

PERU OFFSHORE SAILING — WIND & ROUTE PLANNER

The Operating Manual

Every screen, every number, every color — explained simply, in the same order as the app's menu, with real examples from Wednesday 22 July 2026.

How to read this manual: each section = one app menu. Every technical term appears in full with its abbreviation in parentheses — for example True Wind Speed (TWS) — and every concept comes with a real, dated example you can check inside the app.

TIRRENO · Beneteau First 42 · Offshore Racing Congress (ORC) rating PER-120 · La Punta, Callao

The manual follows the app menu, top to bottom

1	WIND DASHBOARD	the tile grid · days & periods · colors & alarms · the  live correction
2	CURRENT · NEXT · FOLLOWING	the three-station comparison with compasses and trends
3	LOCATION ANALYSIS	the charts: observed wind against every model line
4	WIND FORECAST EVALUATION	best model · active corrections (MOS) · forecast grid cells · update times
5	ROUTE MODEL	the sail crossover chart · the ORC certificate & polars · Monte Carlo & isochrones
6	ROUTE SELECTION	the route library · the P-races · building your own routes
7	ACTIVE ROUTE	the plan: map · leg table · leg cards · boards · odds · the  optimizer · laylines
8	APPENDIX	how it all works underneath · limits · glossary · pre-start checklist

BEFORE YOU START

Three things that are true on every screen

IT REFRESHES ITSELF

New data arrives every 10 minutes automatically. The header shows the time and the build version (for example BUILD 16jul · v186). The ↺ button forces a refresh and resets every view to its default.

ONE WIND TRUTH

Every screen — tiles, charts, comparisons and Route Selection — reads the SAME corrected forecast. If one view says 15 knots (kt), they all mean the same 15 kt. No screen has a private opinion.

LIVE BEATS PREDICTED

Wherever a real instrument exists (an anemometer — a wind sensor), the app shows its reading with a red LIVE badge. Predictions only step in where no instrument exists — and they say so honestly with a BEST MATCH badge.

📍 WHERE THE REAL MEASUREMENTS COME FROM

the six instruments: **Jorge Chávez (SPJC)** · **DHN Chucuito** · **YCP** · **CRL** · **Paracas NW (Estación TP Paracas)** · **Paracas SE**

SECTION 1 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

What am I looking at, and can I trust it? — the home screen, tile by tile

The screen: pick a day, pick a period, read the map of tiles

DAY SELECTOR

Today · tomorrow · up to +14 days. Past periods of today show OBSERVED wind (what really happened) — marked “obs”. The +7d and +14d days exist because the European model — European Centre for Medium-Range Weather Forecasts (ECMWF) — reaches 15 days ahead.

PERIOD SELECTOR


Three sailing windows: 6 am–12 pm (morning) · 12–6 pm (afternoon) · 6 pm–12 am (evening). The window marked “now” is live.

THE TILE GRID IS A MAP

Two columns, north at the top: islands and offshore spots on the LEFT, coastal spots on the RIGHT. Your home club — Yacht Club Peruano (YCP) — carries a ★. TAP ANY TILE to open its menu: Location Analysis plus the five Forecast Evaluation pages, already set to that spot.

Real example (22 Jul, 8:13 am): tapping “5 Wed (+14d) · 6 am–12 pm” shows every tile’s average for that far morning — all forecast by ECMWF, because the German model — Icosahedral Nonhydrostatic (ICON) — ends near day 7. If the 15-day data is still downloading, the legend says “fetching the 15-day forecast...” and fills in by itself.

Reading one tile — a real one, from last night



the location's name

BEST MATCH = no instrument here; this is the app's best corrected estimate. LIVE (red) = a real sensor.

the model in charge here: ICON+ — the German model, corrected (the "+")

wind now, colored by strength (yellow = 11–16 kt) — the live ⚓ correction (+2.2 kt here) is already inside this number; its readout lives in Location Analysis

strongest gust expected — same Beaufort colors (red from 17 kt)

direction now → where it is turning, and by when

Why this tile matters: on 21 July at 2 pm the old app showed 9 kt here while the real breeze blew about 17. Everything on this tile — the corrected model, the silent ⚓ correction, the colors — exists to make sure that never happens silently again.

Who is forecasting each tile — and who checked

THREE FORECASTS COMPETE

ICON (German, fine 7-nautical-mile grid, ~7 days) versus ECMWF (European, 15-nautical-mile grid, 15 days) versus the Blend — the average of the two after correction, which wins when their errors cancel. The “+” means corrected against local sensors — Section 4 explains how.

