

Article

Design and Open-Water Validation of an Autonomous Underwater Robot for Physical Diver Safety Assistance

Min-Gyu Kim ^{1,*}, Hyungjoo Kang ¹, Hansol Jin ¹, Ji-Hong Li ¹, Gun Rae Cho ¹, Gyeongmin Kim ² and Minshik Noh ²

¹ Marine Robotics Research Group, Robotics Application Research Division, Korea Institute of Robotics and Technology Convergence (KIRO), Pohang 37666, Republic of Korea

² KYEONG IN TECH Co., Ltd., Changwon 51395, Republic of Korea

* Correspondence: zxdwa0817@kro.re.kr

Abstract

Diving accidents remain frequently fatal, and most stem from preventable human factors: panic, failed buoyancy control, and disorientation. Underwater human–robot interaction research has focused on observing the diver and offers little help when a diver must be pulled to safety. This paper presents the design and sea validation of NADIA, an autonomous underwater robot that combines diver guidance and following with diver-commanded physical assistance. A rear handle converts the diver’s grip into thrust and springs back to neutral on release; a diver-facing display shows depth, distance, battery state, and warnings. Six vector thrusters, with mounting angles optimized by gradient descent against drag-model thrust-ratio targets, are driven by an operational-space time-delay controller (TDC-OSF) that keeps the robot in the diver’s view; stereo vision and imaging sonar are fused for diver detection. In test-bed, tank, and sea trials, surge and heave thrusts matched design targets within 5–7%; approval-gated guidance held the robot 1–2 m from the diver with position and heading errors below 0.15 m and 3.5°; and the robot towed a diver at about 1 m/s and rose from 5 m to the surface in about 4 s. An underwater robot can thus serve as a physical safety layer for divers.

Keywords: unmanned underwater vehicle; diver safety; underwater human–robot interaction; thrust allocation; time-delay control; sensor fusion; open-water experiment



Academic Editors: Oren Gal,
Alessandro Bucci and
Benedetto Allotta

Received: 11 September 2026

Revised: 23 September 2026

Accepted: 25 September 2026

Published: 2 October 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

1.1. Background and Motivation

Scuba diving is no longer just a water sport. It has grown into a sizeable industry built on direct contact with marine ecosystems, and diving work now ranges from underwater archaeology and marine biology to industrial tasks such as pipeline and aquaculture maintenance [1–3]. The environment itself, however, remains hostile to humans. Water is about 800 times denser than air, hydrostatic pressure rises by roughly 1 atm for every 10 m of depth, and visible light attenuates and scatters so quickly that visibility often shrinks to a few meters.

Underwater, a diver survives on a small set of life-support devices: an air tank, first- and second-stage regulators, and a buoyancy control device (BCD). There is no stepping out of the water when something goes wrong, so a minor equipment fault or a misread current can escalate quickly into cognitive impairment from nitrogen narcosis, decompression sickness, or pulmonary barotrauma caused by a panicked rush to the surface.

Among 310 diving accidents reported in Korea between 2000 and 2013, more than 90% ended in death, and the analyses point not to mechanical failure of the equipment but to preventable human factors: mistakes, insufficient training, panic, failed buoyancy control, and disorientation [4]. More than half of diving fatalities are also reported to occur while a diver is separated from a companion [5]. The diving community has long enforced the buddy system as its basic answer to this risk. In rough water, though, poor visibility, the practical impossibility of voice communication, and down-currents or counter-currents that pull divers apart all erode the mutual-rescue value that the buddy system is supposed to provide.

1.2. Related Work on Underwater Robots and UHRI Systems

Intelligent underwater robots have naturally been proposed as a way to compensate for these human limits [6]. Remotely operated vehicles (ROVs), the most mature class commercially, receive abundant power and return high-definition video through a tether, which suits them to precise underwater work and deep-sea exploration. The same tether, however, disqualifies them from close diver support. Drag on the cable keeps an ROV from following a diver who swims freely in three dimensions, and the cable itself can snag on the diver's equipment or on reefs and wrecks, creating a secondary hazard [7]. Autonomous underwater vehicles (AUVs) dispense with the tether and excel at wide-area survey, but their perception and control architectures are built around pre-programmed waypoint tracks; following an unpredictably moving person in real time, and reacting to that person, lies outside their design [8]. At the other end of the spectrum, commercial diver propulsion vehicles (DPVs) do tow divers, but they are hand-held, manually triggered machines with no perception of the diver, no autonomy, and no link to the surface; they extend a diver's range rather than watch over the diver, and their design literature is accordingly concerned with propulsion and hydrodynamics rather than safety intervention [9].

The EU-funded CADDY (Cognitive Autonomous Diving Buddy) project remains the most comprehensive attempt to date at underwater human–robot interaction (UHRI) at the system scale. An autonomous surface vehicle relayed GPS and acoustic signals while an AUV tracked the diver acoustically and read hand gestures with a stereo camera, taking photographs and guiding the diver on command [10–12]. Yet CADDY stayed in the role of an observer that follows the diver visually and executes passive commands. Its main interface, vision-based gesture recognition, degrades quickly under the turbidity, changing illumination, and particulate noise of real underwater scenes, and the interaction breaks off entirely once the diver drifts out of the camera's view [13,14].

Later work, surveyed in [15,16], has pushed underwater perception considerably further: a multimodal visual–textual gesture-recognition model [17], a large diver-gesture dataset collected with smart gloves [18], gesture classification with explainable AI [19], intent estimation from head motion and throat vibration [20], and diver identification from anthropometric ratios in turbid water [21]. Vision-based diver tracking and person-following by mobile robots have been studied extensively as well [22,23], visual–acoustic sensor fusion has been used to track a diver in open water [24], and, most recently, a robot has been used to estimate a diver's respiration rate as a safety cue [25]. Outside the water, the picture is very different: contactless gesture interfaces have become highly reliable, and recent hybrid deep-learning pipelines report near-perfect recognition for contactless control even under poor lighting and cluttered backgrounds [26]. The contrast is instructive. What makes those interfaces work in air—an unobstructed optical path and stable illumination—is exactly what turbid water takes away, and it is taken away before any classifier is reached. We therefore treat optical perception as one channel among several rather than as the interaction channel itself. On the platform side, researchers have built

robots for aquaculture-net inspection [27], pipeline inspection [28], and diver-replacing inspection of underwater infrastructure [29], along with a cable-driven rescue concept [30] and our own earlier concept design with a diver grip ring [31]. The last two, however, never went beyond CAD models or fragmentary tests in small tanks, and none of these efforts matured into an integrated platform that can hold its own against currents at sea.

Taken together, prior work leaves three gaps on the way to actual diver safety. Most technical progress has gone into individual perception algorithms, with no integrated hardware–dynamics–software system that specifies what the robot should physically do once a hazard is recognized. Optical sensing still dominates, so the systems tend to fail precisely in the turbid conditions where accidents concentrate. Above all, existing autonomous systems only watch, and the machines that do provide thrust do not perceive the diver at all. When a diver is out of air, has lost a regulator, or is being dragged down by a current, what that diver needs is not a photograph or a visual cue but thrust: a machine that can tow, or at least something solid to hold on to on the way up.

1.3. Contributions and Organization

This paper presents NADIA (Nautical Advanced Diving Intelligent Assistance system), an autonomous underwater robot built to provide exactly this kind of physical support, and reports its validation in open water. The aim of the study is to shorten the interval between a diver getting into difficulty and receiving physical assistance, and to keep the robot inside the diver's field of view for as much of a guided dive as possible, while holding the platform light enough for two people to deploy from a small boat. The quantitative design targets were carried over from the concept study [31]: a traction force above 150 N for towing, a diver-localization error below 0.2 m, and an endurance beyond 2 h. We have published parts of the work separately: the support scenarios and concept design [31], the gradient-descent optimization of the thruster arrangement [32], the underlying time-delay tracking control [33], and its operational-space reformulation for diver guidance [34]. What this paper adds is the integration of those elements into a single operational platform and, to the best of our knowledge, the first open-water demonstration of an autonomous diver-support robot that closes the loop from diver perception and autonomous guidance to physical intervention, including diver-commanded towing and emergency ascent. The main features are the following:

- A physical towing interface coupled with an intuitive human–machine interface (HMI). A rear control handle converts the diver's grip displacement directly into fore/aft thrust and springs back to neutral when released, while a diver-facing display presents depth, robot–diver distance, battery state, and emergency warnings. Together, they form a practical underwater human-in-the-loop.
- Scenario-driven hybrid autonomy on an over-actuated propulsion layout. The pitch and yaw mounting angles of six vector thrusters were optimized by gradient descent against thrust-ratio targets derived from a drag model, and the operational-space-formulated time-delay controller (TDC-OSF) keeps the robot inside the diver's visual space under disturbances.
- Multi-sensor diver perception. Stereo vision and imaging sonar are fused with deep-learning detection and pose estimation, so the robot can find and track the diver across visibility conditions without attaching any sensor to the diver [31].
- The rest of the paper is organized as follows. Section 2 derives the operational requirements and describes the platform architecture. Section 3 covers the structural analysis and the thruster-angle optimization. Section 4 presents the control and perception algorithms, and Section 5 the test-bed, water-tank, and open-water experiments. Section 6 discusses the results and their limitations, and Section 7 concludes.

