# Calibrating an Underwater Camera in Air: Removing the Corrective Optic from a Laser–Camera Rig

Christopher L. Crutchfield
University of California San Diego
La Jolla, California, USA
ccrutchf@ucsd.edu

Curt Schurgers
University of California San Diego
La Jolla, California, USA
cschurgers@ucsd.edu

Ryan Kastner
University of California San Diego
La Jolla, California, USA
kastner@ucsd.edu

## Abstract

FishSense Lite measures fish length from a single photograph using a commodity dive camera and a laser pointer, a combination cheap enough to put in the hands of recreational divers. Every part of it is commodity except one: the camera intrinsics are calibrated from a manufactured checkerboard, photographed underwater in a pool, and training citizen scientists to do this reliably is hard. This work removes that step. Modeling the flat port analytically lets the camera be calibrated in air, and in air the calibration target can be a structure of LEGO™ bricks, whose molded pitch needs no measurement since they are famously precisely manufactured. Doing this requires taking off the corrective lens the deployed system relies on underwater, so we first establish what the bare flat port costs: over the region a laser–camera rig can actually measure, the uncorrected length error is 1.2% at the median, an order of magnitude below the free-field figure, because the port's error enters range and angular extent together and cancels in their product. The brick calibration then disagrees with the checkerboard by 2.46% in focal length and 112 px in the principal point, yet delivers fish lengths that differ by 0.03 pp, because the intrinsics it misses are the ones the measurement does not depend on. Every calibration the system needs now runs on objects a diver already owns. The end result is a low-cost fish measurement instrument that is readily available and easily deployed by citizen scientists.

## CCS Concepts

• **Computer systems organization** → *Sensors and actuators.*

## Keywords

Laser triangulation rangefinding, flat port, camera calibration, citizen science, underwater sensing, ecology

**ACM Reference Format:**
Christopher L. Crutchfield, Curt Schurgers, and Ryan Kastner. 2026. Calibrating an Underwater Camera in Air: Removing the Corrective Optic from a Laser–Camera Rig. In *Proceedings of The 20th International Conference on Underwater Networks and Systems (WUWNet '26).* ACM, New York, NY, USA, 8 pages. https://doi.org/XXXXXXX.XXXXXXX

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
*WUWNet '26, Hong Kong, China*
© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

## 1 Introduction

Monitoring open-water fish populations is essential for sustainable management and for protecting one of the world's most significant protein sources. Three billion people worldwide rely on fish and seafood for protein [9]. Fisheries assessment depends on fish length: it informs weight and age distributions, and through these, stock status [3, 14]. Obtaining this information at scale is a key challenge.

One common method of quantifying fish populations is via catch-and-release campaigns. However, these require significant human resources; for example, the California Collaborative Fisheries Research Program (CCFRP) used 1,700 volunteer anglers and the crews of thirty-six vessels to capture 176,594 fish between 2007 and 2020 [11]. These measurement programs are also subject to significant biases [15, 18] and introduce fish mortality issues, as fish brought to the surface suffer barotrauma in air [13]. These programs rely on commercial fishing vessels, which biases the collected data near active commercial harbors and popular commercial fishing sites [23].

To address barotrauma and regional blind spots, scientists have turned to dive surveys to quantify fish populations [7, 20], leveraging the large community of recreational divers (2.5 million in the United States alone [1]). A key element of modern diver-assisted fish population assessment is using underwater cameras to estimate fish size, a significant improvement over the highly error-prone human visual measurement [10, 12].

A recent example of diver-assisted fish measurement technology is the FishSense Lite (FSL) underwater sensing platform [6]. FSL consists of a commodity dive camera equipped with a laser pointer on a 3D-printed mount. The diver points the laser at the fish and takes a single photograph. The laser dot's position in the image gives the range, and the snout and tail-fork pixels, back-projected to that range, give the length.

The FSL system is straightforward to operate *once it is calibrated.* The system needs two calibrations. (1) The camera intrinsics describe the lens; they do not drift, so they are fitted once per camera. (2) The laser extrinsics describe the laser's position and direction relative to the camera; they drift as the mount shifts, so we refit them every dive from photos of a commodity dive slate. The intrinsics are harder for citizen scientists to obtain because they require photographing a manufactured checkerboard underwater in a pool [6]. The waterproof calibration checkerboard costs approximately $300 USD [8]. Our proposed LEGO™ calibration target costs approximately $20 USD and is readily available from LEGO™ stores online or in person.

Camera calibration makes FSL less of an easy-to-use citizen science tool and requires extra effort from the diver. Recreational

divers do not own a calibration checkerboard, and underwater calibration is another hurdle. Because intrinsics are fitted per camera, every prospective citizen scientist must calibrate the lens before their first dive. This calibration friction can significantly hinder citizen scientist participation. Martin et al. [17] found that marine users judge a project's technical demands before joining, and that participants may drop out preemptively over anticipated difficulties.

We propose an inexpensive replacement option using commodity toy bricks, such as LEGO™ bricks, to build calibration structures. They are molded to a published pitch, identical worldwide, and already present in many households [19]. Unlike a printed board, their geometry needs no extra measurements before it can be trusted, and Baronti et al. [5] showed that a camera can be calibrated using a structure built from commodity bricks.

