

Model-free Intelligent Control for Antilock Braking Systems on Rough Terrain

Ricardo de Abreu^{1,*}, Theunis R. Botha^{1,**}, and Herman A. Hamersma^{1,***}

¹Department of Mechanical and Aeronautical Engineering, University of Pretoria, South Africa

Abstract. Advancements have been made in the field of vehicle dynamics, improving the handling and safety of the vehicle through control systems such as the Antilock Braking System (ABS). An ABS enhances the braking performance and steerability of a vehicle under severe braking conditions by preventing wheel lockup. However, its performance degrades on rough terrain resulting in an increased wheel lockup and stopping distance compared to without. This is largely as a result of noisy measurements, and un-modelled dynamics that occur as a result of the vertical and torsional excitation experienced over rough terrain. Therefore, it is proposed that a model-free intelligent technique, which may adapt to these dynamics, be used to overcome this problem. The Double Deep Q-learning (DDQN) technique in conjunction with a Temporal Convolutional Network (TCN) is proposed as the intelligent control algorithm, and straight line braking simulations are performed using a single tyre model, with tyre characteristics approximated by the LuGre tyre model. The rough terrain is modelled after the measured Belgian paving with the normal forces at the tyre contact patch approximated using FTire in ADAMS. Comparisons are drawn against the Bosch algorithm, and results show that the intelligent control approach achieves lateral stability by preventing wheel lockup whilst braking over rough terrain, without deteriorating the stopping distance.

1 Background

Significant advances towards vehicle safety have been made with one such system being the development of the Antilock Braking System (ABS). Originally introduced for trains and later developed for passenger vehicles and light trucks, Kelsey-Hayes initiated the "automatic" braking system exploratory development program, which concluded that the braking system of a vehicle should not only reduce the vehicle's stopping distance but also prevent the loss of vehicle control [1]. ABS is a feedback control system that attempts to maintain controlled braking under all operating conditions. This is accomplished by controlling the slip at each wheel to ensure the optimum forces within the limits of the tyre-terrain interaction are achieved. Reports released by the National Highway Traffic Safety Administration (NHTSA) highlight notable decreases in multi-vehicle crashes and fatal pedestrian crashes with the inclusion of ABS [2], however, the performance of an ABS on rough terrain conditions is found to be suboptimal, resulting in wheel lockup, bad stopping distances and poor

*e-mail: ricardo.deabreu@tuks.co.za

**e-mail: theunis.botha@up.ac.za

***e-mail: herman.hamersma@up.ac.za

lateral control [3]. This has inspired the emerging field of study exploring the influence of several phenomena, such as tyre force generation, that impact the effectiveness of ABS on rough terrain [3].

Braking over rough terrain presents a tough challenge to the performance of an ABS by introducing effects that interfere with its normal operation, leading to subpar performances. These effects include noisy measurements, inconsistent terrain input excitations, vehicle body motions, and tyre oscillations. As a tyre travels along a rough terrain, the undulations in the terrain surface not only change the effective rolling radius of the tyre [3], but also cause changes in the vertical and torsional dynamics of the wheel that are un-modelled and often result in loss of contact between the tyre and terrain. The type of control approach used for controlling an ABS plays a crucial role in its performance, and forms the main focus of this study.

ABS control is a nonlinear problem and several control strategies exist, making use of modern control techniques such as sliding mode control, gain scheduling, and fuzzy logic control [4]. Majority of ABS control approaches rely on a set of rules (a defined model) and are robust on specific smooth terrains (dry, icy, or wet), however this does not hold over rough terrain where mismatches tend to occur between the control design model and the real process [5], and as a result, the un-modelled dynamics and parametric uncertainties lead to underwhelming performances of the ABS. This leads to an interest in model-free control designs. Model-free control techniques, commonly known as data-driven techniques, make use of data (usually input-output) in their control design. Reinforcement Learning (RL), also known as Adaptive Dynamic Programming (ADP), is representative of a model-free technique that exploits the process structure used when designing the control instead of using an identified process model. The field of model-free control techniques and RL in particular remains a relatively new concept, however application of this technique to non linear problems such as the control of ABS have been explored with promising results. In [6] a supervised actor-critic approach is proposed for adaptive cruise control, which approximates the optimal control policy for adaptive cruise control and can be adapted to the control of ABS. In [7] the use of ADP for the optimal control of an ABS over smooth terrain is proposed. This approach relies on penalizing the braking distance of the vehicle, and outperforms many existing solutions in literature. In [5] the design and implementation of a model-free slip control of ABS through Q-learning is implemented over smooth terrain, with the performance of the algorithm proving its capability over a wide operating range.

Majority of ABS controllers are simulated on smooth terrain with focus on improving their models stopping distance. However, excessive wheel lockup occurs as a result of the additional excitation in off-road conditions, this not only results in loss of directional control, but also one of the main focuses of ABS which is to prevent wheel-lockup. Thus developing a model-free ABS controller which prevents wheel-lockup on rough terrain without deteriorating the stopping distance forms the main objective of this study.

2 Implementation

2.1 Overview

The method of implementation is broken up into three sections: 1) modelling of an existing ABS algorithm for comparison, 2) modelling of the brake system dynamics and 3) modelling of the proposed ABS control method.