2. System Design and Hybrid Architecture

The sea combines high hydrostatic pressure, changing currents, and the loss of radio communication. A robot that is supposed to guard a diver must therefore offer more than watertightness; it needs dynamic stability and predictable fail-safe behavior when something breaks [7]. This section derives the requirements imposed by physical interaction and describes the modular architecture we adopted to meet them.

2.1. Operational Requirements Based on Physical Interaction

We derived the operational requirements from a literature survey, consultations with professional divers, and our own diving experience [31]. Around the three basic safety rules of scuba diving—keep a safe travel speed, keep breathing, keep your buddy—we defined 14 robot-support scenarios in four categories: dive-information provision, dive guidance, safety monitoring, and dive assistance [31]. In these scenarios, the robot must react immediately to two emergencies, diver separation and an abnormal diver state such as panic or unconsciousness. Three design requirements follow. The robot needs enough structural strength and physical-interaction capability to tow a diver against gravity and current drag during an emergency ascent. It needs six-degree-of-freedom (6-DOF) maneuverability to stay with a swimming diver. And it needs a heterogeneous optical–acoustic sensor suite for turbid water, combined with an HMI that a diver can still read when visibility is poor.

2.2. Modular Hardware Structure and Passive Restoring Design

NADIA consists of three independent pressure vessels (control, power, and perception modules) on a main frame of AL6061-T6. The modular layout keeps the robot light enough for shipboard launch, about 85 kg in air, and allows each module to be maintained on its own. Static stability comes from the separation between the center of buoyancy (CoB) and the center of mass (CoM): the heavy battery packs sit low on both sides to pull the CoM down, and syntactic-foam buoyancy blocks inside the upper frame push the CoB roughly 80 mm above it, so the hull generates a strong passive restoring moment against roll and pitch disturbances. We ballasted the platform to a net positive buoyancy of about 29 N. If the control system blacks out, the robot floats to the surface instead of sinking. Table 1 lists the general specifications.

Table 1. General specifications of the diver-safety-support underwater robot.

Item	Specification
Weight (in air)	85 kg
Net buoyancy (in water)	+29 N (positive)
Dimensions (L × W × H)	1085 mm × 860 mm × 410 mm
Depth rating	50 m
Thrusters	6 vector thrusters, 105 N forward each
Power	2.4 kW
Endurance	>2 h

2.3. Core Module Design and Integration

Communication and power distribution. Internally, a high-performance single-board computer (NVIDIA Jetson Orin class; NVIDIA Corporation, Santa Clara, CA, USA) serves as the main controller above a set of embedded microcontrollers. The main controller runs the AI vision processing and the navigation, and delivers thrust commands to the motor drivers over a CAN bus at 100 Hz or more, so a bottleneck in video processing cannot disturb the motion-control loop. For positioning and communication where radio waves cannot reach, the robot carries an ultra-short-baseline (USBL) transponder and an acoustic telemetry modem (ATM), both S2C T 42/65 units (EvoLogics GmbH, Berlin, Germany),

linked to a shipboard transceiver hub; at the surface, GPS, Wi-Fi, and RF modules provide position fixes, data backup, and remote access [31]. The platform operates in a hybrid mode: a tether-free AUV mode is the default for diver support, since a cable near a diver is itself a hazard, while a tethered ROV mode with live video is available for a pre-dive survey of the site (turbidity, temperature, and current) before divers enter the water [31].

Multi-sensor perception module. A stereo camera (a pair of Firefly DL machine-vision cameras, Teledyne FLIR LLC, Wilsonville, OR, USA, fitted in a purpose-built housing) sits in a purpose-built rectangular pressure vessel at the front, with a high-output LED and an ultrasonic imaging sonar (Oculus M750d, Blueprint Subsea, Ulverston, UK) mounted on the upper exterior. In clear water, the stereo pair extracts the diver's three-dimensional skeleton and pose at high resolution; when turbidity kills the optical view, the sonar keeps tracking the diver's range and outline, so the two sensors cover each other's blind spots. For navigation, dead reckoning that combines a Doppler velocity log (DVL; Pathfinder 600 kHz, Teledyne RD Instruments, Poway, CA, USA), a precision depth sensor (PAA-36XiW, KELLER AG, Winterthur, Switzerland), and a 9-axis IMU (FI200P, Fiberpro Inc., Daejeon, Republic of Korea) estimates the robot's three-dimensional position where GPS is unavailable [35] (Figure 1).




Figure 1. Configuration and coordinate system of the diver-safety-support underwater robot. USBL, ultra-short-baseline positioning transponder; DCS, digital compass system; DVL, Doppler velocity log.

Physical interaction interface and intuitive HMI. The most distinctive hardware element is the coupled physical–visual interface. The control handle at the rear houses a pressure-sealed potentiometer and a torsion-spring return mechanism. When the diver twists the handle, the displacement becomes a fore/aft thrust command at once; when the diver lets go, the spring snaps the handle back to neutral and the robot stops, so a released handle cannot leave a standing thrust command. The handle is therefore more than a grab bar. It lets the diver drive the robot as a towing vehicle, which is what makes the underwater human-in-the-loop practical (Figures 2 and 3). A 7-inch display, driven by an auxiliary computing unit (LattePanda; DFRobot, Shanghai, China), is mounted on the upper body behind a waterproof acrylic window and tilted 15° toward the diver's line of sight; it shows depth, robot–diver distance, battery level, and warnings such as an ascent advisory in a high-contrast UI: large white numerals on a dark ground, with the robot–diver distance given the largest field because that is the quantity a diver checks most often, and warnings raised as a full-width color banner so that they stay legible at arm's

length through a mask in turbid water. A 50 W underwater speaker, developed in-house, behind a polycarbonate cover adds an auditory channel—in a sea test at 22 m depth, its alert sounds were still audible about 18 m away.




Figure 2. Control handle: cross-section showing the underwater connector, housing, potentiometer, torsion-spring return mechanism and knob. Colors distinguish the components of the CAD model: green, potentiometer board; red and blue, torsion springs; gray, housing and knob.




Figure 3. Manufactured control handle: (a) assembled handle with its mounting bracket; (b) the potentiometer and torsion-spring return assemblies before installation.

3. Structural Analysis and Thruster-Arrangement Optimization

The platform has to survive hydrostatic pressure at its rated depth and, at the same time, the unpredictable loads that come with a diver physically hanging on to it. This section verifies the structural design by finite-element analysis (FEA) and then develops the drag modeling and thrust-allocation optimization for the over-actuated vector propulsion system.

3.1. Structural Robustness of Pressure Vessels and Frame (FEA)

A cracked or even slightly deformed structure floods the electronics and disables the robot. We therefore analyzed the main pressure structures by static FEA in Autodesk Inventor (Autodesk Inc., San Francisco, CA, USA) under a uniform external pressure of 0.5 MPa (5 bar), corresponding to the rated depth of 50 m [36]. The vessels were meshed with linear tetrahedral elements, the cover–housing interfaces were modeled as bonded and separating contacts, and the wall thickness was 1.5 mm for the cylindrical vessels and 8 mm for the side walls of the rectangular perception vessel. The mesh was then refined adaptively until the peak stress changed by less than 5% between successive refinements, a criterion met at 3.2% for the perception vessel and 4.5% for the frame. The two cylindrical vessels are the exception: their stress maxima sit on the bonded flange–endcap edge, where

the linear-elastic solution is singular, so the peak keeps rising with refinement while the displacement field, which is not singular, has converged. With safety factors above 6, this local mesh sensitivity does not affect the design conclusion. Failure of the ductile vessel material (aluminum 6061-T6) was assessed with the von Mises equivalent stress computed from the three principal stresses (σ_1 , σ_2 , σ_3):

$$\sigma_v = \sqrt{\frac{1}{2}[(\sigma_1 - \sigma_2)^2 + (\sigma_2 - \sigma_3)^2 + (\sigma_3 - \sigma_1)^2]} \quad (1)$$

where σ_v is the von Mises equivalent stress and σ_1 , σ_2 and σ_3 are the principal stresses. The safety factor (SF) is the ratio of the material yield strength to the maximum equivalent stress, $SF = \sigma_{\text{yield}}/\sigma_{v,\text{max}}$, with the 6061-T6 yield strength taken from the solver's material library (approximately 275 MPa). The stresses quoted below are rounded, and the safety factors are those reported by the solver.

