TEST REPORT · INSTRUMENT VALIDATION

Stroke-to-stroke repeatability and warm-up behaviour of the PaddleMate Pro strain-gauge paddle-shaft sensor under pneumatic square-wave bending, with a comparison to the PaddleMate Nano

Attila Schulcz, MSc

Electrical engineer; CEO, PaddleMate, www.thepaddlemate.com. Correspondence: contact@thepaddlemate.com

Report number	PM-TR-2026-05
Version	1.0, 15 November 2026
Device under test	PaddleMate Pro test unit, single strain gauge, firmware temperature-offset compensation, BLE data transfer
Software	In-house PaddleMate tools; raw ADC and processed force export
Specimen	Carbon-fibre tube, OD 30 × 1 mm wall, 1.2 m; ends held down, rotation free (rig of PM-TR-2026-02); Pro at L/6, 200 mm from the cylinder on the side opposite the Nano position of PM-TR-2026-02
Loading	Pneumatic cylinder under the shaft at ≈ L/3, upward push; 10 bar supply via regulator, setting unchanged since PM-TR-2026-02; three runs of 301, 308 and 301 strokes, 1.44 s period
Ambient	18 °C laboratory; one run from cold switch-on, two runs after more than 10 min powered
Calibration	In situ: same rig inverted, nominal 10 kg mass (98.07 N) hung at the cylinder position
Measurand	Stroke-to-stroke repeatability and warm-up drift; no in-line load cell during the runs

ABSTRACT

Background. Instrumented paddles are a principal tool for measuring what a kayak or canoe athlete delivers [1]. A shaft-mounted sensor must report a repeated load consistently, and it must do so soon after it is switched on. PM-TR-2026-02 measured the stroke-to-stroke repeatability of the PaddleMate Nano [14]. This report applies the same test to the PaddleMate Pro and adds a cold-start run.

Methods. The rig of PM-TR-2026-02 was used: a 1.2 m carbon tube (OD 30 mm, 1 mm wall) held down at both ends and a pneumatic cylinder pushing upward at about one-third of the span, cycling at 1.44 s. The Pro was mounted 200 mm from the cylinder on the side opposite the Nano position of PM-TR-2026-02, where the bending moment is two-thirds of that at the Nano position. It was calibrated in the same rig, inverted, with a nominal 10 kg mass at the cylinder position. The pressure regulator was not touched after the Nano test. At 18 °C, three runs were recorded: raw ADC from cold switch-on (301 strokes), raw ADC after more than 10 min powered (308 strokes), and the processed force output, also warmed up (301 strokes). For each stroke the trimmed means of the unloaded and loaded plateaus were subtracted to give Δ. The first four strokes of the force run were excluded because the plateau segmentation had not yet locked; no other data were excluded.

Results. The warmed-up force output had a mean Δ of 209.40 N (95% CI ±0.11 N), a standard deviation of 0.97 N (95% CI 0.90–1.05 N) and a coefficient of variation of 0.46%. 98% of strokes were within ±1% of the mean and all within ±1.5%. The warmed-up raw Δ had a CV of 0.61%, with no drift. Both distributions were consistent with a normal distribution and contained no outliers by Grubbs' test ($\alpha = 0.05$). After a cold start, Δ began 1.3% below the warmed-up level and rose steadily; its 25-stroke moving mean came within 0.2% of the warmed-up level after about 6 min. Under the same rig and loading, the Nano force output had shown a CV of 0.99%.

Conclusion. The scatter includes the variability of the pneumatic actuator, so 0.46% is an upper bound for the stroke-to-stroke repeatability of the Pro. This is about half the scatter of the Nano, obtained with two-thirds of the bending moment at the gauge, and more than an order of magnitude below the ±10% class A tolerance of ISO 20957 for rowing training equipment, and at the level of laboratory-grade cycling power

meters. A short warm-up of about 6–7 min after switch-on removes the start-up drift. Absolute accuracy and temperature effects are reported separately.

KEYWORDS kayak; canoe; instrumented paddle; strain gauge; repeatability; coefficient of variation; warm-up drift; dynamic loading

1. Introduction

Strain-gauge paddle shafts infer stroke force from shaft bending. Their accuracy is usually established with known static loads in a defined geometry [2, 3]. Accuracy at a point, however, says little about how consistently a sensor reports the same load when the load is applied and removed many times, as it is during paddling. Sports science expresses this consistency as a typical error or coefficient of variation [4, 5]; metrology calls it repeatability [6].

PM-TR-2026-02 measured the stroke-to-stroke repeatability of the PaddleMate Nano on a pneumatic bending rig and found a CV of 0.99% in force mode [14]. The PaddleMate Pro is the shaft-mounted power and force meter of the same product family. Its static linearity and its temperature behaviour have been reported in PM-TR-2026-03 and PM-TR-2026-04 [15, 16], but its dynamic repeatability has not.

A second practical question is how soon after switch-on the reading can be trusted. Strain-gauge bridges and their electronics warm up when powered, and this can shift both the zero and the sensitivity for some minutes [17]. Firmware offset compensation and the per-stroke difference remove a zero shift, but not a change of sensitivity.

This report therefore asks three questions. How large is the stroke-to-stroke scatter of the Pro under the same repeated load as in PM-TR-2026-02? Does a cold-started Pro drift, and for how long? How does the Pro compare with the Nano on the same rig? As in PM-TR-2026-02, the test was not designed to measure accuracy: the actuator force was not measured independently during the runs.

2. Nomenclature

Table 1. Symbols used in this report.

Symbol	Meaning	Unit
L_k, H_k	trimmed mean of the unloaded (low) and loaded (high) plateau of stroke k	N or counts
Δ_k	stroke result, $H_k - L_k$	N or counts
n	number of strokes analysed in a run, 297–308	–
$\bar{x}$, s	sample mean; sample standard deviation	N or counts
CV	coefficient of variation (relative standard deviation), $100\ s/\bar{x}$	%
CV_d, CV_s	CV after removing a linear trend; CV from successive differences, $100\ \sqrt{(\text{mean}(\Delta_{k+1} - \Delta_k)^2 / 2)} / \bar{x}$	%
r_1	lag-1 autocorrelation between consecutive strokes	–
G	Grubbs statistic, $\max \lvert x_k - \bar{x} \rvert / s$	–

3. Materials and methods

3.1 Device under test and software

The PaddleMate Pro is a paddle-shaft force and power meter with a single strain gauge. Before each stroke its firmware compensates the temperature-dependent offset of the gauge; it does not compensate the sensitivity. It streams data over Bluetooth Low Energy. Recording and export used the same in-house PaddleMate software as PM-TR-2026-02, which can pull both the unprocessed ADC stream and the processed force output and writes per-stroke CSV files. Two output modes were recorded:

- **Force output.** The calibrated force in newtons, as reported by the firmware and the application. The offset is compensated, so the unloaded level is close to zero. No additional client filter was applied.
- **Raw output.** The uncalibrated ADC value, without offset compensation. A client-side moving filter of 10 samples in the in-house software was active during the raw runs.