Toy bricks cannot be used for underwater calibration. They trap air, so an assembled target is buoyant. More importantly, an optical barrier prevents calibration with an object that is not a flat surface. A planar board sits at effectively one depth per view. At a fixed depth, flat-port refraction is well-approximated by radial distortion (Section 2). That is why calibrating a plane underwater works. A three-dimensional target, however, spans multiple depths within a single view, where the refraction varies with depth. A single radial model cannot account for it.

Calibration with a commodity bricks-based target must therefore be performed in air. This impacts the rest of the system, and most importantly, the corrective lens must be removed for both calibration and data collection.

In the current FSL system, a corrective wet-mount lens accounts for underwater distortion [4, 6]. It threads onto the housing's flat-port, and the gap between the two floods with water as it enters. Its domed element cancels the refraction the flat-port would otherwise cause, restoring the camera's topside field of view. The correction depends on water filling that gap. In air, the gap holds air; the assembly is a different optical system, and a calibration taken through it transfers nothing. The lens is made to flood and is not watertight; calibration in air with the lens holding water is not possible. Calibrating dry means taking the optic off and modeling the bare pane instead. We use Łuczyński et al.'s Pinax model [16] unchanged, and investigate the implications of dry calibration on system-level accuracy. The main objective of this work is to determine whether the Pinax model, combined with commodity components for calibration, provides the end-to-end accuracy required for underwater sensing applications such as fish population measurement.

Specifically, our contributions are:

(1) A camera-intrinsics calibration procedure for an underwater laser–camera rig that runs in air on a target built from commodity bricks. It combines the Pinax refraction model, applied unchanged, with Baronti et al.'s brick target, and removes the manufactured board and the pool session (Section 5).
(2) Evidence that a bare flat-port, modeled with Pinax, can replace the corrective optic, so that a calibration taken in air holds underwater across the whole frame, and that the bare port's cost to length is small over the region a laser–camera rig can measure (Section 4).
(3) A self-contained test for whether a port model is working, from the homography residual against field angle, needing no reference object, no known pose, and no known scale (Section 4).
(4) A characterization of what the brick target does and does not settle: the focal-length ratio is far more stably recovered than the absolute scale, and the calibration differences between a brick target and a checkerboard do not reach the delivered length (Section 5).

Our work shows that we can replace expensive, cumbersome underwater calibration with cost-effective in-air calibration while staying within error tolerances. The proposed in-air calibration system makes citizen-science data-collection dives more accessible to non-technical participants.

## 2 Background

A camera is normally treated as a single viewpoint (SVP) device: every ray converges on one point in space. The pinhole model [21] rests on that assumption, and so does nearly every calibration and reconstruction method built on it.

The SVP assumption does not hold underwater, complicating underwater photography and camera calibration. A dive housing presents a flat pane to the water, and rays crossing this pane bend according to Snell's Law of Refraction. The bend increases with the angle between the ray and the pane's normal, meaning rays at wider angles suffer more distortion. There are two consequences of this. Scene directions are compressed toward the axis, which magnifies the image and narrows the field of view. More importantly, the degree of the bend depends on where the ray crosses the pane, so the rays no longer share an origin, violating the key SVP expectation.

Treibitz et al. [22] showed that a camera behind a flat-port, as in a dive housing, is not an SVP system and that correcting it requires a physics-based model of the interface. Agrawal et al. [2] characterized such a camera as *axial*. An axial camera is one whose rays intersect a line through the optical axis rather than a point. Methods that assume a single center of projection do not apply to it.

This leaves two options. The first is to restore the SVP in hardware. A dome port, or a domed accessory lens mounted at the port, presents a curved interface that rays cross near normal incidence, so the camera behaves as a pinhole again. This is the principle of the corrective lens FSL uses [4]. The second is to model the effect of the pane. Łuczyński et al. [16] construct a virtual pinhole camera from the port geometry and the water's refractive index, a model they call Pinax. Because the model is parameterized by the port rather than fitted underwater, a calibration taken in air transfers to the water. We implement their model in this work to support in-air calibration with commodity parts.

The target is the other half of the problem. Zhang's method [24], which remains the standard procedure, fits intrinsics from several views of a planar board. Zhang's calibration method is used underwater in practice, including by FSL [6], and it still serves behind a bare flat-port. The planar calibration board sits at effectively one depth per view, and at fixed depth the port's error is absorbed by the radial distortion terms the fit already estimates (Figure 1). The result is physically approximate but serviceable. A three-dimensional

