

The kite turbine: a search over its design space (draft build)

The kite turbine: a search over its design space

Internal engineering report. Draft v0.87. Date 2026-10-21.

Status DRAFT. Data commit dd3cc6a. Launch commit a76c5e9. The physics era is post-95c385d_wind-authority, with the bank derate applied throughout. Numbers come only from signed register rows, cited by row ID.

Draft status. This is an internal draft. Sections 1 to 6 carry their text. Nomenclature stays incomplete. A section marked [PENDING] waits on a gate.

Abstract. This report presents a search for a 5 kW-class ground-generation rotary kite turbine with a tensile rotary power transmission (TRPT). The search varies ring geometry, polygon count, rotor count, blade bank angles and blade scale. A differential-evolution campaign evaluated 930 candidate designs under one fixed power-law generator load, each scored on the measurement window (NR-001).

The campaign winner is the best feasible candidate under the modeled objective and gates. Its mean power is 5.309 kW at the site wind, 11.01 m/s at 9.41 m (NR-008, NR-032). All checked TRPT rings and twist segments operate below their failure boundaries at the final operating point (NR-021, NR-022). The minimum factor of safety is 14.972 (NR-009).

At the winner operating point, the least segment torque margin is 2.03, at S1, the first of the five segments between the rings (NR-040). The margin is torque capacity over the segment torque it carries, and the least value governs.

All numbers come from one version of the model, the physics era post-95c385d_wind-authority. That version applies the bank derate, which cuts rotor power as the blades bank. The study is simulation-only. Blade integrity is not yet gated, and the outer stiffening bridles of real field turbines sit outside the model. A lifter kite can modulate its tension across wind speeds, and the report treats that ability as an assumption.

The search is a screen, not a survey. The report does not claim a global optimum.

Keywords. airborne wind energy, kite turbine, ground-generation, tensile rotary power transmission, differential evolution, multibody simulation, design search, factor of safety.

Summary for decision. The project completed a 930-design simulation screen of a 5 kW-class ground-generation rotary kite turbine with a tensile rotary power transmission (NR-001). The screen found a feasible candidate inside the gated performance bounds. The decision metrics read as follows.

Decision metric	Value	Rows
Designs evaluated	930	NR-001
Feasible candidates	508 of 930	NR-002
Campaign winner mean power	5.309 kW	NR-008
Winner airborne mass	25.698 kg	NR-018
Minimum factor of safety	14.972	NR-009
Least segment torque margin	2.03, at S1	NR-040
Objective improvement, seed to winner	36.109 to 27.364 kg-equivalent	NR-006, NR-005

Of the 930, 508 designs pass every gate. The other 422 read 301 reject, 110 clearance_reject and 11 reject_twist (NR-002).

All checked TRPT rings and twist segments operate below their failure boundaries at the final operating point (NR-021, NR-022). The ODE path reproduces every recorded winner field (NR-012).

The report shows that a feasible 5 kW-class candidate exists in the searched space under the model and gates. It does not show flight readiness, a global optimum, or full structural safety. Blade integrity, lifter modulation, launch and recovery, and atmospheric realism remain open work items.

1 Introduction

The wind above tower height is stronger and steadier than the wind exploited by traditional wind turbines. Airborne wind energy systems (AWES) are kite-based machines that chase that stronger wind.

Two families of AWES exist. Fly-gen machines carry the generator and turbine aloft. A conducting tether sends the electrical power to the ground. Ground-gen machines keep the generator on the ground. The tether carries force, and the ground station converts tether motion to electrical power.

The field splits a second way, on the wing. Soft wings use fabric. They survive crashes and pack small. Rigid wings use composite airframes. They are more efficient, but brittle and more costly.

A third split matters for this report. It is the mechanical force transmission of ground-gen systems. A pumping (yo-yo) machine sends intermittent tension down a moving tether. A rotary machine sends steady torque down a turning set of tethers.

A rotary machine runs continuously. It needs no reel-in phase. That continuity is where the kite-turbine family sits.

The machine in this report is a kite turbine. Kite turbines exploit multiple, short rigid blades. The blades sit on a ring and form a wide diameter rotor. The rotor and further rings tie together as a tensile rotary shaft. A wheel on the ground station holds the rotating shaft.

When the shaft turns this wheel, the ground station generator converts the rotational power to electrical power.

A tensile rotary power transmission (TRPT) is a shaft made of tensioned lines. Rings sit at fixed distances under tether tension. Torque passes from ring to ring, over multiple tethers. The TRPT shaft is prone to bending and twisting, especially in low tension. It is not a rigid driveshaft. Its capacity to transmit torque comes from the tension in the kite lines.

This report presents a search over the design space of kite turbines. The search varies ring geometry, polygon count, rotor count, blade bank angles and blade scale. The objective is to find the feasible geometries for a 5 kW-class machine, and then characterise them. The generator load law and the safety gates stay fixed. The search asks whether one compact machine can hold the power, mass, structural, clearance and twist constraints at once. The report reads the gate and the winner outcome for each constraint in the results.

The report states what the record supports. It also states several classes of failure in developing this system, each with its lesson. One frame holds throughout. The machine flies only in simulation. The analysis pipeline is the instrument. The report treats the instrument with the same distrust it treats the machine.

The field context rests on the literature. Loyd showed that both crosswind modes share one theoretical power ceiling [1]. Diehl set the classification axes [2]. Trevisi built one power model for both modes and found the glide-ratio flip [3]. Cherubini mapped the ground-gen and fly-gen taxonomy [4]. Van der Burg found no dominant wing design [5].

Both wing families stay in play. Van de Kaa and Kamp found no dominant generation architecture [6]. Pereira ranked the design factors and found that many small units beat few large ones [7].

Flather recorded the rope-drive physics that the TRPT descends from [8]. Benhaïem and Schmehl analysed the torque path of a rotating reel [9]. Tulloch modelled the Daisy TRPT and checked the model against measured flight windows [10]. Daisy is the experimental prototype rotary kite turbine whose measured flight data anchor the model (NR-056). This report extends that line of work.

2 Methods

2.1 The kite turbine and its terms

A line is a tension member, like a tether or a bridle. It carries a load only in tension. A pull at one end passes along the line to the other end. It never pushes.

A ring is a polygon of nodes and tubular beams. A ring with blades works as a rotor. Rings without blades work as line spacers. They hold the TRPT lines apart against the line torsion from torque transfer. Holding the lines apart is what preserves torque transmission and prevents torsional collapse.

The polygon count names the number of lines (n_{lines}). A hexagonal polygon connects six lines. Each line meets the ring at one vertex.

For an active rotor, the polygon count equals the blade count. It also equals the number of TRPT line attachments around the ring. A rotor is a set of blades on a ring. A pentagonal rotor carries five blades.

The topmost ring carries the main rotor. The lowest ring is the ground station power take-off (PTO) ring. The PTO ring anchors the shaft at the ground and turns about its central axis.

Blade bank is the blade angle away from the plane of rotation about its chord axis at the ring. The equivalent term in AeroDyn is pre-cone. Rotor bank is that angle set for the whole rotor, so every blade carries it. Blade pitch is the angle of one blade about its span axis. Rotor pitch is that angle set for the whole rotor, as collective pitch. The machine plates show the blade leaning downward.

Blade scale sets the blade size. The power sizing sets the rotor radius from the power, the wind speed and the power coefficient (NR-050). The scale then multiplies the blade span and chord, and the blade mass follows the span cubed. The base span is 0.75 of the power-sized rotor radius, and the base chord is 0.113 of it.

That radius is not the ring radius. Its own slot sets the ring radius r_{hub} . The blade radius r_{rotor} is the aerodynamic size the blade spans from. The blade runs from the ring radius r_{hub} to the tip radius r_{rotor} . The part outside the ring is its overhang. The power sizing inverts the power equation (NR-050):

$$r_{\text{rotor}} = \sqrt{\frac{P}{f_{\text{proj}} C_p 0.5 \rho \pi v^3}}$$

P is the rotor power share and v is the local wind speed at the ring. C_p is the blade-element table coefficient at the sizing tip-speed ratio. f_{proj} is the bank-and-elevation projection factor. The winner decode gives r_{rotor} 2.434 m (NR-050).

Each of these is one design value the search sets. The ring geometry sets the topmost-ring and ground radii, and the rotor count sets how many rings carry blades. The shaft transition cone holds one fixed half-angle, 22 degrees, by construction. Every machine in the campaign builds that cone at the same angle, so the cone stays out of the comparison. It is not a search value (NR-049).

That 22 measures one side of the cone, from the shaft axis out to the cone wall. A reference machine prints the same 22 as the angle across the full apex, one side against the other. The two numbers name different angles, so this cone runs steeper than that machine. The code change that would match this cone to that machine stays open, awaiting a ruling.

At the top of the shaft the rotor turns in the wind. The whole shaft turns with the rotor. The shaft passes the torque from ring to ring. The lowest ring turns the generator at the ground station.

The shaft works only under tension. Tension gives the shaft its capacity to transmit torque. With low tension the shaft bends and twists. A slack shaft cannot carry the load.

Under the generator load and line drag, the machine settles onto one operating point. In the settled state the shaft runs twisted along its length.

Kite turbine anatomy: the ground-generation rotary machine

Island-3 winner geometry (NR-011): six rings, three TRPT lines, one active rotor. The shaft is drawn at the operating elevation. The dashed outline sits outside the primary render scope.




Figure (roster row 2) · system-anatomy

This figure shows the whole machine and the names of its parts, from the lifting autogyro kite stack in the sky to the ground station that takes the power. The plate shows the campaign winner in a side view. Six triangular rings run from the ground ring (R1, 0.804 m radius) to the topmost ring (R6, 4.580 m radius). Three tension lines span each of the five segments S1 to S5. The main rotor sits on R6 only. Its three blades bank at 7.088°. The lift chain runs from the lifter stack, down the lift line, through the sky anchor and the cyan line, to the lift bearing. Three lifting bridles run from the bearing into the main rotor. The lifter gives lift only. It does not drive the rotor. Torque runs down the TRPT rings to the ground station. The dashed outline marks the ground station: the generator and the power take-off. That outline sits outside the primary render scope. The dashed backline is an altitude limiter, not a load path. The scale bar is 5 m.

Source: data dd3cc6a · script 9241b23 · rows NR-011, NR-017, NR-043

2.2 The simulation analysis pipeline instrument and its terms

2.2.1 The gene

The search uses differential evolution (DE) to assess kite turbine design variants. DE assesses machine genomes. A genome is a design vector that describes the physical character of one machine. The decoder reads it, and the builder builds the machine from it.

The evaluation then feeds the comparison before DE tries a new gene. The genes set the ring geometry, the polygon count, the rotor count, the bank angles and the blade scales. The winner genome names its genes in the decode fields (NR-011).

The genome carries ten slots, x1 to x10. The table lists each slot, its meaning and its search bounds.

Slot	Meaning	Bounds
x1	hub ring radius r_{hub}	0.700 to 7.776 m
x2	ground ring radius r_{bottom}	0.398 to 1.195 m
x3	common length-to-radius target L/r	1.0 to 2.94
x4	polygon side count n_{lines}	3 to 9
x5	ring density bias	-0.8 to 0.8
x6	rotor count	1 to 3
x7	bank at the topmost rotor	0 to 22 deg
x8	bank at the lowest rotor	0 to 22 deg
x9	blade scale at the topmost rotor	0.2 to 1.0
x10	blade scale at the lowest rotor	0.2 to 1.0

The decoder reads x6 as a rotor count. It rounds the gene to an integer between one and three. Inactive slots may hold values with no physical effect. An inactive slot may hold a bank or blade-scale value and still create no physical blade. The winner carries one active rotor, so x8 and x10 stay unused (NR-011). The decoder also rounds x4, the polygon side count, to an integer (NR-011).

The ring count builds separately. A spacing pass lays the intermediate rings along the shaft from the length-to-radius target and the density profile. The builder then adds the ground ring and the hub ring at the ends. The winner reads six rings (NR-011).

2.2.2 The evaluator

The builder builds the machine from the genome. The evaluator runs the machine and measures its performance. Every evaluation has two phases: a relax run and a measurement window. The run starts from a settle state.

That state comes from a simplified power balance (NR-025). The relax run takes 10 seconds. It lets the startup transient decay, and the machine moves to its operating point. The measurement window takes 40 seconds.

It measures the settled machine. The evaluator scores the power over the measurement window only (NR-015). A short window can sample the decay and misstate the power. The team sized the window to outlast the decay.

The evaluator names two scoring quantities. P_score reads the tail-5 power, the mean of the last five seconds of the window (NR-046, NR-041). The draft prints it as the end power. The twist ratio is the absolute segment twist divided by the crossing limit, the worst segment over the window (NR-046). The crossing limit follows (NR-053):

$$\delta\alpha^* = 2 \arcsin \left(\frac{L_{\text{seg}}}{\sqrt{2(L_{\text{seg}}^2 + 2r_{\text{seg}}^2)}} \right)$$

The term r_{seg} is the larger of the two ring radii at the segment. The term L_{seg} is the segment length between the two rings.