3.2 Specimen and loading

The rig was the one described in PM-TR-2026-02 [14]; only the sensor position differed (Figure 1). A 1.2 m carbon-fibre tube with an outer diameter of 30 mm and a 1 mm wall was held down at both ends. The clamps only held the tube down against the upward load, so the tube could rotate at the clamps and bend freely between them; mechanically it is a simply supported span. A pneumatic working cylinder under the tube at $a = L/3 \approx 400$ mm from clamp A pushed it upward. In the retracted position the cylinder rested just clear of the shaft, so the unloaded plateau carried no actuator force.

The Nano of PM-TR-2026-02 sat at $L/2$, 200 mm from the cylinder towards clamp B. The Pro was mounted at the mirror position, 200 mm from the cylinder towards clamp A, i.e. at $x = L/6 \approx 200$ mm. For a point load F at a, the bending moment at the gauge position x is [18]

$$\mathrm{M(x) = F\,(L - a)\,x\,/\,L \text{ for } x \le a, \quad M(x) = F\,a\,(L - x)\,/\,L \text{ for } x \ge a} \tag{1}$$

This gives 0.133 F·m at the Pro position and 0.200 F·m at the Nano position. For the same actuator force, the gauge of the Pro therefore saw two-thirds of the bending moment, and of the strain, that the Nano had seen.

The air supply was 10 bar. A pressure regulator and an air-preparation unit reduced the pressure, and a valve cycled the cylinder. The regulator setting of PM-TR-2026-02 was not touched afterwards, so the nominal actuator force was the same as in the Nano test and the same for all three runs here. The Pro indicated about 209 N, the Nano had indicated 201 N; this difference is discussed in Section 4.6.

The Pro was calibrated in this same rig before the runs. The rig was turned upside down, and a nominal 10 kg mass (98.07 N with $g_0 = 9.80665$ m s^{-2}) was hung at the cylinder position ($L/3$), so that the load acted on the tube in the same direction relative to the clamps and the gauge as the cylinder later did. The Pro calibration was made separately from, and at a different time than, the Nano calibration. The calibration tool did not accept a gram-precise mass value, and it is possible that the Pro had not fully warmed up when it was calibrated. The force output thus estimates the force at the loading point, within the accuracy of this single-point calibration. During the dynamic runs, no in-line load cell measured that force stroke by stroke.

The laboratory temperature was 18 °C, and the shaft was not conditioned. For the cold-start run the Pro was switched on immediately before recording, after being unpowered at room temperature; this run was intended to show the start-up drift. For the two other runs the Pro had been powered for more than 10 min. The actuator cycled at a mean period of 1.437–1.438 s in all runs, from the stroke time stamps.

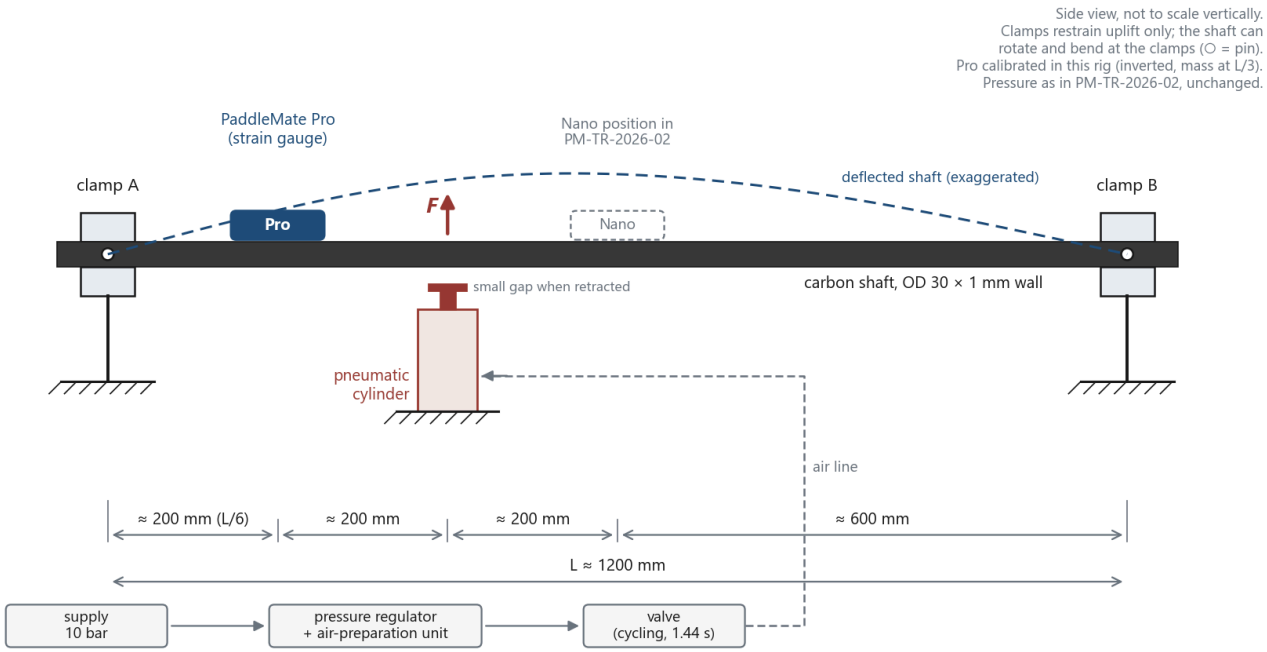


Figure 1. Loading rig, side view (schematic); rig and pressure as in PM-TR-2026-02. Carbon tube, OD 30 × 1 mm wall. The clamps restrain uplift only, so the tube can rotate at the clamps (○); the dashed line shows the deflected shape, exaggerated. The PaddleMate Pro sits at *L*/6, 200 mm from the cylinder towards clamp A; the dashed outline marks the Nano position of PM-TR-2026-02 at *L*/2, 200 mm from the cylinder towards clamp B. The pneumatic cylinder acts upward at *L*/3 and clears the shaft when retracted. Supply: 10 bar through a pressure regulator and air-preparation unit; a valve cycles the cylinder at 1.44 s. For calibration the same rig was inverted and a nominal 10 kg mass hung at *L*/3.

3.3 Stroke detection and per-stroke result

The square wave was segmented by hysteresis thresholds at 30% and 70% of the step height (Figure 2). For each plateau a window of samples was taken, a fraction of it was trimmed to remove the transitions, and the remaining samples were averaged. The stroke result was

$$\Delta_k = H_k - L_k \tag{2}$$

Window and trim settings differed between the output modes and are listed in Table 2. Because Δ_k is a difference within one stroke, a slow drift of the zero that moves both plateaus equally cancels in it. A change of sensitivity does not cancel.

