


Wind Farm Flow Reconstruction and Control Competition

Intuitive Understanding

A wind turbine leaves a slower, messier patch of air behind it, and the machine in the next row has to work in that patch. Turn the front turbine a few degrees off the wind and its wake is pushed to one side, so it misses the one behind: the front machine loses a little and the farm gains more. Doing that for real means knowing what the air between the turbines is doing, from only a handful of sensors.



Steering the wake of the front turbine away from the one behind it.

Project Scope

A wind farm loses a great deal of the energy in the wind to itself. Turbines in the front row take the cleanest inflow and leave a slower, more turbulent wake for everything behind them, and the conventional answer - every turbine greedily maximising its own power - is demonstrably not the answer that maximises the power of the farm. Doing better requires knowing what the flow is actually doing between the turbines, and the only instruments a real farm has are a handful of point sensors.

The aim is to Reconstruct the Flow Field Across a Wind Farm from Sparse Sensor Data and Use It to Control the Turbines - rebuilding the continuous velocity field from a small number of discrete measurements, then using that field to set generator torque and yaw so that the power of the whole array, rather than of any one machine, is maximised. This is the combination of two



tasks that should be completed by sub-teams working closely together.

This project is run with the Department of Aeronautics at Imperial College London, alongside the ICONIC project, and it is unusual among our competitions in that it does not end in a simulation. The finals are run on hardware: six 1:250 scale instrumented turbines in the 10'x5' wind tunnel, in a marine atmospheric boundary layer, with each team's controller installed on the real rig and run live. Teams are also given a large body of existing experimental data - approximately 180 GB of particle image velocimetry measurements of a yawed turbine wake, with synchronous turbine loads - to develop against.

The following sections will break down the requirements and expectations of each task. However, it is important to remember that although these are outlined as separate tasks, they have significant reliance on information from each other - the controller in Subtask 2 can only act on the field that Subtask 1 gives it, at the rate and with the errors that Subtask 1 delivers, and the reconstruction should be built for the control problem rather than for its own sake.

Teams will be guided by members of academia on a regular basis throughout the year, and the wind tunnel work will be supervised. That guidance is part of the project rather than an optional extra: teams are expected to arrive at each session with specific questions and something to show.

Subtask 1 – Flow Field Reconstruction from Sparse Sensor Measurements

1.1.1 This task is to reconstruct the continuous velocity field around a wind turbine from a small number of discrete point measurements. Teams should fix early how many sensors they assume, and where they sit, and justify that choice against what a real installation could actually deliver rather than against what is convenient for the algorithm.

1.1.2 A large experimental dataset is provided and is the foundation of this task. It consists of planar PIV measurements of the wake of a 1:250 scale, three-bladed, 0.5 m diameter turbine at 0.38 m hub height, based on the Senvion 6M126 offshore machine, taken at a hub-height velocity of 7.8 ms⁻¹ in a marine atmospheric boundary layer consistent with ISO 19901-1:2015. Thirteen yaw angles from -30 to +30 degrees in 5 degree steps are included, in both a time-resolved form sampled at 500 Hz and a non-time-resolved form sampled below 10 Hz, on a grid with one vector every 1.847 mm over a field of view of 2D spanwise by 1.2D streamwise centred 2.7D downstream of the rotor. Turbine torque, axial thrust, and rotational speed are recorded synchronously at 1 kHz throughout.

1.1.3 Three further cases are provided and should be used deliberately rather than pooled in: a clean inflow case with no boundary layer and turbulence intensity below 0.1 %, a boundary-layer-only case in the forward-facing configuration, and a case at a reduced tip speed



ratio of 4. The clean case is the natural verification baseline and the reduced tip speed ratio case exists specifically to expose teams that have overfitted to a single operating point.

1.1.4 Strict separation of training and test data is mandatory and will be checked. Teams must hold out whole yaw angles, not random snapshots from within a case, and must report reconstruction error on yaw angles the model has never seen. Error quoted on data the model was fitted to is not a result.

