FishSense Lite: A Camera-Based, Single Laser Hardware Framework with Software System for In Situ Fish Length Measurement

CHRISTOPHER L. CRUTCHFIELD, University of California, San Diego, USA
KYLE HU, University of California, San Diego, USA
NATHAN HUI, University of California, San Diego, USA
OMAIR QAZI, University of California, San Diego, USA
JENNIFER LOCH, Reef Environmental Education Foundation, USA
ALLI CANDELMO, Reef Environmental Education Foundation, USA
BRICE X. SEMMENS, Scripps Institution of Oceanography, USA
CURT SCHURGERS, University of California, San Diego, USA
RYAN KASTNER, University of California, San Diego, USA

Fish stock assessments need fish lengths paired with species from far more places than trained surveyors can reach. Recreational divers already visit many of those places, but measuring a fish underwater has required trained observers or specialist equipment. We present FishSense Lite, which recovers fish length from a single photograph taken with a recreational dive camera, a dive laser pointer, and a 3D-printed mount. The diver aims the laser at a fish and takes a picture; afterward, the dot's position in the image gives the range to the fish, and the range turns the fish's size in the image into a length. A citizen-science program, trained by the system's builders, taught its own divers and collected lengths paired with species for 73 wild fish across seven open-water dives. On seven wild fish also measured by an underwater stereo-imaging system, the two instruments agreed without detectable bias (mean difference +0.8%). On rigid targets of known length, FishSense Lite recovered every target to within 1.3%, and 99% of 1,055 frames fell within a 15% error budget. From the error sources we measure, we derive a field protocol: calibrate against a dive slate on every dive, photograph fish within 15° of perpendicular and beyond 1 m, and take at least ten frames of each.

Additional Key Words and Phrases: Laser triangulation rangefinding, citizen science, underwater sensing, ecology

1 Introduction

Fish are a major source of protein for more than a third of the world's population [20]. Overfishing reduces fish populations, disrupts marine ecosystems, and puts this critical food source at a significant risk [26]. Managing that risk means knowing whether a stock is replacing itself with young fish, which requires knowing how many fish there are, which species they are, and how they are distributed across age classes [30]. The length of an individual fish supports estimates of its weight and its age, and the distribution of lengths within a species supports estimates of the population's age structure [2, 30]. A photograph of a free-swimming fish carries the animal's identity, and researchers have used individual recognition across images to estimate abundance [40]. A photograph does not record the fish's scale, however, since a small fish photographed up close can fill as much of the frame as a large one farther away.

Authors' Contact Information: Christopher L. Crutchfield, ccrutchf@ucsd.edu, University of California, San Diego, La Jolla, California, USA; Kyle Hu, kshu@ucsd.edu, University of California, San Diego, La Jolla, California, USA; Nathan Hui, nthui@ucsd.edu, University of California, San Diego, La Jolla, California, USA; Omair Qazi, omqazi@ucsd.edu, University of California, San Diego, La Jolla, California, USA; Jennifer Loch, jen.loch@reef.org, Reef Environmental Education Foundation, Key Largo, Florida, USA; Alli Candelmo, alli@reef.org, Reef Environmental Education Foundation, Key Largo, Florida, USA; Brice X. Semmens, bsemmens@ucsd.edu, Scripps Institution of Oceanography, La Jolla, California, USA; Curt Schurgers, cschurgers@ucsd.edu, University of California, San Diego, La Jolla, California, USA; Ryan Kastner, kastner@ucsd.edu, University of California, San Diego, La Jolla, California, USA.

Collecting accurate lengths at scale remains difficult even for high-quality programs such as the California Collaborative Fisheries Research Program (CCFRP). Volunteer anglers catch the fish, and trained scientists on deck identify the species, measure its length, and release it [21]. Landed fish are comparatively easy to measure, and phone-based tools already do so on deck [12]. Unfortunately, CCFRP is expensive and requires trained scientists on-site. Because of high assessment costs, many important areas, such as the Indo-Pacific, remain unassessed, creating blind spots in our understanding of the world's fish populations [26]. The RAM Legacy Stock Assessment Database, which compiles stock assessments worldwide, is concentrated in a few regions and omits seven of the ten largest wild-capture fishing nations [26, 45]. Without assessments, there is no basis for setting catch limits, and no way to know whether a fishery is being fished down until it has been.

Two things follow about collecting the missing data. First, scientist-staffed programs cannot easily reach enough places. The regions missing from the assessment databases are many and far apart, and staffing and transporting trained people to each of them is the cost that makes CCFRP hard to replicate. Scaling past that means removing the dependency on the scientist and letting the volunteers CCFRP already has work independently. Second, researchers should not catch fish to measure them. Bringing an animal to the surface exposes it to barotrauma, injury caused by the rapid drop in pressure as it rises [29, 43, 48], and to air [10]. Measuring fish in situ avoids both, pointing to divers rather than anglers. Recreational divers are already in the water, including in the regions the databases miss. In the US alone, 2.5 million people are active divers and 11 million are snorkelers [16]. Volunteer fish surveys are already collected during divers' regular dives [39], so divers go where scientists already want data. A diver who is already there, and who is not asked to catch anything, reaches places a scientist cannot and leaves the fish in the water.

In programs like CCFRP, volunteers contribute labor while a trained scientist is present and owns the measurement. Underwater, volunteers can instead collect observations without a scientist present, and scientists then work from what the volunteers bring back, an arrangement we refer to as independent citizen science. The Reef Environmental Education Foundation (REEF) Volunteer Fish Survey Project already works this way, training divers to survey reef fish on their own and record each species and its abundance [1, 39], and its data support published analyses of fish abundance and distribution [1, 50, 59]. Because it is not limited by the number of scientists available, independent citizen science can grow with the number of volunteers, although someone still needs to process measurements taken from images afterward. The volunteer survey records species and abundance, but not length.

The goal of this work is an instrument that lets a recreational diver measure fish on their own. Scientists already have technologies that measure fish underwater, including underwater stereo camera systems and laser calipers. The question is whether any of them can be adapted for recreational divers, or whether a different instrument is needed. Martin et al. [34] found that marine volunteers judge a project's technical demands before joining, and that participants may drop out preemptively over anticipated difficulty. Each point at which a volunteer must buy something unusual, carry something extra, or hold expertise underwater is a point at which a volunteer may leave. We therefore set three design goals, each aimed at one of those points, for an instrument a volunteer diver uses without a scientist present.

(1) **Deferred judgment**: The volunteer is asked only to do what cannot be done later. Identifying the species and determining the length are judgments, and deferring them to people equipped to make them keeps the diver's burden to the one step that has to happen in the water.
(2) **Commodity components**: The system should use no specialized parts. Any component that must be specified, quoted, or sourced from a specialist supplier creates a stop point, so every part should be something a recreational diver already owns or can order from a consumer retailer.



Fig. 1. The FishSense Lite system adds a laser pointer to the Olympus TG-6 commercial-off-the-shelf (COTS) dive camera. A recreational diver takes a picture of a fish with the laser pointed at it. Our software calculates the fish length using only that information.

(3) **Manual-only training**: The system should require no training beyond a written manual. Training that only a scientist can deliver reintroduces the same bottleneck as having a scientist on the dive, so the system should be usable from written instructions by someone who has not used it before.

Figure 1 shows FishSense Lite, a commercial-off-the-shelf (COTS) dive camera such as the popular Olympus Tough TG-6, modified only by a laser dive pointer rigidly attached by a 3D-printed mount. The diver positions the laser point on the fish and takes a picture, and software recovers metric length from that single image. Their experience is akin to "virtual spearfishing," where each laser tag on a fish is a potential measurement of fish length. The printed mount is the only part that is not bought, and the dive slate the system calibrates against is one most divers already carry. This paper argues that commodity hardware assembled this way measures fish well enough to add volunteer length data alongside existing visual surveys, rather than replace them. We hold it to a 15% error budget per length, a design target set with the fisheries scientists we worked with, and in Section 5 account for how that budget is spent.

Our primary contributions are:

(1) **A fish measurement system built from commodity components**: A handheld device combining a COTS dive camera and a dive laser pointer with a software pipeline that recovers metric fish length from a single image, calibrated against a dive slate.
(2) **A characterization of what that system measures**: An accounting of the system's accuracy against targets of known length and against a state-of-the-art stereo-imaging system, together with the sources of error that consume its 15% budget.
(3) **Evidence that a citizen-science program can operate it without its builders:** We partnered with REEF, providing an initial training to its staff, who then trained their own volunteer divers. REEF's staff and volunteers later collected the field measurements reported here.

Section 2 reviews existing approaches, Section 3 describes FishSense Lite, Section 4 describes the evaluation, Section 5 presents its results, and Section 6 discusses what they mean for the next design.

2 Background

Existing methods of collecting fish length each meet some of the three design goals, and laser calipers come closest, meeting all three in principle.

Catch and release is a common way to obtain fish length, but it is expensive at scale and harmful to the fish. Between 2007 and 2020, the California Collaborative Fisheries Research Program (CCFRP) relied on 1,700 volunteer anglers and the crews of 36 ships to catch 176,594 fish, averaging only 282 per day trip [21]. Catching the fish is itself invasive [4], through air exposure before release [10, 11, 22] and barotrauma that can force a fish's stomach out through its mouth [29, 44, 48], and it may kill the fish the project aims to preserve [4, 43].

Underwater visual surveys avoid handling the animal entirely and can be done anywhere accessible by scuba divers. Volunteer surveys such as REEF's record species and abundance [1, 39]. Length is recorded in scientific visual surveys, where trained divers estimate it by eye [15, 24], which requires divers to identify species and estimate length underwater. Both are judgments that *deferred judgment* would move to after the dive, and the training that supports length estimation is lost without practice [5], which is what *manual-only training* is meant to avoid. These surveys work well for trained divers, and their length data depends on that training.

Instead of estimating length in the water, a diver can record the animal and have the measurement taken afterward, which is where the technologies below come in.

Underwater stereo cameras extract distance and length information from the disparity between two views [18, 19, 33]. This approach is accurate and serves as the reference against which this paper measures performance in Section 5. Commercial underwater stereo systems require purchasing proprietary hardware and software [19]. Conventional systems cost at least US$6,200 for calibration equipment, software, and cameras [15]. Low-cost rigs built from action cameras reduce the price [15], but they still require two cameras held rigidly on a base bar and synchronized at the start of each recording [15, 19], and any movement of the cameras disrupts the stereo calibration and introduces measurement error [19], so the pair must be recalibrated against a checkerboard whenever it shifts [15]. That recurring procedure is what *manual-only training* is meant to avoid, and it has kept stereo with trained scientific divers.

Commercial all-in-one depth cameras offer the same range measurement in a single device. Tueller et al. built a diver-held fish measurement system around an Intel RealSense depth camera behind a flat acrylic port [56, 61], and Wong et al. demonstrated the complexity of calibrating such a camera underwater: calibrated in air, it read depth 25.9% off, and Intel's own calibration tools could not bring both depth and length within acceptable error underwater [61]. The same system produced more data than its on-board computer could write at full frame rate [56]. Its color camera also records only 1280×800 pixels [56], which leaves less detail for the species and individual identifications that *deferred judgment* moves to after the dive. FishSense Lite avoids the flat-port problems Wong et al. observed by correcting its own refraction in hardware (Section 3.2).

Acoustic cameras (imaging sonar) recover range well and have been used to observe fish where cameras cannot [36], but at roughly US$80,000 an instrument in 2006 [36] they fall outside the *commodity components* design goal, and their returns give a shape and a distance without the species a stock assessment pairs with length.

Researchers have attempted to adapt LiDAR for underwater use multiple times, and the most successful efforts have involved adding green lasers to the infrared used on land [17, 52]. This works because green light penetrates much farther than red or infrared [42]. Even with these modifications to typical terrestrial LiDAR, their range is significantly

Technique	Barrier to a volunteer	Accuracy	Range
Stereo Imagery	two calibrated cameras; recalibrated whenever the rig shifts [15, 19]	−0.6 cm mean error [24]	2–10 m [33]
Acoustic Methods	Specialist instrument [36]	±0.5 cm at 1 m to ±6 cm at 12 m [36]	1–12 m [36]
Laser Caliper	Consumer components; lasers kept parallel [25, 51]	5–7% (< 10° off axis) [51]	0.5–5 m [25, 51]
FishSense Lite (this paper)	Consumer components and a printed mount; slate photographed each dive	−2.2% median per frame (pool); 8.9% vs. stereo (Section 5)	0.5–5 m, beyond 1 m recommended (Section 5)

Table 1. Comparison of different underwater ranging techniques. Ranges are those reported or tested in the cited work.

reduced [17]. Such devices remain niche, with only a handful of commercial products [13], which puts them out of reach for the same reason as stereo; like acoustics, they return geometry without appearance. A camera records what the fish looks like, and what it lacks is the distance these methods provide.

