Turkish Journal of Science & Technology
21(2), 351-361, 2026
https://doi.org/10.55525/tjst.1989945

Research Article

Underwater-Mission-Oriented Objective Function Design for Metaheuristic Tuning of PID and Fractional-Order PID Controllers

Hakan ERSOY[1], Aslihan KARTCI[2]*

[1,2] Department of Electronics and Communication Engineering, Faculty of Electrical and Electronics, Yıldız Technical University, İstanbul, Türkiye
[1] hakan.ersoy@yildiz.edu.tr, [*2] aslihan.ozakin@yildiz.edu.tr

(Geliş/Received: 12/07/2026; Kabul/Accepted: 06/09/2026)

Abstract: Metaheuristic tuning of depth controllers for remotely operated vehicles is usually performed against a generic error integral, which does not encode mission priorities such as battery use, actuator stress from chattering, direction-dependent overshoot risk and post-event recovery. This paper treats the objective function itself as the design object. A mission-oriented cost augments a time-weighted error integral with baseline-normalized effort, smoothness, asymmetric-overshoot and recovery terms on a one-degree-of-freedom BlueROV2 heave model with saturation, actuator lag and sensor noise. Adding the terms cumulatively raises the advantage of the fractional controller over PID from 7.0 to 22.7 percent over ten runs per level, and encoding survey, near-bottom inspection and long-endurance missions as weight profiles gives it the lower cost in every run, largest at 35.4 percent on inspection. The conclusion is then tested against the designer's choices: 160 random weight vectors, four normalization schemes, Monte-Carlo sweeps over noise and plant parameters, and eleven held-out scenarios. A linear analysis exposes one negative result: the cost carries no frequency-domain specification, so the tuned loops have small phase margins that saturation hides, and an explicit margin requirement is needed to restore them. The study is simulation-based and one-dimensional; no hardware validation is claimed.

Key words: BlueROV2, fractional-order PID, mission-oriented tuning, objective function design, particle swarm optimization.

Sualtı Görevlerine Yönelik Amaç Fonksiyonu Tasarımı ile PID ve Kesir Dereceli PID Denetleyicilerin Meta-Sezgisel Ayarlanması

Özet: Uzaktan kumandalı su altı araçlarının derinlik denetleyicileri meta-sezgisel yöntemlerle çoğunlukla genel bir hata integraline göre ayarlanır; bu ölçüt batarya kullanımı, denetim salınımlarının oluşturduğu aktüatör yükü, yön bağımlı aşım riski ve akıntı olayı sonrası toparlanma gibi göreve özgü öncelikleri temsil etmez. Bu çalışmada amaç fonksiyonunun kendisi tasarım nesnesi olarak ele alınmıştır. Zaman ağırlıklı hata integrali; denetim çabası, denetim düzgünlüğü, asimetrik aşım ve olay sonrası toparlanma terimleriyle genişletilmiş, tüm terimler referans bir PID tasarımına göre normalize edilmiş ve doygunluk, aktüatör gecikmesi ve gürültü içeren tek serbestlik dereceli bir BlueROV2 düşey hareket modelinde sınanmıştır. Terimler kademeli olarak eklendiğinde kesir dereceli denetleyicinin PID'ye üstünlüğü, seviye başına on koşuda yüzde 7,0'den yüzde 22,7'e yükselmiştir. Tarama, deniz tabanına yakın inceleme ve uzun süreli görevler ağırlık profilleri olarak kodlandığında kesir dereceli denetleyici tüm koşularda daha düşük maliyet üretmiş, en büyük fark yüzde 35,4 ile inceleme görevinde elde edilmiştir. Sonuç, tasarımcı tercihlerine karşı sınanmıştır: 160 rastgele ağırlık vektörü, dört normalizasyon şeması, ölçüm gürültüsü ve model belirsizliği üzerinde Monte-Carlo taramaları ve on bir tutulan senaryo. Doğrusal analiz olumsuz bir sonuç vermiştir: amaç fonksiyonunda frekans bölgesi kısıtı bulunmadığından ayarlanan çevrimlerin faz payları düşüktür ve bu durum itici doygunluğunca gizlenmektedir. Çalışma benzetim tabanlıdır; donanım doğrulaması iddia edilmemektedir.

Anahtar kelimeler: Amaç fonksiyonu tasarımı, BlueROV2, görev odaklı ayarlama, kesir dereceli PID, parçacık sürü optimizasyonu.

1. Introduction

Remotely operated vehicles (ROVs) must maintain accurate depth despite nonlinear hydrodynamic damping, added mass, actuator dynamics, thrust limits, environmental disturbances, and noisy measurements [1, 2]. The meaning of “good” depth control also changes with the mission. A survey or transit task primarily requires fast and well-damped tracking, a near-bottom inspection task places a higher penalty on downward overshoot and oscillatory thrust, and a long-endurance mission must limit sustained actuation demand. A tuning criterion based only on a generic error integral therefore captures only part of the operational problem.

[*] Corresponding author: aslihan.ozakin@yildiz.edu.tr. ORCID Number of authors: [1] 0009-0006-0495-6589, [2*]0000-0001-5690-7574

PID control remains attractive because of its simple structure and low implementation cost. Fractional-order PID ($PI^{\lambda}D^{\mu}$) extends the classical controller through the integral and derivative orders λ and μ [3, 4]. These additional degrees of freedom allow the loop shape to be adjusted more flexibly. Practical realization can be achieved through band-limited rational approximations such as the Oustaloup filter [5], while recent studies have also demonstrated resource-efficient hardware implementations of fractional-order integrators and differentiators using FIR structures on FPGA platforms [6]. However, the benefit observed from $PI^{\lambda}D^{\mu}$ depends not only on the controller structure, but also on what the tuning process is asked to minimize [7]. If the objective is a single tracking-error index, the optimizer may have little reason to exploit the additional fractional degrees of freedom.

