Citation (APA)
Salomone, G., Mohammadikalakoo, B., García Villasol, S., Doan, N. A. K., Greco, C. S., & Kotsonis, M. (2026). An experimental particle swarm optimization-driven real-time control of Tollmien–Schlichting waves on a flat plate. *Physics of Fluids*, *38*(7), Article 074118. https://doi.org/10.1063/5.0340217

RESEARCH ARTICLE | JULY 22 2026

An experimental particle swarm optimization-driven real-time control of Tollmien–Schlichting waves on a flat plate

G. Salomone ; B. Mohammadikalakoo ; S. García Villasol ; N. A. K. Doan ; C. S. Greco ; M. Kotsonis

Physics of Fluids 38, 074118 (2026)
https://doi.org/10.1063/5.0340217

Physics of Fluids ARTICLE pubs.aip.org/aip/pof

An experimental particle swarm optimization-driven real-time control of Tollmien–Schlichting waves on a flat plate

Cite as: Phys. Fluids 38, 074118 (2026); doi: 10.1063/5.0340217
Submitted: 23 April 2026 · Accepted: 9 July 2026 ·
Published Online: 22 July 2026

View Online Export Citation CrossMark

G. Salomone,[1,a)] B. Mohammadikalakoo,[2] S. García Villasol,[3] N. A. K. Doan,[4] C. S. Greco,[1] and M. Kotsonis[2]

AFFILIATIONS

[1]Dipartimento di Ingegneria Industriale, Università degli Studi di Napoli "Federico II", Naples, Italy
[2]Department of Flow Physics and Technology, Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, Delft 2629HS, The Netherlands
[3]Department of Aerospace Engineering, Universidad Carlos III de Madrid, Avenida de la Universidad 30, Leganés 28911, Spain
[4]Department of Aeronautics, Imperial College London, London SW7 2AZ, United Kingdom

[a)]Author to whom correspondence should be addressed: gabriele.salomone@unina.it

ABSTRACT

This experimental study demonstrates the feasibility of integrating particle swarm optimization (PSO) within a hardware-in-the-loop framework for real-time active flow control. The objective is the attenuation of Tollmien–Schlichting (TS) waves developing in an incompressible, two-dimensional laminar boundary layer, using dielectric barrier discharge plasma actuators. Single- and multi-frequency TS waves are artificially excited and amplify before reaching the control region. Control actions are optimized online using PSO, which iteratively updates a population of candidate solutions based on a pressure-based performance metric. In single-frequency scenarios, candidate solutions are parametrized as finite impulse response filters, convolved with an upstream pressure signal to generate the control voltage. In multi-frequency scenarios, the controller constructs a linear superposition of sinusoids at the triggered TS frequencies, identified online. The effectiveness of control is quantified by the attenuation of pressure fluctuations measured by wall-mounted microphones downstream of the actuator. Upon convergence, the best-performing control actions are re-tested in independent experiments, using two-component particle image velocimetry (PIV) to assess the effect of the controller on the velocity field inside the boundary layer. The results show that PSO consistently achieves appreciable TS waves suppression, as evidenced by pressure fluctuation levels reduced by 30%–40% relative to the uncontrolled case and by weakened phase-coherent velocity disturbances. Furthermore, phase-free PIV measurements demonstrate a downstream delay in the growth of velocity fluctuations of up to $15\,\delta_0^*$, where δ_0^* is the displacement thickness at the control location. These findings confirm the feasibility of PSO-driven optimization for real-time flow control in transitional boundary layers.

I. INTRODUCTION

Delaying the laminar-to-turbulent transition in boundary layers developing over aerodynamic surfaces is a well-established strategy for reducing skin friction drag, which can significantly contribute to overall drag reduction. The transition process begins with a receptivity process, where external perturbations, such as acoustic waves, freestream turbulence, or surface roughness,[1,2] introduce small disturbances into the laminar boundary layer. These disturbances initially develop in a linear amplification region before progressing into a nonlinear regime and ultimately breaking down into turbulence.[3,4] The spatial development of these disturbances is governed by growth or decay rates determined by the flow properties of the problem (Reynolds number, frequency and initial magnitude of the disturbances). In low external disturbance environments, transition in two-dimensional subsonic boundary layers is typically dominated by the amplification of Tollmien–Schlichting (TS) waves.[5] Consequently, suppressing or delaying the growth of TS waves facilitates transition delay and skin friction drag reduction.

A. Delay of laminar-to-turbulent transition

Two general approaches can be followed to achieve a delay in laminar-to-turbulent transition. The first option consists of altering the mean flow close to the wall and thereby altering the boundary layer stability, reducing the growth of the instabilities responsible for the transition.[6] Alternatively, the spatial amplification of instabilities is directly targeted through opposition control.[7] Given their role in the laminar-to-turbulent transition, suppression or weakening of TS waves by direct approach has been a widely researched topic, with successful applications of passive[8–10] and active flow control (AFC) techniques.[11–15] A particular class of actuators that has been successfully employed is the dielectric barrier discharge (DBD) plasma actuators.[16,17] DBDs are used as a quasi-steady momentum source to modify the flow stability[18–20] and operated unsteadily to counteract the growing instabilities inside the boundary layer.[21] Recently, the interest toward the application of DBD plasma actuators has increasingly grown thanks to their simple architecture, the absence of moving parts and their short response time, which is suitable for unsteady applications.[16] A further significant benefit is the possibility of modeling their effect in the flowfield as body force models in numerical solvers, leading to developments in both experimental[22–26] and numerical applications.[27,28] By targeting the flow instabilities, DBD plasma actuators can be very energy-efficient. This was verified by Grundmann and Tropea,[18,21] who demonstrated that operating a DBD actuator in pulsed mode[21] to attenuate TS waves in a transitional boundary layer resulted in turbulence reduction levels comparable to those obtained with steady actuation designed to modify the mean flow.[18] However, the pulsed operation resulted in a markedly lower energy consumption, requiring only approximately 12% of the power consumption of the steady mean-flow modification case.[21]

Despite the numerous developments, the classical techniques that have been used to evolve the control strategies are severely challenged when the flowfield exhibits strong non-linearity or multiple temporal and spatial scales.[29,30] In fact, weakly turbulent or transitional flows may no longer be dominated by a small number of coherent harmonic modes, but instead display broadband energy spectra and intermittent dynamics. Under these conditions, control approaches relying on the identification of dominant frequencies or low-dimensional models can lose effectiveness, as the underlying assumptions of modal dominance are no longer valid. Artificial intelligence (AI), particularly machine-learning-based methods, has recently been explored for flow control as a potential means of addressing such challenges. In particular, several studies have reported that these methods can remain effective under more general flow conditions[31] and in the presence of noisy measurements.[32]

B. AI in flow control

Among AI applications to AFC, a prominent research line has explored the use of genetic programming[33] to tackle classical fluid dynamics problems, such as separation control over sharp edges,[34,35] airfoil stall[36] and drag reduction over bluff bodies.[37,38] Recent advances have shown promising results through the application of deep reinforcement learning (DRL)[39–41] and a "lightweight" version of DRL, the single-step deep reinforcement learning (SDRL).[42] These algorithms have been successfully applied to shape optimization in numerical frameworks,[43,44] skin friction reduction,[45,46] and separation control in numerical[47–49] and experimental[50] environments.

Among optimization-based AI approaches, particle swarm optimization (PSO) has gained popularity in engineering applications due to its simplicity and effectiveness in solving optimization problems across various fields.[51] First introduced by Kennedy and Eberhart in the 1990s,[52,53] PSO has been widely applied in several industrial and research contexts.[54]

In the context of AFC, PSO has been primarily investigated within numerical frameworks. For instance, Mohammadikalakoo *et al.*[55,56] employed PSO to optimize unsteady suction and blowing jets for the suppression of TS waves in a two-dimensional boundary layer solved using a harmonic Navier–Stokes formulation. The control action operates at prescribed frequency, at a fixed location derived from a parametric study. The optimization variables consist of a set of physically interpretable variables, namely the actuator amplitude and phase. The objective function is based on the attenuation of downstream wall-pressure fluctuations associated with the TS waves, optionally penalized by the actuator energy expenditure. In a subsequent numerical study, Mohammadikalakoo *et al.*[57] extended the control framework to the time domain by introducing a real-time SDRL controller. In this application, the SDRL optimizes the coefficients of a finite impulse response (FIR) filter that are convolved with pressure signals to generate a voltage signal applied to a wall-mounted plasma actuator modeled as a volume-distributed body force. More importantly, because the actuation signal is generated through convolution with the measured reference signal, the controller does not rely on prescribed frequencies or a predefined harmonic basis. Instead, it naturally inherits the spectral content of the incoming disturbances, making the control framework largely insensitive to variations in the flowfield spectrum characteristics.

These studies demonstrated the effectiveness of AI-based optimization in suppressing TS waves within controlled numerical environments, even when the spectral content is not known *a priori* or in the presence of artificial noise.[57]

Following these numerical investigations,[55–57] the real-time SDRL controller has been experimentally tested to demonstrate control effectiveness as well as robustness to moderate variations in flow conditions.[58]

C. Present work

Building upon the numerical investigations[55–57] and experimental applications,[58] the present study extends PSO-based control to an experimental hardware-in-the-loop framework. Here, TS waves are suppressed in a two-dimensional incompressible laminar boundary layer developing under zero-pressure-gradient (ZPG) conditions over a flat plate.

The disturbances are triggered in the flowfield and controlled with DBD plasma actuators. In the present study, PSO is employed in a real-time, closed-loop feedforward control optimization framework, where control parameters are iteratively adjusted based on sensor feedback. PSO iteratively tunes the actuation signal in terms of amplitude and phase in order to achieve optimal TS waves suppression. The performance of the controllers is first quantified based on the reduction in pressure fluctuations recorded by a wall-mounted microphone downstream of the control actuator. At the end of each optimization run, the most effective controllers, referred to as "designed" controllers, are extracted and re-tested. Alongside wall-mounted microphones, two-component planar particle image velocimetry (PIV) measurements are conducted to visualize the attenuation of coherent flow structures under optimized control conditions. The main contribution is therefore the proof-of-concept hardware implementation of PSO for real-time active flow control, demonstrating practical swarm convergence and reliable operation in the presence of measurement noise, sensor and actuator imperfections, electromagnetic interference, convective delays, voltage constraints, and limited experimental runtime, together with complementary validation of the control effect through wall-pressure and velocity-field measurements.

In the context of the optimization, the choice of PSO is primarily motivated by its algorithmic simplicity and suitability for experimental implementation. In contrast to SDRL-type approaches, PSO does not require training a policy network, computing gradients, or selecting neural-network architectures. The algorithm only requires the experimental evaluation of candidate control laws and the update of the swarm based on a limited number of hyperparameters.

The objective of the present work is therefore not to establish PSO as a superior optimizer, nor to systematically compare multiple algorithms in terms of achievable rewards or robustness.

Such an assessment would require repeated optimization runs under nominally identical flow and actuator conditions, using different random initial particle distributions, together with systematic hyperparameter tuning.

Instead, the main goal of this work is to demonstrate the practical viability of PSO-driven active flow control in a laboratory environment. This is an important step toward real-world implementation, since experimental deployment must overcome challenges that are typically neglected or idealized in simulations, such as measurement noise, actuator and sensor imperfections, latency, and the need for robust real-time optimization under physical constraints.

Regarding the SDRL implementation for TS waves control reported in Ref. 58, a direct quantitative comparison is not attempted here because the two campaigns were conducted independently, with different model insertions and potentially different actuator authority, sensor response, and electromagnetic-interference sensitivity. Moreover, the PSO hyperparameters are not systematically optimized in the present work. A rigorous comparison between PSO and SDRL is therefore left to future studies based on a common flow problem and systematic hyperparameter tuning.

The remainder of this paper is structured as follows. Section II outlines the methodology, detailing the flow problem, the control framework, the experimental setup, and design of the experiments. Section III introduces the implementation of the optimization algorithms and the test cases. Section IV discusses the experimental results. Finally, conclusions are drawn in Sec. V.

II. METHODOLOGY

This section introduces the design of the experiment, starting from an overview of the flow problem, followed by results from a linear stability analysis (LSA). Outputs from LSA provide insights into the amplification of TS waves and guide the experiment design. The characterization of the experiments is completed with the description of the experimental setup and with the design of the framework for triggering and controlling TS waves in the laminar boundary layer.

A. Flow problem

The outline of the flow problem and the layout of input sensors and actuators are included in Fig. 1.

The experimental campaign utilizes flow control to weaken the evolution of TS waves on a flat plate, in zero-pressure gradients conditions. A coordinate reference frame is centered at the control location. The streamwise coordinate x is defined as the distance measured along the surface starting from the controller location ($x_{control} = 0$). To provide local information about the development of the boundary layer, a curvilinear coordinate x_s is also introduced along the flat plate surface. The coordinate x_s has its origin at the leading edge of the flat plate and it is directed in the streamwise direction. The wall-normal coordinate y denotes the distance measured perpendicular to the surface, and the spanwise coordinate z is defined such that (x, y, z) forms a right-handed reference frame.