Other methods keep a single camera to record the animal and add a separate way to measure its distance. Recovering range from where a laser spot falls in a nearby camera image is a well-established technique [3, 35], and a monocular variant instead tracks a handheld laser's pose in each image [53, 54]. Underwater, a single laser pointer offset from a camera has given a vehicle's distance to a wall [37], and laser-line triangulation has been calibrated behind flat [9, 35] and dome [6] ports. Laser triangulation has also been applied to fish length in aquaculture tanks, with line lasers on fixed camera rigs calibrated underwater against a marked plate [32, 58]. Jaquet sized sperm whales from a boat by pairing a camera with a handheld laser rangefinder and scaling each image's pixels by the measured range [28], and Pilgrim et al. scaled ROV survey images of the seabed with five lasers on the camera: four parallel lasers set the image's perspective and scale, and a fifth, set at an angle, gave the range from where its spot fell [41]. Both measure range for a camera rather than replacing it, the same division of labor FishSense Lite uses. The underwater systems among these are mounted on vehicles or fixed rigs, whereas FishSense Lite is held by a diver in the field. An earlier version of this system appears in Hu's master's thesis [27], with initial pool tests on data separate from those reported here, and the measurement method is described in a patent application [31].

Laser calipers are the closest existing approach divers use to measure fish, and in principle they meet all three design goals. These systems position two parallel lasers at a known distance and ask the diver to put both dots on a fish, with the distance between the dots serving as a reference length [25, 46]. The camera records and the components are standard, so the diver estimates nothing. The method has supported long-term monitoring of fish size with minimal training [25, 51]. Its accuracy rests on the two lasers staying parallel at their known separation, so practitioners check them, though how often varies: once on land [25], regularly on land [47], underwater against a reference before measuring [14], before each dive [51], or at the start and end of each encounter, with the encounter's data discarded when the spacing has drifted by the end [62]. The method therefore depends on keeping the lasers parallel and on landing both dots on the fish, which makes the laser baseline the shortest fish it can measure.

Table 1 collects these techniques. Among the methods divers use to measure fish, FishSense Lite is closest to the laser caliper and differs by using a single laser whose pose is measured from photographs of a slate rather than set by alignment, so only one dot needs to land on the fish, and no laser spacing sets the shortest fish it can measure; half the field fish of Section 5.6 were shorter than the 25 cm spacing of one caliper [25]. Unlike the single-laser and

Component	Source	Typically owned?	Share of cost
Dive camera	Consumer retailer	Yes, by underwater photographers	45%
Underwater housing	Consumer retailer	Yes, with the camera	31%
Wide-angle wet lens	Consumer retailer	Sometimes	16%
Dive laser pointer	Dive retailer	Sometimes	7%
Laser mount	3D printed from files released with the paper	No, printed	1%

Table 2. Components of the FishSense Lite system. Every part is available through consumer channels, and none requires a quote or a specialist supplier. The mount is the only part that must be made rather than bought. Shares are of October 2026 retail prices for an equivalent unit, about US$1,300.

fish-length systems above, which are mounted on vehicles or fixed rigs and calibrated before deployment, FishSense Lite is diver-held, recalibrated on every dive in the field from a slate at an unknown pose, and evaluated in a citizen-science deployment.

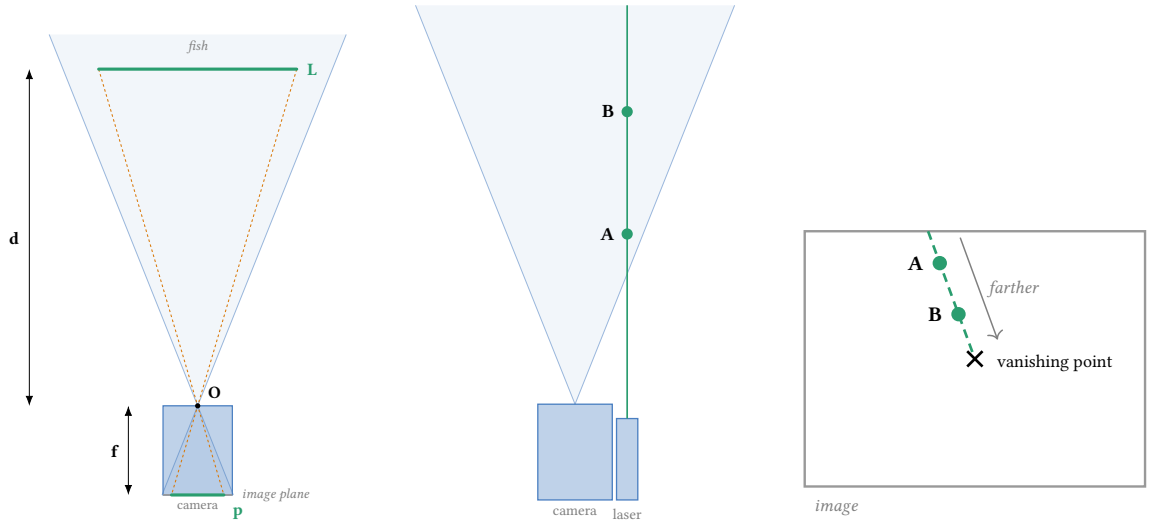
3 FishSense Lite

FishSense Lite is a handheld device built from COTS components, paired with software that recovers metric fish length from a single photograph. The diver photographs the fish, and the system computes the length afterward. This section describes the ideas that make a single-image measurement possible, the device that implements them (Section 3.2), the geometry that recovers a length (Section 3.3), the calibration that geometry depends on (Section 3.4), and the assumptions the whole chain rests on (Section 3.5).

3.1 Core design ideas

Given a photograph of a fish, we want its length. A camera answers part of this immediately, since the fish covers some number of pixels in the image, which we call its *pixel length*. Pixel length alone does not give a metric length, because a small fish close to the camera can cover the same number of pixels as a large fish farther away, so a given pixel length is consistent with many different fish. Knowing the distance to the fish resolves this ambiguity. As illustrated in Figure 2 (a), a fish of true length L at distance d from the camera produces a pixel length p, and by similar triangles $p/f = L/d$, where f is the camera's focal length in pixels. Rearranged, $L = p \cdot d/f$, a first-order relation that Section 3.3 replaces with an exact back-projection. The image gives p, and the camera calibration gives f, leaving only d. For now we take the fish to lie broadside to the camera, an assumption Section 3.5 lists and Section 5.4 tests.

FishSense Lite recovers d with a laser pointer mounted rigidly beside the camera (Figure 2 (b)). When the laser strikes an object, the dot appears somewhere in the image. Move that object farther away, and the dot moves across the frame, because the camera sees the beam from the side. The dot's position therefore indicates how far away the surface is. Figure 2 (c) shows the path it traces. As distance grows, the dot converges toward a fixed point in the image, the vanishing point of the beam, and closer distances put the dot farther from that point. This is *laser triangulation*, which has been used to measure the range from an ROV's camera to the seabed alongside parallel lasers that set the image scale [41]. In FishSense Lite, a single laser provides both range and scale, so no parallel lasers are needed; one photograph of the fish contains everything the measurement needs. We measure the laser's position and direction from photographs of a dive slate during each dive (Section 3.4). Figure 3 shows a frame of a fish: the dot gives the range, and the two endpoints define the pixel length.



(a) Pinhole camera (b) Camera with laser (c) What the camera sees

Fig. 2. The principle of laser triangulation underlying FishSense Lite. **(a)** A pinhole camera projects a fish of length L at distance d to an inverted image of length p. **(b)** A laser mounted offset from the camera; points A and B are at different distances along the laser ray. **(c)** The same points in the image, sliding toward the beam's vanishing point as distance grows.

3.2 Device and Workflow

FishSense Lite consists of a COTS underwater camera with a waterproof laser pointer attached by a custom 3D-printed mount, as shown in Figure 1. Our reference implementation uses an Olympus Tough TG-6 in its standard PT-059 underwater housing, with a Backscatter M52 Underwater 81° Wide Air Lens mounted at the housing's port. Behind a flat port, light bends where it passes from the water into the air inside the housing, so the camera is not strictly a single-viewpoint pinhole underwater, and a pinhole model approximates it well only at a constant range [55]. Divers measure fish across a range of distances, so that approximation does not hold for FishSense Lite. The M52 restores the pinhole behavior assumed by the model in Section 3.3 at every range, without refraction-specific calibration, while keeping the camera's topside field of view. The laser is a Class IIIA waterproof dive pointer, limited to 5 mW of visible output [57]. Divers used both red and green pointers, and because the geometry below depends only on where the dot lands in the image, it is the same for either.

Table 2 lists the system components. The mount clamps the laser to the housing's cold shoe, offset from the camera's optical axis, through a printed taper, body, and cap. The mount is the only part that cannot be bought off the shelf, but it requires no specialist manufacturing. We will release the design files on publication, along with the analysis code. A diver can print the mount in polylactic acid (PLA) on a 3D printer, or order it from an online print service like any other part, so the mount stays within the *commodity components* design goal. An equivalent unit costs about US$1,300 at October 2026 retail prices, three-quarters of it in the camera and housing an underwater photographer already owns (Table 2).

To measure a fish's length, the diver aims the laser at the fish's body and holds the shutter button for a short burst of pictures, recorded as RAW stills rather than video for the reason given in Section 3.3. A usable frame has the laser dot on the fish and both the snout and the tail fork visible. The laser dot also serves as a framing aid: because the beam is

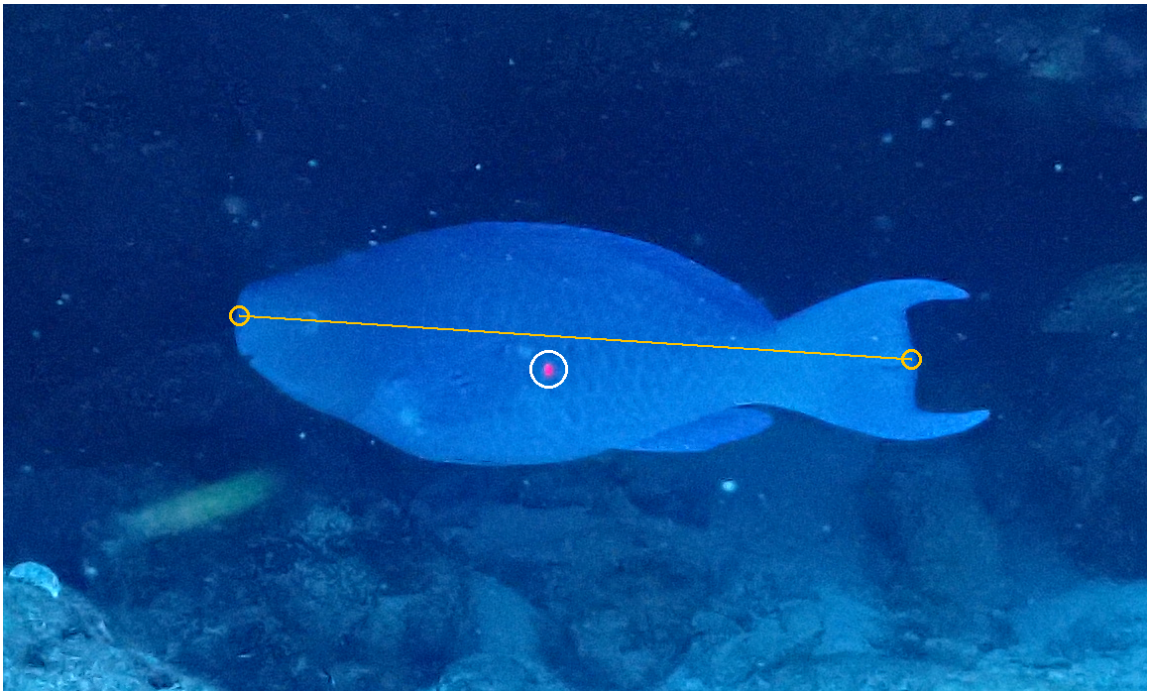


Fig. 3. A FishSense Lite frame. The white ring marks the laser dot; its position in the image gives the fish's range. The amber marks are the snout and tail fork as the labeler placed them, and the length is the distance between those two points back-projected to the dot's depth.

nearly parallel to the optical axis, keeping the dot on the fish keeps the fish near the center of the image. Section 4.1 provides the complete set of instructions divers follow.