Metaheuristic algorithms are widely used for this class of tuning problem because the closed-loop objective is nonlinear, derivative-free, and often includes saturation or simulation-based constraints [8]. Particle swarm optimization (PSO) [9] and grey wolf optimization (GWO) [10] are representative examples. In much of the controller-tuning literature, however, methodological attention is placed on the controller or the optimizer, whereas the objective function is kept as the integral of time-weighted absolute error (ITAE), the integral of absolute error (IAE), the integral of squared error (ISE), or a fixed weighted combination of such conventional indices, which are useful for benchmarking but encode none of the mission priorities above.

Multi-objective formulations of controller tuning are not new: Pareto-based solvers [11] have been applied to PID and fractional-order PID design [12, 13], the properties of weighted-sum scalarization are well characterized [14, 15], and energy-aware objectives for underwater vehicles have been studied inside model predictive control [16] and in variable-structure depth control [17]. What differs here is the question rather than the machinery: every term used below is standard, and the claim is that the measured benefit of the fractional degrees of freedom is a function of which terms the objective prices, established rather than asserted.

This study therefore treats the objective function itself as the primary design object. A normalized five-term cost is constructed for underwater depth control, and the central question is tested directly: do the additional fractional degrees of freedom become more useful as the control objective becomes richer and more conflicting? The analysis is divided into an objective-complexity ablation and a mission-profile study on a BlueROV2 heave model including thrust saturation, actuator lag, and sensor noise.

The main contributions are: (1) a mission-oriented objective function combining tracking, control effort, smoothness, asymmetric overshoot and post-event recovery, with baseline normalization and an explicit procedure for eliciting its weights; (2) an ablation, cross-checked by a randomized weight study, showing that the PID-to-$PI^{\lambda}D^{\mu}$ gap is governed by which terms the objective prices rather than by how many it contains; (3) a mission-profile study with statistical and cross-algorithm comparison; and (4) a verification of the conclusion against the designer's own choices - weights, normalization, evaluation scenario, measurement noise, plant parameters and actuator idealization - together with the linear stability margins of every tuned loop.

2. Plant Model and Scenario

The heave (depth) dynamics of the ROV are modeled in a down-positive frame by the one-degree-of-freedom Fossen-type equation [1, 2]

$$(m - Z_{\dot{w}})\dot{w} = Z_w w + Z_{w|w|} w|w| - B_{\mathrm{res}} + u_{\mathrm{act}} + d(t), \qquad \dot{z} = w. \tag{1}$$

where z is depth, w is heave velocity, B_{res} is the residual positive buoyancy, u_{act} is the applied vertical thrust, and $d(t)$ is an exogenous disturbance force. Following the sign convention of [1], the hydrodynamic coefficients $Z_{\dot{w}}$, Z_w, and $Z_{w|w|}$ are negative (Table 1), giving the effective inertia $m - Z_{\dot{w}} = 26.07$ kg. The commanded thrust u_{cmd} is saturated at $\pm u_{\max}$ and passes through a first-order actuator lag with time constant τ_{act}:

$$\dot{u}_{\mathrm{act}} = \frac{\mathrm{sat}(u_{\mathrm{cmd}}) - u_{\mathrm{act}}}{\tau_{\mathrm{act}}}. \tag{2}$$

The controller operates at 50 Hz on a depth measurement corrupted by zero-mean Gaussian noise with $\sigma_z = 1$ cm; a frozen noise realization is used so that every candidate controller is evaluated on an identical scenario, keeping the optimization deterministic. Model parameters are representative literature values for the BlueROV2 [2].

Table 1. Plant, actuator, and scenario parameters (Fossen sign convention: hydrodynamic derivatives are negative).

Parameter	Value	Parameter	Value
m	11.5 kg	u_{max}	40 N
$Z_{\dot{w}}$	−14.57 kg	τ_{act}	0.15 s
Z_w	−4.57 N s m^{-1}	σ_z	1 cm
$Z_{w\|w\|}$	−34.6 N s^2 m^{-2}	f_s (controller)	50 Hz
B_{res}	2.0 N	T (horizon)	60 s

The evaluation scenario is a 5 m dive step at $t = 1$ s, followed by a finite current event. For $35 \leq t < 40$ s, the disturbance

$$d(t) = F_0 + F_1 \sin\big(2\pi f_1 (t - t_d)\big), \qquad F_0 = 8 \text{ N}, \quad F_1 = 2.5 \text{ N}, \quad f_1 = 0.4 \text{ Hz} \tag{3}$$

is applied and then removed. The event therefore has a well-defined end, so post-event recovery can be measured as a genuine transient rather than as sustained disturbance rejection [18].

The parameters are those of the open-source BlueROV2 benchmark of [2], and independent experimental testing and simulator correction for the BlueROV2 are reported in [19]; no new identification was performed here, so this is a representative BlueROV2-class heave model rather than a validated digital twin, and Section 6.4 re-tests every conclusion under ±30% parameter perturbation. Restricting the model to heave also removes three couplings: roll and pitch moments from vertical-thruster imbalance, projection of surge thrust and drag through a non-zero pitch angle, and off-diagonal added-mass terms. Each enters the heave loop as extra phase lag and as a state-dependent disturbance, and both act against designs relying on high-frequency gain, which is the characteristic of the integer derivative rather than the band-limited fractional one; the one-degree-of-freedom result is therefore expected to be conservative. Settling that expectation requires a coupled six-degree-of-freedom study, which this work does not provide.

2.1. Actuator model and energy

Equation 2 is idealized: a physical thruster also has a dead band, delivers less thrust in reverse, and consumes power non-quadratically. A non-ideal map with dead band 0.6 N and reverse gain 0.8, together with an actuator-disc power law $P = n_{thr}\, c_p\, (|u_{act}| / n_{thr})^{3/2} + P_{hotel}$ over 4 vertical thrusters, is therefore used in Section 6.4 to check that neither idealization carries the result. That law is a model, not a measurement, so conclusions drawn from the quadratic effort term of Equation 13 are statements about actuation demand rather than battery life.

