

PIV-in-the-Loop Optimization of Active Turbulence Grids for Tailored Inflow Generation

Alexander Gehrke¹, Karunmay Aggarwal¹, Rehaan Irani¹, Evan M. diVittorio¹, Kenneth S. Breuer^{1,*}

1: Center for Fluid Mechanics, School of Engineering, Brown University, United States of America

*Corresponding author: kenneth_breuer@brown.edu

Keywords: Particle image velocimetry, active turbulence grid, closed-loop experiments, experimental flow control.

ABSTRACT

Active turbulence grids provide a flexible means of generating tailored turbulent inflow conditions, but their large number of actuation degrees of freedom makes systematic tuning difficult. Conventional approaches often rely on prescribed motion patterns, manual adjustment, or post hoc characterization, which limits their scalability when specific mean-flow and turbulence profiles are desired. In this work, we present a closed-loop experimental framework in which planar particle image velocimetry measurements are used directly as feedback for automated exploration and optimization of an active turbulence grid.

The method combines a reduced-order, mode-based parameterization of grid motion with Gaussian-process regression surrogate models. Instead of treating each grid axis as an independent actuator, the grid kinematics are described using a small set of physically interpretable parameters controlling synchronous motion and vertical gradients in pitch amplitude and offset. The experimental procedure is divided into two stages. First, an active-learning exploration phase selects new actuation states to reduce global uncertainty in the surrogate models of the measured flow statistics. For each candidate state, PIV measurements are acquired downstream of the grid, processed automatically, and reduced to scalar metrics describing turbulence intensity, mean-velocity gradients, turbulence-intensity gradients, homogeneity, and an in-plane anisotropy proxy. These measurements are added to the database and used to update the surrogate models. In a second stage, the learned surrogate models are used to identify actuation parameters that best match prescribed turbulent inflow targets, followed by additional experimental iterations to evaluate and refine the selected conditions.

The framework is demonstrated for three representative turbulent inflow objectives: homogeneous isotropic turbulence with prescribed turbulence intensity, a prescribed mean-velocity gradient at constant turbulence intensity, and a prescribed turbulence-intensity gradient at uniform mean velocity. The target objective separates desired flow quantities from secondary penalty quantities, allowing the optimization to prioritize matching the prescribed flow state before reducing undesirable features such as spatial inhomogeneity or in-plane stress imbalance.

The results show that PIV-in-the-loop active learning and surrogate-guided optimization can identify active-grid actuation strategies that generate distinct prescribed turbulent inflow profiles within a low-dimensional control space. By using spatially resolved velocity-field measurements rather than indirect proxy signals, the approach provides direct control over flow quantities relevant to downstream aerodynamic and fluid–structure-interaction experiments. This methodology offers a scalable route toward on-demand generation of tailored turbulent inflows in wind-tunnel facilities and illustrates the potential of integrating quantitative imaging diagnostics with automated experimental design.

1. Introduction

Controlled turbulent inflow conditions are essential for experimental studies in aerodynamics, fluid–structure interaction, and energy systems, where reproducibility and precise tailoring of flow statistics are required. Applications range from the investigation of turbulence-induced loads on aerodynamic structures to the validation of numerical models and reduced-order descriptions of complex flows. Traditionally, such studies rely on passive turbulence generators or manually tuned inflow devices, which can provide repeatable turbulence conditions but offer limited flexibility once installed.

Active turbulence grids provide a powerful alternative by enabling dynamic control of turbulence properties through prescribed grid motions. Since the introduction of active-grid turbulence generation by Makita (1991), active grids have been used to produce high-Reynolds-number turbulent flows and to study the effect of grid motion on downstream turbulence statistics. For example, Mydlarski & Warhaft (1996) used an active grid to investigate the onset of high-Reynolds-number grid-generated turbulence, while subsequent work has examined how actuation parameters and grid geometry influence the generated flow. Hearst & Lavoie (2015) showed that active-grid operating conditions can strongly affect both the magnitude and structure of the turbulence, highlighting the flexibility of active grids as well as the need for systematic actuation strategies.

Beyond approximately homogeneous turbulence, active grids can also generate more structured turbulent inflow states. Cekli & Van De Water (2010) demonstrated that active grids can be used to generate homogeneous shear turbulence and trigger turbulent boundary-layer development. This capability is particularly attractive for experiments in which aerodynamic models or fluid–structure-interaction systems must be exposed to prescribed mean-flow and turbulence profiles. However, the large number of actuation degrees of freedom inherent to active grids makes the systematic identification of suitable motion patterns difficult. As a result, active-grid studies often rely on prescribed protocols, manual adjustment, or post hoc parametric characterization, which limits scalability when specific target flow states are desired.

In parallel, optimization and control methodologies have increasingly been applied to experimental fluid mechanics. Closed-loop approaches have demonstrated that experimental measurements can be used directly as feedback for flow-control and optimization problems, including aerodynamic optimization based on force measurements (Gehrke & Mulleners, 2021) and separation control using online PIV and machine learning (Gautier et al., 2015). More recently, Stefan-Zavala et al. (2026) introduced a data-driven framework for prescribing on-demand flow states in fan-array wind tunnels, illustrating the potential of combining high-dimensional actuation systems with surrogate modeling and inverse design.

Despite these advances, systematic measurement-driven optimization of active turbulence grids using spatially resolved flow-field data remains less developed. One challenge is that the mapping




Figure 1. (a) Active turbulence grid with 7 horizontal and 8 vertical axes. (b) Schematic of the turbulence grid and PIV field of view in the wind tunnel (not to scale).

from active-grid motion to downstream flow statistics is nonlinear, facility-dependent, and expensive to sample experimentally. A second challenge is that PIV-based feedback requires the acquisition and processing of large image sets at each candidate actuation state. Recent developments in rapid PIV processing algorithms (Bollt et al., 2025) reduce this barrier and enable velocity-field statistics to be extracted on time scales compatible with closed-loop experimental optimization.