The system needs two calibrations. The camera intrinsics are fitted once per camera, underwater against a checkerboard, and do not change. The laser extrinsics are fitted for every dive from photographs of a dive slate the diver already carries, and Section 3.4 describes the procedure. After the dive, images are transferred to a computer, where human annotators label the laser dot, snout, and tail-fork pixel locations. FishSense Lite's software pipeline then extracts a metric length for each pictured fish. This division follows the *deferred judgment* design goal, leaving the diver only the photography that has to happen in the water. Automating the labeling step is a data analysis problem rather than a system one, and it is out of scope for this paper.

3.3 Length Recovery

We define *fish length* as fork length: the distance from the snout to the tail fork. This section turns the snout and fork pixels a labeler marks into that physical length. The similar triangles in Section 3.1 assumed a known distance and a fish centered in the image; here, the distance comes from the laser, and the endpoints may lie anywhere in the frame. Recovery requires the laser's position and direction in the camera frame, which are not known in advance and are recovered by the calibration procedure in Section 3.4.

We work in a coordinate system with the origin at the camera's focal point and the z-axis aligned with the optical axis. We model the camera with the basic pinhole camera model [23], in which a 3D point $\mathbf{P}$ expressed in this frame

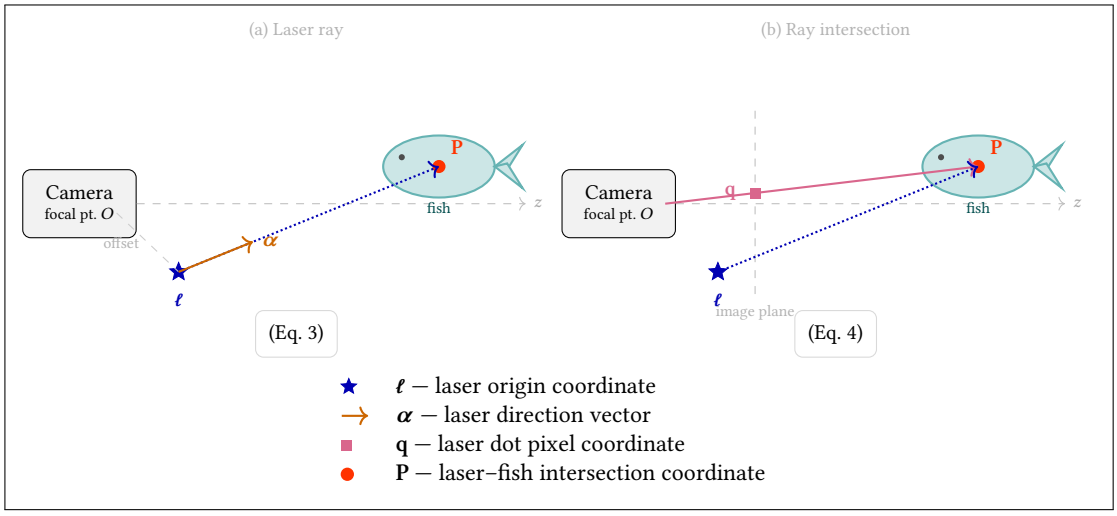


Fig. 4. Laser ray model and ray intersection. **(a)** The laser, with origin $\boldsymbol{\ell}$ and direction $\boldsymbol{\alpha}$, intersects the fish at $\mathbf{P}$. **(b)** The camera ray and laser ray share the terminus $\mathbf{P}$, forming an overdetermined system in λ_1 and λ_2.

projects to a pixel location $\mathbf{q}$. A pixel location fixes a direction from the camera and not a distance, so the projection holds only up to an unknown positive scale. Writing that scale explicitly as λ_1,

$$\lambda_1 \mathbf{q} = K\mathbf{P}, \tag{1}$$

where K, the camera intrinsic matrix, contains the focal lengths f_x and f_y in pixels along each image axis and the principal point (c_x, c_y), the pixel where the optical axis meets the image. Standard intrinsic calibration [8, 63] recovers it once per camera, underwater, using a checkerboard. We fit that calibration to the camera's RAW output, and we use RAW for the measurement for the same reason. The camera's JPEG and video modes apply factory processing that K does not describe, so a length recovered from either would be using a calibration for a different optical path. Rearranging Equation 1 maps each pixel back to its original 3D coordinate,

$$\mathbf{P} = \lambda_1 K^{-1}\mathbf{q}. \tag{2}$$

In an observed image, however, λ_1 is an unknown scale factor that must be recovered, so a pixel alone fixes only a ray through the camera center along which the original 3D point lies.

The laser supplies a second equation for the same point. We define the laser origin as $\boldsymbol{\ell}$ and the laser direction as $\boldsymbol{\alpha}$, the laser's two extrinsic parameters in the camera frame, both illustrated in Figure 4. The point at which the laser intersects the fish is

$$\mathbf{P} = \lambda_2 \boldsymbol{\alpha} + \boldsymbol{\ell}, \tag{3}$$

where λ_2 parameterizes distance along the laser ray. The dot we see in the image is where these two rays meet, so setting Equation 2 equal to Equation 3 gives

$$\lambda_1 K^{-1}\mathbf{q} = \lambda_2 \boldsymbol{\alpha} + \boldsymbol{\ell}. \tag{4}$$

This system is overdetermined, with three scalar equations in two scalar unknowns, so we solve for λ_1 by linear least squares. Let $\mathbf{r} = K^{-1}\mathbf{q}$ denote the back-projected camera ray. Projecting the residual of Equation 4 onto $\mathbf{r}$ and onto $\boldsymbol{\alpha}$

and eliminating λ_2 yields

$$\lambda_1 = \frac{\boldsymbol{\ell}^\top \mathbf{r} - (\boldsymbol{\alpha}^\top \boldsymbol{\ell})(\boldsymbol{\alpha}^\top \mathbf{r})}{\|\mathbf{r}\|^2 - (\boldsymbol{\alpha}^\top \mathbf{r})^2}. \tag{5}$$

The denominator vanishes when $\boldsymbol{\alpha}$ and $\mathbf{r}$ are parallel, which the camera ray approaches as the dot recedes toward the beam's vanishing point, so the scatter in range from imprecisely placed labels, which we call *triangulation noise*, grows with range (Section 5.5). The baseline between the laser and the camera sets how quickly that happens: a wider baseline keeps the two rays further from parallel at a given distance, and so extends the range over which depth is well determined. The reconstructed 3D point is $\mathbf{P} = \lambda_1 \mathbf{r}$, and since $r_z = 1$ we have $P_z = \lambda_1$ directly. We use λ_1 as the fish's metric depth.

Unlike the similar-triangles shortcut of Section 3.1, which is exact only at the image center and requires $f_x = f_y$, a condition that is not guaranteed [8, 23], we back-project the snout and tail fork pixel locations through K^{-1} and place each at the depth the laser gave us:

$$\mathbf{P}_{\text{snout}} = P_z \cdot K^{-1} \mathbf{q}_{\text{snout}} \tag{6}$$

$$\mathbf{P}_{\text{tail}} = P_z \cdot K^{-1} \mathbf{q}_{\text{tail}} \tag{7}$$

The metric fish length is the Euclidean distance between the back-projected points:

$$L = \|\mathbf{P}_{\text{snout}} - \mathbf{P}_{\text{tail}}\|. \tag{8}$$

Handling the endpoints this way keeps the measurement exact for a fish anywhere in the frame and for sensors whose pixels are not square, which the similar-triangles shortcut does not. It recovers the fish's length projected onto a plane at the laser's depth. When the fish lies broadside to the camera, that projection is the fish's length. When it does not, the projected length is shorter, an effect called *foreshortening*, which Section 5.4 measures as its own error term. The laser's depth is itself that of the fish's near flank, closer to the camera than the line from snout to fork, so a thick fish also reads short, which Section 5.2 measures as *body depth*.

From frames to a length. Length comes from several frames of the same fish rather than one, because pose can only shorten a frame's measurement. A fish angled away from the camera, or bent mid-stroke, measures short, so the frame taken closest to broadside gives the longest measurement. Other errors, such as a misplaced endpoint or noise in the dot's position, can push a single frame either way, so we report the p_{90} of an individual's measurements rather than their maximum. Under the nearest-rank definition we use, an individual with fewer than ten frames is assigned its longest frame, which leaves such an estimate exposed to a single frame that reads long; in the pool, estimates from two to eight frames instead read slightly short on average, by 1.4% at two frames (Section 5.2). Section 5.2 also shows that the p_{75} and p_{90} bracket the known length, while the median reads short and the maximum long.

3.4 Calibration

Every length depends on where the laser points, and on FishSense Lite, which changes from one dive to the next. Equation 5 takes the laser's origin $\boldsymbol{\ell}$ and direction $\boldsymbol{\alpha}$ as given, so error in either becomes error in depth, and through depth, error in length. The mount is printed plastic registered to a cold shoe, and the beam it holds shifts between dives by a median of 0.65°. Measuring one dive's fish with another dive's calibration raises the median length error from 3.1% to 25.2% (Section 5.3). The laser's pose is therefore measured anew on every dive.

The calibration recovers the beam's pose from images of the dot landing on a flat object of known geometry, and it assumes only that the object is flat and its pattern known, not that it is held square to the camera. Any planar target

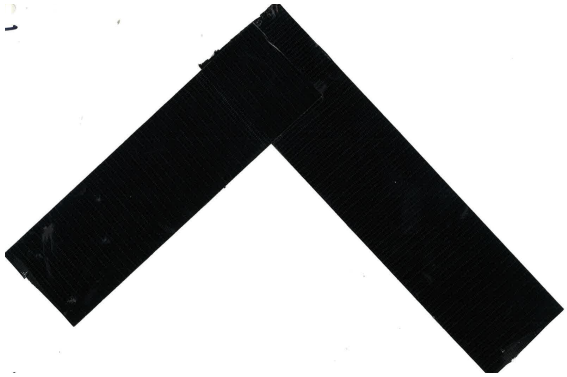


Fig. 5. Field calibration target: a standard underwater dive slate with a chevron of black duct tape.

with labeled features at known spacing will do. A printed camera-calibration checkerboard is the standard choice, and our earliest pool sessions used one, but in the field we use a standard dive slate marked with duct tape, shown in Figure 5, because a slate is already in the diver's kit and asks nothing extra of them. At the start or end of a dive, a buddy holds the slate while the distance between the slate and the camera is swept outward until the laser becomes hard to keep on the slate, then back in, and the diver photographs the slate whenever the dot is on it. The procedure follows the slate check used with some laser calipers [51]. Divers aim for about ten frames, more than the fit needs, because some will be unusable. Sweeping in range is what fixes the laser's direction, and Section 5.3 tests how much range the sweep needs.

For each calibration image, passing the labeled slate corners and the pattern's known metric geometry to OpenCV's `solvePnP` routine [38], using its default iterative solver, recovers a rotation R_i and translation $\mathbf{t}_i$ mapping the slate's body frame to the camera frame. The slate plane Π_i follows, with normal $\mathbf{n}_i$ the rotated slate body-frame z-axis, and any slate corner transformed to the camera frame, $\mathbf{p}_i$, giving a point on the plane. Intersecting the camera ray through the labeled laser dot pixel $\mathbf{q}_i$ with Π_i recovers the laser dot's 3D position:

$$\mathbf{P}_i = \frac{\mathbf{n}_i^\top \mathbf{p}_i}{\mathbf{n}_i^\top K^{-1}\mathbf{q}_i} \cdot K^{-1}\mathbf{q}_i. \tag{9}$$

This produces observed 3D points $\{\mathbf{P}_i\}_{i=1}^{n}$, each lying, up to labeling noise, on the laser ray. The laser direction $\boldsymbol{\alpha}$ is the normalized sum of sign-corrected pairwise differences:

$$\boldsymbol{\alpha} = \frac{\mathbf{d}}{\|\mathbf{d}\|}, \qquad \mathbf{d} = \sum_{i \neq j} \operatorname{sgn}\big((P_i - P_j)_z\big)\,(\mathbf{P}_i - \mathbf{P}_j), \tag{10}$$

and the laser origin $\boldsymbol{\ell}$ is the point at which the fitted ray intersects the $z = 0$ plane:

$$\boldsymbol{\ell} = \bar{\mathbf{P}} - \frac{\bar{P}_z}{\alpha_z}\,\boldsymbol{\alpha}, \tag{11}$$

where $\bar{\mathbf{P}}$ is the mean of the $\mathbf{P}_i$. Equation 10 is a least-squares fit of the beam's lateral position against the rank of each point's range, so widely spaced points, which best fix the direction, count most. When the frames are evenly spaced in range, it equals regressing lateral position on depth. It differs from a total-least-squares fit only by a term that grows with the beam's tilt from the optical axis, and on the field calibrations the two directions agree to within 0.03°. These two parameters complete Equation 5, and together they yield a length from a single labeled frame.