2.3 The operating point and load selection

The generator load follows a power law. The commanded torque reads $\tau = k \cdot \omega^2$, so the drawn power reads $P = k \cdot \omega^3$. The constant k sets the load, and the load sets the speed at which the machine settles.

The operating point uses a power-law constant k of 2.24 N·m/(rad/s)², a relax time of 10 seconds, a window of 40 seconds, a factor of safety (FoS) target and hard floor, both 2.5, and a line length of 18.8 metres (NR-015). The campaign fixed k at 2.24 N·m/(rad/s)² for every candidate design (NR-015). One load applies to the whole search.

A fixed k couples the radius search to the tip-speed ratio. The load curve fixes the torque-speed relation, and the radius sets the tip speed for a given speed. So the radius gene and the operating tip-speed ratio are not independent. Changing the radius with k held fixed moves that ratio.

The search explored radius under that coupling. The load law kept the same form for every candidate, and the speed and radius stay coupled. The team selected the value in a load sweep of the seed machine (2026-08-22). The sweep ran the seed at a range of loads, each over the same measurement window.

At the light end the seed spins up. The k 1.0 row settles at omega 17.6, tip about 85 m/s, near the 100 m/s ceiling. The ceiling is a hard evaluator gate (TIP_SPEED_CEILING_MPS). The selected k 2.24 settles at omega 15.3, tip about 73 m/s (NR-047).

The sweep rows read as follows (NR-047).

k	status	P_mean (kW)	omega_final (rad/s)	FoS_min
0.5	reject	3.49	19.14	51.8
1.0	ok	5.45	17.64	47.8
1.5	ok	6.78	16.59	42.8
1.94	ok	7.56	15.80	40.1
2.24	ok	8.00	15.34	36.4
3.0	ok	8.89	14.42	24.3

At the light end the seed spins up. The k 1.0 row settles at omega 17.6, tip about 85 m/s, near the 100 m/s ceiling. The ceiling is a hard evaluator gate (TIP_SPEED_CEILING_MPS). The selected k 2.24 settles at omega 15.3, tip about 73 m/s (NR-047).

The power is measured, not computed from a mean speed. The evaluator samples the electrical power at one hertz inside the window. At each sample it reads the generator torque times the ground-ring speed. The window mean of those samples is P_mean. The window-mean ground-ring speed is a separate average (NR-029).

The delivered-torque basis (NR-030) is P_mean divided by the window-mean ground-ring speed, a derived ratio, not the instantaneous k times omega squared. The two differ slightly: P_mean 5.309 kW against k times the cube of the window-mean ground-ring speed, 5.366 kW, a 1.07 per cent gap.

The exact source of the gap is open. It may come from a slow speed trend across the window, from ground-ring drag, or from how the instantaneous power is averaged.

At the heavy end the load raises the torque and the line tension, so FoS falls. The FoS_min column reads 36.4 at k 2.24, down to 24.3 at k 3.0 (NR-047). The heavy end also carries a score cost. The objective charges 5 kg-equivalent per excess kW squared above the 5.0 kW ceiling (NR-046).

The load choice also has a structural side. The team sizes the beam for a bounded generator load step: twice the nominal load (NR-048).

A low-wind lull is harmless under the load law. In high wind the load law holds the speed, and easing k removes that hold. The team refuses to ease k for that reason.

The campaign also fixed the shaft length. The search does not vary it: every candidate runs at one length, 18.8 metres (NR-015). The length follows the Daisy anchor. It scales with the machine to this size.

It sets how high the machine stands above the ground station. A longer shaft reaches stronger wind. The length screen history sits in section 3.3 with the era record.

FoS is the design load margin. It is the ratio of what the beam can carry to what the load demands. The model tests the TRPT beam element FoS for buckling risk.

The dynamics run as a coupled multibody model. Every node, line and ring carries state. The evaluator holds the whole set as one matrix of ordinary differential equations (ODEs). The evaluator advances that matrix at a fixed step. The scheme is semi-implicit. Each step advances the velocities before the positions.

The generator brake torque enters implicitly. Read from the last step alone, the load law is positive feedback. Take a ground ring that runs a little fast: the brake pushes too hard, and the next step over-corrects, so the swing grows. The implicit form solves the brake with the new speed, so the step stays stable. The evaluator derives the step per machine from the shortest shaft sub-segment.

The run is deterministic by construction. That determinism is what makes reproduction possible. The search carries the same property. The campaign runner fixes one random seed per island.

NaN and Inf clamps guard the state. A broken line ends the run.

2.4 The campaign

The campaign runs a differential evolution (DE). A DE keeps a design population. Each generation breeds new designs from random members of the population. A gate scores each design and culls the failures.

Three islands run in parallel. An island is an independent population with its own drift. Each island evaluates 310 designs.

The DE holds ten members per island, three islands, and thirty evolution steps after the initial population. The seed is the first member of each population. The other nine members draw at random inside the bounds. Each island runs 31 generations in total, generation indices 0 to 30, ten rows per generation, 310 evaluations per island (NR-045).

Each island starts from its own random seed (42 plus the island number minus one), so the three islands drift apart. No member moves between islands. Each new member forms by a standard DE rule. The mutation takes the difference of two members, scales it by a factor F , and adds it to a third. F draws at random in 0.5 to 0.8 per trial.

The crossover rate CR draws in 0.7 to 0.9 per trial, and one gene always crosses. The trial replaces the parent only if its score is not worse. The campaign evaluates 930 designs in total (NR-001).

Four statuses result: ok, reject, clearance_reject and reject_twist (NR-002). The ok designs pass every gate. The rejects fail a gate. Clearance rejects fail the clearance gate. Twist rejects fail the twist gate.

The campaign applies its gates in one order. The order fixes the status a design carries when it fails more than one gate. The table names each gate, its formulation, its threshold, where it acts, the status it stamps and its register rows.

Gate	Formulation	Threshold	Where it acts	Status on failure	Rows
Genome screens	One active rotor, one or more rings, rotors spaced apart	Above zero, above zero, spacing clear	Runner, on the decoded geometry	reject	NR-046
Clearance	Lowest active rotor tip above the ground	1.5 m	Runner, before any simulation time	clearance_reject	NR-042
Rope break	One line path strained past its break strain	Break strain	Evaluator, measurement window only	reject	NR-013
Twist crossing	Worst segment twist against the geometric crossing limit	Crossing limit	Evaluator, on the window	reject_twist	NR-052, NR-053
Tip-speed ceiling	Rim speed of every ring and rotor tip	100.0 m/s	Evaluator, on the window	reject	NR-047
Betz, per rotor	Each rotor against the Betz limit on its bank-projected annulus	Limit with a margin, at the rating wind	Evaluator, on the window	reject	NR-034
Betz, availability	Power against a coefficient floor on the normal wind area	Below the power floor, at the rating wind	Evaluator, after the window	reject	NR-027, NR-034
Betz, aggregate	Window-mean power against the Betz limit on the total area	Limit with a margin, at the rating wind	Evaluator, after the window	reject	NR-034
Power floor	Window-end power against the power floor	5.0 kW	Score seam	reject	NR-027, NR-039
Safety floor	Window-minimum factor of safety against its floor	2.5	Score seam	reject	NR-015
Charges	Four costs added to the airborne mass	Weights below	Score, after the gates	None, a cost only	NR-039, NR-046

2.4.1 The genome screens and clearance

The runner screens the decoded geometry before it runs any simulation time. It needs one active rotor and one or more rings. It also needs the rotors clear of each other. A rotor tip below 1.5 m of ground clearance fails the clearance gate (NR-042). A failure stamps clearance_reject, and the run never starts.

2.4.2 The rope break

A line path carries the tension from ring to ring. A strain past the break strain ends the run and rejects the design (NR-013). The check arms at the relax boundary and reads the measurement window only. Detection off the window would reject healthy machines, and section 4.1 holds that record.

2.4.3 The twist crossing

The twist gate reads the worst segment twist against the geometric crossing limit (NR-052, NR-053). A segment past its limit stamps reject_twist. The same ratio enters the twist charge, so a design near the limit pays a growing cost before it fails.

2.4.4 The tip-speed ceiling

Every ring rim and every rotor tip carries a speed ceiling of 100.0 m/s (NR-047). The gate reads each of them, and a design past the ceiling rejects.

2.4.5 The three Betz screens

The campaign carries three Betz screens, not one. The per-rotor screen compares each rotor against the Betz limit on its own bank-projected annulus (NR-034). The availability screen compares power against a coefficient floor on the normal wind area, below the power floor (NR-027, NR-034). The aggregate screen compares the window-mean power against the Betz limit on the total area (NR-034).

All three read the rating wind, and that basis is deliberate. They are an anti-cheat tripwire on the rated machine, so they evaluate at the rating wind. They do not evaluate at the wind of a later sweep.

The three screens charge different area bases. The per-rotor screen reads the annulus with no elevation factor. The availability and aggregate screens read the wind-normal area, which carries the elevation factor (NR-034). The aggregate screen charges the window mean, and the power floor beside it charges the window end.

Section 3.3 carries the ceiling value. The per-wind check on the power curve is a figure check, and it is not this gate.

2.4.6 The power floor and the safety floor

The power floor reads the window-end power against 5.0 kW (NR-027, NR-039). The safety floor reads the window-minimum factor of safety against 2.5 (NR-015). The seam names its power parameter P_mean, and the value it carries is the window-end power. The report prints the window end as the power behind the floor.

A torsional safety floor of 1.5 belongs to the design-verification path (NR-052). The two campaign floors read 5.0 kW and 2.5.

2.4.7 The objective and its four charges

The campaign minimizes a kg-equivalent objective score: the airborne mass plus the charges. A lower score is better. The objective score is not the physical airborne mass. One kilogram of airborne mass adds 1.0 to the score (NR-046). The charges read in kg-equivalent.

The score builds along one chain. The chain runs genome, decoder, builder, evaluator, window statistics, gate status, charges, score and rank. The score adds three charges to the airborne mass (NR-046):

$$f = \text{mass} + 5.0 \max(P_{\text{score}} - 5.0, 0)^2 + 20.0 \left(\frac{2.5}{\text{FoS}_{\min}} \right)^2 + 1.0 \left(\frac{r}{1-r} \right)^2$$

The ratio reads $r = \min(\text{twist_ratio}, 1)$. The score then adds the stationarity charge (NR-046):

$$10.0 \max(0, \text{swing} - 0.20)$$

The formula carries three charges. The stationarity charge makes four.

The over-power charge reads P_score, the end power, against the 5.0 kW ceiling (NR-046). The utilisation charge reads the window-minimum factor of safety against the 2.5 floor (NR-046). The twist charge reads the worst segment twist ratio against the crossing limit (NR-046). The stationarity charge reads the window swing against 20 per cent (NR-046). The winner carries all four in its own score: mass plus 0.4933, 0.5577 and 0.6141, and a stationarity charge of zero (NR-039).

The swing reads P_range over P_mean, the window power range over its mean (NR-046). The twist charge clamps r to one. At r equal to one the segment reaches the crossing limit, so the ratio r over (1 minus r) would divide by zero.

The code floors the denominator at 1e-6. The ratio reads one million at the limit and stays finite. The twist gate flags a segment that passes its limit, so a design at the limit exactly reaches the floor (NR-046).

2.4.8 The reject seam

A rejected design carries no score (NR-046). The evaluator carries Inf for a rejected design. The runner stamps the search cost 1e9 on every non-ok design (NR-046). No reject ranks above an ok design. The search cannot trade power or safety against mass.

The differential-evolution campaign fixed the configuration below.

Setting	Value	Source
Islands	3	NR-001

Setting	Value	Source
Designs per island	310	NR-001
Evaluations, total	930	NR-001
Objective	Airborne mass + penalties, kg-equivalent	NR-036, NR-039
Power floor	5.0 kW	NR-027, NR-039
Power ceiling	5.0 kW	NR-046
Structural FoS floor	2.5	NR-015
Clearance floor	1.5 m	NR-042
Tip-speed ceiling	100.0 m/s	NR-047
Torsional FoS floor	1.5	NR-052
Twist gate	geometric crossing check	NR-052, NR-053
Relax run	10 s	NR-015
Measurement window	40 s	NR-015
Load constant k	2.24 N-m/(rad/s) ²	NR-015
Line length	18.8 m	NR-015

2.5 Fidelity paths

Two paths score a design. The ODE path builds the machine, settles it, and runs the full dynamics over the measurement window. That path produces the score the campaign keeps. The rapid path pre-solves a static state and starts from it. It serves as a screening shortcut.

The report names map to the code. The ODE path is :cold. The rapid path is :warm. The campaign scored every design on the cold path (NR-012). The two comparison probes are the only warm runs in the record.

The probes that compare the paths sit in scripts/probe_rapid_vs_ode_winners.jl and scripts/probe_warm_reject_cause.jl. The register header pins their commits.

The rapid path is the less reliable of the two. It once rejected all three winners, through gate placement and not the machine (NR-013). Section 4.1 holds the record, and the fix has landed.