EVERY TILE FOLLOWS ITS “VALIDATOR”

A spot with no sensor is judged by the instrument that behaves most like it. Exposed islands (Horadas, San Lorenzo, Palomino...) follow Jorge Chávez airport (SPJC). Sheltered club spots follow their own club sensors. Southern spots follow Paracas NW (Estación TP Paracas, at the bay entrance) or Paracas SE inside the bay.

CROWNS ARE RE-EARNED EVERY NIGHT

The app scores yesterday’s forecasts against what sensors really measured, and the winner takes each tile’s chip. No model is trusted by reputation — only by last night’s grade.

🔪 REAL EXAMPLE

real chips this morning (23 Jul): **exposed spots “Blend · ICON+” · DHN “ICON+ · Blend” · south spots follow the new Paracas NW station**
meaning: **speed and direction are crowned SEPARATELY, per station — the app splits the job, and the crowns move nightly**

The color language — the exact Beaufort scale in app colors

0–3 kt	grey — Beaufort force (F) 0–1 “calm” — no sailing; the tile border BLINKS grey
4–6 kt	light green — F 2 “light breeze”
7–10 kt	green — F 3 “gentle breeze” — the good sailing band
11–16 kt	yellow — F 4 “moderate breeze”; reef thinking starts
17–27 kt	red — F 5–6 “fresh to strong breeze”; the border BLINKS red when the AVERAGE wind reaches 17
28+ kt	magenta — F 7 “near gale” — STORM; BLINKS magenta at 28+ kt AVERAGE wind. F 7 is the realistic maximum off the Peru coast

THE GREY BORDER

Wind from the northwest–north–northeast sector (315°–045°) is unusual and unstable here — the tile gets a grey border so you look twice.

GUSTS COLOR · AVERAGES BLINK

Every tile carries its own gust line, always from the corrected German model (ICON+), the best gust forecaster here. Gusts color their pill; the BLINKING alarms follow the AVERAGE wind.

THE SEA SHOWS THE FORCE

The animated water at each tile’s foot mirrors its Beaufort force, drawn in the band’s own color — flat in calm, white caps from Force 3, steep streaked seas at Force 7. Faster, bigger waves = more wind.

The correction — when a real sensor corrects the forecast live

THE IDEA IN ONE SENTENCE

If the Jorge Chávez sensor reads 5 kt MORE than the model says right now, spots that behave like Jorge Chávez are probably also windier than their forecast — so the app adds a careful share of that difference, live.

WHY ONLY A SHARE (40%)

Tested on 26 archived days: copying the FULL difference made forecasts worse (–16%); copying 40% between look-alike spots made them 15–20% better. The share shrinks with distance: Ancón 35% · Las Hormigas 25% · Islote Pelado 15%.

IT FADES BY THE HOUR

An error measured now is worth 73% in one hour, 46% in two, 29% in three — fractions the app re-learns from its own archive every day. By evening the correction has faded to zero on its own.

REAL EXAMPLE — THE EXACT TILE FROM PAGE 6

21 Jul, 6 pm — Jorge Chávez: **sensor 18.0 kt · corrected model 12.4 kt → difference +5.6 kt**

Islotes Horadas tile: **model 12.7 + 40% × 5.6 = 14.9 ≈ 15 kt ·  readout in Location Analysis**

Planning a week or two ahead

+7 DAYS VIEW

Uses the LONG-range crown (days 6–10) — whichever model held it after last night's scoring, speed and direction crowned separately (as of 23 Jul: ICON+ speed · Blend direction).

+14 DAYS VIEW

Only ECMWF reaches this far, so every tile wears the ECMWF+ chip. Useful for choosing a regatta weekend, not for picking a sail.

HOW FAR TO TRUST IT

The measured error grows with range: about ± 2.4 kt today-tomorrow, ± 3 kt at a week, ± 3.4 kt beyond ten days. The app never hides this — the Forecast Evaluation pages (Section 4, via any tile's menu) show the error for every range.

🔍 TWO HONEST DETAILS

gusts beyond day 7: **ICON ends** → gusts fall back to the speed model's own gust field

blank tiles? **a rate-limited download retries every 90 s and the legend says "fetching..."**

SECTION 2 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

Three stations side by side — the one-minute briefing before you sail

Three stations side by side — your pre-sail briefing

PICK YOUR THREE COLUMNS

The ▼ picker in Current Conditions chooses any three locations (default: San Lorenzo NW · DHN · Jorge Chávez). The same three carry into Next period and Following period.