The main vessel, which holds the power-distribution and control boards, and the power vessel, which holds the batteries, are cylindrical shell–endcap structures, the shape that spreads hydrostatic load best. Their maximum von Mises stresses came to 30 MPa and 43 MPa, with maximum displacements of 0.05 mm and 0.08 mm, for safety factors of 9.0 and 6.4 (Figures 4 and 5). The front perception vessel is different. To give the stereo camera, the imaging sonar, and the dedicated AI computing board a wide, unobstructed field of view and to pack the boards efficiently, we made it rectangular, and a rectangular vessel concentrates stress far more than a cylinder, so its strength analysis mattered. Under 5 bar, stress concentrated at the corners where the acrylic viewport joins the frame and around the O-ring seats, but it peaked at 160 MPa, with a maximum displacement of 0.45 mm, a safety factor of 1.7 (Figure 6). This is the smallest margin of the four analyses. The peak is a local concentration at an internal corner rather than a stress level carried by the seal face, the sub-millimeter deformation is too small to tilt the optical axis beyond tolerance, and the vessel remains watertight at rated depth; it is nevertheless the component with the least structural reserve and the first one to revisit if the depth rating is extended. We also simulated the main frame under launch-and-recovery (LARS) loading, lifting the full 85 kg from the four upper lifting eyes. The maximum stress was about 31 MPa, a safety factor of 8.92, which means the frame can also absorb the local twisting loads of a diver hanging on the rear handle and impacts from currents (Figure 7). Table 2 collects the numerical results of the four analyses.

Table 2. Summary of the finite-element analyses (Al 6061-T6, yield strength 275 MPa).

Component	Load Case	Wall Thickness (mm)	Nodes/Elements	Max. Von Mises Stress (MPa)	Max. Displacement (mm)	Safety Factor
Main pressure vessel (cylindrical shell-endcap)	0.5 MPa external	1.5	459,230/268,189	30.4	0.051	9.03
Power pressure vessel (cylindrical shell-endcap)	0.5 MPa external	1.5	249,111/129,543	42.9	0.078	6.40
Perception-module vessel (rectangular)	0.5 MPa external	8.0	51,321/27,307	159.8	0.445	1.72
Main frame	85 kg, four-point LARS lift	-	66,555/33,315	30.8	0.109	8.92




Figure 4. FEA of the main pressure vessel under 0.5 MPa external pressure: (a) mesh; (b) von Mises stress; (c) displacement; (d) safety factor.




Figure 5. FEA of the power (battery) pressure vessel under 0.5 MPa external pressure: (a) mesh; (b) von Mises stress; (c) displacement; (d) safety factor.




Figure 6. FEA of the perception-module pressure vessel under 0.5 MPa external pressure: (a) mesh; (b) von Mises stress; (c) displacement; (d) safety factor.




Figure 7. FEA of the main frame under the 85 kg four-point launch-and-recovery lift: (a) mesh; (b) von Mises stress; (c) displacement; (d) safety factor.

3.2. Target-Thrust Modeling Based on Hydrodynamic Drag

Before optimizing the thruster arrangement, we need target speeds, and from them required thrusts, that reflect the mission. For diver guidance and rescue, we set target speeds of $V_x = 1.5$ m/s in surge, $V_y = 0.5$ m/s in sway, and $V_z = 0.5$ m/s in heave [32]. The thrust needed to reach each speed follows from the quadratic drag equation [37]:

$$F_D = \frac{1}{2} C_D \rho v^2 A, \quad (2)$$

where F_D is the drag force, C_D the direction-wise drag coefficient, ρ the water density, v the relative speed of the robot, and A the projected area normal to the direction of travel. Subscripts x , y and z denote the surge, sway and heave directions, so that $C_{D,x}$, $C_{D,y}$ and $C_{D,z}$ are the coefficients quoted below.

We computed the direction-wise drag coefficients by computational fluid dynamics (CFD) in ANSYS Fluent 2024 (Ansys Inc., Canonsburg, PA, USA) [32]. For preprocessing, the internal structures were removed from the computational domain, every opening in the geometry was sealed to form a single solid body, and the thruster blades were left out because they rotate in operation. The surrounding fluid domain was made large enough to resolve the flow around the hull and to keep the solution stable; the discretized domain contained 228,761 elements and 970,265 nodes. We ran steady-state simulations with the shear stress transport (SST) turbulence model [38] to capture near-wall behavior, set the inlet flow perpendicular to the inlet face at 1.5 m/s and the outlet at atmospheric pressure (0 Pa gauge), and swapped the inlet and outlet positions to evaluate each of the robot's motion directions (X , Y , and Z) in turn.

With a freshwater density of 1000 kg/m^3 , the CFD gave forces of 234, 436, and 1116 N at 1.5 m/s on projected areas of $222,472 \text{ mm}^2$ (front), $322,121 \text{ mm}^2$ (side), and $594,006 \text{ mm}^2$ (top), corresponding to drag coefficients of $C_{D,x} = 0.93$, $C_{D,y} = 1.20$, and $C_{D,z} = 1.67$ [32]. A cross-check of the same model in two commercial CFD codes showed that the absolute force levels depend on the solver, so we used the CFD mainly for the non-dimensional coefficients and their ratios and left the physical validation of the targets to the load-cell experiments in Section 5.1. Substituting the coefficients, a seawater density of 1024 kg/m^3 , and the mission speeds into Equation (2) gives required thrusts of 238.6, 49.5, and 127.1 N in surge, sway, and heave, and hence a reference target thrust ratio of $F_x:F_y:F_z = 4.83:1:2.57$ [32]. To limit form drag, we used curved and cylindrical vessels wherever possible and faired the frame corners and thruster struts with streamlined covers.

3.3. Jacobian- and Gradient-Descent-Based Angle Optimization of the Over-Actuated System

To realize this ratio, six thrusters are arranged symmetrically in a vector configuration, three per side. Instead of the usual orthogonal layout, which wastes energy, we treated the pitch angle α and yaw angle β of every thruster as optimization variables [32]. Each thruster (developed and manufactured in-house by KYEONG IN TECH Co., Ltd., Changwon, Republic of Korea) is rated at 500 W (48 VDC, 10 A) and produced about 105 N of thrust at 1000 RPM in still water; allowing a margin for stable power delivery, we set the per-thruster design thrust to 90 N at 920 RPM [32].

The mapping from each thruster's local thrust components to the resultant forces and moments in the body-fixed frame (Figure 8) is written as a 6×6 Jacobian matrix J . On this model, we built a loss function from the squared error between the realized and target thrust ratios ($F_x:F_y:F_z$), normalizing both force vectors by their sway (Y) components so the ratios compare directly, plus a term that rewards total thrust magnitude. In Equations (3) and (4), F is the resultant force vector the six thrusters produce in the body-fixed frame, F_y its sway

component, F_d the target force vector derived from the drag model, and F_{dy} the sway component of that target:

$$L_{\text{ratio}} = \left\| \frac{F}{F_y} - \frac{F_d}{F_{dy}} \right\|^2, \quad (3)$$

$$L = L_{\text{ratio}} + \lambda L_{\text{power}}, \quad L_{\text{power}} = -\|F\|, \quad (4)$$

where $\lambda > 0$ is a hyperparameter weighting the output-magnitude term. It is written λ , not α , to avoid collision with the thruster pitch angle α introduced above. The pitch and yaw angles of the six thrusters, collected in the parameter vector $\theta = [\alpha_1 \dots \alpha_6, \beta_1 \dots \beta_6]^T$, were then updated iteratively by gradient descent:

$$\theta \leftarrow \theta - \eta \nabla L \quad (5)$$

with learning rate η and a numerically computed gradient ∇L [32]. Over 5000 epochs, the gradient norm fell steadily to zero. Since fabrication only allows mounting at integer degrees, we rounded the converged angles, which gives the final three-dimensional geometry in Table 3. In [32], the same algorithm was also checked against three other representative target ratios (4.8:1:3, 4.8:1:2, and 3.3:1:1.7) and reproduced each of them within roughly 0.1.




Figure 8. Body-fixed coordinate system of the platform and thruster locations and numbering.

Table 3. Optimized thruster mounting angles ($v_x = 1.5$ m/s, $v_y = 0.5$ m/s, $v_z = 0.5$ m/s).

