Racing Without a Map: Bearings-Only Guidance and Optical Looming for Single-Course Autonomous Drone Racing under a Blocked-Telemetry Contract

AI-GP Autonomy Team

Abstract

We address autonomous drone racing on a single, fixed, *unsurveyed* course under a competition contract that blocks all pose and gate telemetry. The agent receives a monocular 640×360 camera at 30 Hz, inertial measurements, and its own command history; it emits collective-thrust-and-body-rate (CTBR) commands. This invalidates the dominant paradigm — generate a time-optimal trajectory through surveyed gates, then track it with a state estimate — because neither the survey nor the state estimate is admissible.

We report a structured survey spanning 1958–2026 and roughly 400 works, and three findings. **First**, of the entire modern drone-racing literature, only four published systems are genuinely in-regime for camera+IMU-only operation without odometry; both champion-level monocular systems (Swift, MonoRace) rely on a track map and are therefore inadmissible as-is. **Second**, four mature but mutually disconnected literatures already solve the sub-problems: optical-τ time-to-contact guidance for aircraft; bearings-only gap traversal from the guidance community; image-based visual servoing through windows; and image-space flight through multiple openings. We find no work joining them to autonomous racing, and no pre-2013 vision-only aperture traversal in any language. **Third**, four quantitative facts follow from the competition constants alone, of which one dominates: at the measured 0.109 s control period, a single control step consumes 0.87 m of open-loop travel at $8\,\mathrm{m\,s^{-1}}$ against a 0.61 m aperture half-width, imposing a hard speed ceiling near $5.6\,\mathrm{m\,s^{-1}}$ that no policy can overcome.

We propose a system built on the union of those literatures: a deterministic expert combining bearings-only guidance with optical-τ looming, a drag-aware attitude estimator, gate-corner reprojection injected directly into an error-state filter, and a *speed-conditioned* recurrent policy specialized to the single course by multi-fidelity Bayesian optimization under a three-attempt budget. We give the frozen observation contract, the randomization schedule, the specialization protocol, the metric ladder, and explicit falsifiers.

1 Introduction

Autonomous drone racing is the standard benchmark for agile robot autonomy. Reported top speeds rose from $0.6\,\mathrm{m\,s^{-1}}$ at the inaugural IROS 2016 Autonomous Drone Race to $28\,\mathrm{m\,s^{-1}}$ at the 2025 A2RL×DCL championship, where a monocular system beat three human world champions.

Our setting differs from every published system in one decisive respect. The governing contract blocks not only external localization but *any surveyed knowledge of the course*. The following inbound streams are unavailable: `ATTITUDE`, `LOCAL_POSITION_NED`, `ODOMETRY`, `GATE_INFO`. The deploy policy additionally may not consume `TRACK_INFO`, scorer state, collision flags, the official active-gate ordinal, or surveyed gate coordinates.

Both champion-level monocular systems rely on a track map: Swift fuses gate detections "in combination with a map of the track" via camera resectioning; MonoRace solves Perspective-n-Point against assumed gate locations. Neither transfers. What does transfer is older and comes from outside robotics.

1.1 Contributions

1. A survey organized by lineage (§3), with an explicit record of what does *not* exist (§9) — including verified phantom citations that recur in this literature.
2. Three admissibility filters and a classification of surveyed methods against them (§4), yielding the finding that only four published systems are in-regime.
3. Four quantitative platform results from the competition constants (§2.2), including a control-rate speed ceiling (§2.2) that dominates all policy considerations.
4. Identification of an unoccupied junction between four mature literatures (§5), and a system design that occupies it (§6).
5. A specialization protocol for one fixed course under a three-attempt budget (§7), with metric ladder and falsifiers (§8).

2 Problem formalization

2.1 Admissible observations and actions

The deploy observation admits exactly: JPEG frames (640×360, nominal 30 Hz, UDP 5600, 24-byte header, simulator-nanosecond capture timestamps); allowed inertial measurements; synchronization and heartbeat signals with frame age; the previous outbound command; and internal filter and recurrent state. Blocked-message *absence* must be treated as healthy, not as a timeout.

The action is $u = [f,\ \omega_x^{\mathrm{des}},\ \omega_y^{\mathrm{des}},\ \omega_z^{\mathrm{des}}] \in [-1,1]^4$ via `SET_ATTITUDE_TARGET` with the attitude-ignore mask. That the contract fixes CTBR is fortunate rather than limiting: a controlled benchmark of learned control policies for agile flight found CTBR to dominate both linear-velocity and single-rotor-thrust action spaces on the agility/transfer trade-off, with linear-velocity requiring a full state-estimating stack and single-rotor-thrust transferring poorly. Swift and the pixels-to-control line both adopt CTBR. We treat the action space as settled.

2.2 Four quantitative platform results

Intrinsics are $f_x = f_y = 320$, $(c_x, c_y) = (320, 180)$:

$$\begin{aligned} \mathrm{HFOV} &= 2\arctan\frac{c_x}{f_x} = 90.0^\circ, \\ \mathrm{VFOV} &= 2\arctan\frac{c_y}{f_y} = 58.72^\circ, \end{aligned} \tag{1}$$

with the camera pitched $+20^\circ$ up. (The specification text elsewhere asserts VFOV $= 90^\circ$; the intrinsic matrix is authoritative and the two disagree — a discrepancy worth resolving with the organizers.) Gate outer size is 2.7 m, clear inner opening $W = 1.5$ m, frame depth 0.26 m; the airframe is $0.28 \times 0.28 \times 0.16$ m.

R1: FOV budget. The solid angle is roughly one third that of the A2RL-winning monocular system ($155^\circ \times 115^\circ$), and unlike that system there is no odometry to coast on when the gate leaves frame. Losing the gate is therefore terminal, not merely degrading. Perception-aware objectives are load-bearing.

R2: Aperture volume. Admissible gate-plane crossing points form a square of half-width

$$r_{\mathrm{ap}} = \tfrac{1}{2}\,(W - d_{\mathrm{drone}}) = \tfrac{1}{2}(1.5 - 0.28) = 0.61\,\mathrm{m}, \tag{2}$$

i.e. 1.22 m on a side. Waypoint-through-centre planning discards this volume entirely; gate-as-volume formulations recover it, and the gain compounds over N gates.

R3: Range conditioning. The inner opening projects to $w_{\mathrm{px}}(d) = f_x W/d = 480/d$, so $\delta d/d = -\delta w/w$ (Table 1). This justifies the existing 8 m gate on range-derived position fixes (60 px, 1.7% per pixel) and shows why a 23 m first-gate standoff (21 px, 4.8% per pixel) must remain bearing-only.

Table 1: Apparent inner-opening width and range sensitivity.

d (m)	w_{px}	1 px $\rightarrow$ range err.	subtense
3	160.0	0.62%	28.1°
5	96.0	1.04%	17.1°
8	60.0	1.67%	10.7°
10	48.0	2.08%	8.6°
15	32.0	3.12%	5.7°
20	24.0	4.17%	4.3°
23	20.9	4.79%	3.7°
30	16.0	6.25%	2.9°

R4: The control-rate ceiling. The specification permits command publication to 100 Hz; physics runs at 120 Hz and vision at 30 Hz. The calibrated surrogate, however, uses a control period of 0.109 s, implying a measured live loop near 9.2 Hz. Per control step the vehicle travels $v\Delta t$ open-loop; equating this to the aperture half-width gives

$$v^\star = \frac{r_{\mathrm{ap}}}{\Delta t} = \frac{0.61}{0.109} \approx 5.6\,\mathrm{m\,s^{-1}}. \tag{3}$$

Above $v^\star$ one control decision commits more lateral travel than the whole admissible aperture, so *no* policy can correct a centring error inside the gate. At 3.5 m s^{-1} the step is 0.38 m; at 8 m s^{-1} it is 0.87 m.