THE COMPASS TELLS THE FUTURE

The black needle is a masthead wind vane: the dot sits at the direction the wind comes FROM. The red arc outside the rim shows where it will ROTATE to — hollow dot = now, arrowhead = later, with the bearing and the hour at each end.

PEAKS, LOWS AND TRENDS

Each cell shows a small wind curve (sparkline), a ▲ peak pill and a ▼ low pill (colored by strength band), and a trend arrow (↑ rising · → steady · ↓ falling) against right now.

🔪 REAL EXAMPLE

this morning, 8 am: **Jorge Chávez LIVE 8 kt · DHN LIVE 8 kt · San Lorenzo NW 9 kt (model, ⚓ +1.0)**

afternoon read: **all three climbing to 13-15 kt with gusts 17-20 — classic sea-breeze day**

A one-minute briefing, before you leave the dock

1

READ “NOW”

Live sensors first. If Jorge Chávez already reads 15+ kt at noon, the afternoon will be strong — the airport leads the bay by an hour or two.

2

READ “NEXT PERIOD”

The average, peak and direction for the coming window — with the trend arrow telling you if it builds or fades relative to now.

3

READ THE COMPASS ARCS

A veer (clockwise turn, e.g. 150° → 165°) favors starboard-side strategy on the beat; the arc tells you by WHEN.

Everything here obeys the same one-wind-truth: these are the tile numbers, arranged for comparison.

SECTION 3 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

The charts — where observed wind meets every model line, and models get caught

The charts — where you catch a model lying

WHAT THE LINES ARE

Red dots = OBSERVED wind (the sensor's truth, today only). Dark blue = ICON+ · dark green = ECMWF+ (both corrected). Dashed light lines = the same models RAW, before correction. Light blue = the gust forecast.

THE COMPARE CHIPS

Toggle any combination: corrected vs raw of each model, gusts, both models against each other. [All] and [None] for speed. The app pre-selects whichever lines are winning today.

WHAT A GAP MEANS

When the red observed dots ride ABOVE every model line, the models are under-calling today — expect more wind than forecast everywhere nearby, and trust the ⚓ corrections doing their job.

📌 REAL EXAMPLE — THE NIGHT THE MODELS RAN LOW

21 Jul, evening, Jorge Chávez: **observed 18 kt · ICON+ line 13.6 · ECMWF+ 11.9 — every model low**

what the app did: **bent the next hours up by the ⚓ ladder (+4.1 at +1 h, +2.6 at +2 h...) — visibly**

SECTION 4 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

Everything that makes the WIND numbers honest — five pages of evidence

Tap any tile — its menu holds these five pages

WIND FORECAST EVALUATION — the four pages, opened from any tile’s menu

Active corrections (MOS)	Model Output Statistics — the sensor-trained formulas that complete the “+” (first in the menu, right under Location Analysis)
Best model	ONE page: what the app flies at each range and why — crowns, verified evidence, today’s live grades and the yesterday replay
Forecast grid cells	which SEA points of the model grid feed each location — the first step of the “+”
Model update times	when ICON and ECMWF publish fresh runs, in Peru time

ROUTE MODEL — the racing side (Sail crossover chart · ORC certificate & polars · Monte Carlo & isochrones) — is Section 5.

Rule of the house: every number in the app can show you its exam grade in one of these pages.

Who to trust, by how far ahead — one page: crowns, grades, replay

👑 THE CROWNS AS OF 23 JUL (RE-EARNED NIGHTLY — THEY MOVE)

SHORT · today & tomorrow: **Blend** (speed 2.3 kt · direction 21° mean error)

MEDIUM · 2–5 days: **Blend** for both

LONG · 6–10 days: **speed ICON+** · direction **Blend**

VERY LONG · 11–15 days: **speed ECMWF+** · direction **Blend**

👑 THE GRADE, WORKED — MEAN ABSOLUTE ERROR (MAE)

5 verified hours, misses: **+1, -2, 0, +3, -2 kt**

Mean Absolute Error (MAE): **$(1+2+0+3+2) \div 5 = 1.6$ kt**

read it as: **“this model typically misses by 1.6 kt here”**

TWO HOUSE RULES

≥5 verified days to rule. A model once “won” with one lucky day of data. Never again: small samples lie confidently.

