PIV Pre-Experiment Checklist

A Worked Guide for Wind Tunnel Measurements

Aerodynamics & Flight Mechanics | University of Southampton
M. T. Taylor *(June 20, 2026)*

This guide walks through the critical decisions and checks required before acquiring PIV data, using worked examples in a turbulent boundary layer and a bluff-body wake. Every number is traceable. A one-page summary checklist is provided at the end.

1 Dynamic Range: What Do You Need to Resolve?

Before entering the lab, frame your measurement as a dynamic range problem. Ask three questions:

1. **Largest scale** L: What sets the field of view? (Boundary layer thickness, chord, wake width.)
2. **Smallest scale** l: What must you resolve? (Viscous length, vortex core, shear layer thickness.)
3. **Dynamic range** L/l: Can your sensor deliver enough pixels to span L whilst resolving l?

The dynamic range determines whether your hardware can meet your science goals. If L/l exceeds the number of pixels available (divided by your interrogation window size), you must relax the resolution target, reduce the FOV, tile multiple cameras, or use ensemble PIV to access smaller windows.

2 Worked Example A: Turbulent Boundary Layer

2.1 Flow parameters and scales

Boundary layer thickness	$\delta = 150\,\text{mm}$	(measured or estimated)
Friction velocity	$u_\tau \approx 0.6\,\text{m/s}$	(from $Re_\tau = u_\tau \delta/\nu$)
Freestream velocity	$U_\infty \approx 15\,\text{m/s}$	(facility dependent)
Viscous length scale	$\ell_\nu = \nu/u_\tau \approx 25\,\mu\text{m}$	($\nu = 1.5 \times 10^{-5}\ \text{m}^2/\text{s}$)

Largest scale: $L = 1\delta = 150\,\text{mm}$ (wall-normal FOV to capture the full boundary layer). **Smallest scale:** target $\Delta y^+ \approx 10$, giving $l = 10 \times 25\,\mu\text{m} = 0.25\,\text{mm}$. **Dynamic range:** $L/l = 150/0.25 = 600$.

2.2 Camera and lens: can the hardware deliver?

Sensor	5328×4608 pixels	(long axis streamwise)
Sensor size	$14.6\,\text{mm} \times 12.6\,\text{mm}$	
Pixel pitch	$d_p = 14.6\,\text{mm}/5328 \approx 2.74\,\mu\text{m}$	
Lens	$f = 50\,\text{mm}$	
Working distance	$d_o \approx 550\,\text{mm}$	
Magnification	$M = f/(d_o - f) = 50/500 = 0.1$	

Each pixel images $\Delta x_{\text{px}} = d_p/M = 2.74/0.1 = 27.4\,\mu\text{m}$. Each pixel spans $27.4/25 \approx 1.1\ y^+$. The wall-normal FOV is $4608 \times 27.4 \times 10^{-3} \approx 126\,\text{mm} \approx 0.84\delta$.

Check: an $8\,\text{px}$ window spans $8 \times 27.4 = 219\,\mu\text{m}$, finer than $l = 0.25\,\text{mm}$ ✓. Spanning the full 1δ needs ~ 5500 pixels, which exceeds the 4608 available; accept 0.84δ coverage or relax to $\Delta y^+ \approx 12$.

2.3 Resolution summary

Window (px)	Physical size (mm)	Δy^+	Use case
32×32	0.88×0.88	35	Standard instantaneous
16×16	0.44×0.44	18	Fine instantaneous
8×24	0.22×0.66	9×26	Elongated wall-normal (inst.)
8×8	0.22×0.22	9	Ensemble PIV

Elongated windows: 8×24 explained

In wall-bounded flows, velocity gradients are steep in y but mild in x. An 8×24 window (8 px wall-normal, 24 px streamwise) exploits streamwise homogeneity to maintain correlation signal strength whilst maximising wall-normal resolution. This is standard practice for resolving the near-wall region without ensemble averaging.