In this work, x^*, η, and z^* denote the non-dimensional streamwise, wall-normal and spanwise coordinates, while F indicates the non-dimensional frequency of the TS waves. Non-dimensional parameters are defined as

$$x^* = x/\delta_0^*, \qquad \eta = y/\delta_0^*, \qquad \text{and} \qquad z^* = z/\delta_0^*, \tag{1a}$$

$$F = \omega_{ref}/Re_{\delta_0^*} \times 10^6 \quad \text{and} \quad \omega_{ref} = 2\pi f\, \delta_0^*/U_\infty, \tag{1b}$$

where δ_0^* is the displacement thickness at the control location computed with the Blasius solution, f is the frequency of the disturbance, U_∞ is the freestream velocity, and $Re_{\delta_0^*}$ is a global Reynolds number. $Re_{\delta_0^*} = U_\infty\, \delta_0^*/\nu$ is based on freestream velocity and on the reference

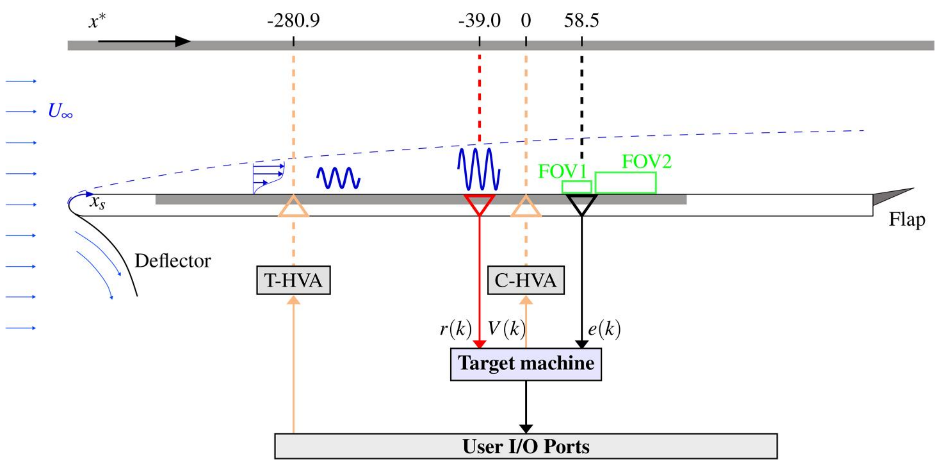

FIG. 1. Sketch of the flat plate model, flow problem, and control framework. Boundary layer edge is represented with a blue dashed line. DBDs are represented as apricot upward triangles, and microphones are denoted by downward triangles.

length δ_0^*, with $\nu = 1.5 \times 10^{-5}$ indicating the kinematic viscosity of air. In subsequent result sections, velocity components are normalized with the freestream velocity,

$$u^* = U/U_\infty \quad \text{and} \quad v^* = V/U_\infty. \tag{2}$$

In the following, the boundary layer thickness is indicated with $\delta_{99}(x^*)$ and it is also estimated with the Blasius solution, while the displacement thickness is indicated as $\delta_0(x^*)$. Finally, pressure measurements and voltage levels are non-dimensionalized based on reference values taken as $p_{ref} = 1$ Pa and $V_{ref} = 7.5$ kV, respectively.

B. Control framework

In order to create a repeatable flow problem for applying flow control, TS waves (blue sinusoidal shapes in Fig. 1) are triggered inside the ZPG boundary layer. Such coherent disturbances are introduced in the boundary layer by a "trigger" DBD plasma actuator (T-DBD, indicated by the upstream upward triangle), mounted far upstream on the flat plate to introduce TS waves close to their neutral point. T-DBD is enabled by a trigger high-voltage amplifier (T-HVA) independent of the flow control. Monochromatic TS waves are introduced by feeding the T-DBD plasma actuator with a sinusoidal voltage signal at the frequency of the target TS wave. A linear superposition of sinusoidal signals is used to trigger multiple TS waves in multi-frequency control scenarios. Second, enabled by a "controller" high-voltage amplifier (C-HVA), a controller DBD plasma actuator (C-DBD, indicated by the downstream upward triangle in Fig. 1) counteracts TS waves in the flowfield. The actuator is enabled with a signal output by the control algorithm. The control strategy utilizes pressure measurements acquired at the wall. A first microphone, namely "reference" microphone (red downward triangle in Fig. 1), acquires a signal upstream of the C-DBD while an "error" microphone (black downward triangle in Fig. 1) records pressure data downstream of the C-DBD. Statistics extracted from the error measurements are sensitive to the effect of control actions and are used to discriminate the effect of different control actions. Locations of T-DBD, C-DBD, and microphones are designed based on linear stability analysis (LSA).

C. Design of the experiments

Depending on the test case, the global Reynolds number is set in the range $Re_{\delta_0^*} \in 1.71 - 1.81 \times 10^3$ ($U_\infty \in$ m/s), yielding $\delta_0^* = 1.29$ and 1.21 mm, respectively. The nominal design condition is $Re_{\delta_0^*} = 1.71 \times 10^3$, corresponding to $U_\infty = 20$ m/s. A numerical boundary layer solver is adopted to obtain the baseline flow for stability computations. The velocity profile outside the boundary layer, referred to as external velocity U_e, is assumed constant with the streamwise direction x ($U_e(x) = U_\infty$), corresponding to ZPG conditions.

Once the boundary layer solution is obtained, the velocity profile at a given station x_{stab} is extracted and input to the LSA solver, which solves numerically the incompressible Orr–Sommerfeld equation,[59] as a polynomial eigenvalue problem.[60] Subsequently, for the selected freestream velocity U_∞ and a given frequency f, non-dimensional growth rates α_i^* output by the LSA solver are used to compute the map of amplification factor N along the flat plate as a function of the disturbance frequency f. The amplification factor quantifies the cumulative amplification of the disturbance at the frequency f and is defined as[61]

$$N(f, x) = \int_{x_0(f)}^{x} -\alpha_i(f, x)dx, \tag{3}$$

where $x_0(f)$ is the neutral point, i.e., the first location for which a negative growth rate (amplification behavior) is observed for frequency f, and $\alpha_i = \alpha_i^*/\delta_0^*$ is the dimensional growth rate computed from stability calculations. Figure 2 illustrates the map of the amplification factor for frequencies in the range F between 24 and 165 ($f \in 100$–700 Hz). Black dashed vertical lines are drawn at the locations of the plasma actuators, at $x^*_{trigger} = -280.9$ ($x_s = 380$ mm) and $x^*_{control} = 0$ ($x_s = 740$ mm), respectively. The coordinate refers to the location of the electrodes' interface. The reference signal is sampled

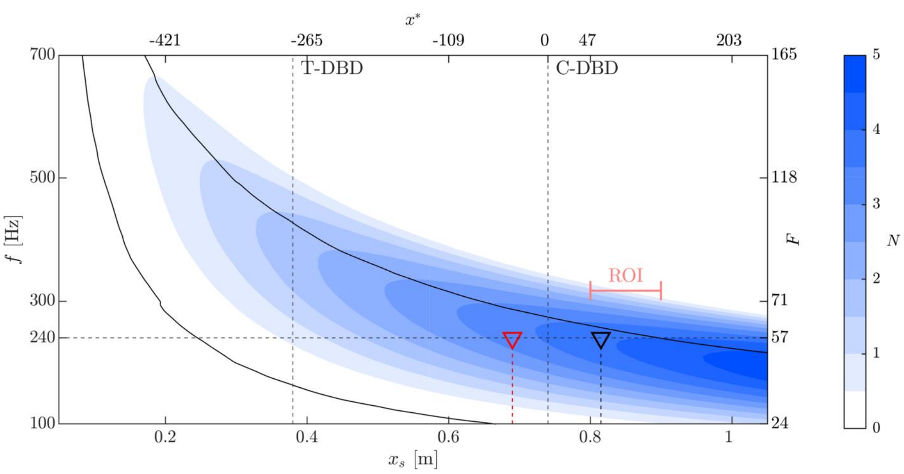

FIG. 2. Contour map of the amplification factor N at $\mathrm{Re}_{\delta_0^*} = 1.71 \times 10^3$ ($U_\infty = 20$ m/s) and neutral stability curve ($\alpha_i = 0$, solid black line). The streamwise locations of the actuators (T-DBD and C-DBD) are indicated by vertical dashed lines. The target TS frequency $F = 57$ ($f = 240$ Hz) is marked by a horizontal dashed line. Locations of the reference and error microphones are marked by red and black triangles. The region of interest (ROI) is indicated by an orange segment.

upstream of the control device, at $x^*_{ref} = -39.0$ ($x_s = 690$ mm). The effect of the control can be appreciated in a region of interest (ROI), defined by $x^* \in 39.0$–136.6 ($x_s \in 790$–915 mm) and placed downstream of the C-DBD, where wall-microphones are located. In this region, disturbances with F between 24 and 71 ($f \in 100$–300 Hz) exhibit the highest amplification and should therefore be targeted in the flow control scenarios. At the chosen error location $x^*_{err} = 58.5$ ($x_s = 0.815$ m), the non-dimensional frequency of the most amplified mode is approximately $F = 57$ ($f = 240$ Hz). TS waves are triggered near the neutral stability point, while the C-DBD and the microphones inside the ROI are located in the region of maximum disturbance amplification.

D. Experimental setup

Experiments take place in the A-Tunnel,[62] located in the Low-speed Wind Tunnel Laboratory of TU Delft. The A-tunnel is an open-jet, closed-circuit, vertical wind tunnel, placed in an anechoic plenum, that is approximately a square room ($6.4 \times 6.4\,\mathrm{m}^2$) with a height of 3.2 m. The high contraction ratio of the settling chamber produces a freefield turbulence intensity smaller than 0.05%, with a velocity non-uniformity at the outlet below 0.6%. Moreover, the inner side of the ceiling and the four walls of the room are covered with wedges made of Flamex[63] acoustic absorbing foam, preventing any reflection in a sound frequency range between 150 and 20 kHz. The flow control setup is sketched in Fig. 1. The wind tunnel model consists of a $860 \times 1350 \times 10\,\mathrm{mm}^3$ aluminum plate, enclosed between four lateral walls. The test model is designed with a superelliptic leading edge and a movable trailing edge flap. This ensures smooth pressure-gradient evolution that maintains a stable laminar boundary layer by preventing early transition of the flow and minimizing flow separation. The streamwise uniformity of pressure is verified using two staggered rows of 128 pressure taps distributed along the flat plate length, at a distance of 110 mm from the lateral walls, that is, $z^* = \pm 248$. ZPG conditions are verified with the order of magnitude of the acceleration parameter[64] $K = \frac{\nu}{U_e^2}\frac{dU_e}{dx}$, where U_e is the velocity outside the boundary layer. In Fig. 3, pressure taps measurements show that ZPG conditions are maintained both downstream of the T-DBD, where TS waves amplify, and in the ROI, where control is applied, with $K \leq 1 \times 10^{-8}$, consistently with the criterion reported in the literature.[64]

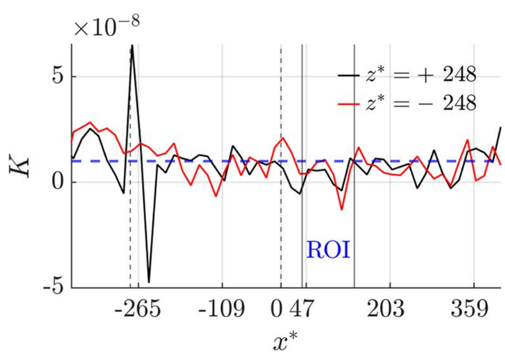

FIG. 3. Acceleration parameter K computed from pressure taps (red and black solid lines) and iso-value $K = 1 \times 10^{-8}$ (blue dashed line).

A two-layer polymethyl methacrylate (PMMA) insertion plate of size $399 \times 719\,\mathrm{mm}^2$ is flush-mounted with the flat plate to lodge the DBD plasma actuators and the microphone array. Both actuators are hand-manufactured with a 1 mm-thick Kapton dielectric layer bonded to a 8 mm-thick PMMA substrate, with exposed and covered electrodes made of self-adhesive copper tape of 50 μm thickness. The exposed electrodes have a streamwise width of 10 mm, while the covered electrode of the control actuator is widened to 25 mm in order to let the plasma expand over a longer distance and increase C-DBD control authority.[65] At both locations, the thickness of the exposed electrodes is less than 5% of the boundary layer displacement thickness. Therefore, their geometrical effect as a forward-facing step on the transition behavior of the boundary layer flow can be neglected.[66] The plasma actuators are enabled by two TREK 20/20C high voltage amplifiers, operated remotely. Only the C-DBD receives the input signal from the control system, while the T-DBD is operated independently. Although the system is capable of reaching voltage levels up to more than 2.67 V_{ref} (20 kV) peak-to-peak, the T-DBD is operated with an amplitude in the range 0.93–1.2 V_{ref} (7–9 kV) while the C-DBD is generally operated below 2.27 V_{ref} (17 kV). These measures are taken to mitigate performance degradation and guarantee longer operational lifetime.[67] The intensity of TS waves is quantified by monitoring pressure fluctuations with calibrated miniature electret condenser microphones (PUI Audio POM-2739L-HD3-LW100-R), with nominal sensitivity $14\,\mathrm{mV\,Pa^{-1}}$. Microphones are inserted in cavities from the back of the insertion plate and read data on the wall through 0.4 mm orifices on the front surface. Microphone cavities are sealed with epoxy to ensure acoustic isolation from the ambient and to prevent leakage through the mounting holes. The resulting Helmholtz-type resonance frequency of the combined tap–cavity system is around 1250 Hz, well above the frequency band of interest $\sim$100–300 Hz and therefore does not interfere with the measured TS wave content. Thirteen microphones are distributed along the streamwise centerline of the plate, from $x^* = -241.9$ ($x_s = 430$ mm) to $x^* = 136.6$ ($x_s = 915$ mm) with a spacing of $\Delta x^* = 19.5$ (25 mm) in the control region and $\Delta x^* = 39.0$ (50 mm) elsewhere. Extra care is given to the wiring and grounding of the microphones in order to minimize the electromagnetic interference coming from the high voltage levels applied to the plasma actuators. Input microphone signals and output control signal are handled by running a Simulink Real Time model on a Speedgoat Performance Target machine.[68]

E. PIV measurements

Two-component PIV measurements are carried out to visualize the effect of control on the evolution of velocity fluctuations inside the boundary layer. Two LaVision Imager sCMOS cameras are mounted laterally to the flat-plate model to acquire two adjacent fields of view (FOV1 and FOV2) downstream of the C-DBD actuator, as sketched in Fig. 1.

FOV1 is recorded using a 200 mm lens, yielding a magnification of 55 pix/mm over the field of view spanning $x^* \in [46.8,\ 78.0]$ ($x_s \in [800,\ 840]$ mm). FOV2 employs a 50 mm lens, providing 16 pix/mm over the field of view spanning $x^* \in [81.9,\ 171.7]$ ($x_s \in [845, 960]$ mm).