Speed crowns are earned on sailing hours. A model that aces midnight calms but misses the 2 pm sea-breeze ramp cannot rule — afternoons and evenings decide.

Blend = the average of ICON+ and ECMWF+ (already corrected — never corrected twice); it wins where their errors cancel. The crown is NOT unanimous: each tile follows its own validator station — e.g. today exposed spots (SPJC) sail on Blend while YCP, CRL and DHN keep ICON+ for speed. In the app the page shows the tables directly; the explanation lives [HERE](#). Colored cell = crowned winner · Δ = short sample · every number is the mean error of hours the model had NOT seen while calibrating. The STATION picker opens on YOUR location's station (the short-range verdict its tiles follow); “Over 2 days · pooled stations” shows the combined verdict that drives all longer-range picks. Every row shows its evidence: mean error, verified hours, verified days — and the runners-up.

Where the “+” begins: ask the model only where it knows water

SEA CELLS ONLY

A model grid point sitting on the coast blows a polluted, half-land wind. The app uses only points flagged as SEA, within 5 nautical miles (nm) of the spot — and shows you exactly which, on a little chart.

TAP TO RE-WIRE

Each used cell (●) and discarded cell (○) is tappable: the navigator can pin or drop cells and the forecast recalculates instantly. Several spots are hand-pinned (for example Islotes Horadas uses its clean South cell — its West cell blows across San Lorenzo island).

DISTANCE DECIDES THE MIX

Closer cells count more — Inverse Distance Weighting (IDW). Direction is mixed as arrows (vectors), never as raw degrees: the average of 350° and 010° must be north, not south!

🔧 INVERSE DISTANCE WEIGHTING (IDW), WORKED


two cells: **A at 4.0 nm reads 12.0 kt** · **B at 8.0 nm reads 9.0 kt**

weights 1/distance: **A: 67%** · **B: 33%** → **$0.67 \times 12.0 + 0.33 \times 9.0 = 11.0$ kt at the spot**

Real cell surgery — YCP and Islotes Horadas, as configured today

YCP — PINNED TO ONE CLEAN CELL

○ deactivated: wind crosses San Lorenzo



The navigator PINNED the north-east sea cell: the forecast never drifts to a different blend, and the western candidates — whose wind blows across the island — stay off.

ISLOTES HORADAS — SOUTH CELL ONLY

○ West cell — wind crosses San Lorenzo

○ East cell · 10.2 nm too far — deactivated



Same discipline offshore: one verified clean SOUTH cell carries the forecast; the far and contaminated candidates are switched off on purpose.

How you do it: tap the location's tile → Forecast grid cells → tap ● to deactivate, ○ to activate, “+ cell...” to add, “⌂ default” to reset — the forecast recalculates instantly and your choice is saved on the device.

Model Output Statistics (MOS) — the “+” in ICON+ / ECMWF+

THE PROBLEM

Every model has habits at every place: “afternoons it runs 2 kt low at YCP”. Habits are predictable — so they can be corrected.

THE FIX — A SIMPLE FORMULA

corrected = $\alpha + \beta \times \text{model}$. Alpha (α) = the fixed offset, learned per morning/afternoon/evening. Beta (β) = the gain: how much reality moves when the model moves 1 kt. Direction gets the same treatment per 45° sector, capped at $\pm 25^\circ$.

SAFETY CAP

The correction can never exceed ± 6 kt — one crazy day is not allowed to bend your race plan.

🔪 REAL EXAMPLE — TODAY’S ACTUAL YCP NUMBERS

YCP · ICON+ (23 Jul live table): $\beta = 0.71$ · $\alpha(\text{morning}) = +3.2$ kt

model says 6 kt → $0.71 \times 6 + 3.2 = 7.4$ kt

model says 12 kt → $0.71 \times 12 + 3.2 = 11.7$ kt ← NOT a flat shift — that is β at work

The nightly exam — never graded on homework it has seen

EVERY NIGHT

Today's forecasts are archived. Yesterday's archived forecasts are graded against what the six sensors really measured. 1,200+ grades so far, growing every night.

OUT-OF-SAMPLE — THE GOLDEN WORD

A forecast is never graded on hours it was tuned on — like never letting a student grade themselves on the exercises they memorized. Every percentage in this app is an unseen-exam grade.