	T1	T2	T3	T4	T5	T6
Pitch α (°)	13	0	−13	−13	0	13
Yaw β (°)	28	−38	28	28	−38	28

The optimized configuration delivers maximum resultant thrusts of $F_{x,\text{max}} = 306.6$ N, $F_{y,\text{max}} = 63.6$ N, and $F_{z,\text{max}} = 167.4$ N. The angles in Table 2 follow the per-thruster Z–Y–X Euler mounting convention of [32]: mirrored port and starboard thrusters share the same nominal angles, and the left–right symmetry of the resultant wrench comes from the thruster position vectors in the allocation Jacobian and from the commanded thrust

directions of the individual thrusters. The per-thruster force components, RPM settings, and the additional validation cases are given in full in [32].

4. Control and Perception Algorithms

Robust hardware alone does not make interaction with a diver safe; currents, sensor noise, and model uncertainty remain. This section describes the motion controller that keeps the robot inside the diver's visual space and the perception pipeline that finds and localizes the diver.

4.1. Operational-Space-Formulated Time-Delay Control (TDC-OSF)

An underwater robot's dynamics carry strong nonlinearity and parameter uncertainty from added mass, nonlinear damping, and restoring forces, and in our mission, the external load changes abruptly whenever a diver grabs or releases the handle. We chose time-delay control (TDC) as the baseline [39]. TDC estimates the lumped unknown dynamics and disturbance from the control input and acceleration measured one sampling step earlier and cancels them, with no explicit dynamic model [33]. On top of this baseline, we reformulated the control law in the operational space of the guidance mission (TDC-OSF) [34]. The reason is visibility. In coastal water, a diver can usually see only 2–3 m, so a robot that heads for its goal in world coordinates alone can slip out of the diver's sight and set off panic. TDC-OSF therefore defines the task workspace as the diver's visual space and keeps projecting the reference trajectory onto a sphere 1–2 m ahead of the diver, regulating position and attitude errors directly in the coordinates in which the robot must remain visible [34]. A companion study extends the same relative-space formulation with a saturation-aware anti-windup design for this class of six-thruster platforms [40]. The controller runs at 50 Hz on the main controller as a TDC loop with PD-type desired error dynamics, and the commanded 6-DOF wrench is distributed to the six thrusters through the pseudo-inverse of the allocation Jacobian of Section 3.3. The full formulation and its experimental validation are reported in [33,34]. It is worth saying explicitly how this controller relates to the PID loops that also run on the platform, because the two are not competing solutions to the same problem. The PID loops serve manual operation: they hold attitude and depth while a diver or a surface operator drives the vehicle, and they deliberately do not fight a steady external load, so thruster activity—and with it the power draw—stays low. TDC-OSF serves autonomous mission execution, where the robot has to hold a commanded position relative to the diver against the current, and it does so by continuously estimating and cancelling the lumped disturbance. That rejection is bought with actuator effort. On a battery-powered platform that must stay down for more than two hours, running disturbance rejection through standby and manual segments would shorten the dive without making it safer, so the separation is a design decision rather than an omission. Section 6.4 returns to this trade-off.

4.2. Guidance Path Management

In the guidance scenario, the robot leads the diver along a planned route, at a pace the diver can hold. A waypoint-based path manager generates smooth reference trajectories between guidance points, slows down when the measured robot–diver distance grows so that the diver is never left behind, and holds position once that distance passes a safety threshold. Its output feeds the TDC-OSF loop described above, closing the chain of perception, path management, and robust control.

4.3. Multi-Sensor Fusion for Diver Detection and Localization

The platform also has to find the diver reliably in real water. The perception stack is not a contribution of this paper: it was developed together with Kongju National University and Zenti, and we describe it here only as far as is needed to interpret the guidance and

following results. Training data came from two sources. More than 100,000 images of divers were recorded with the stereo camera and the imaging sonar in water-tank environments, covering the diver distances, bearings and body attitudes that the robot encounters during guidance, and these were complemented by synthetic scenes rendered in an Unreal Engine 5-based HoloOcean simulation [41], in which turbidity, illumination and diver posture were varied to reach conditions that are difficult to stage in a tank [31]. Imagery recorded during the sea campaign of Section 5.3 was reserved for evaluation and for later retraining, and did not form part of the training set used for the results reported here.

In the configuration used for the trials reported here, the pipeline ran Damo-YOLO [42] for diver detection, a deep body tracker or BoostTrack [43] for frame-to-frame association, stereo disparity estimation (NVIDIA VPI) for three-dimensional localization over consecutive frames, and 6DRepNet [44] for orientation estimation. The detector and the tracker were selected for deployment rather than for benchmark accuracy. Damo-YOLO is designed for real-time inference on embedded accelerators, which the Jetson-class onboard computer requires if detection is to run alongside navigation and control without disturbing the 50 Hz control loop, and BoostTrack strengthens association when detections are weak and intermittent, which is the regime that turbid water produces. We did not benchmark alternative detectors on our own data, and the perception stack has since been revised; the results below correspond to the configuration just described. When visibility drops, a sonar-based position estimate takes over to keep the track alive [31]. Nothing is attached to the diver's body, which makes the approach a natural fit for recreational-diving support.

In tank and sea experiments, the diver's position and orientation were estimated within about ± 0.5 m and $\pm 15^\circ$ (Figure 9). The orientation figure meets the $\pm 15^\circ$ target set in the concept-design phase [31]. The position figure does not yet reach the 0.2 m design target; a quantitative evaluation against sonar-derived ground truth and an AHRS attached to the diver's tank is under way. For context, a recent stereo-vision method for underwater diver pose reports centimeter-level error on two of three axes under controlled conditions [45]; the values quoted here are field values obtained while the vehicle was under closed-loop control in coastal water, and the gap between the two indicates how much room is left in the perception channel. Non-contact estimation of the diver's physiological state from breathing-bubble imagery is being pursued in companion studies, in line with recent robotic respiration-rate estimation for diver safety [25], and lies outside the scope of this paper.




Figure 9. Diver detection and localization: (a) water-tank experiment; (b) open-water experiment, with the detected bounding box and skeletal keypoints overlaid. Red frames mark the detected bounding boxes, green lines and blue dots mark the estimated skeletal keypoints, and the small colored axes indicate the estimated diver pose; the text overlays are the on-screen readouts of the detection software.

5. Experiments and Results

We validated the platform in three stages: thrust measurements on a load-cell test bed, guidance-control experiments in an engineering water tank, and integrated experiments in the sea off Pohang, Republic of Korea.

5.1. Thrust Experiments on the Test Bed

To confirm that the optimized arrangement actually produces the design thrust ratio, we fixed the robot to a load-cell rig in the engineering tank of the Korea Institute of Industrial Technology (KITECH) [32] and measured the resultant force for each commanded direction (Figure 10). At the design operating point of 920 RPM per thruster, the measured maxima were 289.3 N in surge and 176.2 N in heave, within 5.6% and 5.3% of the calculated 306.6 N and 167.4 N; at intermediate RPM settings, the deviations stayed within 7%. Sway came out about 46% below the calculated 63.6 N. We attribute this to jet interference between the inward-angled thrusters and the hull side plates, which the rigid-body allocation model does not capture. The sway axis is not used for towing or emergency ascent, so the deficit does not touch the safety-critical performance envelope, but it is the subject of a planned CFD interaction study (Section 6.6). A related question is how far the allocation can be trusted once the operating condition moves away from the one it was tuned for. The allocation Jacobian of Section 3.3 is purely geometric: it depends on the mounting angles and lever arms, not on payload, depth or flow. A change in vehicle mass and added mass—a diver gripping the handle being the largest such change—therefore does not alter the mapping from thrust to wrench; it alters the acceleration that the same wrench produces, which is exactly the quantity the time-delay controller estimates and compensates. Depth has no first-order effect, because the vessels stay rigid at the rated pressure and the thrusters are not depth-compensated. Vehicle orientation and current do matter, however, because both change the flow entering each duct: in a pitched-up attitude, or in a beam current, the inflow to the inward-canted thrusters is no longer axial, and the jet–hull interference responsible for the sway deficit becomes orientation-dependent. We did not characterize that dependence experimentally, and we regard it as the main open question of the propulsion design (Section 6.6).




Figure 10. Thrust-measurement setup: (a) test rig with the lifting sling and load-cell strut; (b) the robot mounted on the rig in the engineering tank.

