

# Active morphing façades for wind mitigation: Multi-agent reinforcement learning for aerodynamic control

Foad Mohajeri Nav <sup>a</sup>, Reda Snaiki <sup>b</sup>, Pedro L. Fernández-Cabán <sup>c</sup>

<sup>a</sup> *Department of Construction Engineering, École de Technologie Supérieure, Université Du Québec, Montréal, QC, H3C 1K3, Canada, [foad.mohajeri-nav.1@ens.ets.ca](mailto:foad.mohajeri-nav.1@ens.ets.ca)*

<sup>b</sup> *Department of Construction Engineering, École de Technologie Supérieure, Université Du Québec, Montréal, QC, H3C 1K3, Canada, [reda.snaiki@etsmtl.ca](mailto:reda.snaiki@etsmtl.ca)*

<sup>c</sup> *Florida A&M University–Florida State University (FAMU-FSU) College of Engineering, Tallahassee, FL, USA, [plfernandez@eng.famu.fsu.edu](mailto:plfernandez@eng.famu.fsu.edu)*

## SUMMARY

This work presents a data-driven framework for mitigating wind effects on tall buildings. This is based on using a morphing façade system to minimize wind-induced structural response. A surrogate model is first trained on wind-tunnel measurements from an active fin system, mapping various fin angles to the top-level acceleration time history in two directions. The structural response is then reconstructed and calibrated in the time domain to reproduce measured accelerations. Then, a multi-agent reinforcement learning (MARL) controller is used to minimize wind-induced effects by optimizing the fin configuration. Each fin is treated as an agent with a Gaussian policy over its angle, and a global reward equal to the negative RMS of the top-level acceleration is maximized using an update with a running baseline. In addition, an actor–critic formulation based on Proximal Policy Optimization (PPO) is introduced with a multi-objective reward that accounts for resultant and cross-wind responses. For the considered case study, the learned fin configuration achieves approximately 31% reduction in along-wind RMS acceleration using the MARL formulation, while the PPO-based approach further improves performance, achieving reductions of over 60% in resultant acceleration compared to the reference configuration.

**Keywords:** *Wind hazard mitigation, Active flow control, Morphing façade, Reinforcement learning*

## 1. INTRODUCTION

Wind hazards on high-rises are evolving due to more extreme wind events and a changing urban environment. Therefore, traditional wind mitigation strategies cannot provide sufficient solutions for such buildings. Aerodynamic mitigation strategies such as corner modifications, fins, parapets, and openings have been widely investigated through wind-tunnel testing and computational fluid dynamics (CFD). Recent cyber-physical experiments have demonstrated that actively controlled corner fins can significantly reduce along-wind accelerations by coupling an aeroelastic model in a boundary-layer wind tunnel with real-time optimization. In particular, Whiteman et al. (2022) implemented an active fin system driven by a particle swarm optimization (PSO) loop and a controller to search for optimal fin configurations directly in the wind tunnel. However, implementing such techniques in real buildings faces serious cost and technical issues and might not be efficient for real-time structural health monitoring (SHM).

As a solution, surrogate-based aerodynamic shape optimization has emerged as an efficient strategy to approximate the structural response over a high-dimensional design space while limiting the number of physical tests. For example, Lu et al. (2023) proposed a surrogate-based cyber-physical framework for high-rise buildings, in which wind-tunnel measurements are used to train machine-learning models for subsequent optimization. Earlier efforts on aerodynamic tailoring using computational fluid dynamics (Ding et al., 2019), have also highlighted the

effectiveness of geometry modification in improving wind-induced performance of structures. Moreover, the concept of dynamic façades for wind-responsive tall buildings (Ding and Kareem, 2020), further emphasizes the potential of adaptive systems in mitigating wind loads. In addition, reinforcement learning–based aerodynamic optimization frameworks (Li et al., 2021a), have demonstrated the capability of integrating domain knowledge into data-driven shape optimization. However, this field is still challenging from a practical perspective.

