

Continuous Lyapunov Controlled Non-linear System Optimization Using Deep Learning with Memory

Amr Mahmoud^{1,*}, Youmna Ismaeil², Mohamed Zohdy¹

¹Department of Electrical and Computer Engineering, Oakland University, Rochester, MI, USA

²Department of Computer Science, Saarland University, Saarbrücken, Germany

Abstract Initially selected system and controller parameters don't often guarantee continued system stability and performance mainly due to the introduction of unexpected system disturbances or unknown system dynamics. In this research we present a novel approach for detecting early failure indicators of non-linear highly chaotic system and accordingly predict the best parameter calibrations to offset such instability using deep machine learning regression model. The approach proposed continuously monitors the system and controller signals. The Re-calibration of the system and controller parameters is triggered according to a set of conditions designed to maintain system stability without compromise to the system speed, intended outcome or required processing power. The deep neural model predicts the parameter values that would best counteract the expected system in-stability. To demonstrate the effectiveness of the proposed approach, it is applied to the non-linear complex combination of Duffing-Van der pol oscillators. The approach is also tested under different scenarios the system and controller parameters are initially chosen incorrectly or the system parameters are changed while running or new system dynamics are introduced while running to measure effectiveness and reaction time.

Keywords System parameterization, Deep Machine Learning, Complex system, Non-linear controller, Duffing-van der pol, Lyapunov control

1. Introduction

Lyapunov control has been proven successful in controlling highly chaotic non-linear oscillators [1] [2] [3]. One of the fundamentals that contribute to the success or failure of any type of control strategy is the controller and system parameters. Therefore, researchers have explored different methods to find the precise parameters that would lead to achieving the best system results [4] [5].

One of the methods utilized to achieve the previously mentioned goals is Genetic algorithm (GA). GAs have been successful in cases where all the system dynamics are clearly defined and known to some extent or with systems where limited system disturbances are introduced and minor parameter tuning is required [6] [7]. In some cases, several system assumptions are needed in order to allow the GA to run successfully.

Due to some of the limitations found in using GAs such as inability to quickly converge to the final solution or adapt to unknown system dynamics or unknown disturbances. Researchers though after different approaches that wouldn't reduce the system agility and at the same time would be

able to handle unknown system characteristics. A hybrid approach of Fuzzy Control and GAs was researched in [8] but system linearization is a requirement in order to use the previously mentioned method [8].

Another approach that is recently being researched is the use of Machine Learning to enhance the controller performance. For example, through the use of Episodic learning [9] [10]. Most recently, there is the introduction neural Lyapunov control which proposes the use of deep learning to find the control and Lyapunov functions. The approach mentioned in [9] [10] is suitable for find the best system parameters that would initially lead the system to stability and reduced the system error. The problem with approach in [9] [10] is that it assumes that the system is deterministic, time invariant, and affine in the control input. while in real life situation external perturbations might occur resulting in system failure at any moment while the system is running [11] [12]. The approach proposed in [25] [26] and [27] attempts to predict the control and Lyapunov functions that would lead to system stability but under specific conditions where the system dynamics are deterministic in nature. The approach proposed in this research is novel to the best of our knowledge in that it discards the assumption of an ideal environment or fully known system dynamics and seeks continuous enhancement of the controller outcome through continuous monitoring of the system error, reference signal, system dynamics and

* Corresponding author:

amahmoud@oakland.edu (Amr Mahmoud)

Received: Sep. 22, 2020; Accepted: Oct. 16, 2020; Published: Oct. 26, 2020

Published online at <http://journal.sapub.org/control>

control signal and accordingly adjust the system and controller parameters to improve the controller performance without the need to disrupt the system output.

The focus of the research is to allow the Deep Learning Algorithm to learn the system from a continuously improving dataset and according to the slope of the output error the algorithm relearns the system and collects the needed information.