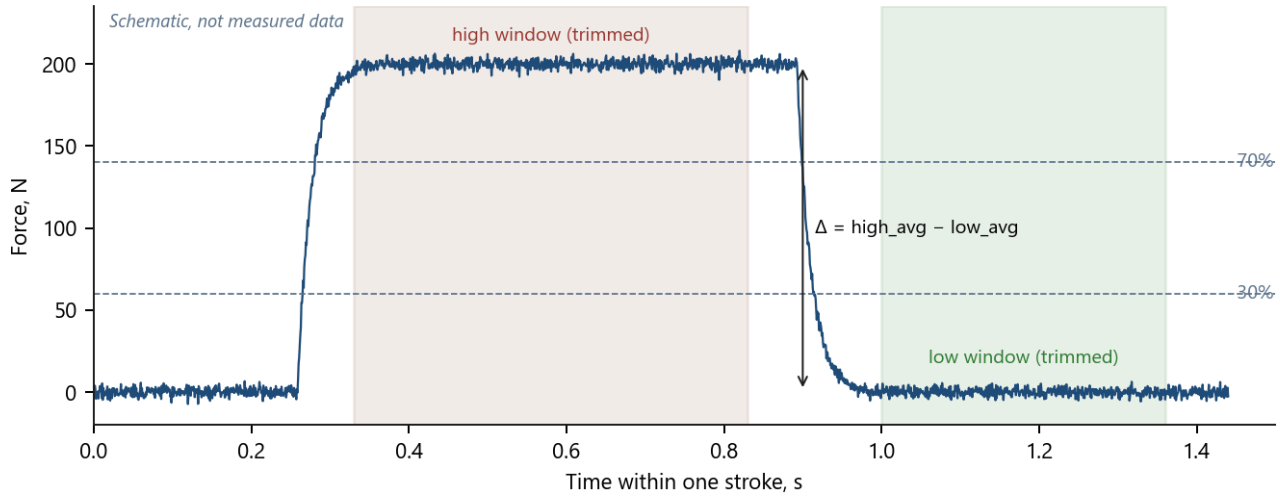


Figure 2. Per-stroke evaluation, schematic (from PM-TR-2026-02). Hysteresis thresholds at 30% and 70% detect the transitions. Trimmed windows on the loaded (high) and unloaded (low) plateau give H_k and L_k, and their difference is the stroke result Δ_k. Illustrative signal, not measured data.

3.4 Runs

Table 2. Recorded runs, 15 November 2026. Times are UTC (local time UTC+1).

Run	Output	Device state	Window / trim	Filter	Start	End	Strokes
1	Raw ADC	cold start	200 / 30%	10-sample	13:18:45	13:25:57	301
2	Raw ADC	warmed up	200 / 30%	10-sample	13:44:26	13:51:48	308
3	Force, N	warmed up	100 / 20%	off	14:03:12	14:10:23	301 (297)

Samples averaged per plateau, unloaded / loaded: run 1, 12–27 / 13–24; run 2, 19–27 / 16–24; run 3, 47–48 / 40–41. In run 3, strokes 1–4 were excluded (see below), leaving 297.

The runs were consecutive, not simultaneous, so raw and force values cannot be paired stroke by stroke. Because the evaluation settings also differed, comparing the CVs of runs 2 and 3 is only indicative.

In run 3 the first four strokes were excluded before analysis. At the start of the recording the plateau segmentation had not yet locked: the unloaded plateau of strokes 2 and 3 was at −208 N, which is physically impossible in this rig, and strokes 2 and 4 gave Δ of 493 N and 163 N. Stroke 1 had an unloaded plateau at +76 N. These are evaluation artefacts of the recording start, a documented cause under the exclusion rule of Section 3.5. From stroke 5 onward both plateaus were stable (unloaded −0.23 ± 1.27 N, loaded 209.17 ± 1.57 N).

3.5 Statistical analysis

The analysis followed PM-TR-2026-02 [14]. For each series the following were computed, with the scatter treated as a Type A evaluation [7]:

- mean, with a 95% confidence interval based on the *t*-distribution;
- standard deviation, with a 95% confidence interval based on the χ^2-distribution;
- CV, range and 2.5th, 50th and 97.5th percentiles;
- skewness and excess kurtosis.

Two tests checked independence. The lag-1 autocorrelation r_1 tested whether consecutive strokes were related. A least-squares trend against stroke number tested for drift, using the *t*-statistic of the slope. Because the runs here are three times longer than in PM-TR-2026-02, two drift-robust measures were added: CV_d, the CV of the residuals after removing the linear trend, and CV_s, computed from successive differences, which is insensitive to slow drift of any shape.

Normality was examined with a normal quantile–quantile plot. The outlier criterion was fixed before evaluation. Grubbs' test was applied two-sided at $\alpha = 0.05$ to the most extreme stroke [8, 9]; the critical value is $G_{crit} = 3.72$ for $n = 297$ and 3.73 for $n = 308$. A stroke would be excluded only if it exceeded G_{crit}, or if a documented physical or evaluation cause was found. Apart from the four start-up strokes of run 3, no stroke met either condition (Section 4.3).

The warm-up was described by expressing each cold-start Δ_k relative to the mean of the warmed-up raw run and smoothing with a 25-stroke moving mean (about 36 s).

4. Results

4.1 Force output

Table 3. Descriptive statistics of the stroke result Δ.

Statistic	Force, warmed up, N	Raw, warmed up, counts	Raw, cold start, counts
Strokes *n*	297	308	301
Mean (95% CI half-width)	209.40 (±0.11)	650.76 (±0.44)	644.79 (±0.53)
Standard deviation *s* (95% CI)	0.968 (0.896–1.053)	3.94 (3.65–4.28)	4.71 (4.36–5.12)
Coefficient of variation	0.46%	0.61%	0.73%
CV_d (trend removed) / CV_s (successive)	0.46% / 0.44%	0.61% / 0.63%	0.59% / 0.59%
Minimum / maximum	207.18 / 212.33	640.56 / 659.96	632.89 / 656.86
Range	5.15	19.39	23.97
Percentiles 2.5 / 50 / 97.5	207.54 / 209.43 / 211.41	642.59 / 650.71 / 658.50	636.26 / 644.13 / 654.60
Skewness / excess kurtosis	0.06 / −0.26	−0.06 / −0.35	0.20 / −0.29
Lag-1 autocorrelation r_1	0.09	−0.10	0.35
Trend per 100 strokes (*t*)	−0.18 N, −0.08% (−2.70)	−0.07, −0.01% (−0.29)	+3.18, +0.49% (12.6)
Grubbs *G* (stroke) vs G_{crit}	3.03 (no. 134) vs 3.72	2.59 (no. 2) vs 3.73	2.56 (no. 297) vs 3.72

|*t*| < 1.97 means the trend is not significant at the 5% level.

The warmed-up force-output Δ averaged 209.40 N with a standard deviation of 0.97 N, giving a CV of 0.46% (Table 3, Figure 3). Consecutive strokes were practically uncorrelated (r_1 = 0.09). A small downward trend was detectable over the 7 min run: −0.18 N per 100 strokes, about −0.5 N (−0.25%) in total. It is statistically significant because the run is long and the scatter small, but it does not contribute to the scatter: CV_d and CV_s are 0.46% and 0.44%. Its source cannot be identified from this run; the actuator supply and the sensor are both possible.