In parallel, reinforcement learning (RL) has shown strong potential for active flow control (AFC) problems, including drag reduction for bluff bodies using jets or rotating cylinders, in both simulations and experiments (Wang et al., 2026; Montalà et al., 2025; Font et al., 2025; Yan et al., 2025; Suárez et al., 2025; Jia and Xu, 2024; Ding et al., 2019; Ding and Kareem, 2020). Furthermore, deep reinforcement learning has also been applied to actively simulate and control transient wind fields in experimental facilities (Li et al., 2021b), highlighting its applicability to complex, time-dependent aerodynamic environments. In particular, multi-agent reinforcement learning (MARL) strategies have been explored recently for coordinated flow control in complex fluid systems, demonstrating improved scalability and control performance in high-dimensional settings (Suárez et al., 2025). However, sufficient studies addressing the challenges related to wind hazard mitigation on tall buildings are still lacking.

This study proposes a surrogate-driven, multi-agent reinforcement learning (MARL) framework tailored to the active façade system used by Whiteman et al. (2022). The goal is to optimize fin angles that minimize the RMS of top-floor acceleration using only a pre-trained surrogate, without repeatedly querying the wind tunnel. This can provide strong practical value for real-time SHM.

## 2. METHODOLOGY

This study proposes a surrogate-driven, multi-agent reinforcement learning (MARL) framework tailored to the active façade system investigated in Whiteman et al. (2022), with the objective of optimizing fin angles to minimize the root-mean-square (RMS) of top-floor acceleration. A key innovation of the framework is that the control policy is trained exclusively using a pre-trained surrogate model, thereby eliminating the need for repeated and costly wind tunnel evaluations during the optimization phase. The adopted control paradigm is consistent with recent advances in distributed flow control using MARL, as demonstrated by Suárez et al. (2025).

The case study is based on a high-fidelity aeroelastic experimental campaign conducted at the Natural Hazard Engineering Research Infrastructure Experimental Facility, where a 1:200 scale multi-degree-of-freedom (MDOF) model of a 76-story benchmark tall building, is tested under realistic boundary layer wind conditions. The model has a total height of 1.53 m and is designed to accurately reproduce the distributed mass and stiffness characteristics of the prototype structure through a central steel spine connected to seven aluminum diaphragms. The structural system exhibits a first-mode natural frequency of approximately 2.30 Hz at model scale with an estimated damping ratio of 2.5%, ensuring dynamic similitude with real tall buildings.

A defining feature of this experimental setup is the integration of an active aerodynamic façade system consisting of twelve independently actuated slotted fins, mounted at three vertical levels along the four building corners. These fins introduce localized perturbations to the flow field, enabling manipulation of vortex shedding, separation, and reattachment mechanisms that govern

aerodynamic loading. Each fin is actuated via a NEMA11 stepper motor, allowing continuous variation of the fin angle within a bounded range of  $[0^\circ, 270^\circ]$ . The fins are geometrically non-uniform along the height, with larger elements near the top region to enhance sensitivity to higher-mode aerodynamic effects. Moreover, a symmetric configuration is assumed, where the angles on one side the building are just mirroring the angles on the other side (Whiteman et al., 2022). This configuration results in a 6-dimensional continuous control space, posing a non-trivial optimization problem well-suited for reinforcement learning.

The wind tunnel experiments provide paired datasets consisting of static fin-angle configurations and the resulting top-floor along-wind acceleration time histories. These data define the input–output mapping that the surrogate model seeks to approximate. The surrogate operates in the wavelet domain, where each acceleration signal is decomposed into multi-resolution coefficients. A feed-forward multilayer perceptron (MLP) is trained to map the scaled 6-dimensional fin-angle vector to the corresponding wavelet coefficients. The predicted coefficients are then reconstructed into the time domain and further refined through a gain–offset calibration step to ensure consistency with the experimental signal statistics. This hybrid representation significantly reduces the dimensionality of the learning problem while preserving temporal fidelity, enabling rapid and deterministic evaluations required for reinforcement learning. Within this framework, two reinforcement learning formulations are investigated: A MARL formulation with REINFORCE, and an actor–critic Proximal Policy Optimization (PPO) formulation with multi-objective reward.