The proposed method is applied to a non-linear chaotic combined system of Duffing and Van der pol oscillators [24]. The aforementioned system was chosen to test the Deep learning algorithm response to unpredicted system disturbances and unknown system dynamics [21] [22] [23]. An algorithm was developed to aid and trigger the Deep Neural Network when needed to adapt to new system dynamics and according to pre-set conditions. The algorithm records and feeds an updated data set to the DNN in order to relearn the system dynamics if certain conditions are detected to be true. Once the process of retraining is complete, the algorithm generates a random array of parameters and the array of parameters is fed to the DNN to predict the output error. If the DNN predicts that the error slope will be reduced from the current value improving the performance of the system, then the suggested controller parameters are allowed to be set as the new controller parameters. Otherwise the algorithm would create a new array set of randomly generated parameters and re-feed them to the DNN to provide its predictions in a continuous loop. Once the controller parameters are updated the algorithm monitors the actual error compared to the predicted error to determine the network viability and accordingly if the error difference is greater than a set value the algorithm would trigger a retraining request of the DNN.

2. Lyapunov Control on Duffing - Vanderpol Oscillator Model

An easy way to comply with the paper formatting requirements of SAP is to use this document as a template and simply type your text into it.

This section provides an overview of the Duffing-Van der pol system dynamics. Duffing-Van der pol mathematical model is

$$\ddot{x} + \delta\dot{x} + \epsilon_1(x^2x + x\dot{x}^2) + \epsilon_2x^3 = \mu(t) + \gamma \cos(\omega t) \quad (1)$$

It combines both non-linear stiffness and non-linear damping Equ.1. Setting

$$\dot{x} = x_2 \quad (2)$$

and

$$x = x_1 \quad (3)$$

Therefore

$$\dot{x}_1 = x_2 \quad (4)$$

$$\dot{x}_2 = \frac{1}{1+\epsilon_1x^2}(\delta\dot{x} - \epsilon_1x\dot{x}^2 - \epsilon_2x^3 + \mu(t) + \gamma \cos(\omega t)) \quad (5)$$

we get the state space model

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ \frac{1}{1+\epsilon_1x^2}(\delta x_2 - \epsilon_1 x x_2^2 - \epsilon_2 x^3) \end{bmatrix} \quad (6)$$

$$y = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (7)$$

3. Lyapunov Controller Design

In this section we give an overview of the controller design. The controller is designed with the purpose of achieving Duffing- van der pol oscillator system stability at an increased forcing frequency. Desired state q_d and actual state q .

$$e = q_d - q \quad (8)$$

A control LYAPUNOV candidate is then chosen to be

$$V = \frac{1}{2}(\gamma_1 e + \gamma_2 \dot{e}) \quad (9)$$

As shown from Equ.11, the transient behavior of the error dynamics can be influenced by a suitable choice of the CLF and gamma parameters, respectively Equ.10. Therefore, the CLF has been chosen to be positive definite for all

$$x, \dot{x} \quad (10)$$

taking the time derivative of V Eqn.9

$$\dot{V} = (\gamma_1 + \gamma_2 e) (\gamma_1 \ddot{e} + \gamma_2 \dot{e}) \quad (11)$$

$$\dot{V} = -kV \quad (12)$$

which is globally exponentially stable as V is globally positive definite. Hence, by substitution in Eqn.11 we find the control law Eqn.13.

$$u = \frac{1+\epsilon_1x^2}{\gamma_2} \left(\frac{k}{2} \gamma_1 e + \frac{k}{2} \gamma_2 \dot{e} + \gamma_1 \ddot{e} + \gamma_2 \ddot{x}_{des} \right) - P \cos(\omega t) + \delta\dot{x} + x + \epsilon_1\dot{x}x^2 + \epsilon_2x^3 \quad (13)$$

4. Simulation Results for Unoptimized System




Figure 1. Reference signal (dotted green) 5th harmonic function with amplitudes of 0.1, 0.5 and 1 and Lyapunov controller output (red) for $t=2.5s$

Section 4 will provide an overview of the results after running the system with the Lyapunov controller and manually selecting the parameters that lead to reduced error across the time frame of 2.5s. The Duffing- van der pol oscillator model was ran with the following selection of

parameters. $\delta=0.5$, $\varepsilon_1=1.6$, $\varepsilon_2=-0.8$, $P=3$, $\omega=10$ and the control parameters $\gamma_1=12$, $\gamma_2=4$, $k=115$. The parameters were manually adjusted and tuned to give low error 1 and high stability 2.