RECENT DAYS COUNT MORE

Each day of age multiplies a record's weight by 0.85 (a one-week “half-life”) — so the corrections follow the season instead of mixing winter with summer.

This is the app's character: it does not ask you to believe it. It shows you last night's grade, every day, for every number.

The same page, below the crowns — one location at a time

WHAT THE % MEANS

Closeness: 100% = perfect. In speed, 71% $\approx$ missing by $\sim 29\%$ of the typical wind (about 3 kt on a 10 kt day). In direction, 86% $\approx$ a 25° mean miss. Below 50%: do not lean on that model today.

INSTRUMENTS FIRST, THEN OPEN SEA

Six instrument stations are graded against real sensors. The fifteen open-sea spots have no sensor, so they are graded against the freshest model analysis — clearly labeled, honestly weaker evidence.

PICKERS + YESTERDAY REPLAY

Two pickers — TYPE (met stations / open-sea spots), then LOCATION — open on the location you tapped. Below the grades, the YESTERDAY REPLAY shows the previous day hour by hour, noon–6 pm: what each of the five models said that morning against what the instrument measured — the day's winner highlighted. One afternoon is one sample; the crowns weigh every verified day.

📌 REAL EXAMPLE

real morning grades: **YCP: ICON+ 70% speed · 85% direction (best of five competitors)**

an open-sea row: **Isla Grande: ECMWF+ 97% — graded against the best-match analysis**

How fresh is the data under your decision?

ICON (GERMAN MODEL)

A new run every 3 hours, arriving ~3–4 hours after its start time. Fine grid, best short-range gusts — your today-and-tomorrow workhorse.

ECMWF (EUROPEAN MODEL)

A new run every 6 hours, slightly later to publish. Coarser but steady — and the only model reaching 15 days.

WHY YOU CARE

A plan made at 11 am may still ride the 6 am run. Before a start, hit ↻ refresh: if a fresh run just landed, targets and odds can shift a few percent.

📍 OBSERVATION FRESHNESS

sensors: **METAR** — Meteorological Aerodrome Report — every 30–60 min · club sensors ~10 min

SECTION 5 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

How the race plan itself is computed — sails, certificate, odds and the optimizer

Three pages: the sail chart, the certificate, the engines

Sail crossover chart

which sail at every True Wind Angle (TWA) × True Wind Speed (TWS) — your Quantum chart, 269/269 cells validated

ORC certificate & polars

TIRRENO PER-120: the rated speed table, beat & gybe angles — and the 100% 🎯 TARGET benchmark

Monte Carlo & isochrones

the odds engine and the manoeuvre optimizer — who decides what, in which order

These three pages ARE the route model: what the boat can do (certificate), what it should wear (chart), and how the plan handles an uncertain future (engines).

The three engines behind every race plan — in one card

①



TARGET

the benchmark: the course at 100% of the ORC polar — what you must beat

②



MONTE CARLO (MC)

the odds: 10,000 possible futures → arrival percentiles and the SIDE of each beat

③




ISOCHRONE OPTIMIZER

the polish: refines every tack and sail change through the evolving wind, on the chosen side

Why this order (tested on 47 archived race-days): picking the beat's side with the odds engine and THEN optimizing the manoeuvres beat the reverse order 6–4 where they disagreed — with 18–55 minutes riding on the average call. Full detail lives in this panel, including the worked P2 example.

The ship's law: which sail, decided by data



True Wind Angle (TWA): 35° close-hauled → 170° dead run

Your real Quantum chart, digitized. For every combination of True Wind Speed (TWS) and True Wind Angle (TWA) it names the sail — restricted to what is aboard: J2 jib · A3 · A2 spinnakers. All 269 chart cells validated against your spreadsheet.

The boundaries move with wind speed: at 11 kt the A3 takes over from TWA 101°; at 22 kt not until ~135°. Every simulated minute — target, odds, optimizer — wears the correct sail, and a peel costs 120 seconds.

REAL EXAMPLE

a real card note: "A3 not on: at 11 kt A3 starts at TWA $\geq 101^\circ$ (82° here)"

18 Jul lesson: a sharp crew CAN beat the chart — races will teach it

The boat's fingerprint — and the benchmark that never moves

THE CERTIFICATE

TIRRENO races under Offshore Racing Congress (ORC) certificate PER-120. Its Velocity Prediction Program (VPP) computes the boat's rated speed at every True Wind Angle (TWA) and True Wind Speed (TWS) — the polar table, auto-ingested by the app.