5.2. Guidance-Control Experiments in the Water Tank

The water-tank campaign in the KITECH tank had two stages (Figure 11). We first characterized the basic closed-loop behavior of the platform with baseline PID loops on heading, depth, and pitch at the 50 Hz control rate, and then ran the TDC-OSF controller with the path manager on a mock guidance course. In the first stage, for step heading commands of 45° and 90° , the robot converged without noticeable overshoot; the over-actuated arrangement redistributes thrust quickly. The depth loop, helped by a feedforward buoyancy-compensation term, held the steady-state depth error under 0.1 m. When we pushed the nose down by hand, the pitch excursion stayed within about $\pm 2^\circ$, which shows the passive metacentric restoring moment of the low-slung battery layout working together with the controller rather than against it. In human-in-the-loop tests, the robot held its commanded depth and heading while a diver drove it fore and aft through the rear handle. Table 4 summarizes these baseline results. In the second stage, the TDC-OSF controller and path manager tracked the mock guidance course, which fixed the control settings carried over to the open-water runs.




Figure 11. The robot in the KITECH engineering water tank during the guidance-control experiments.

Table 4. Summary of baseline closed-loop control performance in the water-tank experiments (PID stage).

Test	Command/Disturbance	Result
Heading step response	45° and 90° step commands	Convergence without appreciable overshoot
Depth keeping	Step depth command	Steady-state error < 0.1 m (buoyancy feedforward)
Pitch disturbance rejection	Manual nose-down push	Attitude deviation within approx. $\pm 2^\circ$
Human-in-the-loop towing	Handle fore/aft commands	Commanded depth and heading held during maneuver

5.3. Integrated Open-Water Experiments

The integrated experiments took place from 26 to 30 May 2025 off Pohang, Republic of Korea: system checks in a sheltered harbor basin, then the main runs in the open sea off Cheongjin 1-ri at depths of 5–15 m, with two instructor-certified divers participating (Figure 12). The water was typical of the Korean east coast in late spring, coastal and turbid with suspended matter, and jellyfish in the water occasionally interfered with visual diver detection. A single scenario tied together diver detection and tracking, guidance along a planned route, autonomous following, handle-commanded towing, and emergency ascent. Every transition between mission phases—descent, guidance, following, ascent—passed through an approval gate: before any high-risk action, the mission manager stopped and waited, proceeding only on approval from the diver (in these trials, given through the surface operator's console on the diver's behalf) and falling back to a call-and-hold behavior on rejection.




Figure 12. Open-water experiment: (a) test site off Cheongjin 1-ri, Pohang (satellite imagery: Google Earth, © 2025 Airbus); (b) the robot and a diver during the sea trial.

In the guidance runs, the robot led the diver around a diamond-shaped route of four guidance points at 7 m depth. It kept the robot–diver distance at 1–2 m and, despite the ambient current, recorded position errors within 0.15 m and heading errors within 3.5°, essentially matching the tank results (Figure 13); the diver-facing display showed depth, distance, and battery state throughout. In the following runs, the robot trailed the diver along arbitrary paths at a separation of 1–5 m. Whenever the diver left the camera's field of view, the robot stopped and waited until it saw the diver again. The wider spread in the following runs is a direct consequence of how detection behaves with range. At the 1–2 m held during guidance, the diver fills a large part of the frame, and the detection is essentially continuous; beyond roughly 4 m in the turbidity of the site, the diver subtends few pixels, contrast against the background drops, and detections start to be missed, so the separation drifts upward until the path manager intervenes. Two failure patterns were visible in the logs: partial framing, when only part of the body was inside the field of view, and false or occluding returns from jellyfish drifting through the scene. Both produced occasional single-frame jumps in the estimated diver position rather than a sustained error, because the tracker enforces temporal continuity and the path manager rate-limits the reference. When a jump persisted, the robot lost the track, held position and waited—the same conservative behavior used when the diver leaves the frame—so the control loop was never driven by a bad estimate. Tightening this further, by fusing the sonar range into the position estimate and by adding automatic camera calibration, is in progress.




Figure 13. Robot and diver trajectories logged during the open-water guidance and following runs: (a) robot horizontal path; (b) diver horizontal path; (c) robot–diver relative position over time.

In the towing runs, the diver gripped the rear handle and commanded fore/aft motion. The spring return stopped the robot reliably every time the diver let go, and the platform towed the diver at up to about 1 m/s, as measured by the onboard DVL, against the ambient current.

For the emergency-ascent test, the ascent command was issued while the robot held at a depth of 5 m. No diver or payload was attached in this test, so the result characterizes the platform's own vertical response. The thrust-allocation logic did not simply climb in a level attitude on heave thrust. It exploited the over-actuation: the platform pitched its nose up and pointed its much larger surge-axis thrust ($F_{x,max} = 306.6$ N, working against the small frontal projected area) along the ascent direction, together with the heave thrusters and the net positive buoyancy. The robot reached the surface from 5 m in about 4 s, an average vertical speed of roughly 1.25 m/s, against the ambient disturbances. With a diver attached, the added mass and drag would lower this rate; what such rates mean for a towed diver is taken up in Section 6.2.

Alongside the robot-side scenario, we validated the entire acoustic communication and shore-monitoring chain. Robot state and diver data traveled over the USBL/ATM link to a shipboard hub, then over an LTE network to an onshore control server, and appeared in real time on a shore-based GUI: the robot's position and status, the diver's position, biometric fields reserved for wearable sensors (e.g., oxygen saturation and heart rate), and emergency alarms on a 3D map (Figure 14). A shore-side safety supervisor can thus watch both the robot and the diver with no tether to the platform.




Figure 14. Data path from the robot to the shore console during the open-water experiment. Arrows indicate the direction of data flow.

6. Discussion

6.1. Physical Assistance as the Missing Layer of Diver Safety

The results back the central claim of this work: an underwater robot can go beyond passive observation and act as a physical safety layer. The surge and heave thrust ratios landed within 5–7% of the drag-model targets, which tells us that a gradient-descent geometric optimization against a CFD-derived drag model survives the transfer to real hardware. The handle proved its value in the sea trials. The two participating divers, neither of whom had used the robot before, operated the towing function immediately, because grip displacement maps to thrust in the most direct way possible and the interface fails safe on release. The emergency ascent, in turn, shows the benefit of over-actuation: by pitching up and redirecting the surge-axis thrust along the ascent direction, the platform reaches a steady ascent speed that a level attitude—or a conventional orthogonal thruster layout—could not.

6.2. Safety Concept and Failure Modes

Because NADIA is meant to be held on to and to pull, what happens when something goes wrong is not a secondary question. We treat diver safety as a layered property of the design rather than as the responsibility of any one component, and Table 5 lists the failure modes we considered together with the mechanism that bounds each of them. Three of those mechanisms are passive, in the sense that they take effect without the robot having to decide anything, and one of them keeps working even when the vehicle has no power at all.

The first is the handle. Its torsion spring carries the knob back to its neutral position the moment the diver lets go, so the zero-thrust command is produced by the mechanism itself: the interface falls back to rest without any action by the diver and without any decision by the control software. The handle is powered and wired like any other sensor, so the neutral reading still travels the usual potentiometer-to-controller-to-thruster path; what the spring provides is a mechanical default rather than a power-independent cut-off. For towing to continue after release, the potentiometer or its wiring would have to fail while holding a non-neutral value, which is one reason the surface operator retains an independent abort. The second is buoyancy. The platform is ballasted to about +29 N, so a total power loss or a flooded vessel ends with the robot rising slowly to the surface rather

than sinking onto the diver or onto the bottom. The third is the approval gate of Section 5.3, which keeps the robot from initiating descent, guidance or ascent on its own; in the sea trials, every high-risk transition waited for an explicit approval, and a rejection put the vehicle into a call-and-hold state.

Table 5. Failure modes considered for physical diver interaction and the mechanisms that bound them.

Failure Mode	Effect on the Diver	Bounding Mechanism	Status
Diver releases the handle	Unintended continued towing	Torsion-spring return to neutral: the zero-thrust command is restored mechanically, with no diver action and no control decision	Verified in tank and sea trials
Stuck or spurious thrust command	Uncommanded motion	Neutral-return potentiometer; saturation-aware anti-windup allocation [40]; surface operator can abort the mission	Partly verified
Control blackout or flooding	Vehicle sinks onto the diver or the seabed	Net positive buoyancy of about +29 N, so the vehicle surfaces	By design; set by ballasting
Loss of diver detection	Robot maneuvers without knowing where the diver is	Hold position and wait for re-acquisition; sonar-based fallback estimate	Verified in sea trials
Loss of the USBL/ATM link	Shore supervisor loses the situational picture	Onboard mission manager continues; approval gate blocks high-risk actions	Verified in sea trials
Collision with the diver or a structure	Impact injury or damage	1–2 m standoff, guidance speed limit, hold-and-wait on track loss	Verified in sea trials
Thruster degradation or partial failure	Reduced or asymmetric thrust during towing or ascent	Not monitored independently at present; appears only as degraded tracking	Open (Section 6.6)
Battery depletion	Loss of propulsion away from the boat	Endurance margin beyond 2 h; state of charge on the diver display and at the shore console	Partly verified

The modes that remain software-dependent are loss of diver detection and loss of the acoustic link. Both are handled by the same conservative rule, which is to hold position and wait rather than to search, because a robot that maneuvers while it has lost the diver is more dangerous than one that stops. Collision risk is bounded by that rule together with the guidance speed limit and the 1–2 m standoff. What the platform does not yet have is an independent monitor of the thrust chain itself: a partially failed thruster would at present show up only as degraded tracking, and the allocation would keep distributing the commanded wrench as though all six units were healthy. Active fault detection and fault-tolerant allocation for underwater vehicles are by now well-developed [46], and a thruster-level monitor is the single most valuable addition we can make to this safety chain.

