

SEMI-ACTIVE STRUCTURAL CONTROL USING VISCOUS DAMPERS AND REINFORCEMENT LEARNING

A. JEDLIŃSKA*, D. PISARSKI, G. MIKUŁOWSKI, B. BŁACHOWSKI, Ł. JANKOWSKI

Institute of Fundamental Technological Research (IPPT PAN)
Polish Academy of Sciences
02-106 Warsaw, Poland
e-mail: ajedlins@ippt.pan.pl | ljank@ippt.pan.pl

Abstract. This contribution presents an approach to structural control based on reinforcement learning. Reinforcement learning, a rapidly developing branch of machine learning, is based on the paradigm of learning through interaction with the environment. Here, it is applied in the context of semi-active structural control, where the considered actuators take the form of viscous dampers with a controllable level of damping. The control forces are thus coupled with the structural response, and the formulation is intrinsically nonlinear. The related optimum control problems are usually more difficult than in the case of active structural control systems, which generate and apply arbitrary external control forces. Analytical derivation of the optimum semi-active control is thus rarely possible, so that many control algorithms applied in practice are suboptimal and/or heuristic in nature. Here, an effective control strategy is developed by means of the Q-learning approach. The control algorithm is determined in interaction with the controlled system, that is, by applying initially random control sequences in order to observe, process, and optimize their effects. Such an approach seems to be new and relatively unexplored in the field of structural control. Verification is performed in a numerical experiment, where the Q-learning procedures interact with an independently simulated finite element model of a structure equipped with a tuned mass damper (TMD) and a controllable viscous damper. The results attest to a performance significantly better than that of an optimally tuned conventional TMD.

Key words: Reinforcement Learning, Semi-active control, Structural control, Damping, Vibration

1 INTRODUCTION

This contribution presents a control strategy for shear-type building structures under seismic excitation by customizing, developing, and applying the machine learning techniques of reinforcement learning (RL). Structural vibration in engineering structures can have detrimental impact on structural conditions and operation. It is also a crucial aspect in engineering safety, as excessive vibrations can negatively impact structural integrity. To mitigate these detrimental effects, it is crucial to prevent the occurrence of harmful vibrations. Various approaches have been developed, including passive, active, and semi-active control methods. This contribution focuses on semi-active control by means of a TMD (tuned mass damper) with a switchable level of viscous damping. TMD is a device used to reduce vibrations in structures such as high-rise buildings and bridges by adding a secondary mass that opposes the motion of the main structure. The TMD is typically attached to the structure via a spring and damper system and is tuned to a specific natural frequency to effectively counteract the vibrations [1]. The TMD method has proven to be highly effective and widely applicable, particularly in high-rise building design [2].

TMD systems, developed in the 1970s as passive vibration control systems for building structures, have been implemented in many high-profile buildings. Accurate calculation of a structure's fundamental vibration frequency is not always possible, and this frequency can change during extreme dynamic events, such as strong ground motion. TMD systems may be partially effective for vibrations under ordinary winds when the fundamental frequency of the structure dominates the response. However, they are less effective for irregular structures under strong ground motion when several vibration modes significantly contribute to the dynamic response of the structure [3].

This contribution applies the concept of reinforcement learning in the context of semi-active structural control systems and analyzes its impact on the control performance. Reinforcement learning has many advantages that make it well-suited for solving problems that are challenging for other structural control techniques. Its ability to learn from experience, adapt to changes in the environment, and take optimal decisions make it a valuable technique in many fields. Although the use of RL in structural control, particularly in semi-active control, is not yet widespread, there is a growing interest in its potential to improve the performance of control systems [4,5]. The main aim of this contribution is to test the application potential of the RL in semi-active structural control. The investigated structure is a numerical model of an 11-story shear-type building equipped with a semi-active TMD. The TMD is controlled by switching its viscous damping coefficient in an on/off manner. The main result is a Q-learning control algorithm that mitigates the vibrations due to random seismic excitation significantly more efficiently than the optimally tuned conventional TMD. Such a result provides initial insights into the potential of reinforcement learning for improving the performance of semi-actively controlled damping systems.