3. Controllers

The PI$^{\lambda}$D$^{\mu}$ control law considered in this study is

$$u_{\mathrm{cmd}}(t) = K_p e(t) + K_i \mathcal{I}^{\lambda}\{e\}(t) + K_d \mathcal{D}^{\mu}\{e\}(t), \tag{4}$$

which reduces to the classical PID controller when $\lambda = \mu = 1$.

For the numerical implementation adopted in this study, the fractional operators are approximated using the band-limited Oustaloup recursive approximation [5, 20, 21]. The approximation parameters are set to $N = 4$ and $\omega \in [0.05{,}50]$ rad/s for all simulations:

$$s^{\alpha} \approx \omega_h^{\alpha} \prod_{k=-N}^{N} \frac{s + \omega'_k}{s + \omega_k}, \tag{5}$$

where

$$\omega'_k = \omega_b \left(\frac{\omega_h}{\omega_b}\right)^{\frac{k+N+\frac{1}{2}(1-\alpha)}{2N+1}}, \qquad \omega_k = \omega_b \left(\frac{\omega_h}{\omega_b}\right)^{\frac{k+N+\frac{1}{2}(1+\alpha)}{2N+1}}. \tag{6}$$

In the Laplace domain, the fractional integral and derivative operators correspond to $s^{-\lambda}$ and s^{μ}, respectively. A direct band-limited Oustaloup approximation of $s^{-\lambda}$ does not retain an exact pole at the origin. Consequently, exact integral action, and therefore zero steady-state error in the implemented closed loop, is not in general guaranteed, as also observed in previous fractional-order controller realizations [8].

To retain an exact integrator, the fractional integral operator is factorized as

$$s^{-\lambda} = s^{-1} s^{1-\lambda}. \tag{7}$$

The factor s^{-1} is implemented as a true integrator, whereas the residual factor $s^{1-\lambda}$ is realized using the Oustaloup approximation. This realization applies to both sides of $\lambda = 1$.

For

$$0 < \lambda < 1{,}1 - \lambda > 0, \tag{8}$$

and the residual term $s^{1-\lambda}$ acts as a fractional differentiating factor. At $\lambda = 1$, the residual factor is unity. For $\lambda > 1$,

$$1 - \lambda < 0, \qquad s^{1-\lambda} = s^{-(\lambda-1)}, \tag{9}$$

so the residual term becomes an additional fractional integrating factor. The Oustaloup approximation is applied to this negative-order factor while the exact integer integrator s^{-1} is retained.

The derivative operator is implemented according to the value of μ. For $0 < \mu \leq 1$, the operator s^{μ} is approximated directly using the Oustaloup method. For $\mu > 1$, it is factorized as

$$s^{\mu} = s\, s^{\mu-1}. \tag{10}$$

The residual fractional factor $s^{\mu-1}$ is approximated using the Oustaloup method, whereas the integer-order derivative factor is implemented in filtered form as

$$s \rightarrow \frac{\omega_f s}{s+\omega_f}, \qquad \omega_f = 50 \text{ rad s}^{-1}. \tag{11}$$

Thus, for $\mu > 1$, the implemented derivative branch consists of the cascade of a filtered integer-order derivative and an Oustaloup-approximated residual fractional operator.

Search bounds are $K_p \in [0.1,600]$, $K_i \in [0.01,60]$, $K_d \in [0.1,250]$, $\lambda \in [0.30,1.10]$, and $\mu \in [0.05,1.10]$; they were widened until the baseline optimum was interior (Section 5).

3.1. Discrete realization and equivalence of the two controllers

Both controllers run at 50 Hz, every band-limited operator is a cascade of nine first-order sections discretized with the bilinear transform, and the exact integrator uses the trapezoidal rule. A difference in numerical realization between the two would be indistinguishable from a difference in control structure, and there is none: the classical PID is the $\lambda = \mu = 1$ member of the same parameterization, evaluated by the same code path. At $\lambda = 1$ the Oustaloup construction places every zero exactly on its pole, so the residual filter is identically unity, and at $\mu = 1$ the derivative branch is the same nine-section differentiator. The two share the derivative roll-off, the discretization error and the anti-windup rule.

That equivalence was measured, not assumed. At $\mu = 1 + 10^{-4}$ the derivative branch becomes the first-order filtered derivative of a textbook PID; tuning with μ pinned there gives a mission cost differing from the band-limited PID by -0.04 to +0.00 percent over ten runs per profile, Mann-Whitney p between 0.28 and 0.97. The integer controller is not handicapped by the shared realization.

3.2. Anti-windup

Because the thrust command saturates for part of every dive, conditional integration [22] is applied identically to both controllers: the trapezoidal update is taken only when the loop is unsaturated, or when it is saturated and the update moves the state away from the active limit. Only the true integrator needs this; the residual factor $s^{1-\lambda}$ and the operator s^{μ} are band-limited approximations with all poles in the open left half plane, so they cannot accumulate an unbounded state.

4. Mission-Oriented Objective Function

Let $e = z_{\text{ref}} - z$, let the simulation horizon be T, and let the finite disturbance event occupy $t_d \leq t < t_e$. The proposed objective combines conventional performance measures with mission-specific terms.

The first term is the classical ITAE, widely used in controller tuning to penalize errors that persist at later times [7, 8]:

$$J_{\text{ITAE}} = \int_0^T t\,|e(t)|\,dt. \tag{12}$$

A quadratic control-effort term is used to penalize sustained actuation demand:

$$J_{\text{CE}} = \int_0^T u_{\text{act}}^2\,(t)\,dt. \tag{13}$$

This term is a proxy for actuation demand rather than a measure of electrical energy, which would require a thrust-to-power map of the specific thruster; Section 2.1 gives such a map and Section 6.4 reports how well the proxy tracks it.

