only a subset of **B** triangles for each **A** triangle, as described by Groth.
 4 Where the lists differ in length, we define list **A** to be the one with the smaller number of
 5 filtered triangles. If more than one triangle from list **B** satisfies these criteria for a single **A**
 6 triangle, only the closest match—i.e. with the smallest value of the sum of the left-hand sides
 7 in eqn 6—is retained.

8 For each pair of **A** and **B** triangles with similar geometry, the relative magnification factor (M)
 9 is computed:

$$10 \quad \log M = \log p_A - \log p_B. \quad (8)$$

11 If the **A** and **B** images contain the same point pattern, corresponding triplets of points will
 12 form many matching triangles all related by a common magnification factor. By contrast, any
 13 falsely matched triangles, i.e. **A** and **B** triangles that coincidentally have similar geometries
 14 but do not arise from the same triplet of points in the two images, will be related by an
 15 arbitrary magnification factor. True matches can therefore be distinguished from false matches
 16 by examining the distribution of magnification values. Figure 5 shows the distribution of
 17 $\log M$ for the sample comparison depicted in Fig. 1. The prominent peak at $\log M$ values near
 18 zero (the expected value when both images have been scaled to unit linear dimensions) is
 19 dominated by true triangle matches, with a smaller contribution within this peak from the
 20 more broadly distributed false matches.

21 Similarly, the orientations (clockwise vs. counter-clockwise) of member triangles among the
 22 matched pairs provide useful information. All true matches should have the same relative
 23 orientation, identical or opposite sense depending on whether the two datasets are mirror-
 24 images of one another. By contrast, the set of false matches should reflect a random mix of
 25 same-sense and opposite-sense triangle pairs. This feature provides a rough estimate of the
 26 number of true, m_T , and false, m_F , matches found in the comparison set. If n_+ and n_- refer to
 27 the number of same-sense and opposite-sense matches, respectively, then

$$28 \quad m_T = |n_+ - n_-| \quad (9)$$

$$29 \quad m_F = n_+ + n_- - m_T. \quad (10)$$

31 To isolate the true matches, Groth describes a simple iterative filter that adapts itself to the
 32 distribution of magnifications and relative orientations for all matches. In each iteration, the
 33 mean and standard deviation of $\log M$ values are computed, and matches are discarded if they
 34 require magnifications more than z standard deviations from the mean value, where

$$35 \quad z = \begin{cases} 1, & \text{if } m_F > m_T \\ 3, & \text{if } m_F < 0.1 m_T \\ 2, & \text{otherwise.} \end{cases} \quad (11)$$

37 Iterations continue until one of the following conditions is met: no matches are discarded in an
 38 iteration, no matches remain in the comparison set, or the number of iterations reaches a preset
 39 limit, e.g. 10. If no matches remain, the two datasets are declared to be different and the
 40 algorithm terminates. Otherwise, it is assumed that, if same-sense matches outnumber
 41 opposite-sense matches, there is no coordinate inversion between the two datasets, and the
 42 remaining opposite-sense matches are discarded. Alternatively, if $n_- > n_+$, coordinate

1 inversion is assumed and the remaining same-sense matches are discarded.

2 **Voting to identify points in common.** At this stage, the algorithm has produced a number of
3 matched triangle pairs, each of which involves three pairs of ostensibly matched vertex points.
4 To determine which points are truly common to both datasets, it is assumed that matching
5 points have a high probability of participating in more than one, likely many, matching
6 triangles. This expectation is quantified through a ‘voting’ scheme: every matched triangle
7 pair casts three votes, one for each vertex pair. When all votes are cumulated, point pairs are
8 ranked according to the number of votes they have received. If no pair receives more than one
9 vote, the datasets are declared to be different. Otherwise, high-ranking pairs are assigned to
10 one another as credibly matched points. These assignments continue until one of three
11 conditions is met: the number of votes drops by a factor of two, a previously assigned point
12 from either dataset reappears in a different pair or, less commonly, the vote count drops to
13 zero.

14 **Iteration.** Finally, the entire algorithm is run a second time, with input restricted only to those
15 points that were matched in the first pass. Groth’s formulation strictly requires that all input
16 points be successfully matched in the second pass as well.

17 THE SPOT-MATCHING ALGORITHM

18 We have tailored Groth’s original astronomical algorithm to the task of identifying whale sharks,
19 with modifications that reflect the properties of typical whale shark spot patterns and the data
20 preparation and extraction procedures we use. These changes increase the algorithm’s robustness
21 but, at the same time, reduce its generality with respect to inversions and arbitrary rotations
22 between the comparison datasets. Building on the description of the original algorithm above, we
23 describe our changes, their motivations, and their implications here.

24 **Formation of triangles.** We supplement the triangle properties considered by Groth— R , C , their
25 tolerances, $\log p$, and orientation—with the following additional quantities:

- 26 • A measure of each triangle’s rotation relative to the image x axis. The rotation
27 angle is defined as a polar coordinate for vertex 1,

$$28 \quad \theta = \tan^{-1} \left[\frac{y_1 - y_c}{x_1 - x_c} \right], \quad (12)$$

29 where the origin (x_c, y_c) corresponds to the triangle centroid,

$$30 \quad x_c = \frac{1}{3}(x_1 + x_2 + x_3) \quad (13)$$

$$31 \quad y_c = \frac{1}{3}(y_1 + y_2 + y_3). \quad (14)$$

32 In other words, rotation is determined by the angle formed between a triangle’s median
33 line for vertex 1 (the line joining vertex 1 with the triangle centroid) and the image
34 horizontal axis. We adopt this ‘local’ measure of rotation over one that encompasses the
35 whole image, e.g. the angle formed by the longest side of each triangle relative to the x
36 axis, because it provides a degree of useful insensitivity to distortions of the image caused
37 by the shark’s curved body and projection effects in photographs not taken at right angles
38 to the shark’s flank.