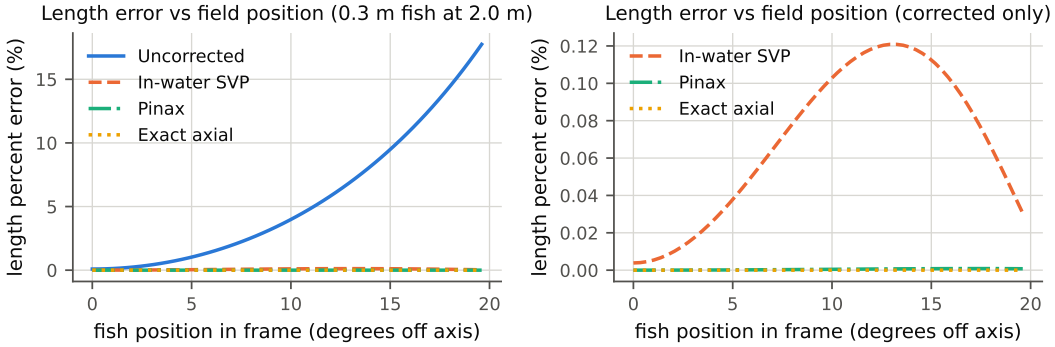



**Figure 1: Length error against field position for a 0.3 m fish at 2.0 m. An in-water SVP fit, a single radial model, holds the error to about 0.1% across the frame at the depth it was fitted for, which is why a planar board calibrates underwater. "Exact axial" is the simulation's camera model, so its error is zero by construction.**

target has no such reprieve, because its features span depths within one view, where the port's error varies and a single radial term cannot absorb it.

A board is nonetheless a fabricated object whose pitch must be trusted before the calibration can be. Baronti et al. [5] proposed commodity LEGO™ bricks in its place, exploiting a molded pitch that is identical worldwide and needs no measurement. Their study is terrestrial. In this work, we apply their findings to the in-air calibration of an underwater camera.

## 3 System

We work with FishSense Lite [6], a handheld underwater fish measurement device built from components recreational divers commonly own. The system consists of an OM System TG-6 in a PT-059 underwater enclosure. An attached waterproof laser pointer on a printed mount is registered to the housing's cold shoe, which places the beam rigidly offset from the optical axis. The diver aims the laser at a fish, frames the shot so the snout and tail-fork are visible, and takes one photograph.

Each measurement depends on two calibrations. The laser extrinsics drift, since the printed mount flexes and the cold-shoe connection has play. They are re-fitted every dive from roughly ten images of a duct-taped dive slate, an object divers already carry. The camera intrinsics do not drift and are fitted once per camera. Obtaining them requires photographing a manufactured checkerboard in a pool.

Length is calculated in two steps. The laser dot's pixel position encodes range, so intersecting the back-projected dot ray with the calibrated laser ray recovers a metric depth. The snout and tail-fork pixels are then back-projected through $K^{-1}$, placed at the depth calculated, and the length is calculated as the Euclidean distance between them. Both steps consume the camera intrinsics $K$, which come from the camera calibration. Length is the product of a range and an angular extent, and both are read through the same camera and the same water. An error that scales the two together divides out of the length, even when it is large in each. This cancellation recurs throughout what follows.

In the baseline FishSense Lite system, hardware restores the pinhole behavior. A wet-mount domed accessory lens threads onto the housing's flat-port [4], and the gap between the two floods with water on entry, canceling the refraction the flat port would otherwise induce. Flooding is a step the diver must perform: the lens is removed underwater and refitted so water fills the gap, meaning the lens does not return to the exact same position dive to dive. It also means air can be trapped under the lens. These trapped air bubbles can be a source of error in human-assisted laser calibration. A bubble that occludes nothing important is recoverable. A bubble over the laser dot costs a measurement, and one over a fish's head or tail leaves an annotator to infer the location. A bubble over the dive slate during calibration is the worst-case scenario, because the scale for the entire dive comes from those frames, and failure to recover it means a lost dive. All of this is invisible at collection time, and surfaces only when lengths are recovered from the survey.

The deployed system's error budget bounds this work. It is held to a 15% error budget on length, a design target set with fisheries scientists [6], and every result here is measured against that bar. The deployed system already operates within that budget [6]. Its random error terms are calibration scale (2.4% standard deviation), laser movement within a dive (2.2%), labeling (1.0% at 2 m), and frame count (1.3%), which combine in quadrature to 3.6%. Two further terms only shorten a measurement: body depth (1.5% at 2 m) and pose, the dominant controllable term. Foreshortening scales as $\cos\theta$, so framing within 15° of broadside holds pose to 3.4%. This paper adds one term, the bare flat-port, modeled with Pinax (Section 4).

This added term is small even before correction. Left uncorrected over the reachable region, the bare port costs 1.2% in length at the median and 3.0% at the ninetieth percentile (Section 4). With Pinax applied, the remaining terms this work introduces are the freshwater index assumption, which changes the delivered length by 0.02% at the median in salt water, and the brick calibration, which delivers lengths within 0.03 pp of the checkerboard's at the median (Section 5). Treating even the uncorrected figure as a random term and adding it in quadrature gives 3.8% at the median and 4.7% at the