3.5 Assumptions

The chain from photograph to length rests on a small number of assumptions, which fall into two kinds.

Two are preconditions, in that the measurement does not exist without them. The laser must produce a single point, and that point must land on the fish. A dot on the reef behind the animal returns the reef's depth instead of the fish's.

The rest are approximations, in that the measurement degrades in a way we can characterize. We treat the camera as a pinhole, which the corrective optic in Section 3.2 makes valid underwater at every range. We treat the fish as flat, perpendicular to the camera axis, and rigid, so both endpoints lie at the laser's depth and neither moves between them. The laser's pose is treated as fixed for the duration of a dive. Each holds only approximately, and Section 5 tests them: the fixed laser pose in Section 5.3, the perpendicular fish in Section 5.4, the flat fish in the body-depth comparison of Section 5.2, and the pinhole model end-to-end through the accuracy on targets of known length in the same section.

4 Evaluation Setup

The evaluation draws on three cohorts that differ in what they can be scored against, all collected with the same instrument and instructions.

4.1 What the diver is asked to do

Divers learned the system from REEF rather than from the group that built it. We trained the program's staff in a series of meetings using a written manual covering assembly, camera settings, the field procedure, and guidelines for taking usable photographs; the staff then trained their divers in an in-person briefing and ran the deployments.

The guidelines reduce to six rules, and each protects something the measurement depends on.

(1) **Point the laser at the fish, as close to its center of body mass as possible.** A dot on the reef behind the animal returns the reef's depth instead of the fish's.
(2) **Do not use the zoom.** Zooming changes the focal length, and K was fitted with no zoom.
(3) **Do not touch or move the laser, and do not use it as a handle.** The calibration holds only while $\boldsymbol{\ell}$ and $\boldsymbol{\alpha}$ hold.
(4) **Hold the camera's optical axis at approximately a right angle to the fish.** This is the perpendicularity assumption, and Section 5 reports what deviation costs.
(5) **Keep the fish between 0.5 and 5 meters away, ideally around 3 meters.** Section 5 revises the near end of this range.
(6) **Make sure the fish is illuminated.** The annotator has to find the snout and the tail fork.

The manual also asks divers to photograph the slate at the start and end of each dive, following the procedure of Section 3.4. None of the rules requires a certification.

4.2 The deployment

We provided FishSense Lite units to REEF, whose staff and volunteer divers had not used the instrument before. Between August 2023 and January 2025, they collected nearly 70,000 images, and this paper's results rest on nearly 2,700 of those that became length measurements, since producing a measurement requires a person to label each frame.

This paper evaluates the instrument on three cohorts, summarized in Table 3. They differ in what they can be scored against. The paired stereo dives have a second instrument measuring the same animals, the pool cohort has targets of known length, and the field cohort has neither.

Cohort	Frames	Targets	Reference
Paired stereo dives	39	7 wild fish	stereo-imaging system
Pool	1,055	5 rigid models	tape or measuring board
Field	162	73 wild fish	none

Table 3. The three cohorts this paper evaluates. The angle experiment's 1,428 frames are reported separately (Section 5.4).

Target	Length	Profile
Box	150 mm	flat
Purple Angelfish	192 mm	flat
Rainbow Trout	313 mm	thick
Grouper	360 mm	flat
Snook	455 mm	thick

Table 4. Rigid targets used in the pool. The flat targets match the planar assumption of Section 3.5, while the thick targets carry the body depth of a real fish.

4.3 Paired stereo dives

The paired stereo dives test the system on live fish in open water, the condition it is built for, against a second instrument that measured the same fish at the same time. Stereo imagery measures length with a mean error of −0.6 cm on targets of known length [24], so it serves as the reference here, though not as ground truth. The stereo-imaging system was a diver-operated SeaGIS stereo camera system built on two GoPro cameras, calibrated against a SeaGIS calibration cube with the CAL software at least once per field trip, and its lengths were measured in EventMeasure [49]. A FishSense Lite unit was carried alongside the stereo-imaging system on the same dives, and the two instruments captured seven fish in common, matched by species, markings, background, and the time each recorded them.

4.4 Pool cohort

The pool cohort establishes what the instrument measures when the answer is known in advance, which is what lets its error sources be separated and measured one at a time. Over several days in 2023, the engineering team and REEF's staff photographed five rigid targets in three pools, using several cameras, housings, and laser units. Table 4 lists the targets. Four are rigid plastic fish models, and one is a box; we measured each once with a tape or measuring board to determine the length the system is scored against. Two of the fish models are flat, and two have the body depth of a real fish, which Section 5.2 uses to measure body-depth cost.

A diver held the targets rather than fixing them to a rig, so pose varied at the holder's discretion. The targets were photographed from 0.25 to 5.5 m away, with a median of 2 m, covering the manual's 0.5 to 5 m working range and extending beyond it at both ends. The lasers were calibrated against a checkerboard in the earlier sessions and against the dive slate in the later ones. Some pool sessions were measured with a calibration taken earlier, as the protocol then allowed. The cohort comprises the sessions whose calibration was judged to fit by checking that a rigid object reads the same length at every range, a test that needs only a fixed size, not a known one, and fails when the 95% interval of the slope of length against range lies entirely beyond ±2% per meter; the calibrations that failed it no longer described the laser, which is why calibration moved to every dive. Section 5.4 reports the five sessions of the angle experiment separately.

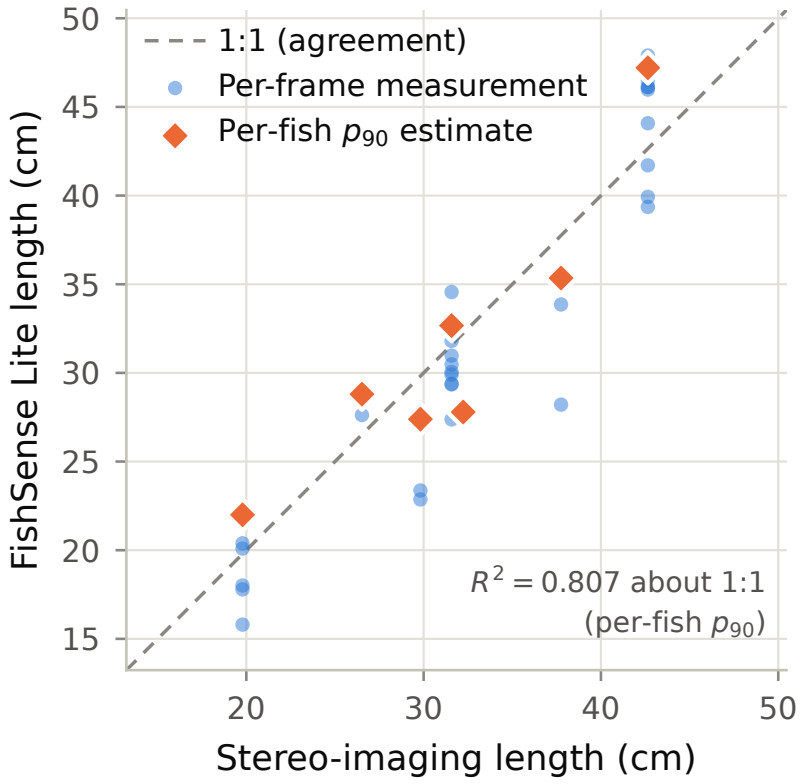


Fig. 6. FishSense Lite against an underwater stereo-imaging system on the same seven individuals. Each orange diamond is one fish's FishSense Lite length, the p_{90} of its frames; the blue points are those frames, 39 in all, between 3 and 11 per fish. The dashed line marks agreement, not truth; neither axis is a known length.

4.5 Field cohort

The field cohort is a realistic deployment of FishSense Lite. Seven open-water dives in the Florida Keys, between November 2023 and December 2024, yielded 162 measurements of 73 wild fish. This cohort has no reference instrument and no known length, so it shows how consistently the system measures wild fish in use, not how accurately it measures them. Of the seven dives, three were calibrated at the start, two at the end, and two at both. Every field calibration spanned at least 0.9 m of range, above the 0.6 m Section 5.3 requires.

5 Results

5.1 Comparison with a stereo-imaging system

We first test FishSense Lite against stereo imagery on an open-water dive, where the two instruments, pointed at the same wild fish, show no detectable bias across the seven fish compared. We measured the wild fish ranging from 20 to 43 cm with both instruments on one day (Figure 6), using each FishSense Lite length as the p_{90} of that fish's frames. The estimates fall on the $1:1$ line with $R^2 = 0.807$ and a residual RMS of 3.0 cm.

Species	Dive	Frames	Difference from stereo
Hogfish	1	11	+3.5%
Hogfish	2	3	+8.6%
Hogfish	1	3	−13.8%
Blue Parrotfish	1	10	+10.7%
Stoplight Parrotfish	1	6	+11.2%
Gray Snapper	2	3	−8.1%
Blue Parrotfish	2	3	−6.4%

Table 5. The seven fish measured by both instruments, from two dives under one morning calibration. Each difference is the FishSense Lite p_{90} minus the stereo length, as a percentage of the stereo length. Within each dive the differences span most of the full range (dive 1: −13.8 to +11.2%; dive 2: −8.1 to +8.6%), while the two dive means differ by 4.9 points.

Individual fish differ by up to 14% (Table 5). The mean difference is +0.8% (95% confidence interval −8.6 to +10.2%), the mean absolute difference is 8.9%, and the 95% limits of agreement [7] run from −19.1% to +20.7%. With seven fish, those limits are themselves uncertain by about ±16 points (95% confidence intervals −35.3 to −2.8% and +4.4 to +36.9%). Each of the seven fish lies within the 15% budget, but seven fish are too few to say how far a single measurement can stray, so we infer per-measurement accuracy from known-length targets (Section 5.2). The difference also shows no detectable trend with fish size (95% confidence interval −0.15 to +0.16% per millimeter) or with the number of frames per fish (Spearman $\rho = 0.57$, $p = 0.18$).

Part of the difference between the two instruments comes from sources the pool cannot reproduce. Four fish read long, by up to 11.2%, and pose only ever shortens a length. The difference is not mainly frame-to-frame noise, since the stereo fish's own frames vary by a median of 6.9%, and that scatter accounts for under half of the variance even if none of it is pose; each fish carries an offset between the two instruments. Nor is it calibration, since the seven fish came from two dives under one morning calibration, and a calibration error shifts every fish on a dive alike, while the spread lies among fish on the same dive rather than between the two dives. Likely contributors include the stereo system's own error, endpoints placed on a live fish that flexes between frames, and poses beyond 15°.

5.2 Comparison with known lengths

Targets of known length let the instrument's own error be measured, rather than compared with another instrument's. Against them, FishSense Lite recovers each target to within 1.3% at the p_{90}, and 99% of individual frames fall inside the 15% budget. The per-target p_{90} is +0.65% for the Box, +0.47% for the Purple Angelfish, −1.00% for the Rainbow Trout, +0.59% for the Grouper, and −1.24% for the Snook, with no trend across a three-fold length range (Figure 7). A single frame reads a median of 2.2% short, and 80% of frames fall within 5% of the known length and 98% within 10%.

The error is one-sided, so we report a high quantile rather than a mean (Figure 8). Each target's distribution has a long tail toward short readings and a short one toward long, as foreshortening requires, so a mean describes how divers held the targets rather than what the instrument measures. The nine frames outside the budget all read short, and all come from two targets. Their shortfall matches the foreshortening of a 32° to 46° pose, a cost Section 5.4 measures against a known angle.