2 REINFORCEMENT LEARNING

2.1 The technique

Reinforcement learning (RL) is a branch of machine learning that focuses on how intelligent agents can learn to make optimal decisions in dynamic and uncertain environments. Unlike supervised learning, which relies on labeled data, and unsupervised learning, which involves identifying patterns in data without explicit feedback, reinforcement learning involves training an agent to learn the best actions through a trial-and-error process based on its interactions with an environment. The agent receives feedback in the form of rewards or penalties, which it uses to improve its decision-making policies over time. Reinforcement learning has shown remarkable successes in a wide range of applications, including robotics, gaming, and finance, and it has the potential to revolutionize many other fields in the future [6]. However, reinforcement learning also presents a number of unique challenges, including the exploration-exploitation dilemma, the credit assignment problem, and sample inefficiency, which continue to be areas of active research in the field [7].

This contribution tests the potential of the RL in semi-active structural control. The RL approach contrasts with supervised learning (which requires optimum control sequences that are rarely known in semi-active control) and even more with unsupervised learning (which relies on the exploration of input data only). It allows learning from interaction and can be seen as a completely new approach to structural control. By allowing the system to learn, adjust and optimize its control strategy based on real-time feedback from the structure, reinforcement learning has the potential to significantly enhance the effectiveness of semi-active damping. The technique applied here involves training an artificial neural network by interacting with a

simulated environment that contains a numerical model of the structure. The structure is subjected to random seismic excitation and its dynamics is simulated using the Newmark integration method. The main result is a Q-learning control algorithm that mitigates structural vibrations significantly more efficiently than a conventional tuned mass damper.

In Q-learning, the agent makes a decision based on observations, by comparing the learned expected reward value for different possible actions. In the context of semi-active control, Q-learning has emerged as a suitable approach for several reasons. First, it is well-suited to problems that involve discrete actions, which is the case in bang-bang type control that is common in semi-active control systems. Second, Q-learning is computationally efficient and can handle large state-action spaces, which is important for real-world applications. Finally, Q-learning is flexible and can be straightforwardly adapted to different control approaches overall.

2.2 Architecture of the RL agent

In recent years, artificial neural networks (ANNs) have been successfully used to approximate the value function in Q-learning, resulting in the development of deep Q-networks (DQNs). DQNs have been shown to outperform traditional Q-learning methods, especially in tasks with large and continuous state or action spaces. The use of neural networks in Q-learning allows the algorithm to learn more complex and abstract representations of the state space, improving the algorithm's ability to generalize to new situations. In addition, DQNs are able to learn directly from raw sensory inputs, removing the need for hand-engineered feature extraction.

The RL agent in this contribution employs a dense ANN to learn and encode the value function. It is implemented in the Python programming language using TensorFlow and Keras, which are two popular open-source libraries for implementing ANNs. TensorFlow is a low-level library for building and training machine learning models, while Keras is a high-level API that simplifies the process of building neural networks. An artificial neural network has been used with 6 hidden sequential dense layers, each consisting of 40 neurons. The input layer provides the network with measurements of structural response. The output layer has two neurons, which correspond to the possible states of the control signal. The activation function used in the neural network is rectified linear unit (ReLU), which is a common choice in deep learning due to its ability to improve convergence during training. Such a network architecture is fully sufficient to effectively encode the dynamics of the employed structure and the expected cumulative future rewards.

3 STRUCTURE

3.1 Shear-type building

The investigated structure is an eleven-story shear-type structure with a TMD as a classical engineering device, consisting of a mass, a spring, and a viscous damper, attached to the last story, as shown in Figure 1. The equation of motion for such a building model experiencing seismic excitation, without a control system, can be expressed as:

$$[M]\{\ddot{u}\} + [C]\{\dot{u}\} + [K]\{u\} = -[M]\{r\}a(t) \quad (1)$$

The total number of the degrees of freedom (DOFs) is $n=12$, which corresponds to the eleven stories and the single TMD. The column vector $\{u\}$ has n rows and represents the absolute displacements of each story and the TMD. The column vector $\{r\}$ also has n rows and denotes the displacement of each DOF resulting from the application of unit horizontal ground