> *Raising the control rate toward the vision rate or above is worth more than any policy improvement for targets above ∼5 m s^{-1}. It is a loop-scheduling change, not a research problem.*

Caveat. 0.109 s is a surrogate calibration constant fitted to live traces, not an instrumented loop period. The first action implied by this paper is to measure the deployed loop. If it is near 30 Hz, $v^\star$ rises to 18 m s^{-1} and this result is moot; if near 9 Hz, it binds the entire speed programme.

Relation to known bounds. Equation (4) is an *aperture-limited* analogue of the only closed-form speed bound in the field: Falanga, Kim and Scaramuzza derive maximum safe speed as a function of perception latency, field of view, sensing range and vehicle agility. Their bound is set by the distance travelled before an obstacle can be reacted to; ours by the distance travelled before a centring error can be corrected within a finite aperture. Both are *sensing-and-actuation* limits, independent of policy quality, and both bind before any aerodynamic limit on this platform.

This matters beyond our system. Our survey found — and a second independent sweep confirmed — that **no work computes a theoretical minimum lap time or performance upper bound for quadrotor racing**. There is no drone analogue of motorsport's minimum-lap-time simulation and g–g diagram tradition. The machinery to build one exists: Falanga et al. supply the perception-latency term, Bauersfeld and Scaramuzza supply a first-principles power and optimal-speed envelope, Abdulrahim et al. measure the g/rate envelope expert FPV pilots actually use, and Majumdar and Pacelli supply rigorous information-theoretic "no policy can beat X" limits for sensor-based control — which have never been instantiated on a race track. Equation (4) is a single-gate special case. Assembling these into a perception-limited lap-time bound for a fixed course is, as far as this survey can establish, an open problem.

2.3 Objective and the legality of specialization

The objective is lexicographic: legal interface $\succ$ repeatable completion $\succ$ zero contacts $\succ$ elapsed time.

We distinguish two senses of specializing to the single course:

Admissible. Specializing the *parameters* of a policy whose *inputs* remain camera-, inertial- and command-derived. A network trained in a replica of this course, or a gain vector tuned on it, satisfies the input contract exactly as a network trained on a thousand courses does. Every champion-level racing system is specialized in this sense.

Inadmissible. Specializing the *inputs*: consuming surveyed gate coordinates, indexing behaviour by a gate ordinal, keying parameters to route position, or triggering on elapsed race time.

Governance flag. The project's standing instructions state that an internal visual passage count "cannot index a route or parameter table" and prohibit "gate-specific control parameters." Read strictly, that self-imposed rule forbids not only inadmissible input specialization but some admissible parameter specialization. The controlling specification bans consuming blocked telemetry and prohibits cheating or simulator manipulation; it does not obviously ban a policy whose weights encode one course. *This determines whether §7 exists at all and must be settled against the specification text before implementation.*

3 Related work

Appendix A gives the full classified bibliography. Here we summarize by lineage and state what each contributes under our constraints.

3.1 Polynomial trajectory generation

Differential flatness for quadrotors makes position and yaw flat outputs, so a trajectory in output space determines the full command sequence. The canonical instance minimizes snap over piecewise polynomials. Richter, Bry and Roy reformulate the joint spline optimization as an *unconstrained* QP over endpoint derivatives — numerically stable for high order and many segments, sparse — and select segment times by gradient descent from a single aggressiveness scalar. Successors repaired each weakness: the endpoint-derivative trick was generalized to a spatial-temporal parameterization with analytic gradients and linear complexity; time allocation was made convex on a fixed path, then bilevel, then learned; and Bernstein, B-spline and MINVO bases replaced dense constraint sampling with control-point containment.

Under our contract none of this is deployable: it needs a map and a state estimate. It remains valuable offline, as teacher, as an upper bound on available lap time, and as the progress term in a reward.

3.2 Time-optimal planning and gates as volumes

Complementary progress constraints eliminate time allocation entirely and exploit single-rotor thrust limits, establishing that smooth polynomials cannot reach the actuator boundary. Most importantly for us, time-optimal *gate-traversing* planners replace point waypoints with convex gate volumes, letting the optimizer choose where within each gate to pass — the formulation that recovers R2. Recent work adds camera field-of-view and perception-quality constraints to that formulation, raising closed-loop success on a Split-S track from 55% to 100%.

3.3 Model-predictive and classical control

Model predictive contouring control fuses time-optimal planning and control in one receding-horizon problem, with a later variant adding a tunnel constraint and terminal set for gate-collision guarantees at $> 80\,\mathrm{km\,h^{-1}}$, tuned by trust-region Bayesian optimization on lap time. All of it presumes full state feedback. At $< 100\,\mathrm{Hz}$ without odometry, the honest classical baselines are drag-aware flatness feedforward with incremental nonlinear dynamic inversion — which needs only gyro measurements, no position and no model — or geometric control on SE(3), whose attitude loop requires no position at all. A carefully controlled comparison finds that with equal tuning effort classical geometric control matches learned policies, implying many reported "RL wins" are tuning artifacts.

3.4 Reinforcement learning and vision-based policies

The progression from state-based racing policies through pixels-to-control is well documented. The estimation-free endpoint is instructive for how little suffices: one system's actor observes only a confidence-valued gate segmentation mask downsampled to 84×84 plus three previous actions — no IMU, no attitude, no velocity — with inner-edge projection used as a render-free training abstraction. Others use gate-corner pixel coordinates or ResNet embeddings, again without IMU; one stabilizes a thrown quadrotor from seven point-feature pixel coordinates, treating the unobservable velocity as a POMDP. The earliest instance maps a single grayscale image to steering and collision probability with no state input.

We nonetheless include inertial measurements. This is a deliberate departure: the contract permits IMU, the gyro is the only signal separating ego-rotation from gate motion in the image, and a drag-aware estimator makes body velocity partially observable. The minimal designs show the mask is *sufficient*, not that it is optimal when an IMU is free.

A relevant null result: a hierarchical optimal-control-initialized racing policy reports that its model-free baseline *never completes the track*. That is the strongest published argument for an expert-first ordering.

3.5 Perception, estimation, and the map problem

Gate perception has converged on corners and inner edges. The reference design uses a five-level U-Net emitting corner heatmaps plus part-affinity fields for corner-to-gate association, with the gate map *estimated online* rather than surveyed. The A2RL-winning system segments with a U-Net, fits line segments, and intersects them for sub-pixel corners — markedly more accurate than heatmap arg max when a gate spans 20–60 px (R3). Zero-shot sim-to-real gate perception is achievable through an edge-like intermediate representation, and label-free adaptation through geometric consistency between successive frames.

On estimation, monocular visual-inertial systems have four unobservable directions (three position, one yaw), so absolute heading must be carried as uncertainty rather than invented. Crucially for a quadrotor, the accelerometer measures specific force dominated at speed by rotor drag, not gravity direction; the naive accelerometer-as-gravity correction is therefore systematically biased, and the standard specific-force-magnitude gate does not remove it. The correct treatment uses the drag model, which additionally makes body velocity observable and is explicitly yaw-independent. The most advanced gate-relative estimator injects gate-corner *pixel reprojection errors* directly as filter innovations, remaining valid with two visible corners where PnP requires four.