- 39 • We quantify each triangle’s size, s , simply adopting the fractional length of its
40 longest side, r_2 in equations 1–2, relative to the maximum value r_2^{\max} of any triangle in
41 the image:

$$s = r_3 / r_3^{\max}. \quad (15)$$

Filtering of triangles. Along the flanks of whale sharks, and especially tailward of the pectoral fin, the distribution of spots typically becomes somewhat regular, falling along curved ventral-dorsal lines. Because the Groth algorithm forms triangles from all possible coordinate triplets, a significant number of narrow, flattened triangles (with C values of nearly 1.0) are generated in which all three vertices lie along a single arc. Such triangles from one image have a high probability of matching a large number of similarly 'flat' triangles from arcs in any arbitrary comparison image. In this circumstance, matched triangles provide little useful information for identification of a unique pattern—the anticipated sharp peak in the distribution of magnifications can be diluted by the many falsely matched triangles. To prevent these uninformative triangle matches from overwhelming correctly matched triangle pairs, we impose a constraint on the C -values of triangles retained for analysis of $C < 0.99$. Our tests have shown that this filter strongly suppresses triangles that produce unwanted false matches. As Fig. 4 demonstrates, our C filter does not overlap significantly with Groth's original R filter in the triangles that are disqualified.

Projection effects related to the photographer's vantage point (see below) distort triangles, making them difficult to match. The distortion is greatest for triangles that span nearly the entire image, i.e. where the vertices lie near opposite edges of the measurement region. We therefore filter out these large triangles by requiring that $s < s_{\max}$, where tests show that a value $s_{\max} = 0.85$ provides a good balance between rejecting distorted triangles and retaining useful ones.

Matching of triangles across lists. The triangle matching criteria in Groth's formulation, equations 6–7, are supplemented by a rotation criterion,

$$\theta_A - \theta_B < \theta_{\max}, \quad (16)$$

where θ_{\max} is a user-selected parameter. The relative rotation between pairs of images in our spot-pattern database is, by construction, small: as described earlier, we rotate each image to align the shark's vertebral column with the horizontal axis. This information is used to match triangles—the rotational invariance of Groth's original algorithm, while useful for astronomical images, unnecessarily diminishes the method's effectiveness when both lists of spot coordinates are known to be based on the same coordinate system. If more than one pair of triangles is deemed a match according to equations 6, 7, and 16, the pair with the smallest quadrature difference δ is retained, where

$$\delta^2 = \frac{(R_A - R_B)^2}{r_{R_A}^2 + r_{R_B}^2} + \frac{(C_A - C_B)^2}{r_{C_A}^2 + r_{C_B}^2} + \frac{(\theta_A - \theta_B)^2}{\theta_{\max}^2}. \quad (17)$$

Similarly, we do not require the original algorithm's insensitivity to inversion of one of the images: we assume that photographs submitted to the database are correctly oriented. We nevertheless track the number of opposite-sense triangle matches, n_- , in applying a variant of Groth's iterative filter on the $\log M$ distribution. In our version, the mean and standard deviation are computed only for same-sense triangles, with the multiplier z determined as follows:

$$z = \begin{cases} 1, & \text{if } n_- > n_+ \\ 3, & \text{if } m_F < 0.5m_T \\ 2, & \text{otherwise.} \end{cases} \quad (18)$$

Iterations continue until no matches remain, none were eliminated in the last iteration, or 20

1 iterations are made. Following this procedure, which effectively isolates true from false
2 matches, the remaining opposite-sense matches are discarded.

3 **Iteration and scoring of encounters.** Voting for spot matches proceeds as in the original
4 algorithm. A second pass through the entire code, with matched spots as input, effectively
5 filters out any points incorrectly identified in the first pass. Our implementation departs from
6 Groth's procedure at this point in allowing single spots to be eliminated during the second
7 pass without disqualifying the comparison pair of images as a potential match.

8 When two spot datasets are compared, a score is computed by summing the votes awarded to
9 each pair of successfully matched spots: if v_i represents the number of votes cumulated for the

10 i th pair of spots, the sum $V = \sum_{i=1}^m v_i$ terminates, in the typical case, when $v_{m+1} < v_m/2$. The vote
11 total V is a useful measure of the similarity between the two input spot patterns.

12 Comparisons across different datasets, i.e. whether the patterns in a pair of images are more
13 closely matched than the patterns in a different pair of images, must however be interpreted
14 with care, because the maximum possible score is not fixed; it is determined specifically by
15 the number of triangles in the smaller filtered dataset. To account in part for this difference,
16 the algorithm also reports the number of triangles that contributed votes for spots at their
17 vertices (see, e.g. Fig. 3), as a fraction f_T of all available (filtered) triangles. We adopt as a
18 final score S for ranking purposes the product of the vote total and this fraction of successfully
19 matched triangles,

$$S = f_T V. \quad (19)$$

21 As we show in the following section, a score value $S > 100$ is a reliable indicator that the
22 algorithm has identified two instances of the same spot pattern, while $S > 10$ provides
23 evidence of a potential match worthy of further review.

24
25 Table 1 summarizes the algorithm's adjustable parameters (e.g. coordinate tolerances and
26 filtering criteria), the values recommended for astronomical images by Groth, and the values we
27 find provide the most robust performance for matching whale shark spot patterns. Optimized
28 values were found in most cases by examining the triangle properties of a handful of comparison
29 pairs in detail, while others were derived by examining the scores of all visually confirmed
30 matches as the parameters were varied. Users of the Library are provided the opportunity to alter
31 these quantities to explore the algorithm's behavior with different input images.