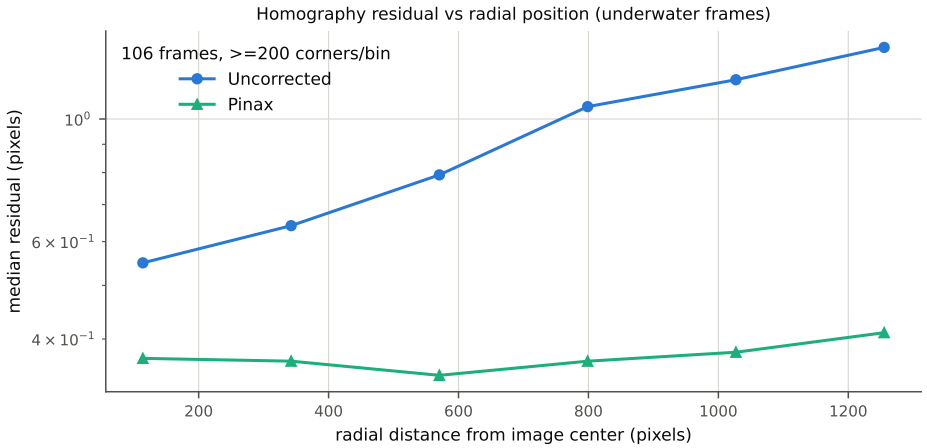



**Figure 2: Homography residual against radial position, over 106 underwater frames of a planar board. An uncorrected pinhole model leaves residual that grows with field angle, the signature an axial camera predicts. Under the analytic correction, the residual is flat across the frame. The test uses only the board's planarity. It requires no known pose, no known scale, and no reference measurement. Residuals are in pixels.**

ninetieth percentile, well inside the 15% budget. Sections 4 and 5 establish these figures.

The rest of the paper asks whether the deployed system's margin can absorb this added source of error.

## 4 Flat Port

Checking whether a camera model accurately describes the optics usually means comparing against something concrete: a known pose, a known scale, or a second instrument. A diver has none of these, and each one brings its own error into the measurement it is meant to validate. The test we use needs none of these baseline measurements.

A plane imaged by a central camera maps to the image by a homography, exactly. Zhang's method rests on this. If the camera is central under the model applied, some homography fits and the residual sits at detection noise. If the camera is axial and the model assumes otherwise, no homography fits, and the misfit is worst where the axial behavior is worst, at large field angles. The residual's size says the model is wrong. Its growth with radius says it is wrong in the way a flat-port is wrong.

Two models are under test. The uncorrected model is an in-air pinhole calibration applied directly underwater, with no account of the interface. The analytic model is Pinax [16], which builds a virtual pinhole from the port's geometry and the refractive indices on either side of it. We take a 6 mm acrylic pane of index 1.49 at a 1 mm standoff. These are nominal, not measured on our housing, and the correction does not depend on them: through a plane-parallel pane the glass terms cancel from Snell's law, so the exit angle is set by the field angle and the water index alone. Varying the pane from 3 to 20 mm, its index from 1.45 to 1.60, and the standoff from 0.5 to 30 mm leaves the median homography residual at 0.363 px throughout. The water index is the parameter that matters, and we assume fresh water at 1.333 as these measurements were taken in a fresh water pool. For each frame, we undistort the detected corners of the checkerboard through the model under test and fit a homography from the board's corner grid to those points. The grid is in meters; a homography absorbs scale, so the units are immaterial. We fit by least squares; a robust fit would discard the corners that carry the signal. What we report is the per-corner reprojection residual.

Under the uncorrected pinhole model, the median residual grows by a factor of 2.4 across the frame, from 0.55 px at the image center to 1.35 px at 1250 px radius (Figure 2). Under the analytic correction it does not grow, holding near 0.37 px throughout and varying by a factor of 1.1 over the same span. The uncorrected model does not reach that floor even on axis, where refraction is near zero. Part of this is the global fit: refitted to the inner bin alone the uncorrected residual falls from 0.55 to 0.39 px. It stays above the corrected model even then.

That floor is consistent with corner-detection scatter: the corrected model reaches it and the uncorrected one does not, so the gap between 0.37 and the uncorrected model's 0.55 to 1.35 px is the port's signature, well above the noise that would otherwise mask it. The brick calibration of Section 5 gives the same result under Pinax, with the residual climbing by a factor of 1.13 across the frame against 1.11 for the checkerboard.

The port is therefore modelable, which leaves the question of what it costs.

The figure usually quoted for an uncorrected flat-port is a free-field bound: the worst length error anywhere in the frame. For a 300 mm fish at 2 m, uncorrected and ranged through the laser, the error reaches 28.7% held horizontal and 36.3% at its worst orientation, both at the frame corner. But the laser has to land on the fish, which rules out most of the frame before any measurement is taken.

The beam sits nearly parallel to the optical axis at a transverse offset of $|O| \approx 104$ mm, so a fish that can be measured is near the axis. The dot itself falls within $\arctan(|O|/Z)$ of the axis, which is