## 2.1. A MARL formulation

The control problem is first formulated as a decentralized stochastic optimization problem. The fin-angle vector is defined as:

$$\theta = [\theta_{1L}, \dots, \theta_{6L}], \quad \theta_j \in [0^\circ, 270^\circ] \quad (1)$$

It should be mentioned that the fin angles on the right side are obtained based on assuming a symmetric configuration. Each fin  $j$  is modeled as an independent agent governed by a Gaussian policy:

$$\theta_j \sim \mathcal{N}(\mu_j, \sigma_j) \quad (2)$$

At each episode, a candidate configuration is sampled and evaluated through the surrogate, which returns the predicted acceleration time history  $a(t; \theta)$ . The reward is defined as:  $\alpha$

$$R(\theta) = -RMS[a(t; \theta)] \quad (3)$$

so that maximizing the reward corresponds to minimizing the along-wind structural response. Policy parameters are updated using the REINFORCE algorithm, with a running baseline introduced to reduce variance. Analytical gradients of the Gaussian policy are computed with respect to both mean and variance parameters, and gradient ascent updates are applied. This formulation provides a simple yet effective approach for exploring the high-dimensional control space.

## 2.2. An actor–critic Proximal Policy Optimization (PPO) formulation

To enhance stability and incorporate additional performance objectives, a second formulation based on an actor–critic architecture is introduced. In this approach, a neural network policy (actor) outputs fin-angle actions, while a value function (critic) estimates the expected return. The policy is trained using PPO, which improves training robustness through clipped policy updates.

The reward function is extended to account for both along-wind and cross-wind responses:

$$r_t = A_0(\omega_t) - A(\omega_t) - \lambda R_y(\omega_t) \quad (4)$$

where:

- $A(\omega_t)$  is the RMS of the resultant acceleration at time window  $\omega_t$ ,
- $A_0(\omega_t)$  is the baseline RMS corresponding to a reference configuration (e.g., all fins at  $270^\circ$ ),
- $R_y(\omega_t)$  is the cross-wind RMS component,
- $\lambda$  is a weighting parameter controlling the penalty on the cross-wind response.

This reward formulation promotes configurations that not only reduce the overall structural response relative to the baseline but also explicitly penalize undesirable cross-wind vibrations, which are often critical for serviceability and occupant comfort. Figure 1 shows the overall process for the PPO-based actor-critic framework.