To discourage rapid command variations and actuator chattering, a control-smoothness term is included:

$$J_S = \int_0^T \dot{u}_{\text{cmd}}^2\,(t)\,dt. \tag{14}$$

To encode the direction-dependent consequences of depth excursions, this study defines an asymmetric overshoot penalty. Let t_c denote the first crossing of the reference after the depth command is applied. The directional excursions are defined as

$$\mathrm{OS}_{\downarrow} = \max\left\{0, \max_{t_c \le t \le T}\left(z(t) - z_{\mathrm{ref}}(t)\right)\right\}, \qquad \mathrm{OS}_{\uparrow} = \max\left\{0, \max_{t_c \le t \le T}\left(z_{\mathrm{ref}}(t) - z(t)\right)\right\}. \tag{15}$$

The asymmetric overshoot cost is then

$$J_{\mathrm{OS}} = k_{\downarrow}\mathrm{OS}_{\downarrow} + k_{\uparrow}\mathrm{OS}_{\uparrow}. \tag{16}$$

Here, $\mathrm{OS}_{\downarrow}$ represents excessive penetration below the commanded depth and therefore increased bottom-collision risk during near-seabed operation, whereas $\mathrm{OS}_{\uparrow}$ represents rebound above the commanded depth. The coefficients $k_{\downarrow}$ and $k_{\uparrow}$ allow the penalty to reflect mission-dependent directional risk.

Finally, to quantify recovery after the finite current event, this study defines the post-event recovery index

$$J_{\mathrm{REC}} = \int_{t_e}^{T} (t - t_e)\,|e(t)|\,dt. \tag{17}$$

Unlike a sustained disturbance-rejection metric, this term is evaluated only after the disturbance has been removed at t_e.

The proposed mission-oriented objective is formed as the weighted sum of the baseline-normalized terms:

$$J = w_1 \frac{J_{\mathrm{ITAE}}}{J_{\mathrm{ITAE}}^{\mathrm{ref}}} + w_2 \frac{J_{\mathrm{CE}}}{J_{\mathrm{CE}}^{\mathrm{ref}}} + w_3 \frac{J_S}{J_S^{\mathrm{ref}}} + w_4 \frac{J_{\mathrm{OS}}}{J_{\mathrm{OS}}^{\mathrm{ref}}} + w_5 \frac{J_{\mathrm{REC}}}{J_{\mathrm{REC}}^{\mathrm{ref}}}. \tag{18}$$

Each reference value comes from a baseline PID tuned by PSO against pure ITAE under the same scenario, with the overshoot reference using that profile's asymmetry coefficients, so every term equals unity for the baseline and the weights read as relative mission priorities. Normalization is a prerequisite for any weighted sum to be meaningful [14], but it defines the scale by a controller of one of the two families compared; Section 6.2 answers that objection empirically.

Because prolonged operation at the thrust limit indicates persistent actuator demand, the saturation percentage is additionally reported as an auxiliary metric:

$$R_{\mathrm{sat}} = \frac{100}{T}\int_0^T \mathbb{I}\left(|u_{\mathrm{cmd}}(t)| \ge u_{\max}\right) dt \quad [\%]. \tag{19}$$

This quantity is reported for every tuned controller but is not included in the optimization objective.

4.1. Choosing the weights

The weights and asymmetry coefficients are design inputs chosen by the operator, so the procedure that produces them is part of the method. Terms are first ranked by mission consequence: in near-bottom inspection a downward excursion risks seabed contact, oscillatory thrust disturbs the sediment being imaged, and a slower approach costs only time, so overshoot outranks smoothness, which outranks speed. The ranking is then placed on the normalized scale of Equation 18, where a weight of 1.5 accepts a 50% rise in every other unit-weighted term in exchange for halving that one. Finally k_{down} is a risk ratio: $k_{\mathrm{down}} = 4$ encodes a judgement that a downward excursion is about 4 times as costly as an upward one. No claim is made that these values are optimal; Section 6.1 shows the conclusions are insensitive to them.

5. Optimization Setup and Results

5.1. Optimization setup and budget adequacy

PSO [9] is used throughout with inertia 0.72, cognitive and social coefficients 1.49 and velocity clamped to 20% of each range; GWO [10] and DE use the same budget. Ten independent runs are performed per cell and significance is assessed with the exact two-sided Mann-Whitney U test [23] and the rank-biserial correlation [24]. Following the reporting guidelines for metaheuristic comparisons [25], p-values are Holm-corrected across the eight primary comparisons and the Vargha-Delaney A measure [26] is reported as effect size: every primary comparison below has $A = 1.00$ and corrected $p \le 8.7 \times 10^{-5}$. The search bounds were widened until the baseline ITAE-PID optimum ($K_p = 193.3$, $K_i = 13.8$, $K_d = 53.2$) lay strictly inside them. The budget is itself a design choice, and an inadequate one does not bias the two controllers equally: the three-dimensional PID search reaches its optimum within a few hundred evaluations, whereas the five-dimensional fractional search keeps improving to roughly 1248 (Table 2). A budget of 24 individuals and 100 iterations is therefore used everywhere, against the 12 by 25 of earlier runs; the change raises the fractional advantage, so the earlier figures were conservative.

5.2. Objective-complexity ablation

To test directly whether richer objectives unlock the value of the fractional degrees of freedom, the cost terms of Equation 18 are added cumulatively with unit weights and symmetric overshoot ($k_{\downarrow} = k_{\uparrow} = 1$):

$$J_0 = \tilde{J}_{\mathrm{ITAE}}, \quad J_1 = J_0 + \tilde{J}_{\mathrm{CE}}, \quad J_2 = J_1 + \tilde{J}_S, \quad J_3 = J_2 + \tilde{J}_{\mathrm{OS}}, \quad J_4 = J_3 + \tilde{J}_{\mathrm{REC}}. \tag{20}$$

where the tilde denotes baseline normalization. PID and $\mathrm{PI}^{\lambda}\mathrm{D}^{\mu}$ are re-tuned from scratch at every stage (three

independent PSO runs each). Table 2 and Figure 1 present the results.