32 Results

33 Our spot pattern matching technique was applied in database 'scans': as new whale shark
34 photographs were submitted to the ECOCEAN Library, spot data were extracted and compared to
35 pattern data from all previously submitted images, separately for left and right flank images. A
36 list of candidate image matches was then produced, usually less than ten in number, ranked
37 according to the score computed by the algorithm. A subset of the Library entries, or
38 'encounters,' represented multiple images of individual sharks. As described below, these
39 instances proved useful in estimating the method's self-consistency: if encounter A matches
40 encounter B, and B matches encounter C, the technique should also provide a match when A is
41 compared directly to C.

1 CORRECT MATCHES: THE METHOD'S EFFICACY

2 To explore the method's success rate and any potential difficulties, spot patterns for each of 27
3 previously identified (i.e. matched by eye) left-side images were scanned across all other
4 available left-side spot datasets. As of Dec. 1, 2004, there were 271 such datasets. Similarly,
5 eight known right-side images were compared to the remaining catalog of 181 right-side datasets.
6 In the vast majority of cases, comparisons involving different sharks produced a zero score; in
7 some cases, however, a small non-zero score resulted. We refer to the latter as 'false positive'
8 matches. When the same shark was imaged in both encounters, a high score typically resulted, an
9 outcome we refer to as a 'correct match.' Occasionally, comparison of two same-shark images
10 produced a very low score, or a 'failed match.'

11 Figure 6 summarizes the results of these tests. The distributions of vote totals V , matched
12 triangle fractions f_T , and product scores S resulting from the comparison of all 27 previously
13 identified pairs of encounters are shown in green. For the same set of comparisons, all false-
14 positive match scores reported by the algorithm were accumulated—the resulting distribution is
15 shown in red. A reliable method for identifying unique patterns should minimize the overlap in
16 the red and green histograms. We find that for vote totals V (top panel of Fig. 6) the distribution
17 of false positives is broad and encroaches, at the high end, on the vote totals garnered by the
18 correct matches. An essential discriminator appears, however, in the triangle fraction (middle
19 panel): when f_T is restricted to values greater than 5%, the number of false positive matches
20 drops from 236 to 11, while just three correct matches are also flagged, two of which had, in any
21 case, the lowest vote totals V . In the bottom panel, the product score S incorporates the additional
22 information contained in f_T : we find that the distribution of S for false positive matches is well-
23 described by a log-normal which drops off rapidly for $S > 10$.

24 The available sample of previously matched pairs of encounters is small but, we believe,
25 representative of the underlying statistical properties of correct and, especially, false-positive
26 match scores. The results shown in Fig. 6 therefore suggest an empirical scheme for classifying
27 the quality of a pattern match as scored by our algorithm:

- 28 • A non-zero score S less than 10 is unlikely to represent a true match, but rather is
29 characteristic of a false-positive.
- 30 • A score between 10 and 100, especially with a fraction f_T greater than 5%, represents a
31 moderately strong likelihood that the two patterns under comparison are truly matched.
- 32 • Any score above 100 represents a strong candidate for a correctly matched pair of spot
33 pattern images. The log-normal distribution of false-positive scores places the $S=100$
34 boundary 3.6 standard deviations above the mean—this implies a formal probability of
35 chance occurrence in this high-confidence category of better than 1 in 6000.

36 Based on these criteria, we can estimate a success rate for the method. From among the 27
37 previously identified pairs of encounters tested, 21 produced scores in the 'strong match'
38 category, another four in the 'moderately strong' category, none were reported as 'weak'
39 candidates, and two failed to match altogether. We combine the two higher-confidence categories
40 to derive a success rate of 25 out of 27, or 92%. Although based on a small sample, this rate is
41 encouraging and may well improve with time, as divers and researchers mindful of the simple
42 requirements of numerical pattern-matching strive to improve their vantage points in obtaining
43 new photographs of whale sharks, as discussed below.

44 FAILED MATCHES AND FALSE POSITIVES: DIFFICULTIES ENCOUNTERED IN

1 APPLYING THE METHOD

2 The performance of pattern-matching techniques is subject to factors beyond the control of any
3 numerical algorithm; our adaptation of the Groth triangle-matching method is no exception. The
4 difficulties that present themselves can be grouped into three categories: image quality, viewing
5 geometry, and spot pattern systematics.

6 Spot extraction from raw whale shark images acquired by divers can be complicated by
7 lighting conditions, shadows, obscuration of spots by other fish, granularity of low-resolution
8 images, and other phenomena. Nevertheless, the triangle-matching algorithm has proven
9 effective even when two images have fewer than half of their spots in common, so that most of
10 these difficulties are simply overcome by careful editing of photographs.

11 The direction from which shark flank images are obtained is important. In photographs taken
12 from an angle away from the perpendicular, projection effects alter the apparent aspect ratio of
13 the spot pattern, changing the geometries of the resulting triangles. Similarly, a camera vantage
14 point too far above or below the center of the measurement region produces altered geometries.
15 The ϵ uncertainty parameter in the triangle-matching algorithm can compensate, in part, for these
16 distortions, and our implementation further mitigates the effects of an oblique image by imposing
17 an upper limit on the size of triangles relative to the image dimensions—large triangles that span
18 nearly the entire image will be most distorted and hardest to match. Nevertheless, an oblique
19 image was responsible for one of the two instances in Fig. 6 in which a previously known match
20 failed to produce a high score. We have experimented with numerical correction of oblique
21 images by trigonometrically adjusting the spacings of spot coordinates in the x -direction
22 immediately following extraction; although dependent on the operator’s estimate of the angle
23 formed between the image plane and the shark’s flank, this technique holds some promise. We
24 note that as the database of encounters grows, the collection of images for a given shark will
25 likely span a range of perspectives, improving the odds that a successful identification will be
26 made. As demonstrated in Fig. 7, photographs obtained from extreme forward or tailward angles
27 will not be correctly matched with each other (we estimate that successful matches can be made
28 for viewing perspectives different by as much as 30°), but each will match other images made at
29 intermediate angles. In the long view, therefore, oblique images of frequently encountered sharks
30 will have a minimal impact on the method’s ability to provide a reliable identification.