For both cameras, the aperture is set to $f = 8$. The images are partially clipped to obtain the reported fields of view. FOV1 provides more spatial resolution inside the boundary layer to quantify the effect of flow control on coherent disturbances via phase-locked measurements. FOV2 serves to detect the effect of control on the later stages of TS waves amplification via phase-free measurements. Illumination is provided by a Quantel Evergreen 200, a dual-cavity Nd:YAG laser releasing 200 mJ per pulse. Laser pulses are separated by a $\Delta t = 25\,\mu$s, so that particles in the near-wall vicinity would displace by about 10 pixels in FOV1 and about 6 pixels in FOV2 between the frames. The acquisition rate is regulated depending on the measurement. In phase-free measurements, it is set to 13.13 Hz avoiding phase-locking with any coherent fluctuations. In phase-locked measurements, one phase per test is captured by choosing a sampling rate of $\frac{f}{20}$, with f being the target TS frequency of the specific test case. A set of lenses (spherical, $f = -100$, spherical $+150$, cylindrical -200) and a mirror are used to focus the beam, to deflect it, and to create the 1 mm-thin laser sheet around the centerline of the flat plate. Flow seeding is dispersed using a SAFEX Fog Generator that uses a water–glycol mixture to create particles of size $1\,\mu$m. The PIV process is based on a multi-pass cross correlation of interrogation windows of size decreasing from 96×96 pixels (with 50% overlap) to 16×16 pixels (with 75% overlap). The uncertainty in PIV measurements is estimated by correlation statistics[69] and the maximum values are quantified as $\epsilon_U \approx 0.08\%\, U_\infty$ and $\epsilon_V \approx 0.05\%\, U_\infty$.

III. CLOSED-LOOP FEEDFORWARD OPTIMIZATION FRAMEWORK

In the present framework, closed-loop feedforward optimization is implemented using PSO. The optimization algorithm is employed to generate and iteratively update a population of candidate solutions, each of which is mapped to a specific control action and evaluated experimentally in the environment described in Subsection II A.

A. Signal preprocessing

The reference and error pressure signals, denoted by $r(k)$ and $e(k)$, respectively, are bandpass filtered using a fourth-order Chebyshev type-I infinite impulse response (IIR) filter with 1 dB pass-band ripple. The cutoff frequencies are set to 200 Hz and $f_{\mathrm{MAX}} + 50$ Hz. Such filtering action prevents signal contamination (actuator noise, electromagnetic interference) from masking the effect of control on the error microphone and keeps the reference signal independent from the state of actuation. As will be detailed in Subsection III B, the magnitude of the control signal depends on the magnitude of the reference oscillations. As a consequence, ensuring that $r(k)$ remains consistent prevents signal contamination from being implicitly incorporated into the optimization and minimizes the influence of time-dependent drift in the T-DBD performance.

B. Control action parametrization and evaluation

The candidate solutions generated by PSO are $D \times 1$ arrays of real-valued coefficients, within a D-dimensional search space $\mathscr{S}$. Each individual parametrizes a FIR filter acting on the phased-forward reference signal $r_A(k)$, which accounts for the convective transport of TS waves between the reference sensor and the actuator location. Assuming a TS waves convective velocity of approximately one-third of the freestream velocity,[22,70] the advection time t_{adv} between reference and C-DBD is estimated as

$$t_{adv} = \frac{x_{\mathrm{control}} - x_{\mathrm{ref}}}{0.33\, U_\infty}, \qquad r_A(k) = r(k - f_s\, t_{adv}), \tag{4}$$

where f_s is the sampling frequency of the signal. The hypothesis on the TS waves convective velocity is verified from LSA calculations, where the phase-velocity is estimated in the range 0.33–0.34 U_∞ for the frequency band of interest and in the space between the reference and the error microphones. At the generic discrete time step k, the "running" controller, denoted by $\mathbf{a}_i$, is convolved with the most recent D samples of the advected reference signal to generate the raw control output $y^*(k)$, defined as

$$y^*(k) = \mathbf{a}_i \cdot \mathbf{r}_A(k), \qquad \mathbf{r}_A(k) = [r_A(k - D + 1),\ \ldots,\ r_A(k)]^\top. \tag{5}$$

The FIR filter $\mathbf{a}_i$ implicitly applies a gain and a phase shift to the advected reference signal. These two quantities represent the physical parameters that are indirectly manipulated and optimized by the PSO-based controller. The key advantage of this parametrization is that the optimization does not rely on *a priori* knowledge of the TS waves frequency and does not require explicit online mode identification. The generality of the approach is only partially constrained by the bandpass IIR filtering applied to the reference signal, whose cutoff frequencies are selected based on the expected range of unstable TS modes.

The raw output $y^*(k)$ is subsequently amplified by a constant gain B to ensure that the actuator input voltage reaches amplitudes sufficient to affect the coherent structures within the boundary layer. To preserve the integrity and operational lifetime of the DBD plasma actuator, the amplified output $y(k) = B\, y^*(k)$ is soft-clipped below a clipping voltage V_c by using a hyperbolic-tangent limiter. The final output voltage is denoted as $V(k)$.

Regarding performance evaluation, when a control action $\mathbf{a}_i$ is activated at time t, its reward is computed using samples of the error signal $e(k)$ collected over a finite observation window, referred to as the episode time T_{ep}, inside the longer time interval during which the same controller is applied. The remaining part of the interval is discarded as a safety margin to avoid transitional effects between two consecutive control actions from biasing the reward evaluation. To account for the finite convective delay between the actuator and the downstream error sensor, the reward evaluation is delayed by an advection time from C-DBD to error computed as in Eq. (4). The reward function to be maximized is defined as

$$R = \frac{\mathrm{std}(e_0) - \mathrm{std}(e)}{\mathrm{std}(e_0)}, \tag{6}$$

where $\mathrm{std}(e)$ denotes the standard deviation of the error signal computed over the episode window of duration T_{ep}, and $\mathrm{std}(e_0)$ is the baseline standard deviation of the error signal evaluated prior to the activation of control, with the C-DBD switched off. Since the microphone signal is treated as a zero-mean fluctuation signal, its standard deviation is practically equivalent to its root mean square value and thus quantifies the amplitude of the pressure oscillations induced by the passage of TS waves. Finally, a penalty $R = -1$ is assigned whenever the root mean square value of the reference signal exceeds its uncontrolled value by more than 4%. This condition is introduced to prevent electromagnetic interference from artificially biasing the optimization process. The threshold is empirically tuned based on preliminary experimental observations. It is selected to preserve sufficient exploration of the control parameter space, including high-voltage actuation levels, while ensuring that spurious signal contamination does not affect the reward evaluation.

The reference signal is also monitored throughout the optimization as an indicator of residual electromagnetic contamination and/or T-DBD drift. In all the optimization runs, the variation of its oscillations remains one order of magnitude smaller than the variation observed at the error sensor under control, indicating that these effects have only a minor impact on the reported reward and attenuation levels.

C. Particle swarm optimization

PSO operates by iteratively improving a population, called swarm, of N_P candidate solutions, or particles, by adjusting their positions in a D-dimensional space $\mathscr{S}$ based on their own experience and the experiences of neighboring particles. Each D-dimensional position in $\mathscr{S}$ uniquely corresponds to an action that, in turn, is awarded with a reward value. The dynamic behavior of a particle, say the ith, with $i \in [1 \ldots N_P]$, at the jth swarm update, is described by the following:

$$\mathbf{v}_i(j+1) = w\,\mathbf{v}_i(j) + c_1 r_1(\mathbf{P}_i^j - \mathbf{a}_i(j)) + c_2 r_2(\mathbf{G}^j - \mathbf{a}_i(j)), \tag{7a}$$

$$\mathbf{a}_i(j+1) = \mathbf{a}_i(j) + \mathbf{v}_i(j+1), \tag{7b}$$

where $\mathbf{v}_i = [v_{i,1} \ldots v_{i,D}]$ is the "displacement" vector of the ith particle; $\mathbf{a}_i = [a_{i,1} \ldots a_{i,D}]$ is the ith particle location, identifying a controller; $\mathbf{P}_i^j = [P_{i,1}^j \ldots P_{i,D}^j]$ is the most rewarding position passed by the ith particle during its own "trajectory" up to the jth population update. Similarly, $\mathbf{G}^j = [G_1^j \ldots G_D^j]$ is the most rewarding position ever recorded by the whole swarm at the time of the jth update. $\mathbf{P}_\mathbf{i}^j$ and $\mathbf{G}^j$ are named "personal best" and "global best" positions, respectively. The three hyperparameters w, c_1, and c_2 that appear in Eq. (7) are named inertia, personal, and social learning parameters. The inertia weight w controls the previous displacement contribution to the updated one. From Eq. (7), the distances $\mathbf{a}_i(j) - \mathbf{P}_i^j$ and $\mathbf{a}_i(j) - \mathbf{G}^j$ drive the personal and social components of the displacement update, weighted with learning parameters c_1 and c_2. These contributions are also scaled with r_1 and r_2. These terms are independent random variables drawn from a uniform distribution on [0, 1] and have the purpose of preserving a certain variety in the population behavior. The mentioned parameters can be tuned to let the swarm rapidly converge to an optimum (low w, high c_1, and/or c_2) or to allow the swarm to explore different regions of the D-dimensional space (higher w). Practical convergence is assumed when particle velocities $\mathbf{v}_i(j^*)$ become sufficiently small and reward improvements stagnate.

Although the inertia and learning terms provide a mechanism to balance exploration and convergence, the performance of PSO remains strongly dependent on the initial random population, and the algorithm may prematurely converge to local maxima.[51] Here, the runtime of each test is constrained to limit the potential influence of actuators degradation; consequently, the range of hyperparameters that can be used is restricted by the need to achieve meaningful convergence within the available experimental time, hence restricting the capability of broad exploration of the variable space $\mathscr{S}$. The selected hyperparameters are reported in Table I. Adapted from previous works,[71] they satisfy the criteria $-1 < w < 1$ and $0 < c_1 + c_2 < 4(1+w)$ to ensure convergent swarm behavior.[72] At the generic time step k, the running controller is the ith particle of the swarm. Once the reward $R(i)$ is calculated with Eq. (6), the personal best position $\mathbf{P}_i$ is updated (if necessary). After all the N_P particles of the swarm are tested, the global best position $\mathbf{G}$ is also updated (if necessary) and the swarm is updated as shown in Eq. (7).

To quantify the convergence of the swarm, a swarm-dispersion metric is introduced by computing the average spread of the particle coordinates in $\mathscr{S}$ across the swarm,

$$\sigma_P = \frac{1}{D}\sum_{d=1}^{D}\sigma_{a_d}, \qquad \sigma_{a_d} = \sqrt{\frac{1}{N_P - 1}\sum_{i=1}^{N_P}(a_{i,d} - \overline{a}_d)^2}, \tag{8}$$

TABLE I. Summary of PSO optimization runs.

Parameter	L-A1	L-A2	H-A3	M-B1	M-B2
U_∞ (m/s)	20	22.5	22.5	20	22.5
δ_0^* (mm)	1.29	1.21	1.21	1.29	1.21
$Re_{\delta_0^*}$ (10^3)	1.71	1.81	1.81	1.71	1.81
f (Hz)	240	240	280	[240, 260]	[240, 260, 280]
F	57	45	52	[57, 61]	[45, 48, 52]
V_t/V_{ref}	1.07	1.07	1.2	1.21	1.13
V_c/V_{ref}	2	2	2.27	2	2
$x_{s,err}$ (m)			0.815		
N_P	50	50	20	50	50
D	10	10	10	4	6
T_{ep} (s)	0.041	0.041	0.041	0.165	0.165
$[w, c_1, c_2]$			[0.7, 1.4, 1.4]		
B	4	3	1.3	3.7	3.7
Tot. No. Episodes	1800	2500	2000	1020	1980
Optim. time (min)	1.90	2.5	2.1	4.2	8.2

where $a_{i,d}$ is the dth coefficient of particle i and $\overline{a}_d$ denotes the swarm-averaged value of that coefficient. When discussing the outcomes of the optimization runs, practical convergence is identified when the swarm-dispersion, normalized for its initial value, satisfies $\sigma_P/\sigma_P(0) < 0.2$. This threshold is used as a practical indicator of swarm clustering, rather than as a strict convergence criterion for each individual particle. In practice, once this condition was reached, the candidate control laws produced similar output-voltage waveforms and the reward evolution showed no further significant improvement. The optimization was therefore stopped to limit the experimental runtime and comply with setup limitations.

D. Multi-frequency TS waves control framework

In multi-frequency control scenarios, the FIR-based optimization framework yielded meaningful results only by driving the swarm toward statistically high voltage levels and frequently reaching the clipping threshold. Under these conditions, electromagnetic interference would significantly affect the sensor measurements in the optimization process. To overcome these limitations, a fast Fourier transform (FFT)-based optimization framework with online frequency identification is developed. In this formulation, PSO evolves a population of $2N \times 1$ parameter arrays whose components represent the amplitudes and relative phase lags of N sinusoidal components. At the sample time t, the raw control signal $y^*(t)$ is constructed as a linear superposition of the N harmonics,

$$y^*(t) = \sum_{j=1}^{N} A_j \sin\left(2\pi f_j t + \phi_{0,j} + \phi_j\right). \tag{9}$$

In Eq. (9), $f_j, j \in [1 \ldots N]$ represent the instability frequencies identified during a frequency-identification stage performed prior to the start of control. For the jth mode, A_j denotes the control amplitude, ϕ_{0j} is a "ground" phase, which will be discussed and ϕ_j denotes the control phase. The optimization parameters are extracted from the particle vector $\mathbf{a}_i$ as

$$A_j = a_i(2j-1), \qquad \phi_j = a_i(2j). \tag{10}$$

In contrast to the FIR-based approach, where amplitude and phase effects are implicitly embedded in the filter coefficients, the present formulation explicitly optimizes physically interpretable quantities. This reduces the dimensionality of the search space and constrains the control signal to a deterministic harmonic wave shape.