In this work, we present an automated, closed-loop framework for controlling an active turbulence grid using planar PIV measurements as direct feedback. The method combines a reduced-order, mode-based parameterization of grid actuation with Gaussian-process surrogate models. The experimental procedure consists of an uncertainty-guided exploration stage, in which new actuation states are selected to improve the surrogate representation of the measured flow metrics, followed by a target-oriented exploitation stage, in which the learned models are used to identify actuation states that best match prescribed turbulent inflow objectives. The approach is demonstrated for three canonical target cases: homogeneous isotropic turbulence with prescribed turbulence intensity, a prescribed mean velocity gradient at constant turbulence intensity, and a prescribed turbulence intensity gradient at constant mean velocity. The results show that PIV-in-the-loop active learning and surrogate-guided optimization can reproducibly identify active-grid actuation strategies that generate tailored turbulent inflow profiles within a low-dimensional control space.

2. Methods

2.1. Experimental Configuration

The active turbulence grid used in this study consists of 15 individually controllable axes, arranged as 7 horizontal and 8 vertical shafts. Each horizontal axis carries 7, and each vertical axis 8, square panels of width $l = 0.087$ m. All axes are actuated by servo motors synchronized through a set of




Figure 2. (left) Data acquisition of time-resolved PIV images, (middle) Time-averaged flow and turbulence characteristic fields, (right) Scalar metrics spatially-averaged over entire field of view.

motion controller boards, allowing for fully prescribed, time-dependent grid motions.

The grid is installed at the entrance of the test section (Figure 1a) of the AFAM wind tunnel at Brown University, which has a square cross-section of 1.20 m. For all experiments presented here, the tunnel was operated at a nominal free-stream velocity of $U_\infty = 5.75$ m/s, leading to a panel width Reynolds number of

$$Re_l = \frac{\rho \bar{U} l}{\mu} = 3.33 \times 10^4. \quad (1)$$

To evaluate the flow states generated by different grid actuation patterns, planar particle image velocimetry (PIV) measurements were performed at a streamwise distance of 3.25 m downstream of the grid. The field of view covers an area of $0.30 \text{ m} \times 0.30 \text{ m}$ and is positioned slightly above the tunnel centerline in the vertical (y) direction, intersecting the fourth horizontal grid row. In the spanwise (z) direction, the field of view is centered within the test section. Illumination is provided by a double-head Nd:YAG laser (DM40-DH, Photonics), and image pairs are recorded using a high-speed camera (Fastcam Nova R2, Photron) at a sampling frequency of $f_{\text{PIV}} = 53$ Hz, ensuring statistical independence between successive velocity fields.

Image processing was performed using the *LaVision* software package with an interrogation window size of $48 \text{ px} \times 48 \text{ px}$ and an overlap of 75 %. For each grid actuation setting, 400 image pairs were recorded, which was found sufficient to obtain converged mean and second-order flow statistics, as verified in preliminary convergence tests.

2.2. Turbulence Metric Extraction

Three primary quantities are used to characterize the turbulent flow. The turbulence intensity TI is defined locally as the root-mean-square of velocity fluctuations normalized by the mean velocity

magnitude,

$$TI(x, y) = \frac{\sqrt{\frac{1}{2} (\overline{u'(x, y, t)^2} + \overline{v'(x, y, t)^2})}}{\sqrt{\bar{u}(x, y)^2 + \bar{v}(x, y)^2}}. \quad (2)$$

Flow homogeneity is quantified through spatial gradients of the mean velocity and turbulence intensity, dU/dy and dTI/dy , as well as by the coefficient of variation of the turbulence-intensity field,

$$CV_{TI} = \frac{\sigma_{\Omega}(TI(x, y))}{\mu_{\Omega}(TI(x, y))}. \quad (3)$$

Here, μ_{Ω} and σ_{Ω} denote the spatial mean and standard deviation over the measurement domain Ω , respectively. Anisotropy is motivated by the normalized Reynolds stress anisotropy tensor

$$a_{ij} = \frac{R_{ij}}{2k} - \frac{1}{3}\delta_{ij}, \quad (4)$$

where $R_{ij} = \overline{u'_i u'_j}$, k is the turbulent kinetic energy, and δ_{ij} is the Kronecker delta. Since the present planar PIV measurements provide only the in-plane velocity components, the scalar metric used for optimization is defined as the normalized imbalance between the two measured normal Reynolds stresses,

$$\mathcal{A}(x, y) = \frac{|\overline{u'^2}(x, y) - \overline{v'^2}(x, y)|}{\overline{u'^2}(x, y) + \overline{v'^2}(x, y)}. \quad (5)$$

The global metric scores are then obtained by spatially averaging the quantities over the measurement domain into a scalar value.

2.3. Grid Actuation Patterns

Active turbulence grids possess a large number of degrees of freedom when each axis is controlled independently. Assuming a sinusoidal pitching motion of each axis,

$$\phi_i(t) = \phi_{0,i} + \hat{\phi}_i \sin(2\pi f_i t + \varphi_i), \quad (6)$$

where $\hat{\phi}_i$ is the pitch amplitude, f_i the pitching frequency, φ_i a phase offset, and $\phi_{0,i}$ an angular bias, each motor introduces four independent parameters. For the present grid with 15 axes, this results in a total of 60 free parameters, rendering direct control impractical, as also noted by Cekli & Van De Water (2010).

To reduce the dimensionality of the control space while retaining expressive actuation patterns, we introduce a mode-based description of the grid motion. Individual motor commands are constructed as linear combinations of prescribed spatial modes,

$$\Phi(t) = \sum_{m=1}^M c_m(t) \psi_m, \quad (7)$$




Figure 3. (a) Synchronous mode and (b) gradient mode pattern used for grid control parameterization.

where $\Phi(t)$ represents the instantaneous angular positions of all grid axes, ψ_m are prescribed mode shapes across the grid, and $c_m(t)$ are time-dependent modal amplitudes. This approach enables the generation of spatially coherent actuation patterns using only a small number of parameters.