To quantify the relative benefit of the fractional-order controller at each objective-complexity level, the relative cost reduction achieved by the PI$^{\lambda}$D$^{\mu}$ controller with respect to PID is defined as

$$\Delta_i = 100\frac{\bar{J}_{i,\mathrm{PID}} - \bar{J}_{i,\mathrm{PI}^{\lambda}\mathrm{D}^{\mu}}}{\bar{J}_{i,\mathrm{PID}}}, \qquad i = 0, \ldots, 4, \tag{21}$$

where $\bar{J}_{i,\mathrm{PID}}$ and $\bar{J}_{i,\mathrm{PI}^{\lambda}\mathrm{D}^{\mu}}$ denote the mean optimized costs obtained over the independent PSO runs for objective J_i. Accordingly, $\Delta_i > 0$ indicates a lower optimized cost for the PI$^{\lambda}$D$^{\mu}$ controller relative to PID, with its magnitude representing the corresponding percentage reduction.

Table 2. Objective-complexity ablation, mean ± standard deviation over ten PSO runs.

Objective	PID J	PI$^{\lambda}$D$^{\mu}$ J	PID R_{sat} [%]	PI$^{\lambda}$D$^{\mu}$ R_{sat} [%]	Improvement [%]
J0 (ITAE)	1.000 ± 0.000	0.930 ± 0.020	17.8	16.0	7.0
J1 (+CE)	1.949 ± 0.006	1.826 ± 0.010	12.3	10.2	6.3
J2 (+S)	2.386 ± 0.000	2.051 ± 0.104	8.0	7.3	14.1
J3 (+OS)	3.359 ± 0.000	2.634 ± 0.101	8.9	6.8	21.6
J4 (+REC)	4.491 ± 0.002	3.470 ± 0.319	12.2	7.9	22.7

Three observations follow from Table 2. Under plain ITAE the two extra parameters yield only 7.0%, a margin that would not justify implementing two fractional operators. The reduction then grows to 22.7% on the mean and 30.9% on the best run under the full cost, but not monotonically: adding the effort term alone (6.3%) leaves the margin where the tracking objective left it, because a quadratic effort penalty is minimized by lowering the loop gain, which both families can do. The margin opens only from J_2 onward, once the objective prices high-frequency command activity and directional excursions. Third, saturation time falls from 17.8% for the baseline to 7.9% for the full-cost PI$^{\lambda}$D$^{\mu}$.

The mechanism is structural. The optimizer selects a derivative order between 0.45 and 0.82, supplying phase lead comparable with an integer derivative while its magnitude rises as ω^{μ} rather than ω, so its gain at the top of the sensor-noise band is smaller by roughly a factor of 4: it buys damping without turning the 1 cm depth noise into thruster chatter. This is not simply a lower gain. On the inspection profile the fractional design matches the tuned PID tracking error to within one percent, yet halves the smoothness term and improves post-event recovery by 44%, and recovery needs more loop authority, not less.

5.3. Mission profiles

Three representative missions are encoded as weight profiles of Equation 18 (Table 3): survey/transit (fast, balanced), near-bottom inspection (smooth actuation, strong aversion to downward overshoot, fast recovery), and long-endurance (effort-dominated).

Table 3 gives the three weight profiles, Table 4 the tuned controllers with their statistics, and Figure 1 the depth responses. The PI$^{\lambda}$D$^{\mu}$ controller attains the lower cost on every profile in every run, so the samples are completely separated, with rank-biserial correlation 1.0; the largest margin, 35.4% (bootstrap 95% confidence interval 34.2 to 36.7%), occurs on inspection, whose cost simultaneously demands smoothness, asymmetric overshoot aversion and fast recovery. The optimizer selects an integral order near unity on all three profiles, so the useful fractional freedom here is the derivative one, and saturation time falls on every profile.

Table 3. Mission profiles: weights of Equation 18 and overshoot asymmetry.

Profile	w1 (ITAE)	w2 (CE)	w3 (S)	w4 (OS)	w5 (REC)	k↓	k↑
Survey	1.0	0.3	0.1	0.3	0.3	1.0	1.0
Inspection	0.6	0.2	0.8	1.5	1.0	4.0	1.0
Endurance	0.5	1.5	0.4	0.3	0.3	1.0	1.0

Table 4. Tuned controllers, ten-seed statistics and selected metrics per mission profile.

Profile	Ctrl	J (mean ± std)	K_p	K_i	K_d	λ	μ	R_{sat} [%]	Impr. [%]
Baseline	PID	—	193	13.8	53	—	—	17.8	—
Survey	PID	1.893 ± 0.000	151	44.5	53	—	—	16.6	—
Survey	$PI^{\lambda}D^{\mu}$	1.612 ± 0.028	210	5.0	97	0.99	0.76	11.1	14.8
Inspection	PID	3.760 ± 0.009	121	22.3	45	—	—	12.3	—
Inspection	$PI^{\lambda}D^{\mu}$	2.430 ± 0.076	169	0.4	111	1.10	0.65	9.0	35.4
Endurance	PID	2.715 ± 0.001	99	28.8	35	—	—	9.1	—
Endurance	$PI^{\lambda}D^{\mu}$	2.196 ± 0.042	69	1.2	61	1.00	0.72	7.3	19.1

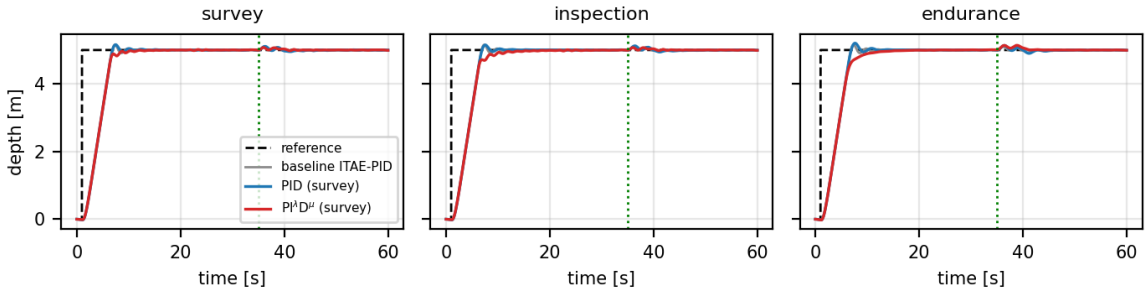


Figure 1. Depth responses per mission profile. Dashed: reference; grey: baseline ITAE-PID; blue: mission-tuned PID; red: mission-tuned $PI^{\lambda}D^{\mu}$. The dotted line marks the current event.