As anticipated, an online frequency identification phase is performed with the C-DBD switched off. During this preliminary stage, TS waves frequencies $[f_1, \ldots, f_N]$ are extracted from the reference signal $r(k)$ via spectral analysis. N peaks are identified as the N maxima of the spectrum calculated with 1024 samples of $r(k)$, yielding a frequency resolution $\approx$5 Hz. These frequencies are then stored and kept constant throughout the subsequent optimization run. As a consequence of this identification stage, the FFT-based formulation is only partially less general than the FIR-based one, since the harmonic frequencies are not prescribed *a priori* by the user but are automatically identified online from the reference signal before the optimization. Because the optimization variables include phase lags ϕ_j, a consistent phase reference must be defined for each identified mode. To this end, the reference signal is first bandpass filtered around each frequency F_j (bandwidth 5 Hz), yielding $r_j(t)$. The analytic representation of each narrowband component is then computed using the Hilbert transform,

$$z_j(t) = r_j(t) + i\,\mathscr{H}\{r_j(t)\}, \tag{11}$$

from which the ground phase is unwrapped as

$$\phi_{0j}(t) = \arg\left(z_j(t)\right), \tag{12}$$

and used in Eq. (9). Since the ground phase is updated, the optimization parameters ϕ_j effectively represent phase offsets relative to the current incoming TS wave phase and ensure that the optimization remains coherent with the evolving instability dynamics. In particular, small phase drifts of the incoming TS waves during runtime do not degrade performance, since the control law is continuously referenced to the instantaneous phase of each mode.

The FFT-based formulation introduced here should therefore be regarded as a further demonstration of the developed PSO-driven, real-time active flow control framework. The FIR-based approach is not intrinsically limited to single-frequency disturbances, since the generated control signal inherits the spectral content of the measured reference signal. In this regard, future work will further assess the FIR-based framework for multi-frequency and broadband disturbances and determine its practical operating envelope.

E. Test cases

This section introduces the experimental test cases and identifies the key outputs used to assess the effectiveness of the AFC optimization. The PSO-based optimization tests considered in this study are summarized in Table I. The first two tests, denoted L-A1 and L-A2, investigate flow control of low-amplitude monochromatic TS waves excited at non-dimensional frequencies $F = 57$ and 45 ($f = 240$ Hz) for $Re_{\delta_0^*} = 1.71 \times 10^3$ ($U_\infty = 20$ m/s) and 1.81×10^3 ($U_\infty = 22.5$ m/s), respectively. The T-DBD is driven by a sinusoidal signal with amplitude $V_t = 1.07\,V_{ref}$ (8.0 kV) peak-to-peak, while the C-DBD signal is soft-clipped below $V_c = 2V_{ref}$ (15 kV) peak-to-peak. In both cases, a swarm composed of $N_P = 50$ particles is employed to optimize $D = 10$ FIR filter coefficients. These values are tuned during the experimental campaign in order to find a trade-off between design space exploration, convergence and control overall algorithm performance. A longer FIR filter would increase the control-law flexibility, but also the number of optimization variables and the experimental time required for convergence. However, the selected FIR length spans approximately half of the TS waves period, and it is found to provide sufficiently scattered gain/phase-shift combinations while keeping the optimization problem low-dimensional.

In order to choose the episode time, a base frequency is set to $f_{base} = 240$ Hz and the T_{ep} window is sized to contain $[T_{ep} \cdot f_{base}] = 10$ oscillation periods of f_{base}, ensuring statistically meaningful performance estimates.

The test H-A3 represents a variation of the single-frequency scenario. The global Reynolds number is fixed to the highest value, the non-dimensional frequency is set to $F = 52$ ($f = 280$ Hz), and the trigger voltage is increased to $1.2\,V_{ref}$ (9.0 kV) peak-to-peak. The control framework is tested with the task of controlling higher-amplitude TS waves that would otherwise promote saturation and boundary layer transition within the ROI defined in Subsection II C. Given the higher disturbance, the C-DBD clipping level is brought to

$V_c = 2.27V_{ref}$ (17 kV) and $N_P = 20$ particles are used to speed up convergence given the increased actuator loading on the control actuator.

Test cases M-B1 and M-B2 address multi-frequency TS waves control scenarios, where two (M-B1) and three (M-B2) modes are triggered.

For these optimization runs, the FFT-based optimization framework with online frequency identification described in Subsection III D is developed and applied. This approach proved more effective than the FIR-based approach, likely owing to the reduced complexity of the optimized controllers and to the deterministic nature of the control signal driving the C-DBD actuator, despite the reduced clipping voltage $V_c = 2V_{ref}$ (15 kV). In the presence of multi-frequency disturbances, and building on previous works, the reward is evaluated over a time window corresponding to $[T_{ep} \cdot f_{base}] = 40$ repetitions of the base frequency f_{base}, while the amplification gain is fixed to $B = 3.7$.

For all flow-control scenarios, the PSO hyperparameters are kept constant and selected according to the guidelines described in Subsection III C, resulting in the parameter set $[w, c_1, c_2] = [0.7, 1.4, 1.4]$. Although the optimal balance between exploration and convergence may depend on the search-space dimensionality and the reward landscape of the specific flow problem, a systematic hyperparameters optimization or sensitivity analysis would require repeated runs that were not compatible with the available experimental runtime and actuator-lifetime constraints. The reported results should therefore be interpreted as demonstrating practical convergence with a stable parameter set, rather than as the best attainable PSO performance for each framework. For single-frequency control scenarios, the gain B is determined empirically by computing the root mean square voltage values of a set of randomly generated controllers. B is selected such that a prescribed fraction (typically 50%) of the resulting control signals exhibits a root mean square value larger than that of a reference sinusoid with peak-to-peak voltage $V_{pp} = 1.33V_{ref}$ (10 kV).

When discussing the outcomes of the optimization runs, the evolution of the control strategies is illustrated through the time history of the reward and the progressive convergence of the swarm toward the best-performing controllers.

The effectiveness of the converged control strategies is quantified using microphone measurements analyzed in both the time and frequency domains. In the time domain, the reference and error signals under converged control are presented together with the corresponding actuator voltage signal that achieves the observed attenuation.

Moreover, to complement the local reward evaluated at the error microphone, the spatial extent of the control effect is assessed using the microphone array. In the frequency domain, the effect of the optimized control is assessed through the FFT spectra reconstructed from pressure measurements along the streamwise direction. This analysis highlights the attenuation of the dominant instability tones and, where present, the modification of broadband fluctuations.

Additionally, for each microphone downstream of the C-DBD ($x^* > 0$), the pressure spectral energy is integrated over the unstable frequency band $200 \le f \le 300$ Hz,

$$E_b(x^*) = \int_{200}^{300} |\hat{p}(f, x^*)|^2 \, df. \tag{13}$$

The resulting quantity, $E_b(x^*)$, provides a streamwise distribution of the disturbance energy in uncontrolled and controlled conditions. A further spatial integration of $E_b(x^*)$ along the downstream array is used to discuss an overall band-energy reduction due to the effect of converged control,

$$\overline{E_b} = \int_{x^* \in \mathrm{ROI}} E_b(x^*) \, dx^*, \tag{14}$$

allowing the effectiveness of the controller to be evaluated over the entire monitored region rather than only at the reward location.

Although the present reward is evaluated at a single error microphone, future implementations could incorporate measurements from multiple downstream microphones or a spatially integrated band-energy metric $\overline{E_b}$. This would quantify the level of control with the fluctuations attenuation over the entire monitored region rather than only at one location.

Tests L-A1, H-A3, and M-B2 are complemented with phase-locked and phase-free PIV measurements to visualize the effect of control on velocity fluctuations within the boundary layer. First, a baseline laminar boundary layer field is computed by averaging PIV fields acquired with both the T-DBD and C-DBD actuators switched off, yielding $\overline{u^*_{\mathrm{baseline}}}$ and $\overline{v^*_{\mathrm{baseline}}}$. Phase-locked PIV measurements are then used to extract the phase-coherent disturbance fields in the uncontrolled case (T-DBD on, C-DBD off, abbreviated in subscripts as "T") and in the controlled case (both actuators on, abbreviated in subscripts as "TC"). On the other hand, in phase-free measurements, the effect of control is assessed with the evolution of velocity fluctuations in the uncontrolled and controlled cases, while the mean fields $\overline{u^*_T}$, $\overline{v^*_T}$, $\overline{u^*_{TC}}$, and $\overline{v^*_{TC}}$ are statistically equal to the baseline mean flow, apart from some potential effects due to the actuator operation, which are discussed in Subsection IV A.

For each control scenario (T or TC), the phase-coherent disturbances are extracted by subtracting the baseline mean velocity field from the phase-averaged velocity field,

$$u' = \langle u^* \rangle - \overline{u^*}, \qquad v' = \langle v^* \rangle - \overline{v^*}, \tag{15}$$

where the brackets $\langle\,\rangle$ denote the time average operator applied on phase-locked fields, while the overbar refers to the mean computed for the specific control scenario in phase-free measurements ($\overline{u^*_T}$, $\overline{v^*_T}$, $\overline{u^*_{TC}}$, and $\overline{v^*_{TC}}$).

In phase-free measurements, the fluctuation levels are quantified via the standard deviation of the velocity fields u_{std} and v_{std} and by extracting a pseudo-turbulent kinetic energy (TKE),

$$u_{\mathrm{std}} = \mathrm{std}(u^*) \quad \text{and} \quad v_{\mathrm{std}} = \mathrm{std}(v^*), \tag{16a}$$

$$TKE = \frac{u_{std}^2 + v_{std}^2}{2}. \tag{16b}$$

Differences in the downstream growth and spatial distribution of these fluctuations provide a quantitative measure of the control-induced delay in the amplification of TS waves.

IV. RESULTS AND DISCUSSION

The present section introduces the outcomes of the explored test cases. Tests L-A1, H-A3, and M-B2 are here thoroughly discussed. The remaining tests (L-A2 and M-B1) are reported in the Appendix.

A. Low-amplitude single-frequency TS waves control

The time history of the extracted reward during the PSO run for case L-A1 is reported in Fig. 4. The optimization rapidly identifies controllers achieving promising rewards and drives the swarm toward progressively more effective solutions. Over the first $\sim$1000 episodes, the swarm-averaged reward $\overline{R}$ increases approximately linearly, before saturating at $\overline{R} \approx 0.4$, with practical convergence reached after approximately 1200 episodes. In parallel, the normalized dispersion of the reward across the swarm, quantified by $\sigma(R)/\overline{R}$, decreases markedly and reaches a residual level of $\sigma(R)/\overline{R} \approx 0.20$. Residual dispersion reflects incomplete convergence caused by the stochastic nature of the swarm updates. Such variability is quantified by the swarm-spread metric [Eq. (8)], which decreases to approximately $\sigma_P = 16\%$ of its initial value, consistent with partial, but not complete, concentration of the swarm.

The effect of the converged control strategy is illustrated by the time histories in Fig. 5. The shown controlled segment, recorded around episode ≈ 1728 ($t \approx 108$ s from the start of the optimization), is compared to an uncontrolled segment acquired prior to the start of the optimization run. The shown realization comprises 12 periods T of the main mode $F = 57$ ($f = 240$ Hz), that is, $\tau = t/T = t/(1/f) \leq 12$. The uncontrolled error signal $e(k)$ oscillates with an amplitude of $1.5\,p_{ref}$, reduced by 50% by the converged control. The corresponding actuation voltage and reference signals indicate that the optimized FIR filter produces a control signal that is nearly in phase with the incoming TS waves at the C-DBD position, represented by the phased-forward reference $r_A(k)$.

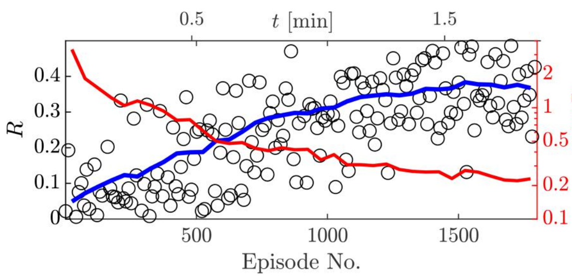

FIG. 4. L-A1 case (optimization run)—time history of the reward R (black circles, displayed with a downsampling factor of 10), swarm-averaged reward $\overline{R}$ (blue line), and swarm-dispersion of reward $\sigma(R)/\overline{R}$ (red line).

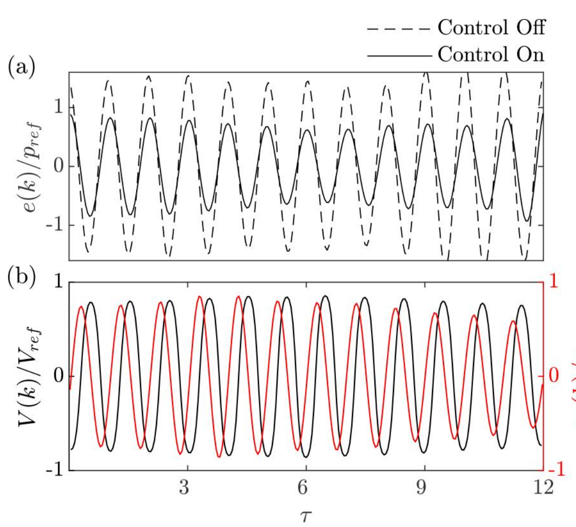

FIG. 5. L-A1 case (optimization run)—effect of converged control: (a) Error signal $e(k)$ and (b) applied control voltage $V(k)$ (black line) and phased-forward reference $r_A(k)$ (red line).

Notably, despite the larger clipping threshold available (Table I), the optimization converges to a solution that does not reach the maximum voltage range. This is consistent with the early convergence limitations of PSO as well as with the imposed penalty on the reference root mean square value, so that some high-voltage controllers might have been discarded due to the electromagnetic interference.

The frequency-domain effect of converged control is shown in Fig. 6, where the spectra of the non-dimensional pressure signals are reported. Spectra are calculated inside a 80-episodes (5 s) window. The controlled spectra are calculated at the convergence of the optimization run, in a window centered around episode 1728. In addition to displaying a small peak at $F = 59$ ($f = 250$ Hz), due to electromagnetic interference, both uncontrolled and controlled spectra exhibit a distinct peak at $F = 57$ ($f = 240$ Hz). By tracking the spectral peak along the flat plate, it can be seen that the $F = 57$ mode continuously amplifies in the control region. However, due to the control effect, the peak magnitude downstream of the C-DBD is reduced by approximately 50%. This is then confirmed by the reduction of the metric introduced in Eq. (14), similarly demonstrating an overall reduction of approximately 59% over the unstable frequency band, confirming that the attenuation observed at the error microphone is not limited to the reward location.