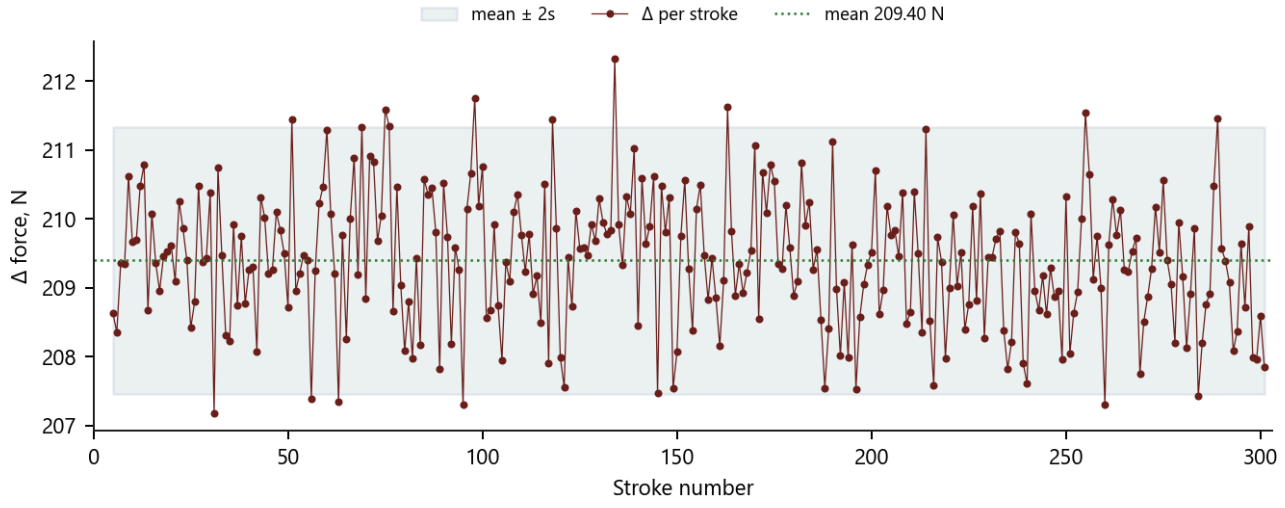


Figure 3. Force-output stroke result Δ_k, warmed up, strokes 5–301. Dotted line: mean. Shaded band: mean ± 2*s*.

4.2 Distribution and tolerance bands

Figure 4 shows the deviations from the series mean in 0.5 N bins, and Figure 5 the cumulative fraction of strokes within a given deviation. Skewness was 0.06 and excess kurtosis −0.26, and the quantile–quantile plot (Figure 6) follows the identity line. The distribution is therefore close to normal. 290 of 297 strokes (98%) were within ±1% of the mean, and all were within ±1.5%.

Table 4. Fraction of strokes within a deviation band around the series mean, force output, warmed up, *n* = 297.

Band, % of mean	±0.5%	±1.0%	±1.5%	±2.0%	±2.5%	±3.0%
Band, N	±1.05	±2.09	±3.14	±4.19	±5.24	±6.28
Strokes	208	290	297	297	297	297
Fraction	70%	98%	100%	100%	100%	100%

Mean absolute deviation 0.77 N (0.37%). Largest absolute deviation 2.93 N (1.40%), stroke 134.

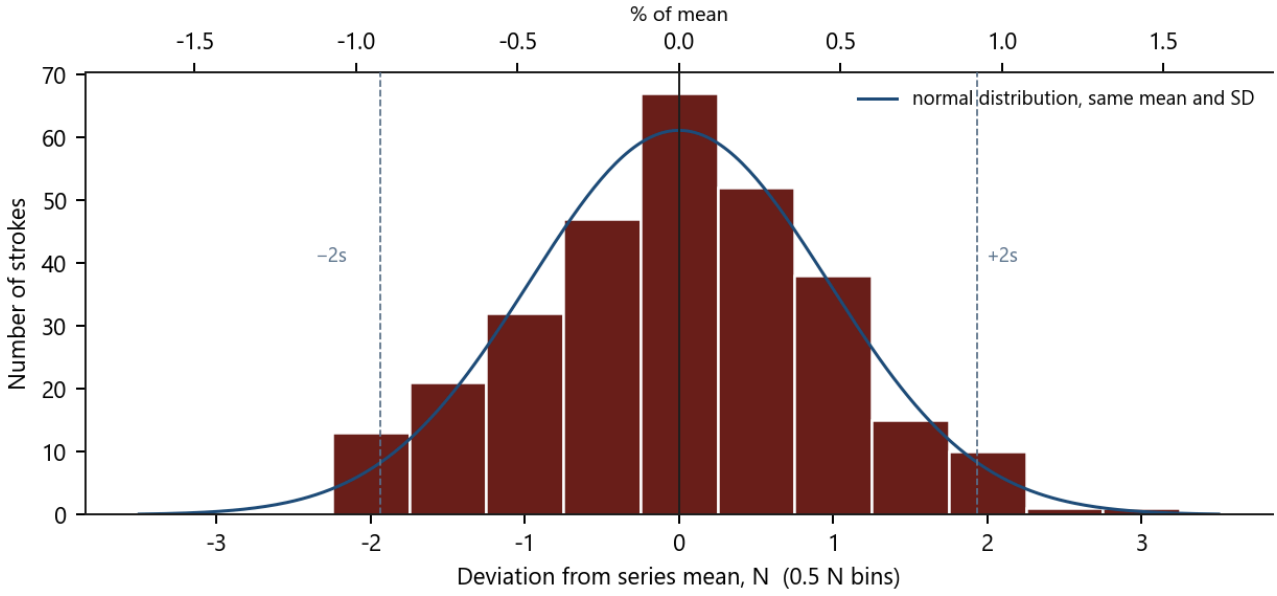


Figure 4. Deviation of the force-output Δ from the series mean, in 0.5 N bins. Blue curve: normal distribution with the sample mean and standard deviation, scaled to *n*. Dashed lines: ±2*s*. Top axis: deviation as a percentage of the mean.