6.3. Ascent-Rate Safety Considerations

A maximum-thrust ascent far exceeds the rates that diver-training organizations recommend, about 9–18 m/min or 0.15–0.30 m/s, to avoid pulmonary barotrauma and decompression stress [47]—the same injuries listed in Section 1 among the principal accident outcomes. The unloaded ascent in Section 5.3 averaged roughly 1.25 m/s from 5 m, four to eight times that envelope, and even with the added mass and drag of a diver, the platform could clearly exceed the recommended rates. We therefore read the demonstration as the platform’s maximum dynamic response, justified when the immediate risk of drowning of an unconscious diver at shallow depth outweighs the barotrauma risk, and not as a default rescue profile. For operational use, a rate-limited ascent controller will be added on top of the TDC-OSF framework, enforcing depth-dependent safe ascent profiles and reserving the

maximum-thrust mode for shallow-water drowning prevention under explicit operator authorization. Any physically assistive diver robot will, we believe, have to draw this same line between what the vehicle can do and what a medically safe rescue procedure should do.

6.4. Controller Selection and the Energy-Robustness Trade-Off

A natural question is how TDC-OSF compares with the other controllers commonly applied to underwater trajectory tracking, and why the platform also carries PID loops. The two controllers on board are not competing candidates for the same task. PID is used where the vehicle is driven by a person, that is, for joystick operation and station keeping, and it is deliberately tuned not to fight a steady external load, so the thrusters stay near idle when the vehicle is merely holding station in a mild current. TDC-OSF is used where the robot has to hold a commanded position relative to a moving diver against the current, and it achieves that by estimating and cancelling the lumped disturbance at every sample. The difference matters in operation, because rejecting a disturbance means acting against it continuously, which raises thruster duty and power draw; on a platform that has to remain available for more than two hours, a controller that fights the sea during standby buys nothing. Table 4 reports the PID stage because it characterizes the vehicle, not because it is an alternative to TDC-OSF for the guidance task.

Against the other candidates, the argument for time-delay control is practical rather than theoretical. Model predictive control would give explicit constraint handling, which is attractive for a vehicle that must respect thrust saturation, but it needs a dynamic model and an optimization at every step, and the model is exactly what is hardest to obtain here, because the effective mass and damping change the moment a diver grips the handle. Adaptive control accommodates parameter uncertainty but converges over time, whereas the load change at handle engagement is a step. Sliding-mode control rejects matched disturbances robustly but chatters unless it is carefully smoothed, and chattering on a vehicle that a diver is holding is undesirable for reasons beyond actuator wear. Time-delay control needs no explicit model, costs little computation, and reacts to an abrupt load change within one sampling interval, which is why it was chosen; a saturation-aware extension has since been developed for the same class of platform [40].

What we cannot offer is a controlled comparison. The trials reported here were scenario runs with divers in the water rather than a controller benchmark, and sea state, current and diver behavior were not repeatable between runs, so setting PID, MPC, adaptive and sliding-mode control against TDC-OSF under identical disturbances was not possible within this campaign. Such a benchmark, ideally in the engineering tank with a programmable current and a towed diver surrogate, with actuator energy logged alongside tracking error, would settle both the tracking question and the power question, and we regard it as the natural next study.

6.5. Comparison with Prior Diver-Support Systems

What separates this platform from CADDY [10–12] and the perception-oriented studies that followed it [17–21] is that the interaction channel is physical as well as visual. Table 6 places the platform beside the families of system it draws on, along the axes that matter for diver safety: whether the machine perceives the diver at all, whether it can guide and follow autonomously, whether it can act on the diver physically, whether a person authorizes high-risk actions, and whether any of it has been shown at sea with divers in the water. No previous system occupies all five columns, and that combination, rather than any single one of the five, is what we claim as the contribution. What separates it from a DPV [9] is the reverse: the towing function sits inside an autonomous system that perceives the diver,

leads or follows, and reports to the surface, rather than being the whole product. Gesture- and vision-only interaction collapses exactly where accidents concentrate—turbidity, panic, loss of line of sight—whereas a graspable handle and a towing-capable thrust envelope keep working there. The diver-approval gating demonstrated in the sea trials adds a safeguard the earlier systems lacked: the robot never starts a high-risk action on its own. Our perception pipeline is deliberately conservative (detection, tracking, localization) and complements the physical channel rather than replacing it, though the quantitative perception results reported here remain preliminary.

Table 6. Capability comparison with the families of system most closely related to this work.

System	Diver Perception	Autonomous Guidance and Following	Physical Assistance to the Diver	Human Approval of High-Risk Actions	Open-Water Validation with Divers
CADDY [10–12]	Acoustic tracking and stereo gesture recognition	Follows the diver and executes gesture commands	None	Gesture command only	Yes
Diver-following AUVs [22–24]	Vision, or vision fused with acoustics	Following	None	No	Partly
Diver propulsion vehicles [9]	None	None	Towing under a manual trigger	Not applicable (hand-held)	Commercial practice
Cable-driven rescue concept [30]	None	None	Tethered retrieval (concept)	None	No (concept and CAD)
NADIA (this work)	Stereo vision fused with imaging sonar	Guidance along a planned route and autonomous following	Handle-commanded towing and pitch-up emergency ascent	Approval gate on every high-risk transition	Yes, at 5–15 m off Pohang

6.6. Limitations and Future Work

Four limitations bound these results. First, the sway-thrust deficit of Section 5.1 shows that the rigid-body allocation model misses jet-hull interaction; a CFD study of thruster-hull interference and an allocation-matrix correction are planned. Second, the open-water experiments stayed at recreational depths of 5–15 m in moderate sea states, so the full 50 m envelope and stronger currents remain untested. Third, non-contact estimation of the diver’s physiological state is still exploratory in companion studies and needs large-scale quantitative evaluation before anything operational rests on it; extending the platform toward autonomous anomaly response therefore remains future work. Fourth, the sea trials exposed practical problems of routine operation. Launching and recovering an 85 kg platform with a standard vessel winch endangers the bottom-mounted sonar and camera; shipboard connectors and the Wi-Fi module were intermittently unreliable under salt exposure; and electrical noise from the vessel’s 220 VAC generator disturbed the USBL positioning until we switched to an independent battery power bank. We plan a sealed shipboard control architecture, an independent power system, a dedicated launch-and-recovery cradle, and a lighter frame and vessels in composite materials such as CFRP.

Beyond these, it is worth being explicit about the conditions in which the platform should not yet be relied upon, because the claim that a robot can act as a physical safety layer is only as good as the envelope in which that has been shown. The vehicle is rated to 50 m but has been exercised to 15 m, and the ascent discussion of Section 6.3 becomes more constraining with depth, not less, since a fast ascent from 40 m carries a decompression risk that a shallow-water rescue does not. In strong currents, the limiting factor is not surge thrust, which exceeds the drag at the design speed by a wide margin, but sway and yaw authority, which Section 5.1 shows to be weaker than the model predicts; a beam current is therefore the worst case for holding station beside a diver. In very low visibility, the optical channel fails first, and although the sonar keeps a range estimate alive, the

present pipeline cannot recover diver orientation from sonar alone, so guidance degrades to following at a larger standoff. If the acoustic link is lost, the robot continues its mission but the shore supervisor loses sight of it, which is acceptable for the short guidance legs demonstrated here and would not be acceptable for an unsupervised rescue. Battery depletion is handled by display and by margin only, with no automatic return-to-surface at a low-charge threshold. Finally, a completely incapacitated diver lies outside the present design: every assistance function reported here, towing included, requires the diver to grip the handle, so an unconscious diver cannot be moved by the platform at all. Closing that gap needs a mechanical attachment the robot can engage by itself, which raises its own safety questions and is the subject of the next design iteration.