The estimate needs few frames, and its accuracy holds at the frame counts divers produce in the field. Drawn at random from the 15 sessions in which the Box or the Rainbow Trout was photographed at least 30 times, the p_{90} of four frames lands within 2% of the all-frames value 90% of the time, of ten frames within 1.5%, and of thirteen within 1%

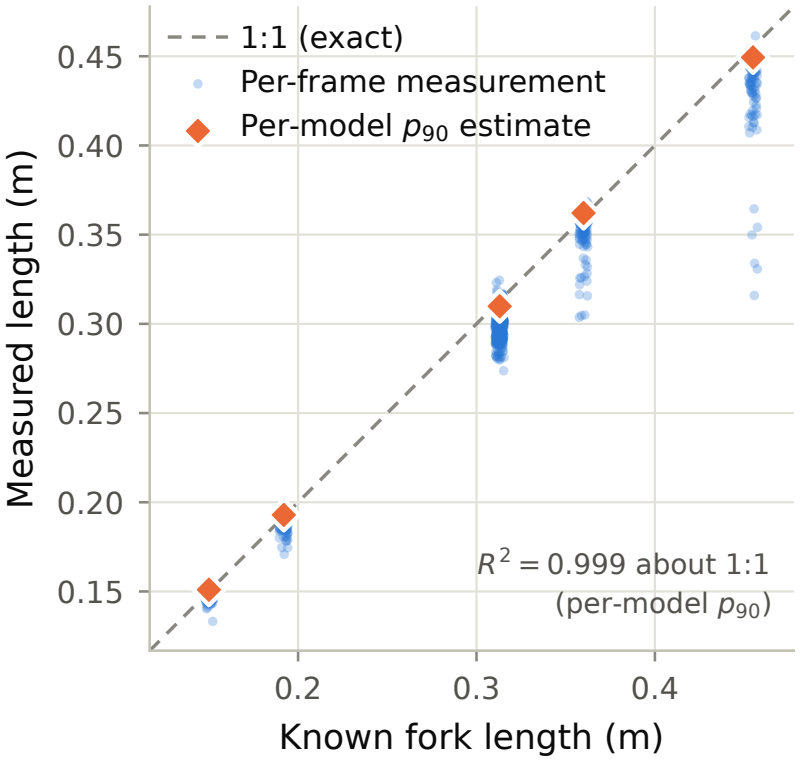


Fig. 7. Measured against known length over the pool cohort, 1,055 frames of five rigid targets. Diamonds are the per-target p_{90}. $R^2 = 0.999$ about the $1:1$ line on those five estimates, and 0.976 on the individual frames.

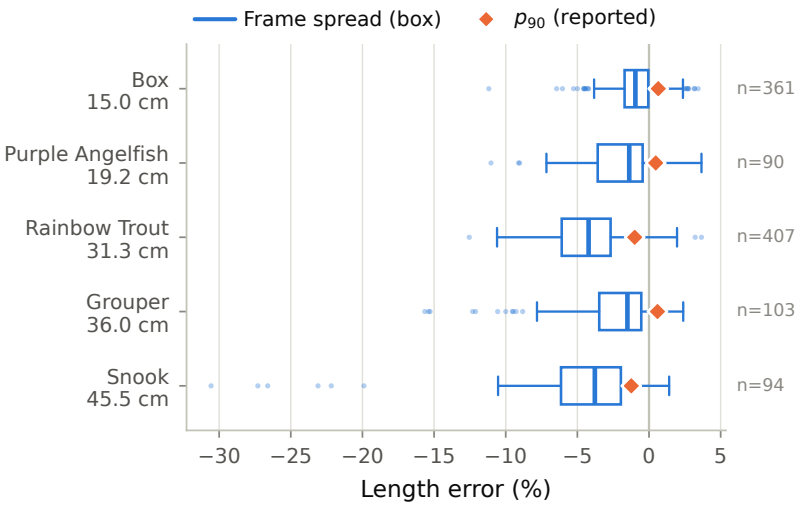


Fig. 8. Percent length error by target. Box and whiskers give the frame-level spread, the diamond the reported p_{90}. All 1,055 frames are shown, out to the worst at −30.6%.

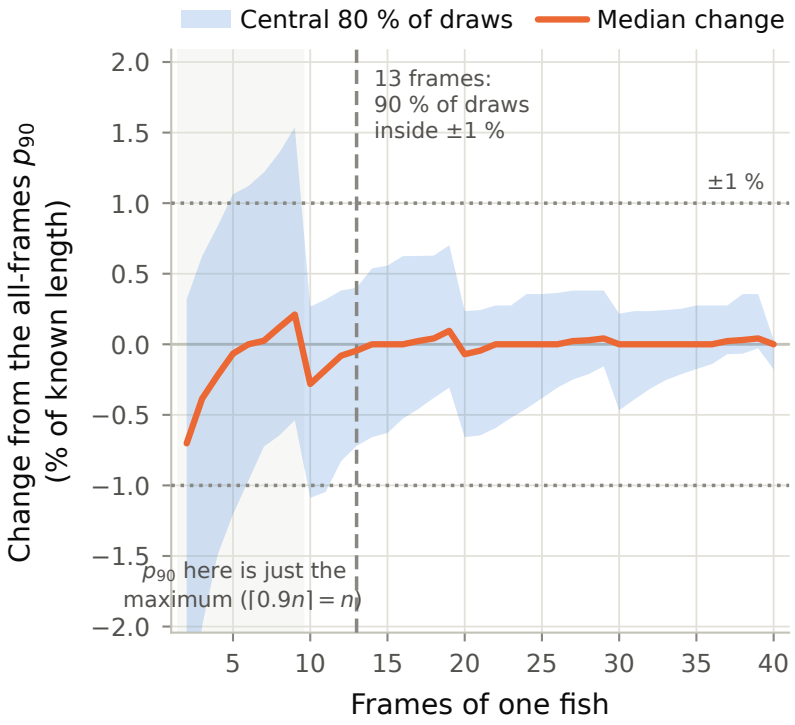


Fig. 9. How far a pool target's p_{90} moves when estimated from n randomly drawn frames rather than all its frames, across the 15 sessions in which the Box or the Rainbow Trout was photographed at least 30 times. The shaded band holds 80% of draws. Below ten frames, the p_{90} is the longest frame.

(Figure 9). Scored against the known length instead, the p_{90} for two frames is 1.4% short, and for eight frames within 0.1%, since fewer than ten frames can only raise it. The p_{90} is also a balanced choice of quantile: across the five targets the median reads 1.5% short and the maximum 3.4% long, while the p_{75} and p_{90} bracket the known length. We therefore recommend photographing each fish at least ten times, though fewer frames still give a usable length.

Thick models read slightly shorter than flat ones, and body depth (Section 3.3) explains part of the difference. Under each of the seven checkerboard calibrations shared by the trout and the box, the trout read 3.0 points lower on average (95% confidence interval 1.4 to 4.5). In each of the five sessions that photographed the snook alongside the angelfish and the grouper, the snook read 1.8 points lower on average (95% confidence interval 0.8 to 2.8). For the trout, half its 58.7 mm body depth predicts a 1.7% shortfall at its median range of 1.7 m, about half its observed gap.

5.3 Laser calibration stability

Every length assumes the laser calibration is current, but a calibration records the laser's position and direction only when it was taken. If the laser moves afterward, the measurement frames reveal some of that movement and hide the rest. The camera's center and the laser beam together define a plane, which the camera views edge-on, the way a sheet of glass seen edge-on appears as a single line, so the dots the beam makes all fall along one line in the image. A beam that tilts out of the plane moves the dots off that line, and the measurement frames show it. A beam that rotates within

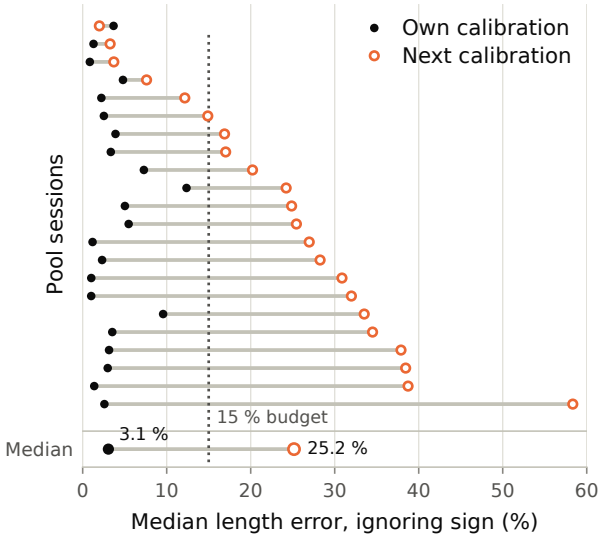


Fig. 10. Each row is one pool session measured with its own calibration and, on the same frames, with the same unit's next calibration. A neighboring calibration moves 21 of 22 sessions further from their targets' known lengths and 16 of them past the 15% budget, by an amount that tracks how far the beam moved between the two calibrations.

the plane keeps the dots on the line. The reprojection error, the distance between where the calibration predicts each dot and where it appears, is therefore unchanged. Each position along the line now corresponds to a different range, though, so every length is rescaled. Only a calibration can see this rotation, because the slate's known geometry gives each dot a range independent of the laser. Recalibrating on every dive is therefore worth its cost if the laser rotates between dives by enough to matter, which the rest of this section measures.

Between dives, the beam rotates within the plane by more than the error budget allows. At 2 m, a rotation of about half a degree (−0.53° to +0.39°) puts a length outside the 15% budget; between successive calibrations of a unit, the median rotation is 0.65°, exceeding that tolerance in 13 of 18 cases. Between the pool tests and the 2024 deployments, three of the four units we can check rotated by about 3°. The laser moves far less within a dive, where no new calibration is taken. On two field dives, the slate was photographed at both the start and the end of the dive. The calibration from one end matched the ranges measured at the other to within 3.7%, well inside the 15% budget.

Substituting one session's calibration for another's confirms that the laser moved, not that the fits were noisy. If a calibration still described the laser at the next session, using it there would leave the error unchanged. Instead, re-measuring each pool session with the same unit's next calibration raises the median error from 3.1% to 25.2% and puts 16 of 22 sessions outside the budget (Figure 10). A borrowed calibration costs what the beam moved between the two calibrations (Spearman $\rho = 0.83$ between a session's error and that rotation), so calibration has to happen on every dive.

The change between calibrations comes from the laser, not from the calibration target. A unit calibrated against both should report the same baseline, the laser's offset from the camera's optical center, and across the six units calibrated against both the checkerboard and the slate, the baselines differ by +0.66% (95% confidence interval −0.23% to +1.69%), within the variation between repeat calibrations against one target.

The movement comes from the joints in the mount, which stacks a cold shoe, a taper, the mount, a cap, and the pointer (Section 3.2). Flex in the printed mount tilts the beam up or down, roughly within the plane because the laser sits above the lens; a twist at the cold shoe or the taper's screwed joint turns it sideways, mostly out of the plane; and the pointer turning in its clamp moves the beam around a small cone, because the laser diode is not exactly on the pointer's axis. Between calibrations, the beam moves both ways but more in-plane than out of plane, with a median of 0.65° in-plane vs. 0.35° out-of-plane, and the in-plane component is larger in 11 of 18 steps, though a single step can be almost entirely one or the other. Over deployments, the printed parts also crept, cracked, and absorbed water; screws pulled through them; and a published laser-caliper mount has likewise been reported to shift with temperature [60]. An alignment mark on the pointer only detects the pointer turning in its clamp, so recalibrating at the dive is the only step that accounts for every joint.

The largest steps within a session came in the pool, across pauses in recording while the camera was out of use. In one, the beam rotated 0.85° within the plane between two consecutive calibrations. In the other, the beam rotated 0.29° partway through a session, after which that session's trout read 9.6% short instead of 4.4%. A field diver, by contrast, holds one camera throughout the dive, without such pauses, and the two field dives calibrated at both ends held to within 3.7% from start to finish, so on the dives we could check, one calibration per dive was enough.