2.6 Reproducibility

Every dataset carries a provenance stamp. The stamp names the commit, the physics era and the geometry fingerprint. The report names the commits in the register header. The dataset sits at data commit dd3cc6a on the bankderate2-results branch.

The launch commit is a76c5e9. The Phase 0 audit sits at master 31ca00d. The campaign runner script, scripts/run_v13_5kw_masslift.jl, pins the island seeds.

3 Results

3.1 The search at scale

The campaign evaluated 930 designs (NR-001). The status mix reads 508 ok, 301 reject, 110 clearance_reject and 11 reject_twist (NR-002). The reject class splits three ways (NR-055). 96 ran below the 5 kW power floor. Four ran at or above 5 kW with a factor of safety below 2.5. The remaining 201 carried no logged cause at campaign time.

A re-run probe now names those 201 (NR-061). They split six ways. 93 rotors cannot reach 80 per cent of the power floor on their swept area and power coefficient. 37 settle at the torsional realisability cliff. 36 refuse on rotor spacing.

30 break a line at the rope-break gate, not a beam FoS breach or a twist crossing (NR-013). Three cross the Betz ceiling. Two diverge at the hub tip speed. Only the last class is a numerical divergence. The others are geometry or design refusals that the telemetry left blank.

Each island starts its population from the same seed genome. The recorded common seed genome scores 36.109 kg-equivalent under the campaign objective (NR-006). The other members of each population draw at random under a fixed seed per island. This campaign is the first search campaign of the wind-authority era. Larger and wider campaigns sit ahead.

3.2 Convergence

The best recorded objective score on island 1 is 30.578 kg-equivalent at generation 30 (NR-003). Island 2 reaches 30.637 kg-equivalent at generation 30 (NR-004). Island 3 reaches 27.364 kg-equivalent at generation 26 (NR-005). The physical airborne mass of the winner is 25.698 kg (NR-018). The objective score adds 1.665 kg-equivalent of penalties (NR-039). The surviving ok designs hold a mean power between 5.02 and 6.32 kW (NR-007).

Convergence in this section reads the search, not the window. The three results above read the best score per island against generation. The window test is a separate reading.

It compares P_mean and P_end. The two agree below 0.5 per cent at every row of the load sweep (NR-041). The heavy-end rows of that sweep sit in a band, 8.9475 to 8.9511 kW at 4 dp. The register reads that band as a plateau.

3.3 The winner

The island-3 winner at its operating point: plan, side and downwind

Built geometry with the recorded wind-up: six rings, three lines, one rotor (NR-011); per-segment twist at the final operating step (NR-022).
Ring stations read at the same step (NR-021). One scale holds across the three views. Blades: three, span 1.499 m (NR-043).




Figure (roster row 10) · machine-winner-hero

The campaign winner, drawn three ways: down the shaft, from the side, and along the wind. The machine stands at its settled operating point. The six rings sit at their recorded stations. The shaft carries its recorded wind-up. Each ring turns against the ring below it. All three views share one scale, so the drawn size reads true.

The island-3 winner at its final operating step. The machine reads six rings, three lines and one active rotor (NR-011). The ring stations and the per-segment twist come from the signed capture (NR-021, NR-022). The main rotor carries three blades of span 1.499 m and a bank of 7.088 degrees (NR-043, NR-017). The plan view looks along the shaft. It draws each ring as its node polygon. The side view uses the shaft frame. Distance along the shaft is vertical, and the operating elevation is not drawn. The downwind view looks along the wind, so the shaft reads foreshortened. The swept annulus draws as one truncated cone in the side view. It draws as a ring band in the plan view and as a flattened band in the downwind view. One scale holds across the three views. A 5 m scale bar sits in the footer.

Source: pair-extract-2026-10-08-winner-ops · data dd3cc6a · script d0e410d · rows NR-011, NR-017, NR-021, NR-022, NR-043

Island 3 carries the campaign winner, the best feasible candidate under the modeled objective and gates. Its mean power reads 5.309 kW and its end power 5.314 kW (NR-008). Its factor of safety reads 14.972 (NR-009). That value is the window minimum. The capture topmost ring reads a factor of safety of 15.033 at the final step (NR-021).

Its operating point reads as follows.

Quantity	Value	Source
Mean power, measurement window	5.309 kW	NR-008
End-of-window power	5.314 kW	NR-008
Minimum factor of safety	14.972	NR-009
Window power swing	0.029 kW	NR-014
Settle angular speed	10.170 rad/s	NR-016
Window-mean speed	13.38 rad/s	NR-029
Delivered torque	396.76 N·m	NR-030
Airborne mass	25.698 kg	NR-018
Topmost ring radius	4.580 m	NR-011
Ground radius	0.804 m	NR-011
Ring count	6	NR-011
Polygon count	3	NR-011
Active rotors	1	NR-011
Main-rotor bank, ring 6	7.088 deg	NR-017
Lift-line elevation	70 deg	NR-026
Lift-line tension	324.1 N	NR-010
Reference wind speed	11.01 m/s at 9.41 m	NR-008, NR-032
Shear law	Hellmann power law, exponent 1/7, from the measured Daisy pair 10.0 m/s at 4.8 m	NR-056
Shaft elevation	30 deg, hub at 9.4000 m	NR-057
Ring speed ratio	5.566	NR-031

Quantity	Value	Source
Tip-speed ratio	6.8415	NR-032

Its window swing reads 0.029 kW on 5.309 kW, half a per cent (NR-014). The reported trace runs steady wind, so the swing reads as a steady-wind response, not a stability property (NR-058). The machine held a steady operating point through the measurement window. Steady and drifting describe two observables on two timescales. The power settles: the window swing reads half a per cent of the mean (NR-014). The extension speed spread reads 0.009 per cent (NR-059).

The bow creeps slowly: each five-second step reads smaller than the last (NR-059). At the hub the extension steps fall from 0.0254 m to 0.00053 m (NR-059). So the machine holds its operating point while the shaft still creeps toward its bowed shape. The settle state starts the machine at 10.170 rad/s (NR-016). The machine accelerates through the relax run.

The shaft axis bows off its design line, sagging straight down. At the midpoint ring the bow reads 0.002 m at settle, and 0.105 to 0.310 m through the measurement window (NR-059). The maximum sits at the hub and reads 0.209 to 0.613 m through the window (NR-059).

An extension run past the window ends at 0.422 m at the midpoint and 0.833 m at the hub (NR-059). Every measured state at the midpoint ring and the hub reads below the 1.5 m comparison bar (NR-059). The record ends at 200 seconds, so the report claims no bound past the measured span (NR-059).

The settle hands off a straight stack while the run bows it. The two do not agree on the bowed shape. The disagreement reads as an instrument finding, not a design question (NR-059).

The cause of the bow is open. Gravity, frame drag and line drag all act in the wind-height plane. The model carries each. The straight-down direction does not by itself attribute the split between them.

The swept area reads 45.9733 square metres raw and 39.4912 square metres on the wind-normal basis (NR-033). The air density reads 1.225 kilograms per cubic metre (NR-034).

The rotor Cp reads from the AeroDyn v5.0.0 blade element momentum (BEM) table. The value reads 0.2944 at the ring speed ratio, and 0.2258 at the tip-speed ratio (NR-031, NR-032). The Betz ceiling reads 19.1304 kW, and the winner delivers 27.75 per cent of it (NR-034).

The winner clears each checked gate.

Gate	Threshold	Winner outcome	Margin	Source
Power floor	5.0 kW	pass, P_end 5.314 kW. Margin +0.314 kW / +6.28 per cent	NR-008, NR-027, NR-039	
Structural FoS floor	2.5	pass, 14.972 minimum. 14.972 on 2.5, margin +12.472 / +498.9 per cent, 5.99 times	NR-009, NR-015	
Clearance	blade-tip gate	pass, 5.468 m, margin +3.968 m	+264.5 per cent, 3.65 times	NR-042
Twist	per-segment limits	pass, all five segments within limit	S1 binds: 33.946 deg against the 72.25 deg realisability floor, 2.13 times (NR-053). Against 98.859 deg δ_{crit} , 2.91 times (NR-022). Crossing ratio 0.4393, 2.28 times (NR-052)	NR-022, NR-052, NR-053
Torsional capacity	capacity/demand	least margin 2.03 at S1	+102.9 per cent above unity, 2.03 times	NR-040
Rope break	none on the window	pass on the cold path	the warm path rejects healthy winners	NR-012, NR-013

The gate set reads per class. The structural members and the lines hold a factor of safety of 2.5 or higher (NR-015). Torsion keeps its own criteria. The collapse factor of safety reads 1.5 or higher, and the twist stays within the crossing limit (NR-052). The realisability floor reads 72.25 degrees (NR-053). As run, the 2.5 gate does not cover torsion (NR-052).

A lifter kite holds the machine aloft. The lifter gives lift only. It gives no torque and no drive. Its load path runs from the lifter kite down the lift line to the sky anchor, the three-way knot. Then it runs down the cyan line to the lift bearing, and down the lifting bridles to the main rotor.

In the field the lifter launches first. It lifts the rig into tension. The rotor turns after that. The lift acts throughout operation, not only at spin-up.

The team sizes the lift line up front. The lifter sizing holds a 1.5 margin on the vertical load at the lift bearing. The sizing mass is the airborne mass without the lifter itself, 20.698 kilograms. The kite carries its own weight (2026-08-21 ruling).

The vertical load is $1.5 \times 20.698 \times 9.81 = 304.6$ newtons. The lift line runs at 70 degrees of elevation (NR-026). The line tension is that load over $\sin 70$ degrees: 324.1 newtons (NR-010).

The lifter is a stack of autogyros. The report assumes the stack can modulate to meet the lift requirement at all wind speeds. A fixed kite cannot: its pull grows as the square of the wind speed. Without modulation, each weather band would need its own single kite. That kite would reach giant sizes and be too hard to handle.

The report applies the lift force the machine needs and tests the machine on that basis. Designing a lifter that delivers that force across wind speeds is outside this report. The report states the presumption as one (2026-08-05). The report also assumes a staged launch of the stack. A staged launch is easier to handle than one giant kite.

The winner genome contains two blade-scale slots. The main-rotor slot, `blade_scale_top`, reads 0.8213 and decodes to the single active rotor. It is quoted 0.82. The second-rotor slot, `blade_scale_bottom`, reads 0.20 and stays inactive because the rotor count is one (NR-011).

The winner genome contains two bank-angle slots. Ring 6 decodes the active main-rotor bank angle as 7.088 degrees (NR-017). The second bank slot holds `bank_bottom`. It reads 22 degrees at the top of its search box. It does not decode to a physical rotor at a rotor count of one (NR-011, NR-017).

This result differs from earlier eras for recorded reasons. The bank derate, the measurement window and the tightened gates moved the search space. Numbers across the eras would compare models, not machines.

The three movers carry their records. The bank derate is a calculated depowering of the main-rotor disc (NR-019, NR-020). The window basis sits at section 2.2.2, and the gate set at section 2.4.

Era	Effect on the search space	Data status
Before the bank derate	Earlier winners read different bank-angle physics, so they are not directly comparable.	Superseded as era record.
The bank-derate basis	The measured bank derate applies to main-rotor disc power throughout (NR-019, NR-020).	Current campaign basis.
The cp fix and the folded class	The span shrink and the cp fix moved the screen readings. The pre-fold rows failed the floor (NR-027).	Recorded: the folded class passes at 18.8 metres (NR-028).
The wind-authority era (post-95c385d_wind-authority)	One wind authority reads the scored results.	Current physics era.

This run also trades breadth for fidelity: fewer evaluations, each on the measurement window. It stopped early: a screen, not a survey. This report covers the wind-authority era campaign and the recorded corrections that shaped it.

The length screen ran 2026-10-05 on the pre-fold island-3 winner. Two terms name the two states of that design. The pre-fold winner is the island-3 genome as it then stood, `r_hub` 3.8347 m and `blade_scale_top` 0.6994.

The fold is the 2026-10-05 re-centering that pinned `r_hub` to 4.32 m and `blade_scale_top` to 1.0 (the S2-fold seed). The folded class is the set of designs after that re-centering.

Two further terms name the two code changes in play. The cp fix is the 2026-10-05 power-coefficient correction (era a82cafd). The span shrink is the D1 re-decode that reduced the blade span, which cut the pre-fold power below the floor. The screen covered 18.8, 21.2 and 25.0 metres, twice, before and after the cp fix. All six rows failed the power floor (NR-027).

The cause is recorded. The screen measured the pre-fold winner under the post-landing decode, whose span shrink took the power below the floor. The folded class passes at 18.8 metres (NR-028). The campaign searched the folded class, so the launch carries the passing physics.