1.1.5 The sensor model is a specific and mandatory sub-topic, and is where this task is most often done badly. A full-scale LiDAR returns a line-of-sight velocity component at roughly 1 Hz, sampled in a cone upstream of the rotor, which at 1:250 scale corresponds to 250 Hz on the rig. Teams must degrade the PIV data into what such an instrument would genuinely see - sparse locations, line-of-sight projection rather than true components, measurement noise, and latency - and must state the model they used. Selecting arbitrary points from a dense PIV grid and calling them sensors is not acceptable.

1.1.6 The reconstruction method itself is open: modal approaches (e.g. gappy POD), stochastic estimation, state estimation and Kalman filtering, and machine learning approaches (convolutional networks, physics-informed networks, operator learning) are all in scope. Whatever is chosen, at least one classical baseline must be implemented and reported alongside it, so that the benefit of the sophisticated method is visible rather than assumed.

1.1.7 Real-time capability is a requirement, not a refinement. The reconstruction has to run inside the control loop that Subtask 2 builds, so teams must report wall-clock inference time, the hardware it was measured on, and the resulting achievable update rate. Computational cost is explicitly part of the score: a reconstruction that is more accurate but too slow to act on is worth less than a coarser one that runs.

1.1.8 Uncertainty must be carried through to the output. Teams should attach a stated confidence to the reconstructed field, ideally spatially resolved, and pass it to Subtask 2 so that the controller can distinguish a region it knows about from one it is guessing at.

1.1.9 We would like to see evidence of verification of the reconstruction against published work on flow field estimation from sparse measurements, and against the analytical and empirical wake models in the literature (e.g. Jensen and Gaussian wake models) as a sanity check on the mean field before any unsteady claims are made.

1.1.10 The primary figure of merit is reconstruction accuracy on unseen yaw angles per unit of computational cost. However, there is also significant weighting on other things such as:

a) Modelling approach taken – the thoroughness, quality, and compliance with good coding practice such as PEP8.

- b) Project management and process – planning, recording of information, communication, risk capture etc.
- c) The consideration and justification of the level of modelling fidelity used. What assumptions have been used and how have they been captured? This is a powerful skill.
- d) Key points of learning during the task and how this is recorded. We want to see what you discover, how you explore options, what you learnt and how you then applied this.
- e) Novelty of the reconstruction approach or of the sensor placement strategy.
- f) First phase presentation.
- g) Final presentation and summary report.
- h) Cross-functionality working and integration between teams.
- i) Depth of research.
- j) Realism of the sensor model, and of its noise and latency assumptions.

1.1.11 The scope outlined isn't expected to be completed by the first deadline – this is for the finals. Progress up to the first phase is what will be assessed in the initial project down-selection. Students should utilise their relevant skills and interests to complete tasks they think will be useful and interesting, even if not explicitly listed here.

Subtask 2 – Turbine Control Optimisation and the Wind Tunnel Final

2.1.1 This task turns the reconstructed field into a controller that sets generator torque and yaw angle across the array, and it ends with that controller running on real hardware. The team should work closely with the Subtask 1 team throughout, and should treat the reconstruction's update rate and error as design constraints rather than as someone else's problem.

2.1.2 A baseline is mandatory. The rig's turbines operate in region 2 under maximum power point tracking, with optimum torque set as the square of angular velocity times a constant obtained from a tip speed ratio ramp. Teams must implement that greedy baseline, measure it, and express every subsequent result as a gain over it. A controller that cannot be compared to greedy control cannot be credited.

2.1.3 The physical envelope of the rig is fixed and must be respected in the design. Torque is commanded through the generator winding current of a 200 W brushed DC motor in generator configuration; the rotary torque sensor is rated to 0.5 Nm and the load cell to 30 N; the yaw device is a stepper-driven worm gear capable of full 360 degree rotation at up to 40 degrees per second. A control strategy that commands outside these limits will not be run.

