13th International Symposium on Turbulence and Shear Flow Phenomena (TSFP13)
Montreal, Canada, June 25–28, 2024

OPTIMIZATION OF TOLLMIEN-SCHLICHTING WAVES CONTROL: COMPARISON BETWEEN DEEP REINFORCEMENT LEARNING AND PARTICLE SWARM OPTIMIZATION APPROACH

Babak Mohammadikalakoo
Department of Flow Physics and Technology
Delft University of Technology
Kluyverweg 1, Delft 2629HS, The Netherlands
b.mohammadikalakoo@tudelft.nl

Marios Kotsonis
Department of Flow Physics and Technology
Delft University of Technology
Kluyverweg 1, Delft 2629HS, The Netherlands
m.kotsonis@tudelft.nl

Nguyen Anh Khoa Doan
Department of Flow Physics and Technology
Delft University of Technology
Kluyverweg 1, Delft 2629HS, The Netherlands
n.a.k.doan@tudelft.nl

ABSTRACT

This work focuses on the control of Tollmien-Schlichting (TS) through optimized unsteady suction and blowing jets with the main aim of inducing a delay in laminar-turbulent transition and reducing the resulting skin friction drag. The suppression of TS waves via this Active Flow Control (AFC) system is enabled through two artificial intelligence-based optimization methodologies: Single-Step Deep Reinforcement Learning (SDRL) and Particle Swarm Optimization (PSO). The primary aim of this research is to assess the performance of these methods in optimizing the AFC parameters with respect to convergence rate, computational efficiency, and ability to find an optimum control. The findings demonstrate the success of both methods to find appropriate control parameters leading to effective suppression of the TS waves. While both methods are found to be effective in optimizing the AFC parameters for TS waves' suppression, SDRL outperforms PSO in terms of convergence rate and computational efficiency and exhibits a relatively better performance in finding an improved optimum.

INTRODUCTION

Controlling complex, transitional, and turbulent flows plays a pivotal role in enhancing the overall aerodynamic performance of modern transport and energy systems, enabling the design of more efficient aircraft and wind farms. Active flow control (AFC) methods have gained prominence in various applications, encompassing drag reduction, wind energy power extraction, noise reduction, and more (Pino *et al.*, 2023). Controlling transition from laminar to turbulent flow is particularly crucial, as delaying this process can lead to aerodynamic drag reduction and fuel savings for aircraft and other high speed vehicles, along with reduced greenhouse gas emissions (Niu & Li, 2022). The intricate nature of boundary layer transition however makes this challenging due to the diverse and often contradicting nature of underlying instabilities, such as crossflow vortices, Tollmien-Schlichting waves, Kelvin-Helmholtz, and Rayleigh modes (Luo & Zhou, 1987).

Tollmien-Schlichting (TS) waves, crucial in laminar-to-turbulent transition for two-dimensional boundary layers in low-speed and low-disturbance flows, can be suppressed to delay the transition, extend laminar flow, and reduce skin friction drag making them an attractive target for both passive and active flow control methods (Cossu C., 2002; Kotsonis *et al.*, 2015). However, the control of naturally occurring TS waves in real flows is challenging due to their unpredictable phase and frequency spectrum. Thus, closed-loop strategies based on wave superposition (Thomas, 1983; Liepmann *et al.*, 1982) and advanced adaptive control strategies (Tol *et al.*, 2019; Kotsonis *et al.*, 2015) work well for deterministic, artificially introduced TS waves but not for naturally occurring ones (Tol *et al.*, 2019). This limitation has led to the exploration of modern flow control techniques involving Artificial Intelligence (AI), such as Genetic Programming (GP) (Gautier N. *et al.*, 2015) and Deep Reinforcement Learning (DRL) algorithms (Rabault *et al.*, 2019; Ren *et al.*, 2020).