7. Conclusions

This paper presented the integrated design and open-water validation of NADIA, an autonomous underwater robot that gives divers physical safety assistance. FEA verified the modular pressure-vessel architecture and its passive restoring design at the 50 m rated depth. The six-thruster vector arrangement, optimized by gradient descent against a CFD-derived drag model, realized the target thrust ratio within 5–7% in surge and heave. The operational-space time-delay controller held guidance position errors within 0.15 m and heading errors within 3.5° while keeping the robot 1–2 m ahead of the diver, and the stereo-camera/imaging-sonar pipeline localized the diver within about ± 0.5 m and $\pm 15^\circ$. In the sea trials off Pohang, divers commanded the robot through the rear handle for towing, and followed robot-led guidance gated by explicit approval, and the platform climbed from a 5 m holding depth to the surface in about 4 s in a pitch-up emergency ascent, while robot and diver data reached an onshore console over a USBL/ATM–LTE communication chain. As far as we know, this is the first open-water demonstration of an autonomous diver-support robot that combines diver perception and autonomous guidance with diver-commanded physical intervention, and we hope it serves as a practical foundation for moving underwater human–robot interaction from observation to intervention.

Author Contributions: Conceptualization, H.K. and G.R.C.; methodology, M.-G.K., H.K. and G.R.C.; software, H.J. and G.R.C.; validation, M.-G.K., H.J. and H.K.; formal analysis, M.-G.K.; investigation, M.-G.K., H.K. and H.J.; resources, G.K. and M.N.; data curation, H.J.; writing—original draft preparation, M.-G.K.; writing—review and editing, H.K., G.R.C. and J.-H.L.; visualization, M.-G.K.; supervision, J.-H.L.; project administration, H.K.; funding acquisition, J.-H.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Korea Institute of Marine Science & Technology Promotion (KIMST), funded by the Ministry of Oceans and Fisheries, Republic of Korea (RS-2022-KS221668 and RS-2023-00256122).

Institutional Review Board Statement: Ethical review and approval by an external ethics committee were waived for this study in accordance with institutional policy, because the sea trials were engineering tests performed by instructor-certified professional divers of the project team under the internal safety review and approval of the Korea Institute of Robotics and Technology Convergence (KIRO). All participating divers were briefed on the procedures and risks and gave their consent prior to participation.

Informed Consent Statement: Informed consent was obtained from all subjects involved in this study. Written informed consent has been obtained from the individuals appearing in the photographs (Figures 9, 11 and 12) to publish identifiable images in this paper.

Data Availability Statement: The data presented in this study are available from the corresponding author upon reasonable request. The data are not publicly available due to the confidentiality requirements of the funded project.

Acknowledgments: The authors thank Kongju National University and Zenti for the joint development of the diver-recognition algorithms, and the Marine Robotics Center of the Korea Institute of Industrial Technology (KITECH) for providing the engineering-tank facility used in the thrust and control experiments.

Conflicts of Interest: Authors Gyeongmin Kim and Minshik Noh are employed by the company KYEONG IN TECH Co., Ltd., which participated in the mechanical design and fabrication of the platform. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest. The funders had no role in the design of this study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation	Definition
ATM	Acoustic telemetry modem
AUV	Autonomous underwater vehicle
BCD	Buoyancy control device
CAD	Computer-aided design
CFD	Computational fluid dynamics
CFRP	Carbon-fiber-reinforced plastic
CoB/CoM	Center of buoyancy/center of mass
DCS	Digital compass system
DOF	Degrees of freedom
DPV	Diver propulsion vehicle
DVL	Doppler velocity log
FEA	Finite-element analysis
GPS	Global positioning system
GUI	Graphical user interface
HMI	Human-machine interface
IMU	Inertial measurement unit
LARS	Launch-and-recovery system
LED	Light-emitting diode
LTE	Long-term evolution (cellular network)
MPC	Model predictive control
PID	Proportional-integral-derivative (control)
RF	Radio frequency
ROV	Remotely operated vehicle
SF	Safety factor
SST	Shear stress transport (turbulence model)
TDC/TDC-OSF	Time-delay control/operational-space-formulated time-delay control
UHRI	Underwater human-robot interaction
USBL	Ultra-short baseline

References

1. Korea Maritime Institute. *Marine Leisure Statistical Yearbook*; KMI: Busan, Republic of Korea, 2019.
2. Kang, C.-W.; Seo, D.-W.; Kwon, T.-S. Structural Relationships among Enjoyment, Tour Value and Tourism Behavioral Intention of Scuba Diving Experience. *Korea J. Sports Sci.* **2018**, *27*, 231–245. [[CrossRef](#)]
3. Kang, H.; Lee, M.-J.; Cho, G.R.; Ki, G.; Kim, M.-G.; Li, J.-H. Development of ROV Trencher URI-T and Its Sea Trial. *J. Ocean Eng. Technol.* **2019**, *33*, 300–311. [[CrossRef](#)]
4. Kwon, Y.T.; Pyo, H.G. A Study on Safety Measures Proposals through Analysis of Diving Accident Types. *J. Korean Phys. Educ. Sport Sci.* **2015**, *24*, 135–148.
5. Nađ, Đ.; Ferreira, F.; Mišković, N. Using Autonomous Underwater Vehicles for Diver Tracking and Navigation Aiding. *J. Mar. Sci. Eng.* **2020**, *8*, 413. [[CrossRef](#)]