31 As described earlier, whale shark spots sometimes fall along curved, neatly arrayed arcs. In
32 rare cases, spots are found to lie, within each arc, at quasi-regular intervals, so that they form a
33 loose grid. We find that, when one image in a comparison pair exhibits such a gridded spot
34 pattern, our triangle-based matching algorithm can produce a relatively high score even when the
35 images correspond to different sharks. Such behavior is not surprising given the large number of
36 geometrically similar triangles that can be produced from spots arranged in a loose grid. Spots
37 arrayed in grids account for the very highest scores ($S \approx 20$) we have found among the false
38 positives, and generally also produce $f_T > 0.05$. In addition to high false-positive scores, gridded
39 spots may also be responsible for failed matches. Indeed, this is the case for the remaining failed
40 match from our previously-identified test dataset, because falsely matched similar triangles from
41 the two images overwhelm those that are correctly matched.

42 In instances where our algorithm fails to establish a strong match, visual inspection or some
43 other method is needed to identify the imaged shark. False positive outcomes are undesirable, but
44 even in cases where the algorithm cannot provide an unambiguous identification for a given
45 submitted image, it reduces dramatically (by a factor of between ten and 100) the number of

1 images that a user need examine visually to uncover a successful identification.

2 'BLIND' MATCHES: THE METHOD'S SUCCESSES

3 To date, a total of 111 image pairs not previously known to be associated have been matched
4 through the use of our algorithm and, of these, 96 had scores $S > 10$. Typically, database scans
5 produced a list of candidate matches, the most highly ranked of which were examined visually
6 for spot-pattern compatibility and unrelated identification markers such as scars. Confirmed
7 matches were noted and tabulated, resulting in the black histograms of score distributions shown
8 in Figure 6. As expected, most of the successful matches have scores in the high-confidence
9 range, with decreasing numbers in the moderate- and low-confidence ranges. We note that not all
10 of the blind matches constitute new *identifications*: in cases where three or more encounters were
11 available for a single shark, all possible image pairs, e.g. three pairs for the shark shown in Fig. 7,
12 were included in the category of blind matches, forming a rough self-consistency test of the
13 method. The high-scoring fraction of $96/111 = 86\%$ among blind matches provides supporting
14 evidence for the method's efficacy. We emphasize that these results have been obtained with a
15 dataset that is not prejudiced against moderately oblique images; it reflects, in other words, a
16 collection of encounter photographs that were acquired under real-world diving conditions.

17

18 Discussion

19 The use of natural spot patterning from digital (or digitized) images and of computer-based
20 pattern recognition offers several new benefits to whale shark mark/recapture studies. A large
21 number of photographs of whale sharks has been acquired by researchers, management agencies,
22 dive operators, and tourists over the past two decades. Photo-based pattern recognition allows for
23 'data mining' of these archives—our method has uncovered verifiable matches among images
24 (photographs as well as still frames captured from video footage) that predate the development of
25 our algorithm. In essence, we have been able to go back and 'mark' sharks that had not
26 previously been physically tagged, thereby increasing the number of sharks that can be
27 'recaptured' in the future. In addition, computer-based pattern recognition solves the problem of
28 scalability inherent in photo-identification by eye. Already, we have found that the number of
29 photographs and pattern samples in the ECOCEAN Library exceed the ability of any single
30 individual to efficiently match new photographs by visual comparison, with each new photograph
31 increasing the amount of time required for positive identification. Rather, visual comparison now
32 simply serves as a final validation of computer-executed scans that can sift through hundreds of
33 patterns with high accuracy in less than one hour. Built into the Web-based framework of the
34 ECOCEAN Library, this automated system allows geographically dispersed researchers to access
35 the most current data, and to make rapid identifications from a communal body of data and
36 research.

37 An added benefit of the photo-identification methodology is that the research community can
38 readily train tourists and dive operators to gather, during recreational encounters with whale
39 sharks and while maintaining a safe and non-threatening distance from the sharks, photographs
40 suitable for use in mark/recapture studies. Underwater cameras and housings are within the
41 budgets of most divers, and the large quantity of data that could in principle be gathered through
42 cooperation with the dive tourism industry can now be efficiently managed and processed using
43 automated pattern-recognition and image management software such as those found in the

1 ECOCEAN Library.

2 Among the highlights of our algorithm's pattern-matching successes to date are high-scoring
3 comparisons of images acquired as much as eight years apart. Future additions to the Library
4 should allow matching across steadily longer time baselines. We find marginally significant
5 evidence for a degradation in pattern-matching fidelity with increasing time spans, as might be
6 expected if spot patterns evolve as sharks grow. It is possible that straightforward recognition of
7 spot patterns may apply only to sharks larger than a certain minimum size, below which rapid
8 growth may shift spot locations in juvenile sharks. To date, the algorithm has made successful
9 multi-year matches with sharks as small as 4.5 meters. A detailed study of this and other
10 biological implications of new identifications will be presented in a subsequent paper (Norman et
11 al., in preparation).