This study focuses on optimizing the parameters of an AFC system using unsteady suction and blowing jets with a spatial sinusoidal velocity distribution, targeting TS wave suppression. The study aims to evaluate the efficacy of the Single-Step Deep Reinforcement Learning (SDRL) algorithm (Viquerat *et al.*, 2022) and Particle Swarm Optimization (PSO) method (Kennedy & Eberhart, 1995) in optimizing key AFC parameters through a comparative study in terms of convergence rate, computational efficiency, and the capability to discover an enhanced optimum. The investigation involves multiple test cases of increasing complexity, starting with a linear single-frequency TS wave (Herbert, 1994), moving to weakly nonlinear ensembles of multiple TS wave modes, and finally, considering highly nonlinear scenarios with nonlinear multi-frequency control cases.

AI-BASED CONTROLLER

The single-step deep reinforcement learning method as a degenerative DRL and policy-based optimization (PBO)

method is utilized for the optimization of the operating parameters of unsteady suction and blowing jets. The performance of this method in maximizing TS wave attenuation is compared with the PSO technique. The operating parameters include the center location of the jet (C_x), the spatial wavelength (width) (λ_{AFC}), the amplitude (J_A), and the phase (J_ϕ). The latter is measured relative to the phase of the streamwise velocity component of TS wave at the center location of the jet (C_x). The primary goal of this study is to achieve the highest possible reduction in amplitude at a selected downstream location for TS waves, all while minimizing the energy expenditure of the jets.

To simulate the TS wave instability, a Harmonic Navier-Stokes (HNS) solution framework is employed. This solver employs harmonic expansions in both the spanwise and temporal directions for the perturbations. The solution to the HNS equations is obtained through an in-house solver (Westerbeek *et al.*, 2023). The gray box in figure 1 represents the computational domain, where the time-invariant boundary layer profile is established at the inflow using the Falkner-Skan-Cooke equations. This boundary layer commences at the virtual origin of $X_0 = -400\delta_0$ to ensure a developed boundary layer at the computational domain's inflow.

A single-frequency TS wave, introduced by Herbert (1994) with frequency $\omega_1 = 0.0344$ at $Re_0 = 400$, is considered as a reference scenario for all linear control cases. Furthermore, two types of objective functions are used for the optimisation. A first group of cases focuses on directly reducing the maximum amplitude of TS waves at a specified streamwise coordinate. A second group shares the same goal but also includes penalties for the energy expended by unsteady suction and blowing jets. A collection of 16 distinct test cases introduced in Table 1 serves as a versatile and broad foundation for evaluating the performance of SDRL when compared to PSO. The notation of these cases follows a specific format: the first two letters denote the optimization method, namely "SD" for SDRL and "PS" for PSO. The initial number in the code signifies the number of parameters under optimization. If the letter "E" is present in the name, it indicates that the reward function is penalized based on the energy expenditure of unsteady suction and blowing jets. Additionally, the number of attempts for each case is denoted by t1, t2, and t3 in the event of three different trials. For instance, SD3Et2 corresponds to the second trial of single-step deep reinforcement learning, focusing on optimizing three parameters while penalizing the reward function with the energy consumption of the actuators.

RESULTS

The test cases introduced in Table 1 are explored to find the optimal set of operating parameters for the respective AFC system and to compare the performance of the SDRL with the PSO technique. A typical example of successful suppression of TS waves is presented in Figure 2.

Figure 2-a shows the perturbation velocity field in the $x-y$ plane for three different situations. Comparing the velocity field of the TS wave and the wave generated by the unsteady suction and blowing jets, the same amplitude and opposite phase of these waves is observed indicating the opposition control strategy. Near-complete elimination of TS wave as illustrated in Figure 2-a is achievable through this control strategy.

Considering the control case studies defined in Table 1, it is possible to have a flexible and general enough baseline to compare the performance of a machine learning algorithm with a metaheuristic algorithm in terms of convergence rate, computational efficiency, and ability to find the optimum parameters for the AFC system. The general control law extracted from all the obtained results of the SDRL and PSO shows that these two methods discover an opposition control-like law as an effective method of suppressing TS waves similar to the results observed in Figure 2.

Table 1: Linear control cases, Amplitude reduction in TS wave streamwise velocity is denoted with $A_{red.}$ after reaching convergence at $Episode$.