3.6 The four in-regime systems

Filtering the whole corpus by “monocular camera + IMU, no odometry, no map” leaves four published systems: pixels-to-CTBR without state estimation; bearing-based traversal of moving gates; INDI driven directly by optical flow; and visual model-predictive localization on a 72 g platform. Every other racing system assumes motion capture or VIO.

Table 2: Method families against the three filters.

Family	F1	F2	F3
Min-snap / MINCO tracked with a state estimate	✗	✗	—
CPC / TOGT as an *online* controller	✗	✗	—
CPC / TOGT as an *offline* teacher	✓	✓	✓
MPCC / MPCC++ online	✗	✗	—
VIO + gate-corner EKF + track map	✗	✓	✓
Gate-corner reprojection into an ESKF, no map	✓	✓	✓
Pixels → CTBR, no state estimation	✓	✓	✓
Bearings-only / proportional-navigation guidance	✓	✓	✓
Optical-τ looming control	✓	✓	✓
Image-moment IBVS	✓	✓	✓
Learned inertial odometry (IMU-only)	✓	✓	✓
INDI inner loop (gyro-only)	✓	✓	✓
Multi-fidelity BO over a low-dim. parameter vector	✓	✓	✓
Model-free RL from scratch on the live simulator	✓	✓	✗

4 Three filters

F1 Legality. Does the deploy-time input set fall within §2? This eliminates every method consuming a surveyed map at run time, including Swift’s resectioning and MonoRace’s PnP against assumed gate locations. It does not eliminate training-time privileged information.

F2 Observability. Is the required quantity recoverable from a monocular camera and IMU with no map? Methods needing absolute heading or absolute position fail. Methods needing only gate-relative bearing, apparent scale, and their rates survive.

F3 Attempt budget. Can it be validated within ∼3 official attempts per hypothesis, each bounded to 90 s? This eliminates tuning loops needing hundreds of real rollouts and favours methods whose expensive search happens offline.

5 An unoccupied junction

Four mature literatures each solve part of the problem, and we find no work joining them to racing.

Optical τ. Time-to-contact τ, the ratio of apparent size to its rate of expansion, is computable from the image alone with no range, no speed and no scale. It grounds a family of control laws — constant-τ, constant-$\dot{\tau}$, τ-coupling — that were formalized for aircraft guidance with stability analysis and flight tests, including for *docking*, i.e. aperture passage. Flight-test evidence shows human helicopter pilots execute decelerations, turns and pull-ups by τ-coupling.

Bearings-only gap traversal. A guidance law exists that traverses a gap using only bearing measurements to the gap edges — no range, no gap pose, no state estimate — with guidance-theoretic guarantees, extended to moving gaps and to windows. It is published in the aerospace guidance community and is essentially invisible to robotics-side searches.

Image-based servoing through apertures. Window traversal under pure image-space feedback has a rigorous nonlinear-control treatment. Image moments of a planar target give decoupled, well-conditioned features: centroid for bearing, area$^{-1/2}$ as a scale-free depth proxy, second-order moments for tilt. An experimental comparison establishes that spherical and moment features give the best-conditioned quadrotor servo dynamics.

Image-space multi-opening flight. High-speed flight through *consecutive* openings has been demonstrated with estimation, planning and control entirely in image space — a race course solved without a map.

The nearest modern robotics instance adapts proportional navigation to vision-based gate traversal with a closed-form optimal solution, without a track map or VIO. It is effectively a rediscovery of the guidance-community result.

Narrowing the claim. An earlier draft of this paper asserted that the junction of optical flow, aperture traversal and aerial autonomy was unoccupied. A dedicated non-English sweep falsified that. Daïni, Raffin, Raharijaona and Ruffier (*Advanced Robotics*, 2026, `10.1080/01691864.2026.2667369`) head toward and traverse apertures using *panoramic optic flow*, with no gate detector and no metric state estimate — a direct occupant of the optic-flow→aperture cell, descending from the Aix-Marseille bio-inspired guidance line rather than from robotics racing. It is written in English but sits entirely outside the drone-racing citation graph, which is why conventional racing-literature search does not surface it.

What remains, to our knowledge, unoccupied is narrower and we state it accordingly: the use of bearings-only and looming guidance for *competitive time minimization* through an *ordered sequence* of gates, under a contract that forbids both a map and a state estimate. Daïni et al. optimize exploration, not lap time; Midhun and Ratnoo treat a single gap; Tang et al. treat a single window. Our design occupies that narrower cell, and Daïni et al. should be read as the closest prior art rather than as unrelated work.

6 Proposed system

6.1 Deterministic expert

Two channels, both closed-form, neither requiring a state estimate:

$$\text{lateral/yaw:} \quad \dot{\lambda} \text{ (bearing rate)} \rightarrow \text{PN law,} \tag{4}$$

$$\text{thrust/pitch:} \quad \tau = 1/\dot{s},\ s = \log w_{\text{px}} \rightarrow \text{looming law.} \tag{5}$$

Commit is triggered on a τ threshold rather than a size threshold, which is the biologically and aeronautically attested form and is invariant to gate distance and approach speed. Aperture fit is evaluated as a body-scaled image ratio, following the human and avian aperture literature, not in metric units.

6.2 Estimation

An error-state Kalman filter on SO(3) with gyro-bias state, using the *drag model* rather than accelerometer-as-gravity; yaw carried as integrated relative heading with growing uncertainty. Gate-relative state via direct corner reprojection innovations, valid at two visible corners. Between detections, propagate with a learned inertial prior trained on the racing motion distribution.

6.3 Perception

Segment with a small U-Net; fit lines; intersect for sub-pixel inner corners. Train on `MonoRaceGate` (inner-corner labels, MIT) and TII-RATM (COCO keypoints, CC BY 4.0), plus semi-synthetic composites of real backgrounds with randomized gate renders. Adapt label-free by geometric consistency.

6.4 Offline reference (training only)

Gate-as-volume time-optimal solve over the 1.22 m crossing squares with single-rotor thrust limits; re-time with a rigid-body time-optimal path parameterization respecting actual angular-rate bounds; clip wherever the next gate would leave the FOV. Use as reward shaping and behaviour-cloning labels, never as an online reference — the state to track it is unobservable, and plan-and-track decomposition is known to underperform at the limit.

6.5 Policy

A single-layer GRU (hidden 128) over the frozen observation, emitting tanh-bounded CTBR, with auxiliary phase, time-to-contact and uncertainty heads. Asymmetric actor-critic: privileged critic, contract-restricted actor. Bootstrap by imitation before reinforcement. Deploy first as a bounded residual around the expert, $u = \text{clamp}(u_{\text{expert}} + g\delta)$ with $g = 0$ on stale or ambiguous perception.

6.6 Safety filter

Contacts are ranked above time, so enforce them outside the learned policy. A control-barrier-function filter keeping gate corners inside the field of view is directly applicable, and CBF filters have been demonstrated on racing airframes at $> 100\,\mathrm{km\,h^{-1}}$ on a $10\,\mathrm{g}$ microcontroller. A filter also survives the parameter freeze that online adaptation violates.

7 Specialization protocol

Contingent on §2.3.