The simplest mode employed here is a synchronous mode, in which all axes follow the same periodic pitch trajectory (Figure 3a). For practical implementation and optimization, this trajectory was parameterized in terms of its angular endpoints rather than directly in terms of the sinusoidal angular bias and amplitude. The upper endpoint was denoted by α , while the lower endpoint was defined as

$$\beta = r_\beta \alpha, \quad 0 \leq r_\beta \leq 1, \quad (8)$$

where r_β sets the lower endpoint relative to the upper endpoint. The corresponding angular excursion is

$$2\hat{\phi} = \alpha - \beta = \alpha(1 - r_\beta). \quad (9)$$

Thus, the base pitch motion is bounded by

$$\phi_{\min} = \beta = r_\beta \alpha, \quad \phi_{\max} = \alpha. \quad (10)$$

Equivalently, if the endpoint motion is written in the sinusoidal form of Equation (6), the corresponding angular bias and pitch amplitude are

$$\phi_0 = \frac{\alpha + \beta}{2} = \frac{\alpha}{2} (1 + r_\beta), \quad \hat{\phi} = \frac{\alpha - \beta}{2} = \frac{\alpha}{2} (1 - r_\beta). \quad (11)$$

Therefore, α is not the sinusoidal pitch amplitude, but the maximum pitch angle reached during the cycle. The parameter r_β controls the balance between angular bias and oscillatory motion. For




Figure 4. Endpoint parameterization of the pitch actuation. The control parameters α and r_β define the upper and lower pitch endpoints, with $\beta = r_\beta \alpha$. The equivalent sinusoidal representation has angular bias $\phi_0 = (\alpha + \beta)/2$ and pitch amplitude $\hat{\phi} = (\alpha - \beta)/2$.

$r_\beta = 0$, the axis pitches between 0 and α , giving the largest angular excursion for a given α . For $r_\beta = 1$, the angular excursion is zero and the grid remains at the static pitch angle $\phi = \alpha$.

A second mode introduces a linear gradient in the angular excursion and angular bias across the horizontal axes only, while the vertical axes retain the base synchronous motion (Figure 3b). Let $\eta_i \in [-1, 1]$ denote the normalized vertical position of horizontal axis i , with $\eta_i = 0$ at the center of the grid. The implemented gradient parameterization is

$$\Delta\phi_i = \Delta\phi + g_A\eta_i, \quad o_i = \beta + g_0\eta_i, \quad (12)$$

where g_A and g_0 are the amplitude-gradient and offset-gradient parameters, respectively. These quantities represent center-to-edge deviations; the corresponding end-to-end changes across the horizontal axes are therefore $2g_A$ and $2g_0$. The synchronous actuation mode is recovered when $g_A = g_0 = 0$, for which all axes follow the same pitch trajectory. Nonzero values of g_A and g_0 enable controlled vertical variations in the oscillatory excursion and mean pitch of the horizontal axes, allowing the grid to generate mean-velocity or turbulence-intensity gradients across the test section.

Using this modal decomposition, the control space is reduced to four parameters for the synchronous mode, and gradient mode is included, $(\alpha, r_\beta, g_A, g_0)$. The actuation frequency is kept fixed at $f = 6$ Hz. The control variables, bounds, and the most important derived physical quantities are summarized in Table 1.

This parameterization samples only one representative branch of sign-mirrored base actuation states. Because the base endpoints are defined with $0 \leq \beta \leq \alpha$, trajectories whose angular range straddles zero, for example -10° to 20° , are not treated as independent base motions. In a symmetric facility and inflow, such motions are expected to be equivalent to their mirrored counterparts after reflection about the grid centerline. The reduced parameterization therefore avoids redundant exploration of mirror-symmetric configurations while retaining simple optimization bounds.




Figure 5. Schematic representation of the Gaussian-process surrogate model integration. The active-grid control parameters are first mapped to the normalized input vector ξ_* , which is used as the common input to a set of independent scalar-output Gaussian-process models. Each model is trained from the accumulated experimental data X, y_m and predicts the posterior mean $\mu_m(\xi_*)$ and uncertainty $\sigma_m(\xi_*)$ for one flow metric.

2.4. Gaussian-process surrogate modeling

The relationship between the active-grid control parameters and the resulting flow statistics is nonlinear and not known a priori. We therefore approximate this relationship using Gaussian-process regression (GPR) surrogate models (Rasmussen & Williams, 2006). GPR is well suited for the present measurement-driven framework because it provides both a predicted mean response and a predictive uncertainty for each unmeasured actuation state. The former is used for target-oriented optimization, while the latter is used to guide the active-learning exploration described in Section 2.5.

Let $p \in \mathbb{R}^d$ denote the vector of active grid-control parameters which are described in Section 2.3. Before model training, each physical input coordinate is mapped to a normalized design coordi-

Table 1. Control variables used to parameterize the active-grid actuation. The gradient parameters g_A and g_0 are center-to-edge deviations; the corresponding end-to-end changes across the horizontal axes are $2g_A$ and $2g_0$, respectively.

Quantity	Description	Bounds or range
α	Upper endpoint of the base motion	$0^\circ \leq \alpha \leq 59.5^\circ$
r_β	Relative lower endpoint, β/α	$0 \leq r_\beta \leq 1$
g_A	Centre-to-edge gradient in angular excursion	$-45^\circ \leq g_A \leq 45^\circ$
g_0	Centre-to-edge gradient in offset angle	$-45^\circ \leq g_0 \leq 45^\circ$
$\hat{\phi}$	Base sinusoidal amplitude, $\hat{\phi} = (\alpha - \beta)/2$	$0^\circ \leq \hat{\phi} \leq 29.75^\circ$
ϕ_0	Base angular bias, $\phi_0 = \alpha(1 + r_\beta)/2$	$\alpha/2 \leq \phi_0 \leq \alpha$

nate,

$$\xi_k = \frac{x_k - x_{k,\min}}{x_{k,\max} - x_{k,\min}}, \quad 0 \leq \xi_k \leq 1, \quad (13)$$

where $x_{k,\min}$ and $x_{k,\max}$ are the lower and upper bounds of the k -th control parameter. The normalized input vector is denoted by ξ . This normalization removes differences in physical units and parameter ranges, so that all input dimensions are treated on comparable scales during model training and exploration.

For the m -th scalar flow metric, the available training data after n measurements are written as the input matrix $X = [\xi_1, \dots, \xi_n]^T$ and the measured output vector $\mathbf{y}_m = [y_{1,m}, \dots, y_{n,m}]^T$. In the present study, separate surrogate models are trained for the mean turbulence intensity, the mean in-plane anisotropy proxy, the coefficients of variation of mean velocity and turbulence intensity, and the signed vertical slopes of the mean-velocity and turbulence-intensity profiles. Training one independent scalar-output model for each quantity keeps the surrogate-model structure simple and allows the prediction uncertainty to be evaluated separately for each flow metric (Figure 5).