Regarding the PIV results, a first outcome is the verification of the mean flow as a validation of the initial hypothesis underlying the stability calculations used in the experiment design. The phase-free mean profiles extracted from the masked PIV fields are reported in Fig. 7. The profiles are shown at two streamwise stations downstream of the C-DBD (one for each field of view) and compared with the

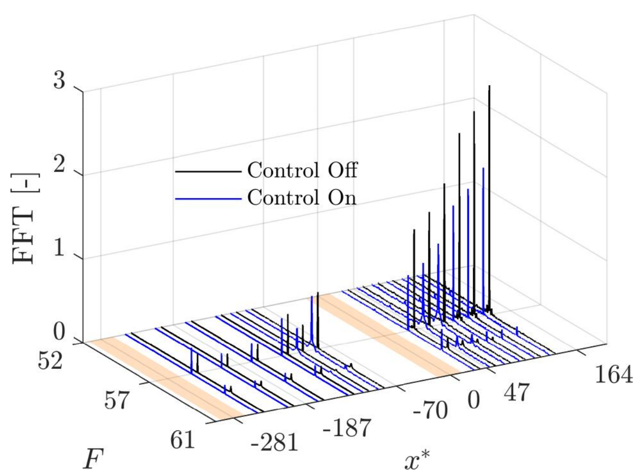

FIG. 6. L-A1 case (optimization run)—FFT amplitude of the microphone signals as a function of the non-dimensional streamwise coordinate x^*. DBD plasma actuators are represented with copper-like patches.

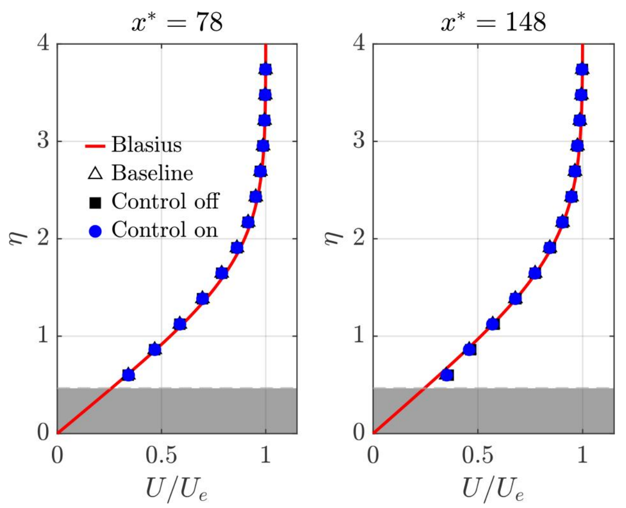

FIG. 7. L-A1 case (converged retest)—extracted wall-normal profiles of the mean streamwise velocity U, for baseline, uncontrolled and controlled conditions. Red line represents the Blasius profile used for the stability calculations.

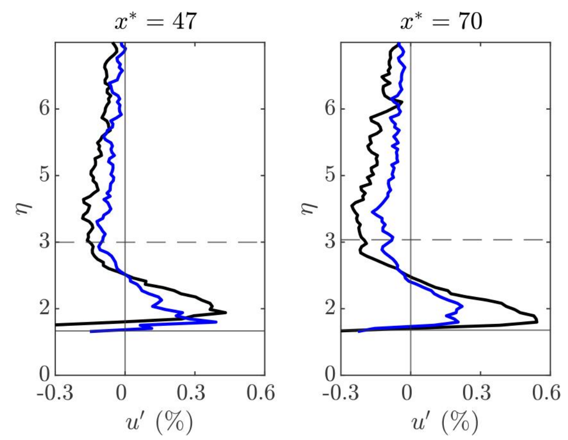

FIG. 8. L-A1 case (converged retest)—extracted wall-normal profiles of the phase-averaged streamwise disturbance u', in uncontrolled (black) and controlled (blue) conditions. Black lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$.

Blasius zero-pressure-gradient boundary-layer solution. The region affected by near-wall reflections and/or lack of data due to partial shading is masked and excluded from the interpretation. Outside this region, the measured baseline, uncontrolled, and controlled profiles nearly collapse and remain close to the Blasius profile. This indicates that, within the measured fields of view, the activation of the trigger and control actuators does not produce a measurable modification of the mean boundary layer. The reduction observed in the phase-locked disturbance fields is therefore interpreted primarily as attenuation of the coherent TS wave component rather than as a consequence of mean-flow modification. The agreement with the Blasius profile further supports its use as the base flow for non-dimensionalization and for the linear stability calculations adopted in the present work.

Phase-locked PIV measurements are shown in Figs. 8 and 9. As shown in Fig. 8, in FOV1, the uncontrolled TS waves have a relatively low amplitude (0.4%–0.6% U_∞), within the linear growth range of TS amplification ($\leq$1% U_∞). From Fig. 9, the estimated wavelength is $\lambda \approx 21\delta_0^*$ (28 mm), which is in line with the LSA estimates. Moreover, TS waves feature a dual-peak shape in the streamwise disturbance, placed approximately at a wall-normal distance equal to the local $\delta_0(x^*)$ $(y/\delta_{99}(x^*) = 0.34)$ and close to the local $\delta_{99}(x^*)$ $(y/\delta_{99}(x^*) = 1.0)$, respectively.[22,73] The whole development of TS waves across FOV1 and FOV2 is shown in Fig. 9. In FOV2, TS waves display a more regular structure and amplify up to 0.9% U_∞, but remain inside the linear growth regime.

The effect of converged control inside FOV1 is illustrated in Fig. 8. For both extracted profiles, both peaks of the TS wave seem reduced and the disturbance rapidly goes to zero outside the boundary layer. Close to the wall, u' is generally reduced and the bottom peak is reduced by half of its uncontrolled value at $x^* = 70$, in agreement with pressure measurements in time domain and frequency domain (Figs. 5 and 6). The effect of control on phase-locked disturbances is higher[26] or in the same order of magnitude[22] as observed in previous experimental applications, for generally more intense TS waves.

The spatial evolution of non-coherent velocity fluctuations inside the boundary layer, extracted by phase-free PIV measurements inside FOV2, is reported in Fig. 10. These fluctuations are quantified by the TKE, which generally increases in the streamwise direction as well as inward from the external flow toward the wall. Fluctuations are generally small and reach 0.1% of the squared freestream velocity U_∞^2 close to the TS waves bottom peak ($\eta \approx 1.5$) and 0.05% of U_∞^2 in the outer peak. At $\eta = 2$, the levels are even lower ($\leq$0.05% U_∞^2). This may be expected from the structure of TS waves observed in Fig. 8, since the streamwise disturbance profiles reach zero slightly above $\eta = 2$. With regard to the effect of control, the effect on the

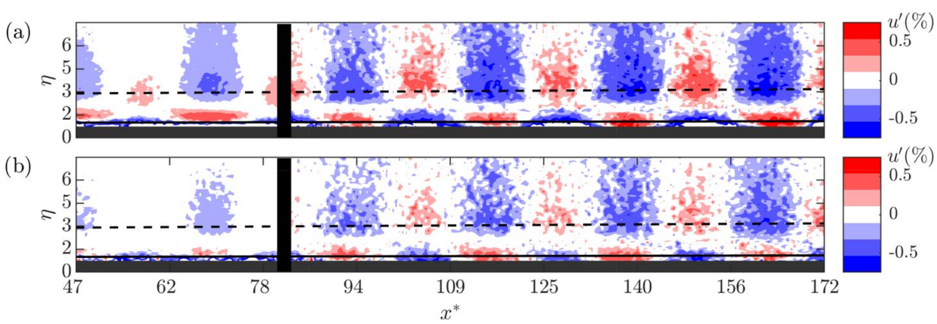

FIG. 9. L-A1 case (converged retest)—phase-averaged streamwise disturbances u', in uncontrolled (a) and controlled (b) conditions. Black lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$. Black rectangles mask the gap between FOV1 and FOV2.

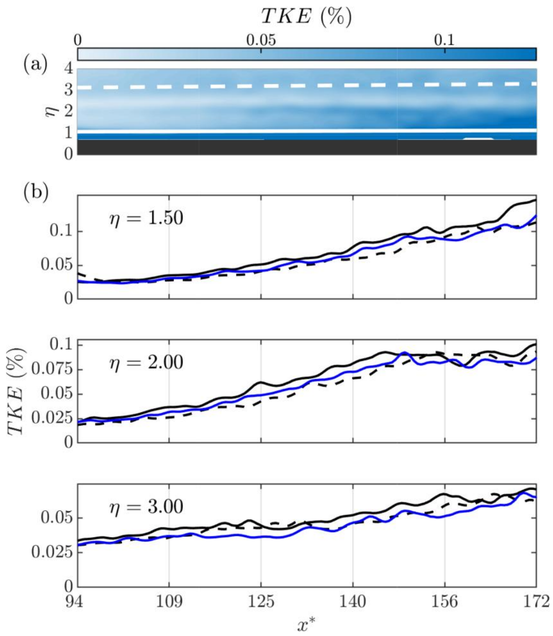

FIG. 10. L-A1 case (converged retest)—velocity fluctuations in phase-free PIV measurements. (a) Contour map of *TKE* computed in the uncontrolled case. White lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$ (dashed). (b) From top to bottom, *TKE* streamwise profiles computed in uncontrolled (black solid), uncontrolled and shifted by 8.5 δ_0^* (black dashed) and controlled (blue) conditions, at $\eta = 1.5$, 2.0, and 3.0.

fluctuations magnitude is only clear on the most downstream region $x^* > 145$. Despite the magnitude of fluctuations not being significantly reduced, the converged control can delay their growth inside the boundary layer. The quantification of the lag of the controlled curve is not accurate due to sensitivity limitations. However, especially in later stages of boundary layer development, the controlled evolution approximately lags behind the uncontrolled trend of $\approx 8.5\,\delta_0^*$, as shown by black dashed lines in Fig. 10 for three stations inside the boundary layer. These lines represent the streamwise evolution of TKE, shifted by 8.5 δ_0^* for comparison with controlled (blue) TKE profiles.

B. High-amplitude single-frequency TS waves control

Owing to the combined effect of higher freestream velocity, higher excitation frequency, and increased actuation voltage, the H-A3 test case is characterized by the generation of TS waves with significantly larger amplitudes. Despite these more demanding flow conditions, PSO successfully converges to reward values of approximately 30%, comparable to those obtained in the linear L-A1 case, and within a similar number of episodes [Fig. 11(a)]. Moreover, the practical convergence is reached significantly earlier (after around 400 episodes), while at the end of the optimization run both the swarm-spread metric and the fluctuations of the reward are approximately one order of magnitude lower than those observed in the linear TS waves control scenario. This behavior is likely attributable to the higher voltage levels applied to the C-DBD and to the reduced number of particles employed in the swarm, promoting faster convergence. Around episode 500, the dispersion metric reaches values comparable to those observed in the linear case, after which the population continues clustering around the best-performing controller. In Fig. 11(b), signals extracted around episode $\approx$1730 are displayed to represent the effect of converged control. As previously mentioned, the converged control signal has higher amplitude, reaching 2 V_{ref} peak-to-peak, and it is in phase with the incoming TS waves at $x_{control}$.

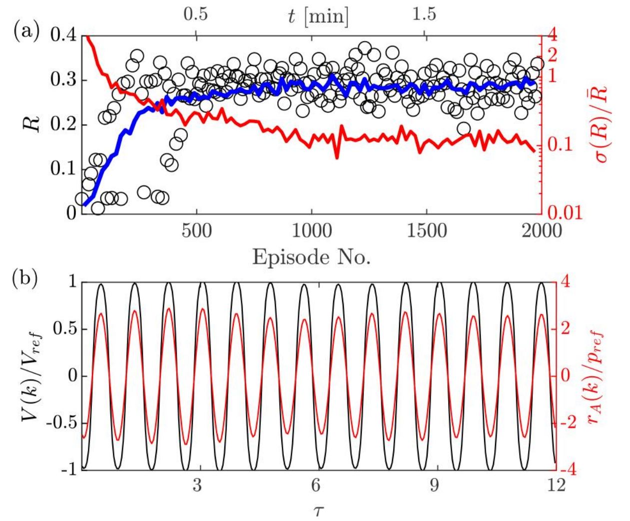

FIG. 11. H-A3 case (optimization run)—summary of PSO optimization and converged control. (a) time history of the reward R (black circles, displayed with a downsampling factor of 10), swarm-averaged reward $\overline{R}$ (blue line) and swarm-dispersion of reward $\sigma(R)/\overline{R}$ (red line). (b) At convergence, applied control voltage $V(k)$ (black line) and phased-forward reference $r_A(k)$ (red line).

Further insight into the dynamics of the controlled flow is provided by the spectral analysis of the microphone measurements, shown in Fig. 12. In contrast to the linear cases, spectra in H-A3 exhibit significantly broader frequency content, reflecting an increased nonlinearity of the disturbances. As expected from LSA, the main mode at $F = 52$ ($f = 280\,\text{Hz}$) does not show a monotonic behavior, as it peaks inside the control region at $x^* = 103.5$ ($x_s = 865\,\text{mm}$), before decreasing.

Although only the main mode is targeted by the control, broadband fluctuations are reduced as well downstream of the C-DBD, indicating weaker non-linear effects in the controlled conditions. This is clarified in Fig. 13, showing the streamwise profile of $\Delta\text{std}(p) = \text{std}_T(p) - \text{std}_{TC}(p)$, that is, the difference in the standard deviation of microphone signals after filtering out the main mode. In particular, Fig. 13 demonstrates that the broadband fluctuations are systematically lower in the controlled condition, particularly in the downstream region where nonlinear interactions become significant. As a consequence, the band-integrated spectral-energy distribution $E_b(x^*)$ is reduced under controlled conditions, and spatial integration along the microphone array gives an overall reduction of approximately 42%.