Figure 2. Phase portrait showing a stable system with no interference under the harmonic reference signal $t=2.5s$

5. Lyapunov Controller Parameter Optimization

5.1. Overview

The process of selecting the proper Sigmas to reduce the error of the model can be modelled as an exhaustive search problem. Where several Sigmas are tried and the one that results in a reduction in the system error is selected. Though this process results in the best Sigmas, yet it is time consuming. To overcome this problem, we developed a deep learning model that researches different system and controller parameters.

A given sigma is good if it reduces the error of the system. The pipeline is presented in Fig. 5. Initially, the model is run, and the system error is monitored. If the error surpasses a certain threshold the neural network is queried for a sigma that would reduce the system error given the current state of the system. The algorithm for the entire process is defined below. The network proposed in Fig.3 is initially trained and then used to predict new Sigmas; Sigma 1 (s_1), and Sigma 2 (s_2) if the system error (sys err) exceeds a given threshold (0.8).

The network is then given the current state of the system in terms of time (t), position (p), velocity (v), and the randomly generated (s_1) and its corresponding (s_2) and it predicts the system error. If the predicted error is less than the current system error, then the system is updated with the newly generated Sigmas that are speculated to reduce the error. Every iteration the newly generated Sigmas are their corresponding system parameters are stored. After each 100 iterations the system is probed for perturbations by

computing the average of the system error throughout the previous 100 iterations (avg sys err) if this error is larger than the previous system average (prev sys avg) the network is retrained on a portion of the old data (Memo) used in previous training as well as the data from the previous 100 iterations (new data). The motive behind using the old data as well as the new data is to avoid catastrophic forgetting as mentioned in [14].

The idea of using a memory is usually known as self-refreshing memory and was initially introduced and tested in [15] [16] [17]. The idea here is that even when there are perturbations they might not last for even so long and consequently we don't want the network to lose its ability to predict the System error if the system parameters return back to norm after the network was retrained several times.

Memory = 10% of the old training data (prev sys avg = system error average) of the training data initialization.

```

while n > 2000 do
    sys err = current system error;
    if sys err > 0.8 then
        while True do
             $s_1$  = randomly generated;
             $s_2$  = -  $s_1/8$ ;
            predicted sys err = predict (t, V, p, v,  $s_1$ ,  $s_2$ );
            if predicted sys err < sys err then
                Break;
            end
        end
        end
        update controller ( $s_1$ ,  $s_2$ );
        new_data = Save (t, V, p, v,  $s_1$ ,  $s_2$ , new-sys err);
        avg sys err = avg sys err + new sys err;
    end
    if n == 0 then
        if avg sys-err / 100 > prev avg then
            retrain network (Memo + new data);
            Memo = Memo + 10% of the new data;
            prev sys avg = avg sys err;
            Clear new data;
            Clear avg sys err;
        end
    end
end
end

```

Algorithm 1. Adaptive Algorithm Steps for Sigma prediction

5.2. Dataset

The Time, Error, velocity, position and control signal were collected every 1 millisecond and added into an array of values. The dataset was split into 60% training and 40% test. A continuous update to the dataset is done according to the calculated error slope every 10ms (100 iterations) as explained above.

5.3. Network Architecture

The network architecture is in Fig. 3 The network is composed of 5 blocks (15 layers). The number of layers 15 was chosen for its efficiency and high performance. It was found that if a lower number of layers is used the DNN

performance was affected and when a higher number of layers was used it had no effect on the performance but reduced the DNN efficiency. A block is composed of 3 layers: a fully connected convolution layer followed by batch normalization layer and rectified linear unit (RELU) in Eqn.14.