Each measured output is represented as

$$y_m(\xi) = f_m(\xi) + \epsilon_m, \quad (14)$$

where f_m is the latent response function of the m -th metric and ϵ_m represents measurement noise and unresolved experimental variability. The latent response is modeled as a Gaussian process,

$$f_m(\xi) \sim \mathcal{GP}(\mu_m(\xi), k_m(\xi, \xi')). \quad (15)$$

This formulation follows the common use of stochastic-process surrogate models for response prediction in expensive experimental or numerical design spaces Sacks et al. (1989).

The covariance function is chosen as an automatic-relevance-determination Matérn 5/2 kernel,

$$k_m(\xi, \xi') = \sigma_{f,m}^2 \left(1 + \sqrt{5}r_m + \frac{5}{3}r_m^2 \right) \exp(-\sqrt{5}r_m), \quad (16)$$

with

$$r_m^2 = \sum_{k=1}^d \left(\frac{\xi_k - \xi'_k}{\ell_{m,k}} \right)^2. \quad (17)$$

Here, $\ell_{m,k}$ is the characteristic length scale of the m -th output metric in the k -th input direction, and $\sigma_{f,m}$ is the signal amplitude. The Matérn class is commonly used in Gaussian-process surrogate modeling because it provides a flexible family of covariance functions with controlled smoothness (Rasmussen & Williams, 2006). The Matérn 5/2 kernel represents a pragmatic choice for the present experimental response surfaces: it assumes that the measured flow metrics vary smoothly with the actuation parameters, but it is less restrictive than the squared-exponential kernel, which

imposes infinitely smooth sample functions. This behavior is appropriate for PIV-derived turbulence metrics, which are expected to change continuously with the grid motion but may contain measurement noise, finite-sample variability, and nonlinear dependence on the actuation parameters. The automatic-relevance-determination length scales allow each model to identify the input directions to which a given flow metric is most sensitive. We do not expect the main conclusions to depend sensitively on this particular kernel choice; other stationary kernels with comparable smoothness would likely provide similar surrogate models for the present low-dimensional design space. The hyperparameters are estimated from the experimental measurements each time the surrogate models are updated.

For a new candidate actuation state ξ_* , the m -th surrogate model returns a posterior mean and standard deviation (Figure 5),

$$\mu_m(\xi_*), \quad \sigma_m(\xi_*). \quad (18)$$

The posterior mean μ_m is used as the predicted value of the corresponding flow metric during target-oriented optimization. The posterior standard deviation σ_m quantifies the local model uncertainty and is used during active-learning exploration.

Because the output metrics have different physical units and numerical ranges, the predictive standard deviation of each model is normalized by an output scale,

$$\sigma_{m,\text{norm}}(\xi_*) = \frac{\sigma_m(\xi_*)}{s_m}, \quad (19)$$

where s_m is taken as the standard deviation of the measured training outputs for the m -th metric. This normalization allows uncertainties associated with turbulence intensity, homogeneity, anisotropy, and profile gradients to be combined in the active-learning criterion without introducing unit-dependent weights.

2.5. Exploration strategy

The purpose of the exploration phase is to identify how the measured flow metrics vary over the reduced active-grid control space before prescribing a specific target flow. Rather than optimizing a target flow condition during this stage, the algorithm selects new actuation states that are expected to be informative for the Gaussian-process surrogate models. This follows the general principle of active data selection, in which new measurements are chosen based on the current model uncertainty (MacKay, 1992; Cohn et al., 1996). Here, the exploration objective over a sequence of experiments is to reduce the global uncertainty of the surrogate models across the admissible actuation space.

The exploration loop starts with an initial bootstrapping stage, during which actuation states are sampled uniformly at random within the active design bounds until a minimum number of mea-




Figure 6. Closed-loop framework used for active-learning exploration and target-oriented exploitation. At each iteration, the accumulated input–output measurements are used to update the Gaussian-process surrogate models. The next actuation state is selected by an inner numerical optimization performed only on the surrogate models: during exploration, the acquisition function maximizes normalized predictive uncertainty with an anti-clustering correction, whereas during exploitation it minimizes the surrogate-predicted target objective. Only after this in-silico selection step is complete is the selected grid motion executed in the wind tunnel, followed by PIV acquisition, processing, and metric extraction.

measurements $n_{\min}$ is available for model training. These initial samples provide a minimal set of independent measurements from which the first surrogate models can be constructed.

After the bootstrapping stage, the models described in Section 2.4 are retrained after each new measurement. For a candidate input p in physical control coordinates, the corresponding normalized design coordinate ξ is obtained using Equation (13). Each trained surrogate model returns a posterior standard deviation $\sigma_m(\xi)$ for the m -th output metric. Following Equation (19), this uncertainty is normalized by the output scale s_m . The raw uncertainty score is defined as the root-mean-square normalized predictive uncertainty across all active output models,

$$\mathcal{I}(\xi) = \left[\frac{1}{N_{\text{out}}} \sum_{m=1}^{N_{\text{out}}} \left(\frac{\sigma_m(\xi)}{s_m} \right)^2 \right]^{1/2}, \quad (20)$$

where N_{out} is the number of trained output models. This definition gives equal importance to each modeled flow metric after normalization by its measured output scale. Thus, turbulence intensity, homogeneity measures, in-plane anisotropy, and profile gradients can all contribute to the exploration without introducing unit-dependent weights.

Maximizing only the predictive uncertainty can repeatedly select points close to previously tested actuation states, especially near boundaries of the design space where Gaussian-process uncer-




Figure 7. Schematic illustration of the distance-based anti-clustering factor used during active-learning exploration.

