12th International Symposium on Turbulence and Shear Flow Phenomena (TSFP12)
Osaka, Japan (Online), July 19-22, 2022

GRADIENT-ENRICHED MACHINE LEARNING CONTROL FOR SHEAR FLOW STABILIZATION

Guy Y. Cornejo Maceda
School of Mechanical Engineering and Automation
Harbin Institute of Technology
Shenzhen 518055, People's Republic of China
guy.cornejo@limsi.fr

Yiqing Li
School of Mechanical Engineering and Automation
Harbin Institute of Technology
Shenzhen 518055, People's Republic of China
anne.yiqing.li@outlook.com

François Lusseyran
Laboratoire Interdisciplinaire des Sciences du Numérique
Université Paris-Saclay, CNRS
91400, Orsay, France
francois.lusseyran@limsi.fr

Marek Morzyński
Department of Virtual Engineering
Poznań University of Technology
PL 60-965 Poznań, Poland
morzynski@virtual.edu.pl

Bernd R. Noack
School of Mechanical Engineering and Automation
Harbin Institute of Technology
Shenzhen 518055, People's Republic of China
bernd.noack@hit.edu.cn

ABSTRACT

In this work, we propose the learning of feedback control laws with gradient-enriched machine learning control (Cornejo Maceda *et al.*, 2021, gMLC) algorithm. Gradient-enriched machine learning control builds upon machine learning control (Duriez *et al.*, 2016) and combines the explorative capability of genetic programming and the fast convergence of the downhill simplex gradient-descent method. Gradient-enriched MLC is demonstrated on the stabilization of two shear flows: a DNS of the flow past a cluster of three rotating cylinders—the fluidic pinball—and an open cavity flow experiment. For both cases, gMLC optimized feedback control laws outperforming previous open-loop and closed-loop controls with minimum actuation power. Moreover, the need of feedback has been demonstrated, revealing that feedback is a major feature for flow stabilization. Key enablers are automated machine learning algorithms augmented with gradient search: explorative gradient method (Li *et al.*, 2022) for the open-loop parameter optimization and a gradient-enriched machine learning control for the feedback optimization. Finally, gMLC learns the control law significantly faster than previously employed genetic programming control.

INTRODUCTION

Flow control is at the heart of many engineering applications. However, control design is challenged by the high-dimensionality of the flow, the nonlinearities inherent to the Navier-Stokes equations, and the time-delays between actuation and sensing. Hence, most closed-loop control studies of turbulence resort to a model-free approach. Genetic programming control (GPC) has been pioneered by Dracopoulos (1997) over 20 years ago and has been proven to be particularly successful for nonlinear feedback turbulence control in experiments. Machine learning control (Duriez *et al.*, 2016), based on GPC, has consistently outperformed existing optimized control approaches, often with unexpected frequency crosstalk mechanisms (Noack, 2019). GPC has a powerful capability to find new mechanisms (exploration) and populate the best minima (exploitation). Yet, the exploitation is inefficient leading to increasing redundant testing of similar control laws with poor convergence to the minimum. This challenge is well known and is addressed in this study with our new algorithm gradient-enriched machine learning control (Cornejo Maceda *et al.*, 2021, gMLC).

In this work, we report on the stabilization of two shear flows, a 2D numerical simulations of the fluidic pinball and the open cavity flow experiment.

Our first demonstrator is the fluidic pinball. It consists of three rotating cylinders in an incoming flow. The surprisingly rich dynamics of the fluidic pinball motivate our choice. Indeed, with increasing Reynolds number, the steady wake undergoes a Hopf, then a pitchfork bifurcation until reaching a chaotic regime (Deng *et al.*, 2020). Moreover, the many actuators allow to reproduce most common wake stabilization approaches , like Coanda forcing, base bleed, low- and high-frequency forcing, phasor control and circulation control (Ishar *et al.*, 2019). Finally, the rich unforced and controlled dynamics mimics nonlinear behavior of turbulence while the computation of the two-dimensional flow is manageable on workstations. To summarize, the fluidic pinball is an attractive all-weather plant for non-trivial multiple-input multiple-output control dynamics. In this work, the goal is to stabilize the flow to the symmetric steady solution by learning feedback controllers directly from the plant thanks to gradient-enriched machine learning control.