Dataset	TS	Convergence
Case	$A_{red.}(\%)$	Episode
SD3t1	98.23	348
SD3t2	99.96	553
SD3t3	98.29	494
PS3t1	98.31	824
PS3t2	96.99	728
PS3t3	98.28	727
SD3Et1	97.69	477
SD3Et2	99.03	456
SD3Et3	98.68	485
PS3Et1	99.96	801
PS3Et2	99.95	748
PS3Et3	99.95	748
SD2Et1	98.87	286
PS2Et1	99.82	808
SD2t1	98.35	256
PS2t1	99.82	748

The overall suppression of TS waves of up to 99.96% (40 dB) based on results of Table 1 show the success of both the PSO and SDRL algorithms in finding a proper control strategy to suppress TS waves. Comparing the convergence episodes in Table 1, it becomes evident that SDRL consistently outperforms PSO in terms of convergence speed when searching for optimal parameters. The convergence behavior of these two methods is illustrated in Figure 3. The convergence criterion for these methods is considered by defining $\pm 0.2\%$ bandwidth for the rate of variation of the reward's moving average. The observing window for the moving average is over 20 episodes.

By comparing the convergence trends of different optimization parameters in Figure 4 for SDRL and PSO, two different strategies of optimization can be seen which finally leads to a similar control performance. In the PSO, all the parameters to be optimized are converging at almost the same rate and pass the convergence criterion after a similar number of episodes while in SDRL the rate of convergence is different for the three considered operating parameters of the AFC system. For instance, in case SD3Et1 shown in Figure 4, the phase is the first parameter that is converged while amplitude

and width of AFC are the next parameters with a slower convergence rate. It appears that the SDRL narrows down the exploration space dimension as it approaches convergence. This potential reduction in dimensionality is hypothesized to be the underlying factor behind SDRL's faster overall convergence when compared to the PSO. Nevertheless, more extensive investigations into SDRL's optimization strategy are necessary to gain a comprehensive understanding of the primary mechanism behind its faster convergence.

The dashed vertical lines indicated in Figure 3 are representing the minimum number of episodes required to reach convergence for both the SDRL and PSO algorithms. The number of episodes for SDRL and PSO is 477, and 748 episodes respectively, which shows a 57% faster performance of SDRL compared to the PSO algorithm. This difference was observed to be even higher for the cases with only two optimization parameters where SDRL was 3 times faster than PSO. The faster convergence of SDRL compared to PSO is a noticeable advantage of this machine-learning algorithm, especially for fluid dynamic problems with high computational cost per episode.

The ability of SDRL in finding an improved optimum is also evidence of SDRL outperforming the PSO method: when comparing the maximum suppression of TS waves in the different test cases presented in Table 1, SDRL shows an increased suppression. Although the suppression of the TS waves in all case studies is high, slightly different results can still be seen by comparing the SDRL and PSO algorithms. For instance, the maximum amplitude suppression for the PSO method (case: PS3t1) is 38 dB compared to the 40 dB in the SDRL method (case: SD3t2).

By comparing two different optimization scenarios with and without energy penalization (SD3Et2, and SD3t2) in Table 1, it becomes evident that the SDRL strategy is able to compromise overall TS wave attenuation as a result of a trade-off with energy expenditure of unsteady jets. Therefore, the suppression of TS waves is slightly lower when defining the target of optimization with energy penalization.

Nonlinear effects on controller performance

Two nonlinear test cases are considered in this part to study the effect of multi-modal interactions on the control performance of both SDRL and PSO methods. A nonlinear multi-frequency control case with a combination of the first three harmonics of the TS wave is considered while the frequency of the third harmonic is the same as the previously considered linear control cases ($\omega_1 = 0.0344$). The initial amplitude of the first two modes is the root mean square amplitude of 0.5% while the third mode is set at 0.25% (Herbert (1994)). The spatial wavelength of the unsteady suction and blowing jets is fixed for both cases ($\lambda_{AFC} = \lambda_{TS}$). The fixed upstream ($C_x = 660\delta_0$) and downstream ($C_x = 1653\delta_0$) locations of the actuator are selected for cases *NL*1 and *NL*2, respectively. This positioning will help to assess the link between the controller performance and the nonlinear evolution of the TS wave.