(a) Candidate actuation state ξ in the normalized input space, with the nearest-neighbor distance $d_{nn}(\xi)$ and anti-clustering length scale d_0 indicated relative to previously sampled states. (b) Anti-clustering factor P_{ac} as a function of d_{nn}/d_0 , showing the suppression of candidates close to existing samples.

tainties are often large. To discourage such clustering, the raw uncertainty score is multiplied by a distance-based anti-clustering factor. Let $\mathcal{X}_n = \{\xi_i\}_{i=1}^n$ denote the set of existing samples in normalized coordinates (Figure 7a). For a candidate ξ , the distance to the nearest existing sample is

$$d_{nn}(\xi) = \min_{\xi_i \in \mathcal{X}_n} \|\xi - \xi_i\|_2. \quad (21)$$

The anti-clustering factor is defined as

$$P_{ac}(\xi) = 1 - \exp \left[- \left(\frac{d_{nn}(\xi)}{d_0} \right)^2 \right], \quad (22)$$

where d_0 is an adaptive length scale in normalized input space (Figure 7b). At each iteration, d_0 is computed from the median nearest-neighbor distance among the existing samples,

$$d_0 = \text{clip}(\gamma \text{median}[d_{nn}(\xi_i)], d_{\min}, d_{\max}). \quad (23)$$

Here, γ is a scale factor and the clipping bounds $d_{\min}$ and $d_{\max}$ prevent the anti-clustering length scale from becoming unphysically small or overly restrictive. If too few existing samples are available to estimate the median nearest-neighbor distance reliably, a fallback value is used instead. The anti-clustering factor approaches zero near previously sampled points and approaches one far from them. Because it is multiplicative, it suppresses redundant samples but cannot turn a candidate with low surrogate uncertainty into a high-scoring exploration point.

The final exploration score is

$$E(\xi) = \mathcal{I}(\xi) P_{ac}(\xi). \quad (24)$$

The next experiment is selected as

$$p_{n+1} = \arg \max_{p \in \mathcal{B}} E(\xi(p)), \quad (25)$$

Table 2. Algorithmic settings used for active-learning exploration. Distances are computed in normalized input coordinates. The active exploration inputs for the results presented here are α , r_β , g_A , and g_0 ; the same framework can include f as an additional active input.

Quantity	Value	Description
$n_{\min}$	6	Minimum samples before GPR training
N_{out}	6	Active output metrics modeled by GPR
γ	0.5	Anti-clustering scale factor
$d_{\min}$	0.02	Minimum anti-clustering scale
$d_{\max}$	0.25	Maximum anti-clustering scale
d_{fallback}	0.10	Fallback scale for small sample sets
N_c	10 000	Candidate points for global uncertainty estimate
N_{pop}	$80 + 10d$	Genetic-algorithm population size (here: $d = 4$)
N_{gen}	100	Maximum genetic-algorithm generations

where $\mathcal{B}$ denotes the physical bounds of the active control parameters (Table 1).

The maximization of E is a secondary numerical optimization performed entirely on the surrogate models. During this step, the genetic algorithm evaluates many candidate actuation states, but each evaluation only requires Gaussian-process predictions and distances to the previously measured samples. No grid motion, PIV acquisition, or image processing is performed inside this inner optimization. Consequently, the computational cost of selecting the next point is small compared with the experimental iteration itself. For the results presented here, the actuation frequency was held fixed and the active exploration variables were α , r_β , g_A , and g_0 .

In addition to selecting individual high-uncertainty points, the overall progress of the exploration is monitored using a global uncertainty diagnostic. A fixed set of N_c candidate points is sampled uniformly in the normalized design domain. The raw uncertainty score $\mathcal{I}(\xi)$ is evaluated at each candidate point, and the mean global uncertainty is defined as

$$\bar{\mathcal{I}}_{\text{global}} = \frac{1}{N_c} \sum_{j=1}^{N_c} \mathcal{I}(\xi_j). \quad (26)$$

This quantity is independent of the single point selected by the genetic algorithm and provides a measure of how well the accumulated experiments have reduced the surrogate-model uncertainty over the full active design space. In practice, the exploration is continued until these global uncertainty measures stabilize or until the prescribed experimental budget is reached. The numerical settings used for the exploration procedure are summarized in Table 2.

2.6. Target objective and exploitation strategy

After the exploration phase, the accumulated measurements are used to train the surrogate models described in Section 2.4. These models are then used to select actuation states that are predicted to match prescribed flow targets. In contrast to the exploration strategy, which seeks to reduce model uncertainty, the exploitation strategy minimizes a target objective constructed from the surrogate-predicted mean flow metrics.

The measured and modeled outputs are separated into target quantities, penalty quantities, and diagnostic quantities. Target quantities prescribe desired values, such as a turbulence intensity or profile gradient. Penalty quantities are lower-is-better metrics, such as measures of spatial inhomogeneity or in-plane anisotropy. Diagnostic quantities may be recorded and modeled, but are not included in the target objective. In the present implementation, the target set is

$$\mathcal{T} = \left\{ \overline{TI}, \frac{dU}{dy}, \frac{dTI}{dy} \right\}, \quad (27)$$

while the penalty set is

$$\mathcal{P} = \{\mathcal{A}, CV_U, CV_{TI}\}. \quad (28)$$

The target values assigned to the quantities in $\mathcal{T}$ depend on the desired inflow case. The same objective formulation is used for all target cases.

For a candidate actuation state $\mathbf{p}$, the corresponding normalized input vector $\boldsymbol{\xi}$ is obtained from Equation (13). Each surrogate model then provides a posterior mean prediction $\mu_m(\boldsymbol{\xi})$ for the m -th output metric. For each target quantity $j \in \mathcal{T}$, the deviation from its prescribed target value y_j^* is normalized by an absolute tolerance τ_j ,

$$e_j(\boldsymbol{\xi}) = \frac{\mu_j(\boldsymbol{\xi}) - y_j^*}{\tau_j}. \quad (29)$$

The combined target score is then defined as the root-mean-square normalized target error,

$$S_T(\boldsymbol{\xi}) = \left(\frac{1}{N_T} \sum_{j \in \mathcal{T}} e_j(\boldsymbol{\xi})^2 \right)^{1/2}, \quad (30)$$

where N_T is the number of target quantities. With this definition, $S_T \approx 1$ corresponds to the tolerance boundary in an RMS sense. To distinguish candidates that are inside the target region from candidates that remain outside at least one tolerance band, a target-violation score is also defined,

$$V_T(\boldsymbol{\xi}) = \left[\frac{1}{N_T} \sum_{j \in \mathcal{T}} \max(|e_j(\boldsymbol{\xi})| - 1, 0)^2 \right]^{1/2}. \quad (31)$$

This quantity is zero only when all target quantities lie within their prescribed tolerance bands.