3.4 The levers of performance

Every checked ring sits far below its fail boundary

Beam-column interaction per checked ring · island-3 winner operating point





Figure (roster row 6) · ring-utilisation

Every airborne ring passes its structural check with a wide margin. The topmost rotor ring R6 governs the structural check at FoS 15.033. All other airborne rings hold much larger margins (NR-021). Each bar shows one checked ring. The bar end reads the ring utilisation: the beam-column interaction on the worst beam, the axial share N/N_{crit} plus the bending share $\sqrt{(M_{ip}^2 + M_{oop}^2)}/M_{el}$. The rows run from R2 at the base to R6 at the top. Utilisation is demand over capacity. The axis is linear from zero to the unit value. The margin of every ring reads as a direct fraction of capacity. The red dashed rule is the fail boundary at the unit value, the physical interaction limit at FoS 1. The purple rule is the campaign structural gate. It draws at utilisation 0.40, the reciprocal of FoS 2.5 (NR-015). The readout at the row end gives the ring FoS: capacity over demand, the reciprocal of the interaction. The model tracks torsion outside the interaction formula. The figure covers every airborne ring. The check does not cover the ground ring, which sits on ground support. The figure includes and marks the topmost ring. The values come from the campaign winner at its final operating point. R5 reads far below the other rows, so it draws a tension badge instead of a bar. It carries net axial tension from the cone flare (NR-060). The check scores compression only, so the axial share enters the interaction as zero. The badge readout prints the true value, FoS 326,974 (NR-021). For every other checked ring, the bending share sits below 0.05 per cent of the ring utilisation.

Source: pair-extract-2026-10-08-winner-ops · data dd5a9dd · script 9c4646a · rows NR-021 · NR-015 · NR-060

Every segment of the winner runs well inside its binding limit

Transmitted twist against its binding limit, the realisability floor · shaft profile S1-S5, ground to top · island-3 winner operating point



source: data commit dd3cc6a · extract extract-winner-operating-point.csv · capture pair-extract-2026-10-08-winner-ops

s = 1724859 (t = 49.99999483862295 s) · rows NR-022 · NR-023 · NR-040 · NR-052 · NR-053 · script commit e3d664c · authority commit f8114be

Figure (roster row 7) · twist-limit

The shaft passes its twist and torque checks along its whole length.

All five TRPT segments remain inside their over-twist limits at the winner's final operating point. S1 carries the least segment torque margin, 2.03 (NR-040). Panel A draws the transmitted twist of each segment as a magnitude, against its limits: the over-twist limit δ_{crit} and the realisability floor. δ_{crit} comes from the committed authority `trpt_twist_limit`. The floor follows $\sin \Delta\alpha \leq 1/1.05$ and holds the geometry only to 72.25° (NR-053). The crossing limit $\delta\alpha^*$ is a per-segment collapse-detector limit. It is not the over-twist limit. A no limit segment carries no torsional ceiling. Its tension lines cannot reach the axis, so no over-twist limit exists. The figure draws no ratio for such a segment. Tether break strain and ring compression carry its limit instead. Panel B pairs segment torque against torsional capacity at the final operating point of the campaign winner. Each segment draws the torsional FoS gate at capacity divided by 1.5 (NR-052). Pair a capacity only with the load of the same state.

Source: pair-extract-2026-10-08-winner-ops · data dd5a9dd · script 58f0022 · rows NR-022 · NR-023 · NR-040 · NR-052 · NR-053

The topmost rotor ring R6 governs the structural check at the winner operating point (NR-021). The window-minimum factor of safety is 14.972 (NR-009). At the final step, R6 reads 15.033 (NR-021). R2 to R4 read 29.8 to 35.2 (NR-021).

R6 binds because its beams carry compression. The ring check charges compressive axial load, not tension. R6 reads the largest axial share, 0.0665 (NR-021). Compression drives the buckling term, so the most compressed ring binds.

R5 carries net axial tension, not no load (NR-060). Its beams read -157.6 to -142.8 N of tension. The check does not charge tension, so its axial share reads zero and its factor of safety reads far above the others (NR-060, NR-021).


The beam load and the line load differ. The lift line carries 324.1 N of tension (NR-010). The ring beams hold compression at R6 and at R2 to R4, and tension at R5 alone. The ground ring sits on ground support and stays outside the ring check.

All five TRPT segments remain inside their over-twist limits at the winner operating point (NR-022). S1 to S4 transmit 33.5 to 33.9 degrees against a 98.9 degree limit (NR-022). S5 transmits 22.8 degrees against its 92.6 degree limit (NR-022). The least segment torque margin is 2.03, at S1 (NR-040). S5 carries the largest segment torque margin, 2.69 (NR-040).

The winner is not close to ring buckling, over-twist or segment torque saturation at this operating point. The margins are conditional. They cover the checked TRPT rings and segments only. They do not yet include blade integrity or lifter loads across wind bands. The next steps add the blade checks and the gust cases.

V13 campaign — parameter-space landscape

the 930-design screen · three DE islands · PC1-PC2 projection of the ten-gene search · colour is the search objective (kg-equivalent)



THE THREE JOURNEYS

930 designs · three islands · 31 generations each
one common seed at 36.109 (kg-equivalent)
winner: island 3 at 27.364 (kg-equivalent)
each line: the island best-so-far path, one step per generation

source: island telemetry · data commit dd3cc6a · run a76c5e9 · era post-95c385d_wind-authority · rows NR-001, NR-005, NR-006, NR-044, NR-045
projection: z-scored 10-gene matrix, SVD · variance shares NR-044 · script commit d0e410d

Figure (roster row 20) · pca-landscape

One seed, three islands: this figure maps the search journeys in gene space and shows where the light machines live. The figure projects the 930-design screen of the v13 5 kW campaign on its first two principal components (NR-001). The axes come from a z-scored matrix of the ten genes. Colour is the cell mean of the search objective in kg-equivalent. Bright reads lower. Three curves trace the island best-so-far values, one step per generation. The common seed starts at 36.109 (NR-006). The winner ends at 27.364 (NR-005). The PC1 inset reads larger main-rotor blades, wider ring spacing and a wider ground ring to the right, and more TRPT lines and more rotors to the left. The PC2 inset reads a deeper second-rotor bank, smaller second blades and rings biased toward the ground upward. The variance shares read PC1 22.4 per cent and PC2 17.3 per cent (NR-044).
Source: data dd3cc6a · script d0e410d · rows NR-001, NR-005, NR-006, NR-044, NR-045

V13 campaign — parameter atlas

the same PC1-PC2 projection, coloured once per gene · PC1 → horizontal · PC2 ↑ vertical · star = common seed, diamond = winner · thin lines = island journeys




Figure (roster row 21) · pca-atlas

One map of the search space, coloured ten times: each panel shows where one design gene runs high or low across the search. The figure repeats one PC1-PC2 projection of the 930-design screen ten times, once per gene slot (NR-001). The projection and the journey marks match the levers landscape plate. Each panel colours the cell mean of that gene over all 930 designs. The star marks the common seed (36.109, NR-006). The diamond marks the winner (27.364, NR-005). The thin curves trace the island journeys. The colour bars print each gene at three significant figures. The projection reads NR-044: PC1 22.4 per cent, PC2 17.3 per cent.
Source: data dd3cc6a · script d0e410d · rows NR-001, NR-005, NR-006, NR-044

The corner machines: one typical design from each corner of the PC1-PC2 field




Figure (roster row 23) · machine-corner-grid

The four corner designs show how widely the machine changes across the search field, from multi-line two-rotor stacks to single-rotor cones. Four panels show the nearest feasible design to each corner of the [1,99] percentile window of the certified PC1-PC2 projection (NR-044, NR-051). The picks read i1-g13-i1, i1-g7-i7, i2-g21-i4 and i3-g26-i8. Their search objectives read 60.646, 58.864, 33.704 and 33.812 kg-equivalent at 3 dp. Each panel draws the built state of that genome at one shared scale. The drawing shows the ring segments in their section colours, the TRPT lines between rings, the blade segments, and the swept annulus of each rotor as one truncated cone. Distance along the shaft is vertical, and the operating elevation is not drawn. The picks are typical machines for their corners, not optima. All four objectives sit above the winner value (NR-005). The plate claims geometry and the search objective. It prints no power value and no airborne mass value.

Source: data dd3cc6a · script d0e410d · rows NR-044, NR-051

The two plates map the same search in two views. They project the 930 designs of the campaign onto the first two principal components of the ten-gene search (NR-001). Those axes hold 22.4 and 17.3 per cent of the variance in the search, 39.7 per cent together (NR-044).

The horizontal axis is PC1: blade scale against unit count. Its inset reads larger main-rotor blades, wider ring spacing and a wider ground ring to the right. It reads more TRPT lines and more rotors to the left. The top loadings read -0.5385 for the TRPT line count, -0.5222 for the rotor count and +0.5180 for the main-rotor blade scale (NR-044).

The vertical axis is PC2: second-rotor and clustering settings. Its inset reads a deeper second-rotor bank, smaller second blades and rings biased toward the ground upward. It reads even rings, a shallow second bank and larger second blades downward. The top loadings read +0.5621 for the second-rotor bank, -0.5605 for the second-rotor blade scale and -0.5018 for the ring density profile (NR-044).

The second-rotor bank and blade-scale slots are inactive at a rotor count of one (NR-011, NR-017). The winner carries one rotor, so those two slots hold no physical blade. PC2 therefore reads mostly unused genes for the winner and for every one-rotor design in the field. The axis is real, but the inactive second-rotor slots dominate it. The active ring-density gene also loads it, at -0.5018 (NR-044).

The colour is the mean search objective in each cell, in kg-equivalent. Bright reads lower, and only feasible designs shade the field. The low-objective cells gather toward the upper right, where the winner sits.

Three curves trace the island best-so-far values, one step per generation (NR-045). Each island ran 31 generations. All three start at the common seed, 36.109 kg-equivalent (NR-006). Island 3 reaches the winner value, 27.364, at generation 26 and holds it to the end (NR-005). Islands 1 and 2 finish lower in the field.


The atlas repeats the projection ten times, once per gene slot. Each panel colours the mean value of one gene across the field. The star marks the common seed and the diamond marks the winner. The atlas lets a reader follow one gene at a time across that ground.

The two plates fix the search frame. They show where the search went, how light the designs get and how each gene tracks across that ground. The gene-level effects wait on the landscape extracts.

3.5 The evolution story

Island 3 best-so-far: seven machines from the seed to the winner

Built geometry at one shared scale - each machine prints its search objective (kg-eq, lower is better), its decoded mass (kg) and its window-end power (kW) - generations 0 - 6 - 10 - 18 - 22 - 25 - 26



source: data commit dd3cc6a - extract strip-extract-2026-10-09-best-so-far.csv - island 3 best-so-far (0, 6, 10, 18, 22, 25, 26)

rows NR-005 - NR-006 - NR-011 - NR-036 - NR-062 - NR-066 - script commit ce9c5d8 - engine CairoMakie - PDF date normalised to 1970-01-01

shaft frame: distance along the shaft is vertical - the operating elevation is not drawn

Figure (roster row 9) - machine-evolution-strip

The search improves the machine at every recorded step, from the common seed to the campaign winner. Island 3's best-so-far objective falls from 36.109 to 27.364 kg-equivalent. The built geometry converges to the register winner (NR-006, NR-005). Seven panels show the best machine of island 3 at generations 0, 6, 10, 18, 22, 25 and 26. The labels match the recorded generations. Each panel draws the built state of that genome. The drawing shows the ring segments in their section colours, the TRPT lines between rings, the blade segments, and the swept annulus of the main rotor as one truncated cone. Each machine prints three signed values. The first is the search objective in kg-eq. The second is the decoded mass in kg. The third is the window-end power in kW (NR-036, NR-066, NR-062). NR-062 signs the window mean and the window end as a pair. The power on each card is the window-end value the score uses (NR-062). The drawings use the shaft frame. Distance along the shaft is vertical, and the operating elevation is not drawn. One fixed scale holds across the set. A 5 m scale bar sits in the footer. The set reads in two rows: four panels above, three below. The machines differ in drawn size only where their built size differs. The final panel carries the campaign winner. It matches the register machine: three lines, six rings, topmost ring radius 4.580 m, ground radius 0.804 m, blade scales 0.82 and 0.20. The strip claims geometry, the search objective, the decoded mass and the window-end power. The objective is mass plus penalties. The penalties score the window-end power, the beam utilisation and the twist margin. The winner is not the lightest machine. Gen 10 reads 25.466 kg against the winner's 25.698 kg. The winner pays 0.233 kg of mass to shed 4.515 kg of penalty. From gen 10 the mass moves inside 1.04 kg while the objective falls 4.297 kg. The later gains come from penalty, not from mass. Three decimals on the mass are a decomposition against a three-decimal objective. They are not a measurement.

Source: data dd3cc6a - script ce9c5d8 - rows NR-005, NR-006, NR-011, NR-036, NR-062, NR-066

The machine evolution strip traces the search on island 3 from the common seed to the campaign winner (NR-036). The objective falls from 36.109 kg-equivalent at generation 0 to 27.364 kg-equivalent at generation 26 (NR-036). The largest recorded drop falls between generations 0 and 6, from 36.109 to 32.293 (NR-036). The final recorded step reads 28.315 to 27.364, a drop of 0.951 (NR-036).