THE TARGET = 100% OF IT

Every plan's target is this table at exactly 100% — the rating benchmark you must beat, every time. It never changes with mood, mode or model.

WHAT READS IT

The kt column of every leg (polar speed at the sailed angle) · the beat and gybe angles for the two-board plans · the optimizer's speed at every simulated step · the odds engine's boat.

REAL NUMBERS FROM THE PANEL

sample rated speeds: **TWA 45°: 4.8 kt @6 · 6.1 @10 · 6.6 @14** | **TWA 120°: 5.4 @6 · 7.2 @10 · 8.3 @14**

one leg · 4.2 nm @ 6.3 kt: **TARGET = $4.2 \div 6.3 \times 60 = 40$ min — the number to beat**

SECTION 6 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

Your races, ready to plan — and how to build new ones

The route library — your races, ready to plan

THE LIST IS A CHART TABLE

Routes are grouped North (Ancón–Huacho) · Central (Callao–Lima) · South (Paracas), each with distance and a rough sailing time. Regatta routes carry their P-number on the right.

THE P-RACES

P1 Camotal–Hormigas–Isla Grande–YCP (71.9 nm) · P2 Camotal–Horadas–Pan de Azúcar–Camotal–YCP (26.6) · P4 Paracas–Chincha–San Marino (47.7) · P5 Paracas–Isla Blanca–Ballestas–Tres Marías (~31) · P6 Paracas–Ballestas–San Gallán (~31). All drawn from your own Navionics tracks, with the real rounding rules built in.

MAKE IT YOURS

“Solo P” filters to your races. Swipe a route: ← delete · → copy as custom. Build any route waypoint by waypoint and ★ save it to the library. Tap a route → the app plans it immediately.

🔪 REAL EXAMPLE

defaults: **P-races only** · next start at 12:00 noon (today, or tomorrow if past noon)

SECTION 7 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

The plan itself: map, legs, boards, odds, and the ⚡ optimizer

The command bar at the top of every plan

THE GREEN P PILL IS A MENU

Tap it to jump between all P-races — the plan reloads and always lands on the full race summary first (never buried in a sub-screen).

🏁 RACE — THE PACE TOGGLE

ON (default): odds and expected times use your measured 113% pace. OFF (cruise): planning at 97% — the honest number for “when do we tie up”, never for racing. The 🎯 TARGET stays at 100% of the rating either way.

DAY · TIME · EDIT · REVERSE

Pick any start day/hour (defaults to the next noon). EDIT opens the waypoint builder; REVERSE flips the course — and any change replans automatically. You never press “calculate”.

🔪 REAL EXAMPLE

the red line below: “← go back to Route Selection” — always one tap from the library

The chart — everything the boat must respect, drawn

THE MAGENTA LINE IS THE PLAN

Numbered waypoints (WPT) in circles; “sail this” magenta versus the thin grey straight-line reference. Focused legs zoom laser-tight with every time label visible.

GATES AND CORRIDORS

Flag markers are certified water gates ( S Camotal,  Costa Verde, the San Gallán ring...). Where geometry is dangerous the plan follows these verified corridors — never a straight line through trouble.

RED MEANS NEVER

Red dashed areas are charted no-go zones: El Camotal bank (2 m of water), the San Lorenzo–Frontón gap, Pan de Azúcar’s rocky apron — digitized from your own Navionics charts. The planner cannot cross them; neither should you.

REAL EXAMPLE

real P2 route today: **Camotal 2 (W)** →  **S Camotal** →  **W Horadas** → **Islotes Horadas (port!)** → **Pan de Azúcar**

the return: →  **Costa Verde** →  **S Camotal** → **Camotal 2 (W)** → **YCP** — exactly your GPX line

The leg table — the real P2 plan, 22 July, noon start

1-2	Camotal 2 (W) → S Camotal	0.7	⚡ 0h11	6.3	12:11 pm
2-3	S Camotal → W Horadas	4.2	0h45	6.3	12:56 pm
4-5	Islotes Horadas → Pan de Azúcar	6.8	1h08	6.4	2:15 pm
5-6	Pan de Azúcar → 🚩 Costa Verde	4.3	⚡ 0h45	6.8	3:00 pm
6-7	🚩 Costa Verde → S Camotal	8.0	1h08	7.9	4:08 pm
1-9	TOTAL · full route	26.6	⚡ 4h42	6.9	4:42 pm

nm — nautical miles of the leg (rhumb line).

target — the leg at 100% of the Offshore Racing Congress (ORC) polar: the benchmark. ⚡ = the isochrone optimizer found a faster way to sail it.

kt — the ORC polar speed at the angle you will sail — the speed to steer to, never a diluted average.

eta — Estimated Time of Arrival, cumulative, at the benchmark pace.