$$\text{output} = \max(0.0, \text{input}) \quad (14)$$

The batch normalization layer allows us to use higher learning rates as it makes ensures that activations are not too high or low. The RELU is a nearly linear function in the sense that it is a piece wise linear function with two linear pieces. This feature increases its capability to preserve many of the properties that make linear models easy to optimize with gradient based methods as well as the properties that make linear models generalize well. The last block has a dropout layer aside from the fully connected convolution layer and a regression layer to predict the error. Dropout [18] is a regularization method that was developed to solve the problem of overfitting. The overfitting problem occurs when the neural network learns every minor detail in the training data. The effect of such a behaviour is that the network will have high accuracy on the training data but very low accuracy on the test data. This means that the network is unable to generalize on unforeseen data. The proposed solution for this problem is the dropout technique. The dropout technique ignores a number of layer outputs randomly. Such a behaviour makes the layer appear as and be treated-as a layer with a different number of nodes and connectivity to the prior layer. Consequently, each update to a layer during training is performed with a different “perspective” of the configured layer. The regression layer computes the mean squared error loss. The network is trained to learn a function F in Eqn.15.

$$F : (t, x, v, s1, s2) \rightarrow e \quad (15)$$

Where $s1$ and $s2$ are the two Sigmas, t is the time and e is the error of the model.

5.4. Network Training and Testing

The network is initially trained for 5 epochs using Adam optimization algorithm. The learning rate is initially set to 0.001 and is reduced by a factor of 0.2 every 5 epochs. This allows large weight changes in the beginning of the learning process and small changes or fine-tuning towards the end of the learning process. This gives more time for fine-tuning. While the model is running if the system error passes a given threshold the Neural network is used to choose two Sigmas that would lower the error. This is done by selecting a uniformly distributed random number in the interval $[-50, 50]$ for $s1$.

The restriction of $s1$ and $s2$ was added later in the research to avoid the genetic algorithm and DNN going into a continuous loop of changing the Sigmas to find the perfect candidate. It was also found that the smaller the sigma the better the system outcome but in order to maintain controller

flexibility the range was set to $[-50, 50]$ to account for unexpected changes in the system behaviour. The network does not pass the proposed Sigma 1 and Sigma 2 to the model and controller unless Sigma 1 and Sigma 2 are predicted to reduce the system error. This helps in tuning the entire pipeline to take small/large steps to change the behaviour of the duffing van der pol model given that sudden large changes in the Sigmas may result in system failure. The idea behind predicting the error and not the Sigmas is the fact that we can have more than one error for a given Sigma, which gives more flexibility for choosing the range of sigmas to consider. The root of the mean square error graph for the first 5 epochs are plated in Fig. 4. The results show that the neural network presented in Fig. 5 was able to predict the error with an error of 0.03564.

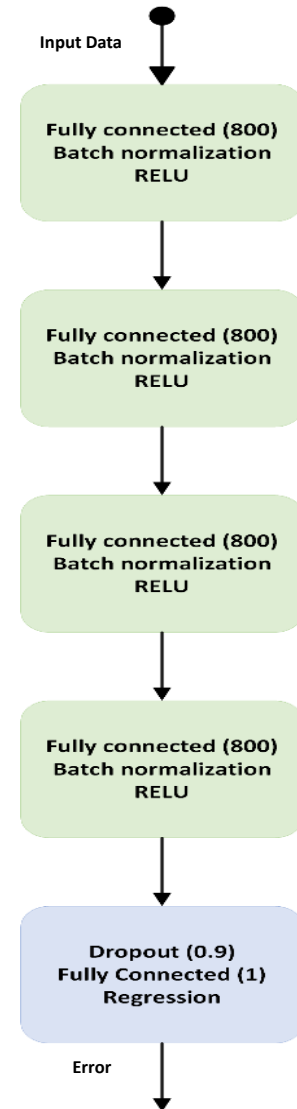


Figure 3. Neural network used for predicting the best Sigmas that would result in a decrease in the system error. The input to the network is the time (t), displacement (x), velocity (v), Sigma 1 ($s1$), Sigma 2 ($s2$) and the output is the expected system error (Error) for the system. $s2 = -s1/8$




Figure 4. Neural network used for predicting the best Sigmas that would result in a decrease in the system error. The input to the network is the time (t), displacement (x), velocity (v), Sigma 1 (s1), Sigma 2 (s2) and the output is the expected system error (Error) for the system




Figure 6. The system error after using the neural network in Fig. 7 to get a good estimate of s1 and s2

6. Results and Discussion

In Fig. 5 the system is shown to go into instability and fails 0.3 seconds from start due to the use of unfit Sigma 1 = 100 and Sigma 2 = 0.6 while in Fig. 5 we show that using the proposed deep neural network to predict the appropriate Sigma 1 and Sigma 2 lead to stabilizing the system and maintaining the system error under the set threshold.