For each penalty quantity $\ell \in \mathcal{P}$, the surrogate-predicted mean is normalized by the corresponding output scale s_ℓ ,

$$q_\ell(\boldsymbol{\xi}) = \frac{\max[\mu_\ell(\boldsymbol{\xi}), 0]}{s_\ell}. \quad (32)$$

The maximum operation prevents small negative surrogate predictions of inherently positive lower-is-better metrics from being rewarded. Unless an explicit physical scale is prescribed, s_ℓ is taken as the training-output scale defined in Equation (19). The combined penalty score is

$$S_P(\boldsymbol{\xi}) = \left(\frac{1}{N_P} \sum_{\ell \in \mathcal{P}} q_\ell(\boldsymbol{\xi})^2 \right)^{1/2}, \quad (33)$$

where N_P is the number of penalty quantities.

The target and penalty scores are combined using a smooth gate,

$$G(S_T) = \frac{1}{1 + \exp[\kappa(S_T - 1)]}, \quad (34)$$

where κ controls the sharpness of the transition. For $S_T > 1$, the candidate is outside the target region and G is small, so the objective is dominated by target matching. Once the target region is approached, $S_T \lesssim 1$, the gate approaches unity and the lower-is-better penalty metrics become increasingly important.

The surrogate-predicted target objective is then

$$J(\boldsymbol{\xi}) = S_T(\boldsymbol{\xi})^2 + \lambda V_T(\boldsymbol{\xi})^2 + G(S_T) S_P(\boldsymbol{\xi})^2, \quad (35)$$

where λ penalizes target errors that remain outside the prescribed tolerance bands. This formulation prioritizes reaching the desired target quantities before optimizing secondary quantities such as homogeneity or in-plane anisotropy, while still discouraging solutions that remain outside the target tolerances.

The next targeted experiment is selected by minimizing the surrogate-predicted objective over the bounded actuation space,

$$\mathbf{p}_{n+1} = \arg \min_{\mathbf{p} \in \mathcal{B}} [J(\boldsymbol{\xi}(\mathbf{p}))], \quad (36)$$

where $\mathcal{B}$ denotes the physical bounds of the active control parameters. As in the exploration phase, this inner optimization is performed using a genetic algorithm and requires only surrogate-model predictions. The wind-tunnel experiment is run only after the candidate actuation state has been selected. The newly measured PIV metrics are then used to evaluate the same objective in measured form, appended to the input-output training set, and used to update the surrogate models before the next targeted iteration. The target-objective and exploitation settings used in the present study are summarized in Table 3.




Figure 8. Active-learning exploration of the reduced actuation space. (a) Global normalized surrogate-model uncertainty, $\bar{\mathcal{I}}_{\text{global}}$, evaluated over a fixed set of candidate points in the normalized design space. (b) Mean normalized predictive standard deviation, $\bar{\sigma}_{m,\text{norm}}$, for the individual Gaussian-process models. (c) Mean fractional contribution of each output model to the squared normalized exploration uncertainty, computed over the same fixed candidate set.

3. Results

The results are presented in three steps. First, we examine the active-learning exploration of the reduced actuation space and assess how the surrogate-model uncertainty evolves as additional PIV measurements are incorporated. We then use the explored data to characterize the range of flow responses generated by the modal actuation parameterization. Finally, we evaluate the surrogate-guided exploitation strategy for three prescribed turbulent inflow targets.

Table 3. Settings used for the target objective and surrogate-guided exploitation. Target values y_j^* are prescribed separately for each inflow case. The explicit tolerances are given in physical output units. Penalty scales are taken from the training-output scale s_m unless an explicit physical scale is provided.

Quantity	Value	Description
τ_{TI}	0.01	Tolerance for TI_{mean}
$\tau_{dU/dy}$	0.30	Tolerance for dU/dy
$\tau_{dTI/dy}$	0.03	Tolerance for dTI/dy
r_{tol}	0.05	Fallback relative tolerance for nonzero targets
λ	10	Target-violation penalty weight
κ	10	Target-gate sharpness

3.1. Exploration of the reduced actuation space

The active-learning exploration was performed in the four-dimensional control space defined by $\mathbf{p} = (\alpha, r_\beta, g_A, g_0)^T$, with the actuation frequency held fixed. Each exploration iteration consists of one physical experiment: the selected grid motion is executed, PIV image pairs are acquired and processed, and the resulting velocity fields are reduced to the scalar metrics defined in Section 2.2. These measurements are then used to update the Gaussian-process surrogate models before the next actuation state is selected based on the highest uncertainty of the system.

A total of 150 closed-loop exploration iterations were completed, requiring approximately 12 h 30 min of unattended run time. The wall-clock time per iteration increased only modestly over the course of the run, from 4 min 25 s at $n = 1$ to 4 min 51 s at $n = 150$. The dominant cost was the automated PIV processing, which required approximately 3 min 13 s per iteration, followed by execution of the wind-tunnel experiment, including image-frame download. Surrogate-model retraining and candidate selection required less than one second during the initial iterations and increased to approximately 25 s by the final iteration as the training set grew. Thus, the closed-loop iteration rate was limited primarily by image processing and data acquisition rather than by the Gaussian-process optimization.

The evolution of the surrogate-model uncertainty during exploration is shown in Figure 8. The global uncertainty $\bar{\mathcal{I}}_{\text{global}}$ (Equation (26)) decreases over the course of the exploration, indicating that the active-learning strategy progressively reduces the predictive uncertainty over the admissible actuation space (Figure 8a). After the initial random bootstrapping stage, the global uncertainty is of order one and subsequently decreases as additional measurements are incorporated, reaching values of approximately 0.6–0.7 by the end of the exploration run. Because this quantity is computed from predictive standard deviations normalized by the measured output scales, its absolute value should be interpreted relative to the variability of the training data. Thus, $\bar{\mathcal{I}}_{\text{global}} \approx 1$ indicates that the typical predictive standard deviation is comparable to the characteristic spread of the measured outputs, whereas values below one indicate a reduction relative to that scale.

The individual output models do not converge at the same rate (Figure 8b). The uncertainty associated with the mean turbulence intensity, $\bar{TI}$, decreases rapidly during the early iterations and reaches the lowest final uncertainty among the modeled outputs. This suggests that $\bar{TI}$ varies comparatively smoothly and consistently with the selected actuation parameters. In contrast, the uncertainties associated with CV_U , CV_{TL} , $\mathcal{A}$, dU/dy , and dTI/dy decrease more gradually and remain higher at the end of the run. These metrics are expected to be more sensitive to local variations in the measured velocity field and to finite-sample variability, because they depend on spatial gradients, spatial non-uniformity, or second-order statistics rather than on a domain-averaged turbulence level alone. They may also exhibit more structured response surfaces over the explored actuation space, particularly when small changes in the gradient-mode parameters alter the vertical distribution of mean velocity and turbulence intensity.