As for the open cavity, the control objective is to reduce the oscillations of the shear layer. Indeed, mitigation of the shear layer oscillations is of great engineering interest as the oscillations are related to aerodynamic drag and noise for aerospace and transport applications. Stabilization of these oscillations has been possible thanks to model-based methods for incompressible and compressible regimes. We refer to Rowley & Williams (2006) and Cattafesta III *et al.* (2008) for reviews of past successes of active flow control on the open cavity. A limit of linear control approaches for the open cavity, is the shift of oscillations of the cavity to other Rossiter modes (Cabell *et al.*, 2002; Williams *et al.*, 2000) resulting in multi-frequency or mode-switching regimes that compete for the available energy. Controlling such regimes has been possible in the past by Samimy *et al.* (2007), combining several models in linear quadratic optimal controllers. However, there is no general method for building control-oriented methods that includes the nonlinear frequency crosstalk between modes. Hence, in this work we employ a model-free approach based on machine learning to stabilize the shear layer oscillations and in particular a mode-switching regime.

PLANTS

The fluidic pinball—A MIMO control benchmark

As control benchmark problem, we chose the fluidic pinball, the flow around three cylinders located at the tips of an equilateral triangle pointing downstream. The fluidic pinball is studied here as a multiple-input multiple-output (MIMO) system. The actuation is performed by rotating each cylinder independently and 9 velocity probes located downstream feed back the flow velocity, see figure 1a. The velocity probes outside the centerline measure the streamwise velocity to follow the passing of the vortices. As for the probes on the centerline, they measure the spanwise velocity to characterize the symmetry of the flow. For this study, the Reynolds number based, on the diameter of one cylinder, is set to $\mathrm{Re} = 100$. In this regime, the flow is beyond the first two bifurcations: A first Hopf bifurcation at $\mathrm{Re} \approx 18$ enabling the vortex shedding and a pitchfork bifurcation at $\mathrm{Re} \approx 68$ deflecting the jet-like flow upwards or downwards (Deng *et al.*, 2020). Figure 1a shows a snapshot of the fluidic pinball in the post-transient regime. The mean flow (figure 1b), computed by averaging 1000 time units, highlights the strong asymmetry of the flow.

The control goal is then the stabilization of the unstable symmetric steady Navier-Stokes solution $\boldsymbol{u}_s$ (figure 1c). The cost function J_{FP} to minimize is the time-averaged L_2-norm of the difference between the instantaneous flow field and the objective flow field:

$$\begin{aligned} J_{\mathrm{FP}} &= \frac{1}{T_{\mathrm{FP}}} \int_{t_0}^{t_0+T_{\mathrm{FP}}} j_{fp}(t)\,\mathrm{d}t \\ j_{fp}(t) &= \|\boldsymbol{u_b}(t) - \boldsymbol{u}_s\|_\Omega^2 \end{aligned} \tag{1}$$

where t_0 is the time when the actuation is turned on, $T_{\mathrm{FP}} = 1000$ time units, $\mathrm{u_b}$ is the instantaneous flow field and $||\cdot||_\Omega^2$ denotes the spatial integration. The control transient is included in the evaluation window to reach the steady solution faster. The direct numerical simulations are carried out with an in-house solver using a finite-element method third order accurate in time and space. We refer to (Deng *et al.*, 2020) for more details on the solver.