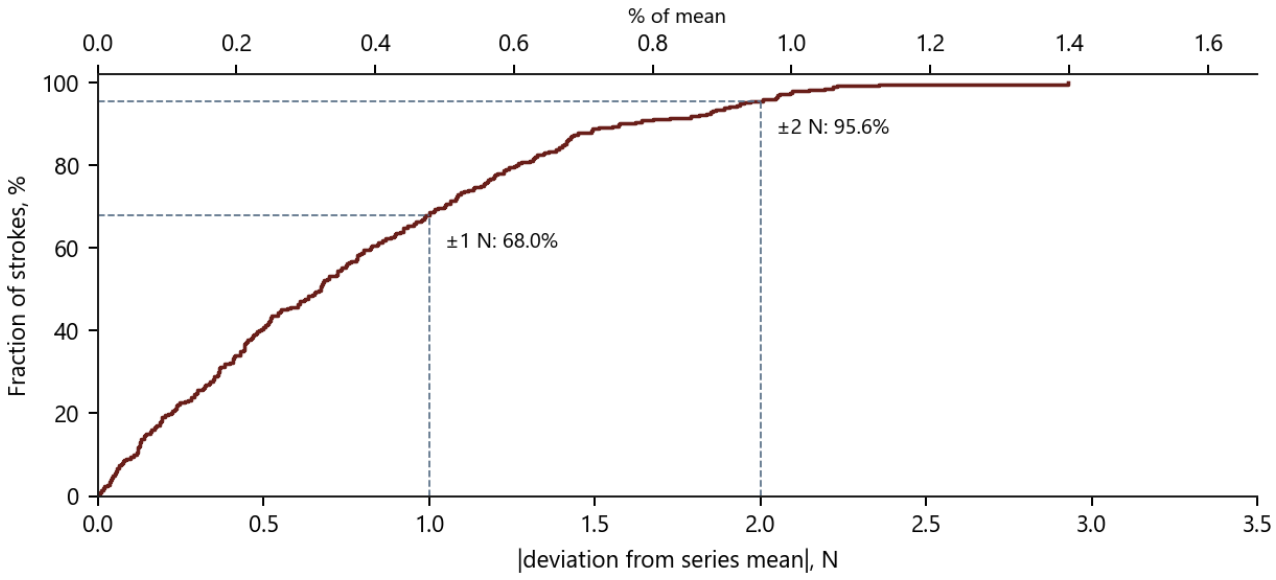


Figure 5. Cumulative fraction of strokes whose absolute deviation from the mean is at most the abscissa value, force output. Top axis: the same deviation as a percentage of the mean.

4.3 Outliers

The most extreme force-output stroke (no. 134, 212.33 N) gave $G = 3.03$, below $G_{crit} = 3.72$. It was the only stroke beyond 3*s*, about as expected for 297 strokes from a normal distribution (0.8). Sixteen strokes fell outside ±2*s*, compared with 13.5 expected. In the warmed-up raw run the most extreme stroke (no. 2) gave

G = 2.59, with no stroke beyond 3s and 17 beyond ±2s (14.0 expected). No physical anomaly was observed during the runs. Apart from the four start-up strokes of run 3 (Section 3.4), all strokes were retained.

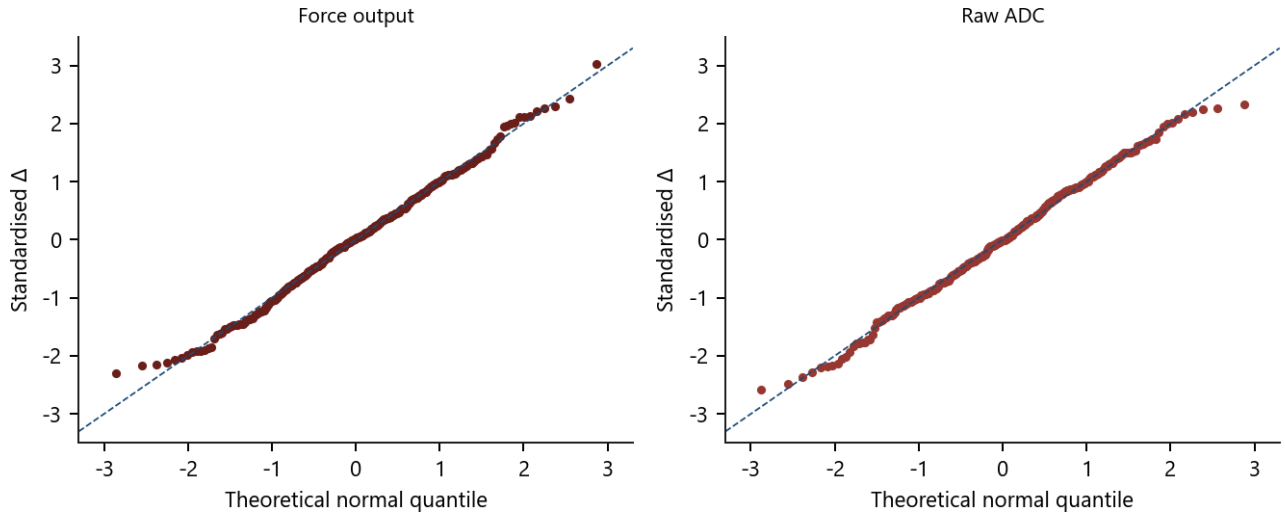


Figure 6. Normal quantile–quantile plots of the standardised stroke result, warmed up. Left: force output. Right: raw ADC. Dashed line: identity, i.e. perfect normality.

4.4 Raw output

The warmed-up raw Δ averaged 650.76 counts with s = 3.94 counts, a CV of 0.61% (Table 3, Figure 7). There was no drift (−0.01% per 100 strokes, t = −0.29), and consecutive strokes were nearly uncorrelated (r_1 = −0.10). An indicative scale of 209.40 N / 650.76 counts = 0.322 N per count follows from the means of runs 2 and 3. On that scale the raw scatter corresponds to about 1.27 N. The Pro thus delivers about eight times more counts per newton at the loading point than the Nano (2.58 N per count [14]), although its gauge saw only two-thirds of the bending moment.

The raw scatter is larger than the 0.97 N of the force run. The runs differ in time and in evaluation settings: the force run averaged about twice as many samples per plateau (Table 2). The difference therefore cannot be attributed to the force processing itself.

The absolute levels of both plateaus drifted during the warmed-up raw run, by about −9 counts (≈ 3 N equivalent) over 7 min, but they moved together: the correlation between the unloaded and loaded plateau was 0.82, and Δ did not drift. This is the offset drift that the per-stroke difference, and in normal use the firmware offset compensation, removes.

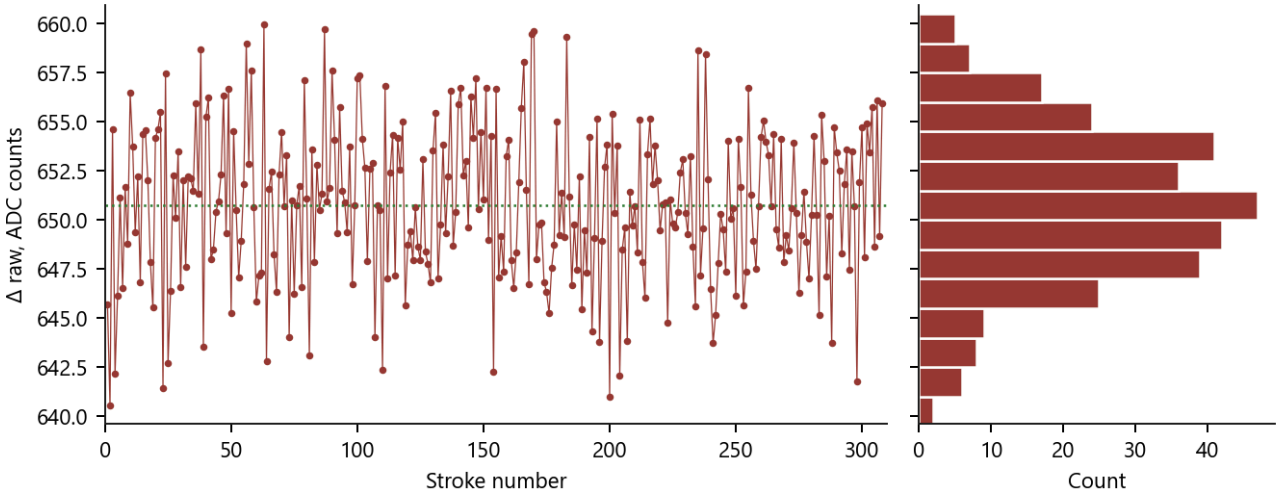


Figure 7. Raw-output stroke result Δ_k in ADC counts, warmed up, 308 consecutive strokes, with its distribution on the right. Dotted line: mean.