Board rows (“L→ 2-3·A · stbd”) split each beat into its starboard (STBD) and port tacks, with a fetch when one board lays the mark.

Below the TOTAL row: 🏁 the RACE TARGET summary (depart · arrival · duration · nm) and the odds — today: P10 3:50 pm · promise the P90.

One leg card — a real ⚡ board from today's P2

↳ ⚡ 1-2-1 · Camotal 2 (W) → T1

HEADING	TWD	TWA
120°	172°	±52°
BTW	TWS	SAIL
139°	11.0 kt	J2
DTW	ORC SPD	VMG $\angle$
0.7 nm	7.0 kt	±40°
TARGET	START	ETA
0h08	12:00 pm	12:07 pm

T1 is the tack point the optimizer chose — this board sails to T1, the next card sails T1 → S Camotal.

Sail colors are your real cloth: J2 grey · A3 white · A2 bright red — you can match the card to the bag in one glance.

The card explains itself: when the spinnaker is NOT called, the SAIL cell says why — e.g. “A3 not on: chart crossover at 11 kt starts at True Wind Angle (TWA) $\geq 101^\circ$ (here 82°)”.

Wind @ mark on full leg cards shows the wind WHERE YOU ARE GOING and whether it is veering (turning clockwise) or backing — with the bearing at each end of the compass arc.

Moving through the plan — legs, boards, peels

◀ ▶ WALK THE COURSE

The arrows step board by board through the whole race. Each screen shows **ONLY** the active piece, laser-zoomed — previous and next legs appear as faint ghosts so you keep context.

DIVIDED LEGS

A beat shows its tack boards (2-3·A starboard, 2-3·B port...). A sail change mid-leg creates **PEEL** screens the same way — each with its own card, shaded darker on the chart.

THE MAP OBEYS YOU

Background refreshes never yank you off the board you are studying; the **REFRESH** button (or picking a race) returns you to the full summary.

🔪 REAL EXAMPLE

tip: **tap any row of the leg table → the chart zooms to exactly that board**

The odds — what Monte Carlo (MC) adds under your plan

10,000 FUTURES

The forecast will be wrong by SOME amount. The app sails the whole course 10,000 times, each with realistic wind errors (sticky, hour-to-hour) and a slightly different crew day ($113 \pm 3\%$). The arrivals form your odds.

P10 · P50 · P90

Percentiles of those arrivals: P10 = a fast day (only 1 in 10 beats it) · P50 = the honest middle · P90 = 9 of 10 futures have arrived — THE number you promise the crew and the club.

DNF %

Did Not Finish: the share of futures that died in a wind hole. 0% today. When it is not zero, you want to know before the start, not at sunset.

📌 REAL EXAMPLE

real P2 today: **benchmark 4h42 · expected $\approx$ 4h10 · promise (P90) 4h40**

real P1 (71.9 nm): **the evening stretches the tail: P90 - P50 = 70 minutes. Promise accordingly.**

The ⚡ amber rows — verified faster ways to sail a leg

WHAT ⚡ MEANS

After every plan, the isochrone optimizer re-sails each leg through the CHANGING forecast, trying tacks, gybes and sail changes every 10 minutes. When it beats the textbook plan, the amber ⚡ time and its boards replace it. If not, the textbook stands — a slower “optimization” is never accepted.

WHY IT FINDS TIME

The wind of 3 pm is not the wind of noon. Tacking BEFORE a forecast veer turns a header into a lift — no static plan can do that. Every manoeuvre pays its cost: 45 seconds per tack, 120 per sail peel.

WHAT IT FOUND TODAY (P2)

An early first tack off Camotal (−3 min) · a set gybe into the Horadas rounding (−1) · a two-gybe A2 run to Costa Verde riding the veer (−1) · a hot A3 angle home (−7). Across all five P-races: 26–49 minutes, every race.