Additional experiments could help distinguish between these possible sources of uncertainty. For example, repeated measurements at a fixed actuation state, or longer acquisitions with substantially larger sample counts, would quantify the intrinsic experimental variability of the extracted metrics. By contrast, additional exploration points in regions of persistently high surrogate uncertainty would reveal whether the remaining uncertainty is primarily caused by an under-resolved response surface. If the latter is the dominant source, the global uncertainty should continue to decrease as additional data are incorporated and the surrogate models resolve the local variation more accurately. Finally, if certain gradient-related targets remain difficult to reach even after additional sampling, this would suggest that the present four-parameter actuation basis could be enriched with additional spatial modes to provide finer control over the vertical structure of the generated flow.

The non-monotonic variations in the uncertainty curves are expected for an iteratively retrained surrogate model. Although the overall trend is downward, an individual measurement can reveal that a region previously represented by a smooth response requires a shorter length scale, larger noise level, or a more complex local variation. When the Gaussian-process hyperparameters are re-estimated, this can temporarily increase the predicted uncertainty, even though the new sample adds information to the training set. These increases therefore do not indicate a failed exploration step, but rather reflect the updated model interpretation of the measured response surface.

The relative contribution of each output metric to the remaining exploration uncertainty is summarized in Figure 8c. For this diagnostic, the squared normalized predictive uncertainty of each output model is evaluated at every candidate point and divided by the corresponding sum over all output models. The plotted values are then averaged over the candidate set, so that each bar represents the mean fraction of the normalized uncertainty budget associated with a given output metric. Consistent with the rapid reduction seen in Figure 8b, $\overline{TI}$ contributes only weakly to the final uncertainty budget. The gradient metrics, especially dTI/dy , contribute more strongly, indicating that they remain among the more difficult quantities to predict from the present exploration data. This behavior is important for the targeted optimization stage: while the exploration produces a useful surrogate representation of the actuation space, target cases involving profile gradients are expected to be more challenging than cases involving only the mean turbulence intensity.

3.2. Surrogate-model uncertainty and learned response range

The sampled input distributions in Figure 9a show how the active-learning strategy populates the reduced actuation space. The resulting samples are not uniformly distributed, because the algorithm preferentially selects regions where the surrogate models remain uncertain while suppressing redundant samples through the anti-clustering factor. Especially in the beginning of the exploration process, this will favor regions at the edges of the parameter space, while over time the




Figure 9. (a) Distribution of sampled input parameters during exploration. The histogram ranges over the same space as the input parameter ranges in Table 1. (b) Distribution of measured output metrics obtained from the explored actuation states. Here the absolute counts are normalized by the total number of iterations n .

center is further mapped out. The measured output distributions in Figure 9b provide the corresponding learned response range of the grid over the explored parameter space. This demonstrates the capability of the grid for the given actuation modes and their ranges at the selected actuation frequency f (Table 4).

Together, these results show that the exploration stage builds a surrogate description of both the accessible flow states and the remaining model uncertainty, which is then used for the target-oriented exploitation discussed below.

3.3. Target case optimization results

The optimization targets three canonical turbulent flow states (Figure 10a):

- I. Homogeneous, isotropic turbulence with a prescribed target turbulence intensity TI_{target}
- II. A prescribed mean velocity gradient dU/dy at constant turbulence intensity TI_{target}
- III. A prescribed turbulence intensity gradient dTI/dy at constant mean velocity with a prescribed target turbulence intensity TI_{target}

The grid actuation is defined by the input parameter vector $\mathbf{p}$, comprising the amplitudes and offsets of the synchronous and gradient mode shapes (Section 2.3). Optimization progress was evaluated using the target objective defined in Section 2.6.

Each optimization case consisted of 50 closed-loop iterations and required approximately 4 h 45 min of unattended run time. Compared with the exploration stage, the target-oriented optimization introduced an additional computational cost of approximately 54 s per iteration due to the genetic-algorithm search for the minimum of the surrogate-predicted target objective.

For all three target cases, the largest improvement occurs during the first 5–10 iterations, after which the objective decreases more gradually and only minor refinements are obtained (Figure 10c). The two independent trials produce similar convergence histories and comparable final objective

Table 4. Summary of active turbulence grid solution space, including turbulence intensities, coefficients of variation, and spatial gradients.

Parameter name	Variable	Min value	Max value
Mean turbulence intensity	$\overline{TI}$	0.03	0.32
Coefficient of variation of stream-wise velocity	CV_U	0.005	0.10
Coefficient of variation of turbulence intensity	CV_{TI}	0.02	0.14
Mean anisotropy	$\mathcal{A}$	0.04	0.43
Mean stream-wise velocity gradient	dU/dy	−4.17 (1/s)	6.02 (1/s)
Mean turbulence intensity gradient	dTI/dy	−0.34 (1/m)	0.17 (1/m)




Figure 10. (a) Velocity and turbulence-intensity profiles defining the optimization targets. (b) Measured velocity and turbulence-intensity profiles of the best-performing actuation states. The dashed lines indicate the average slopes of the measured profiles. (c) Convergence of the target objective J for all three optimization scenarios (I)–(III). The light and dark markers in panels (b) and (c) represent two independent optimization trials.

values, particularly for cases (II) and (III). This repeatability is also reflected in the measured velocity and turbulence-intensity profiles of the best-performing actuation states (Figure 10b), indicating that the closed-loop procedure converges reliably toward similar solutions.

Overall, the prescribed targets for $\bar{TI}_{\text{target}}$ and the profile gradients are captured well by the optimization. For all three cases, the measured streamwise velocity profiles remain close to their prescribed forms, with only minor deviations from the mean or target gradient. The turbulence-intensity-gradient case also achieves the desired average dTI/dy , although the resulting TI profiles exhibit more pronounced spatial variation about the mean trend. This variation occurs over a length scale comparable to the grid-panel width l , suggesting that the finite spatial structure of the grid actuation remains visible in the downstream turbulence-intensity field. Although this behavior is partly reflected in the penalty metric CV_{TI} , future implementations could penalize deviations from the target gradient profile more directly, for example by including a profile-residual metric in the objective.