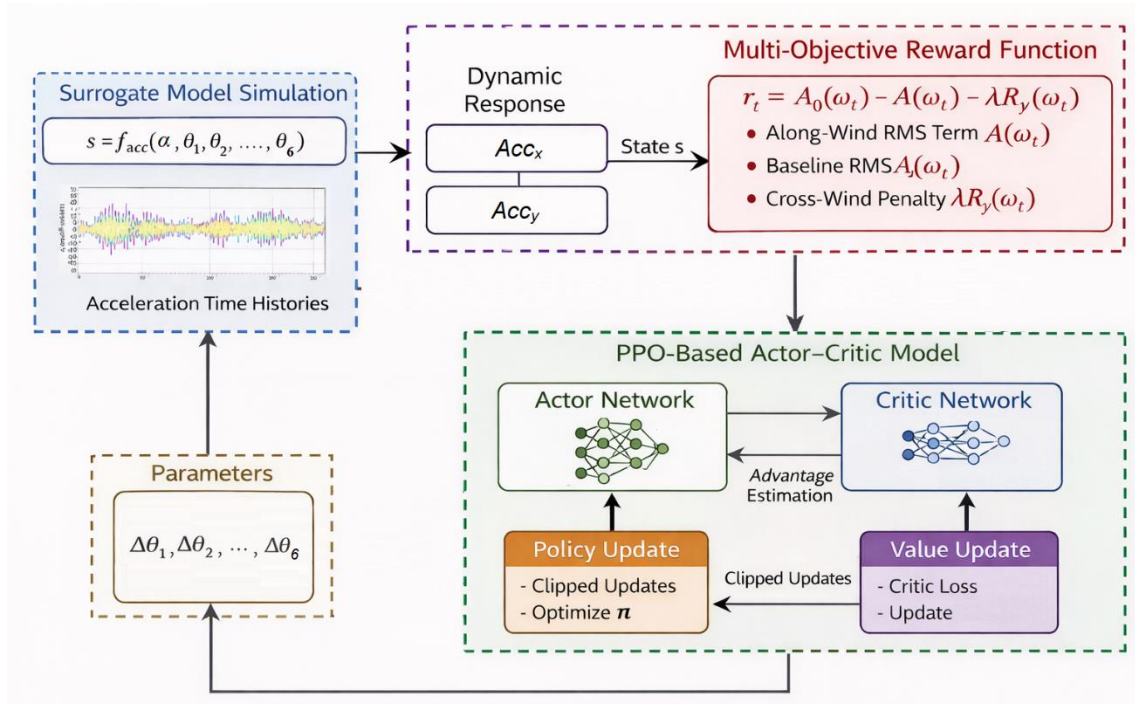


Figure 1: Reduced acceleration at the top-level vs the original response (MARL controller).

Training is performed over multiple episodes using rollout buffers, where trajectories of states, actions, and rewards are collected. The PPO objective then updates the actor network while constraining policy deviations, ensuring smooth and reliable convergence. After training, the learned policy is evaluated deterministically by selecting the mean action outputs, and the resulting fin configuration is assessed using the surrogate model. Performance is quantified through RMS-based metrics and directly compared against the baseline configuration.

### 3. RESULTS AND DISCUSSION

The performance of the proposed surrogate-driven reinforcement learning framework is evaluated by comparing optimized fin configurations against a baseline case where all fins are set to  $270^\circ$ . Two control formulations are considered: (i) a decentralized MARL approach trained using the REINFORCE algorithm with a scalar reward based on along-wind RMS, and (ii) a PPO-based actor-critic formulation with a multi-objective reward incorporating resultant and cross-wind responses. The inclusion of two reward formulations enables a comprehensive evaluation of control strategies. The REINFORCE-based approach provides a lightweight and interpretable baseline for decentralized control, while the PPO-based actor-critic formulation introduces improved stability, sample efficiency, and the ability to incorporate multi-objective performance criteria.

#### 3.1. Performance of REINFORCE-Based MARL Controller

The REINFORCE-based MARL controller achieves an approximately 31% reduction in along-wind RMS acceleration at the top floor relative to the baseline configuration. The corresponding time histories (Figure 2) show a consistent decrease in peak amplitudes and reduced fluctuation intensity throughout the record. This indicates that the optimized fin configuration effectively alters the flow-structure interaction, reducing both resonant amplification and broadband aerodynamic excitation.




Figure 2: Reduced acceleration at the top-level vs the original response (MARL controller).

From a control perspective, this result demonstrates that even a relatively simple policy-gradient formulation, operating on a static action space and a scalar objective, is capable of identifying non-

trivial aerodynamic configurations in a high-dimensional design space. The use of a surrogate model enables rapid policy evaluation, making the training process computationally efficient and independent of the wind tunnel.

However, because the reward function is defined solely in terms of along-wind response, the optimization is inherently biased toward reducing drag-related excitation. As a result, potential improvements in cross-wind response (often governed by vortex shedding) are not explicitly targeted in this formulation.

### **3.2. Performance of PPO-Based Actor–Critic Controller**

Figure 3 presents the results obtained using the PPO-based actor–critic formulation with the multi-objective reward. In contrast to the REINFORCE case, this approach yields substantially larger reductions across all response components: Along-wind RMS, reduced by approximately 45.5%; cross-wind RMS, reduced by approximately 64.7%; and resultant RMS, reduced by approximately 61.2%. The time-history comparisons show a pronounced attenuation of both peak responses and high-frequency fluctuations, with particularly strong improvements in the cross-wind direction. This is a critical observation, as cross-wind vibrations are typically associated with vortex shedding and are often the dominant contributor to occupant discomfort in tall buildings.

The improved performance of the PPO-based controller can be attributed to two key factors. First, the multi-objective reward function explicitly penalizes cross-wind response while also maximizing reduction relative to the baseline, leading to a more balanced control strategy. Second, the actor–critic architecture enables coordinated optimization across all fins through a shared policy network, capturing coupling effects that are not accessible in a fully decentralized formulation.