Optimal TS wave attenuation and convergence of these control cases is provided in Table 2 which indicates similar control performance of SDRL and PSO methods when the actuator is positioned upstream (*NL*1). However, when the actuator is placed downstream (*NL*2), the SDRL-based controller is outperformed by the PSO-based method. Specifically, the SDRL-based controller is not able to effectively mitigate the second harmonic of the TS wave when compared to the PSO-based controller.

These findings indicate that the extended exploration space in multi-frequency nonlinear control cases adversely affects the convergence rate of the SDRL-based controller, bringing it on par with that of the PSO-based controller. Moreover, the higher complexity of the optimization problem due to the presence of developed nonlinear interactions downstream poses a challenge for the SDRL-based controller resulting in decreased performance.

Table 2: Multi-frequency nonlinear control cases, *NL*1: upstream control, *NL*2: downstream control

Cases #	**Modes**	Δu_{max} **(%)** *PSO*	*SDRL*
*NL*1	M1	99.93	99.22
	M2	98.67	92.12
	M3	95.3	98.97
*NL*2	M1	86.43	82.17
	M2	98.89	77.64
	M3	95.07	97.02

It should be noted that the results presented in this section are based on a small subset of nonlinear case studies. Drawing a definitive conclusion regarding the two methods performance in nonlinear control cases requires conducting additional control cases and sensitivity analyses to evaluate how the SDRL-based controller's initial randomization affects its performance.

CONCLUSION

The application of Single-step Deep Reinforcement Learning (SDRL) and Particle Swarm Optimization (PSO) algorithms to find an optimum unsteady suction and blowing jets as active flow control is studied with the main aim of Tollmien-Schlichting waves' suppression. The utilization of SDRL and PSO-based controllers can suppress linearly developing Tollmien-Schlichting waves by up to 99.96% (40 dB). This indicates the effectiveness of these methods in devising suitable control strategies. The strategy identified by the SDRL and PSO involves an opposition control method, which generates an active flow control wave of the same amplitude but in the opposite phase to the TS wave.

A comparative analysis of SDRL and PSO methods is performed across 16 different test cases. These cases vary in the number of optimization parameters and the definition of the reward function, which either includes or excludes the penalization of performance based on the energy expenditure of the actuator. Faster convergence of SDRL by up to 3 times when compared to the PSO method was observed in single-frequency linear control cases resulting in notable computational power savings. The SDRL-based controller demonstrated superior performance in identifying improved optima compared to PSO. These findings underscore the effectiveness of the SDRL-based controller in suppressing linearly developed TS waves.

Two additional nonlinear test cases were conducted in a multi-frequency nonlinear control scenario to investigate if the compelling features of SDRL-based controller still mani-

fest when dealing with nonlinearly developing TS waves. The findings reveal that the expanded exploration space in multi-frequency nonlinear control cases adversely affects the convergence rate of the SDRL-based controller, making it converge at a slower pace similar to the convergence rate of the PSO-based controller. The increased complexity resulting from downstream nonlinear interactions poses additional challenges for the SDRL-based controller, resulting in a decline in its performance, compared to PSO. Based on the initial investigation into multi-frequency nonlinear control cases, it becomes apparent that the SDRL-based controller no longer demonstrates an advantage over the PSO-based controller in highly nonlinear control scenarios. To fully assess the SDRL-based controller's capabilities in suppressing nonlinearly developed TS waves, more comprehensive studies involving a greater variety of test cases with different levels of nonlinearity are necessary.

REFERENCES

Cossu C., Brandt L. 2002 Stabilization of tollmien-schlichting waves by finite amplitude optimal streaks in the blasius boundary layer. *Phys. Fluids* **14** (L57).

Gautier N., Aider J. L, Duriez T. *et al.* 2015 Closed-loop separation control using machine learning. *Journal of Fluid Mechanics* **770**, 442–457.

Herbert, T. 1994 Special course on progress in transition modelling. *J. AGARD R793* **5**.

Kennedy, J. & Eberhart, R. 1995 Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, , vol. 4, pp. 1942–1948 vol.4.