4.5 Warm-up after a cold start

In the cold-start run, Δ rose steadily (Figure 8). The first 25 strokes averaged 1.3% below the warmed-up level. The trend was +0.49% per 100 strokes (t = 12.6), about +1.5% over the 7 min run, and it made consecutive strokes correlated (r_1 = 0.35). The 25-stroke moving mean came within 0.2% of the warmed-up level at about stroke 262, 6.3 min after switch-on. The last 50 strokes averaged 0.15% below the warmed-up level.

The drift is a change of sensitivity, not only of zero. Both plateaus moved down, but the unloaded plateau moved more (−21 counts between the first and last 25 strokes) than the loaded plateau (−13 counts), so their difference grew. The offset part cancels in Δ and is in any case removed by the firmware before each stroke; the sensitivity part does not cancel.

Once the trend is removed, the cold-start scatter (CV_d = CV_s = 0.59%) is the same as in the warmed-up run (0.61–0.63%). The cold start therefore shifted the level of the reading but did not make single strokes less repeatable.

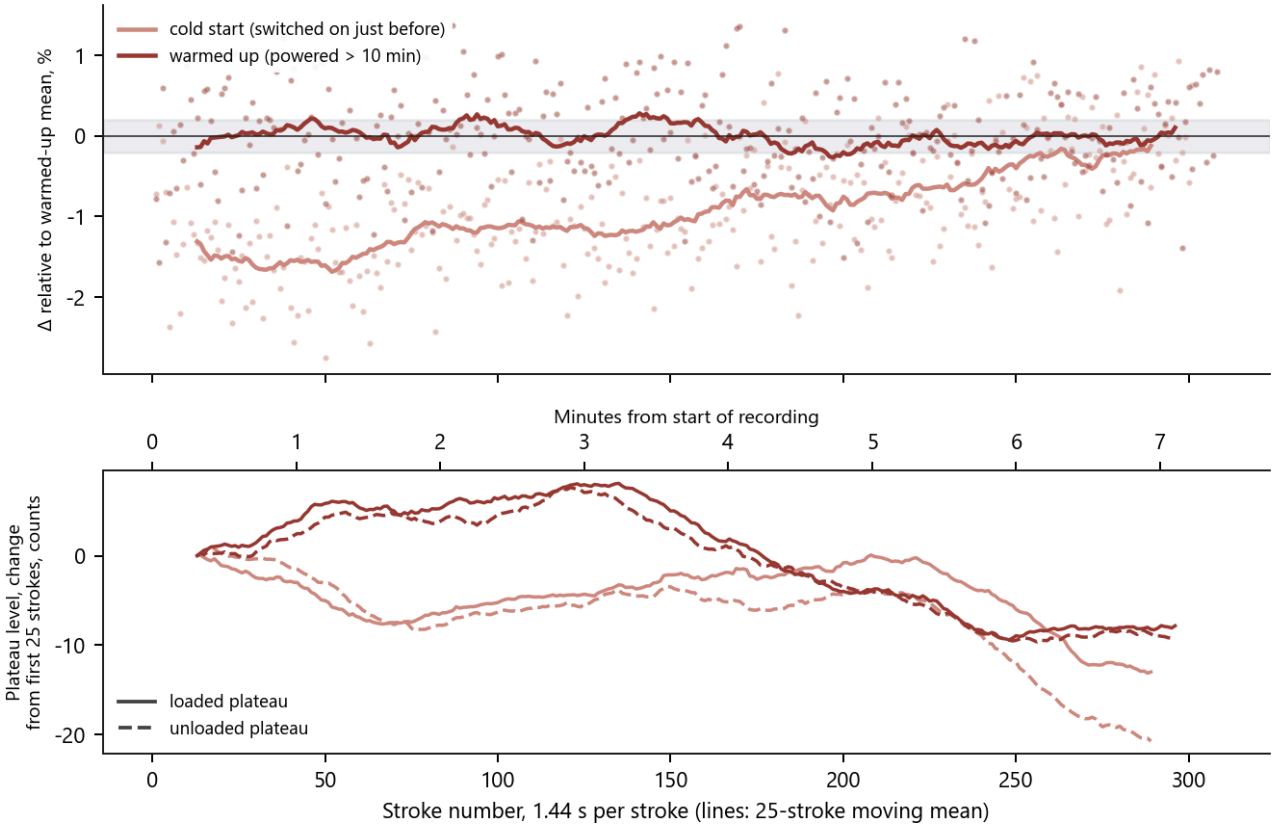


Figure 8. Cold start (switched on immediately before recording) compared with the warmed-up run (powered for more than 10 min), raw output, same rig and pressure. Top: Δ_k relative to the mean of the warmed-up run; dots are single strokes, lines are 25-stroke moving means; grey band ±0.2%. Bottom: 25-stroke moving means of the loaded (solid) and unloaded (dashed) plateau levels, as change from the first 25 strokes. Top axis: minutes from the start of recording.

4.6 Comparison with the PaddleMate Nano

The Nano was tested on the same rig, with the same regulator setting, at the same laboratory temperature and with the same analysis, one day earlier [14]. Its gauge sat at the mirror position on the other side of the cylinder, where the bending moment is 1.5 times larger, equation (1). Table 5 and Figure 9 compare the warmed-up runs.

Table 5. Stroke-to-stroke repeatability of the PaddleMate Nano (PM-TR-2026-02, 14 November 2026) and the PaddleMate Pro (this report, 15 November 2026), warmed up, same rig.

Series	*n*	Mean Δ	*s* (95% CI)	CV	r_1	Within ±1%	Within ±2%
Nano, raw	101	77.95 counts	0.683 (0.600–0.793)	0.88%	−0.03	72%	99%
Nano, force	101	201.19 N	1.98 N (1.74–2.30)	0.99%	0.02	70%	95%
Pro, raw	308	650.76 counts	3.94 (3.65–4.28)	0.61%	−0.10	90%	100%
Pro, force	297	209.40 N	0.97 N (0.90–1.05)	0.46%	0.09	98%	100%

CV_d / CV_s: Nano raw 0.88 / 0.88%, Nano force 0.99 / 0.97%, Pro raw 0.61 / 0.63%, Pro force 0.46 / 0.44%. Window / trim and client filter: Nano raw 200 / 30%, 20-sample; Nano force 200 / 30%, off; Pro raw 200 / 30%, 10-sample; Pro force 100 / 20%, off.

In force mode the standard deviation of the Pro was half that of the Nano, 0.97 N against 1.98 N, for the same nominal actuator force. The confidence intervals of the two standard deviations do not overlap. In raw mode the CV was 0.61% against 0.88%. In both sensors the scatter was random, normally distributed and without outliers, and the drift-robust CVs agree with the plain CVs.