A calibration also needs slate photographs taken across a range of distances. Fitting the beam is like drawing a line through points: points close together let the line pivot, while points far apart hold it in place. Mislocating the dot by one pixel at range z moves it sideways by z/f, where f is the focal length in pixels, and across a spread Δz between the nearest and farthest photographs that error tilts the fitted beam by about $(z/f)/\Delta z$, so halving the spread doubles the tilt. In one pool session, all the calibration photographs were taken within 3 cm of the same distance, so the calibration pinned the beam down only in that narrow neighborhood. It reproduces ranges there to within 0.12%, but reads 17% short at a distance 2.2 m farther away. None of the usual checks catches this, because the session's known-length targets sat near the calibration distance and measured well at −2.4%, reprojection error is blind to in-plane rotation, and leaving out any one photograph leaves the same narrow neighborhood. The distance between the slate and the camera is therefore swept outward and back during calibration (Section 3.4), and we require the calibration photographs to span at least 0.6 m of range.

5.4 Fish pose

Pose is the error source a diver controls most directly, and it costs length in a way the geometry predicts. A fish angled away from the camera, so that its broadside is not perpendicular to the camera's axis, appears shorter than it is. A fish angled θ from perpendicular projects to $\cos\theta$ of its length, as in Section 3.3. Because FishSense Lite places the snout and the fork at the laser's depth, it measures this projected length and reads $\cos\theta - 1$ short. Setting that shortfall equal to an allowable error ϵ gives the largest angle a measurement can tolerate,

$$\theta = \cos^{-1}(1 - \epsilon), \tag{12}$$

which is 31.8° at the 15% budget, and within 15° of perpendicular foreshortening costs at most 3.4%.

The pool data follow this curve over five sessions in which a Snook model was presented from 0° to 45° in 5° steps, with the angle read off a protractor card held in frame, on five camera units at about 2 and 4 m, for 1,428 frames (Figure 11). The sessions agree at every angle and in both ranges, and the pooled median follows the shape of $\cos\theta - 1$, crossing the budget at 30° versus the predicted 31.8°.

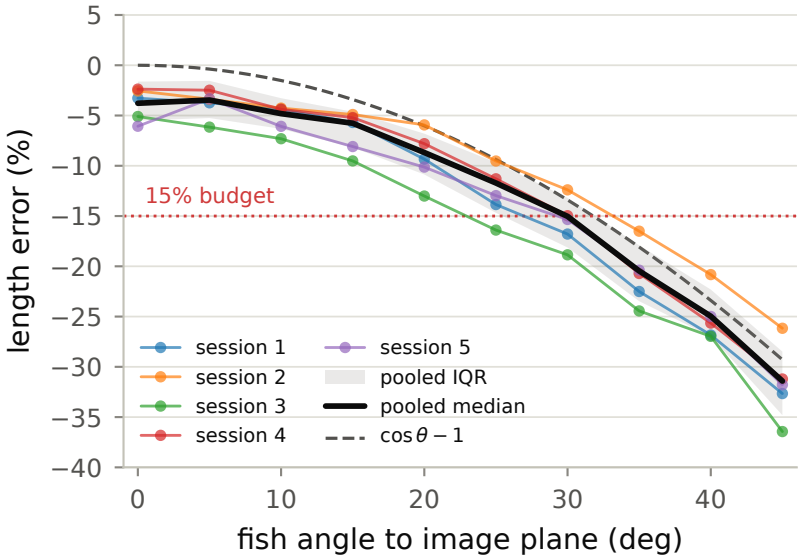


Fig. 11. Percent length error against the fish's angle from perpendicular, from five sessions of one target stepped through 0° to 45°. Thin lines are per-session medians, black is the pooled median and interquartile range, dashed is $\cos\theta - 1$, and dotted is the 15% budget.

The pooled curve sits below $\cos\theta - 1$ at every angle, by 3.8 points at 0° and 2.4 at 15°; an offset pose cannot produce this, since a protractor card misaligned enough to explain the 0° reading would need to be 15.9° out, and would then put 15° at −14.2% rather than the −5.8% observed. Part of the offset is the Snook model's body depth, which reads it about 1.8 points below the flat models measured alongside it (Section 5.2). The rest is calibration, since a scale error fits the data best and differs by session, from −6.5% to +0.2% across the five, a spread that body depth, a wrong reference length, or a consistent habit in where labelers click could not produce, as each would shift all five sessions by similar amounts. Of the −5.8% at 15°, $\cos 15° - 1 = -3.4$ points is pose.

Within 15° of perpendicular, every frame stays within budget, with room to spare. Over the 542 frames inside 15°, the median error is −4.7%, the p_{90} is −0.8%, and the worst is −10.7%. The first frames outside the budget appear at 25°, and the pooled median does not cross until 30°. We therefore recommend photographing fish within 15° of perpendicular, which keeps foreshortening small enough to leave room in the budget for the other sources of error. Laser-caliper studies draw a similar line, reporting errors of 5–7% within 10° of perpendicular and 10–20% beyond [51], and rejecting frames past 10° [25].

5.5 Working range

Range decides how far from a fish a diver can work, and across most of it, length error does not grow with distance. Beyond 0.82 m the median error of the pool cohort stays between −2.4% and −1.5%, and its spread, measured as the interquartile range, stays between 3.0 and 4.7 points. This holds out to 5 m, although we took only 30 frames beyond 4 m. The geometry does not set that limit, since at 5 m the usual error in placing the dot shifts the range by only 1.5%. In practice, beyond about 5 m the dot became hard both to see on a fish and to keep on it.

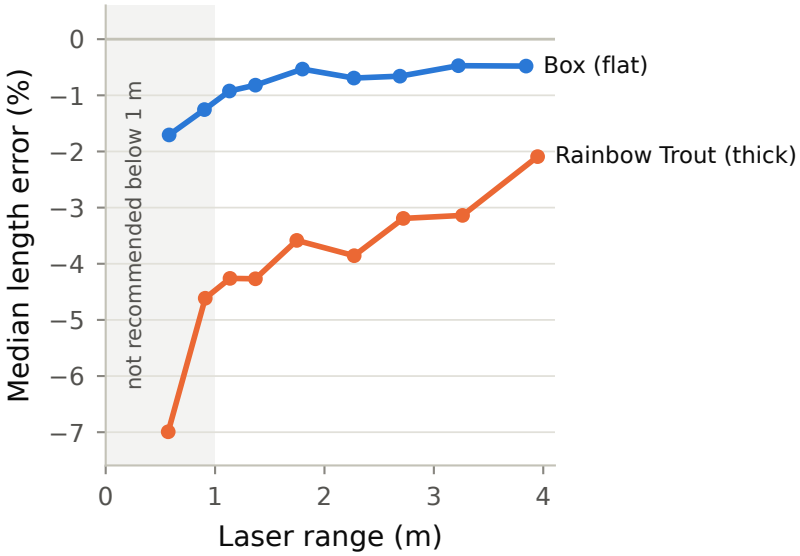


Fig. 12. Median length error against laser range for the Rainbow Trout, a thick model 58.7 mm deep, and the Box, a flat face, in fixed range bins over the pool cohort. Shaded: ranges under 1 m.

Triangulation noise (Section 3.3) does grow with range, but too slowly to show in these data. From the typical error in placing the labels, 0.9 px for the dot and 2.3 px for the snout and fork, it stays inside the observed spread out to about 5 m and would reach the 15% budget only beyond 18 m.

Close to the camera, length error depends on the target rather than on the range (Figure 12). Below 0.82 m, the trout reads −7.0% and the box only −1.7%. That gap is body depth (Section 3.3), because the laser lands on the trout's near flank, a fixed distance in front of the snout and fork, so the shortfall grows as a fraction of the range as the fish comes closer, and it predicts 5.0% at those frames' median range of 0.58 m. Close in, the gap between the two targets matches that prediction, though farther out the trout still reads about two points shorter than depth alone gives, from a source the data do not identify. For a real fish, then, length error falls with distance rather than rising, because the body-depth shortfall shrinks with range faster than triangulation noise grows.

Every real fish has body depth, so we recommend working beyond 1 m. None of the paired stereo frames were taken closer than 1 m, so the agreement of Section 5.1 does not depend on that range.

5.6 The field cohort

These open-water dives produced what Section 1 argues a stock assessment needs, a length paired with a species for each individual fish. Over seven dives, divers measured 73 wild fish from nine named species and a nontarget category, ranging from 13.0 to 84.7 cm with a median of 25.5 cm (Figure 13).

The field cohort has no known lengths, so it cannot measure accuracy, but it can measure repeatability, which needs no reference. We measure repeatability as the within-fish coefficient of variation, the spread of one fish's frames relative to their mean. A calibration error is common to every frame of a fish, so it cancels out in this measure; therefore, it cannot detect a calibration shift shared by every fish on a dive, a term worth 2.4% standard deviation in the pool (Table 6). For the 25 wild fish photographed in three or more frames, the median coefficient of variation is 2.9% (95%

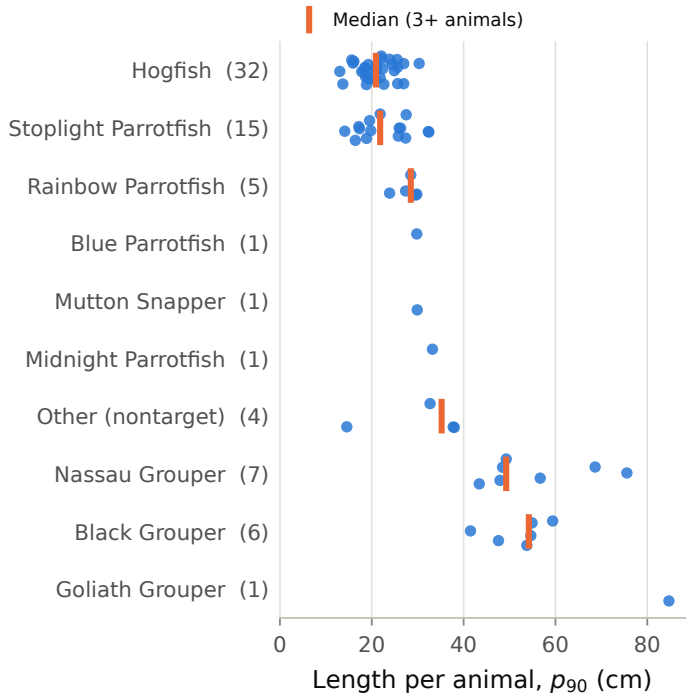


Fig. 13. Length per animal by species over the seven dives, one point per animal at its p_{90}. The bar marks the median for species with three or more animals.

confidence interval 1.6 to 4.1%). For a pool target photographed repeatedly within one session, the median over 33 such cases is 1.5% (95% confidence interval 1.3 to 2.0%), as Figure 14 shows. A wild fish therefore varies about twice as much between frames as a posed model, though with 25 animals the two intervals overlap at their edges. The likely sources are conditions the pool lacks, since the animal moves between frames, the water is turbid, and the snout and fork are harder to place.

Three properties of the cohort limit the precision of its lengths. Each wild animal carries at most eight frames, so its p_{90} is its longest, least foreshortened frame; in the pool, estimates from so few frames read slightly short, by 1.4% at two frames (Section 5.2). Four of the seven field calibrations rest on only two or three usable slate frames, enough to span range but too few for the fit to check itself against, though dot noise alone puts at most 1.2% uncertainty on their lengths at 3 m. And 35 of the 162 measurements were taken closer than 1 m, where body depth makes them read several percent short (Section 5.5).

6 Discussion

The error sources measured in the pool are each small, and together they keep a pool measurement within its budget. Table 6 collects them. The terms that vary at random from one session, dive, fish, or frame combine to a standard deviation of 3.6%, and two more only ever shorten a length: foreshortening, which the p_{90} offsets by favoring the least foreshortened frames, and body depth, which nothing in the current pipeline offsets. Against stereo imagery, the spread across seven fish was 10.1%, which leaves about 9 points of standard deviation that no source in Table 6 explains.