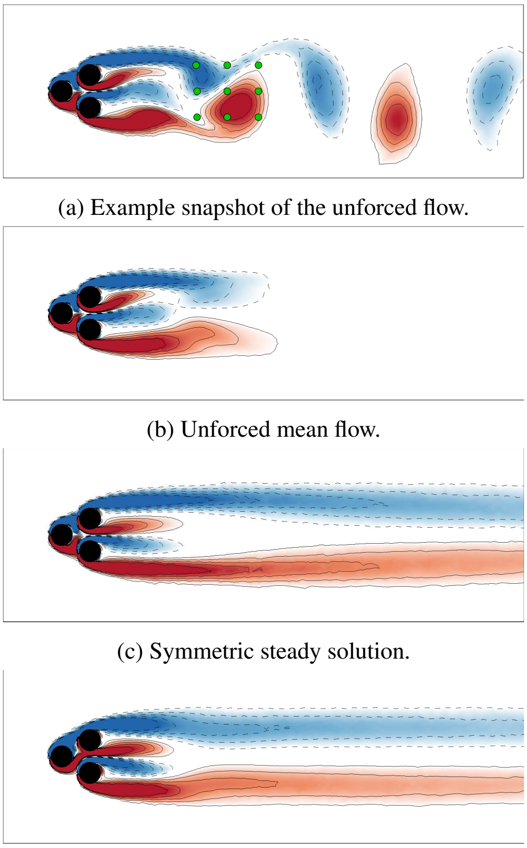

(a) Example snapshot of the unforced flow.

(b) Unforced mean flow.

(c) Symmetric steady solution.

(d) Mean field of the flow controlled with the gMLC law.

Figure 1: Vorticity fields of the unforced and controlled flow for the fluidic pinball. The position of the velocity sensors are denoted by green dots in (a).

The open cavity experiment

For the experimental demonstration, we have chosen a flow configuration gathering most of the mechanisms responsible for nonlinear interactions but keeping the self-organization of the spatial structures still highly coherent, i.e. the flow above an open cavity. We are interested in the stabilization of the oscillating shear layer resulting of the interaction between an incoming boundary layer and the cavity. For compressible flows, the well-known flow-acoustic resonance or Rossiter mechanism describes the self-sustained oscillations and frequencies of the flow (Rowley & Williams, 2006). However, for incompressible flows, there is no consensus on the mechanism leading to self-sustained oscillations. Tuerke *et al.* (2020) shows that Stuart-Landau type amplitude equation including two distinct delay lines, associated with a backward wave time and an intra-cavity overturning time, is able to predict the two frequencies of a multi-frequency regime of the open cavity. In this study, the control goal is to stabilize the flow by mitigating the self-oscillations of the shear layer with a single actuator and a single sensor (figure 2). Our open cavity experiment is then a single-input single-output (SISO) system. Our cavity is $D = 0.05\,\mathrm{m}$ deep, $W = 0.30\,\mathrm{m}$ wide and its length and incoming velocity are chosen to enter two distinct regimes. The first regime, referred in the following by *narrow-bandwidth regime* displays one dominant mode ($f_a \approx 28.81\,\mathrm{Hz}$). This

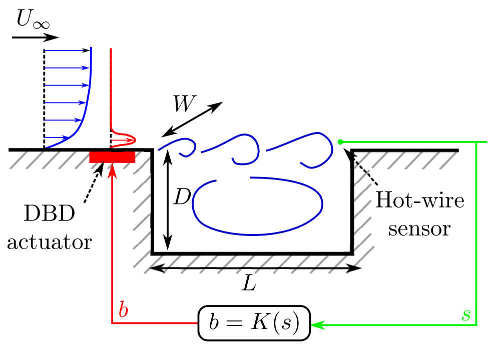

Figure 2: Schematic of the open cavity experiment and the control loop. The incoming velocity profile is denoted in blue and the profile of the body force generated by the DBD actuator in red. L, D and W denote, respectively, the length, depth and width of the cavity. The width is not represented. b, s and K are, respectively, the actuation command, the sensor signal and the control law.

regime is achieved for $L = 0.075\,\mathrm{m}$, $U_\infty = 2.13\,\mathrm{m\,s^{-1}}$). The second mode, referred in the following by *mode-switching regime*, includes two competing modes ($f_a = 27.31\,\mathrm{Hz}$ and $f^+ = 36.91\,\mathrm{Hz}$). A spectral analysis over time shows that the existence time for each mode is between 15 and 20 s. This regime is achieved for $L = 0.0875\,\mathrm{m}$, $U_\infty = 2.23\,\mathrm{m\,s^{-1}}$). We refer to Basley *et al.* (2013) for more details on the narrow-bandwidth and mode-switching regimes.