5.4. Cross-algorithm comparison

Earlier runs compared PSO with a single GWO run. Here GWO and differential evolution (DE) [27] are each run ten times per cell at the same budget. The three optimizers agree on the ranking on every profile, their improvement figures agree to within 2.3 percentage points, and the convergence histories of Figure 2 show the same picture. This is not a comparison of optimization algorithms; it establishes that the ranking is not an artefact of the metaheuristic.

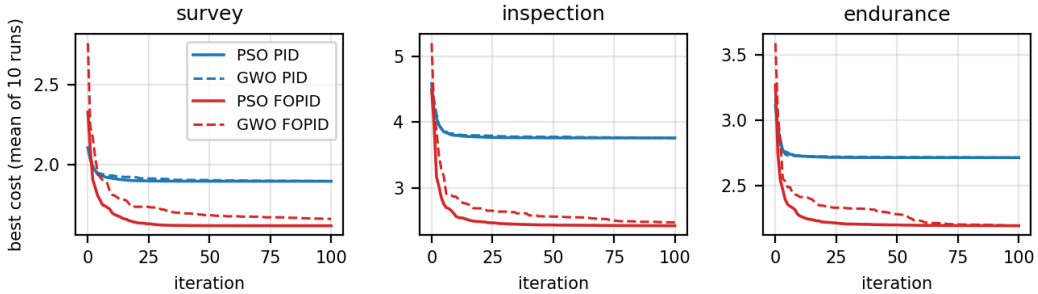


Figure 2. Convergence of PSO and GWO per mission profile, mean over ten independent runs.

6. Robustness of the Conclusions

6.1. Weight sensitivity

The results above rest on designer choices - weights, normalization, scenario, plant and actuator idealizations - and each is varied below with both controllers re-tuned or re-evaluated under the change. Perturbing each of the five weights of each profile by ±50% and re-tuning both controllers from scratch for all 30 resulting weight vectors

leaves the fractional controller ahead in every one, by 11.2 to 39.0%. Drawing the weights at random over more than a decade and re-tuning for each of 160 draws leaves it ahead in 159 of them (99.4%), median improvement 25.9%. Sweeping k_{down} alone over [1, 8] moves the margin only between 34.1 and 38.6%, so the value 4 is not load-bearing.

The random sweep also identifies where the advantage comes from, which a cumulative ablation cannot. The Spearman correlations between each weight and the resulting margin are -0.55 for tracking, -0.29 for effort, +0.39 for smoothness, +0.53 for overshoot and +0.06 for recovery, whereas the correlation with the number of active terms, +0.03, is not significant. Weight on the plain error integral works against the fractional controller and weight on command smoothness and directional overshoot works for it; this also explains the single adverse draw, whose weight vector is almost pure ITAE. Because rank correlations can miss interactions, total-effect Sobol indices were also estimated over the same domain with a Saltelli design [28] of 512 re-tuned weight vectors: 0.49 for the tracking weight, 0.32 and 0.43 for the smoothness and overshoot weights, every other factor below 0.21, so the variance decomposition and the correlations tell the same story. Repeating the ablation with the terms added in three different orders confirms it: no ordering gives a monotone curve, the largest jump always follows the smoothness or overshoot term, and adding the effort term last makes the margin fall from 32.3% to 22.7%. The gap is governed by which terms are priced, not by how many.

6.2. Normalization

Normalizing by a baseline PID defines the cost scale using a member of one of the two families compared. The study was therefore repeated under four schemes: the baseline ITAE-PID; a reference PID tuned on a different index; a specification-based scheme using a 5 cm depth tolerance, the hover thrust, the fastest useful command slew and a 10 cm excursion, with no reference controller at all; and a min-max scheme from the 5th-to-95th-percentile range over 4000 random controllers. Absolute costs and tuned gains move, the ranking does not: the fractional controller wins under all four schemes on all three profiles, by 12 to 69%.

6.3. Stability margins and what the objective does not constrain

A low cost from a metaheuristic is not evidence of a well-damped loop. Linearizing about the depth-holding equilibrium $w = 0$, where the quadratic drag contributes no damping, and closing the loop with the exact discrete controller gives the margins of Table 5, and they are not what the time responses suggest: the phase margins (PM) lie between 7 and 22 degrees and the peak sensitivities M_s between 2.7 and 9.1, against the 30 to 60 degrees and M_s below 2 conventionally required, and on two of the three profiles the fractional design has the smaller margin. The reason the responses of Figure 1 look well damped is thrust saturation: a 5 m dive commands hundreds of newtons from a 40 N actuator, so the transient is rate-limited rather than linear. A 5 cm step, which never saturates, peaks at 110 to 164% of the command.

This is a defect of the objective, not of either controller family, and one a metaheuristic will exploit: Equation 18 prices behaviour on one large step and carries no frequency-domain specification, so it cannot tell a design that is fast because it is well damped from one that is fast because it saturates. The remedy is to put the specification inside the tuning problem. Re-tuning both controllers with a hard requirement of $PM \geq 40$ degrees and $M_s \leq 2.0$, imposed as a relative-violation penalty, makes both families meet it by construction, removes the small-signal overshoot entirely, and costs 2 to 35% of mission performance (Table 5). The constraint narrows the fractional advantage and removes it on the one profile whose cost is dominated by control effort, while leaving it intact on the two that price command smoothness and directional overshoot, so the size of the fractional advantage is contingent on the constraint set as well as on the weights and should be quoted together with both. Under transport delay the fractional designs are also the more fragile of the two on the survey and inspection profiles, tolerating 40 to 60 ms against 140 ms, which has the same cause: they place their gain crossover higher.