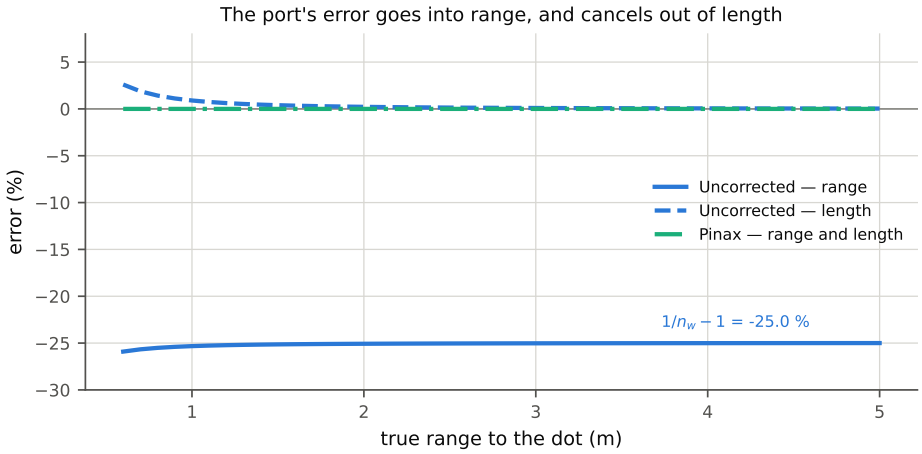



**Figure 3: The flat-port's error enters range and length differently. Through an uncorrected port the triangulated range reads about 25% short at every distance, but the length barely moves: the same refraction that shrinks the range magnifies the apparent extent, and length is their product. Pinax removes both.**

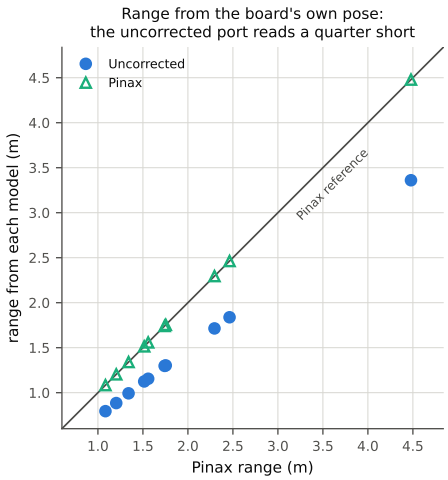



**Figure 4: Range from each camera model on real pool frames. The uncorrected model reads 25.7% short of Pinax at the median, against the 25% that water's refractive index predicts. The port's range error is real and measurable, even though it does not survive into the length.**

5.9° at 1 m and 1.2° at 5 m. On 1,499 pool frames of fish models, 96% of dots land between 0.40 and 0.65 of body length, so a 300 mm fish carrying the dot in that band reaches no further than 16.65° at 1 m, or 3.42° at 5 m.

This is the cancellation of Section 3, and the flat-port is its first instance. Figure 3 shows it directly: the uncorrected range runs about 25% short across the whole working distance, while the length it produces stays near zero error. The port's error is almost entirely common to range and angular extent, so it leaves the product nearly untouched.

The short range is not a simulation artifact. Figure 4 shows it on real pool frames: the uncorrected model reads consistently short of Pinax, by the margin water's index predicts. The error is there in the range; it simply does not reach the length.

What survives the cancellation is small but not zero, and it is worst near the frame edge. Figure 5 shows the residual length error against field angle, with the reachable region shaded. Over that region, at the worst bearing for each fish, it is 1.2% at the median and 3.0% at the ninetieth percentile: a harder case than the held-horizontal figure, and an order of magnitude below the free-field maximum. The rig cannot take the measurement that maximum describes.

The three curves in Figure 5 are three lenses spanning more than a twofold change in focal length, and they are indistinguishable. The error depends on field angle alone, so the correction is a property of the port and carries to other lenses behind a flat pane.

## 5 Calibration

With the port modeled, the camera can be calibrated in air, and the target no longer has to be a printed plane. We calibrate the intrinsics from a structure of commodity LEGO™ bricks, molded to a published pitch that needs no measurement. The laser extrinsics already come from a dive slate every dive. The camera intrinsics now come from bricks at home, and no specialized manufactured object is needed.

The detector finds brick faces in each frame and matches them to the known model (Figure 6). From 21 views we fit the intrinsics and one radial distortion term from a generic starting point, using nothing from a board; higher-order terms are not stable on this capture. Against a checkerboard calibration of the same camera, the brick result differs by 2.46% in focal length and 112 px in the principal point.

The two parameters are not equally well pinned. Absolute focal length trades against target distance: a larger object farther away images the same, and a shallow object shot over a narrow range of

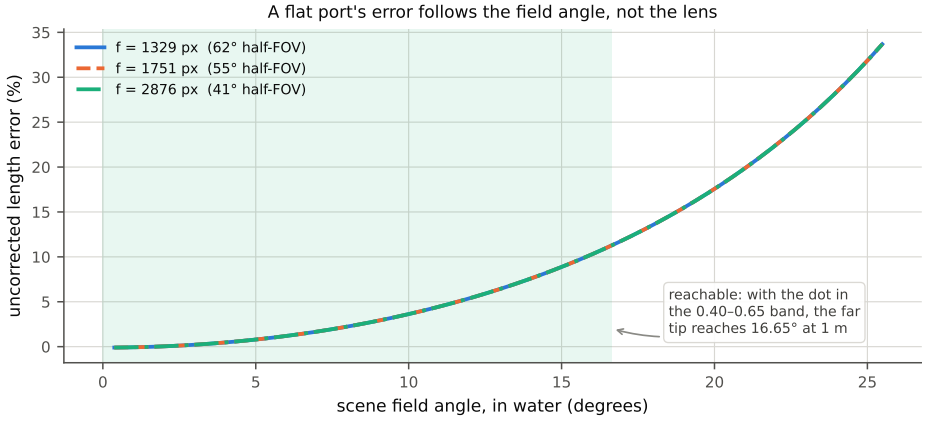



**Figure 5: Uncorrected length error against field angle, for three lenses spanning 41° to 62° of half-field. The curves are indistinguishable: the error depends on field angle alone. The shaded region is what the rig can reach, given that the laser dot must land on the fish at a transverse offset of 104 mm and between 0.40 and 0.65 of body length; the bound is set by 0.65, which leaves the most body extending away from the axis. Over the reachable region the uncorrected error is 1.2% at the median and 3.0% at the ninetieth percentile, against 36.3% at the frame corner.**