The results show the efficiency of the neural network in predicting the error given the sigma, as well as the efficiency of the algorithm in preventing the system from failing likewise in continuously enhancing its performance by keeping the system error as low as possible. On the other hand, maintaining a constant sigma results in large system error and the system might eventually fails.




Figure 5. The system error goes to infinity as shown when the algorithm is not applied

6.1. Algorithm Reaction to Parameter Sabotage




Figure 7. Error uptick at $t = 0.8867$




Figure 8. The Algorithm reacting to the sudden change by adjusting Sigma 1

To elaborate on the ability of the proposed solution to tackle mid system instability or sudden changes. The system parameters are manually overwritten while the system is running to measure the NN and Algorithm effectiveness in returning the system to stability.

As demonstrated by Figs 7 and 8 the deep network was able to predict the best parameters to re-establish stability of the system and controller within 0.4 ms.

6.2. Improving the System Performance

In this section we show the ability of the proposed method in finding the best system parameters while the system and controllers are running.




Figure 9. The Algorithm reacting to the sudden change by adjusting s_2




Figure 10. Phase diagram of the algorithm before and after finding the optimal system parameters for the conditions which shows major improvement compared to Fig.4

6.3. Network Retraining

During the re-training phase of the neural network we

noticed that without the use of the memory it might take longer and, in some cases, go into an infinite state without finding the optimum solution. The speed with which the system returns back to its stable state affects the retraining phase as we only consider the change in the average of the system error. If the difference between the previous system error average and the new system error average is not large enough then the condition for retraining won't be met and the system will get stuck trying to find a pair of Sigmas that would result in error reduction. This won't be possible without memory because the moment the network is retrained the new data overwrites the information learned by the network. The use of memory allows the network to remember the stable and perturbed history of the system parameters.




Figure 11. s_1 and s_2 are changed while the system is running to improve performance and reduce error




Figure 12. Error is shown to be reduced from the starting point by updating the system parameters and the algorithm is successful in maintaining the error within bounds




Figure 13. Re-training to improve the deep neural network performance in order to adapt to the introduction of new system dynamics

7. Conclusions and Future Work

Lyapunov control was applied to Duffing- van der pol oscillator model that is experiencing chaotic behaviour. The study shows the effectiveness of deep learning combination with nonlinear Lyapunov control in finding the best parameters to maintain system stability. The study shows that deep learning with the proposed algorithm enables the user to effortlessly find the best parameters for the controller and the system initially and recalibrate the parameters if any disturbances or new dynamics are introduced to the system. In future work we would like to investigate the effect of changing the control strategy according to the type of instability detected. We speculate that depending on the type of perturbations and system dynamics, changing the control strategy might be more effective than changing the parameters only. Future considerations include a network that can determine which control strategy would have the highest impact on returning the system to stability and cause error reduction. The proposed approach can be aided with the use of generative adversarial networks (GAN).