### **3.3. Discussions**

A direct comparison between the two approaches highlights several important insights: first, while the REINFORCE approach achieves a meaningful 31% reduction in along-wind response, the PPO formulation delivers significantly larger reductions across all metrics, exceeding 60% in resultant response. This demonstrates the advantage of incorporating additional physical objectives into the reward function. Second, the decentralized nature of the REINFORCE-based MARL limits interaction between individual fins, potentially leading to suboptimal global configurations. In contrast, the PPO-based actor–critic framework implicitly captures interdependencies among actuators, enabling more effective manipulation of the global flow field. Finally, the PPO-controlled response exhibits smoother time histories with reduced intermittency and lower peak amplitudes. This suggests that the controller not only reduces overall energy but also modifies the spectral content of the response, likely through improved control of vortex shedding mechanisms.

From an engineering standpoint, the achieved reductions in acceleration are highly significant. A 30–60% decrease in RMS acceleration can substantially improve occupant comfort, as serviceability criteria in tall buildings are often governed by acceleration thresholds rather than strength limits. Moreover, reduced dynamic response leads to lower cyclic demand on both structural and non-structural components, which can decrease fatigue damage accumulation,

extend the service life of façade systems and connections, and reduce maintenance requirements over the building lifecycle. The strong reduction in cross-wind response observed in the PPO-based formulation is particularly important, as vortex-induced vibrations are often the governing design consideration for slender high-rise structures.




Figure 3: Original (initial) vs. reduced acceleration time histories for along wind ( $a_x$ ) and across wind ( $a_y$ ) directions (PPO-based actor-critic controller)

The present evaluation is conducted for a single wind direction ( $0^\circ$ ) under a specific turbulence condition derived from the wind tunnel dataset. While the results demonstrate the effectiveness of the proposed framework, further validation is required to assess its robustness and generalization capability.

#### 4. CONCLUSIONS

The present research introduces a surrogate-driven reinforcement learning framework for optimizing a morphing façade system on a tall building subjected to wind loading. By integrating a physics-informed surrogate model with advanced control strategies, the proposed approach enables efficient exploration of a high-dimensional aerodynamic design space without reliance on repeated wind tunnel experiments.

A wavelet-based neural surrogate, calibrated in the time domain, can efficiently map 6-dimensional fin-angle vectors to top-floor acceleration time histories, providing a fast, stable, and deterministic environment for controller training. A decentralized MARL formulation, in which each fin is modeled as an agent with a Gaussian policy and trained using a scalar reward based on the negative RMS of along-wind acceleration, is capable of identifying improved aerodynamic configurations with minimal algorithmic complexity. For the considered case study, the REINFORCE-based MARL controller achieves an approximately 31% reduction in along-wind RMS acceleration relative to the reference configuration with all fins at zero angle, demonstrating the effectiveness of surrogate-driven policy-gradient optimization.

An extended actor–critic formulation based on PPO, incorporating a multi-objective reward that accounts for resultant and cross-wind responses, significantly enhances control performance. The PPO-based controller achieves reductions of approximately 45% in along-wind, 65% in cross-wind, and over 60% in resultant RMS acceleration, indicating a more balanced and physically meaningful mitigation of aerodynamic loading. The improved performance of the PPO-based approach highlights the importance of coordinated control and multi-objective reward design, enabling effective suppression of both buffeting and vortex-induced responses through implicit coupling of fin actions.

The proposed framework is fully compatible with existing cyber-physical experimental platforms and offers strong potential for extension to more advanced scenarios, including multi-directional wind conditions, time-varying control policies, and real-time adaptive deployment. Overall, the results demonstrate that surrogate-driven reinforcement learning provides a powerful and scalable approach for aerodynamic control of tall buildings. The ability to achieve substantial reductions in structural response, exceeding 60% in resultant acceleration, suggests significant potential for improving occupant comfort, reducing fatigue demand, and enabling next-generation adaptive façade systems in high-rise structures.

#### **ACKNOWLEDGEMENTS**

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) [Grant Number CRSNG RGPIN 2022-03492].