🔪 REAL EXAMPLE

the guard: ⚡ appears ONLY when the optimized time beats the baseline outright (test T5)

Which side of the beat? The layline call, with odds

THE DUEL

For 1,500 of the futures, each beat is sailed TWICE — starting on port, starting on starboard — through IDENTICAL wind. The only difference is the choice, so the win percentage is pure tactics.

READING THE CALL

“STBD 63%” on a leg card = starboard-first wins 63 of 100 futures. 70%+ = protect that side. Near 50% = genuinely no edge: sail the wind shifts and stay between the fleet and the mark.

SIDE FIRST, THEN POLISH

The layline call SEEDS the optimizer’s first tack (strategy before tactics). Tested on 47 archived race-days: this order won the disagreements 6–4, with 18–55 minutes at stake per call.

📌 REAL EXAMPLE

real extremes, 22 Jul: **P6 San Gallán beats: PORT 96% · 100% · 98% — one side simply right**

and the honest flip: **P1 long beat: 53/47 — nobody knows; sail the shifts**

SECTION 8 OF 8

1 · WIND DASHBOARD

2 · CURRENT · NEXT · FOLLOWING

3 · LOCATION ANALYSIS

4 · WIND FORECAST EVALUATION

5 · ROUTE MODEL

6 · ROUTE SELECTION

7 · ACTIVE ROUTE

8 · APPENDIX

Under the hood · limits · glossary · the five-minute routine

Under the hood: seven layers, one loop

1 • RAW MODEL

sea cells only, mixed by Inverse Distance Weighting (IDW)

2 • MOS

Model Output Statistics: $\text{corrected} = \alpha + \beta \times \text{model}$, capped ± 6 kt

3 • CROWN

the verified best model for this spot at this range

4 • NOWCAST

today's live sensor delta, fading by the learned ladder

5 • TARGET

the 100% ORC benchmark plan — real sails, real coastline

6 • MONTE CARLO

10,000 futures → odds + the side of every beat

7 • OPTIMIZER

manoeuvres refined through the evolving wind, on that side

THE DAILY LOOP

Night: archive today · grade yesterday (out-of-sample).

Every load: re-fit MOS · re-crown models · re-learn the  ladder.

Every 10 min: re-measure live deltas at six sensors.

Every logged race: the boat model itself learns.

Nothing is hand-tuned.

What the app cannot do — read this twice

IT TRUSTS ITS WIND

Every minute and every probability assumes the corrected forecast is roughly right. On a broken forecast day, all engines are wrong TOGETHER. The nightly grades are your guardrail, not a guarantee.

NOT MODELED

The Humboldt current (0.3–0.7 kt), waves, leeway, kelp, traffic, the start-line bias. The chart soundings are approximate — NEVER sail into less than 2 metres on the app's word.

ODDS ARE NOT PROPHECY

A 73% side call loses 27 times in 100 — and was still the right call. Judge your decisions by process, not by one result. That is what professionals do.

An instrument that states its limits is one you can bet a race on.

Glossary — every abbreviation in this app

TWS / TWD / TWA	True Wind Speed / Direction / Angle — the real wind, not the moving-boat illusion	MAE	Mean Absolute Error — the average miss; the grade
AWA / AWS	Apparent Wind Angle / Speed — the wind you feel while moving	IDW	Inverse Distance Weighting — closer grid cells count more
BTW / DTW	Bearing / Distance To Waypoint	MC	Monte Carlo — many random futures instead of one guess
ETA	Estimated Time of Arrival	P10/P50/P90	arrival percentiles: fast day · median · the promise
VMG	Velocity Made Good — progress toward the wind or the mark	DNF	Did Not Finish — futures that died in a wind hole
ORC / VPP	Offshore Racing Congress / Velocity Prediction Program — the rating and its math	METAR	Meteorological Aerodrome Report — the airport weather bulletin
MOS	Model Output Statistics — the sensor-trained correction (the “+”)	STBD	starboard — wind over the right side; PORT — over the left

The five-minute pre-start routine

- 1 Wind Dashboard** colors and alarms for your window — gusts, calms, direction turns.
- 2 Current · Next · Following** the three-station briefing; the compass arcs tell you the shift and its hour.
- 3 Active Route** tap your race in the P menu — read the 🎯 benchmark, then the ⚡ legs (that is where you win).
- 4 The odds** promise the P90, not the target. Check Did Not Finish (DNF) % and the layline call $\geq 70\%$? protect that side.
- 5 Refresh at the dock** 🔄 once before casting off — a fresh model run can move things a few percent.