7.1 The premise: champion policies are already track-specialized

The strategy is not a departure from practice; it is a description of it. The Nature-level racing result identified its residual dynamics model from real flight data collected *on the competition track itself* — the authors liken this to pilots training for weeks on the specific course. Two independent papers from the same group state the consequence plainly: champion-level policies "fail in unseen track configurations, always requiring complete retraining," and reinforcement learning yields "specialized policies…even human-champion-level performance in single-task scenarios" that "struggle with novel tasks." Generalization has a measured price: the most careful study of the trade-off reports a 7.4× generalization improvement while explicitly noting that prior generalization methods "impose a substantial cost on flight speed."

The generalist/specialist trade-off in racing was named and measured two decades before any of this, in evolutionary computation, and is essentially uncited by the modern literature.

7.2 Speed conditioning replaces a policy ladder

The champion precedent trains a family of policies at different speed/robustness points and selects at the event. We instead condition one policy on a scalar $\beta \in [0, 1]$ scaling the target closing rate, sampled uniformly during training and used as a dial at deployment. The mechanism is a feature-wise linear modulation layer, demonstrated for quadrotors to span an envelope to $60\,\mathrm{km\,h^{-1}}$ and $4.5\,g$ under a user-set conditioning variable.

The budget argument is decisive: conditioning turns specialization into a **one-dimensional search over β** rather than a search over network weights. It yields the whole ladder from one training run, and a failed fast attempt does not invalidate the slow policy because they are the same weights.

7.3 How much is specialization worth?

The one directly comparable measurement we found: iterative-learning MPC applied to drone racing improves initial trajectories by up to 60.85% in lap time, and still yields 6.05% *on top of a well-tuned* MPCC++ baseline on a real drone. That residual few percent, obtained by specializing an already-good general controller to the specific track, is the entire competitive case for spending official attempts.

7.4 Randomization substitutes for calibration; it does not add to it

Two controlled ablations appear to contradict each other, and reconciling them gives the design rule.

One study randomizes motor limits, thrust coefficient, roll/pitch effectiveness and actuator time constant at 0/10/20/30% and reports a *complete rank inversion* between simulation and reality: simulated reward orders 0% > 10% > 20% > 30%, while real reward orders 10% > 20% > 30% ≫ 0%. The 0% policy had the *lowest* simulated crash rate (0.6%) yet passed only 5 gates in reality. Mean speed falls monotonically with randomization ($7.13 \to 6.61 \to 6.40\,\mathrm{m\,s^{-1}}$).

A second study reports the opposite: with careful static identification of the motor time constant and thrust coefficient, adding randomization monotonically *hurts* — real tracking error $0.028\,\mathrm{m}$ with no randomization, 0.036 at ±10%, 0.053 at ±30% — yet with a deliberate 30% parameter offset, the un-randomized policy fails outright (∞) while ±30% randomization recovers it to $0.066\,\mathrm{m}$.

> *The two agree once conditioned on calibration quality: domain randomization is a substitute for calibration, not a complement. It buys robustness exactly in proportion to residual mis-calibration, and charges lap time regardless.*

The operational rule is therefore: **identify hard, randomize only what cannot be pinned down**, and treat every randomized parameter as a confession of ignorance with a measurable price in speed. Concretely: calibrate the motor time constant, thrust coefficient and rate-loop response against the target; randomize the perception and timing channels (motion blur, exposure, JPEG artifacts, dropout, corner noise, latency jitter) that the surrogate cannot faithfully reproduce.

A caution against relying on randomization alone. The Swift authors' own ablation deploys policies into shifted *simulated* settings — idealized versus blade-element aerodynamics, crossed with ground-truth versus identified-noise observations. Zero-shot reinforcement learning, domain randomization, *and* time-optimal trajectory tracking with MPC all "cannot complete even a single lap" under that shift; only the approach using empirically identified residual models succeeds across all settings. Blanket randomization did not survive a dynamics-and-observation domain shift; identification did. A plan resting primarily on randomization should treat this as its strongest counterexample.

7.5 Our transfer is sim-to-sim, and that has a methodology

The transfer here is not sim-to-real: it runs from a calibrated surrogate to a *second simulator* (the official scoring environment). This is a named and studied setting. Ligot and Birattari's *pseudo-reality* argues formally that evaluating a design in a second, different simulator reproduces reality-gap effects and is a legitimate proxy for reality; the terminology precedent is James et al.'s "Sim-to-Real via Sim-to-Sim." Both sit outside the aerial-robotics literature and are, as far as this survey can establish, never cited by it.

This changes one expectation materially. The documented failure of simulated crash rate to predict real crash rate was attributed by its authors to *unmodelled vision-based state-estimation error*. In a sim-to-sim transfer that channel is not absent — both ends render images — but it is *modellable and under our control*, because we can inject measured detector noise into the surrogate rather than hoping it emerges. We therefore predict better crash-rate predictiveness than any sim-to-real study reports, and we treat that as a falsifiable claim rather than an assumption: if surrogate contact rate still fails to rank candidates after detector-noise injection, the surrogate is not fit for contact screening and Algorithm 1's rejection step must be removed.

Note also that no published work reports a correlation statistic between simulated and real lap times across policies; the evidence is qualitative or must be recomputed from paired tables. Reporting one is cheap here and would be a contribution.

7.6 Multi-fidelity Bayesian optimization

With $\sim$3 official attempts per hypothesis, use multi-fidelity black-box optimization with a feasibility classifier, fusing analytic, surrogate and rare real evaluations. Optimize $\dim\theta \leq 8$: β, aperture margin, commit τ threshold, terminal centring tolerance, clear duration, and the vertical and lateral gains. Do not include per-gate parameters — besides legality, they multiply search dimension by gate count and guarantee budget exhaustion. Crash-constrained and early-stopping BO variants model failed evaluations explicitly; the existing eight-second no-progress watchdog already emits the right signal for the latter.

7.7 Lap-over-lap refinement

Where repeated runs on the identical course are available, iterative learning control gives provably non-increasing lap time across laps, and a spatial variant within a virtual tube keeps the refined line inside a safety corridor. An independent iterative-learning MPC result adds a shifted local safe set that prevents the optimizer from shortcutting the course. ILC refines the geometric line; Bayesian optimization refines policy parameters.

Algorithm 1 Course specialization under an attempt budget

Fit surrogate dynamics from all existing live traces
Train $\pi_\theta(\cdot;\beta)$ with the asymmetric randomization schedule
Initialize GP over θ from surrogate rollouts only
while official attempts remain **do**
  $\theta^\star \leftarrow$ acquisition maximizer over the surrogate-feasible set
  Evaluate $\theta^\star$ in surrogate at high fidelity
  if predicted contact probability $> \epsilon$ **then**
    reject; **continue**
  end if
  Spend *one* official attempt on $\theta^\star$
  Update GP with lap time **and** contact outcome
  if contact occurred **then**
    shrink the feasible β ceiling
  end if
end while

8 Experimental protocol

Metric ladder. Strictly ordered: interface audit clean (0 blocked types configured, 0 observed) $\rightarrow$ actuation calibrated (sign, scale, saturation, safe stop measured) $\rightarrow$ detector meets held-out recall/precision/latency $\rightarrow$ attitude RMS and tracker false-handoff rate met $\rightarrow$ one official completion $\rightarrow$ two in three attempts $\rightarrow$ zero contacts $\rightarrow$ elapsed time.

Falsifiers. Declared before each attempt. Any new contact class refutes a speed increase; visual passage phase leading official progress refutes the passage detector; a detector-confidence distribution shifted from held-out refutes perception transfer; persistent residual saturation refutes the residual bound.

First actions, by expected value per unit effort.