The forcing is achieved by a dielectric barrier discharge (DBD) actuator located upstream at the receptivity point of the cavity. The actuation command b is then the amplitude level of a high-frequency carrier signal (3 kV). The minimum value of the actuation command is set to the ionization threshold of the DBD actuator and the maximum value to the maximum amplitude level that keeps a trace of dynamics in the spectrum. The effect of the forcing is similar to a body force extended in the spanwise direction. A hot-wire sensor downstream feeds back the streamwise velocity enabling closed-loop control (see figure 2). The spectra of the velocity downstream for the two unforced regimes are depicted in black in figure 3a and 3b.

The control objective is then to reduce the energy associated to the main frequencies f_a and f^+ of the flow. A constant forcing study at different actuation levels shows that a high amplitude steady forcing is enough to reduce the amplitude of the frequency peaks, hence we decide the penalize the actuation power. The cost function J_{OC} is then the sum of two terms: Ja defined as the ratio between the maximum values in the power spectral density (PSD) of the unforced and controlled cases, and J_b the actuation penalization term:

$$\begin{aligned} J_{\mathrm{OC}} &= J_a + \gamma J_b \\ J_a &= \frac{\max \mathrm{PSD}(u_{\mathrm{HW}})}{\max \mathrm{PSD}(u_{\mathrm{HW},0})} \\ J_b &= \frac{\langle (b+1)^2 \rangle}{4} \end{aligned} \tag{2}$$

where γ, the penalization parameter, is set to 1 as J_a and J_b are both normalized by the unforced case and the maximum actuation level respectively. $\langle \cdot \rangle$ denotes a time average quantity. The maximum of the PSD is detected in a limited frequency window including only the two main frequencies of the flow f_a and f^+. The evaluation of each individual is set to 40 s; This

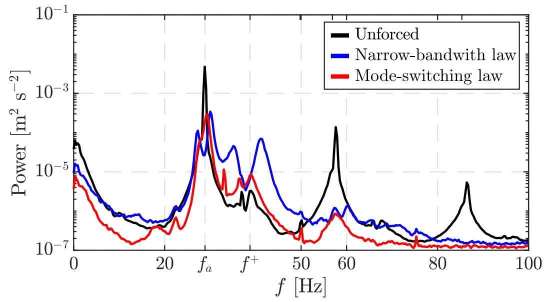

(a) Narrow-bandwidth regime.

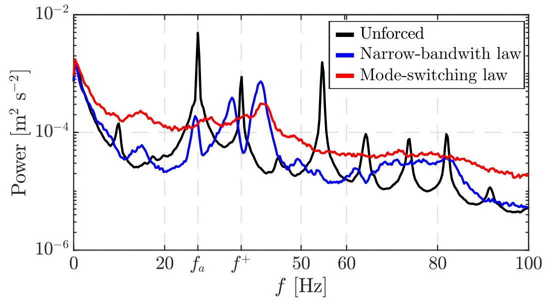

(b) Mode-switching regime.

Figure 3: Power Spectra of the open cavity experiment for the narrow-bandwith and the mode-switching regimes. Spectra associated to the unforced flow are depicted in black. The spectra resulting of the control by the control law learned in the narrow-bandwidth (mode-switching) are depicted in blue (red).

value balances good convergence of the statistics and practicality as most experiments have a limited testing budget. Moreover with 40 s, we assure that the trace of both frequencies (f_a and f^+) are present in the spectrum of the mode-switching regime.

CONTROL PROBLEM AND METHODOLOGY

The flow control problem is to derive a mapping $\boldsymbol{K}$ between the outputs (sensors signals $\boldsymbol{s}$) and the inputs of the system (actuation commands $\boldsymbol{b}$) to achieve a control goal: $\boldsymbol{b} = \boldsymbol{K}(\boldsymbol{s})$. $\boldsymbol{K}$, also referred as *control law* is a scalar function for SISO control (case of the open cavity experiment) and a vectorial function for MIMO control (case of the fluidic pinball). Leveraging machine learning techniques for flow control relies on the reformulation of the flow control problem as a functional regression problem where the goal is to derive the optimal control $\boldsymbol{K}^*$ that minimize the cost function J.

$$\boldsymbol{K}^* = \underset{\boldsymbol{K} \in \boldsymbol{K}}{\arg\min}\, J(\boldsymbol{K}) \tag{3}$$

with Λ being the control law space. In general, equation (3) is a challenging non-convex optimization problem including several minima.