2.1.4 A specific and mandatory sub-topic is the cost of yaw. Steering a wake is not free: the



manoeuvre costs energy and actuator wear, it costs power at the turbine being yawed, and it changes the unsteady thrust and buffeting loads the structure sees, which the load cell measures directly. Teams must include an explicit penalty for yaw activity and for load excursions in their objective function, state it, and justify the weighting rather than tuning it until the answer looks good.

2.1.5 Computational cost must sit in the same objective. Teams should state the loop rate their controller achieves and the latency between a measurement arriving and an actuator command leaving, and should show what performance they gave up to meet it. Power that a controller could theoretically capture but cannot compute in time is not power.

2.1.6 Array-level rather than turbine-level performance is what is being optimised. The finals rig is six turbines on a 5.6 m turntable, arranged three deep in the streamwise direction and two across, spaced 5 rotor diameters streamwise and 4 spanwise, so wake steering and induction control across the array are both in scope. The reported figure must be total array power.

2.1.7 Robustness is assessed explicitly and is not treated as a bonus. Teams must show how the controller behaves when a sensor drops out, when the reconstruction is confidently wrong, and when the inflow moves outside the conditions it was developed against, and must define a safe fallback state the system reverts to rather than allowing an unbounded command.

2.1.8 The task requires a literature review of existing wind farm flow control work (e.g. the FLORIS framework, published wake steering field trials, and reinforcement learning approaches to farm control) to be completed, followed by development of the controller in stages as before, in simulation first and then against the provided experimental data.

2.1.9 Delivery for the finals is a hard requirement. The controller must be handed over as a defined, documented interface that can be installed on the rig's LabVIEW and NI PXIe host by the handover deadline, must have deterministic timing, and must run unattended for the duration of a tunnel run. Teams should design for this interface from the start; a controller that cannot be installed in time will be assessed on its analysis alone and cannot score on hardware.

2.1.10 The main scoring criteria for this subtask is the measured gain in total array power over the greedy baseline in the wind tunnel, net of the yaw, load, and computational penalties - measured on the rig, not predicted in simulation. A modest gain that survives contact with the hardware is scored more highly than a large one that only exists in a model. However, the other criteria listed for Subtask 1 are also of importance.



Timeline

This competition has two phases: the first ends at the beginning of January, wherein each team presents their reconstruction and control approach to judges online; from here, 2 teams are selected from each university for the in-person finals. Members from teams which did not get selected can apply to join the two selected teams for the remainder of the competition.

In January a mid-project review will be held, and for this the only deliverable is a presentation outlining the progress made thus far. Alongside this, teams on this project meet their academic mentors on a regular basis throughout the year; these sessions are part of the project rather than an optional extra, and teams are expected to come to them with specific questions and something to show.

The finals are run live in the 10'x5' wind tunnel at Imperial College London, and are the reason the handover deadline a week earlier is a hard one - a controller that cannot be installed on the rig in time cannot be run, and cannot be scored on hardware. Each team gets an equal, fixed allocation of tunnel time under identical inflow conditions, and the controller runs unattended for the duration of each run.

The finals also consist of a 30 minute presentation, as well as the submission of a report; this has a limit of 8 pages (title, contents, and reference pages not included, appendix not included but will not be included in the scoring) and should cover the concept, the method of development, evidence of initial V&V, and key lessons learned. This should be delivered 2 weeks before the final presentation. We recommend including short snippets of code in the appendix for sections of interest. The report should cover both tasks.

21 September 2026	Applications Open
25 October 2026	Applications Close
25 October 2026	Teams selected and formed
From 2 November 2026	Fortnightly academic mentor sessions begin
23-27 November 2026	Academic Mentor Check-In (Feedback given)
11-15 January 2027	Phase 1 Presentations (Feedback given)
1-5 February 2027	Academic Mentor Check-In (Feedback given)
12 February 2027	Controller handover deadline
19 February 2027	Finals, in the 10'x5' wind tunnel at Imperial

2026/2027



	College London
--	----------------