Phase-locked PIV measurements provide more insights into TS waves structure and evolution. Wall-normal profiles of u' extracted in FOV1 (Fig. 14) are significantly higher than the L-A1 case, exceeding the 1.5% U_∞ amplitude inside the FOV1 and displaying an estimated

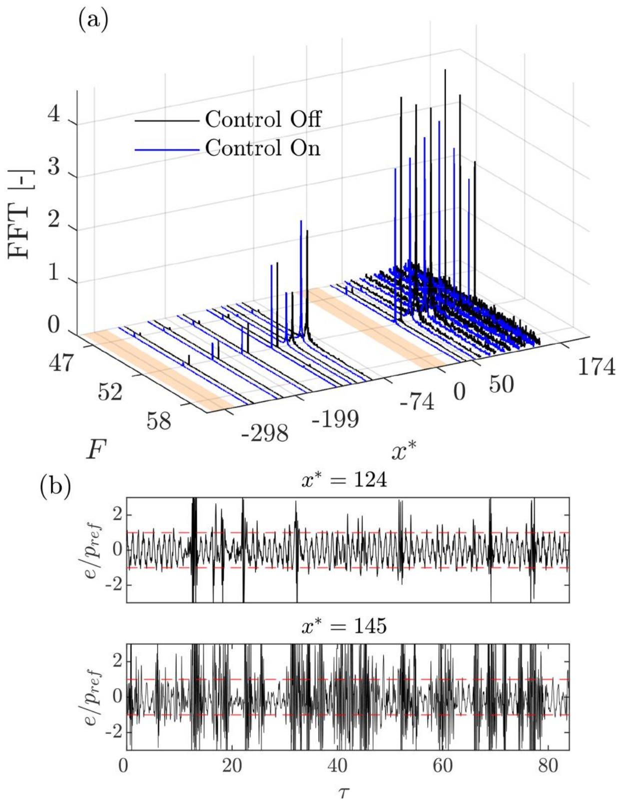

FIG. 12. H-A3 case (converged retest)—(a) FFT amplitude of microphone signals as a function of the non-dimensional streamwise coordinate x^*. DBD plasma actuators are represented with copper-like patches. (b) Time signal of downstream microphones, characterized by pressure spikes consistent with turbulent spots appearing in the flowfield.

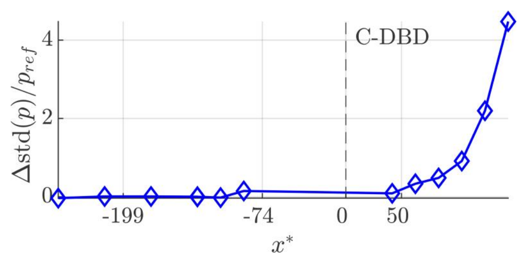

FIG. 13. H-A3 case (optimization run)—Drop in the standard deviation of wall-pressure signals after filtering out the 280 Hz mode.

wavelength of 23 δ_0^* (29 mm). Similarly to the L-A1 case, the effect of control is more appreciable at later stages of TS development. In fact, while the upper peak appears reduced at both stations $x^* = 50$ and 74, the bottom peak is reduced by 55% only in the latter station.

Also, the whole evolution displayed in Fig. 15 shows that the uncontrolled TS waves exhibit nonlinear behavior inside the FOV2, with positive and negative u' regions near the wall increasingly overlapping and distorting, ultimately leading to turbulent transition. This behavior is partially delayed by the C-DBD. In fact, the

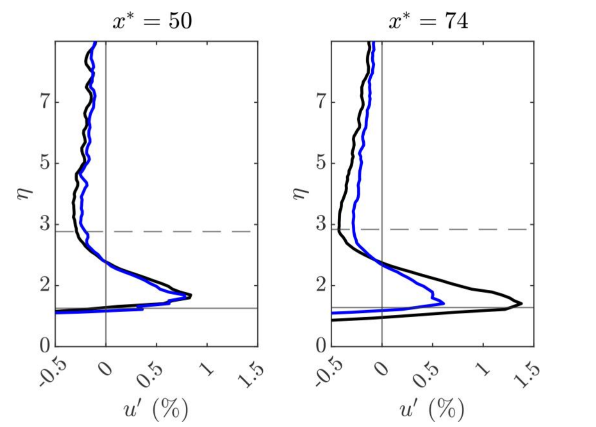

FIG. 14. H-A3 case—extracted wall-normal profiles of the phase-averaged streamwise disturbance u', in uncontrolled (black) and controlled (blue) conditions. Black lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$.

u' disturbances appear reduced within the linear growth regime and a more regular structure of the disturbances close to the wall persists up to $x^* = 120$.

On the other hand, as displayed in Fig. 16, phase-free velocity fluctuations recorded in FOV2 are larger, up to 1% U_∞^2 inside the boundary layer. The growth of disturbances in the uncontrolled case along the streamwise direction is mostly linear at the wall distance $\eta = 1.5$, while a change of the slope can be observed at $\eta = 2.0$ and 3.0 around $x^* = 150$. With regard to the effect of control, similarly to L-A1 case, it causes a delay in the growth of the fluctuations of $\approx 15\,\delta_0^*$.

C. Multi-frequency TS waves control

In preliminary tests addressing multi-frequency flow control, TS waves are triggered inside the boundary layer by linear superposition of sine signals at multiple tones with random relative phase shifts. Flow control optimization is performed with the FFT-based framework (Subsection III D). In test case M-B2, three tones $F = 45$, 48, 52 ($f = 240$, 260, and 280 Hz) are triggered. The multi-tonal behavior of the pressure signals is displayed in the FFT spectra along the flat plate (Fig. 17). Figure 17 shows that the interaction between the triggered modes introduces an additional mode at $F = 41$ ($f = 220$ Hz), likely resulting from nonlinear modal interactions. Starting with smaller amplitudes, it grows to amplitudes comparable to those of the triggered frequency modes downstream of the C-DBD.

Compared to low-amplitude single-frequency TS waves control, the optimized controller results in similar levels of reward and swarm-dispersion after an equivalent number of episodes (Fig. 18), and practical convergence is reached after 1550 episodes. Here, the swarm-dispersion of the reward is defined with the absolute value of the mean reward in the denominator as its values at the early stages of the optimization are negative (fluctuations increase). Figure 19 represents time signals extracted around episode 1930, with $\tau = t/T = t/(1/f)$ based on the $F = 52$ ($f = 280$ Hz) mode. The control signal occurs as "bursts" of the $F = 52$ mode, in opposition to the incoming bursts represented by the phased-forward reference $r_A(k)$, which

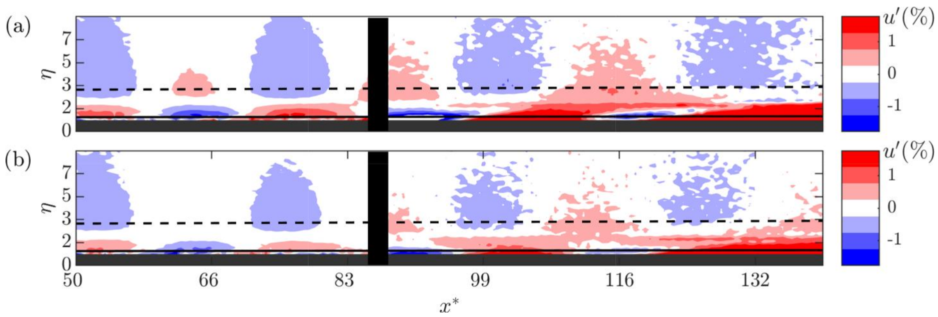

FIG. 15. H-A3 case (converged retest)—phase-averaged streamwise disturbances u', in uncontrolled (a) and controlled (b) conditions. Black lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$. Black masks fill the gaps between FOVs.

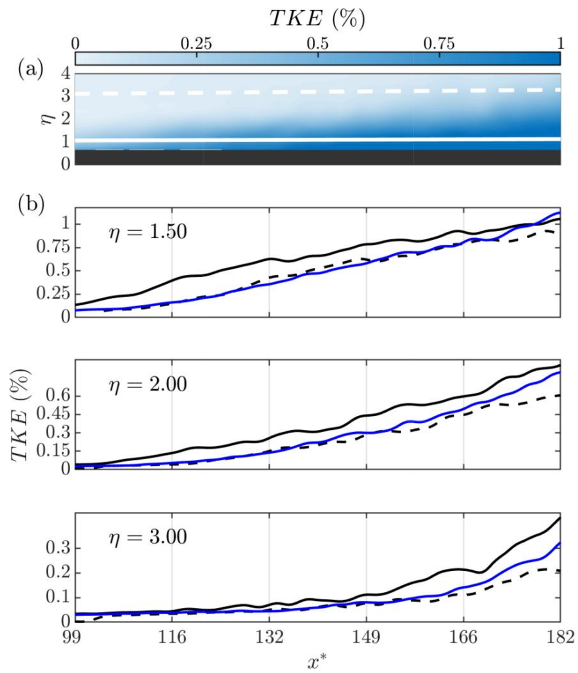

FIG. 16. H-A3 case (converged retest)—velocity fluctuations in phase-free PIV measurements. (a) Contour map of *TKE* computed in the uncontrolled case. White lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$ (dashed). (b) From top to bottom, *TKE* streamwise profiles computed in uncontrolled (black solid), uncontrolled and shifted by 15 δ_0^* (black dashed) and controlled (blue) conditions, at $\eta = 1.5$, 2.0, and 3.0.

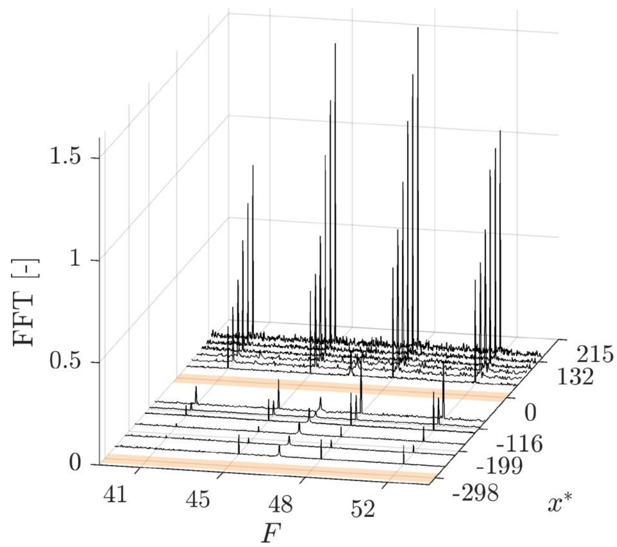

FIG. 17. M-B2 case (optimization run)—FFT amplitude of the microphone signals as a function of the non-dimensional streamwise coordinate x^*, in uncontrolled conditions. DBD plasma actuators are represented with copper-like patches.

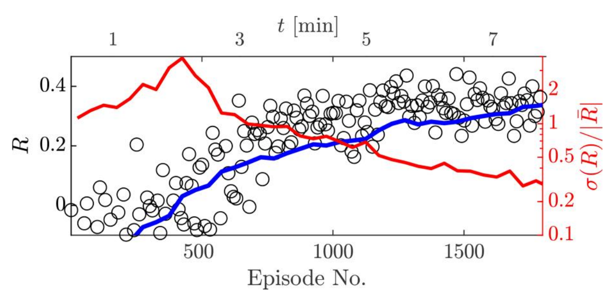

FIG. 18. M-B2 case (optimization run)—time history of the reward R (black circles, displayed with a downsampling factor of 10), swarm-averaged reward $\bar{R}$ (blue line) and swarm-dispersion of reward $\sigma(R)/|\bar{R}|$ (red line).

are produced from the fact that the triggered frequencies are closely spaced. The PSO therefore converges to a control signal whose largest envelope amplitudes are approximately opposed to the largest incoming disturbance modes.

The overall effect of the controller in the frequency domain is displayed in Fig. 20, where the difference between the FFT in uncontrolled and controlled conditions is represented. Blue peaks indicate attenuation, while red ones amplification. Error oscillations are more effectively reduced at $F = 48$. Additionally, modes at $F = 45$ and 52 are also weakened, while only the additional mode at $F = 41$ appears to be amplified by the control action.

For this reason, it is likely that the PSO optimization converged to a local optimum in which the dominant targeted components are attenuated at the cost of a moderate increase in this weaker, non-targeted mode. Future studies should address this aspect more systematically by performing multiple independent optimization runs in the multi-frequency case, in order to outline the envelope of PSO performance in terms of convergence and/or disturbance attenuation.

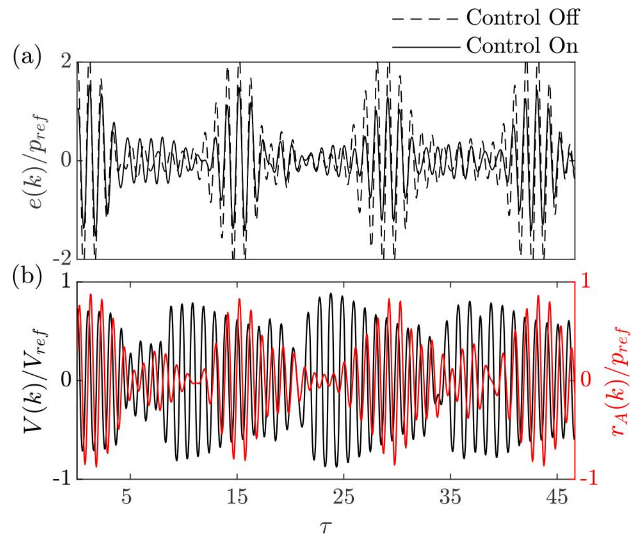

FIG. 19. M-B2 case (optimization run)—effect of converged control: (a) Error signal $e(k)$ and (b) applied control voltage $V(k)$ (black line) and phased-forward reference $r_A(k)$ (red line).

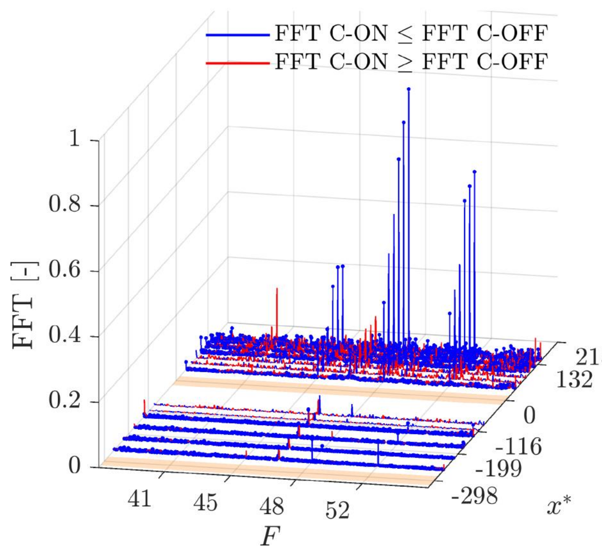

FIG. 20. M-B2 case (optimization run)—streamwise evolution of the FFT amplitude difference of the microphone signals between uncontrolled and controlled conditions. Blue lines represent FFT amplitude reduction, and red ones indicate an increase. DBD plasma actuators are represented with copper-like patches.