Kotsonis, M., Shukla, R. K. & Pröbsting, S. 2015 Control of natural tollmien-schlichting waves using dielectric barrier discharge plasma actuators. *International Journal of Flow Control* **7**, 37–54.

Liepmann, H. W., Brown, G. L. & Nosenchuck, D. M. 1982 Control of laminar-instability waves using a new technique. *Journal of Fluid Mechanics* **118**, 187–200.

Luo, J. S. & Zhou, H. 1987 On the generation of tollmien-schlichting waves in the boundary layer of a flat plate by the disturbances in the free stream. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* **413**, 351–367.

Niu, Xiaotian & Li, Jie 2022 Investigation and design of the transonic laminar flow characteristics in a laminar aircraft. *Applied Sciences* **12** (22).

Pino, F., Schena, L., Rabault, J. & Mendez, M. A. 2023 Comparative analysis of machine learning methods for active flow control. *Journal of Fluid Mechanics* **958**.

Rabault, J., Kuchta M., Jensen, A. *et al.* 2019 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of Fluid Mechanics* **865**, 281–302.

Ren, F., Hu, H. & Tang, H. 2020 Active flow control using machine learning: A brief review. *Journal of Hydrodynamics* **32** (2), 247–253.

Thomas, A. S. W. 1983 The control of boundary-layer transition using a wave-superposition principle. *Journal of Fluid Mechanics* **137**, 233–250.

Tol, H. J., de Visser, C. C. & Kotsonis, M. 2019 Experimental model-based estimation and control of natural tollmien–schlichting waves. *AIAA Journal* **57** (6), 2344–2355.

Viquerat, J., Duvigneau, R., Meliga, P., Kuhnle, A. & Hachem, E. 2022 Policy-based optimization: single-step policy gradient method seen as an evolution strategy. *Neural Computing and Applications* .

Westerbeek, S., Sumariva, J. A. F., Michelis, T., Hein, S. & Kotsonis, M. 2023 Linear and nonlinear stability analysis of a three-dimensional boundary layer over a hump. *AIAA SCITECH Forum* .

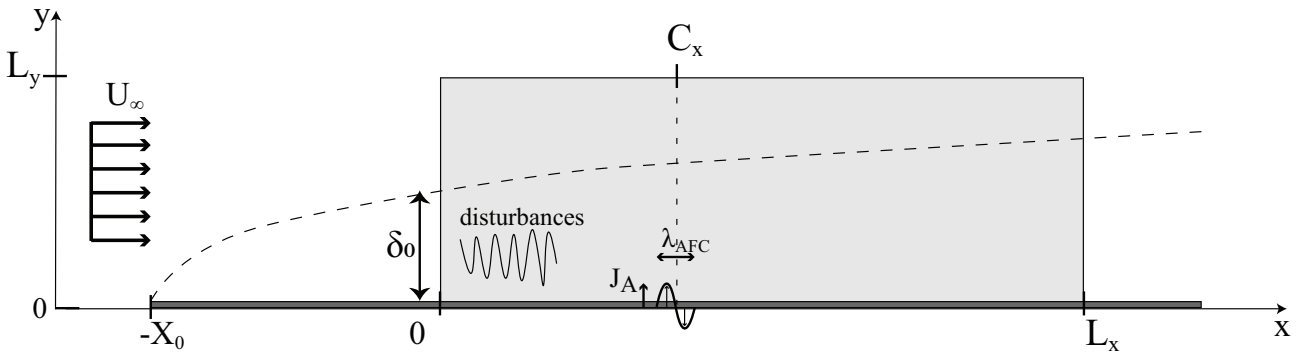

Figure 1: Schematic view of the computational domain, $\Omega = (0, L_x = 2480) \times (0, L_y = 82)$ denoted by the gray region. $\lambda_{AFC} = L \times \lambda_{TS}, L \in [0.5, 2]$: Width of AFC. The AFC center is located at $C_x = 660, 0$. The inlet streamwise velocity U_∞, has a value of $10\frac{m}{s}$ corresponding to Reynolds number $Re_0 = \frac{U_\infty \delta_0}{\nu_0} = 400$, where δ_0 corresponds to the Blasius length scale