Together, these results demonstrate automated, closed-loop control of a 15-axis active turbulence grid capable of generating tailored turbulent inflow states with prescribed mean and fluctuating

properties. Such controlled flow environments are well suited for systematically exposing aerodynamic structures, such as aircraft components or wind-turbine models, to realistic and repeatable turbulent inflow conditions.

4. Conclusion

This study demonstrates an automated, measurement-driven framework for controlling an active turbulence grid to generate tailored turbulent inflow states. The approach combines a reduced-order, mode-based parameterization of the grid actuation with planar PIV feedback and Gaussian-process surrogate models. Instead of tuning the grid manually or relying on fixed actuation protocols, the framework first explores the reduced actuation space using an uncertainty-guided active-learning strategy and then uses the learned surrogate models to select actuation states for prescribed target flow conditions.

The active-learning exploration was performed in a four-dimensional control space defined by the synchronous and gradient-mode parameters. Over 150 closed-loop iterations, the global normalized surrogate-model uncertainty decreased substantially, indicating that the accumulated PIV measurements progressively improved the surrogate representation of the accessible flow states. The exploration also showed that different output metrics were learned at different rates. The mean turbulence intensity reached comparatively low predictive uncertainty, whereas spatial gradients, coefficients of variation, and the in-plane anisotropy proxy retained higher uncertainty over the explored design space.

The target-oriented optimization was demonstrated for three canonical turbulent inflow objectives: homogeneous isotropic turbulence with prescribed turbulence intensity, a prescribed mean velocity gradient at constant turbulence intensity, and a prescribed turbulence intensity gradient at constant mean velocity. For all three cases, the optimization identified actuation states that closely matched the prescribed mean turbulence intensity and profile-gradient targets within a limited number of iterations. Independent optimization trials converged toward similar objective values and comparable measured profiles, demonstrating the repeatability of the closed-loop procedure.

The timing measurements show that the additional surrogate-based optimization step does not substantially slow the closed-loop procedure compared with the experimental and PIV-processing costs. Instead, its main benefit is that it can reduce the number of physical experiments required to identify useful actuation states by selecting new measurements and target candidates from the learned surrogate models. This is particularly important for PIV-in-the-loop optimization, where each iteration requires image acquisition, image-frame download, velocity-field processing, and metric extraction.

Several next steps follow naturally from the present framework. First, repeated measurements

at fixed actuation states, longer acquisitions, and additional samples in regions of high surrogate uncertainty could be used to distinguish intrinsic experimental variability from under-resolved response surfaces. Second, the objective function could be extended to include additional flow metrics, such as temporal spectra, integral time scales, or profile-residual penalties, and could be combined with complementary diagnostics such as hot-wire measurements. Third, the actuation basis could be enriched by including additional spatial modes, varying the actuation frequency, or allowing more independent control of horizontal and vertical grid axes. These extensions would provide finer control over both the spatial and temporal structure of the generated turbulence.

Overall, the presented framework provides a scalable route toward on-demand generation of turbulent inflow conditions in wind-tunnel experiments. By using spatially resolved velocity-field measurements directly as feedback, the method enables active turbulence grids to be tuned for prescribed mean and fluctuating flow properties without exhaustive manual parameter studies. This capability is particularly useful for experiments in aerodynamics, fluid–structure interaction, and wind-energy applications, where models must be exposed to repeatable turbulent inflow conditions with controlled mean-flow and turbulence profiles.

Acknowledgements

The authors would like to acknowledge support from the Brown University Institute of Sustainable Energy (ISE) and the Office of Naval Research, grant ONR/NUWC GR5250066.

References

- Bollt, S. A., Foxman, S. H., & Gharib, M. (2025, October). RapidPIV: full flow-field kHz PIV for real-time display and control. *Measurement Science and Technology*, 36(10), 105116. doi: 10.1088/1361-6501/ae124d
- Cekli, H. E., & Van De Water, W. (2010, August). Tailoring turbulence with an active grid. *Experiments in Fluids*, 49(2), 409–416. doi: 10.1007/s00348-009-0812-5
- Cohn, D. A., Ghahramani, Z., & Jordan, M. I. (1996). Active Learning with Statistical Models. *Journal of Artificial Intelligence Research*, 4, 129–145. doi: 10.1613/jair.295
- Gautier, N., Aider, J.-L., Duriez, T., Noack, B., Segond, M., & Abel, M. (2015, May). Closed-loop separation control using machine learning. *Journal of Fluid Mechanics*, 770, 442–457. doi: 10.1017/jfm.2015.95

- Gehrke, A., & Mulleners, K. (2021, March). Phenomenology and scaling of optimal flapping wing kinematics. *Bioinspiration & Biomimetics*, 16(2), 026016. (Number: 2) doi: 10.1088/1748-3190/abd012
- Hearst, R. J., & Lavoie, P. (2015, October). The effect of active grid initial conditions on high Reynolds number turbulence. *Experiments in Fluids*, 56(10), 185. doi: 10.1007/s00348-015-2052-1
- MacKay, D. J. C. (1992, July). Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 4(4), 590–604. doi: 10.1162/neco.1992.4.4.590
- Makita, H. (1991, October). Realization of a large-scale turbulence field in a small wind tunnel. *Fluid Dynamics Research*, 8(1-4), 53–64. doi: 10.1016/0169-5983(91)90030-M
- Mydlarski, L., & Warhaft, Z. (1996, August). On the onset of high-Reynolds-number grid-generated wind tunnel turbulence. *Journal of Fluid Mechanics*, 320, 331–368. doi: 10.1017/S0022112096007562
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning* (3. print ed.). Cambridge, Mass.: MIT Press. doi: 10.7551/mitpress/3206.001.0001
- Sacks, J., Welch, W. J., Mitchell, T. J., & Wynn, H. P. (1989, November). Design and Analysis of Computer Experiments. *Statistical Science*, 4(4). doi: 10.1214/ss/1177012413
- Stefan-Zavala, A., Scherl, I., Mandralis, I., Brunton, S. L., & Gharib, M. (2026). Data-driven modelling for on-demand flow prescription in fan-array wind tunnels. *Flow*, 6, E3. doi: 10.1017/flo.2025.10034