Such studies could also assess whether improved solutions can be obtained by modifying the reward function or by tuning the PSO hyperparameters.

However, although the reward function is spatially local, it is not constrained to one prescribed frequency. Consequently, the optimization can reduce the dominant contribution to the local pressure fluctuation while still allowing amplification of weaker off-target modes.

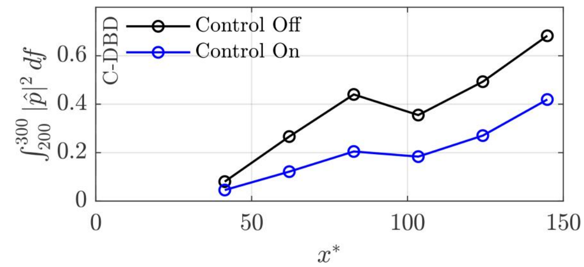

FIG. 21. M-B2 case (optimization run) streamwise evolution of the pressure spectral energy integrated over the unstable band $200 \leq f \leq 300$ Hz for the multi-frequency case, in uncontrolled and controlled conditions.

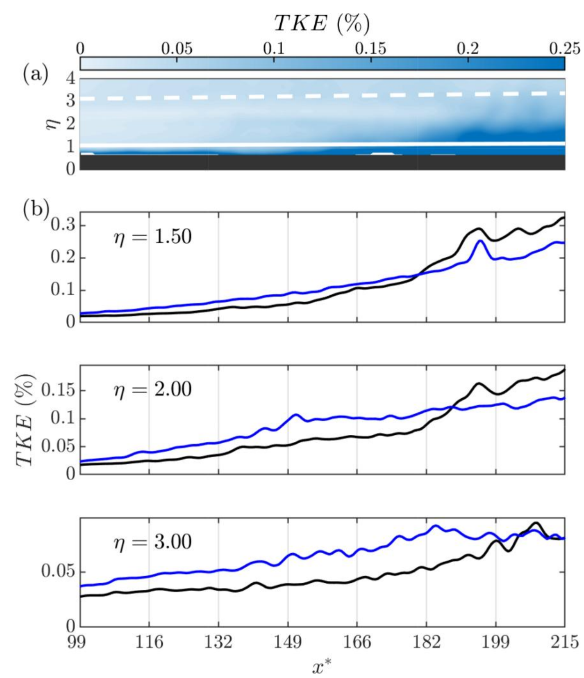

FIG. 22. M-B2 case (converged retest)—velocity fluctuations in phase-free PIV measurements. (a) Contour map of *TKE* computed in the uncontrolled case. White lines denote the displacement thickness $\delta_0(x^*)$ (solid) and boundary layer thickness $\delta_{99}(x^*)$ (dashed). (b) From top to bottom, *TKE* streamwise profiles computed in uncontrolled (black solid) and controlled (blue) conditions, at $\eta = 1.5$, 2.0, and 3.0.

The resulting distribution is reported in Fig. 21 for the uncontrolled and controlled conditions. Despite the local nature of the reward and the observed amplification of some off-target components, the controlled case shows a clear decrease in $E_b(x^*)$ along the downstream microphone array. A further integration in the streamwise direction gives an overall reduction of approximately 48%, indicating that the controller reduces the integrated spectral energy over the relevant unstable band across the downstream array.

Given the inherent multi-modal physics, only phase-free measurements are carried out and the results are shown in Fig. 22. The observed effect on velocity fluctuations is less straightforward than that observed for monochromatic TS waves. Near downstream of the

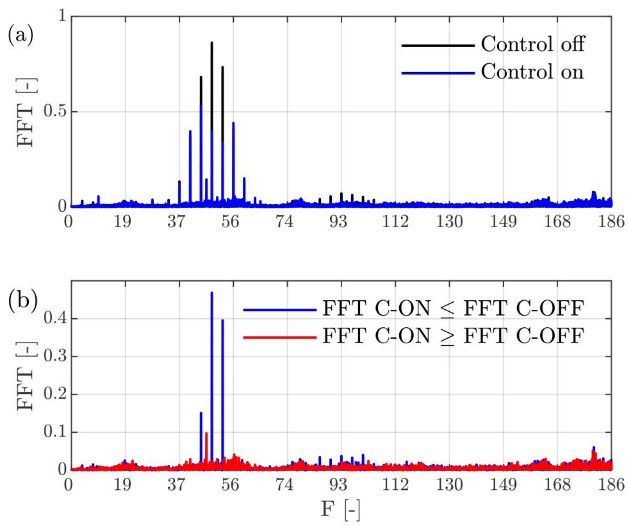

FIG. 23. M-B2 case (converged retest): (a) FFT amplitude of the microphone signals at $x^* = 78$ and (b) FFT amplitude difference of the microphone signals at $x^* = 78$.

C-DBD, during the early stages of instability development, where near-wall levels of $TKE \leq 0.1$, the applied control seems to induce larger velocity fluctuations that are not captured by microphones likely due to the filtering and the low intensity. Although the *TKE* in controlled conditions exhibits a milder streamwise growth, quantified by a smaller streamwise slope at $\eta = 1.5$, the measured fluctuation levels remain very low over a significant portion of the field, particularly for $x^* \leq 180$. Under these conditions, the PIV outcomes are less reliable than the microphone measurements shown in Fig. 20. Additionally, in Fig. 23, unfiltered microphone measurements at $x^* = 78$, sampled during the PIV test further clarify the low amplitude of disturbances and the effectiveness of control on the targeted peaks. It is also important to remark that the broadband increase observed in Fig. 23 for controlled conditions cannot be conclusively linked to a physical broadband excitation of the flow, as the contribution of electromagnetic interference induced by the plasma actuator cannot be fully decoupled from the velocity measurements.

V. CONCLUSIONS

A framework for an experimental application of particle swarm optimization (PSO)-based real-time active flow control of Tollmien–Schlichting (TS) waves on a flat plate is developed and successfully tested. The results demonstrate that candidate control laws can be evaluated and updated in real time, and that the swarm can reach practical swarm convergence to significant solutions despite measurement noise, actuator and sensor imperfections, electromagnetic interference, convective delays, voltage constraints, and limited experimental runtime.

The global Reynolds number $Re_{\delta_0^*}$ is set in the range 1.71–1.81×10^3 ($U_\infty \in 20$–22.5 m/s), and TS waves of non-dimensional frequency F in the range 41–66 ($f \in 240$–280 Hz) are triggered by using a dielectric barrier discharge (DBD) plasma actuator located upstream in the flowfield, in direct mode. The control device is also a DBD plasma actuator (C-DBD), and the control signal is generated in real time with a Speedgoat Performance[68] target machine. In single-frequency scenarios, PSO is employed to optimize FIR filters that are convolved with a reference microphone signal sampled upstream of the C-DBD to generate the control actions. In multi-frequency scenarios, PSO optimizes control actions built by a linear superposition of sinusoids at TS frequencies identified online. The control strategy achieves a 30%–40% reduction in pressure fluctuations, quantified by a microphone signal, called error signal, placed downstream of the control.

The performance of the controllers designed by PSO is further validated through complementary PIV measurements, which quantify the modification of coherent and non-coherent velocity fluctuations within the boundary layer.

In a field-of-view spanning $x^* \in [46.8,\ 78.0]$ ($x_s \in [800, 840]$ mm), uncontrolled coherent disturbances reach amplitudes 0.4%–1.5% U_∞. Once the converged controller is applied, these coherent disturbances are observed to be weakened up to 50% of their uncontrolled peak value. Within the field of view spanning $x^* \in [81.9, 171.7]$ ($x_s \in [845, 960]$ mm), non-coherent velocity fluctuations are estimated via a metric of variance. Non-coherent fluctuations reach up to 1% U_∞^2, and the effect of control is visualized by a downstream lag of the growth of fluctuations, up to 15 δ_0^*.

Given the experimental context, some setup limitations constrained the test run time and the number of control episodes that could be deployed. To ensure safe actuator operation throughout the experimental campaign, the C-DBD voltages are clipped or the run time of the optimization process is restricted, hence bringing some constraints on the variables space exploration with PSO. Future developments of the framework could address some of the setup limitations and allow control runs with more control episodes, allowing a more systematic sweep of hyperparameters for assessing performance in terms of convergence and TS waves suppression.

ACKNOWLEDGMENTS

The work has received financial support from the TU Delft AIFluids lab and the European Research Council (Grant No. 101125132). The first author expresses gratitude toward the Department of Flow Physics and Technology of TU Delft, the hosting institution where experiments took place. Special thanks to the technical supporting staff of the Low Speed Laboratory Emiel Langedijk and Stefan Bernardy for their role in supporting the work.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

G. Salomone: Conceptualization (equal); Data curation (lead); Formal analysis (lead); Investigation (lead); Methodology (equal); Software (lead); Visualization (equal); Writing – original draft (lead); Writing – review & editing (lead). **B. Mohammadikalakoo:** Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Supervision (equal);

Writing – review & editing (equal). **S. García Villasol:** Formal analysis (equal); Methodology (equal); Software (supporting). **N. A. K. Doan:** Conceptualization (equal); Formal analysis (equal); Methodology (equal); Supervision (equal); Writing – review & editing (equal). **C. S. Greco:** Conceptualization (equal); Formal analysis (equal); Methodology (equal); Supervision (equal); Writing – review & editing (equal). **M. Kotsonis:** Conceptualization (equal); Formal analysis (equal); Funding acquisition (lead); Methodology (equal); Project administration (lead); Resources (lead); Supervision (lead); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX: FURTHER TEST CASES

In the present appendix, further results regarding the remaining test cases are briefly commented. These test cases are run in similar scenarios as the ones already presented.

1. Single-frequency TS waves control: L-A2 test case

In L-A2 test case, TS waves are triggered at $F = 45$ ($f = 240\,\mathrm{Hz}$), at $Re_{\delta_0^*} = 1.81 \times 10^3$ ($U_\infty = 22.5\,\mathrm{m/s}$), with the same trigger voltage $V_t = 1.07 V_{\mathrm{ref}}$ used in L-A1 test case. The uncontrolled TS amplitude is approximately 30% higher in both reference and error microphones than in L-A1, and, as displayed in Fig. 24(a), PSO optimization reduces pressure fluctuations amplitudes of similar order to L-A1 case, in both time domain and frequency domain. Interestingly, Fig. 24(b) displays that the converged controller yields a comparable phase lag between $r_A(k)$ and the voltage $V(k)$ to the phase-lag displayed in Fig. 5(b).

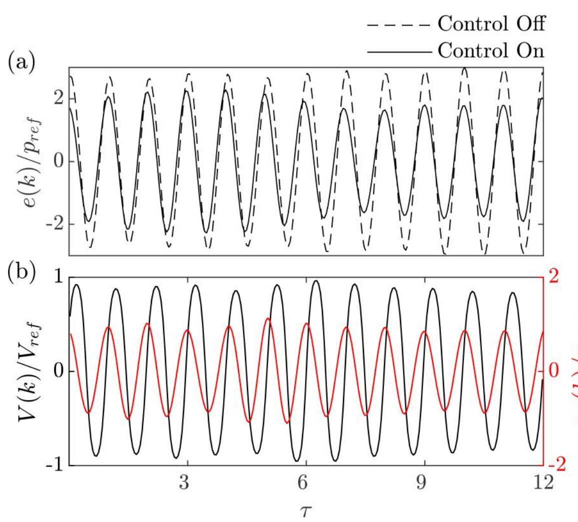

FIG. 24. L-A2 case (optimization run)—effect of converged control: (a) Error signal $e(k)$ and (b) applied control voltage $V(k)$ (black line) and phased-forward reference $r_A(k)$ (red line).

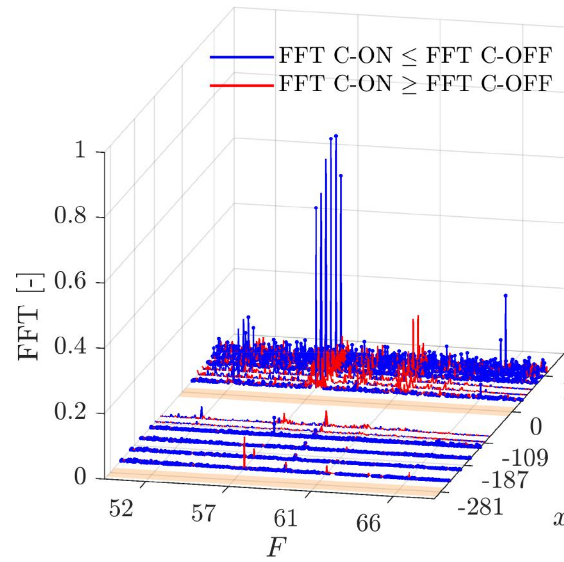

FIG. 25. M-B1 case (optimization run)–streamwise evolution of the FFT amplitude difference of the microphone signals between uncontrolled and controlled conditions. Blue lines represent FFT amplitude reduction, red ones indicate an increase. DBD plasma actuators are represented with copper-like patches.

2. Multi-frequency TS waves control: M-B1 test case

In test case M-B1, TS waves are triggered inside the boundary layer by linear superposition of sine signals at $F = 57$ and 61 ($f = 240$ and $260\,\mathrm{Hz}$), with random phase delay. Similarly to test case M-B2, Fig. 25 attests that the interaction between the triggered modes has triggered strong modes at $F = 52$ and 66 ($f = 220$ and $280\,\mathrm{Hz}$). Moreover, the converged controller does not focus on the triggered modes equally, as it found more convenient to control the main frequency $F = 57$ at the expense of the second mode $F = 61$. Despite not being able to intervene on further modes, acting on the two main modes partially relieves inter-modes interactions as both peaks at $F = 52$ and 66 are effectively reduced.

REFERENCES

[1] E. Kerschen, "Boundary layer receptivity," in 12th Aeroacoustic Conference, 1989.

[2] E. Reshotko, "Boundary layer instability, transition and control," in 32nd Aerospace Sciences Meeting and Exhibit, 1994.