REFERENCES

- [1] A. Harb, A. Zaher and M. Zohdy, "Nonlinear Recursive Chaos Control", Proc. of ACC, Anchorage, AL, USA, pp 2251-2254.
- [2] A. Harb, A. Zaher and, A. Al-Qaisia and M. Zohdy, "Estimation based Control of Chaotic Duffing Oscillators".
- [3] Kovacic, I.; Brennan, M.J., eds. (2011), The Duffing Equation: Nonlinear Oscillators and their Behavior, Wiley, pp. 392.
- [4] K. Tomita (1986): "Periodically forced nonlinear oscillators". In: Chaos, Ed. Arun V. Holden. Manchester University Press, ISBN 0719018110, pp. 213-214.
- [5] T. Alain, K. Timoleon "Chaos generalized synchronization of coupled Mathieu-Van der Pol and coupled Duffing-Van der Pol systems using fractional order derivative" Chaos, Solitons and Fractals Volume 98, May 2017, pp. 88-100, ELSevier.
- [6] S. S. Ge, T. H. Lee and G. Zhu, "Genetic algorithm tuning of Lyapunov-based controllers: an application to a single-link flexible robot system," in IEEE Transactions on Industrial Electronics, vol. 43, no. 5, pp. 567-574, Oct. 1996, doi: 10.1109/41.538614.
- [7] V. Giordano, D. Naso and B. Turchiano, "Combining Genetic Algorithms and Lyapunov-Based Adaptation for Online Design of Fuzzy Controllers," in IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), vol. 36, no. 5, pp. 1118-1127, Oct. 2006, doi: 10.1109/TSMCB.2006.873187.
- [8] R. Wai and C. Tu, "Development of Lyapunov-Based Genetic Algorithm Control for Linear Piezoelectric Ceramic Motor Drive," in IEEE Transactions on Industrial Electronics, vol. 54, no. 5, pp. 2566-2582, Oct. 2007, doi: 10.1109/TIE.2007.899934.
- [9] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue and A. D. Ames, "Episodic Learning with Control Lyapunov Functions for Uncertain Robotic Systems*," 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 2019, pp. 6878-6884, doi: 10.1109/IROS40897.2019.8967820.
- [10] St'ephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In Proceedings of the fourteenth international conference on artificial intelligence and statistics, pages 627-635, 2011.
- [11] St'ephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. No-regret reductions for imitation learning and structured prediction. CoRR, abs/1011.0686, 2010.
- [12] Aaron D Ames, Kevin Galloway, Koushil Sreenath, and Jessy W Grizzle. Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics. IEEE Transactions on Automatic Control, 59(4): 876-891, 2014.
- [13] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. IEEE Transactions on Automatic Control, 62(8): 3861-3876, 2017.
- [14] McCloskey, M. and Cohen, N. (1989) Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (ed.) The Psychology of Learning and Motivation, 24, 109-164.
- [15] Robins, A. (1995). Catastrophic Forgetting, rehearsal and pseudorehearsal. Connection Science, 7, 123-146.
- [16] Robins, A. (1996). Consolidation in Neural Networks and in the Sleeping Brain. Connection Science, 8(2), 259-276.
- [17] Ans, B., and Rousset, S. (2000). Neural networks with a self-refreshing memory: Knowledge transfer in sequential Learning tasks without catastrophic forgetting. Connection Science, 12, 1-19.
- [18] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting", J. Mach. Learn. Res., vol. 15, no. 1, pp. 1929-1958, 2014.

- [19] Alghassab, M., Mahmoud, A. and Zohdy, M.A. (2017) Nonlinear Control of Chaotic Forced Duffing and Van der Pol Oscillators. *International Journal of Modern Nonlinear Theory and Application*, 6, 26-37.
- [20] Thompson, J.M.T. and Stewart, H.B. (2002) *Nonlinear Dynamics and Chaos*. John Wiley & Sons, New York.
- [21] Zaher, A.A., Harb, A.M. and Zohdy, M.A. (2004) Recursive Backstepping Control of Chaotic Duffing Oscillators. *Proceedings of the 2004 American Control Conference*, 5, 4302-4306.
- [22] Harb, A., Zaher, A. and Zohdy, M. (2002) Nonlinear Recursive Chaos Control. *Proceedings of the 2002 American Control Conference*, 8-10 May 2002. <https://doi.org/10.1109/acc.2002.1023974>.
- [23] Zaher, A.A. (2011) Secure Communication Using Duffing Oscillators. *IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*, KualaLumpur, 563-568. <https://doi.org/10.1109/ICSIPA.2011.6144150>.
- [24] Li, J. and Shen, Y. (2009) The Study of Weak Signal Detection Using Duffing Oscillators Array. *IEEE Circuits and Systems International Conference on Testing and Diagnosis*, Chengdu, 28-29 April 2009, 1-4.
- [25] Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. In *Advances in Neural Information Processing Systems*, 2019.
- [26] Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning control lyapunov functions from counterexamples and demonstrations. *Autonomous Robots*, 43(2): 275–307, 2019.
- [27] Spencer M. Richards, Felix Berkenkamp, and Andreas Krause. The lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems. In *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 466–476, 29–31 Oct 2018.