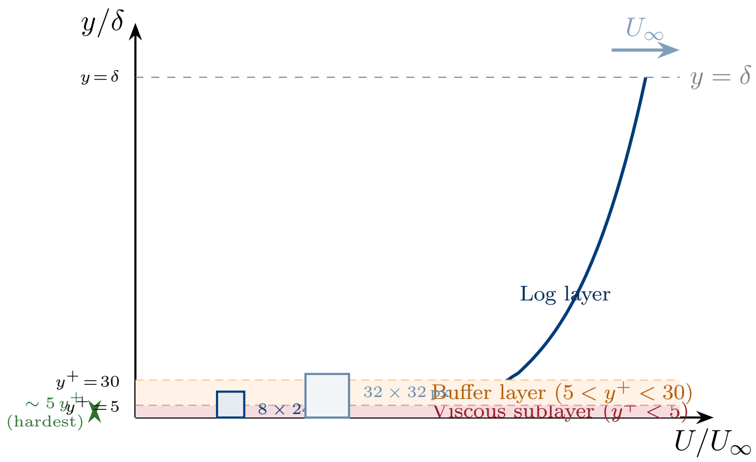

Figure 1: Schematic mean velocity profile with near-wall regions and representative interrogation windows.

3 Worked Example B: Bluff-Body Wake

Body width	$W = 100\,\text{mm}$	Desired FOV	$L = 5W = 500\,\text{mm}$
Smallest scale	$l = 0.02W = 2\,\text{mm}$	Dynamic range	$L/l = 250$

(50 independent vectors across a vortex of diameter $\sim W$)

Using the 25 MP CX1: a 16×16 window spanning $l = 2\,\text{mm}$ gives $125\,\mu\text{m/px}$. Spanning $L = 500\,\text{mm}$ requires $500/0.125 = 4000$ pixels, which fits within 5328. This measurement is feasible with a single camera.

Using a 4 MP high-speed camera (2560×1600, $d_p = 10\,\mu\text{m}$): the same magnification requires 4000 pixels, exceeding the sensor. Reducing L to $\sim 4W$ and relaxing l to $2.5\,\text{mm}$ ($156\,\mu\text{m/px}$) gives $400/0.156 \approx 2560$ pixels, fitting exactly on the long axis. This illustrates why the dynamic range check must happen *before* entering the lab.

4 Particle Image Diameter and f-Number

With $d_p = 2.74\,\mu\text{m}$, **particle images must span at least 2–3 pixels** in diameter on the sensor. If particles appear smaller than $\sim 2\,\text{px}$, the sub-pixel Gaussian peak-fit is aliased, producing **peak-locking**: displacement histograms cluster around integer pixel values, corrupting velocity statistics.

(sometimes this is unavoidable) The particle image diameter on the sensor is

$$d_\tau = \sqrt{(M\, d_{\text{phys}})^2 + d_{\text{diff}}^2}\,, \quad \text{where} \quad d_{\text{diff}} = 2.44\,\lambda\,(1+M)\,f_\# \,. \tag{1}$$

For $\lambda = 532\,\text{nm}$ (Nd:YAG), $M = 0.1$:

$f_\#$	d_{diff} (µm)	d_τ/d_p (pixels)	**Verdict**
$f/2.8$	4.0	1.5	Too small: peak-locking
$f/4$	5.7	2.1	Marginal
$f/5.6$	8.0	2.9	Good: $d_\tau \approx 3\,\text{px}$
$f/8$	11.4	4.2	Safe but costs light

Peak-locking is invisible unless you check

Peak-locking does not produce obviously wrong vector fields. It corrupts PDFs and turbulence statistics silently. Always check displacement histograms for spikes at integer values. For 2.74 µm pixel cameras, **do not use a wide-open aperture**: $f/5.6$ is a sensible starting point.