[3] D. Bushnell, M. Malik, and W. Harvey, "Transition prediction in external flows via linear stability theory," in *Symposium Transsonicum III: IUTAM Symposium Göttingen* (Springer, 1989), pp. 225–242.

[4] Y. S. Kachanov, "Physical mechanisms of laminar-boundary-layer transition," Annu. Rev. Fluid Mech. **26**, 411–482 (1994).

[5] H. Schlichting and J. Kestin, *Boundary Layer Theory* (Springer, 1961), Vol. 121.

[6] J. H. Fransson, A. Talamelli, L. Brandt, and C. Cossu, "Delaying transition to turbulence by a passive mechanism," Phys. Rev. Lett. **96**, 064501 (2006).

[7] A. S. Thomas, "The control of boundary-layer transition using a wave-superposition principle," J. Fluid Mech. **137**, 233–250 (1983).

[8] J. H. Fransson, L. Brandt, A. Talamelli, and C. Cossu, "Experimental study of the stabilization of Tollmien–Schlichting waves by finite amplitude streaks," Phys. Fluids **17**, 054110 (2005).

[9] S. Bagheri, J. H. Fransson, and P. Schlatter, "Research on the interaction between streamwise streaks and Tollmien-Schlichting waves at KTH," ERCOFTAC Bull. **74**, 37–43 (2007).

[10]T. Michelis, A. Putranto, and M. Kotsonis, "Attenuation of Tollmien–Schlichting waves using resonating surface-embedded phononic crystals," Phys. Fluids **35**, 044101 (2023).
[11]H. Liepmann, G. Brown, and D. Nosenchuck, "Control of laminar-instability waves using a new technique," J. Fluid Mech. **118**, 187–200 (1982).
[12]N. R. Losse, R. King, M. Zengl, U. Rist, and B. R. Noack, "Control of Tollmien–Schlichting instabilities by finite distributed wall actuation," Theor. Comput. Fluid Dyn. **25**, 167–178 (2011).
[13]S. Walther, C. Airiau, and A. Bottaro, "Optimal control of Tollmien–Schlichting waves in a developing boundary layer," Phys. Fluids **13**, 2087–2096 (2001).
[14]R. Erdmann, A. Pätzold, M. Engert, I. Peltzer, and W. Nitsche, "On active control of laminar–turbulent transition on two-dimensional wings," Philos. Trans. R. Soc., A **369**, 1382–1395 (2011).
[15]M. Baumann and W. Nitsche, "Experiments on active control of Tollmien–Schlichting waves on a wing," in *New Results in Numerical and Experimental Fluid Mechanics: Contributions to the 10th AG STAB/DGLR Symposium Braunschweig, Germany 1996* (Springer, 1997), pp. 56–63.
[16]T. C. Corke, C. L. Enloe, and S. P. Wilkinson, "Dielectric barrier discharge plasma actuators for flow control," Annu. Rev. Fluid Mech. **42**, 505–529 (2010).
[17]M. Kotsonis, "Diagnostics for characterisation of plasma actuators," Meas. Sci. Technol. **26**, 092001 (2015).
[18]S. Grundmann and C. Tropea, "Experimental transition delay using glow-discharge plasma actuators," Exp. Fluids **42**, 653–657 (2007).
[19]A. Seraudie, O. Vermeersch, and D. Arnal, "DBD plasma actuator effect on a 2D model laminar boundary layer. Transition delay under ionic wind effect," AIAA Paper No. 2011-3515, 2011.
[20]A. Duchmann, S. Grundmann, and C. Tropea, "Delay of natural transition with dielectric barrier discharges," Exp. Fluids **54**, 1461 (2013).
[21]S. Grundmann and C. Tropea, "Active cancellation of artificially introduced Tollmien–Schlichting waves using plasma actuators," Exp. Fluids **44**, 795–806 (2008).
[22]S. Grundmann and C. Tropea, "Experimental damping of boundary-layer oscillations using DBD plasma actuators," Int. J. Heat Fluid Flow **30**, 394–402 (2009).
[23]A. Kurz, C. Tropea, S. Grundmann, N. Goldin, and R. King, "Development of active wave cancellation using DBD plasma actuators for in-flight transition control," AIAA Paper No. 2012-2946, 2012.
[24]A. Kurz, C. Tropea, S. Grundmann, M. Forte, O. Vermeersch, A. Seraudie, D. Arnal, N. Goldin, and R. King, "Transition delay using DBD plasma actuators in direct frequency mode," AIAA Paper No. 2012-2945, 2012.
[25]H. Tol, C. De Visser, and M. Kotsonis, "Experimental model-based estimation and control of natural Tollmien–Schlichting waves," AIAA J. **57**, 2344–2355 (2019).
[26]M. Kotsonis, R. K. Shukla, and S. Pröbsting, "Control of natural Tollmien-Schlichting waves using dielectric barrier discharge plasma actuators," Int. J. Flow Control **7**, 37 (2015).
[27]M. Kotsonis, R. Giepman, S. Hulshoff, and L. Veldhuis, "Numerical study of the control of Tollmien–Schlichting waves using plasma actuators," AIAA J. **51**, 2353–2364 (2013).
[28]M. Riherd and S. Roy, "Damping Tollmien–Schlichting waves in a boundary layer using plasma actuators," J. Phys. D **46**, 485203 (2013).
[29]F. Ren, H.-B. Hu, and H. Tang, "Active flow control using machine learning: A brief review," J. Hydrodyn. **32**, 247–253 (2020).
[30]F. Ren, J. Rabault, and H. Tang, "Applying deep reinforcement learning to active flow control in weakly turbulent conditions," Phys. Fluids **33**, 037121 (2021).
[31]C. Vignon, J. Rabault, and R. Vinuesa, "Recent advances in applying deep reinforcement learning for flow control: Perspectives and future directions," Phys. Fluids **35**, 031301 (2023).
[32]R. Castellanos, G. Cornejo Maceda, I. De La Fuente, B. Noack, A. Ianiro, and S. Discetti, "Machine-learning flow control with few sensor feedback and measurement noise," Phys. Fluids **34**, 047118 (2022).
[33]T. Duriez, S. L. Brunton, and B. R. Noack, *Machine Learning Control—Taming Nonlinear Dynamics and Turbulence* (Springer, 2017), Vol. 116.
[34]N. Gautier, J.-L. Aider, T. Duriez, B. R. Noack, M. Segond, and M. Abel, "Closed-loop separation control using machine learning," J. Fluid Mech. **770**, 442–457 (2015).
[35]A. Debien, K. A. Von Krbek, N. Mazellier, T. Duriez, L. Cordier, B. R. Noack, M. W. Abel, and A. Kourta, "Closed-loop separation control over a sharp edge ramp using genetic programming," Exp. Fluids **57**, 40 (2016).
[36]P.-Y. Passaggia, A. Quansah, N. Mazellier, G. Y. C. Maceda, and A. Kourta, "Real-time feedback stall control of an airfoil at large Reynolds numbers using linear genetic programming," Phys. Fluids **34**, 045108 (2022).
[37]E. Amico, D. Di Bari, G. Cafiero, and G. Iuso, "Genetic algorithm-based control of the wake of a bluff body," J. Phys.: Conf. Ser. **2293**, 012016 (2022).
[38]A. Scala, G. Paolillo, C. S. Greco, T. Astarita, and G. Cardone, "Genetically-based active flow control of a circular cylinder wake via synthetic jets," Exp. Therm. Fluid Sci. **162**, 111362 (2025).
[39]J. Rabault, F. Ren, W. Zhang, H. Tang, and H. Xu, "Deep reinforcement learning in fluid mechanics: A promising method for both active flow control and shape optimization," J. Hydrodyn. **32**, 234–246 (2020).
[40]P. Garnier, J. Viquerat, J. Rabault, A. Larcher, A. Kuhnle, and E. Hachem, "A review on deep reinforcement learning for fluid mechanics," Comput. Fluids **225**, 104973 (2021).
[41]J. Viquerat, P. Meliga, A. Larcher, and E. Hachem, "A review on deep reinforcement learning for fluid mechanics: An update," Phys. Fluids **34**, 111301 (2022).
[42]J. Viquerat, R. Duvigneau, P. Meliga, A. Kuhnle, and E. Hachem, "Policy-based optimization: Single-step policy gradient method seen as an evolution strategy," Neural Comput. Appl. **35**, 449–467 (2023).
[43]H. Ghraieb, J. Viquerat, A. Larcher, P. Meliga, and E. Hachem, "Single-step deep reinforcement learning for two- and three-dimensional optimal shape design," AIP Adv. **12**, 085108 (2022).
[44]P. Scavella, G. Paolillo, and C. Greco, "Deep reinforcement learning-based airfoil design and optimization: An aerodynamic analysis," Aerosp. Sci. Technol. **167**, 110638 (2025).
[45]T. Sonoda, Z. Liu, T. Itoh, and Y. Hasegawa, "Reinforcement learning of control strategies for reducing skin friction drag in a fully developed turbulent channel flow," J. Fluid Mech. **960**, A30 (2023).
[46]L. Guastoni, J. Rabault, P. Schlatter, H. Azizpour, and R. Vinuesa, "Deep reinforcement learning for turbulent drag reduction in channel flows," Eur. Phys. J. E **46**, 27 (2023).
[47]J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, "Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control," J. Fluid Mech. **865**, 281–302 (2019).
[48]J. Rabault and A. Kuhnle, "Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach," Phys. Fluids **31**, 094105 (2019).
[49]H. Tang, J. Rabault, A. Kuhnle, Y. Wang, and T. Wang, "Robust active flow control over a range of Reynolds numbers using an artificial neural network trained through deep reinforcement learning," Phys. Fluids **32**, 053605 (2020).
[50]S. Shimomura, S. Sekimoto, A. Oyama, K. Fujii, and H. Nishida, "Closed-loop flow separation control using the deep Q network over airfoil," AIAA J. **58**, 4260–4270 (2020).
[51]S. S. Aote, M. Raghuwanshi, L. Malik *et al.*, "A brief review on particle swarm optimization: Limitations & future directions," Int. J. Comput. Sci. Eng. **14**, 196–200 (2013).
[52]J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of International Conference on Neural Networks (ICNN'95)* (IEEE, 1995), Vol. 4, pp. 1942–1948.
[53]R. Eberhart and J. Kennedy, "A new optimizer using particle swarm theory," in *Proceedings of the Sixth International Symposium on Micro Machine and Human Science (MHS'95)* (IEEE, 1995), pp. 39–43.
[54]Y. Zhang, S. Wang, and G. Ji, "A comprehensive survey on particle swarm optimization algorithm and its applications," Math. Probl. Eng. **2015**, 931256.
[55]B. Mohammadikalakoo, M. Kotsonis, and N. A. K. Doan, "Optimization of Tollmien-Schlichting waves control: Comparison between a deep reinforcement learning and particle swarm optimization approach," in Proceedings of TSFP13, Montréal, Canada, 2024.
[56]B. Mohammadikalakoo, M. Kotsonis, and N. Doan, "Control of Tollmien–Schlichting waves using particle swarm optimization," Phys. Fluids **36**, 124130 (2024).

[57]B. Mohammadikalakoo, M. Kotsonis, and N. Doan, "Real-time single-step deep reinforcement learning framework for control of Tollmien-Schlichting waves," Phys. Rev. Fluids **10**, 124902 (2025).
[58]B. Mohammadikalakoo, S. G. Villasol, G. Salomone, M. Kotsonis, and N. A. K. Doan, "Experimental demonstration of SDRL controller for TS wave suppression with DBD actuator," arXiv:2604.19326 (2026).
[59]L. M. Mack, "Boundary-layer linear stability theory," AGARD Repor No. 709, 1984.
[60]T. Bridges and P. J. Morris, "Differential eigenvalue problems in which the parameter appears nonlinearly," J. Comput. Phys. **55**, 437–460 (1984).
[61]M. Kotsonis and J. Van Ingen, "A two-parameter method for e^N transition prediction," AIAA Paper No. 2011-3928, 2011.
[62]R. Merino-Martínez, A. R. Carpio, L. T. L. Pereira, S. van Herk, F. Avallone, D. Ragni, and M. Kotsonis, "Aeroacoustic design and characterization of the 3D-printed, open-jet, anechoic wind tunnel of Delft University of Technology," Appl. Acoust. **170**, 107504 (2020).
[63]See https://www.acousticshop.co.uk/flamex-basic-absorption for "Flamex basic acoustic panels" (accessed June 01, 2024).
[64]M. Schultz and K. Flack, "The rough-wall turbulent boundary layer from the hydraulically smooth to the fully rough regime," J. Fluid Mech. **580**, 381–405 (2007).
[65]M. Forte, J. Jolibois, J. Pons, E. Moreau, G. Touchard, and M. Cazalens, "Optimization of a dielectric barrier discharge actuator by stationary and non-stationary measurements of the induced flow velocity: Application to airflow control," Exp. Fluids **43**, 917–928 (2007).
[66]Y. Wang and M. Gaster, "Effect of surface steps on boundary layer transition," Exp. Fluids **39**, 679–686 (2005).
[67]R. E. Hanson, N. M. Houser, and P. Lavoie, "Dielectric material degradation monitoring of dielectric barrier discharge plasma actuators," J. Appl. Phys. **115**, 043301 (2014).
[68]See https://www.speedgoat.com/products-services/real-time-target-machines/performance-real-time-target-machine for "Performance real-time target machine" (accessed June 01, 2024).
[69]B. Wieneke, "PIV uncertainty quantification from correlation statistics," Meas. Sci. Technol. **26**, 074002 (2015).
[70]A. J. Dietz, "Boundary-layer receptivity to transient convected disturbances," AIAA J. **36**, 1171–1177 (1998).
[71]W. Zhang, D. Ma, J.-J. Wei, and H.-F. Liang, "A parameter selection strategy for particle swarm optimization based on particle positions," Expert Syst. Appl. **41**, 3576–3584 (2014).
[72]M. Jiang, Y. P. Luo, and S. Y. Yang, "Stochastic convergence analysis and parameter selection of the standard particle swarm optimization algorithm," Inf. Process. Lett. **102**, 8–16 (2007).
[73]A. Widmann, A. Duchmann, A. Kurz, S. Grundmann, and C. Tropea, "Measuring Tollmien–Schlichting waves using phase-averaged particle image velocimetry," Exp. Fluids **53**, 707–715 (2012).