Table 5. Linear stability margins of the implemented discrete loops - gain margin GM, phase margin PM, peak sensitivity M_s - and the effect of re-tuning with the margin requirement inside the objective.

Controller	PM [deg]	M_s	delay [ms]	step peak [%]	PM [deg]	M_s	delay [ms]	step peak [%]	J rise [%]
Survey PID	18.1	3.19	117	146	40.0	1.63	250	84	18
Survey $PI^{\lambda}D^{\mu}$	7.8	7.53	37	160	40.0	1.61	306	60	14
Inspection PID	20.3	2.86	147	134	44.7	1.50	316	75	21
Inspection $PI^{\lambda}D^{\mu}$	6.5	9.05	32	164	39.9	1.58	332	56	35
Endurance PID	15.8	3.65	132	149	55.1	1.31	510	38	2
Endurance $PI^{\lambda}D^{\mu}$	22.4	2.72	163	110	40.0	1.65	294	100	34

6.4. Generalization, uncertainty and actuator realism

The remaining idealizations are examined together in Table 6, which lists what was varied, how many cases were evaluated and how the margin moved. Every candidate was tuned on one scenario with one frozen noise sequence, so the mission-profile controllers were re-evaluated without re-tuning on 500 fresh noise realizations, on 500 plants with every hydrodynamic and actuator parameter perturbed by up to ±30%, and on eleven held-out scenarios the optimizer never saw, covering step amplitudes of 2 to 8 m, an ascent, a staircase, a ramp, sinusoidal tracking, three disturbance amplitudes and a shifted event. Nothing diverged, and the ranking holds in 31 of the 33 scenario-profile combinations.

Three checks remain. The non-dominated sorting genetic algorithm NSGA-II [11], run on each family with the conflicting terms as a vector objective, instead of scalarizing them, gives a fractional front dominating 97% of the integer front in two objectives and 78% in three, the fronts nearly coinciding at the extremes and separating in the interior. The non-ideal thruster of Section 2.1 raises every cost but leaves the margin unchanged, and the effort term tracks modelled electrical energy with a Spearman correlation of 0.999. Representing the neglected axes as a demand-proportional loss of vertical thrust and a velocity-dependent coupled force, swept as admitted uncertainties, leaves the ranking intact up to a thrust loss of 30% with a coupled force of 30% of the thrust limit: a bound on tolerable coupling, not a six-degree-of-freedom study.

Table 6. Verification summary: what was varied, how many cases were evaluated, and the resulting range of the improvement of $PI^{\lambda}D^{\mu}$ over PID.

What was varied	Cases	Improvement range [%]	Ranking held
Weights: ±50% one-at-a-time + random draws (5th-95th pct)	190	10.5 to 42.7	189 of 190
Normalization scheme	12	12.0 to 68.5	all
Measurement noise / plant parameters (±30%)	2 × 500	11.4 to 34.1	all
Held-out scenarios	33	-8.3 to 43.4	31 of 33
Non-ideal thruster	3	19.3 to 34.4	all
Unmodelled axis coupling	48	-18.7 to 54.5	46 of 48
Third optimizer (DE)	60	16.9 to 37.3	all
Pareto fronts (NSGA-II)	2 problems	97% / 78% front dominance	both

7. Discussion and Limitations

Three strands bound this work: fractional-order control of underwater vehicles [8, 29], which shows the extra freedom can help but evaluates it against a fixed criterion; multi-objective tuning [12, 13], which solves a given problem rather than asking where it comes from; and energy-aware underwater control [16, 17], which motivates terms but treats the criterion as given. The transferable finding here is negative in form: a comparison against a

single generic index understates the fractional controller by roughly a factor of 5 here, because such an index does not price what the fractional operator trades - high-frequency command activity against phase lead.

Four rules follow for practice: do not expect a fractional controller to repay its implementation cost when the criterion is a single error integral; include a command-rate term whenever the measurement is noisy and the actuator bandwidth is finite; set the weights on the normalized scale and do not refine them beyond a factor of about two; and put the stability specification inside the objective rather than checking it afterwards.

The limitations bound the scope of the conclusion rather than qualifying it. The dynamics are one-dimensional, and neither Section 2 nor the coupling bound of Section 6.4 replaces a coupled six-degree-of-freedom study with identified coefficients. The plant parameters are representative literature values, so all costs are relative, and no result has been validated on hardware. The objective is a scalarization; a decision-maker who needs the trade-off surface should solve the multi-objective problem directly. The thruster model omits the PWM-to-thrust characteristic and supply droop, so the energy figures are comparative only.

8. Conclusion

This paper treated the tuning objective, rather than the controller or the optimizer, as the primary design object for underwater depth control. On a one-degree-of-freedom BlueROV2 heave model with saturation, actuator lag and sensor noise, the advantage of $PI^{\lambda}D^{\mu}$ over PID rose from 7.0% under plain ITAE to 22.7% under a five-term mission cost, and encoding survey, inspection and endurance missions as weight profiles gave a fractional advantage on every profile and in every run, largest at 35.4% on the conflict-rich inspection mission. What unlocks the fractional degrees of freedom is not the number of terms in the cost but their identity: the margin follows the terms pricing command activity and directional excursions, while an effort penalty works mildly the other way.

Because that conclusion depends on designer choices, it was tested against them, surviving 160 random weight vectors, four normalization schemes, a Pareto analysis, 500 noise realizations, 500 parameter draws at ±30%, eleven held-out scenarios and a non-ideal thruster. One negative result is reported rather than hidden: the objective carries no frequency-domain requirement, so the designs it selects have small phase margins that saturation masks; an explicit margin requirement restores them at a cost in mission performance and narrows where the fractional advantage holds. Future work is ordered accordingly: coupled six-degree-of-freedom dynamics, hardware-in-the-loop and in-water experiments on the FPGA architecture of [6], and online adaptation of the weight profile to the mission phase.