1. Instrument the control loop; confirm or refute 9.2 Hz (§2.2). If confirmed, raise it before anything else.
2. Settle the specialization-legality question (§2.3).
3. Implement the drag-aware attitude ESKF. It does not exist yet and everything downstream depends on it.
4. Build the perception corpus from `MonoRaceGate`, TII-RATM and semi-synthetic renders — not from live attempts.

5. Replace the hand-tuned mode machine with the closed-form PN + τ expert of §6.

9 What does not exist

Negative results, stated so they can be checked.

- **No pre-2013 vision-only aperture traversal that we could substantiate.** Targeted Crossref, DBLP, OpenAlex and Semantic Scholar sweeps over window/gap/aperture/opening/doorway × UAV/MAV/quadrotor/helicopter, date-restricted and not, found none; the earliest substantiable instance is 2018, and earlier work is mocap-based, corridor-centring rather than aperture passage, or landing rather than traversal. *This negative must be stated with a specific caveat.* Crossref is structurally blind to Chinese aerospace literature: the major venues (航空学报, 宇航学报, 控制理论与应用, 北航学报) register through CNKI rather than Crossref and return zero DOIs, so a Crossref negative is not evidence of absence. A dedicated sweep reaching CQVIP and the open-access layer confirmed the pre-2013 negative for Chinese, but **CNKI and Wanfang were never successfully queried**. Aggregator queries in CJK are also actively misleading — Semantic Scholar returns 104 hits for 无人机穿越 with none on-topic, dominated by the film *Interstellar* and time-travel fiction. We therefore state this as "we found none" and identify the exact index that would settle it.

- **A Chinese-language task literature exists that English racing work does not cover.** It is uniformly downstream of the Western methodological line — there is no Chinese precursor to minimum-snap or polynomial trajectory generation — but it addresses aperture variants the English literature has not: irregular non-parametric openings in rubble, apertures containing a rotating obstacle requiring speed-timed passage, formation squeeze-through with inter-vehicle downwash, and a fixed-wing gate-racing competition ecosystem (无人争锋) absent from every English survey, which are uniformly multirotor.

- **No work joining τ/optic-flow guidance to *time-optimal* racing.** The τ →aircraft link exists in guidance; the τ →aperture link exists in biology, in human locomotion, and — as a dedicated non-English sweep established, correcting our own earlier claim — in robotics, via panoramic-optic-flow aperture traversal for aerial exploration. What we could not find is any work using these image-only quantities to minimize time through an ordered gate sequence.

- **A confirmed 2005 Japanese-language result that is nonetheless not hidden prior art.** A Chiba University group demonstrated autonomous aggressive helicopter flight in 2005 — including a complete 360° roll executed by feedforward sequence control *learned from expert pilot flight data*, six years before the canonical minimum-snap work (`10.1299/jsmemecjo.2005.7.0_271`, in Japanese). It was subsequently republished in English as a Springer chapter, so it is discoverable; we record it because a Japanese-only 2005 aggressive-flight paper is exactly the artifact a literature search of this kind is expected to miss, and it does exist.

- **A genuinely English-invisible Japanese line does exist, but it is not racing.** The Chiba aerobatics programme (flip, hammerhead turn, Immelmann turn; 2016–2023) is published only in *Transactions of the JSME (in Japanese)* with no English counterpart, verified by exhaustive author-level search. It concerns single-rotor helicopters and maneuver-library reference-trajectory design, not minimum-time quadrotor racing. Cite as contrast, not precursor.

- **No Korean-language autonomous drone racing research exists.** The major Korean index returns literally zero for the query; all Korean drone-racing publications are virtual-reality and game content. In Russian, all 320 hits for "drone racing" concern sport pedagogy and human factors, not autonomy. The one German-language hit on time-optimal racing lines concerns cars.

- **No A2RL drone-race overview or organizer paper.** The complete set of A2RL drone team papers is five.

- **No papers on this competition.** Verified across arXiv, Crossref and GitHub. The only public artifact is an open practice simulator.

- **No minimum-lap-time theory for quadrotor racing.** Double-confirmed by two independent sweeps across arXiv full text and DBLP on nine bound-related phrasings. No drone analogue of motorsport's minimum-lap-time simulation exists. Relatedly: no paper is dedicated to quantifying the optimality gap between learned policies and time-optimal references (it is reported only incidentally inside method papers), and no survey or benchmark exists specifically on time-optimal quadrotor planning.

- **No diffusion or flow-matching planner for drone racing**, and no state-space-model (Mamba-class) quadrotor policies. The nearest neighbours are in aerobatics and CBF-guided diffusion.

- **No work studies a hard budget of ∼3 real attempts per hypothesis.** The nearest formal treatments are multi-fidelity Bayesian optimization, which *minimizes* real evaluations without fixing a count, and safe Bayesian optimization, which constrains failures rather than trial count. The closest methodological ancestor is a 2010 quadrotor multi-flip result that parameterizes the maneuver by five numbers and obtains both the initial values *and the corrective gradient* from a coarse model, so only a handful of real trials are needed — precisely the structure a three-attempt budget requires.

- **No study measures the lap-time delta between a track-generalist and a track-specialist policy on the same hardware.** One side reports the generalization gain and notes it costs speed; the other reports the specialization gain over a fixed controller. The clean two-sided comparison does not exist.
- **The ETH iterative-learning-control lineage released no code at all**, across roughly a dozen papers spanning 2009–2017. Neither did MPCC++, weighted-maximum-likelihood controller tuning, or the spatial iterative-learning racing work. This is a reproducibility problem for anyone trying to adopt these methods under time pressure.
- **No FPGA gate detector** and no published Safe Robot Learning Competition report.
- **Phantom citations** that recur and should not be cited: "Hierarchical Reinforcement Learning for Drone Racing"; "Boosting RL with HJB" as a distinct publication; "Learning to Race in Head-to-Head" in the drone domain.

10 Conclusion

The blocked-telemetry contract makes most of the drone-racing literature inadmissible, but it makes an older and largely forgotten literature exactly right. Bearing, bearing rate, and optical looming are the quantities the contract leaves observable, and closed-form guidance laws over precisely those quantities have existed — with stability analysis and flight tests — for decades, in communities that do not read robotics venues. Combined with a speed-conditioned policy and multi-fidelity specialization, they form a design that fits both the contract and the attempt budget. The binding constraint, however, is neither perception nor policy: it is a control loop running an order of magnitude below the permitted rate. Measure it first.

A Classified bibliography

Identifiers were verified against Crossref, arXiv, DBLP or the source PDF. Items we could not independently resolve are marked UNVERIFIED and should be checked before publication.