12 **Conclusion**

13 We have developed a method for the automated identification of individual whale sharks from
14 digital images of their spot patterning. The technique, adapted from an algorithm designed for
15 comparing star patterns in astronomical images, is based on the geometric properties of triangles
16 formed by joining all possible combinations of three spot coordinates within a well-defined
17 measurement region. This essentially unique and archivable digital 'fingerprint' can be used as a
18 natural marker to track individual fish over wide geographic areas and time spans much longer
19 than can be achieved with other tracking techniques, provided that a large number of
20 photographic encounters are organized and stored in a single data repository. The ECOCEAN
21 Whale Shark Photo-identification Library, created and maintained by the authors, serves this
22 purpose. At the time of this writing, the Library holds over 1500 images, with more than 270 left-
23 side and 180 right-side spot pattern datasets available for automated identification.

24 Although its performance is susceptible to degrading factors such as image quality,
25 photographic perspective, and the highly organized nature of spot patterns found on a small
26 number of sharks, tests of the method using real-world data show that it identifies pairs of
27 matched images with reliability nearing 90%, while producing a small number of false positive
28 matches that are easily discounted by visual inspection. The algorithm is thus a useful element in
29 a toolbox of research technologies, such as satellite and data logging tags, used to study whale
30 sharks. For long-term population monitoring, virtual tagging eclipses plastic visual-identification
31 tags, as these typically have a life span of less than one year.

32 We continue to work on refinements to the method to improve its robustness and ease of use,
33 and to explore the limits of its capabilities, along several lines. Our implementation currently
34 requires that a trained operator extract spot coordinates from submitted images (the procedure
35 takes about 10 minutes), and inspect the results of the automated scan across the image library.
36 The latter is an unavoidable, and ultimately desirable, check on the method's scoring of image
37 comparisons, but the former task can in principle be further automated to improve efficiency and
38 minimize the possibility of operator error. For example, a more sophisticated filtering scheme for
39 triangle matches could restore the original algorithm's insensitivity to rotations of the image. We
40 are also investigating techniques for extracting pattern information from the locations and shapes
41 of lines that often accompany the spots on whale shark surfaces, in order to broaden the
42 automated pattern-matching capabilities of the ECOCEAN Whale Shark Photo-Identification
43 Library.

44 We are indebted to Dr. G. Nelemans for bringing the Groth algorithm to our attention. We

1 extend special thanks to the Western Australian Department of Conservation and Land
2 Management, Australian Department of Environment and Heritage, ecotourism operators at
3 Ningaloo Marine Park in Western Australia and in many other locations around the world,
4 Murdoch University, the Thyne Reid Education Trust, the Australian Marine Conservation
5 Society, the Fielman Foundation, Rhiannon Bennett, Suzy Quasnichka, Allison Richards, Ed
6 Stastny, and members of the general public who have assisted via submissions to the ECOCEAN
7 Whale Shark Photo-Identification Library.

8 **References**

- 9 Compagno, L.J.V. (1984) *FAO species catalogue. Vol. 4. Sharks of the world. An Annotated and*
10 *illustrated catalogue of shark species known to date*, FAO Fish Synopsis 125, Vol. 4, Pt. 1,
11 Hexanchiformes to Lamniformes, Rome.
- 12 Groth, E.J. (1986) A Pattern-Matching Algorithm for Two-Dimensional Coordinate Lists. *The*
13 *Astronomical Journal*, **91**, 1244–1248.
- 14 Last, P.R. & Stevens, J.D. (1994) *Sharks and Rays of Australia*. CSIRO, Hobart, 513 pp.
- 15 Lettink, M. & Armstrong, D.P. (2003) An introduction to using mark-recapture analysis for
16 monitoring threatened species. New Zealand Department of Conservation *Technical Series*
17 28A, 33 pp.
- 18 Norman, B.M. (1999) *Aspects of the biology and ecotourism industry of the whale shark*
19 *Rhincodon typus in northwestern Australia*. MPhil thesis, Murdoch University, Western
20 Australia.
- 21 Norman, B.M. (2000) In: *2000 IUCN Red List of Threatened Species*. IUCN, Gland, Switzerland
22 and Cambridge, UK. xviii+61 pp. (Book & CD). Also on <http://www.redlist.org/>
- 23 Norman, B.M. (2004) Review of the current conservation concerns for the whale shark
24 (*Rhincodon typus*): A regional perspective. *Coast & Clean Seas Project 2127* Final Report to
25 the Australian Government Department of the Environment and Heritage, 74 pp.
- 26 Norman, B.M., Holmberg, J. & Arzoumanian, Z. (in preparation). Whale shark photo-
27 identification: Unique skin patterning 'tags' individuals.
- 28 Schmidt, B.P., Suntzeff, N.B., Phillips, M.M., Schommer, R.A., Clocchiatti, A., Kirshner, R.P.,
29 Garnavich, P., Challis, P., Leibundgut, B., Spyromilio, J., Riess, A.G., Fillipenko, A.V.,
30 Hamuy, M., Smith, R.C., Hogan, C., Stubbs, C., Diercks, A., Reiss, D., Gilliland, R., Tonry,
31 J., Maza, J., Dressler, A., Walsh, J., Ciardullo, R. (1998) The High-Z Supernova Search:
32 Measuring Cosmic Deceleration and Global Curvature of the Universe Using Type IA
33 Supernovae. *The Astrophysical Journal*, **507**, 46–63.
- 34 Taylor, G. (1994) *Whale Sharks*. Angus & Robertson Publishers, Sydney.
- 35 Thompson, W.L., White, G.C. & Gowan, C. (1998) *Monitoring vertebrate populations*.
36 Academic Press, Inc., San Diego.
- 37 Wolfson, F.H. (1986) Occurrences of the whale shark, *Rhincodon typus*, Smith. *Indo-Pacific*
38 *Fish Biology: Proceedings of the Second International Conference on Indo-Pacific Fishes*. Pp.
39 208–226 (eds. Uyeno, T., Arai, R., Taniuchi, T. & Matsuura, K.) Ichthyological Society of
40 Japan, Tokyo.