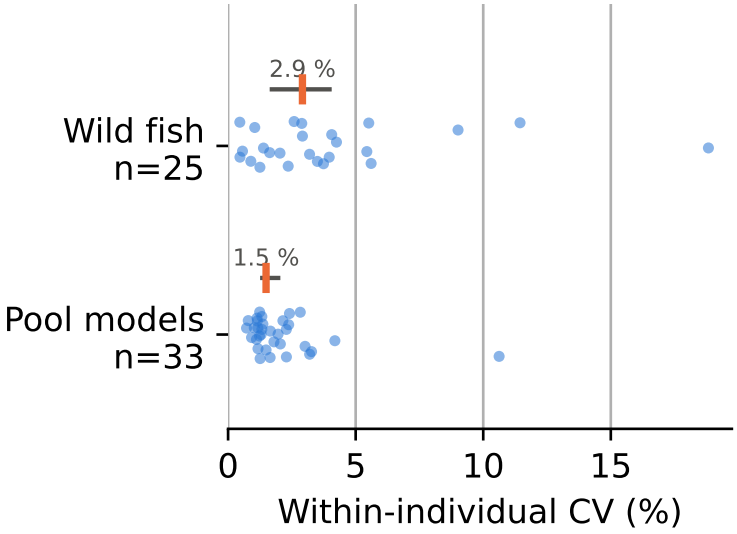


Fig. 14. Within-individual repeatability, wild fish against posed models. Each point is one wild individual with three or more frames, or one target photographed repeatedly in one pool session. Bars give the median and its 95% confidence interval.

Error source	Size (% of length)	Behavior	Section
Calibration scale between sessions	2.4 (SD)	random per session	5.4
Laser movement within a dive (two dives)	2.2 (SD)	random per dive	5.3
Labeling noise at 2 m	1.0 (SD)	random per frame	5.5
Few frames per fish	1.3 (SD)	random per fish	5.2
Pose within 15°	up to 3.4	always shortens	5.4
Body depth	1.5 at 2 m, 5.0 at 0.58 m	always shortens	5.2, 5.5

Table 6. The error sources measured in this paper, as magnitudes in percent of length. Each random term is a standard deviation: calibration scale, the spread of the scale errors of the five angle-experiment sessions (−3.0, +0.2, −6.5, −2.3 and −3.3%; Section 5.4); laser movement within a dive, the largest change observed between a dive's two ends, 3.7%, treated as the edge of a uniform distribution; labeling noise, the 0.9 px dot and 2.3 px endpoint errors of Section 5.5 for a 30 cm fish at 2 m; and few frames per fish, the spread a three-frame p_{90} adds over an all-frames p_{90} across the pool targets of Section 5.2. These combine as independent errors, as the square root of the sum of their squares, to 3.6%. The last two rows always shorten a length, so they add rather than cancel.

Calibration is not the cause, since a calibration error shifts every fish on a dive alike, while the spread lies among fish on the same dive (Section 5.1); the remainder lies in sources that differ fish by fish, and separating them will take a larger paired study.

Two of these terms point at the design rather than the diver. The mount is the largest source of error a design can remove, since its joints move the beam between dives by more than the budget allows (Section 5.3), and a stiffer mount would shrink the movement that per-dive calibration exists to catch. Calibration is the one step divers must repeat on every dive, and the manual's closing calibration was the part most often left out, with only two of the seven field dives

calibrated at both ends (Section 4.5), so it is the part of the protocol most worth making easier. Body depth, the one bias left uncorrected, could be removed with a per-species correction wherever depth-to-length ratios are known.

These results have three limits. The comparison with stereo imagery covers seven fish on one day, too few to bound error for a single wild fish, so the pool provides per-measurement accuracy evidence and the field cohort shows repeatability rather than accuracy. Laser movement within the camera–laser plane during a dive cannot be seen in the photographs, and was measured on two field dives. REEF briefed its divers in person, so the deployment shows that a citizen-science program can teach the system without its builders, but not yet that the written manual alone is enough.

7 Conclusion

FishSense Lite gives a citizen-science program fish lengths, paired with species, from commodity hardware and a single photograph. On targets of known length, it recovers each to within 1.3%, and on live fish it agrees with stereo imagery without detectable bias. From the error sources these results measure, we derive a field protocol, in which the diver calibrates against a dive slate on every dive, photographs fish within 15° of perpendicular and beyond 1 m, and takes at least ten frames of each. REEF, trained once by the system's builders, then taught its own divers and collected lengths paired with species for 73 wild fish across seven dives.

The design meets the *commodity components* and *deferred judgment* design goals set in Section 1, since every part is owned, ordered from a consumer retailer or printed, and species and length are judged afterward by people equipped to make them. It meets *manual-only training* in content, since everything a diver must know fits in a written manual, though in this deployment divers were also briefed in person. That combination lets length paired with species come from divers wherever they already dive, adding to the existing surveys rather than replacing them.

Acknowledgments

Funding and resources. This work was funded in part by the NSF Research Experiences for Undergraduates (REU) Site: Engineers for Exploration at UC San Diego (Award 2244123) and the Matthew Henson Fellowship Program supported by the UC HBCU Program. We also acknowledge the Computer Science and Engineering department at UC San Diego and the UCSD IEEE Supercomputing Club at San Diego Supercomputer Center for hosting compute resources for this research. We thank the REEF volunteer divers who collected the field data.

Competing interests. Christopher L. Crutchfield, Brice X. Semmens, and Ryan Kastner are co-founders of Aqua3D, a company that may commercialize technology described in this paper. The measurement method is the subject of a pending patent application [31].

References

[1] Elizabeth A. Ashley, Christy V. Pattengill-Semmens, James W. Orr, Janna D. Nichols, and Joseph K. Gaydos. 2022. Documenting fishes in an inland sea with citizen scientist diver surveys: using taxonomic expertise to inform the observation potential of fish species. *Environmental Monitoring and Assessment* 194, 3 (2022), 227. doi:10.1007/s10661-022-09857-1

[2] Jerald S. Ault, Steven G. Smith, Jiangang Luo, Mark E. Monaco, and Richard S. Appeldoorn. 2008. Length-based assessment of sustainability benchmarks for coral reef fishes in Puerto Rico. *Environmental Conservation* 35, 3 (2008), 221–231. doi:10.1017/s0376892908005043

[3] Saulo Vinicius Ferreira Barreto, Remy Eskinazi Sant'Anna, and Marcílio André Félix Feitosa. 2013. A method for image processing and distance measuring based on laser distance triangulation. In *2013 IEEE 20th International Conference on Electronics, Circuits, and Systems (ICECS)*. 695–698. doi:10.1109/ICECS.2013.6815509

[4] Aaron Bartholomew and James A. Bohnsack. 2005. A Review of Catch-and-Release Angling Mortality with Implications for No-take Reserves. *Reviews in Fish Biology and Fisheries* 15, 1-2 (2005), 129–154. doi:10.1007/s11160-005-2175-1

[5] J. D. Bell, G. J. S. Craik, D. A. Pollard, and B. C. Russell. 1985. Estimating length frequency distributions of large reef fish underwater. *Coral Reefs* 4, 1 (1985), 41–44. doi:10.1007/BF00302203

[6] Magnus Bjerkeng, Trine Kirkhus, Walter Caharija, Jens T. Thielemann, Herman B. Amundsen, Sveinung Johan Ohrem, and Esten Ingar Grøtli. 2021. ROV Navigation in a Fish Cage with Laser-Camera Triangulation. *Journal of Marine Science and Engineering* 9, 1 (2021), 79. doi:10.3390/jmse9010079

[7] J. Martin Bland and Douglas G. Altman. 1986. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet* 327, 8476 (1986), 307–310. doi:10.1016/S0140-6736(86)90837-8

[8] Gary R. Bradski and Adrian Kaehler. 2008. *Learning OpenCV: Computer Vision with the OpenCV Library* (1 ed.). O'Reilly Media, Sebastopol, CA.

[9] M. J. Chantler, James Clark, and Manickam Umasuthan. 1997. Calibration and operation of an underwater laser triangulation sensor: the varying baseline problem. *Optical Engineering* 36, 09 (1997), 2604–2611. doi:10.1117/1.601500

[10] Katrina V. Cook, Robert J. Lennox, Scott G. Hinch, and Steven J. Cooke. 2015. Fish Out of Water: How Much Air is Too Much? *Fisheries* 40, 9 (09 2015), 452–461. doi:10.1080/03632415.2015.1074570

[11] S. J. Cooke and H. L. Schramm. 2007. Catch-and-release science and its application to conservation and management of recreational fisheries. *Fisheries Management and Ecology* 14, 2 (2007), 73–79. doi:10.1111/j.1365-2400.2007.00527.x

[12] Christopher L. Crutchfield, Nykolas Rekasius, Anu Jajodia, Theodore Darci-Maher, Sijan Shrestha, Aayushka Budhathoki, Annabelle Chen, Brendon Mwamba, Zach Merson, Shawn Doyle, Grace Turner, Jakob Tallman, Cole Wilson, Jack Elstner, Brice X. Semmens, Curt Schurgers, and Ryan Kastner. 2025. FishSense Mobile: A Mobile Device App for on-Deck Fisheries Management Operations. In *OCEANS 2025 - Great Lakes*. IEEE, 1–10. doi:10.23919/OCEANS59106.2025.11245071

[13] Laura Davis and Jacques Malan. 2024. Developing a Prototype Underwater LiDAR System. In *OCEANS 2024 - Halifax*. IEEE, 1–9. doi:10.1109/oceans55160.2024.10754530

[14] Mark H. Deakos. 2010. Paired-laser photogrammetry as a simple and accurate system for measuring the body size of free-ranging manta rays Manta alfredi. *Aquatic Biology* 10, 1 (2010), 1–10. doi:10.3354/ab00258

[15] C. R. Delacy, A. Olsen, L. A. Howey, D. D. Chapman, E. J. Brooks, and M. E. Bond. 2017. Affordable and accurate stereo-video system for measuring dimensions underwater: a case study using oceanic whitetip sharks Carcharhinus longimanus. *Marine Ecology Progress Series* 574 (2017), 75–84. doi:10.3354/meps12190

[16] Diving Equipment and Marketing Association. 2024. *2024 Diving Fast Facts*. Technical Report. https://www.dema.org/store/download.aspx?id=7811b097-8882-4707-a160-f999b49614b6

[17] Andrew Filisetti, Andreas Marouchos, Andrew Martini, Tara Martin, and Simon Collings. 2018. Developments and applications of underwater LiDAR systems in support of marine science. In *OCEANS 2018 MTS/IEEE Charleston*. IEEE, 1–10. doi:10.1109/oceans.2018.8604547

[18] Jason Gedge, Minglun Gong, and Yee-Hong Yang. 2011. Refractive Epipolar Geometry for Underwater Stereo Matching. In *2011 Canadian Conference on Computer and Robot Vision*. IEEE, 146–152. doi:10.1109/crv.2011.26

[19] Jordan S. Goetze, Todd. Bond, Dianne L. McLean, Benjamin J. Saunders, Tim J. Langlois, Steve Lindfield, Laura. A. F. Fullwood, Damon Driessen, George Shedrawi, and Euan S. Harvey. 2019. A field and video analysis guide for diver operated stereo-video. *Methods in Ecology and Evolution* 10, 7 (2019), 1083–1090. doi:10.1111/2041-210X.13189

[20] Jordi Guillen, Fabrizio Natale, Natacha Carvalho, John Casey, Johann Hofherr, Jean-Noël Druon, Gianluca Fiore, Maurizio Gibin, Antonella Zanzi, and Jann Th. Martinsohn. 2019. Global seafood consumption footprint. *Ambio* 48, 2 (2019), 111–122. doi:10.1007/s13280-018-1060-9

[21] Scott Hamilton, Rick Starr, Dean Wendt, Benjamin Ruttenberg, Jennifer Caselle, Brice Semmens, Lyall Bellquist, Steven Morgan, Tim Mulligan, Joe Tyburczy, Shelby Ziegler, Rachel Brooks, Grant Waltz, Erica Mason, Chris Honeyman, Sadie Small, and Jay Staton. 2021. *CCFRP Final Report*. Technical Report. UC San Diego. https://caseagrant.ucsd.edu/sites/default/files/CCFRP_Final_Report.pdf

[22] K. C. Hanson, S. J. Cooke, C. D. Suski, and D. P. Philipp. 2007. Effects of different angling practices on post-release behaviour of nest-guarding male black bass, Micropterus spp. *Fisheries Management and Ecology* 14, 2 (2007), 141–148. doi:10.1111/j.1365-2400.2007.00534.x

[23] Richard Hartley and Andrew Zisserman. 2004. *Multiple View Geometry in Computer Vision* (2 ed.). Cambridge University Press, Cambridge, UK. doi:10.1017/CBO9780511811685