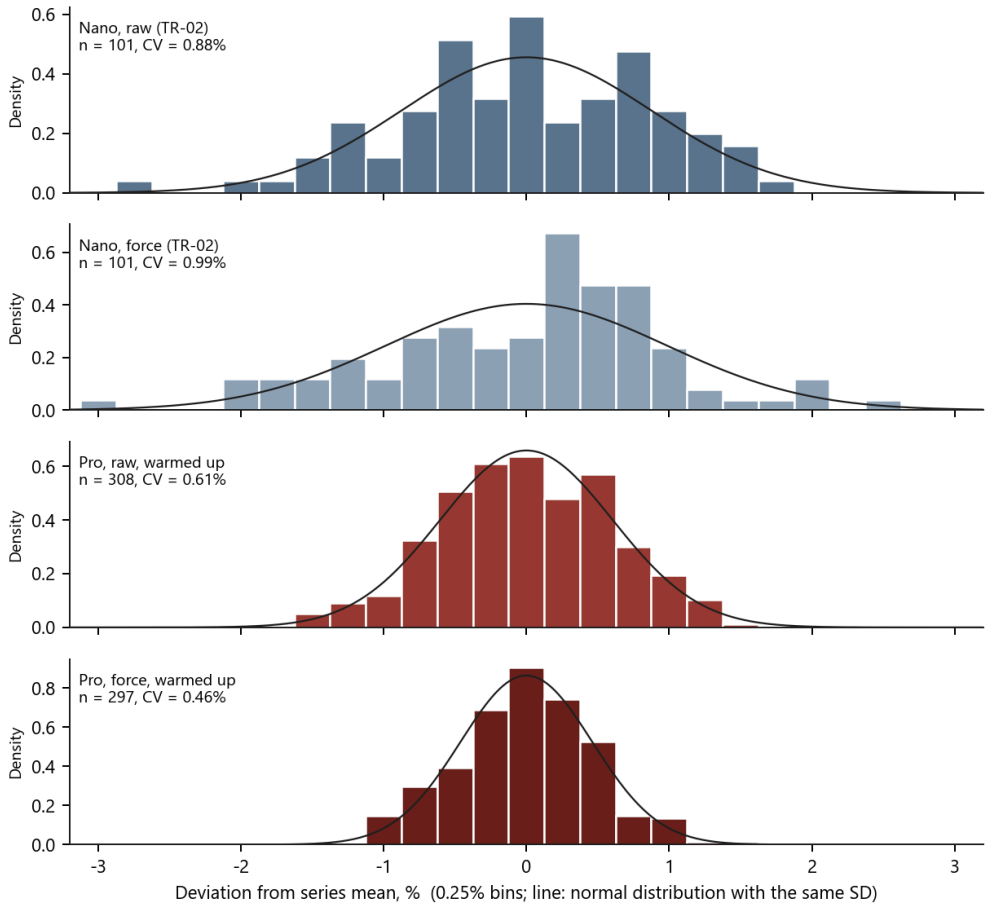


Figure 9. Distribution of the stroke result, deviation from the series mean in per cent, 0.25% bins, normalised to unit area. From top: Nano raw and force (PM-TR-2026-02), Pro raw and force (this report), all warmed up. Line: normal distribution with the same standard deviation.

The evaluation settings were not identical. The Pro force run used a shorter window with less trimming than the Nano force run. A shorter window averages fewer samples and, if anything, adds scatter, so the smaller scatter of the Pro does not come from the settings. The Nano force output also has 1 N resolution [14], which adds a small quantisation component. The sensor position works against the Pro: with two-thirds of the strain, the same electronic noise would be a larger fraction of the signal.

Different mean readings. With the regulator unchanged, the Pro read 209.40 N and the Nano 201.19 N, a difference of 4.1%. Both sensors were calibrated separately with a single hanging mass, and the difference reflects the accuracy of those calibrations, not the repeatability. Three sources are identified:

- the two calibrations were made at different times, each with its own placement of the mass;
- the Pro calibration tool did not accept a gram-precise mass value;
- the Pro may not have been fully warmed up during calibration. Section 4.5 shows that a cold Pro gives up to 1.3% fewer counts; a slope set in that state makes the warmed-up Pro read correspondingly high.

The placement of the mass matters in this geometry. By equation (1), if the mass hangs a distance δ from the axis of the cylinder towards clamp B, the calibration moment changes by $-\delta/(L - a) = -0.125\%$ per mm at the Pro and by $+\delta/a = +0.25\%$ per mm at the Nano, with opposite effects on their readings. A placement difference of about 10 mm, or a combination of the three sources, is enough to explain 4%. Because the repeatability is expressed relative to each sensor's own mean, it is not affected.

5. Discussion

Under a repeated pneumatic load, the warmed-up PaddleMate Pro reported the stroke result with a CV of 0.46% in force mode and 0.61% in raw mode. The scatter was random: consecutive strokes were practically uncorrelated, the distribution was close to normal, and the pre-defined outlier test flagged no stroke. In sports-science terms, the typical error of a single stroke is about 1 N at a 209 N load [4, 5]. Averaging over ten strokes reduces it to about 0.3 N.

What the scatter contains. As in PM-TR-2026-02, the observed standard deviation combines the stroke-to-stroke variation of the actuator force, the sensor noise after averaging and, in force mode, the quantisation of the output. The test cannot separate them, so 0.97 N (0.46%) is an upper bound for the repeatability of the Pro itself. It also bounds the actuator: the rig cannot have contributed more than about 0.46% on 15 November. The regulator was not changed between the two tests, so the actuator is expected to have behaved similarly one day earlier. Most of the 0.99% of the Nano then came from the Nano and not from the rig.

Warm-up. From a cold switch-on, the sensitivity of the Pro rose by about 1.5% over 6–7 min and then settled. The cause cannot be identified here: self-heating of the gauge, the bridge and the measuring electronics are all possible [17], and the electronics temperature was not measured. In practice, a paddler's warm-up usually lasts longer than 7 min, so a Pro switched on at the start of the warm-up is settled before the main work. For testing and calibration, the Pro should be powered for at least 10 min beforehand, as it was for the warmed-up runs here. A Pro calibrated immediately after a cold switch-on would read about 1% high once warmed up.

Comparison with sport-specific equipment. For rowing machines, the closest standardised equipment to paddling, ISO 20957 class A requires the displayed power to be within ±10% above 50 W [10, 11]. Mobile cycling power meters showed a mean precision of 1.2 ± 0.9% CV in a calibration study of 54 units, and about 0.8% for the SRM and PowerTap brands [12]. The best performance tests in sport science reach a CV of about 1–2% [4]. With a CV of ≤ 0.46%, including the actuator, the stroke-to-stroke scatter of the Pro is more than twenty times smaller than the class A tolerance and at or below the level of laboratory-grade cycling power meters.

Laboratory test and paddling. A bench test is not the same as paddling on water. On water the load profile is not a square wave, the centre of pressure on the blade moves during the stroke [13], and the paddle is exposed to water, wind and temperature changes. The athlete's own stroke-to-stroke variation is also much larger than the scatter found here. The value of the bench test is that it is controlled and reproducible: the same load is applied hundreds of times in the same geometry, so the scatter of the instrument can be measured on its own, which is impossible on water.