41

1 **Tables**

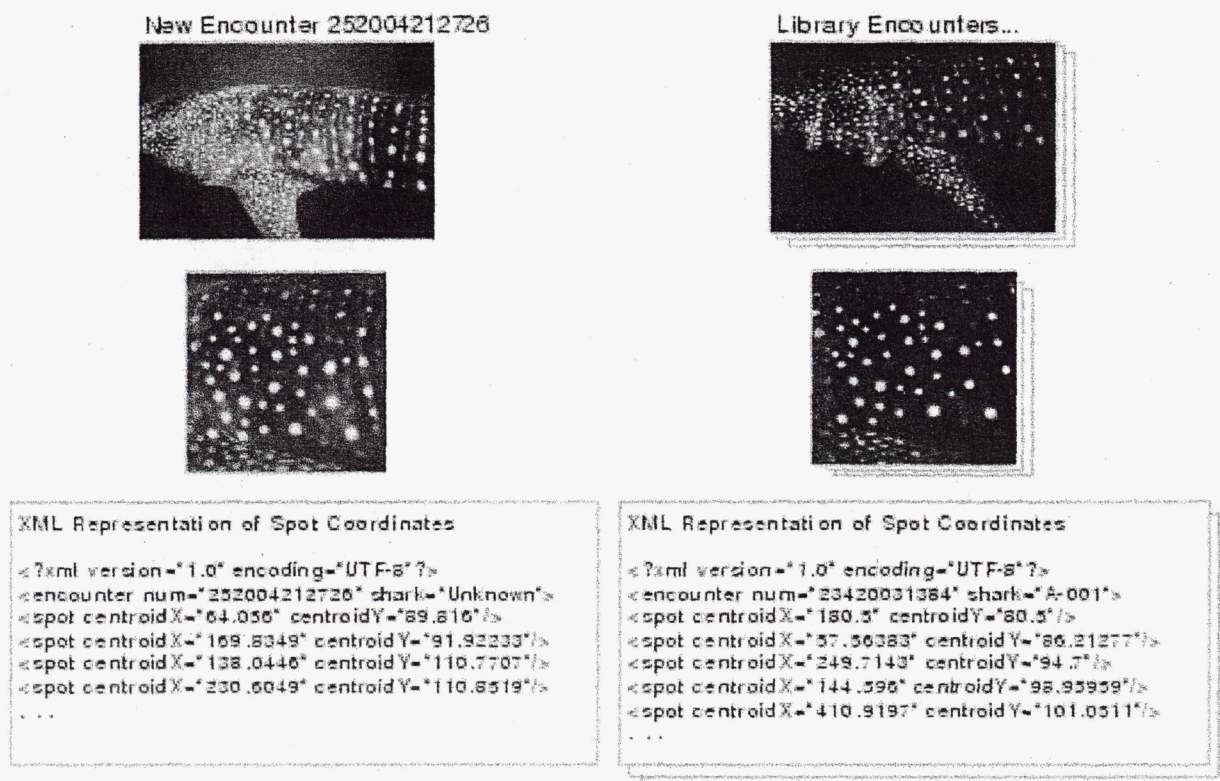
2 Table 1: Adjustable parameters that govern the performance of the triangle-based pattern
 3 matching algorithm. Length units are normalized to the largest distance between two points in an
 4 image.

5

Parameter	Adopted Value		Description
	Groth	This work	
ϵ	0.001	0.01	One-dimensional coordinate uncertainty
R_{\max}	10	8	Maximum triangle-side length ratio
C_{\max}	N/A	0.99	Maximum cosine of angle at vertex 1
s_{\max}	N/A	0.85	Maximum triangle size
θ_{\max}	N/A	10°	Maximum relative triangle rotation

6


1 Figures




2

3 Figure 1: A sample of the ECOCEAN Whale Shark Photo-identification Library's spot pattern
4 dataset. Raw images (*top row*) from newly submitted (*left*) and cataloged (*right*) encounters are
5 processed (see text) to highlight the naturally occurring spots (*middle row*), and a commercial
6 software package is used to extract their coordinates within the image frame. The resulting lists
7 of coordinates (*bottom row*) are then stored and input to the pattern-matching algorithm for
8 identification and virtual 'tagging' of individual sharks.


9



1
2 Figure 2: In a raw image submitted to the Library, the shark's orientation may be arbitrary. In
3 these cases, the image is rotated so that the vertebral column is aligned with the image horizontal
4 and the forward and rear boundaries of the measurement region are vertical. The image is then
5 cropped to isolate the correctly oriented pattern of spots and lines.
6