A. Optical τ, looming, and time-to-contact (the ancestral line)

Authors	Title / venue	Year	DOI
Lee	A theory of visual control of braking based on information about time-to-collision. *Perception* 5(4)	1976	10.1068/p050437
Lee & Reddish	Plummeting gannets: a paradigm of ecological optics. *Nature* 293	1981	10.1038/293293a0
Warren & Whang	Visual guidance of walking through apertures. *JEP:HPP* 13(3)	1987	10.1037/0096-1523.13.3.371
Nelson	Using flow field divergence for obstacle avoidance. *ICCV*	1988	10.1109/CCV.1988.589990
Srinivasan et al.	Range perception through apparent image speed in honeybees. *Vis. Neurosci.* 6(5)	1991	10.1017/S095252380000136X
Gabbiani et al.	Computation of object approach by a wide-field neuron. *J. Neurosci.* 19(3)	1999	10.1523/JNEUROSCI.19-03-01122.1999
Tresilian	Visually timed action: time-out for 'tau'? *TiCS* 3(8)	1999	10.1016/S1364-6613(99)01352-2
Coombs et al.	Real-time obstacle avoidance using central flow divergence. *T-RA* 14(1)	1998	10.1109/70.660840
Srinivasan et al.	How honeybees make grazing landings. *Biol. Cybern.* 83(3)	2000	10.1007/s004220000162
Padfield, Lee & Bradley	How do helicopter pilots know when to stop, turn or pull up? *JAHS* 48(2)	2003	10.4050/JAHS.48.108
Ruffier & Franceschini	Optic flow regulation: the key to aircraft automatic guidance. *RAS* 50(4)	2005	10.1016/j.robot.2004.09.016
Zufferey & Floreano	Fly-inspired visual steering of an ultralight indoor aircraft. *T-RO* 22(1)	2006	10.1109/TRO.2005.858857
Horn, Fang & Masaki	Time to contact relative to a planar surface. *IEEE IV*	2007	10.1109/IVS.2007.4290093
Beyeler et al.	Vision-based control of near-obstacle flight. *Auton. Robots* 27(3)	2009	10.1007/s10514-009-9139-6
Hérisse et al.	Landing a VTOL UAV on a moving platform using optical flow. *T-RO* 28(1)	2012	10.1109/TRO.2011.2163435
Izzo & de Croon	Landing with time-to-contact and ventral optic flow. *JGCD* 35(4)	2012	10.2514/1.56598
Kendoul & Ahmed	Bio-inspired TauPilot for 4D docking and landing. *IROS*	2012	10.1109/IROS.2012.6385586
Kendoul	**Four-dimensional guidance and control using time-to-contact. *IJRR* 33(2)**	**2013**	**10.1177/0278364913509496**
Schiffner et al.	Minding the gap: in-flight body awareness in birds. *Front. Zool.* 11	2014	10.1186/s12983-014-0064-y
Ravi et al.	Gap perception in bumblebees. *JEB* 222(2)	2019	10.1242/jeb.184135
Ravi et al.	Bumblebees perceive spatial layout relative to body size. *PNAS* 117(49)	2020	10.1073/pnas.2016872117
de Croon et al.	Accommodating unobservability to control attitude with optic flow. *Nature* 610	2022	10.1038/s41586-022-05182-2

B. Image-based visual servoing and aperture traversal

Authors	Title / venue	Year	DOI
Hamel & Mahony	Visual servoing of an under-actuated dynamic rigid-body system. *T-RA* 18(2)	2002	10.1109/TRA.2002.999647
Chaumette	Image moments: a general and useful set of features. *T-RO* 20(4)	2004	10.1109/TRO.2004.829463
Tahri & Chaumette	**Point- and region-based image moments for servoing of planar objects. *T-RO* 21(6)**	**2005**	**10.1109/TRO.2005.853500**
Chaumette & Hutchinson	Visual servo control I / II. *RAM* 13(4), 14(1)	2006/07	10.1109/MRA.2006.250573
Bourquardez et al.	IBVS of the translation kinematics of a quadrotor. *T-RO* 25(3)	2009	10.1109/TRO.2008.2011419
Mellinger, Michael & Kumar	Trajectory generation and control for precise aggressive maneuvers. *IJRR* 31(5)	2012	10.1177/0278364911434236
Serra et al.	Landing on a moving target using dynamic IBVS. *T-RO* 32(6)	2016	10.1109/TRO.2016.2604495
Falanga et al.	Aggressive flight through narrow gaps using active vision. *ICRA*	2017	10.1109/ICRA.2017.7989679
Loianno et al.	Aggressive flight with a single camera and IMU. *RA-L* 2(2)	2017	10.1109/LRA.2016.2633290
Sanket et al.	GapFlyt: structure-less gap detection. *RA-L* 3(4)	2018	10.1109/LRA.2018.2843445
Tang et al.	Going through a window and landing using optical flow. *ECC*	2018	10.23919/ECC.2018.8550193
Tang et al.	**Quadrotor going through a window: an IBVS approach. *Control Eng. Practice* 112**	**2021**	**10.1016/j.conengprac.2021.104827**
Guo & Leang	**Image-based estimation, planning and control through multiple openings. *IJRR* 39(9)**	**2020**	**10.1177/0278364920921943**
Midhun & Ratnoo	**Gap traversal guidance using bearing information. *JGCD* 45(12)**	**2022**	**10.2514/1.G006898**
Midhun & Ratnoo	Quadrotor guidance for window traversal: a bearings-only approach. *JGCD* 48(10)	2025	10.2514/1.G009084
Qin et al.	Perception-aware IBVS of aggressive quadrotor UAVs. *T-Mech* 28(4)	2023	10.1109/TMECH.2023.3251954 (UNVER.)
Römer, Emmert & Schoellig	**Flying through moving gates without full state estimation. *ICRA***	**2025**	**arXiv:2410.15799**
Wu et al.	Precise aggressive aerial maneuvers with sensorimotor policies. *Sci. Robotics* 11(115)	2026	10.1126/scirobotics.aeb0180

C. Polynomial and time-optimal trajectory generation

Authors	Title / venue	Year	DOI / arXiv
Fliess et al.	Flatness and defect of non-linear systems. *Int. J. Control* 61(6)	1995	10.1080/00207179508921959
Martin, Devasia & Paden	A different look at output tracking: control of a VTOL aircraft. *Automatica* 32(1)	1996	10.1016/0005-1098(95)00099-2
van Nieuwstadt & Murray	Real-time trajectory generation for differentially flat systems. *IJRNC* 8(11)	1998	10.1002/(SICI)1099-1239(199809)8:11
Mellinger & Kumar	Minimum snap trajectory generation and control. *ICRA*	2011	10.1109/ICRA.2011.5980409
Hehn, Ritz & D'Andrea	Performance benchmarking using time-optimal control. *Auton. Robots* 33(1–2)	2012	10.1007/s10514-012-9282-3

Authors	Title / venue	Year	DOI / arXiv
Richter, Bry & Roy	**Polynomial trajectory planning for aggressive quadrotor flight. *ISRR* → STAR 114**	**2013/16**	**10.1007/978-3-319-28872-7_37**
Bry et al.	Aggressive flight of fixed-wing and quadrotor aircraft. *IJRR* 34(7)	2015	10.1177/0278364914558129
Mueller, Hehn & D'Andrea	A computationally efficient motion primitive. *T-RO* 31(6)	2015	10.1109/TRO.2015.2479878
Hehn & D'Andrea	Real-time trajectory generation for quadrocopters. *T-RO* 31(4)	2015	10.1109/TRO.2015.2432611
Liu et al.	Search-based motion planning for aggressive flight in SE(3). *RA-L* 3(3)	2018	10.1109/LRA.2018.2795654
Foehn, Romero & Scaramuzza	Time-optimal planning for quadrotor waypoint flight. *Sci. Robotics* 6(56)	2021	10.1126/scirobotics.abh1221
Wang et al.	Geometrically constrained trajectory optimization (MINCO/GCOPTER). *T-RO* 38(5)	2022	10.1109/TRO.2022.3160022
Han et al.	Fast-Racing: SE(3) planning baseline. *RA-L* 6(4)	2021	10.1109/LRA.2021.3113976
Qin et al.	**Time-optimal gate-traversing planner. *ICRA***	**2024**	**10.1109/ICRA57147.2024.10610148**
Mao et al.	TOPPQuad: dynamically-feasible time-optimal path parametrization. *IROS*	2024	arXiv:2309.11637
Spasojevic, Murali & Karaman	Perception-aware time-optimal path parameterization. *ICRA*	2020	10.1109/ICRA40945.2020.9197157
Qin et al.	**Perception-aware time-optimal planning for waypoint flight**	**2026**	**arXiv:2603.04305**
Ryou, Tal & Karaman	Multi-fidelity black-box optimization for time-optimal maneuvers. *RSS/IJRR*	2020/21	10.1177/02783649211033317