Acknowledgement

This work was supported by The Scientific and Technological Research Council of Turkiye (TUBITAK), Department of Science Fellowships and Grant Programmes (BIDEB) through the 2232-B International Fellowship for Early Stage Researchers Program under Grant Number 121C126.

References

[1] Fossen TI. Handbook of Marine Craft Hydrodynamics and Motion Control. Chichester, UK: Wiley, 2011.
[2] von Benzon M, Sørensen FF, Uth E, Jouffroy J, Liniger J, Pedersen S. An open-source benchmark simulator: Control of a BlueROV2 underwater robot. J Mar Sci Eng 2022; 10(12): 1898.
[3] Podlubny I. Fractional-order systems and $PI^{\lambda}D^{\mu}$-controllers. IEEE Trans Autom Control 1999; 44(1): 208-214.
[4] Shah P, Agashe S. Review of fractional PID controller. Mechatronics 2016; 38: 29-41.
[5] Oustaloup A, Levron F, Mathieu B, Nanot FM. Frequency-band complex noninteger differentiator: characterization and synthesis. IEEE Trans Circuits Syst I Fundam Theory Appl 2000; 47(1): 25-39.
[6] Akgul B, Güler E, Ersoy H, Kartci A, Ayten UE. Resource-efficient design of high-speed FIR filter: Fractional-order integrator/differentiator on FPGAs. Circuits Syst Signal Process 2026; 1-24.
[7] Ersoy H, Akgül B, Akpinar E, Kartci A, Ayten UE. A robust speed controller design for PMSM using $PI^{\lambda}D^{\mu}$ controllers. In: 2024 47th International Conference on Telecommunications and Signal Processing (TSP). New York, NY, USA: IEEE, 2024. pp. 289-293.
[8] Ersoy H, Akgül B, Akpınar E, Kartci A, Ayten UE. Design and implementation of $PI^{\lambda}D^{\mu}$ controller for ROVs: Thruster modeling, controller parameter optimization, and FPGA realization. Eng Sci Technol Int J 2026; 73: 102261.
[9] Kennedy J, Eberhart R. Particle swarm optimization. In: Proceedings of ICNN'95 - International Conference on Neural Networks. New York, NY, USA: IEEE, 1995. pp. 1942-1948.
[10] Mirjalili S, Mirjalili SM, Lewis A. Grey wolf optimizer. Adv Eng Softw 2014; 69: 46-61.
[11] Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans Evol Comput 2002; 6(2): 182-197.
[12] Joseph SB, Dada EG, Abidemi A, Oyewola DO, Khammas BM. Metaheuristic algorithms for PID controller parameters

tuning: review, approaches and open problems. Heliyon 2022; 8(5): e09399.
[13] Tian T, Huang Q, Lv J, Zhou M, Tang J. Application of multiobjective optimization based on modified SPEA2 method for tuning FOPID plus second-order derivative controller in hydraulic turbine governing system. J Electr Comput Eng 2024; 2024: 4217192.
[14] Marler RT, Arora JS. Survey of multi-objective optimization methods for engineering. Struct Multidiscip Optim 2004; 26(6): 369-395.
[15] Das I, Dennis JE. A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. Struct Optim 1997; 14(1): 63-69.
[16] Yao F, Yang C, Zhang M, Wang Y. Optimization of the energy consumption of depth tracking control based on model predictive control for autonomous underwater vehicles. Sensors 2019; 19(1): 162.
[17] Pinto JB, Carneiro JF, de Almeida FG, Cruz NA. Variable structure depth controller for energy savings in an underwater device: proof of stability. Actuators 2025; 14(7): 340.
[18] Türetken AE, Ersoy H, Kartci A. Energy-efficient and actuator-friendly control under wave disturbances: Model reference vs. PID for thruster surge. In: 2025 16th International Conference on Electrical and Electronics Engineering (ELECO). New York, NY, USA: IEEE, 2025. pp. 1-5.
[19] Ng P, Krieg M. Modifications to ArduSub that improve BlueROV SITL accuracy and design of hybrid autopilot. Appl Sci 2024; 14(17): 7453.
[20] Öztürk Ö, Çavdar B. Otomatik gerilim regülatörü sistemi denetleyici tasarımı için meta-sezgisel algoritmaların performansı. Karadeniz Fen Bilimleri Dergisi 2024; 14(4): 2258-2289.
[21] Li J, Liu Y, Shan C, Dai C. Implementation of simplified fractional-order PID controller based on modified Oustaloup's recursive filter. J Shanghai Jiaotong Univ Sci 2020; 25(1): 44-50.
[22] Bohn C, Atherton DP. A SIMULINK package for comparative studies of PID anti-windup strategies. In: Proceedings of the IEEE/IFAC Joint Symposium on Computer-Aided Control System Design. New York, NY, USA: IEEE, 1994. pp. 447-452.
[23] Mann HB, Whitney DR. On a test of whether one of two random variables is stochastically larger than the other. Ann Math Stat 1947; 18(1): 50-60.
[24] Cureton EE. Rank-biserial correlation. Psychometrika 1956; 21(3): 287-290.
[25] Derrac J, Garcia S, Molina D, Herrera F. A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms. Swarm Evol Comput 2011; 1(1): 3-18.
[26] Vargha A, Delaney HD. A critique and improvement of the CL common language effect size statistics of McGraw and Wong. J Educ Behav Stat 2000; 25(2): 101-132.
[27] Storn R, Price K. Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. J Glob Optim 1997; 11(4): 341-359.
[28] Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M, Tarantola S. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. Comput Phys Commun 2010; 181(2): 259-270.
[29] He L, Zhang Y, Xie M, Yuan Z, Bai C. Three-dimensional trajectory tracking control of underactuated AUV based on fractional-order PID and super-twisting extended state observer. Fractal Fract 2025; 9(9): 580.