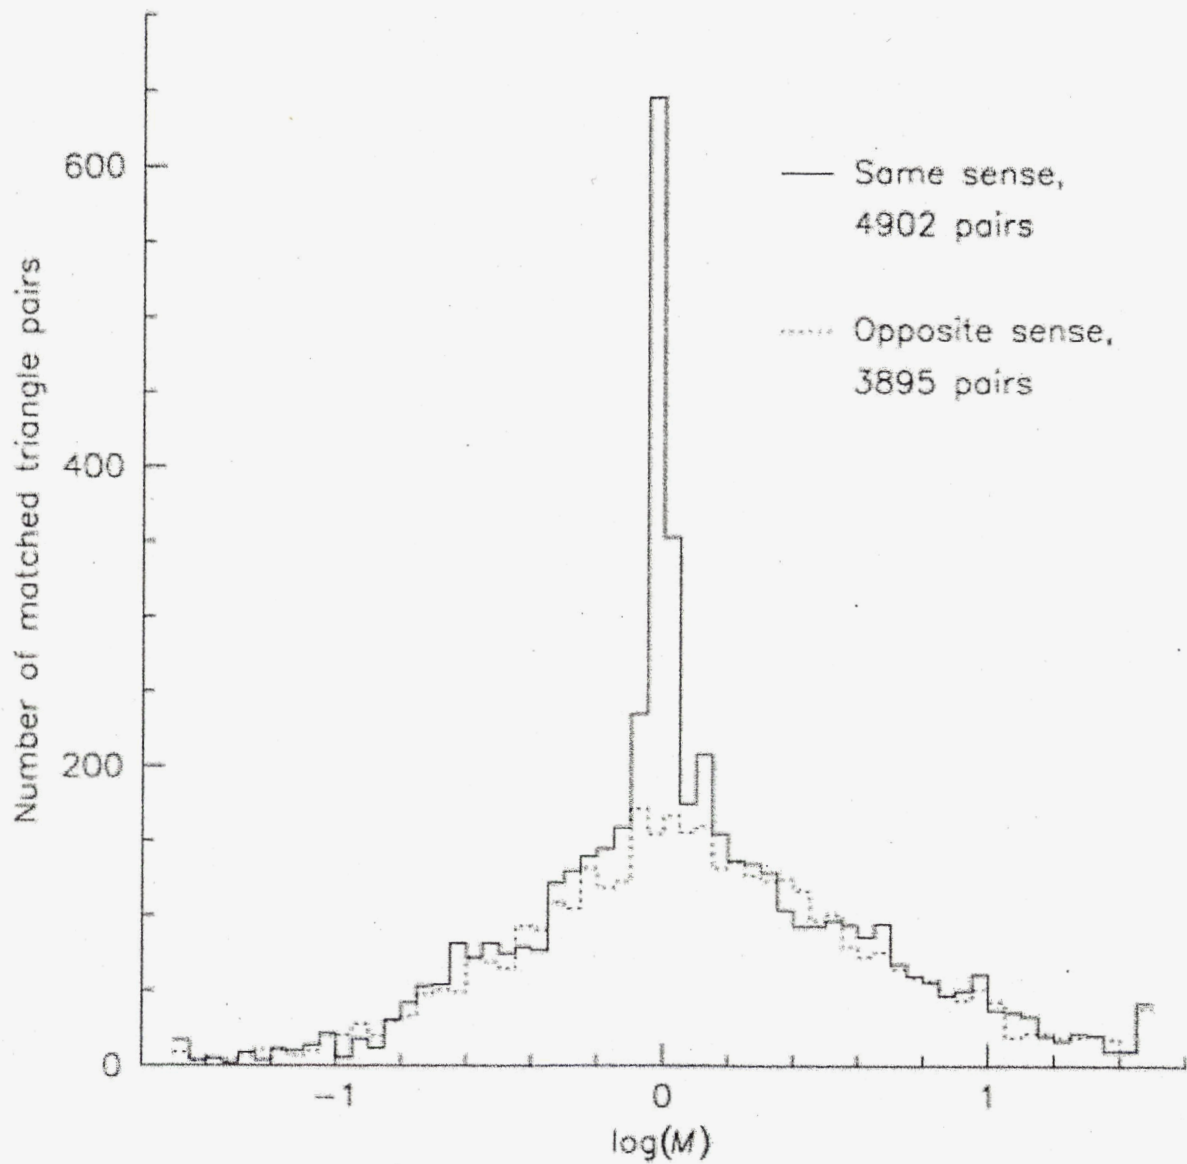
1
 2 Figure 3: A sketch of the basic pattern-comparison process based on the formation of triangles
 3 from triplets of points.