D. Control: MPC, flatness, INDI, geometric

Authors	Title / venue	Year	DOI / arXiv
Lee, Leok & McClamroch	Geometric tracking control of a quadrotor on SE(3). *CDC*	2010	10.1109/CDC.2010.5717652
Smeur, Chu & de Croon	Adaptive INDI for attitude control of micro air vehicles. *JGCD* 39(3)	2016	10.2514/1.G001490
Faessler, Falanga & Scaramuzza	Thrust mixing, saturation, and body-rate control. *RA-L* 2(2)	2017	10.1109/LRA.2016.2640362
Faessler, Franchi & Scaramuzza	Differential flatness subject to rotor drag. *RA-L* 3(2)	2018	10.1109/LRA.2017.2776353
Falanga et al.	PAMPC: perception-aware MPC for quadrotors. *IROS*	2018	arXiv:1804.04811
Brescianini & D'Andrea	Tilt-prioritized quadrocopter attitude control. *T-CST* 28(2)	2020	10.1109/TCST.2018.2873224
Tal & Karaman	Accurate tracking using INDI and differential flatness. *T-CST* 29(3)	2021	10.1109/TCST.2020.3001117
Torrente et al.	Data-driven MPC for quadrotors. *RA-L* 6(2)	2021	10.1109/LRA.2021.3061307
Romero et al.	Model predictive contouring control for time-optimal flight. *T-RO* 38(6)	2022	10.1109/TRO.2022.3173711
Sun et al.	Comparative study of NMPC and differential-flatness-based control. *T-RO* 38(6)	2022	10.1109/TRO.2022.3177279
Ho & Zhou	INDI-based optical flow control for flying robots. *RSS*	2023	arXiv:2307.02702

Authors	Title / venue	Year	DOI / arXiv
Krinner et al.	MPCC++: time-optimal flight with safety constraints. *RSS*	2024	10.15607/RSS.2024.XX.109
Kunapuli et al.	Leveling the playing field: classical vs learned controllers. *RSS*	2025	arXiv:2506.17832
Singletary et al.	Onboard safety guarantees for racing drones (CBF geofencing). *RA-L* 7(2)	2022	10.1109/LRA.2022.3144777
Trimarchi, Schiano & Tron	A CBF candidate for quadrotors with limited field of view. *CDC*	2024	arXiv:2410.01277

E. Learning-based racing, perception, and estimation

Authors	Title / venue	Year	DOI / arXiv
Kaufmann et al.	Deep drone racing. *CoRL*	2018	arXiv:1806.08548
Kaufmann et al.	Beauty and the beast: optimal methods meet learning. *ICRA*	2019	arXiv:1810.06224
Loquercio et al.	Deep drone racing: sim to reality with domain randomization. *T-RO* 36(1)	2020	10.1109/TRO.2019.2942989
Foehn et al.	AlphaPilot: autonomous drone racing. *RSS* / *Auton. Robots* 46	2020/22	10.1007/s10514-021-10011-y
De Wagter et al.	Sensing, state-estimation and control behind the AIRR winner. *Field Robotics* 2	2022	10.55417/fr.2022042
Kaufmann, Bauersfeld & Scaramuzza	Benchmark comparison of learned control policies. *ICRA*	2022	arXiv:2202.10796
Kaufmann et al.	Champion-level drone racing using deep RL (Swift). *Nature* 620	2023	10.1038/s41586-023-06419-4
Song et al.	Reaching the limit: optimal control vs RL. *Sci. Robotics* 8(82)	2023	10.1126/scirobotics.adg1462
Cioffi et al.	Learned inertial odometry for autonomous drone racing. *RA-L* 8(5)	2023	10.1109/LRA.2023.3252342
Geles et al.	**Demonstrating agile flight from pixels without state estimation. *RSS***	**2024**	**arXiv:2406.12505**
Xing et al.	Bootstrapping RL with imitation for vision-based agile flight. *CoRL*	2024	arXiv:2403.12203
Heeg, Song & Scaramuzza	Learning quadrotor control from visual features. *ICRA*	2025	arXiv:2410.15979
Bahnam et al.	MonoRace: winning champion-level drone racing with monocular AI	2026	arXiv:2601.15222
Novák et al.	Vision-only UAV state estimation (A2RL finalist)	2026	arXiv:2602.01860
Azhari et al.	Robust tightly-coupled filter-based monocular VI estimation (ADR-VINS). *IROS*	2026	arXiv:2603.02742
Pham et al.	PencilNet: zero-shot sim-to-real gate perception. *RA-L* 7(4)	2022	10.1109/LRA.2022.3207545
Bosello et al.	Race against the machine (TII-RATM dataset). *RA-L* 9(4)	2024	10.1109/LRA.2024.3371288
Hanover et al.	Autonomous drone racing: a survey. *T-RO* 40	2024	10.1109/TRO.2024.3400838

F. Estimation foundations and the drag/gravity problem

Authors	Title / venue	Year	DOI / arXiv
Mahony, Hamel & Pflimlin	Nonlinear complementary filters on SO(3). *TAC*	2008	10.1109/tac.2008.923738
Martin & Salaün	**The true role of accelerometer feedback in quadrotor control. *ICRA***	**2010**	**10.1109/robot.2010.5509980**
Martinelli	Vision and IMU data fusion: closed-form solutions. *T-RO*	2012	10.1109/tro.2011.2160468
Leishman et al.	Quadrotors and accelerometers: state estimation with an improved model. *CSM*	2014	10.1109/mcs.2013.2287362
Hesch et al.	Consistency analysis and improvement of vision-aided inertial navigation. *T-RO*	2014	10.1109/tro.2013.2277549
Solà	Quaternion kinematics for the error-state Kalman filter	2017	arXiv:1711.02508
Forster et al.	On-manifold preintegration for real-time VIO. *T-RO*	2017	arXiv:1512.02363
Svacha, Loianno & Kumar	**Inertial yaw-independent velocity and attitude estimation. *RA-L***	**2019**	**10.1109/lra.2019.2894220**
Delmerico et al.	Are we ready for autonomous drone racing? (UZH-FPV). *ICRA*	2019	10.1109/ICRA.2019.8793887
Collins & Bartoli	Infinitesimal plane-based pose estimation (IPPE). *IJCV*	2014	10.1007/s11263-014-0725-5