The employed algorithm for control law optimization is gradient-enriched machine learning control (Cornejo Maceda *et al.*, 2021, gMLC). The algorithm is inspired by the explorative gradient method (Li *et al.*, 2022, EGM) that combines exploration to discover new minima and exploitation with gradient-based approach for a fast convergence. Starting point is machine learning control (Duriez *et al.*, 2016, MLC) based genetic programming control (Dracopoulos, 1997). The principle of MLC mimics the Darwinian natural selection to optimize control laws, by generating a set, or a population of random control laws that evolve through generations. Following the evolutionary terminology, the control laws are also referred as *individuals*. The process of evolution relies on the recombination of most performing individuals from one generation to build the next generations of individuals. The recombination is carried out by genetic operators: crossover and mutation (Li *et al.*, 2019). MLC has been successfully employed in dozens of experiment outperforming previous controllers often by deriving nonlinear mechanisms (Noack, 2019). Like EGM, MLC can be accelerated by including intermediate gradient steps to exploit the local gradients in the search space. The resulting algorithm is the gradient-enriched machine learning control. Gradient-enriched MLC (figure 4) departs in two aspects from MLC. First, the concept of evolution from generation to generation is not adopted. The genetic operations include all tested individuals. Second, the exploitation is accelerated by downhill simplex iterations. The algorithm begins with a broad exploration of the search space thanks to a Monte Carlo sampling like MLC, generating N_{MC} random individuals. Then N_{G} new individuals are generated thanks to crossover and mutation. This step is referred as an exploration step as it is mainly expected to discover new minima. The algorithm continues with an exploitation step carried out by the downhill subplex algorithm (Rowan, 1990). The downhill subplex algorithm is a variant of the downhill simplex algorithm (Nelder & Mead, 1965) for infinite dimensions spaces. In this step, the N_{sub} best individuals evaluated so far are selected to form a subspace of finite dimension. New individuals are then generated thanks to downhill simplex steps performed in this subspace. The new individuals are linear combinations of the N_{sub} bests. The downhill simplex steps are repeated until N_{G} individuals are generated to balance the exploration and exploitation phase. The algorithm then continues with new exploration-exploitation phases until a stopping criterion is reached. In the exploration phases, the individuals to be recombined are selected among all the individuals evaluated thanks to a tournament method. The subspace basis is updated with new individuals when those are better than ones in the current basis. Figure 4 summarizes the gradient-enriched algorithm. We refer to (Cornejo Maceda *et al.*, 2021) for more details on the gMLC algorithm.

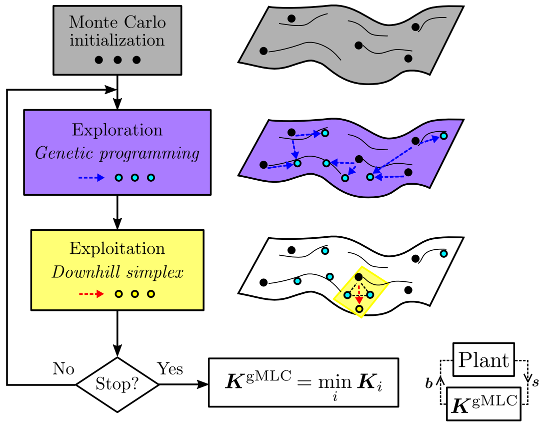

Figure 4: Gradient-enriched machine learning algorithm and principle sketch of the learning process in the search space. The blue and red arrows denote how the new individuals are generated. Blue is for genetic operators and red for linear combination. The fully colored search space for the Monte Carlo initialization and the exploration phase indicate that these phases allow to discover control laws in the whole space. The yellow rectangle symbolizes the subspace formed by the best individuals where the exploitation is performed. The plant stands for the flow system to control, i.e., in this study, the fluidic pinball or the open cavity experiment.