5 Pulse Separation Δt

The time between laser pulses A and B determines particle displacement in pixels. **Target ∼10 pixels** in the dominant flow direction ($\approx 1/4$ of a 32×32 window):

$$\Delta t = \frac{N_{\text{target}} \cdot d_p}{M \cdot U_{\max}}\,. \tag{2}$$

For $N_{\text{target}} = 10\,\text{px}$, $d_p = 2.74\,\mu\text{m}$, $M = 0.1$, $U_{\max} = U_\infty = 15\,\text{m/s}$:

$$\Delta t = \frac{10 \times 2.74 \times 10^{-6}}{0.1 \times 15} = \frac{27.4 \times 10^{-6}}{1.5} \approx 18\,\mu\text{s}\,.$$

Out-of-plane pair loss

Particles moving through the laser sheet between frames are lost. Out-of-plane displacement should not exceed $\sim 25\%$ of the sheet thickness. For a 1 mm sheet: $\Delta z_{\max}/\Delta t = 0.25/18 \approx 14\,\text{m/s}$. In a boundary layer at 15 m/s, spanwise fluctuations ($w' \approx 0.05 U_\infty \approx 0.75\,\text{m/s}$) are well within tolerance. For highly three-dimensional flows, revisit this carefully.

6 Instantaneous vs. Ensemble PIV

These are complementary, not competing, techniques.

Instantaneous PIV computes a displacement from each image pair independently, capturing time-resolved dynamics: ejections, sweeps, and coherent structures. For wall-bounded turbulence, $8 \times 24\,\text{px}$ windows represent a practical limit, giving $\Delta y^+ \approx 10$ in this facility.

Ensemble PIV averages correlation planes across many image pairs before extracting the displacement peak, suppressing noise and enabling windows down to $8 \times 8\,\text{px}$ with high confidence. Temporal information is sacrificed. For mean profiles and Reynolds stress estimation, ensemble PIV recovers energy that instantaneous methods miss at equivalent window sizes.

Use both

Acquire time-resolved data. Process it twice: once with instantaneous cross-correlation for dynamics, once with ensemble averaging for converged mean fields at higher resolution. The same raw images serve both purposes.

7 How Many Images to Acquire

Statistical convergence determines your dataset size. The number of independent image pairs N must be large enough that statistical uncertainty falls below measurement uncertainty (~ 0.1 px, i.e. $\sim 1\%$ at 10 px displacement). At 95% confidence:

Quantity	Criterion (95% confidence)	Typical N
Mean velocity	$1.96 \dfrac{\sigma_U}{\sqrt{N}} < 0.01 \overline{U}$	~ 400
Variance / Reynolds stress	$1.96 \dfrac{\sqrt{\overline{u'^4}}}{\sqrt{N}} < 0.01 \sigma_U$	$\sim 2\,500$
Higher-order statistics	$1.96 \dfrac{\sigma_B}{\sqrt{N}} < 0.01 B$	$> 10\,000$

Rule of thumb: 400 pairs for mean profiles, 2 500 for Reynolds stresses, 10 000+ for skewness or spectra. If using ensemble PIV, the same N applies to the ensemble. Always acquire more than you think you need; you cannot go back.

8 Laser Sheet

The $1/e^2$ sheet thickness is typically 1 mm–2 mm for boundary layer work. The sheet must be thick enough to retain particles between frames (see §5) but thin enough that the measurement volume is well defined.

Verify that **both A and B pulses** fire with even intensity across the sheet. A weak B pulse or misaligned dual-cavity laser halves your effective seeding density and produces systematically biased displacements. Fire both pulses at a card to visually confirm.

9 Camera Alignment and Reflections

Keep the sensor orthogonal to the laser sheet for in-plane (2C) PIV. Off-axis viewing introduces perspective distortion that biases displacement measurements.

Saturation destroys data three ways

(1) Camera health: sustained specular reflections can damage the sensor. **(2) Wall location:** a saturated band makes it impossible to determine the wall position; wall-normal profiles depend on this. **(3) Particle detection:** saturated regions eliminate particle signal entirely.

Offset the camera 1–2° above or below the wall plane so that specular reflections are directed away from the lens. For rough surfaces, fluorescent particles with optical filters can help. Find the minimum offset that controls reflections without sacrificing near-wall FOV.

10 Calibration

The board must be in the plane of the laser sheet, not merely near it.