G. Sim-to-sim methodology, randomization, and identification

Authors	Title / venue	Year	DOI / arXiv
Ligot & Birattari	**Simulation-only experiments to mimic the reality gap (*pseudo-reality*). *Swarm Intell.***	**2019**	**10.1007/s11721-019-00175-w**
Ligot & Birattari	On using simulation to predict the performance of robot swarms. *Sci. Data*	2022	10.1038/s41597-022-01895-1
James et al.	Sim-to-real via sim-to-sim: randomized-to-canonical adaptation. *CVPR*	2019	arXiv:1812.07252
Bauersfeld et al.	NeuroBEM: hybrid aerodynamic quadrotor model. *RSS*	2021	10.15607/RSS.2021.XVII.042
Sun, de Visser & Chu	Quadrotor gray-box model identification from high-speed flight. *J. Aircraft*	2019	10.2514/1.C035135
Eschmann, Albani & Loianno	Data-driven system identification subject to motor delays. *IROS*	2024	arXiv:2404.07837
Ferede et al.	One net to rule them all: domain randomization across platforms. *ICRA*	2025	arXiv:2504.21586
Chen et al.	What matters in learning a zero-shot sim-to-real RL policy (SimpleFlight). *RA-L*	2025	10.1109/LRA.2025.3575011
Aljalbout et al.	The reality gap in robotics: challenges, solutions, best practices	2025	arXiv:2510.20808
Huang et al.	DATT: deep adaptive trajectory tracking. *CoRL*	2023	arXiv:2310.09053
O'Connell et al.	Neural-Fly enables rapid learning for agile flight in strong winds. *Sci. Robotics* 7	2022	10.1126/scirobotics.abm6597
Song et al.	Flightmare: a flexible quadrotor simulator. *CoRL*	2020	arXiv:2009.00563
Miao et al.	FalconGym: photorealistic simulation for zero-shot vision-based navigation. *IROS*	2025	arXiv:2503.02198
Low et al.	SOUS VIDE: cooking visual drone policies in a Gaussian splatting vacuum. *RA-L*	2025	10.1109/LRA.2025.3553785

H. Specialization to a fixed track

Authors	Title / venue	Year	DOI / arXiv
Togelius & Lucas	**Evolving robust and *specialized* car racing skills. *IEEE CEC***	**2006**	**10.1109/cec.2006.1688444**
Cardamone, Loiacono & Lanzi	Learning to drive in TORCS using online neuroevolution. *T-CIAIG* 2(3)	2010	10.1109/tciaig.2010.2052102
Lupashin et al.	A simple learning strategy for high-speed multi-flips. *ICRA*	2010	10.1109/ROBOT.2010.5509452
Schoellig, Mueller & D'Andrea	Optimization-based iterative learning for precise quadrocopter tracking. *Auton. Robots* 32(4)	2012	10.1007/s10514-012-9283-2
Hehn & D'Andrea	Frequency-domain iterative learning for periodic quadrocopter maneuvers. *Mechatronics* 24(8)	2014	10.1016/j.mechatronics.2014.09.013
Rosolia & Borrelli	Learning how to autonomously race a car (LMPC). *T-CST* 28(6)	2020	arXiv:1901.08184
Vallon & Borrelli	Task decomposition for iterative learning MPC. *ACC*	2020	arXiv:1903.07003
Wurman et al.	Outracing champion Gran Turismo drivers with deep RL. *Nature* 602	2022	10.1038/s41586-021-04357-7
Bauersfeld, Kaufmann & Scaramuzza	User-conditioned neural control policies (FiLM). *ICRA*	2023	arXiv:2211.12181
Wang et al.	Environment as policy: learning to race in unseen tracks. *ICRA*	2025	arXiv:2410.22308
Trumpp et al.	Attenuated residual policy optimization for real-world racing	2026	arXiv:2603.12960
Jain & Morari	Computing the racing line using Bayesian optimization. *CDC*	2020	10.1109/CDC42340.2020.9304147
Loquercio, Saviolo & Scaramuzza	AutoTune: controller tuning for high-speed flight. *RA-L* 7(2)	2022	10.1109/LRA.2022.3146897
Fröhlich et al.	Contextual tuning of MPC for autonomous racing. *IROS*	2022	arXiv:2110.02710
Romero et al.	Weighted maximum likelihood for controller tuning. *ICRA*	2023	arXiv:2210.11087
Bauersfeld, Kaufmann & Scaramuzza	User-conditioned neural control policies. *ICRA*	2023	—
Lv et al.	**Time-optimal spatial iterative learning control within a virtual tube. *ICRA***	**2023**	**10.1109/icra48891.2023.10161383**
Cheng et al.	DiffTune: auto-tuning through auto-differentiation. *T-RO* 40	2024	arXiv:2209.10021
Zhao et al.	Improving drone racing performance through iterative learning MPC. *IROS*	2025	arXiv:2508.01103
Nam et al.	Track-centric iterative learning for global trajectory optimization	2026	arXiv:2601.21027
von Rohr et al.	Local Bayesian optimization with crash constraints. *at-Automatisierungstechnik*	2024	arXiv:2411.16267

I. Performance bounds (the open problem)

Authors	Title / venue	Year	DOI / arXiv
Falanga, Kim & Scaramuzza	**How fast is too fast? The role of perception latency. *RA-L* 4(2)**	**2019**	**10.1109/LRA.2019.2898117**

Authors	Title / venue	Year	DOI / arXiv
Abdulrahim et al.	Flight envelope requirements for FPV quadrotor racing. *AIAA SciTech*	2019	10.2514/6.2019-0825
Bauersfeld & Scaramuzza	Range, endurance and optimal speed estimates for multicopters. *RA-L* 7(2)	2022	10.1109/LRA.2022.3145063
Majumdar & Pacelli	Fundamental performance limits for sensor-based robot control. *RSS*	2022	10.15607/RSS.2022.XVIII.036

J. Non-English literature (selected; see §9)

Authors	Title / venue	Year	DOI
许勇 et al.	Fixed-wing UAV high-speed gate-traversal visual navigation. 飞行力学	2025	10.13645/j.cnki.f.d.20250619.003
刘坤达 et al.	Safe robust control for multi-UAV formation narrow-passage traversal. 航空学报	2023	10.7527/S1000-6893.2023.29768
徐力昊 et al.	Agile flight control via bidirectional-thrust quadrotor. 北航学报	2021	10.13700/j.bh.1001-5965.2020.0221
段三星	Learning helicopter aerobatic trajectories from demonstrations. MSc, HIT	2011	10.7666/d.D243238
杨佳辉	Rotor UAV narrow-passage traversal (irregular apertures). MSc, DLUT	2021	10.26991/d.cnki.gdllu.2021.000364
陈文杰	Indoor localization and traversal of multiple gates. MSc, UESTC	2021	10.27005/d.cnki.gdzku.2021.004738
邸健	Autonomous flight control for racing UAVs. PhD, USTC	2023	10.27517/d.cnki.gzkju.2023.000001

B Verified resources

- `tudelft/MonoRaceGate` — inner-corner-labelled gate dataset, MIT.
- `tii-racing/drone-racing-dataset` — TII-RATM, COCO keypoints, CC BY 4.0.
- `elodin-sys/ai-grand-prix` — practice simulator, matching intrinsics, Betaflight SITL.
- `learnsyslab/lsy_drone_racing` — sim-to-real racing benchmark, MIT.
- `FSC-Lab/TOGT-Planner` — gate-as-volume time-optimal planner.
- `ZJU-FAST-Lab/GCOPTER` — MINCO spatial-temporal optimizer.
- `uzh-rpg/rpg_time_optimal` — CPC reference.
- `uzh-rpg/minimum_jerk_trajectories` — closed-form primitives, no solver.
- `uzh-rpg/rpg_quadrotor_control` — drag-aware flatness feedforward.
- `uzh-rpg/LAFR` — real-world agile flight adaptation, GPLv3.
- `maokat12/lbTOPPQuad` — rigid-body time-optimal path parameterization.
- `ntnu-arl/PX4-CBF`, `dev10110/gatekeeper` — deployable safety filters.