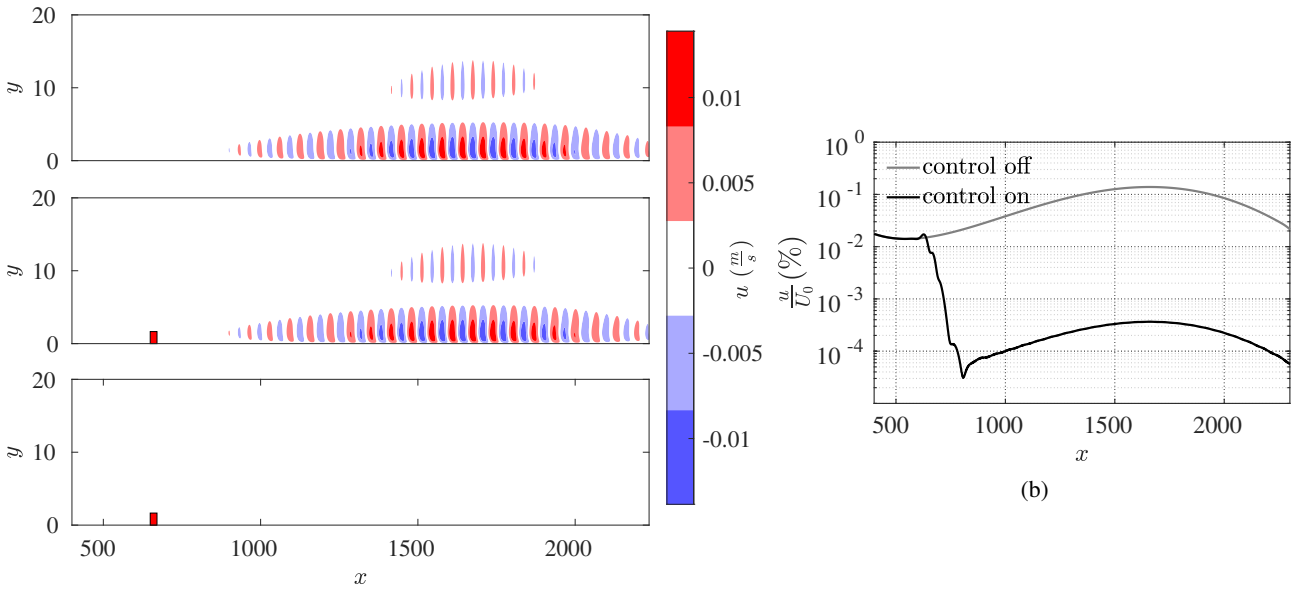

(a) (b)

Figure 2: a) Streamwise velocity of perturbation, top: velocity field when the only disturbance is triggered at the inflow, middle: velocity field with only control implemented upstream, bottom: velocity field of the controlled TS waves, the position of the AFC is illustrated with a red rectangle ($x_s = 660$). b) Amplitude evolution of streamwise (u) disturbance velocity, Nondimensionalized data using Blasius length scale (δ_0)

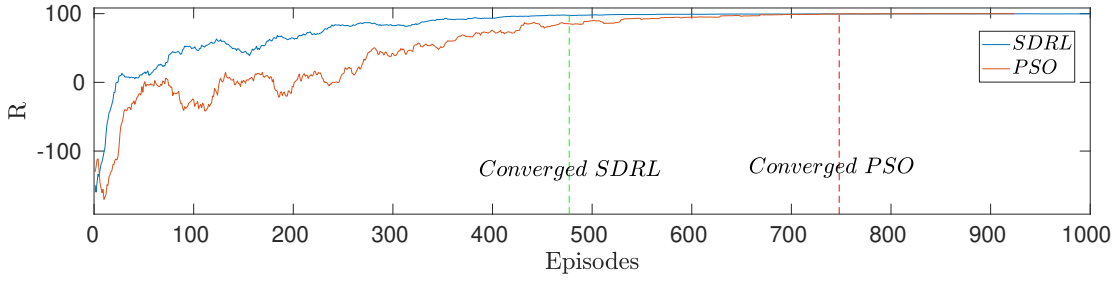

Figure 3: Evolution of reward function for SDRL and PSO algorithm (SD3Et1 vs PS3Et3), the solid lines are referring to moving average of the reward function