The winner value is the search objective, not a mass reading. The objective reads the airborne mass 25.698 kg plus 1.665 kg-equivalent of penalties (NR-018, NR-039). The sum at three decimals reads 27.363, not the 27.364 the register quotes. The full-precision values add to 27.36354, which rounds to 27.364. The independent three-decimal rounding of mass and penalty is the 0.001 gap.

The final panel matches the register winner. The panel reads three lines and six rings (NR-011). The topmost ring radius is 4.580 m and the ground radius is 0.804 m. The blade scales read 0.82 and 0.20 (NR-011). The winner genome carries inactive slots at a rotor count of one (NR-011, NR-017).

The shaft transition cone holds one fixed half-angle, 22 degrees, in every machine of the campaign (NR-049). A later campaign can vary it.

The strip shows built geometry and three signed values on each card (NR-036, NR-062, NR-066). The values are the search objective, the decoded mass and the window-end power. The power on each card is the window-end sample that the score charges (NR-062). The strip shows that the winner is not the lightest machine (NR-066). The later fall in the objective comes from penalty, not from mass (NR-036, NR-066).

The strip reads one island of the campaign. It does not prove a global optimum. A longer or broader run would extend the search.

Window-end power drives the over-power penalty




Figure (roster row 30) · machine-power-trend

The score charges window-end power, squared, above a 5 kW floor. The last machines get lighter. Their power comes down with them. Island 3 falls from 5.89 kW at generation 10 to 5.31 kW at the winner. That fall is 3.480 of the 4.515 kg-eq of penalty the winner sheds. The rest is beam utilisation, less a small twist cost. No machine can cross the floor. The score rejects a design below 5 kW, so the one step that lies on the rule is the tightest pass of the campaign and not a limit that got close.

The running-best series draws as steps. The running-best machine changes only when a better design lands, so the power holds flat between changes. No step interpolates between generations. Each generation carries three points, one running-best machine per island. The plate draws 93 values from 930 recorded designs. The plate prints four window-end values on the island 3 running-best series. The seed reads 5.86 kW. Generation 6 reads 5.4 kW. Generation 10 reads 5.89 kW. The winner reads 5.31 kW. Each printed value is one member of the pair NR-062 signs. The strip cards carry all seven. The 5 kW rule is a gate and not a datum. The fitness function returns infinity below the power floor (objective_v12.jl:143), and the evaluator passes the window-end power into that slot (objective_evaluator.jl:1153). Of 508 ok evaluations (NR-002), none reads below 5.00 and exactly one reads 5.00 (NR-069). The next lowest reads 5.02. That one machine holds the running-best place from generation 8 to generation 14, and its window end sits below its own window mean. The plate labels it the tightest pass, one design of 508. Island 3 peaks at generation 10. The machine is lighter than the seed there and it still draws more power. The score improves at that step because the mass falls 4.7955 kg, and the penalty rises 4.1548 kg against it. The generation 10 penalty, 6.1670 kg, is the largest of the seven. From generation 10 to generation 26 the machine gets heavier by 0.2328 kg and the penalty falls 4.5146 kg. The fall splits three ways: over-power 3.4800, utilisation 1.1682, twist up 0.1336. The over-power charge is $5 \times (P_{\text{end}} - 5)^2$, so most of that relief is the window-end power coming down from 5.89 to 5.31. The split reads the three penalty columns of extract-mass.csv at ce9c5d8. NR-066 signs the mass and the totals, not the split. Each component is rounded to four decimals and the total is computed unrounded, so a row's parts can sum one digit from its printed total. Panels 1 and 6 read one digit low: 4.5796 against 4.5797, and 2.1265 against 2.1266. The plate therefore does not claim that power is independent of the score. The power path is one of the scored terms, and the last relief comes mostly from it. The strip caption states the same trade as 0.233 kg of mass for 4.515 kg of penalty. No objective curve prints here. The convergence plate (row 16) carries it. No slope numeral prints, because no register row carries one. Two decimals on the power are the recorded precision. Generation 6 prints 5.4 as recorded: two significant digits and not a defect.

Source: data dd3cc6a · script e66dd94 · rows NR-002, NR-005, NR-006, NR-062, NR-066, NR-069

The power trend reads the same search on all three islands. It draws the window-end power of the running-best machine on each island at every generation. Each island carries 31 values, one for each generation (NR-045). A step holds until a better design lands, so the power stays flat between changes. No step interpolates between generations.

The 5 kW floor is a gate and not a target. The score returns infinity below the floor, so no ok design sits below 5.00 kW. Of 508 ok evaluations, exactly one reads 5.00 kW and the next reads 5.02 kW (NR-069, NR-002). That design is the tightest pass of the campaign, one design of 508 (NR-069).

The plate prints no objective curve, no mass curve and no slope numeral. The convergence plate (row 16) carries the objective curve. The trend shows where the search held flat and where it moved.

3.6 Design implications

The record supports four implications. First, the searched TRPT family contains a compact feasible candidate under the modeled objective and gates. Second, the winner holds large structural and twist margins at the final operating point. The margins suggest room for further mass reduction. Any reduction waits on blade integrity and lifter validation.

Third, future searches should treat clearance as a design constraint in its own right. The status mix reads 110 clearance_reject designs (NR-002). Future searches can steer away from low-clearance geometry with a clearance map or a shaped penalty.

Fourth, part of the winner genome stays inactive at one active rotor (NR-011, NR-017). Conditional genes can improve the search efficiency without changing the machine problem. The implications hold under the modeled objective and gates. They do not prove an optimal gene set.

Recommended next steps:

1. Use the rapid path for screening under its fixed rope-break gate (NR-013 ext2).
2. Add blade root factor-of-safety checks to the structural gate set.
3. Validate lifter modulation across the wind bands.
4. Extend the wind cases beyond the steady sheared profile.
5. Extend the model to the launch and recovery sequence.
6. Extract the lever statistics from the campaign telemetry.
7. Re-run a short campaign with conditional genes and a clearance-aware decoder.

4 Discussion

4.1 Evaluation fidelity

The ODE path reproduces every recorded winner field at the precision the record stores. Island 3 agrees at full precision (NR-012). The rapid path once rejected all three winners at the rope-break gate (NR-013). The cause was gate placement, not the machine.

The rapid start injects ring velocities from the static pre-solve. The evaluator ran the rope-break gate across the relax run as well as the window. The injected transient over-strain tripped the gate at once.

The team has solved this pattern twice on record. On 2026-08-14 the rule was: the check judges only real operation. The settle transients must not break healthy machines. On 2026-09-16 the ruling was: breaks on the measurement window only, with break detection off during the relax run. The fix here is the same pattern applied to the evaluator.

The fix has landed. The evaluator arms the rope-break check at the relax boundary (NR-013 ext2). ext2 names the second extension note of the row. Detection runs on the measurement window only.

The fix removes false rejections. It never allows breaks. A rope break in the measurement window always rejects the design. No design passes with a broken rope at its operating point.

The lesson stands. The rapid path is the evaluator default. Any consumer that does not pass the cold flag measures the instrument, not the machine. Every healthy winner reads reject on the default path. Pass the cold flag, and every recorded winner field reproduces (NR-012).

4.2 The failure ledger

Candidate entries. Each failure enters with its lesson.

Date / event	Class	What happened	Detection	Correction / lesson	Status
2026-08-13 k-alignment fault	Gain misalignment	The evaluator ran 5-kilowatt designs at k_mppt 10.0, the 50-kilowatt default gain. That gain is about five times the scaled system's own gain, near 1.94.	Acceptance test	One constant, one source file.	Caught and corrected.
2026-08-14 inverted-taper divergence	Geometry/decoder failure	An invalid genome diverged within seconds. The ground instruments still read normal power.	Divergence review	The decoder must refuse what the builder cannot make. The refusal lands at gate time.	Lesson applied.
2026-08-22 blade-mass-law correction	Model correction	The correction voided prior campaign winners.	Physics review	Void the affected data. Keep the record.	Data voided, record kept.
2026-08-22 FoS=Inf guard gap	Gate near-miss	The guard never reached several scoring paths. The re-gate screens winners.	Telemetry review	A gate must pin a measured number. Drift must fail the suite.	Near-miss, no corrupted survivor found.
2026-08-24 stale-gain re-gate	Stale configuration	The winner re-gate ran at a superseded gain. The check caught it.	Pre-push review	The re-gate at the campaign gain returned 5.18 kilowatts.	Corrected.
2026-08-26 rotor-to-ring off-by-one	Decoder mapping error	The middle rotor landed on the topmost ring. Every expansion rotor shifted one ring.	Visual geometry check (preview_genome_geometry.jl)	Validate the rotor-to-ring map against the built geometry.	Corrected.

- The dead geometry genes. The search proposes genomes that cannot build. An off-by-one in the rotor-to-ring map put the middle rotor on the topmost ring (2026-08-26).

An inverted-taper genome diverged within seconds, while the ground instruments still read normal power (2026-08-14).

Lesson: the decoder must refuse what the builder cannot make, at gate time.

- The voided campaigns. A later finding can break a model or a geometry. The data from before it then counts no more.

The blade-mass-law correction landed on 2026-08-22. It voided the campaign winners. The rotor-count data was doubly void: it predated the FoS fix, and its rotors sat one ring toward the top.

Superseded is the milder class, not void. The pre-derate campaign stands as era record, because the bank derate moved the search space.

Lesson: void the data, keep the record.

- The k-alignment faults. Twice, a check ran at a different generator gain than the machine it scored.

On 2026-08-13 the evaluator ran 5 kW designs at the 50 kW default gain. On 2026-08-24 a winner re-gate ran at a stale gain, the superseded 20-second-window value.

A misaligned gain scores a different machine, and the data voids. Lesson: one constant, one source file.

- The retired load anchors. The anchor 3.55 was a mass-scaled value with no anchor. The anchor 5.39 was the 3-blade anchor chosen on the broken machine. That load forces about 9 kW from the seed (NR-047).

An earlier round swept the generator gain k on earlier machines. At the top of that sweep the TRPT shaft entered a destructive resonance. That round shares no number with this campaign, because the machine and the model differ.

- The stale-gate near-miss. The FoS floor gate fell behind the model it guarded.

The FoS=Inf exploit fix landed in one scoring path only. The mass-minimum objective and the older fitness functions never received the guard. The campaign ran through them unguarded (2026-08-22).

Telemetry shows no surviving design carried the infinite value. That is a near-miss, not a corrupted dataset. Lesson: a gate must pin a measured number, so drift fails the suite.

4.3 Model boundaries

The report states these boundaries plainly. They are part of the valid-science case, not footnotes.

- The model does not check blade integrity. The blade root factor of safety sits at infinity in the model. The team deferred the full beam-on-supports blade model (PRD 0006:147). The rotor blades are not yet structurally gated in this analysis.
- Flown blades carry tether-to-blade bridling that resists thrust and expansion better than the unbridled model assumes. That bridling also adds drag. The model does not include it for the main rotor.
- The model does not simulate launch or recovery. Each run starts with the machine already aloft at its operating point. The full launch sequence and the recovery are outside this report.

The lift system must also carry the machine during hoisting, after a TRPT break and in high wind. The record does not quantify those cases.

- The campaign wind is steady and sheared. Its speed rises with altitude on the site wind profile. The profile comes from wind measured at the Daisy machine. Gusts and turbulence do not occur in a scored run. The wind direction stays fixed. The scored profile is a base case, not a full weather envelope.

5 Conclusion

The report screened 930 rigid-wing ground-generation rotary kite turbine designs with a tensile rotary power transmission (NR-001). The campaign scored each design on the ODE path over the measurement window (NR-015). One fixed power-law generator load and one fixed gate set held for every design (NR-015).

The search found 508 feasible candidates in the searched space (NR-002). The 301 rejects split three ways (NR-055). 96 ran below the 5 kW power floor. Four ran at or above 5 kW with a factor of safety below 2.5. The remaining 201 carried no logged cause at campaign time.

The re-run probe (NR-061) names them. 93 rotors cannot reach 80 per cent of the power floor on their swept area and power coefficient. 37 settle at the realisability cliff.

36 refuse on rotor spacing. 30 break a line at the rope-break gate, not a beam FoS breach or a twist crossing (NR-013). Three cross the Betz ceiling. Two diverge at the hub tip speed.

The campaign winner, the best feasible candidate under the modeled objective and gates, reads 5.309 kW of mean power (NR-008). All checked TRPT rings and twist segments operate below their failure boundaries at the final operating point (NR-021, NR-022).

The result shows that a compact TRPT machine can satisfy the power, clearance, twist and structural gates in simulation. That result holds under the modeled load and wind. It does not show flight readiness, a global optimum, or full structural safety. Blade integrity is not yet gated, and lifter modulation remains an assumption. The rapid path takes the rope-break gate placement fix (NR-013 ext2).

The Daisy anchor ties the model to a measured flight reference. The next work closes the model gaps. The recommended next steps sit in section 3.6.

Where this goes.

The instrument sizes and scores kite-turbine designs. Its physics model follows the published TRPT lineage, and the cross-check names the matches and the mismatches. Each dataset carries its era, its commits and its geometry fingerprint.