displacement, that is, it consists of 12 ones. The seismic excitation (ground acceleration) is denoted by $a(t)$. The matrices $[M]$, $[C]$, and $[K]$ are $n \times n$ and represent the mass, damping, and stiffness of the structure, respectively. The mass matrix is diagonal, and the mass of each story and the TMD are listed on the diagonal. This assumes that the masses are lumped at the floor levels. The damping matrix is proportional to the stiffness matrix. The proportionality coefficient is chosen to achieve 2% of critical damping for the first mode of vibration of the structure without the TMD. Building specifications, including the number of stories, their masses, and stiffness, are taken from literature [8].




Figure 1. The investigated 11-DOF structure with a semi-active TMD on the top level

This report focuses on analyzing the effectiveness of a semi-active control system in mitigating the responses of the investigated shear structure. To achieve this goal, a semi-active tuned mass damper (TMD) system, illustrated in Figure 1, is installed on the top floor. Additional TMDs can be incorporated if necessary. The three essential parameters in a TMD system are the TMD mass, TMD stiffness, and TMD damping coefficient. In the considered example, the mass of the TMD constitutes 3% of the total mass of the building, while its stiffness is selected as 2,509,600 N/m. The damping coefficient is discussed in the following subsection.

3.2 Seismic load and feedback signal

Structures subjected to seismic loads can experience significant damage due to the severity and unpredictable nature of ground motions. To test the effectiveness of a control system in mitigating dynamic response (and thus the damage), it is usually necessary to evaluate its performance under a range of different ground motions. In this contribution, to prevent the RL agent from learning to respond to a specific set of ground motions only (which would correspond to overfitting in supervised learning), the seismic load $a(t)$ in equation (1) is assumed to be white Gaussian noise and is simulated anew for each training and evaluation episode. This ensures that the measurements used to train and evaluate the proposed control system are based on a variety of ground motions and not biased towards any particular motion.

The measurements used as a feedback in training and control are linearly transformed full state vectors, consisting of the relative displacements and velocities between the ground, successive floors, and the TMD. This means that the RL agent has full information about the state of the structure. Such a choice allows evaluating the ultimate potential of the control with the proposed RL agent.

3.3 Control and evaluation

The considered control is of the bang-bang type and affects the viscous damping coefficient of the TMD. This results in the control signal that directly affects the damping matrix $[C]$ in equation (1) in a linear manner, which corresponds to bilinear control. In such control systems, the open loop optimal control is often of the bang-bang type. Consequently, in the controlled structure considered here, the damping of the TMD can be switched between two states: no damping and very high damping. In the latter case, the damping coefficient is large enough to model (in transient analysis) an effective merging of the TMD mass with the mass of the top floor.

The objective of the control is to minimize the vibrations of the top floor of the building, as quantified by the root mean square (RMS) of its displacement. During testing of the RL agent, three RMS values can be compared, two of which may be used as references:

- 1) RMS in the RL-controlled structure,
- 2) RMS in the structure equipped with the optimally tuned passive TMD (reference 1: passively controlled system),
- 3) RMS in the passive original 11-story structure without the TMD (reference 2: passive uncontrolled system).

To generate these three RMS values, exactly the same white noise excitation is generated and applied to the structure.

Figure 2 presents plots of the amplitudes of three frequency response functions (FRFs) of the considered system (last floor displacement related to ground acceleration). Three cases are depicted, which correspond to the three following levels of the TMD damping:

- the optimally tuned passive TMD (green line),
- the TMD without damping (blue line) and
- the highly damped TMD (yellow line).

The considered control amounts to dynamic switching between the blue and yellow characteristics, while the reference controlled system (optimally tuned passive TMD) corresponds to the green line.