1

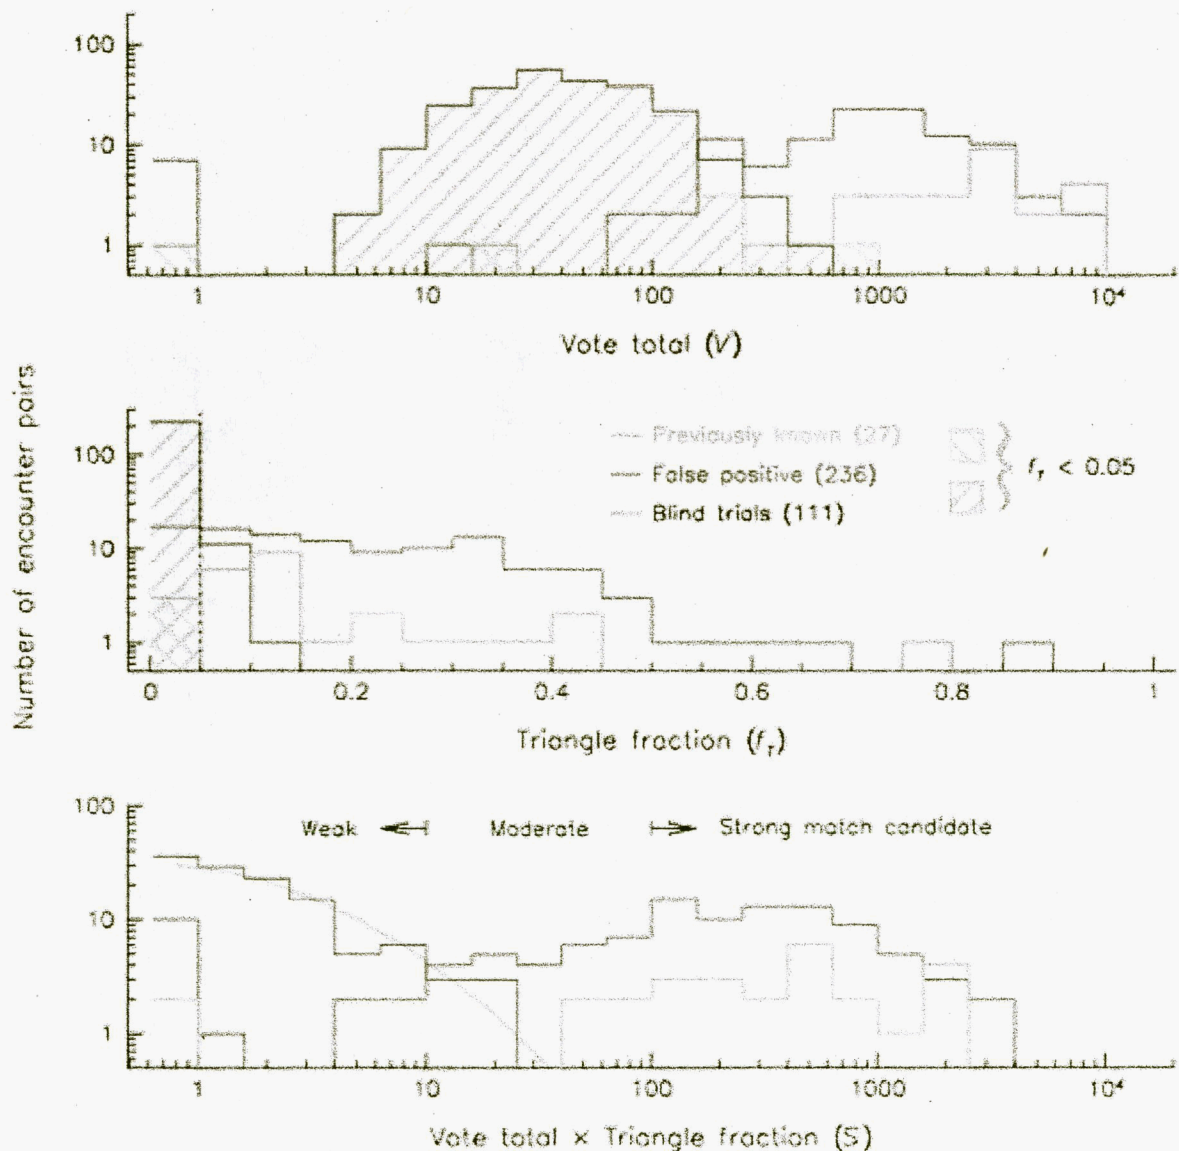
2 Figure 4: Distributions of ratio (R) and cosine at vertex 1 (C) values for triangles derived from
 3 spot coordinates of the whale shark image shown in Fig. 1. Note the logarithmic vertical axes.
 4 Filtering criteria for maximum R and C values are shown by dashed vertical lines, and hatched
 5 regions show the resulting distributions of filtered triangles not used for matching.

6




1

2 Figure 5: Distributions of magnifications for pairs of similar triangles derived from the spot
 3 coordinate lists depicted in Fig. 1. The narrow central peak for same-sense matches is evidence
 4 that portions of the two images contain the same point pattern. An iterative filter is used to isolate
 5 the matches contained in this peak.
 6




1
 2 Figure 6: Quantitative measures of match quality provided by the pattern-matching algorithm:
 3 vote total (*top panel*), fraction of triangles contributing votes (*middle panel*), and their product,
 4 our preferred ranking criterion (*bottom panel*). Right- and left-side trials have been combined.
 5 Distributions for correct (*green*) and false-positive (*red*) matches among previously identified
 6 images are shown, as well as those for new ‘blind’ matches (*black*) made by our algorithm. Trials
 7 resulting in zero votes are shown in the leftmost bin of the top and bottom panels. The mean
 8 ($\log S = -0.22$) and standard deviation ($\sigma_{\log S} = 0.61$) of the false-positive scores, in the log, are
 9 represented by a Gaussian curve in the bottom panel (*blue*). Hatched regions reflect trials in
 10 which fewer than 5% of triangles contributed to the vote total. A qualitative assessment of
 11 matches is suggested by the empirical scoring thresholds shown in the bottom panel, with weak,
 12 moderate, and strong candidates corresponding to high, medium, and low probability,
 13 respectively, of a false positive.

Encounter 231120040554
Ningaloo Marine Park, 30 March 1996




A




Encounter 15200475215
Ningaloo Marine Park, 18 April 1999




B



Encounter 1520048513
Ningaloo Marine Park, 3 May 1999



C



Comparison	Vote Total	Triangle Fraction	Product Score	Match Quality
A - B	3144	14.3%	450.3	High
B - C	1589	7.2%	115.0	High
A - C	93	0.3%	0.3	Low

1

2 Figure 7: The effects of photographic perspective on scoring of numerical spot pattern
 3 comparisons. In the sequence of images A through C, the shark's head moves progressively away
 4 from the camera, so that image A is obtained from a vantage point essentially normal to the flank,
 5 while C's perspective is oblique. Contrast-enhanced spot patterns are shown in the lower row of
 6 images. Scores for the three comparisons quoted in the table demonstrate that adjacent pairs, with
 7 small angular displacements, produce reliable matches, but the comparison of A against C fails to
 8 produce a match because of distortion.