The register, the test suite and the pipeline hold every number. Findings, corrections and retractions stay in the record with their lessons.

The search already moves ten genes, so the same pipeline can probe wider parameter regimes. The expansion rotor model ships with tests, so the pipeline can study multi-rotor and expansion-class machines.

The programme targets larger machines, 7 to 50 kW. Each one is a target, not a result.

Wider parameter regimes, flying higher on longer shafts, and the scale where multi-rotor machines win are programme targets too. The axial-length study runs to 60 m as a plan probe. That figure is not a signed row.

The signed length screen varied length alone on one pre-fold genome. Every screened row read below the power floor. The folded class passed at the base length. The record sits in section 2.3.

The programme takes one validation step at the Daisy scale. The mission target is a displaced-cost metric, set out in phase G of the programme strategy.

6 Potential extension work

The report answers the questions it set. It also opens questions it cannot answer yet. This section sets out three of them. Each one states the question, the reason the current model cannot close it, and the test that would. None of the three carries a scored result.

6.1 Cyclic pitch on the ring rotor

The question is whether cyclic pitch or cyclic banking can act on one blade. It could also act on the whole rotor together. A cyclic input changes the pitch or the bank around the revolution. The Some AWE group studies cyclic pitch on a rotary kite [18].

Today the bridles attach each blade to the network, and that attachment sets the pitch and the bank. Both angles stay fixed all the way round, so the rotor has no cyclic input today. The model has no cyclic input either, so the test starts by adding both.

6.2 Cost of energy with scale

The campaign scores power, mass and load. It does not score cost. The report cannot state the cost of energy at any scale, because no signed row for it exists.

A cost-of-energy number needs two inputs the project does not hold. The first is build cost from field-tested prototypes at each scale. The second is the performance data from this model at that scale.

The enabling step comes before any figure. The value-claims ledger classes every economics reading the project holds today as either held or unsigned. Its rule of use governs this section as it governs section 5. No value prints here until a fresh read produces a signed row for it.

The report does not know the cost of energy at scale, and it says so.

The test: one method, applied at each scale. Each printed point needs a signed row.

6.3 Failure modes and effects

The model runs a healthy machine. The failure ledger records campaign failures: a run that stops, a row that reads below the floor. It does not record how the system fails. It does not record what each failure does to the machine, the ground equipment or the site.

The prior art exists. The Daisy machine carries an FMEA and a fault tree to 5 kW, held in `DaisyKiteTurbineControl`. That work covers the rotor, the control set and the ground equipment on that machine. It does not cover this one.

The extension lifts the Daisy method onto the full system: the airborne rotor, the TRPT shaft and the ground station. Each failure mode takes a severity score, an occurrence score and a detection score. Each mitigation points at a test.

The test: a fault tree. Every top event traces to a component, a mitigation and a test that exercises the mitigation. A score without a test does not enter the record.

7 Nomenclature [PENDING. Track D ratifies the glossary.]

8 References

The references below list the sources behind this report. Each entry prints the DOI where the source carries one. A bracketed number in the body points at its entry here. Entries 1 to 10 sit in the order the body first names them.

Eighteen sources follow. Seventeen carry verified sources. The Hancock deck cites as unpublished. The founder dropped the Ranneberg entry.

1. Loyd, M. L. "Crosswind kite power". *Journal of Energy* 4(3), 106-111, 1980. DOI 10.2514/3.48021.
2. Diehl, M. "Airborne Wind Energy: Basic Concepts and Physical Foundations". In Ahrens, U., Diehl, M. and Schmehl, R. (eds), *Airborne Wind Energy*. Springer, Berlin, Heidelberg, 3-22, 2013. DOI 10.1007/978-3-642-39965-7_1.
3. Trevisi, F., Gaunaa, M. and McWilliam, M. "Unified engineering models for the performance and cost of Ground-Gen and Fly-Gen crosswind Airborne Wind Energy Systems". *Renewable Energy* 162, 893-907, 2020. DOI 10.1016/j.renene.2020.07.129.
4. Cherubini, A., Papini, A., Vertechy, R. and Fontana, M. "Airborne Wind Energy Systems: A review of the technologies". *Renewable and Sustainable Energy Reviews* 51, 1461-1476, 2015. DOI 10.1016/j.rser.2015.07.053.

5. van der Burg, S., Jurg, M. F. M., Tadema, F. M., Kamp, L. M. and van de Kaa, G. "Dominant Designs for Wings of Airborne Wind Energy Systems". *Energies* 15(19), 7291, 2022. DOI 10.3390/en15197291.
6. van de Kaa, G. and Kamp, L. "Exploring design dominance in early stages of the dominance process: The case of airborne wind energy". *Journal of Cleaner Production* 321, 128918, 2021. DOI 10.1016/j.jclepro.2021.128918.
7. Pereira, A. F. C. and Sousa, J. M. M. "A Review on Crosswind Airborne Wind Energy Systems: Key Factors for a Design Choice". *Energies* 16(1), 351, 2023. DOI 10.3390/en16010351.
8. Flather, J. J. *Rope-Driving: A Treatise on the Transmission of Power by Means of Fibrous Ropes*. John Wiley and Sons, New York, 1895. <https://archive.org/details/ropedrivingtreat00flatrich>
9. Benhaïem, P. and Schmehl, R. "Airborne Wind Energy Conversion Using a Rotating Reel System". In Schmehl, R. (ed.), *Airborne Wind Energy: Advances in Technology Development and Research*. Springer, Singapore, 539-577, 2018. DOI 10.1007/978-981-10-1947-0_22.
10. Tulloch, O., Yue, H., Kazemi Amiri, A. M. and Read, R. "A Tensile Rotary Airborne Wind Energy System: Modelling, Analysis and Improved Design". *Energies* 16(6), 2610, 2023. DOI 10.3390/en16062610.
11. Read, R. "Kite Networks for Harvesting Wind Energy". In Schmehl, R. (ed.), *Airborne Wind Energy: Advances in Technology Development and Research*. Springer, Singapore, 515-537, 2018. DOI 10.1007/978-981-10-1947-0_21.
12. Joshi, R., von Terzi, D. and Schmehl, R. "System design and scaling trends in airborne wind energy demonstrated for a ground-generation concept". *Wind Energy Science* 10(4), 695-718, 2025. DOI 10.5194/wes-10-695-2025.
13. Aull, M., Stough, A. and Cohen, K. "Design Optimization and Sizing for Fly-Gen Airborne Wind Energy Systems". *Automation* 1(1), 1-16, 2020. DOI 10.3390/automation1010001.
14. Joshi, R., Kruijff, M. and Schmehl, R. "Value-Driven System Design of Utility-Scale Airborne Wind Energy". *Energies* 16(4), 2075, 2023. DOI 10.3390/en16042075.
15. Weber, J. and eleven others. "Challenges and Opportunities for Airborne Wind Energy in the United States". Report to Congress, National Renewable Energy Laboratory, Golden, CO, November 2021. <https://www.energy.gov/sites/default/files/2021-12/report-to-congress-challenges-opportunities-airborne-wind-energy-united-states.pdf>
16. Jamieson, P. and Branney, M. "Multi-rotors: a solution to 20 MW and beyond?". *Energy Procedia* 24, 52-59, 2012. DOI 10.1016/j.egypro.2012.06.086.
17. Hancock, M. "Multi-Rotors Seen Through a Blade Lens". Presentation at the 3rd International Seminar on Multi Rotor Wind Technology, Glasgow, 13 June 2024. Unpublished.
18. Unterweger, D. "Cyclic Pitch Control of a Rotary Kite". Master thesis presentation, 31 January 2022. <https://someawe.org/?p=196>

9 Data availability and reproducibility

The model is `KiteTurbineDynamics.jl`, a Julia package under the MIT licence. One tree holds the solver, the campaign runner, the evaluator and the test suite. The repository sits at <https://github.com/rodWindswept/KiteTurbineDynamics.jl>. Two further interfaces drive the same model. A desktop dashboard built on Makie runs the simulator by hand and plots the response. A Pluto notebook sets a design from sliders and draws it in three dimensions.

The dataset commits to the repository at `dd3cc6a` on `bankderate2-results`. The launch commit is `a76c5e9`. The Phase 0 audit sits at master `31ca00d`. The fidelity probes sit on the `bank-derate-cos2p65` branch. The campaign runner script, `scripts/run_v13_5kw_masslift.jl`, pins the island seeds.

Figure 12 names no dataset extract. It prints the tools of this repository at one commit, and the kit holds a hash for each frame.





Figure (roster row 31) · instrument-gallery

Three panels, each a whole frame from the tool it names, drawn in one pass at one commit. The panels show instruments that this repository tracks. They show what a designer meets when setting a run and inspecting a design: presence, arrangement and state, not results. No register row signs a numeral in it. Where a panel shows a quantity, that number belongs to the tool that prints it at that commit. No other value in this report comes from it.
Source: one render pass, repository at 87e6b10 · script 20365cb

Every figure will ship with its generating script and its data commit. Every public post will cite register rows by ID.

Document control

The version history, the working notes, the open requests to the analysis room and the seat notes sit in docs/reporting/draft-tracker.md. This document carries only reader-facing material.

Appendix: Claims and evidence

The table ties each headline claim of this report to its evidence and its signed register rows.

Claim	Evidence	Report location	Register rows
Campaign scale and status mix	930 designs. The status mix reads 508 ok, 301 reject, 110 clearance_reject and 11 reject_twist.	\$3.1	NR-001, NR-002
Winner power	Mean power 5.309 kW over the measurement window.	\$3.3	NR-008
Winner mass	Airborne mass 25.698 kg.	\$3.3	NR-018
Structural margin	Window-minimum factor of safety 14.972. The checked ring utilisations stay below the failure boundary.	\$3.3, \$3.4	NR-009, NR-021
Twist and torque margin	All five segments stay inside their over-twist limits. The least segment torque margin is 2.03 at S1.	\$3.4	NR-022, NR-040
Search improvement	The search improved the best objective from the common seed to the winner. It reads 36.109 to 27.364 kg-equivalent.	\$3.2	NR-003 to NR-006, NR-036
Instrument fidelity	The cold path reproduces the recorded winner fields. The rapid path once rejected all three winners at the rope-break gate. The placement fix has landed.	\$4.1	NR-012, NR-013 ext2
Model boundaries	Blade integrity is not gated. The report assumes lifter modulation. Launch and recovery sit outside this report. Gusts and turbulence are not modelled.	\$4.3	PRD 0006:147, assumptions

Appendix: Figure provenance

One row per built exhibit. Nine read dataset commit dd3cc6a and the current physics era, post-95c385d_wind-authority, so the era sits here once instead of in a table column.

The tenth names no extract of that dataset. It prints the tools of this repository, and its row names a commit in place of an extract.

Exhibit	Source extract	Cited rows	Script commit
Ring utilisation plate	pair-extract-2026-10-08-winner-ops	NR-021, NR-015, NR-060	e3d664c
Twist limit plate	pair-extract-2026-10-08-winner-ops	NR-022, NR-023, NR-040, NR-052, NR-053	e3d664c
Machine evolution strip	strip-extract-2026-10-09-best-so-far	NR-005, NR-006, NR-011, NR-036, NR-062, NR-066	ce9c5d8
System anatomy plate	scratch/fig1_decode_winner.jl · best_vector.csv	NR-011, NR-017, NR-043	9241b23
Levers landscape	scripts/results/v13_5kw_masslift_len18.8_rotorcount_bankderate2/island_{1,2,3}/telemetry.csv	NR-001, NR-005, NR-006, NR-044, NR-045	d0e410d
Parameter atlas	scripts/results/v13_5kw_masslift_len18.8_rotorcount_bankderate2/island_{1,2,3}/telemetry.csv	NR-001, NR-005, NR-006, NR-044	d0e410d
Machine winner	strip-extract-2026-10-09-best-so-far · pair-extract-2026-10-08-winner-ops	NR-011, NR-017, NR-021, NR-022, NR-043	d0e410d
The corner machines	corner-grid-extract-2026-10-15-picks	NR-044, NR-051	d0e410d
The power trend	scripts/results/v13_5kw_masslift_len18.8_rotorcount_bankderate2/island_{1,2,3}/telemetry.csv	NR-002, NR-005, NR-006, NR-062, NR-066, NR-069	e66dd94
The instrument gallery	instrument-gallery/captures/, three frames hashed in stamp.json	none. No value in this report is read from a panel	87e6b10

Output pins.

Each figure kit under docs/reporting/figures/<slug>/ carries the canonical pins for its exhibit.

The kit manifest.json holds the rendered-output pins (png, svg, pdf, ledger). The check-log pin reads as the sha256 of the kit checks.log.

This appendix cites the kits once instead of reprinting the 64-character hashes.

Figure 12 admission.

Figure 12 prints only the frames that pass four conditions. A panel must add an instrument that the set lacks. It must fit an 85 × 59.89 mm box and measure at least 79 mm wide. It must come from a tracked generator at one commit. Each panel carries a receipt that hashes the frame and the renderer.