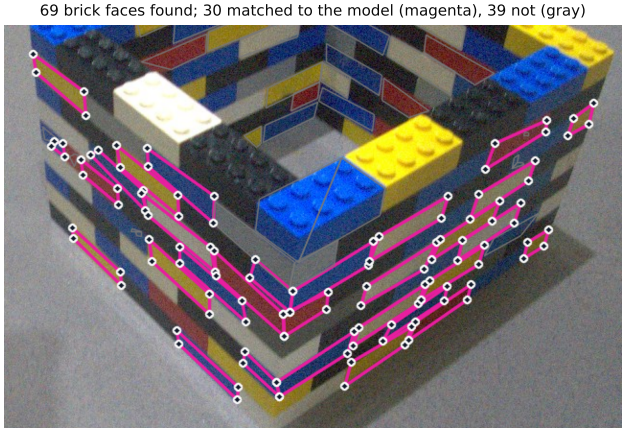



**Figure 6: The detector on one raw frame of the brick target: 69 brick faces found, 30 matched to the model. Matched faces supply the image-to-model correspondences the calibration is fitted from.**

poses does not break that trade strongly. Across random subsets of the views the focal length varies with a standard deviation of 2.6%, and removing any single view shifts it by −1.4 to +1.0 percentage points (Figure 7). That reflects the capture. A checkerboard cannot recover absolute scale from its images either; it takes scale from the known square size, the way the bricks take it from the known pitch. The focal-length ratio does not trade against distance, and its standard deviation across the same subsets is 0.33%.

The ratio is better conditioned than the scale, but this capture does not fully observe it either. The two calibrations disagree on it by 0.75%, and the disagreement lies almost entirely in a direction the images have no leverage on. Scaling the target model along its depth axis by ±1.5% moves the recovered ratio across the whole board-to-brick span while the fit residual stays flat: the data cannot distinguish those solutions. The brick pitch itself, fixed at 9.6 mm, is not in question, but with none of the views rolled and their poses spanning only about 15°, the anisotropy that separates $f_x$ from $f_y$ is barely seen. This is why the ratio, though far tighter than the scale across subsets, cannot yet be trusted over the checkerboard's on this data.

The delivered length does not inherit the scatter, because a focal error cancels. The laser is fitted through the same intrinsics the length uses, so a scale error shrinks the fitted range by the factor it inflates the angular extent, and length is their product. We can see how much cancels by decomposing the checkerboard-to-brick difference and applying each part on its own to the same frames. The 2.46% scale error costs 0.018 pp. Anisotropy and centering do not cancel, at 0.129 and 0.159 pp applied alone with the checkerboard's distortion. The brick calibration as a whole delivers 0.027 pp at the median; the parts measure larger because the checkerboard's higher-order distortion terms do not suit the brick intrinsics.

A single decoy frame shows the cancellation directly. In one pool frame of the calipered 312.5 mm target, the laser dot falls at pixel (1944, 1070). Back-projected through the Pinax-corrected brick calibration and triangulated against the per-dive beam, it gives a range of $Z = 1.085$ m. The snout and tail-fork, labeled at (2229, 1325) and (1292, 751), are back-projected through the same calibration and placed at that depth; the distance between them is 308.0 mm, within 1.5% of the caliper. Through the checkerboard calibration the range is 1.113 m, 2.5% longer, and the length is the same 308.0 mm, because the range and the angular extent change together.

Figure 8 shows this end to end. Run through the same correction on the same twelve pool frames, the brick and checkerboard calibrations deliver median lengths of −2.66% and −2.67%, a difference of 0.03 pp. A calibration that disagrees with the checkerboard by