RESULTS

Stabilization of the fluidic pinball

For the stabilization of the fluidic pinball, first, steady symmetric forcing is optimized. The front cylinder does not rotate and the two back cylinders rotate at constant speed in opposite directions. A parametric study to derive the remaining parameter reveals that a base bleeding solution leads to a flow which is 49% closer to the symmetric solution than the unforced attractor. Second, a general steady actuation is optimized by allowing the independent rotation of all the cylinders at constant speed. The three velocity parameters are optimized with the explorative gradient method. Surprisingly, an asymmetric actuation reduces the average distance between the flow and the symmetric steady target solution further to 28% of the unforced value. Third, gMLC is employ to optimize a feedback control law. The gMLC parameters chosen for the optimization are $N_{\mathrm{MC}} = 100$, $N_{\mathrm{G}} = 50$ and $N_{\mathrm{sub}} = 10$. The exploration and exploitation phases are repeated until 1000 individuals are evaluated. These parameters are discussed in Cornejo Maceda *et al.* (2021). The feedback control optimized with gMLC brings the flow even closer to the steady target solution reducing the cost function to 20% of the unforced value with small actuation power. The actuation is a combination of asymmetric steady forcing and phasor control. The resulting mean flow (figure 1d) is similar to the steady target solution, as the two vorticity branches reach the end of the computation domain and the jet between the two back cylinders is almost statistically symmetric. The slight deflection of the near jet may be related to the asymmetry of the control law. The vorticity branches are also vectored towards the centerline.

Moreover, feedback plays a decisive role in the stabilization of the fluidic pinball with the gMLC control law. Indeed, the actuation commands associated to the gMLC control include low-amplitude unsteady components. However, despite being small, the unsteady component is a key feature for stabilization as an averaged actuation failed to bring the flow close to the symmetric steady solution, achieving a relative cost of only 59%.

Finally, a comparison with MLC reveals that gMLC learns more performing individuals and with lesser evaluated individuals.

Feedback control of the open cavity experiment

In this section, the gMLC algorithm is employed to stabilize the narrow-bandwidth and the mode-switching regime. The same gMLC parameters as for the fluidic pinball are employed. For the narrow-bandwidth regime, gMLC optimized a feedback control reducing the cost J_{OC} by 98%, which corresponds to a 99% reduction of the fluctuation energy with less than 1% of the maximum actuation level. We note that the reduction of the main peak of the spectrum f_a comes with the increase of the frequency f^+ associated with the second mode of the flow, see the spectrum blue in figure 3a. For the mode-switching regime, gMLC derived a feedback control law that reduce the cost by 94%, which corresponds to 97% reduction of the fluctuation energy with 2% of the maximum actuation level. The latter control law has been able to reduce the power associated to both frequencies f_a and f^+, see the spectrum red in figure 3b.

The effectiveness of the control is also tested by evaluating each control in the other regime. Expectedly, the law learned in the narrow-bandwidth regime is only able to partially control the mode-switching regime; the amplitude of the main peak at f_a is reduced but the peak at f^+ remains, see the spectrum blue in figure 3b. Surprisingly, the law learned in the mode-switching regime was able to reduce the main peak at f_a while preventing the growth of the peak at f^+, see the spectrum red in figure 3a. Therefore, this study shows the benefit of learning feedback control laws in complex regimes as the richness of the regime is reflected in the control laws efficiency.

Moreover, like for the fluidic pinball, the need of feedback has been demonstrated to be an essential feature to mitigate the oscillations of the cavity. The actuation commands employed during the control of the narrow-bandwidth regime haven been recorded and employed as open-loop control commands. The resulting spectra show that such control is enable to control the main mode of the flow at frequency f_a. This tests show that both control laws effectively operate in a closed-loop manner.

Finally, like for the fluidic pinball, gMLC outperforms MLC both in terms of performance of the final solution and convergence speed.

CONCLUSIONS

In this study, we employ the gradient-enriched machine learning control for an automated learning of feedback control laws stabilizing the flow past a cluster of three cylinders—the fluidic pinball—and the open cavity flow in two regimes: one dominated by one single mode and one where two modes compete. The optimized feedback control laws achieve not only the best performance so far compared to previously employed steady actuation but also with small actuation power. Moreover, the necessity of feedback for the control has been demonstrated. The presented stabilizations are expected to be independent of the employed optimizer as different approaches lead to very similar results. The chosen optimizer balances exploration (search for better minima) and exploitation (downhill descend of found minima). The automated learning of feedback control laws has been significantly accelerated by intermittently adding gradient-based descends, outperforming MLC. Building on this success, we believe that gradient-enriched MLC will greatly accelerate the optimization of control laws for MIMO control experiments including a large number of actuators and sensors. Recent experimental applications of gMLC include successful drag reduction of a generic truck model under yaw and lift increase of a high-Reynolds number airfoil.