Figure 2. FRF amplitudes of the structure depicted in Figure 1 for three levels of TMD damping

4 RESULTS

4.1 Displacement RMS

Response root mean square (RMS) is widely used in structural control as a performance metric. In reinforcement learning, RMS is often used as a measure of the error or distance between the controlled system trajectory and the desired trajectory. It can be used to evaluate the performance of an RL agent and the effectiveness of its actions. Figure 3 plots the top floor displacement RMS per test episode, which is simulated (with zero exploration rate) every 10 training episodes. A relatively quick convergence to 75% of the RMS of the reference passively controlled structure can be observed, which attests to the high effectiveness of the proposed RL control.




Figure 3. RMS response ratio per test episode (RL-controlled system to reference passively controlled system). Test episodes are simulated every 10 training episodes. Zero exploration rate is used

4.2 Convergence of rewards

Training of the agent is performed based on the rewards obtained from each training interaction episode. Rewards per episode are calculated based on the vibration level of the top floor. At the end of each episode, the agent receives a numerical reward signal that reflects the quality of its behavior during that episode and quantifies the cumulative distance from the equilibrium point in all time steps. The rewards per episode are used to update the value function, so that the agent can improve its behavior in subsequent episodes. Figure 4 presents the characteristic increase of the total reward per training episode. The value of 1000 corresponds to a completely still structure. The results on the chart include the effect of a 10% exploration rate.




Figure 4. Total rewards per training episode

4.3 Time-domain responses

Figure 5 presents examples of time-domain displacement responses of the top floor in three cases: the RL-controlled structure, the reference passively controlled structure, and the uncontrolled original structure (without TMD). A clear decrease in amplitudes can be noted: first by applying the optimum passive TMD, and then by using the RL agent.

5 SUMMARY

This contribution considered an application of reinforcement learning in semi-active structural control. An 11-story shear-type building equipped with a TMD was considered. The control signal was of the bang-bang type and related to the viscous damping of the TMD. The results show that the proposed control system, using the full state vectors as measurement feedback, achieved a significant improvement in reducing the structural response to seismic excitations compared to the optimally tuned passive TMD.




Figure 5. Examples of time-domain displacements

ACKNOWLEDGEMENT

This research was funded in whole or in part by the National Science Centre, Poland, under the grant agreement 2020/39/B/ST8/02615. For the purpose of Open Access, the authors have applied a CC-BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

REFERENCES

- [1] A.K Rathi, A. Chakraborty, "Reliability-based performance optimization of TMD for vibration control of structures with uncertainty in parameters and excitation." *Structural Control & Health Monitoring*, vol. 24, id. e1857, 2017.
- [2] L. Jin, B. Li, S. Lin, G. Li, "Optimal Design Formula for Tuned Mass Damper Based on an Analytical Solution of Interaction between Soil and Structure with Rigid Foundation Subjected to Plane SH-Waves.", *Buildings*, pp. 13-17, 2023.
- [3] N.R. Fisco, H. Adeli, "Smart structures: Part I—Active and semi-active control", Department of Civil and Environmental Engineering and Geodetic Science, The Ohio State University, *Scientia Iranica*, vol. 18(3), pp. 275-284, 2011.
- [4] A. Bernard, I.F. Smith, "Reinforcement learning for structural control.", *Journal of Computing in Civil Engineering*, 22.2, pp. 133-139, 2008.
- [5] Z. Dworakowski, K. Mendrok, "Reinforcement learning for vibration suppression of an unknown system.", *Mechanisms and Machine Science*, vol. 73, pp. 4045-4054, 2019.
- [6] D. Silver, T. Hubert, "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play.", *Science*, vol. 362.6419, pp. 1140-1144, 2018.
- [7] F.L. Lewis, D. Vrabie, K.G. Vamvoudakis, "Reinforcement learning and feedback control: Using natural decision methods to design optimal adaptive controllers.", *IEEE Control Systems Magazine* vol. 32(6), pp.76–105, 2012.
- [8] S. Pourzeynali, H.H. Lavasani, A.H. Modarayi, "Active control of high rise building structures using fuzzy logic and genetic algorithms", *Engineering Structures*, vol. 29, pp. 346-357, 2007.