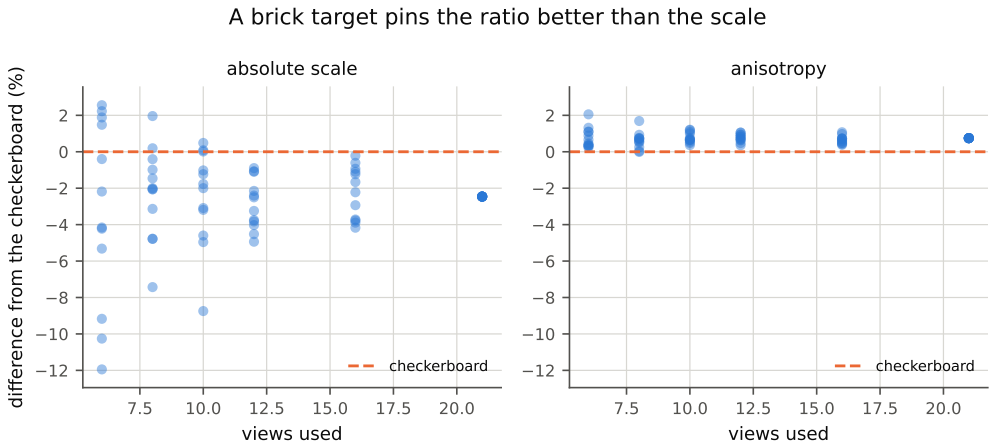



**Figure 7: Calibrated parameters across random subsets of the views, each plotted as its difference from the checkerboard calibration. Left, the absolute focal length varies with a standard deviation of 2.6% across subsets (−12.0 to +2.6%); right, the focal-length ratio $f_x/f_y$ varies by 0.33% (−0.01 to +2.05%). The two are recovered from the same images, so the difference is one of conditioning: absolute scale trades against target distance and is weakly pinned by a shallow object over a narrow range of poses, while the ratio does not trade against distance and is fixed about eight times more tightly.**

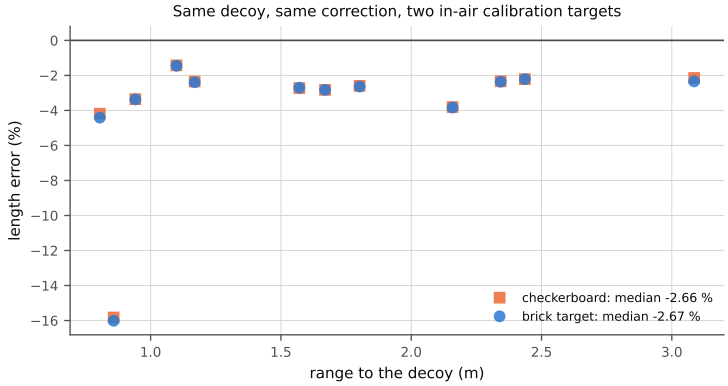



**Figure 8: Delivered length from the brick and checkerboard calibrations, run through the same correction on the same 12 frames of the calipered 312.5 mm decoy, with the laser beam fitted on the pool checkerboard session for both. The calibrations differ by 2.46% in focal length and 0.75% in $f_x/f_y$, yet the lengths they deliver differ by 0.03 pp at the median and 0.23 pp at worst (Bland–Altman 95% limits of agreement −0.23 to +0.11 pp).**

2.46% in focal length and 112 px in the principal point delivers, for this measurement, the same answer. The intrinsics the brick target gets wrong are the ones the measurement does not depend on.

## 6 Limitations

This capture leaves two questions open. The first is the focal-length ratio: the brick and checkerboard calibrations disagree on it by 0.75%, and, as Section 5 shows, this capture barely constrains the direction that disagreement lies in, so we cannot yet say which is closer to the sensor's true ratio. Neither possibility reaches the delivered length.

The absolute focal length is weakly conditioned for the same reason: it trades against target distance, and this capture is shallow and narrow in pose. The known brick pitch supplies scale regardless, as a board's known square size does, so the length is unaffected. The end-to-end agreement in Section 5 is measured on one decoy at ranges from 0.80 to 3.04 m, where the brick and checkerboard lengths agree to within −0.23 to +0.11 pp (95% limits of agreement). The scale part of the calibration difference nearly cancels at any target size or range, since it shrinks the fitted range by the factor it inflates the angular extent. The anisotropy and principal-point parts do not cancel, and their effect depends on where the fish sits in the frame; this decoy samples field angles up to 19.8°.

Our evidence is one camera in one housing, calibrated against one build of the target and validated on pool frames. The water index is assumed fresh, matching the pool; the deployment corpus

is salt water, where the index is marginally higher. In simulation, with laser ranging, using the fresh index in salt water changes the delivered length by 0.02% at the median and 0.2% at most over the reachable region.

## 7 Conclusion

FishSense Lite measures fish from a single photograph, using gear a recreational diver already owns. The one part that did not fit was calibration: the camera intrinsics came from a manufactured checkerboard, photographed underwater in a pool. This paper removes that step by modeling the flat-port analytically, which lets the camera be calibrated in air, where the target can be built from commodity bricks. The system's calibrations now run on objects a diver has: bricks at home, a dive slate at the dive site.

Taking the corrective optic off costs little. The analytic model describes the bare port well: its homography residual stays flat across the frame, while the uncorrected model's grows toward the edge. And over the region a laser–camera rig can actually measure, the uncorrected error runs 1.2% at the median, an order of magnitude under the free-field maximum. It is small because the port distorts range and angular extent together, and length is their product, so most of the error divides out.

The LEGO™ brick calibration does not match the checkerboard exactly. It differs by 2.46% in focal length and 112 px in the principal point, and its focal-length ratio, though far better conditioned than its absolute scale, we do not yet trust over the checkerboard's. None of this changes the length the system reports. The same cancellation that bounds the port's cost bounds this one: a focal error fitted through the beam's own geometry nearly divides out, and the two calibrations deliver lengths that differ by 0.03 pp at the median.