One of the three frames needs a flag to print what the plate shows. The cartography prints under KTD_WINNER_TITLE=geometry, and the renderer default is full.

The commit and the three receipts fix the set. The flag fixes one invocation. A reader who holds the commit, the generators and the three input hashes reproduces two frames without a flag and the cartography only with it.

Data register (cited rows)

Every number in this draft traces to a signed row in the numbers register. A missing row here is a release blocker.

Row	Quantity	Value	Units	Source	S
NR-001	Campaign size	930 (3 islands × 310)	evals	island_*/telemetry.csv	S (v2C)
NR-002	Status mix	508 ok / 301 reject / 110 clearance_reject / 11 reject_twist	designs	telemetry status column	S (v2C)
NR-003	Island 1 best fitness	30.578 (found gen 30)	kg-equivalent	island_1/convergence.csv	S (v2C)
NR-004	Island 2 best fitness	30.637 (found gen 30)	kg-equivalent	island_2/convergence.csv	S (v2C)
NR-005	Island 3 best fitness	27.364 (found gen 26)	kg-equivalent	island_3/convergence.csv	S (v2C)
NR-006	All-island start point	36.109 (gen 1, common seed)	kg-equivalent	convergence.csv	S (v2C)
NR-007	P_mean range, ok rows	5.02-6.32	kW	telemetry, ok only	S (v2C)

Row	Quantity	Value	Units	Source	S (v2C)
NR-008	Winner P_mean (island 3)	5.309 / P_end 5.314	kW	island_3 telemetry / re-eval (\$4, Phase 0 audit)	S (v2C)
NR-009	Winner FoS (island 3)	14.972	—	island_3 telemetry / re-eval (\$4, Phase 0 audit)	S (v2C)
NR-010	Winner T_lift (island 3)	324.1	N	island_3 telemetry / re-eval (\$4, Phase 0 audit)	S (v2C)
NR-011	Winner geometry (island 3)	n_lines 3, rings 6, n_active 1, r_hub 4.580 m, r_bot 0.804 m, bank_bot 22°, blade scales 0.821326/0.20 (quote 0.82/0.20)	—	decode of best_vector.csv	S (v2C)
NR-012	Cold-path reproduction	:cold re-eval matches recorded P_mean, P_end, FoS, fitness, T_lift to 3 dp on all three winners	—	probe_rapid_vs_ode_winners.jl	S (v2C)
NR-013	Rapid-path result on winners	:warm rejects all three winners — rope-break gate (broken=true), transient line over-strain during relax	—	probe_rapid_vs_ode_winners.jl + probe_warm_reject_cause.jl	S (v2C)
NR-014	Winner window swing	P_range 0.029 on 5.309 kW (0.5%)	kW / %	re-eval (\$4, Phase 0 audit)	S (v2C)
NR-015	Operating point	k 2.24, relax 10 s + window 40 s, FoS gate 2.5/2.5, length 18.8 m	—	runner provenance	S (v2C)
NR-016	Winner ω (ODE)	10.170	rad/s	re-eval (\$4, Phase 0 audit)	S (v2C)
NR-017	Winner main-rotor bank (island 3)	7.088° at ring 6. Slot 8 = 22° reaches no rotor at n_active 1	deg	decode probe d219b6c. Threading src/objective_v10.jl:339	S (v2C)
NR-018	Winner airborne mass (island 3)	25.70	kg	audit \$4 + probe 0c4046b	S (v2C)
NR-019	Bank-derate exponent in use (main-rotor disc power)	2.65 (power x cosd(bank)^2.65, the same exponent the elevation factor uses)	—	one authority BEM.projection_factor (src/bem.jl, landed b835ebc, record 779fa55). Application sites src/ring_forces.jl:234-235, src/initialization.jl:1033-1037, src/sim_frame.jl:158-159, 480-481 (94223eb, Rod ruling 2026-10-01). Guard test/test_ring_forces.jl "main-rotor bank derate"	S (v2C)
NR-020	Measured bank retention at the Cp peak	0.854 at 20°. cos^2.65 predicts 0.848 (a match to 0.7 per cent). cos^3 predicts 0.830	—	10-precone-sweep.md sections 3 to 5. The precone rows of 10-precone-sweep.csv at commanded lambda 4.0 (Cp 0.22540 vs 0.26386)	S (v2C)
NR-021	Ring check at the winner operating point (island 3): interaction, shares and FoS per checked ring	R2 0.03324, FoS 30.087. R3 0.03355, FoS 29.807. R4 0.02838, FoS 35.236. R5 3.058e-6, FoS 326974. R6 hub 0.06652, FoS 15.033	—	Extract docs/reporting/figures/pair-extract/extract-winner-operating-point.csv, ring block. Emitted by capture_extended; max_util re-derived via ring_element_analysis; A1 identity checked, residual 0.0 at every ring	S (v2C)
NR-022	Segment twist and over-twist limit at the winner operating point	Δα : S1 33.95°, S2 33.81°, S3 33.68°, S4 33.55°, S5 22.82°. 6_crit: S1 to S4 98.86°, S5 92.64°. has_limit true on all five	deg	Extract, segment block. Raw reads from the stored alpha block; trpt_twist_limit on the rope_forces C1 call pattern (node radii, mean line rest length)	S (v2C)
NR-023	Segment torque pair at the winner operating point	segment torque: S1 396.8, S2 397.2, S3 397.5, S4 397.7, S5 397.2. torque_capacity: S1 805.0, S2 808.7, S3 812.2, S4 815.7, S5 1069.9	N·m	Extract, segment block. Capacity = trpt_torque_capacity_axial at trpt_axial_force of the same state. One instant for both series	S (v2C)
NR-025	Cold-start protocol (campaign setting)	Settle settle_to_operational_state (n_op 150000 steps, about 6 s simulated; ω_rated_max 60.0 rad/s), then kickstart off (kickstart_s 0.0; the legacy kick is -60.0 N·m·s²/rad², about 115 times the MPPT torque), then one combined relax+window run, 10 s + 40 s, rope breaks enabled across both. Winner run: 1724859 steps	s, steps	src/objective_evaluator.jl:774, 840-842, 856-892, 978-989; src/initialization.jl:2184-2190; scripts/run_v13_5kw_masslift.jl:157	S (v2C)

Row	Quantity	Value	Units	Source	S
NR-026	Lifter sizing margin (campaign regime)	margin 1.5 (vertical force at the lift bearing = $1.5 \times m_{\text{airborne}} \times g$); line tension = $F_{\text{vert}} / \sin(70^\circ)$, flat at all wind speeds; m_{airborne} excludes the lifter's own 5 kg (the kite carries itself)	—	src/lift_kite.jl:197-228; regime line scripts/run_v13_5kw_masslift.jl:111	S (v 2 C
NR-027	10-05 length screen on the pre-fold island-3 genome (both rounds, arrow = round 1 to round 2)	P_mean 18.8 m: 3.7955 → 4.2332; 21.2 m: 3.8645 → 4.3055; 25.0 m: 3.9747 → 4.4122. All six rows FAIL (P<5.0)	kW	scripts/results/length_variants_island3/variants.csv (f225e95, 2026-10-05T10:10Z); scripts/results/length_variants_island3-cpfix/variants.csv (a82cafd, 2026-10-05T11:25Z)	S (v 2 1
NR-028	18.8 m fold comparison at the cpfix era	pre-fold genome 4.2332 (fail), folded class P_end 5.3691 (ok)	kW	Pre-fold: NR-027 round 2. Folded: scripts/compute_seeds.jl:50-56; docs/validation/2026-10-05-s2-fold-flight.md	S (v 2 1
NR-029	Winner operating speed, window-mean ground ring	13.38 (quote 13.38, 127.772 rpm)	rad/s	scripts/probe_window_trace.jl (b665e6e), parity against the recorded cold re-eval (NR-012)	S (v 2 1
NR-030	Winner delivered torque, window-mean basis	396.8 (quote 396.76)	N·m	P_mean (NR-008) divided by NR-029. Cross-reads: S1 panel 396.8 (0.0075 % above), S5 hub 397.2	S (v 2 1
NR-031	Winner operating coefficients: tip-speed ratio, rotor Cp, system coefficient	lambda_gnd 5.566 (quote 5.566, hub value 5.566 agrees inside 1.1e-4), Cp 0.2944 (quote 0.2944), system coeff 0.1016 (quote 0.1016)	-	lambda from NR-029, r_ring 4.580 m, v_ref 11.01 m/s. Cp at lambda 5.566 from the measured n = 3 surface (src/aerodynamics.jl:263, anchor a82cafd). Coeff from NR-008 divided by 0.5 rho A v^3, A 63.9 m², rho 1.225	S (v 2 1
NR-032	Winner tip-speed ratio and rotor Cp, tip reference	lambda_tip 6.842 (quote 6.8415, hub value 6.841 agrees inside 5.5e-5), Cp 0.2258 (quote 0.2258)	-	lambda from the NR-029 window trace, r_out 5.630 m, v_ref 11.01 m/s. Cp at the hub-basis lambda from the measured n = 3 surface (src/aerodynamics.jl:263, anchor a82cafd). Receipts: scratch/sv_probe_winner_areas.jl, scratch/av_probe_tipref_counter.jl	S (v 2 1
NR-033	Winner swept-area set	A_raw 45.97 (quote 45.9733), A_bank 45.60, A_zn 39.49 (quote 39.4912), A_zy 39.81 (elevation-only cross-check), A_hub_bot 63.87 (retired basis)	m²	Built radii through annulus_area (src/bem.jl). A_zn is the ruled wind-normal basis (objective_evaluator.jl). Receipts: scratch/sv_probe_winner_areas.jl, scratch/av_probe_tipref_counter.jl	S (v 2 1
NR-034	Winner Betz ceiling and re-based coefficients	ceiling (16/27) x 0.5 rho A_zn v^3 = 19.13 kW (quote 19.1304), delivered share 0.2775 (quote 27.75 %), system coeff 0.1413 on A_raw and 0.1644 on A_zn	kW / - / -	v_ref 11.01 m/s, rho 1.225. The code gate charges the rounded 0.593 form: 19.14 kW. Arithmetic route: scratch/sv_rebase_arithmetic.py	S (v 2 1
NR-036	Island 3 running-best search objective at the seven strip milestones	36.109 (gen 0), 32.293 (gen 6), 31.661 (gen 10), 30.929 (gen 18), 29.854 (gen 22), 28.315 (gen 25), 27.364 (gen 26)	kg-equivalent	docs/reporting/figures/machine-evolution-strip/extract-best-so-far.csv (certified 2026-10-09). This seat re-checked all seven against island_3/telemetry.csv: each is an ok row and the least ok fitness of its window. Figure label: search objective	S (v 2 1
NR-039	Island-3 fitness decomposition, NR-005 against NR-018	fitness 27.36 = mass 25.70 + over-power 0.4933 + utilisation 0.5577 + twist 0.6141. Penalty total 1.665 (quote 1.665)	kg-equivalent	island_3/convergence.csv gen-26 row + scratch/sv_p0_winner_reeval.txt + evaluator chain at launch a76c5e9	S (v 2 1
NR-040	Winner segment torque margin	S1 2.029, S2 2.036, S3 2.044, S4 2.051, S5 2.694. S1 governs (quote 2.03). The S1 to S4 band reads 2.029 to 2.051.	—	NR-023 strings, segment by segment: torque_capacity / segment_torque. Recipe: 805.0 / 396.8 = 2.029. Cross-read: docs/reporting/figures/pair-extract/extract-winner-operating-point.csv, segment block, bit for bit	S (v 2 1
NR-041	Honest-window k sweep of the 5 kW seed (2026-08-22): window-mean power by load	P_mean at 2 dp: 3.49 reads reject. Then 5.45, 6.78, 7.56, 8.00, 8.89. Then 8.95 at the last four.	kW	scripts/results/k_sweep_daisy_5kw.csv, blob d17d1377b191dde5b44df7f5534d0cdc4c48d0f3, landed adea1a3 (2026-08-22 14:11:49 +0100). Column line and ten data rows	S (v 2 1
NR-042	Winner blade-tip clearance and the gate floor	floor 1.5, winner 5.468 at decode. Margin +3.968 (3.65 times, +264.5 per cent)	m	pair-extract log line 11 (docs/reporting/figures/pair-extract/extract-winner-operating-point.log, sha256 7579e84a5724a2a41f9ada08200b0dc59e0a03177ce9193c4978abc1d71474f8). Runner scripts/run_v13_5kw_masslift.jl:76 and :254-257. Knob probe scratch/av_probe_clearance_knob.jl	S (v 2 1
NR-043	Winner blade span (island 3, main rotor)	1.499 (quote 1.499)	m	scratch/sv_probe_winner_areas.log, the span (built) line, sha256 in the manifest. R_main 5.630 minus blade_hub_radius 4.130. The decode convention src/objective_v10.jl:374-376. Figures decode log scratch/fig1_decode_winner.log. Counter-read: the aero tip-reference receipt (2026-10-11)	S (v 2 1