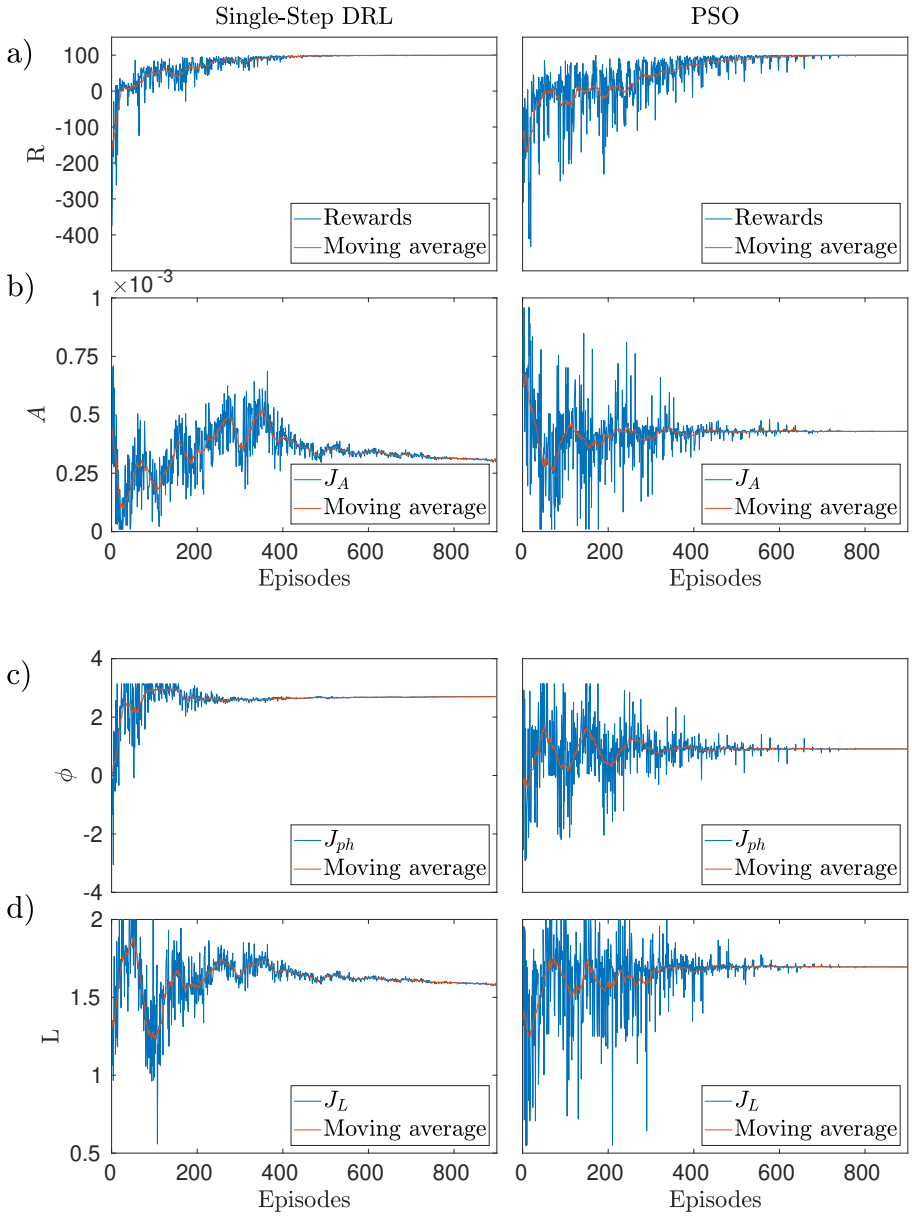

Figure 4: Convergence of SDRL and PSO algorithm towards optimum control (SD3Et1 vs PS3Et3), a) Reward function, b) Amplitude of unsteady suction and blowing jets (AFC), c) Phase of AFC, d) Relative Width of AFC