#### **REFERENCES**

- Ding, F. and Kareem, A., 2020. Tall buildings with dynamic facade under winds. *Engineering*, 6(12), pp.1443-1453. <https://doi.org/10.1016/j.eng.2020.07.020>
- Ding, F., Kareem, A. and Wan, J., 2019. Aerodynamic tailoring of structures using computational fluid dynamics. *Structural Engineering International*, 29(1), pp.26-39. <https://doi.org/10.1080/10168664.2018.1522936>
- Fan, D., Yang, L., Wang, Z., Triantafyllou, M.S., Karniadakis, G.E., 2020. Reinforcement learning for bluff body active flow control in experiments and simulations. *Proc. Natl. Acad. Sci. U.S.A.* 117(42), 26091–26098. <https://doi.org/10.1073/pnas.2004939117>
- Font, B., Alcántara-Ávila, F., Rabault, J., Vinuesa, R. and Lehmkuhl, O., 2025. Deep reinforcement learning for active flow control in a turbulent separation bubble. *Nature communications*, 16(1), p.1422. <https://doi.org/10.1038/s41467-025-56408-6>
- Jia, W., & Xu, H., 2024. Effect of synthetic jets actuator parameters on deep reinforcement learning-based flow control performance in a square cylinder. *Physics of Fluids*, 36(8). <https://doi.org/10.1063/5.0220149>
- Li, S., Snaiki, R. and Wu, T., 2021a. A knowledge-enhanced deep reinforcement learning-based shape optimizer for aerodynamic mitigation of wind-sensitive structures. *Computer-Aided Civil and Infrastructure Engineering*, 36(6), pp.733-746. <https://doi.org/10.1111/mice.12655>
- Li, S., Snaiki, R. and Wu, T., 2021b. Active simulation of transient wind field in a multiple-fan wind tunnel via deep

- reinforcement learning. *Journal of Engineering Mechanics*, 147(9), p.04021056. [https://doi.org/10.1061/\(ASCE\)EM.1943-7889.00019](https://doi.org/10.1061/(ASCE)EM.1943-7889.00019)
- Lu, W.T., Phillips, B.M., Jiang, Z., 2023. Surrogate-based cyber-physical aerodynamic shape optimization of high-rise buildings using wind tunnel testing. *J. Wind Eng. Ind. Aerodyn.* 242, 105586. <https://doi.org/10.1016/j.jweia.2023.105586>
- Montalà, R., Font, B., Suárez, P., Rabault, J., Lehmkuhl, O., Vinuesa, R. and Rodriguez, I., 2025. Deep reinforcement learning for active flow control around a three-dimensional flow-separated wing at  $re=1,000$ . arXiv preprint arXiv:2509.10195. <https://doi.org/10.48550/arXiv.2509.10195>
- Suárez, P., Alcántara-Ávila, F., Rabault, J., Miró, A., Font, B., Lehmkuhl, O., & Vinuesa, R., 2025. Flow control of three-dimensional cylinders transitioning to turbulence via multi-agent reinforcement learning. *Communications Engineering*, 4(1), 113. <https://doi.org/10.1038/s44172-025-00446-x>
- Wang, Y., Zhang, H.N., Wang, C.Y., Zheng, X., Li, X.B. and Li, F.C., 2026. Model-based reinforcement learning for active flow control: escalating performance and accelerated training. *Engineering Applications of Computational Fluid Mechanics*, 20(1), p.2628971. <https://doi.org/10.1080/19942060.2026.2628971>
- Whiteman, M.L., Fernández-Cabán, P.L., Phillips, B.M., Masters, F.J., Davis, J.R., Bridge, J.A., 2022. Cyber-physical aerodynamic shape optimization of a tall building in a wind tunnel using an active fin system. *J. Wind Eng. Ind. Aerodyn.* 220, 104835. <https://doi.org/10.1016/j.jweia.2021.104835>
- Yan, L., Cai, H., Wang, Q., Chen, L., Li, C., & Hu, G., 2025. Deep reinforcement learning-based active flow control for a tall building. *Physics of Fluids*, 37(4). <https://doi.org/10.1063/5.0267175>