Row	Quantity	Value	Units	Source	S
NR-044	PCA projection summary (v13 5 kW screen, 930 draws): variance shares, top loadings, bin rule	PC1 22.434667 (quote 22.4), PC2 17.263044 (quote 17.3), pair total 39.697710 (quote 39.7) per cent. Loadings at four decimals, the abs(v) >= 0.2 set: PC1 n_lines -0.5385, rotor_count -0.5222, blade_scale_top +0.5180, target_Lr +0.2517, r_bottom +0.2036. PC2 bank_bottom +0.5621, blade_scale_bottom -0.5605, density_profile -0.5018, bank_top -0.2444. The field bins at 18 x 18 cells over the 1 to 99 percentile rectangle	per cent, -, cells	Recipe: z-score the 930 x 10 matrix of x1..x10 (three island telemetry files, data commit dd3cc6a), one SVD, shares from the squared singular values. Reproduction below	s (v 2 1
NR-045	Island generation counts (v13 5 kW screen)	31 generations per island (gen index 0 to 30), 10 rows per generation (idx 0 to 9), 310 rows per island, 930 in total	generations / rows	island_*/telemetry.csv, gen and idx columns, data commit dd3cc6a	s (v 2 1
NR-046	The objective seam forms, the weights, the thresholds and the reject placement	f = mass + 5.0 * max(P_score - 5.0, 0)^2 + 20.0 * (2.5 / FoS_min)^2 + 1.0 * (r / (1 - r))^2 with r = min(twist_ratio, 1), plus the stationarity add-on 10.0 * max(0, swing - 0.20) with swing = P_range / P_mean. P_score reads the tail-5 power (P_end). Winner closure: 25.70 + 0.4933 + 0.5577 + 0.6141 = 27.36 (NR-039). Rejects carry no score: the evaluator returns Inf and the runner stamps the search cost 1e9 on every non-ok design (422 of 930 rows)	kg-equivalent: 5.0 per kW², 1.0 per twist-ratio², 20.0 per FoS-ratio², 10.0 per swing fraction. Mass reads kg. The 1e9 stamp is a search cost, not an objective unit	Source read at 7b8c101, the tip at read time. Byte-stable from launch a76c5e9. Sites: src/objective_v12.jl:114-116 (weights), :145-168 (seam body). src/objective_evaluator.jl:90-91, 1148, 1161, 1192-1194, 1200-1203. scripts/run_v13_5kw_masslift.jl:153-156, 256-257, 266, 272, 281-283. Census over island_*/telemetry.csv. Checker: scratch/sv_nr046_objective_seam_check.py and its log	s (v 2 1
NR-047	The k-sweep display rows for section 2.3: six loads of the honest-window sweep, one row per k	P_mean at 2 dp: k 0.5 reads reject. Then 5.45, 6.78, 7.56, 8.00, 8.89. omega_final at 2 dp: 19.14, 17.64, 16.59, 15.80, 15.34, 14.42. FoS_min at 1 dp: 51.8, 47.8, 42.8, 40.1, 36.4, 24.3	kW, rad/s, -	scripts/results/k_sweep_daisy_5kw.csv, blob d17d1377b191dde5b44df7f5534d0cdc4c48d0f3, landed adea1a3 (2026-08-22). Column line and ten data rows. The six rows read at omega_final	s (v 2 1
NR-048 (amended)	The sizing pair, after the twist-limit amend	The load-step case stands at k_mppt x 2 and F_helix = 1278 N/vertex. The twist pair is withdrawn: 87.9° has no resolved object and 78.8° has no derivation at HEAD	N/vertex, degrees	DECISIONS.md [2026-09-06]; docs/plans/2026-09-06-closed-form-beam-sizing.md section 7.1; src/trpt_twist_limit.jl; src/objective_evaluator.jl:448; check scripts/check_nr048_twist_limit.py	a (v 2 2
NR-049	The shaft transition cone angle, one fixed half-angle across the campaign	cone_slope_deg 22.0	degrees	scripts/run_v13_5kw_masslift.jl:172, src/objective_v10.jl:207 and :259, src/objective_evaluator.jl:148. Census: 41 cone_slope_deg sites across src/, scripts/, test/ and the figure scripts, each a literal 22 or a cfg plumb. Probe scratch/sv_fb5_source_reads_probe.jl	s (v 2 1
NR-050	The blade-scale base: the sized rotor radius, the base span at 0.75 of it, the base chord at 0.113 of it	Winner decode: r_rotor 2.434, base span 1.826, span 1.499 (NR-043), chord 0.2259	m	src/objective_v10.jl:355-377, src/bem.jl:156-179, probe scratch/sv_fb5_source_reads_probe.jl and its log	s (v 2 1
NR-051	The corner-grid picks (v13 5 kW screen): four machines, one per PC1-PC2 window corner, at the least feasible distance	The picks stand in the blocks below	kg-equivalent, PC units	scratch/fig_cg_picks.py and its log. The island telemetry files (data commit dd3cc6a, run a76c5e9). Projection per NR-044. Checker scratch/sv_r6_count_nr051_check.jl and its log	s (v 2 1
NR-052	The torsion criteria as implemented, and the 2.5 boundary	Criterion 1: the geometric crossing check is the campaign's twist gate (form in NR-053). Criterion 2: the torsional-stability FoS hard constraint reads 1.5, on the design-verification path. Winner: least segment torque margin 2.03 at S1 (NR-040), 1.35 times the 1.5 floor. Worst crossing ratio 0.4393, margin 2.28	deg, -	Constant src/trpt_optimization.jl:39 (OPT_TORSION_FOS_REQUIRED = 1.5), per segment :872-882, feasibility and message :1100-1123. Crossing detector src/objective_evaluator.jl:454, 913-920, hard reject :1024. Policy boundary per software-validator-2026-10-11-item13-torsion-gate.md. Winner cell: island_3/telemetry.csv, gen 26, idx 9	s (v 2 1
NR-053	The realisability floor and its 1.05 margin	sin(Δα) <= 1/1.05 = 0.9524 (quote 0.9524). The continuum floor angle reads 72.25 deg (quote 72.25). The discrete crossing limit reads Δα* = 2asin(L_seg / sqrt(2(L_seg^2 + 2*r_seg^2))), r_seg = max(r_i, r_{i+1}). Winner cross-reads: S1 twist 33.95 deg, margin 2.13 to the floor and 2.91 to Δ_crit (NR-022)	deg, -	Constant src/initialization.jl:1691 (TRPT_REALISABILITY_TENSION_MARGIN = 1.05), basis :1663-1691. Crossing form src/objective_evaluator.jl:444-462 and src/initialization.jl:1414-1456. Ruling DECISIONS [2026-09-20]	s (v 2 1

Row	Quantity	Value	Units	Source	S
NR-055	The reject class of the status mix, split three ways for print	The plain reject class reads 301. Subgroup A, ran with power below the 5 kW floor: 96. Subgroup B, ran at or above 5 kW with FoS below 2.5: 4. Subgroup C, fields zeroed and the cause not logged: 201	designs	Read of the three campaign telemetry files, scripts/results/v13_5kw_masslift_len18.8_rotorcount_bankderate2/island_{1,2,3}/telemetry.csv. Subgroup rule and counts: docs/reporting/figures/figure-generation-notes.md:135-147. Checker scratch/sv_301_split_check.py and its log	s (v 2 1
NR-056	The shear law and its exponent	Hellmann power law, exponent 1/7 (0.1429). One authority: the measured Daisy pair, 10.0 m/s at 4.8 m. Reads: 11.01 m/s at 9.412 m (the table pair reads 11.01 at 9.41), 11.01 m/s at the 9.400 m design hub, 10.92 m/s at the settled hub 8.914 m	-, m/s, m	src/wind_profile.jl:50,54 (SITE_DAISSY, SITE_ANCHOR), :110-112 (site_shear), :130-137 (wind_at_altitude). Guard test/test_wind_authority.jl. Landing record docs/validation/2026-10-03-d1-wind-landing-verification.md	s (v 2 1
NR-057	The shaft-elevation angle and the hub altitude	30.0 degrees (pi/6), the model convention. Angles do not scale. The design hub reads 18.8 * sin 30 deg = 9.4000 m. The rung reference reads 9.4117 m (gap 0.0117 m, immaterial). The flown settled read: mean elevation 28.64 deg at mean z 8.914 m	deg, m	src/parameters.jl:294,369,606. Campaign runner scripts/run_v13_5kw_masslift.jl:73 (ELEV = pi/6). Register boundary note 5 (docs/reporting/numbers-register.md:618-622)	s (v 2 1
NR-058	The steady-wind basis of the reported trace	The evaluator wind closure reads altitude only, so the reported trace runs steady wind. The steady_wind constructor returns a constant. Ramp, gust and turbulence tools exist and none is engaged. The NR-014 window swing (0.029 kW on 5.309 kW, 0.5 per cent) is a steady-wind response, not a stability property. The register reads 0.54 per cent on the same quotient	-, kW	src/objective_evaluator.jl:792-795. src/wind_profile.jl:160-164. Harris review docs/reporting/feedback/2026-10-10-harris-review.md:184. Register :826	s (v 2 1
NR-059	Winner bow (island 3)	midpoint ring settle 0.002 m, window 0.105 to 0.310 m. Hub max settle 0.002 m, window 0.209 to 0.613 m. Extension end midpoint 0.422 m, hub 0.833 m. Settle and run do not agree on the bowed shape	m	probe scratch/av_bow_probe.jl (sha256 24a1601718dcb2d8536872ca975b6a823fed640e72686486c434fdc6216c73a4), its log scratch/av_bow_probe.log (sha256 eacc0ce3acbcdeb66b6dd9db975b177ceb401a2929bee2c40fd258e22f2bf987), both first landed at c6d2885. Re-run scratch/sv_bow_probe_recheck.log. Checker scratch/sv_bow_sign_check.py and its log	s (v 2 1
NR-060	R5 signed axial force at the winner capture	-157.6 to -142.8 (tension), three beams, N/Ncrit -0.09240 to -0.08371	N	Own re-run at the science seat, scratch/sv_r5_axial_sign/sv_r5_axial_sign.jl, same evaluator chain as scripts/extract_pair_capture.jl, parity bit-exact on nine recorded cold re-eval fields	s (v 2 1
NR-061	Six-class split of the 201 FoS=Inf rejects, classified with the producing code	betz_cp_floor 93, settle_fail 37, spacing 36, rope_break 30, betz_ceiling 3, hub_tip 2 (total 201)	count	Own sharded re-derivation at a76c5e9 (scratch/sv_reject_sixclass_a76c5e9/sv_reject_sixclass_a76c5e9.jl, eight shards over the 201 recorded rejects). Merge scratch/sv_reject_sixclass_a76c5e9/sv_reject_sixclass_merge.py, 317 checks, 0 fail	s (v 2 1
NR-062	Window power at the seven strip milestones (the NR-036 steps)	The seven pairs sit in the block below	kW	docs/reporting/figures/machine-evolution-strip/extract-best-so-far.csv (certified 2026-10-09). Parity against island_3/telemetry.csv bit for bit on all seven rows	s (v 2 1
NR-066 (amended)	The decoded mass of each machine-evolution strip panel on the wind the campaign scored, and the two label forms	The table below	kg, kg-eq	docs/reporting/figures/machine-evolution-strip/extract-mass.csv at ce9c5d8. Decoding commit 3307af6. Generator scripts/strip_mass_decode.jl. Input extract-best-so-far.csv at dd3cc6a	a (v 2 2
NR-069	Window-end power bound over the ok evaluations, three-island campaign	508 ok of 930. 0 below 5.00. Exactly 1 records 5.00. Next lowest 5.02. Highest 6.33	kW, count	The three telemetry files at data commit dd3cc6a: scripts/results/v13_5kw_masslift_len18.8_rotorcount_bankderate2/island_1..3/telemetry.csv, 310 rows each. Read of the status and P_end columns. Recount script scratch/sv_floor_bound_2026-10-20.py	s (v 2 2

Build record (builder output)

This block is builder output. It sits outside the report proper and does not pass the STE gate.

Draft: 2661194 (docs/reporting/drafts/2026-10-08-report-first-draft.md)

Register: v1 ROWS SIGNED (2026-10-08, science-validator). See the sign-off — file at 52d6769

Citation check: 59 cited, every one signed (82 signed rows total), uncited: NR-024, NR-035, NR-037, NR-038, NR-054, NR-063, NR-064, NR-065, NR-067, NR-068, NR-070, NR-071, NR-072, NR-073, NR-074, NR-075, NR-076, NR-077, NR-078, NR-079, NR-080, NR-081, NR-082, draft file alone: 59 cited — checker exit 0

Seats resolved: 10 of 10 — system-anatomy, machine-winner-hero, ring-utilisation, twist-limit, pca-landscape, pca-atlas, machine-corner-grid, machine-evolution-strip, machine-power-trend, instrument-gallery

Embed route: SVG vector-first, PNG fallback; kits embedded as shipped, nothing re-rendered

Toolchain: pandoc 3.1.3 → HTML → google-chrome headless PDF (A4)

Built: 2026-10-21T21:13:19Z