Relation to accuracy claims. Repeatability and accuracy are separate properties [6]. A sensor can repeat to 0.5% and still read several per cent off if its calibration slope is wrong. The Pro was calibrated at 98 N and tested at 209 N, so its mean reading also depends on its linearity, which is reported in PM-TR-2026-03 [15]; the repeatability result does not. The 4.1% difference between the mean readings of the Pro and the Nano under the same actuator setting (Section 4.6) shows that a single hanging-mass calibration, with the placement of the mass, the entered mass value and the warm-up state of the sensor, limits the absolute accuracy far more than the repeatability does. A calibration jig with a defined load point, the exact mass entered and a warmed-up sensor would reduce it [2]. The effect of shaft temperature on the sensitivity is reported in PM-TR-2026-04 [16].

Limitations.

- One device, one shaft, one load level and one laboratory temperature were tested. The serial number of the Pro unit is not reported here.

- The three runs were consecutive, not simultaneous, and the force and raw runs used different evaluation settings in the in-house software.
- The Nano and the Pro were tested on different days and at mirror positions about the cylinder, not side by side. The regulator was unchanged, but the actuator force was not measured on either day.
- The calibration was a single-point hanging-mass calibration, made with a mass value that could not be entered to the gram and possibly before full warm-up. It limits the absolute accuracy of the force reading (Section 4.6) but not the repeatability.
- The load waveform was a square wave, not an on-water force profile, and the loading point was fixed.
- No in-line load cell measured the actuator force during the runs, so the actuator's own stroke-to-stroke variation, and the source of the small trend in the force run, remain part of the result.
- The cold-start behaviour was recorded once, at 18 °C. The electronics temperature was not measured, so gauge and electronics warm-up cannot be separated.
- The four start-up strokes of the force run were excluded as evaluation artefacts; this rule should be built into the recording software.

A follow-up test on the same rig is planned with the shaft temperature-conditioned: about 1000 strokes at 18 °C, continuous recording while the shaft is warmed to 35 °C, and about 1000 strokes at 35 °C. This extends the temperature result of PM-TR-2026-04 without re-placing a calibration mass.

6. Conclusion

Over 297 consecutive pneumatic strokes at about 209 N, the warmed-up PaddleMate Pro force output repeated with a standard deviation of 0.97 N (95% CI 0.90–1.05 N), a CV of 0.46%. 98% of strokes were within ±1% of the mean and all within ±1.5%. The raw ADC difference repeated with a CV of 0.61% and did not drift. The scatter was random and normally distributed, with no outliers. Because the scatter includes the actuator, these figures are upper bounds for sensor repeatability.

After a cold switch-on, the reading started 1.3% low and settled within 0.2% after about 6 min; the stroke-to-stroke scatter was unaffected. On the same rig and with the same actuator setting, the Pro scattered half as much as the PaddleMate Nano (0.46% against 0.99% CV in force mode), although its gauge saw two-thirds of the bending moment. The 4.1% difference between their mean readings comes from their separate hanging-mass calibrations. Both sensors are well within the ±10% class A tolerance of ISO 20957 for rowing training equipment [11], and the Pro matches or exceeds the precision of laboratory-grade cycling power meters [12]. Absolute accuracy and temperature effects are addressed by the other reports of this series.

Conflict of interest

The author is the CEO of PaddleMate, which develops and sells the devices tested and the in-house recording software. The measurements were made by the manufacturer and have not been independently verified.

Data availability

The per-stroke CSV files of all runs and the analysis script are available on request from contact@thepaddlemate.com.

References

1. Cruz, M.I.; Sarmento, H.; Amaro, A.M.; Roseiro, L.; Gomes, B.B. Advancements in performance monitoring: a systematic review of sensor technologies in rowing and canoeing biomechanics. *Sports* **2024**, 12, 254. doi:10.3390/sports12090254
2. Rosdahl, H.; Aitken, D.; Osborne, M.; Willén, J.; Nilsson, J. A new versatile jig for the calibration and validation of force metrics with instrumented paddles in sprint kayaking. *Sensors* **2024**, 24, 4870. doi:10.3390/s24154870
3. Goreham, J.A.; Ladouceur, M. The One Giant Leap commercial wireless power meter can be used for sprint kayaking with the appropriate calibration. *Frontiers in Physiology* **2025**, 16, 1461644. doi:10.3389/fphys.2025.1461644

4. Hopkins, W.G. Measures of reliability in sports medicine and science. *Sports Medicine* **2000**, 30(1), 1–15. doi:10.2165/00007256-200030010-00001
5. Atkinson, G.; Nevill, A.M. Statistical methods for assessing measurement error (reliability) in variables relevant to sports medicine. *Sports Medicine* **1998**, 26(4), 217–238. doi:10.2165/00007256-199826040-00002
6. JCGM 200:2012. *International vocabulary of metrology — Basic and general concepts and associated terms* (VIM), 3rd ed.; Joint Committee for Guides in Metrology, 2012.
7. JCGM 100:2008. *Evaluation of measurement data — Guide to the expression of uncertainty in measurement* (GUM); Joint Committee for Guides in Metrology, 2008.
8. Grubbs, F.E. Procedures for detecting outlying observations in samples. *Technometrics* **1969**, 11(1), 1–21. doi:10.1080/00401706.1969.10490657
9. ASTM E178-21. *Standard Practice for Dealing With Outlying Observations*; ASTM International: West Conshohocken, PA, 2021.
10. ISO 20957-1:2024. *Stationary training equipment — Part 1: General safety requirements and test methods*; International Organization for Standardization: Geneva, 2024.
11. ISO 20957-7:2024. *Stationary training equipment — Part 7: Rowing equipment — Additional specific safety requirements and test methods*; International Organization for Standardization: Geneva, 2024.
12. Maier, T.; Schmid, L.; Müller, B.; Steiner, T.; Wehrlin, J.P. Accuracy of cycling power meters against a mathematical model of treadmill cycling. *International Journal of Sports Medicine* **2017**, 38(6), 456–461. doi:10.1055/s-0043-102945
13. Tullis, S.; Galipeau, C.; Morgoch, D. Detailed on-water measurements of blade forces and stroke efficiencies in sprint canoe. *Proceedings* **2018**, 2, 306. doi:10.3390/proceedings2060306
14. Schulcz, A. *Stroke-to-stroke repeatability of the PaddleMate Nano strain-gauge paddle-shaft sensor under pneumatic square-wave bending*. PaddleMate Technical Report PM-TR-2026-02, 2026.
15. Schulcz, A. *Static linearity of the PaddleMate Pro strain-gauge paddle-shaft sensor on a carbon dragon-boat paddle, 5–20 kg*. PaddleMate Technical Report PM-TR-2026-03, 2026.
16. Schulcz, A. *Effect of shaft temperature on the sensitivity of the PaddleMate Pro strain-gauge sensor on a water-conditioned carbon tube, 2–35 °C*. PaddleMate Technical Report PM-TR-2026-04, 2026.
17. Hoffmann, K. *An Introduction to Measurements using Strain Gages*; Hottinger Baldwin Messtechnik: Darmstadt, 1989.
18. Gere, J.M.; Goodno, B.J. *Mechanics of Materials*, 8th ed.; Cengage Learning: Stamford, CT, 2013.