The world origin is a deliberate, recorded decision. Calibration places the origin wherever you anchor it, and every velocity, every wall-normal position, every y^+ is reported in that frame. Tie it physically to whatever your science is referenced to (the wall, the model leading edge, the tunnel centreline) and **write the offset down**. In practice, include a reference marker at a known distance from the wall in the calibration image and record that distance. An origin you cannot locate in the physical model is an origin you cannot publish against.

Ruler in the image

Place a small ruler or scale bar in the calibration images alongside the board. It provides an independent cross-check of the calibration magnification.

Fill the frame to avoid silent extrapolation

Calibration is only trustworthy *inside the cloud of detected dots*. Outside it the model invents coordinates. A polynomial fit returns a finite, plausible, **wrong** value with no warning (never a `NaN`, never a flag), and a pinhole model degrades nearly as fast once you leave the constrained region. Spread the board so the dots cover most of the field of view, and ensure the **union of dot positions across all views encloses your entire measurement region**. Coverage is usually lost in the corners and at the near-wall edge, exactly where extrapolation does the most damage.

For multiple cameras, ensure sufficient overlap and identify a feature visible in both views for registration.

10.1 Camera model: pinhole vs. polynomial

Two model families map image pixels to world coordinates. A *polynomial* model fits a smooth function (typically 3rd order) directly from pixel position to world position. A *pinhole* model is a physical camera model (focal length, principal point, lens distortion, and the camera's position and orientation) constrained to obey perspective projection.

A pinhole model cannot be built from a single view. One image of a planar target constrains only a homography: it fixes the in-plane mapping *on that plane* but cannot separate focal length from working distance, nor recover how the projection changes with depth. The intrinsics are under-determined and out-of-plane (w) reconstruction is unsupported. **At least three views at different depths and orientations are required**, each one constraining the perspective geometry until the full 3D model is pinned down. Spreading the board to the edges of the frame matters here too: the points near the image periphery are what constrain lens distortion, so a board clustered in the centre leaves the radial terms under-determined even across many views.

The first view should sit on the datum, the plane of the laser sheet, at $z = 0$. This anchors the world origin and the plane your velocities are referenced to. Subsequent views, translated along z and ideally slightly tilted, supply the depth information that gives the model its third dimension.

One pose is not enough

A single board pose yields only a 2D dewarp of that plane. If that is all you have, **use a polynomial model**, an honest 2D mapping. Do not fit a pinhole model to one view and trust its 3D output: the depth behaviour is invented, not measured. A model built from one plane is geometrically flat. It cannot represent how the imaging changes through the sheet, so w-velocity and self-calibration both rest on nothing. For any w-resolving measurement, acquire multiple views.

10.2 Stereo PIV: self-calibration

In stereo PIV the laser sheet position is part of the measurement. Calibration places world $z = 0$ wherever the board sat, but the sheet is never exactly there. It is offset and slightly tilted, so both cameras dewarp onto the wrong plane and disagree about where each particle lies. That residual is the *disparity*.

Self-calibration removes it using the particle images themselves (Wieneke): it measures the cross-camera disparity, solves for the true sheet position and tilt, and bakes the correction into the projection. This needs a model that knows how world position changes with depth, so **self-calibration is a pinhole operation**, relying on the multi-view perspective geometry the pinhole model carries. A single-plane polynomial has no depth sensitivity: it can detect the disparity but cannot interpret it as a sheet offset, so it cannot be self-calibrated.

Stereo is therefore pinhole. Build each camera from multiple views, run self-calibration to lock the sheet, and confirm the disparity has collapsed. Getting that depth model right is the single biggest lever on w-accuracy.

11 Image Preprocessing

Apply a **temporal minimum subtraction**: compute the pixel-wise minimum over 10–20 images and subtract it from each frame. This removes static background (fixed-pattern noise, reflections, hot pixels) without altering particle signal, and is particularly effective at reducing wall glare.

12 Seeding Density

Target: 4–6 clean particles per interrogation window. Fewer particles produce weak or spurious correlation peaks.