REFERENCES

Basley, J., Pastur, L. R., Delprat, N. & Lusseyran, F. 2013 Space-time aspects of a three-dimensional multi-modulated open cavity flow. *Phys. Fluids* **25** (6), 064105.

Cabell, R., Kegerise, M., Cox, D. & Gibbs, G. 2002 Experimental feedback control of flow induced cavity tones. In *8th AIAA/CEAS Aeroacoustics Conference & Exhibit*.

Cattafesta III, L. N., Song, Q., Williams, D. R., Rowley, C. W. & Alvi, F. S. 2008 Active control of flow-induced cavity oscillations. *Progress in Aerospace Sciences* **44** (7-8), 479–502.

Cornejo Maceda, G. Y., Li, Y., Lusseyran, F.and Morzyski, M. & Noack, B. R. 2021 Stabilization of the fluidic pinball with gradient-enriched machine learning control. *J. Fluid Mech.* **917**, A42.

Deng, N., Noack, B. R., Morzyski, M. & Pastur, L. R. 2020 Low-order model for successive bifurcations of the fluidic pinball. *J. Fluid Mech.* **884**, A37.

Dracopoulos, D. C. 1997 *Evolutionary Learning Algorithms for Neural Adaptive Control*. London, etc.: Springer-Verlag.

Duriez, T., Brunton, S. L. & Noack, B. R. 2016 *Machine Learning Control — Taming Nonlinear Dynamics and Turbulence*, *Fluid Mechanics and Its Applications*, vol. 116. Springer-Verlag.

Ishar, R., Kaiser, E., Morzynski, M., Albers, M., Meysonnat, P., Schröder, W. & Noack, B. R. 2019 Metric for attractor overlap. *J. Fluid Mech.* **874**, 720–752.

Li, R., Borée, J., Noack, B. R., Cordier, L. & Harambat, F. 2019 Drag reduction mechanisms of a car model at moderate yaw by bi-frequency forcing. *Phys. Rev. Fluids* **4** (3), 034604.

Li, Y., Cui, W., Jia, Q., Li, Q., Yang, Z., Morzyński, M. & Noack, B. R. 2022 Explorative gradient method for active drag reduction of the fluidic pinball and slanted ahmed body. *J. Fluid Mech.* **932**, A7.

Nelder, J. A. & Mead, R 1965 A simplex method for function minimization. *Computer Journal* **7** (4), 308–313.

Noack, B. R. 2019 Closed-loop turbulence control—From human to machine learning (and retour). In *Fluid-Structure-Sound Interactions and Control. Proceedings of the 4th Symposium on Fluid-Structure-Sound Interactions and Control* (ed. Y. Zhou, M. Kimura, G. Peng, A. D. Lucey & L. Hung), pp. 23–32. Springer.

Rowan, T. 1990 The subplex method for unconstrained optimization. PhD thesis, PhD thesis, Department of Computer Sciences, University of Texas.

Rowley, C. W. & Williams, D. R. 2006 Dynamics and control of high-Reynolds number flow over open cavities. *Annu. Rev. Fluid Mech.* **38**, 251–276.

Samimy, M., Kim, J.-H., Kastner, J., Adamovic, I. & Utkin, Y. 2007 Active control of high-speed and high-reynolds-number jets using plasma actuators. *J. Fluid Mech.* **578**, 305–330.

Tuerke, F., Lusseyran, F., Sciamarella, D., Pastur, L. R. & Artana, G. 2020 Nonlinear delayed feedback model for incompressible open cavity flow. *Phys. Rev. Fluids* **5**, 024401.

Williams, D., Drazen, F., Iwanski, K. & Morrow, J. 2000 Closed-loop control in cavities with unsteady bleed forcing. In *38th Aerospace Sciences Meeting and Exhibit*.