6. Choi, H.-T.; Yuh, J. Underwater Robots. In *Springer Handbook of Robotics*, 2nd ed.; Siciliano, B., Khatib, O., Eds.; Springer: Cham, Switzerland, 2016; pp. 595–622.
7. Capocci, R.; Dooly, G.; Omerdić, E.; Coleman, J.; Newe, T.; Toal, D. Inspection-Class Remotely Operated Vehicles—A Review. *J. Mar. Sci. Eng.* **2017**, *5*, 13. [\[CrossRef\]](#)
8. Wynn, R.B.; Huvenne, V.A.I.; Le Bas, T.P.; Murton, B.J.; Connelly, D.P.; Bett, B.J.; Ruhl, H.A.; Morris, K.J.; Peakall, J.; Parsons, D.R.; et al. Autonomous Underwater Vehicles (AUVs): Their Past, Present and Future Contributions to the Advancement of Marine Geoscience. *Mar. Geol.* **2014**, *352*, 451–468. [\[CrossRef\]](#)
9. Qin, H.; Li, Z.; Xu, S.; Liu, X.; Cao, X. Review of Diver Propulsion Vehicle: A Review. *Phys. Fluids* **2024**, *36*, 101305. [\[CrossRef\]](#)
10. Mišković, N.; Pascoal, A.; Bibuli, M.; Caccia, M.; Neasham, J.A.; Birk, A.; Egi, M.; Grammer, K.; Marroni, A.; Vasilijević, A.; et al. CADDY Project, Year 3: The Final Validation Trials. In Proceedings of the OCEANS 2017—Aberdeen, Aberdeen, UK, 19–22 June 2017; pp. 1–5.
11. Stilinović, N.; Nađ, Đ.; Mišković, N. AUV for Diver Assistance and Safety—Design and Implementation. In Proceedings of the OCEANS 2015—Genova, Genoa, Italy, 18–21 May 2015; pp. 1–4.
12. Egi, S.M.; Balestra, C.; Pieri, M.; Cialoni, D.; Thomas, G.; Marroni, A. Cognitive Autonomous Diving Buddy (CADDY): AUV Design Requirements for Divers’ Safety. In Proceedings of the OCEANS 2015—Genova, Genoa, Italy, 18–21 May 2015; pp. 1–5.
13. Chavez, A.G.; Ranieri, A.; Chiarella, D.; Zereik, E.; Babić, A.; Birk, A. CADDY Underwater Stereo-Vision Dataset for Human–Robot Interaction (HRI) in the Context of Diver Activities. *J. Mar. Sci. Eng.* **2019**, *7*, 16. [\[CrossRef\]](#)
14. Chiarella, D.; Bibuli, M.; Bruzzone, G.; Caccia, M.; Ranieri, A.; Zereik, E.; Marconi, L.; Cutugno, P. A Novel Gesture-Based Language for Underwater Human–Robot Interaction. *J. Mar. Sci. Eng.* **2018**, *6*, 91. [\[CrossRef\]](#)
15. Birk, A. A Survey of Underwater Human–Robot Interaction (U-HRI). *Curr. Robot. Rep.* **2022**, *3*, 199–211. [\[CrossRef\]](#)
16. Aldhaheri, S.; Renda, F.; De Masi, G. Underwater Human-Robot and Human-Swarm Interaction: A Review and Perspective. In Proceedings of the OCEANS 2024—Singapore, Singapore, 15–18 April 2024; pp. 1–9.
17. Zhang, Y.; Jiang, Y.; Qi, H.; Zhao, M.; Wang, Y.; Wang, K.; Wei, F. An Underwater Human–Robot Interaction Using a Visual–Textual Model for Autonomous Underwater Vehicles. *Sensors* **2023**, *23*, 197. [\[CrossRef\]](#) [\[PubMed\]](#)
18. Kvasić, I.; Orbaugh Antillon, D.; Nađ, Đ.; Walker, C.; Anderson, I.; Mišković, N. Diver-Robot Communication Dataset for Underwater Hand Gesture Recognition. *arXiv* **2025**, arXiv:2506.08974.
19. Mangalvedhekar, S.; Nahar, S.; Maskare, S.; Mahajan, K.; Bagade, A. Interpretable Underwater Diver Gesture Recognition. *arXiv* **2023**, arXiv:2312.04874.
20. Guo, Y.; Zhang, R.; Qiu, W.; Asada, H.; Wan, F.; Song, C. Underwater Intention Recognition Using Head Motion and Throat Vibration for Supernumerary Robotic Assistance. In Proceedings of the 2023 IEEE 19th International Conference on Automation Science and Engineering (CASE), Auckland, New Zealand, 26–30 August 2023; pp. 1–6.
21. Hong, J.; Enan, S.S.; Sattar, J. Diver Identification Using Anthropometric Data Ratios for Underwater Multi-Human–Robot Collaboration. *IEEE Robot. Autom. Lett.* **2024**, *9*, 3514–3521. [\[CrossRef\]](#)
22. Sattar, J.; Dudek, G. Where Is Your Dive Buddy: Tracking Humans Underwater Using Spatio-Temporal Features. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Diego, CA, USA, 29 October–2 November 2007; pp. 3654–3659.
23. Islam, M.J.; Hong, J.; Sattar, J. Person-Following by Autonomous Robots: A Categorical Overview. *Int. J. Robot. Res.* **2019**, *38*, 1581–1618. [\[CrossRef\]](#)
24. Remmas, W.; Chemori, A.; Kruusmaa, M. Diver Tracking in Open Waters: A Low-Cost Approach Based on Visual and Acoustic Sensor Fusion. *J. Field Robot.* **2021**, *38*, 494–508. [\[CrossRef\]](#)
25. Kutzke, D.T.; Dupati, V.; Sattar, J. Robotic Estimation of Single Scuba Diver Respiration Rate for Safety in Underwater Human–Robot Collaboration. *Int. J. Robot. Res.* **2026**, early access. [\[CrossRef\]](#)
26. Hung, N.; Huynh, P.; Tung, M.; Vu, N.; Dat, N. CLVM: A Hybrid Deep Learning Framework for Contactless Virtual Mouse Control. *Inform. Autom.* **2026**, *25*, 378–409. [\[CrossRef\]](#)
27. Wu, Y.; Liu, J.; Wei, Y.; An, D.; Duan, Y.; Li, W.; Li, B.; Chen, Y.; Wei, Q. Intelligent Control Method of Underwater Inspection Robot in Netcage. *Aquac. Res.* **2022**, *53*, 1928–1938. [\[CrossRef\]](#)
28. Kartal, S.K.; Cantekin, R.F. Autonomous Underwater Pipe Damage Detection Positioning and Pipe Line Tracking Experiment with Unmanned Underwater Vehicle. *J. Mar. Sci. Eng.* **2024**, *12*, 2002. [\[CrossRef\]](#)
29. Yang, Y.; Hirose, S.; Debenest, P.; Guarnieri, M.; Izumi, N.; Suzumori, K. Development of a Stable Localized Visual Inspection System for Underwater Structures. *Adv. Robot.* **2016**, *30*, 1415–1429. [\[CrossRef\]](#)
30. Fattah, S.A.; Abedin, F.; Ansary, M.N.; Rokib, M.A.; Saha, N.; Shahnaz, C. R3Diver: Remote Robotic Rescue Diver for Rapid Underwater Search and Rescue Operation. In Proceedings of the 2016 IEEE Region 10 Conference (TENCON), Singapore, 22–25 November 2016; pp. 3280–3283.
31. Kang, H.; Cho, G.R.; Jin, H.; Kim, M.-G.; Li, J.-H.; Jin, S.; Lee, J.; Min, J. Robot-Based Diver Safety Support Scenario and Concept Design for NADIA. *J. Ocean Eng. Technol.* **2025**, *39*, 235–247. [\[CrossRef\]](#)

32. Jin, S.; Cho, G.R.; Kang, H.; Kim, M.-G.; Choi, Y. Optimal Thruster Arrangement Design for Underwater Drones Based on Gradient Descent. *InUS Trans.* **2025**, *1*, 84–97. [\[CrossRef\]](#)
33. Cho, G.R.; Kang, H.; Kim, M.-G.; Lee, M.-J.; Li, J.-H.; Kim, H.; Lee, H.; Lee, G. An Experimental Study on Trajectory Tracking Control of Torpedo-like AUVs Using Coupled Error Dynamics. *J. Mar. Sci. Eng.* **2023**, *11*, 1334. [\[CrossRef\]](#)
34. Cho, G.R.; Kang, H.; Kim, M.-G.; Park, S.; Jin, H.; Jin, S. Operational Space Formulation of Time Delay Control for Underwater Vehicles and Its Extension to Diver Guidance. In Proceedings of the 2025 IEEE Underwater Technology (UT), Taipei, Taiwan, 2–5 March 2025; pp. 1–8. [\[CrossRef\]](#)
35. Paull, L.; Saeedi, S.; Seto, M.; Li, H. AUV Navigation and Localization: A Review. *IEEE J. Ocean. Eng.* **2014**, *39*, 131–149. [\[CrossRef\]](#)
36. Kim, G.; Rho, M.; Choi, C. Structural Analysis of a Pressure Vessel for Mounting a Diver Safety Support Robot Platform. In Proceedings of the 19th Korea Robotics Society Annual Conference (KRoC 2024), Pyeongchang, Republic of Korea, 21–24 February 2024.
37. Fossen, T.I. *Handbook of Marine Craft Hydrodynamics and Motion Control*; John Wiley & Sons: Chichester, UK, 2011.
38. Menter, F.R. Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications. *AIAA J.* **1994**, *32*, 1598–1605. [\[CrossRef\]](#) [\[PubMed\]](#)
39. Youcef-Toumi, K.; Ito, O. A Time Delay Controller for Systems with Unknown Dynamics. *J. Dyn. Syst. Meas. Control* **1990**, *112*, 133–142. [\[CrossRef\]](#)
40. Cho, G.R.; Kang, H.; Park, S.; Kim, M.-G.; Jin, H.-S.; Bae, C.; Jin, S.; Li, J.-H. Saturation-Aware Anti-Windup Design for Robust Relative-Space Motion Control of Underwater Vehicles Using Time-Delay Estimation. *InUS Trans.* **2026**, *2*, 96–111.
41. Potokar, E.; Ashford, S.; Kaess, M.; Mangelson, J.G. HoloOcean: An Underwater Robotics Simulator. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23–27 May 2022; pp. 3040–3046.
42. Xu, X.; Jiang, Y.; Chen, W.; Huang, Y.; Zhang, Y.; Sun, X. DAMO-YOLO: A Report on Real-Time Object Detection Design. *arXiv* **2022**, arXiv:2211.15444.
43. Stanojević, V.D.; Todorović, B.T. BoostTrack: Boosting the Similarity Measure and Detection Confidence for Improved Multiple Object Tracking. *Mach. Vis. Appl.* **2024**, *35*, 53. [\[CrossRef\]](#)
44. Hempel, T.; Abdelrahman, A.A.; Al-Hamadi, A. 6D Rotation Representation for Unconstrained Head Pose Estimation. In Proceedings of the IEEE International Conference on Image Processing (ICIP), Bordeaux, France, 16–19 October 2022; pp. 2496–2500.
45. Wang, K.; Yang, Y.; Wei, Q.; Zhou, X.; Hu, Z.; Zheng, Q. An Underwater 6-DoF Position and Orientation Estimation Method for Divers Based on the VideoPose5CH Model. *Sensors* **2026**, *26*, 1335. [\[CrossRef\]](#) [\[PubMed\]](#)
46. Liu, F.; Tang, H.; Qin, Y.; Duan, C.; Luo, J.; Pu, H. Review on Fault Diagnosis of Unmanned Underwater Vehicles. *Ocean Eng.* **2022**, *243*, 110290. [\[CrossRef\]](#)
47. Scuba Diving International. *SDI Open Water Scuba Diver Manual*; Scuba Diving International: Stuart, FL, USA, 2019.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.