25 MP sensors are deceptive

A full-frame 25 MP image *looks* well seeded at full zoom-out. Always zoom into your smallest region of interest and overlay the interrogation grid. Count particles per window at multiple locations, particularly near the wall.

Seeding drift: Concentration changes during a run. Check periodically, not just at the start.

Diagnostic: If seeding appears low, rule out: laser/FOV misalignment; weak A or B pulse; lens aperture too far closed; sheet too thin.

13 Correlation Quality Checks

Process a test set (50–100 image pairs) and inspect results before committing to a full run.

Peak ratio $Q = P_{cc}/P_{ac}$ (cross-correlation peak over autocorrelation peak): $Q > 1.3$ is a common threshold. Examine the spatial distribution of low-Q vectors; clustering indicates a local problem.

NaN fraction should be below $\sim 5\%$ in the freestream. Some increase near the wall is expected.

Infilled vectors are the default

PIV software interpolates over rejected vectors before displaying the final field. A smooth-looking field can contain large regions of fabricated data. **Always** inspect the pre-infill field, the NaN mask, and the peak-ratio map.

Spatial uniformity: patchy failures, streaks of low Q, or asymmetric NaN distributions indicate optical or seeding problems to resolve before production data.

50% overlap between adjacent windows is standard, satisfying the Nyquist criterion and doubling output grid density. 75% overlap looks smoother but adds no independent information.

PIV Pre-Experiment Summary Checklist

Print this page and take it to the tunnel.

---

1. Before the Lab

☐ Define largest scale L and smallest scale l
☐ Check dynamic range L/l fits your sensor
☐ Calculate FOV from sensor, lens, d_o
☐ Calculate Δt for $\sim 10\,\text{px}$ displacement
☐ Choose $f_\#$ for d_τ
☐ Confirm out-of-plane loss $< 25\%$ of sheet
☐ Estimate N for target statistics

2. Optics Set-Up

☐ Camera sensor $\perp$ to laser sheet
☐ Camera offset 1–2° to reject wall glare
☐ No saturated pixels on test image
☐ Both A and B pulses fire evenly
☐ Sheet thickness set ($\sim 1\,\text{mm}$ for BL)
☐ Sheet aligned to camera FOV

3. Calibration

☐ Board **in the plane** of the laser sheet
☐ First calibration view on the datum ($z = 0$)
☐ Origin located & recorded (offset to wall / model)
☐ Dots span most of FOV; measurement region inside dot cloud
☐ Ruler visible in calibration image
☐ Pinhole: ≥ 3 views; 1 pose $\rightarrow$ polynomial only
☐ Polynomial: planar 2C only (no w, no self-cal)
☐ Multi-cam: overlap confirmed
☐ Stereo (pinhole): self-calibration run, disparity minimised
☐ Calibration residual acceptable

4. Seeding and Illumination

☐ 4–6 particles per smallest window
☐ Zoom in to check (do not trust full-frame)
☐ Check seeding at multiple FOV locations
☐ Re-check seeding periodically during run
☐ If seeding looks low: rule out laser/alignment

5. Image Quality

☐ Apply temporal min-filter (10–20 images)
☐ Particles visible near the wall after filtering
☐ No fixed bright artefacts remaining
☐ Particle images $\approx$ 2–3 px diameter

6. Processing Verification

☐ Process 50–100 test pairs
☐ Peak ratio $Q = P_{\text{cc}}/P_{\text{ac}} > 1.3$ spatially
☐ NaN fraction $< 5\%$ in freestream
☐ Inspect pre-infill vectors (not interpolated)
☐ Displacement histogram: no peak-locking
☐ Spatially uniform correlation quality
☐ 50% window overlap applied

7. Dataset Size

☐ ≥ 400 pairs for mean velocity
☐ $\geq 2\,500$ pairs for Reynolds stresses
☐ $\geq 10\,000$ pairs for higher-order stats

If any check fails, diagnose before acquiring production data. Time spent on set-up is always recovered in data quality.