The correction depends on the flat-port alone, so it carries to other lenses behind a flat pane. Two things it does not remove: the laser calibration still drifts and is re-fitted each dive, and the water's refractive index is assumed fresh, which in simulation changes the delivered length by 0.02% at the median in salt water. Both aside, every calibration the system needs now runs on commodity objects.

## Acknowledgments

This work was funded in part by the Research Experiences for Undergraduates Site Engineers for Exploration at UC San Diego (NSF 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.

The authors used generative AI tools in preparing this paper. Anthropic's Claude helped draft and revise the text, including the abstract and every section, and reviewed drafts for consistency of numbers and claims. Claude Code assisted in writing analysis scripts. Claude and Google's Gemini suggested references during the literature search, and the authors checked each cited source. The authors reviewed all text, code, figures, and numbers, and take full responsibility for the content of this paper.

## References

[1] 2024. *2024 Diving Fast Facts.* Technical Report. Diving Equipment and Marketing Association. https://www.dema.org/store/download.aspx?id=7811b097-8882-4707-a160-f999b49614b6

[2] Amit Agrawal, Srikumar Ramalingam, Yuichi Taguchi, and Visesh Chari. 2012. A theory of multi-layer flat refractive geometry. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. 3346–3353. doi:10.1109/CVPR.2012.6248073

[3] 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 (sep 2008), 221–231. doi:10.1017/s0376892908005043

[4] Backscatter. [n. d.]. Backscatter M52 Wide Angle Air Lens Review. https://www.backscatter.com/reviews/post/Best-Olympus-TG-5-Underwater-Lens-Backscatter-M52-Wide-Angle-Air-Lens-Review. Accessed 2026-09-21.

[5] Luca Baronti, Matteo Dellepiane, and Roberto Scopigno. 2010. Using Lego pieces for camera calibration: a preliminary study.. In *Eurographics (Short Papers)*. 97–100.

[6] Christopher Crutchfield, Kyle Hu, Nathan Hui, Omair Qazi, Jennifer Loch, Allison Candelmo, Brice Semmens, Curt Schurgers, and Ryan Kastner. 2026. FishSense Lite: A Camera-Based, Single Laser Hardware Framework with Software System for In Situ Fish Length Measurement. doi:10.5281/zenodo.22882398

[7] Graham J Edgar and Rick D Stuart-Smith. 2014. Systematic global assessment of reef fish communities by the Reef Life Survey program. *Scientific data* 1, 1 (2014), 140007.

[8] foamcoreprint. 2026. Checkerboard Calibration Targets. https://www.foamcoreprint.com/checkerboard.

[9] World Wildlife Fund. 2026. Sustainable seafood. https://www.worldwildlife.org/our-work/oceans/sustainable-seafood/.

[10] Xavier Grane-Feliu, Scott Bennett, Bernat Hereu, Eneko Aspillaga, and Julia Santana-Garcon. 2019. Comparison of diver operated stereo-video and visual census to assess targeted fish species in Mediterranean marine protected areas. *Journal of Experimental Marine Biology and Ecology* 520 (2019), 151205.

[11] 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

[12] Euan Harvey, David Fletcher, and Mark Shortis. 2001. Improving the statistical power of length estimates of reef fish: a comparison of estimates determined visually by divers with estimates produced by a stereo-video system. *Fishery bulletin-national oceanic and atmospheric administration* 99, 1 (2001), 72–80.

[13] 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 (2008), 1286–1296.

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

[15] Peter T Kuriyama, Trevor A Branch, Allan C Hicks, John H Harms, and Owen S Hamel. 2019. Investigating three sources of bias in hook-and-line surveys: survey design, gear saturation, and multispecies interactions. *Canadian Journal of Fisheries and Aquatic Sciences* 76, 2 (2019), 192–207.

[16] Tomasz Łuczyński, Max Pfingsthorn, and Andreas Birk. 2017. The Pinax-model for accurate and efficient refraction correction of underwater cameras in flat-pane housings. *Ocean Engineering* 133 (mar 2017), 9–22. doi:10.1016/j.oceaneng.2017.01.029

[17] Victoria Y. Martin, Leslie Christidis, David J. Lloyd, and Gretta 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

[18] D Parker, H Winker, ATF Bernard, ER Heyns-Veale, TJ Langlois, ES Harvey, and A Götz. 2016. Insights from baited video sampling of temperate reef fishes: how biased are angling surveys? *Fisheries Research* 179 (2016), 191–201.

[19] David Robertson and Bill Breen. 2013. *Brick by brick: How LEGO rewrote the rules of innovation and conquered the global toy industry.* Crown Currency.

[20] E Schmitt, R Sluka, and K Sullivan-Sealey. 2002. Evaluating the use of roving diver and transect surveys to assess the coral reef fish assemblage off southeastern Hispaniola. *Coral Reefs* 21, 2 (2002), 216–223.

[21] Peter Sturm. 2021. Pinhole camera model. In *Computer Vision: A Reference Guide*. Springer, 983–986.

[22] 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

[23] Noëlle Yochum, Richard M Starr, and Dean E Wendt. 2011. Utilizing fishermen knowledge and expertise: keys to success for collaborative fisheries research. *Fisheries* 36, 12 (2011), 593–605.

[24] Zhengyou Zhang. 2000. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence* 22, 11 (2000), 1330–1334.