[24] Euan Harvey, David Fletcher, and Mark Shortis. 2001. A comparison of the precision and accuracy of estimates of reef-fish lengths determined visually by divers with estimates produced by a stereo-video system. *Fishery Bulletin* 99, 1 (2001), 63–71. https://spo.nmfs.noaa.gov/sites/default/files/pdf-content/2001/991/Harvey.pdf

[25] Scott A. Heppell, Brice X. Semmens, Stephanie K. Archer, Christy V. Pattengill-Semmens, Philippe G. Bush, Croy M. McCoy, Selina S. Heppell, and Bradley C. Johnson. 2012. Documenting recovery of a spawning aggregation through size frequency analysis from underwater laser calipers measurements. *Biological Conservation* 155 (oct 2012), 119–127. doi:10.1016/j.biocon.2012.06.002

[26] Ray Hilborn, Ricardo Oscar Amoroso, Christopher M. Anderson, Julia K. Baum, Trevor A. Branch, Christopher Costello, Carryn L. de Moor, Abdelmalek Faraj, Daniel Hively, Olaf P. Jensen, Hiroyuki Kurota, L. Richard Little, Pamela Mace, Tim McClanahan, Michael C. Melnychuk, Cóilín Minto, Giacomo Chato Osio, Ana M. Parma, Maite Pons, Susana Segurado, Cody S. Szuwalski, Jono R. Wilson, and Yimin Ye. 2020. Effective fisheries management instrumental in improving fish stock status. *Proceedings of the National Academy of Sciences* 117, 4 (2020), 2218–2224. doi:10.1073/pnas.1909726116

[27] Kyle Sirong Hu. 2024. *Methods for In-Situ Fish Length Measurement with a Single Laser and Camera to Support Citizen Science*. Master's thesis. University of California San Diego, San Diego, CA, USA. https://escholarship.org/uc/item/41x511ht

[28] Nathalie Jaquet. 2006. A SIMPLE PHOTOGRAMMETRIC TECHNIQUE TO MEASURE SPERM WHALES AT SEA. *Marine Mammal Science* 22, 4 (2006), 862–879. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1748-7692.2006.00060.x doi:10.1111/j.1748-7692.2006.00060.x

[29] Erica T. Jarvis and Christopher G. Lowe. 2008. The effects of barotrauma on the catch-and-release survival of southern California nearshore and shelf rockfish (Scorpaenidae, Sebastes spp.). *Canadian Journal of Fisheries and Aquatic Sciences* 65, 7 (jul 2008), 1286–1296. doi:10.1139/f08-071

[30] Simon Jennings, Michel J. Kaiser, and John D. Reynolds. 2001. *Marine fisheries ecology.* Blackwell Science, Oxford.

[31] Ryan Kastner, Christopher L. Crutchfield, Kyle Hu, Nathan Hui, and Brice X. Semmens. 2025. Imaging systems and methods for dimension measurement independent of laser alignment. U.S. Patent Application Publication US 2025/0315981 A1. https://patents.google.com/patent/US20250315981A1/en

[32] Jiakang Li, Shengmao Zhang, Penglong Li, Yang Dai, and Zuli Wu. 2024. Research on Measuring the Bodies of Underwater Fish with Inclined Positions Using the YOLOv8 Model and a Line-Laser System. *Fishes* 9, 6 (2024), 206. doi:10.3390/fishes9060206

[33] Delphine Mallet and Dominique Pelletier. 2014. Underwater video techniques for observing coastal marine biodiversity: A review of sixty years of publications (1952-2012). *Fisheries Research* 154 (2014), 44–62. doi:10.1016/j.fishres.2014.01.019

[34] Victoria Y. Martin, Leslie Christidis, David J. Lloyd, and Gretta T. Pecl. 2016. Understanding Drivers, Barriers and Information Sources for Public Participation in Marine Citizen Science. *Journal of Science Communication* 15, 2 (2016), A02. doi:10.22323/2.15020202

[35] Gabriel Matos, Pedro Buschinelli, and Tiago Pinto. 2020. Underwater Laser Triangulation Sensor Model With Flat Refractive Interfaces. *IEEE Journal of Oceanic Engineering* 45, 3 (2020), 937–945. doi:10.1109/JOE.2019.2891863

[36] Robert P. Mueller, Richard S. Brown, Haakon Hop, and Larry Moulton. 2006. Video and acoustic camera techniques for studying fish under ice: a review and comparison. *Reviews in Fish Biology and Fisheries* 16, 2 (sep 2006), 213–226. doi:10.1007/s11160-006-9011-0

[37] K. Muljowidodo, Mochammad A. Rasyid, SaptoAdi N., and Agus Budiyono. 2009. Vision based distance measurement system using single laser pointer design for underwater vehicle. *Indian Journal of Marine Sciences* 38, 3 (2009), 324–331. https://nopr.niscpr.res.in/handle/123456789/6209

[38] OpenCV. 2026. *Perspective-n-Point (PnP) pose computation.* https://docs.opencv.org/4.13.0/d5/d1f/calib3d_solvePnP.html

[39] Christy V. Pattengill-Semmens and Brice X. Semmens. 2003. Conservation and Management Applications of the Reef Volunteer Fish Monitoring Program. In *Coastal Monitoring through Partnerships*. Springer Netherlands, 43–50. doi:10.1007/978-94-017-0299-7_5

[40] AT Pettit, M Haggerty, MR Morse, R Freedman, F Joyce, C Jainese, K Seeto, JE Caselle, and DJ McCauley. 2025. Using community science to assess the population dynamics and spatial ecology of the giant sea bass in Southern California. *Marine Ecology Progress Series* 760 (May 2025), 151–169. doi:10.3354/meps14843

[41] D. A. Pilgrim, D. M. Parry, M. B. Jones, and M. A. Kendall. 2000. ROV Image Scaling with Laser Spot Patterns. *Underwater Technology* 24, 3 (2000), 93–103. doi:10.3723/175605400783259684

[42] Robin M. Pope and Edward S. Fry. 1997. Absorption spectrum (380–700 nm) of pure water. II. Integrating cavity measurements. *Applied Optics* 36, 33 (1997), 8710–8723. doi:10.1364/AO.36.008710

[43] Alena L. Pribyl, Michael L. Kent, Steven J. Parker, and Carl B. Schreck. 2011. The Response to Forced Decompression in Six Species of Pacific Rockfish. *Transactions of the American Fisheries Society* 140, 2 (2011), 374–383. doi:10.1080/00028487.2011.567858

[44] Polly S. Rankin, Robert W. Hannah, Matthew T.O. Blume, Timothy J. Miller-Morgan, and Jerry R. Heidel. 2017. Delayed effects of capture-induced barotrauma on physical condition and behavioral competency of recompressed yelloweye rockfish, Sebastes ruberrimus. *Fisheries Research* 186 (feb 2017), 258–268. doi:10.1016/j.fishres.2016.09.004

[45] Daniel Ricard, Cóilín Minto, Olaf P Jensen, and Julia K Baum. 2011. Examining the knowledge base and status of commercially exploited marine species with the RAM Legacy Stock Assessment Database. *Fish and Fisheries* 13, 4 (2011), 380–398. doi:10.1111/j.1467-2979.2011.00435.x

[46] C. A. Rohner, A. J. Richardson, A. D. Marshall, S. J. Weeks, and S. J. Pierce. 2011. How large is the world's largest fish? Measuring whale sharks Rhincodon typus with laser photogrammetry. *Journal of Fish Biology* 78, 1 (2011), 378–385. doi:10.1111/j.1095-8649.2010.02861.x

[47] Christoph A. Rohner, Anthony J. Richardson, Clare E.M. Prebble, Andrea D. Marshall, Michael B. Bennett, Scarla J. Weeks, Geremy Cliff, Sabine P. Wintner, and Simon J. Pierce. 2015. Laser photogrammetry improves size and demographic estimates for whale sharks. *PeerJ* 3 (2015), e886. doi:10.7717/peerj.886

[48] Jodie L. Rummer and Wayne A. Bennett. 2005. Physiological Effects of Swim Bladder Overexpansion and Catastrophic Decompression on Red Snapper. *Transactions of the American Fisheries Society* 134, 6 (2005), 1457–1470. doi:10.1577/T04-235.1

[49] SeaGIS. 2026. *Hardware.* https://www.seagis.com.au/hardware.html

[50] Christopher D. Stallings. 2009. Fishery-Independent Data Reveal Negative Effect of Human Population Density on Caribbean Predatory Fish Communities. *PLoS ONE* 4, 5 (2009), e5333. doi:10.1371/journal.pone.0005333

[51] Brian C Stock, Scott A Heppell, Lynn Waterhouse, India C Dove, Christy V Pattengill-Semmens, Croy M McCoy, Phillippe G Bush, Gina Ebanks-Petrie, and Brice X Semmens. 2021. Pulse recruitment and recovery of Cayman Islands Nassau Grouper (Epinephelus striatus) spawning aggregations revealed by in situ length-frequency data. *ICES Journal of Marine Science* 78, 1 (2021), 277–292. doi:10.1093/icesjms/fsaa221

[52] Anna Szafarczyk and Cezary Toś. 2023. The Use of Green Laser in LiDAR Bathymetry: State of the Art and Recent Advancements. *Sensors* 23, 1 (2023), 292. doi:10.3390/s23010292

[53] M. Takatsuka, G.A.W. West, S. Venkatesh, and T.M. Caelli. 1999. Low-cost interactive active monocular range finder. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*. IEEE Comput. Soc, Fort Collins, CO, USA, 444–449. doi:10.1109/CVPR.1999.786976

[54] Masahiro Takatsuka, Geoff A.W. West, Svetha Venkatesh, and Terry M. Caelli. 2003. Low-cost interactive active range finder. *Machine Vision and Applications* 14, 3 (2003), 139–144. doi:10.1007/s00138-003-0129-y

[55] Tali Treibitz, Yoav Schechner, Clayton Kunz, and Hanumant Singh. 2012. Flat Refractive Geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 1 (2012), 51–65. doi:10.1109/TPAMI.2011.105

[56] Peter Tueller, Raghav Maddukuri, Patrick Paxson, Vivaswat Suresh, Arjun Ashok, Madison Bland, Ronan Wallace, Julia Guerrero, Brice Semmens, and Ryan Kastner. 2021. FishSense: Underwater RGBD Imaging for Fish Measurement. In *OCEANS 2021: San Diego – Porto*. IEEE. doi:10.23919/oceans44145.2021.9705929

[57] U.S. Food and Drug Administration. 2026. Laser Products. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-J/part-1040

[58] Shuxian Wang, Shengmao Zhang, Yongchuang Shi, Zuli Wu, and Tianfei Cheng. 2026. Integrated Laser Imaging for Fusiform Fish Measurement in Aquaculture. *Fishes* 11, 5 (2026), 298. doi:10.3390/fishes11050298

[59] Christine A. Ward-Paige, Camilo Mora, Heike K. Lotze, Christy Pattengill-Semmens, Loren McClenachan, Ery Arias-Castro, and Ransom A. Myers. 2010. Large-Scale Absence of Sharks on Reefs in the Greater-Caribbean: A Footprint of Human Pressures. *PLoS ONE* 5, 8 (2010), e11968. doi:10.1371/journal.pone.0011968

[60] Trudi Webster, Steve Dawson, and Elisabeth Slooten. 2010. A simple laser photogrammetry technique for measuring Hector's dolphins (Cephalorhynchus hectori) in the field. *Marine Mammal Science* 26, 2 (2010), 296–308. doi:10.1111/j.1748-7692.2009.00326.x

[61] Emily Wong, Isabella Humphrey, Scott Switzer, Christopher Crutchfield, Nathan Hui, Curt Schurgers, and Ryan Kastner. 2022. Underwater Depth Calibration Using a Commercial Depth Camera. In *Proceedings of the 16th International Conference on Underwater Networks & Systems*. Association for Computing Machinery, 1–5. doi:10.1145/3567600.3568158

[62] Joanna B. Wong and Marie Auger-Méthé. 2018. Using laser photogrammetry to measure long-finned pilot whales (Globicephala melas). *Proceedings of the Nova Scotian Institute of Science (NSIS)* 49, 2 (2018), 269–291. doi:10.15273/pnsis.v49i2.8164

[63] Z. Zhang. 2000. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11 (2000), 1330–1334. doi:10.1109/34.888718