2.2 Bosch Algorithm

The best-documented ABS algorithm was published by Robert Bosch GmbH in 1999 [8], and is known as the Bosch algorithm. The Bosch algorithm is a bang bang type, rule-based control strategy that takes the wheel angular acceleration and longitudinal slip as inputs, and makes use of three different pressure control modes to control the brake pressure at each wheel: *pump* (increase pressure), *dump* (reduce pressure), and *hold* (maintain pressure). Dictated by two upper and one lower threshold for the angular acceleration and an upper longitudinal slip threshold, the Bosch algorithm combines these three pressure control modes in various configurations resulting in a repetitive ABS control cycle consisting of eight phases. Studies [10, 11] have shown that by fine tuning the parameters of the Bosch algorithm, reasonable results can be obtained when braking over rough terrain, thus it is a suitable algorithm to be used as a baseline.

2.3 Brake System Modelling

A mathematical model of the brake system dynamics is discussed and this includes the tyre model. This model provides a link between the tyre braking dynamics and the brake hydraulic dynamics, and although simple, retains the essential characteristics of the actual system.

2.3.1 Brake Model

In order to reduce solution times and determine all the initial aspects of the controller, a simplified model is used in the initial stages of this research, after which more complexity is added. A single tyre model is considered and demonstrated in Figure 1, where the vehicle velocity acts in the positive direction.

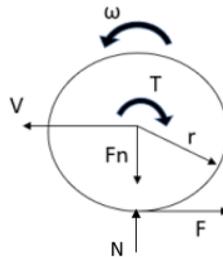


Figure 1: A single tyre free-body diagram with lumped friction

Applying first principles to the free-body diagram, with the sum of the forces acting positive in the left direction and the sum of moments acting positive in the anticlockwise direction, a system of the form is obtained:

$$m\dot{v} = -F, \tag{1}$$

$$I\dot{\omega} = rF - T_b, \tag{2}$$

where m is a 1/4 of the vehicle mass, I the moment of inertia around the centre of the wheel and r is the radius of the wheel. The vehicle velocity, angular velocity, brake torque and tyre-road friction force are represented by v , ω , T and F respectively. For the sake of simplicity, the vertical force is assumed to act in line with the centre of the wheel neglecting effects of rubber

hysteresis, whilst effects such as the torsional deflection of the sidewall and rolling resistance are not considered. This simplified approach also neglects rolling distance or pneumatic trail as this component is small compared to the brake torque. The tyre-road friction force is described by the lumped LuGre model [9], defined in Equations (3)-(5).

$$\dot{z} = v_r - \frac{\sigma_0 |v_r|}{g(v_r)} z. \tag{3}$$

$$F = (\sigma_0 z + \sigma_1 \dot{z} + \sigma_2 v_r) F_n, \tag{4}$$

where

$$g(v_r) = \mu_c + (\mu_s - \mu_c) e^{-|v_r/v_s|^\alpha}. \tag{5}$$

The brake torque is the force applied at the brake disc to stop its motion, and is defined as a linear relationship with the brake pressure in Equation (6) [10].

$$T_b = 271 \times P, \tag{6}$$

$$\dot{P} = \frac{P_{set} - P}{\tau}. \tag{7}$$

The pressure as a function of time \dot{P} is equal to the difference between the backpressure (P_{set}) and the current pressure in the valve, divided by a time delay constant τ . P_{set} and τ are set according to the specifications used in [10]. Therefore considering the four nonlinear, differential equations [\dot{P} , \dot{z} , \dot{v} , $\dot{\omega}$], the relationship of the system dynamics and how they interact with each other is established.

2.4 Model of Controller

The purpose of the controller model is to actuate and perform the operation of the ABS, through regulating the brake pressure of the vehicle. Literature reveals several different processes can be used in the control of the ABS [4], however results from [5, 7] indicate the possibilities that exist in model-free control methods. In an attempt to overcome the unmodelled dynamics and parametric uncertainties that lead to underwhelming performances of the ABS over a rough terrain, an intelligent model-free control is proposed.

RL is a method that learns through the interaction between an autonomous active decision-making agent and its dynamic environment, with the objective of the agent to maximize a reward signal despite uncertainty about its environment [12]. Model-free RL methods make use of empirical evidence from past experiences hereby enabling the agent to learn from its environment through trial and error, despite no prior knowledge of the environment. Basic model-free approaches update according to Temporal Difference (TD) learning rule, and can be separated into two types: (1) Q-Learning and (2) Policy Optimization [12].

Q-Learning is a action-value method that makes use of Q-values $Q(s,a)$ (action-value functions) that choose an optimal policy based on the values of the actions of the agent. This approach is almost exclusively achieved through off-policy methods, meaning that each update is not dependent on the agent, and data obtained at any stage during training can be used [12]. Q-learning is easy to implement, has a high sample efficiency and examples include the Deep Q-learning (DQN) and Double Deep Q-learning (DDQN) algorithms.

Policy Optimization methods learn a parameterized policy that can select actions without consulting a value function [12]. The value function may still be used to learn the policy parameter, however it is not required to select the action. These methods are referred to as policy gradient methods. Actor-critic methods are methods that learn approximations to both policy and value functions, where the policy structure is known as the actor because

it is used to select actions, and the estimated value function is known as the critic which criticizes the actions made by the actor. Examples of the policy gradient methods include the Asynchronous Advantage Actor-Critic (A3C) and Advantage Actor-Critic (A2C), as well as the Soft Actor-Critic (SAC) algorithm.

The Q-learning and policy gradient methods present advantages and disadvantages for each separate task/problem, thus it is proposed that the different Q-learning (DQN and DDQN) and policy gradient methods (A2C, A3C and SAC) be compared as alternative ABS control algorithms, in order to find the best suitable method for this problem. Due to the availability of baseline implementations for RL algorithms, the GitHub repository of Deep-Reinforcement-Learning-Algorithms-with-PyTorch as provided by [13] is proposed as the main source of implementation and comparison of the above RL algorithms.

The explored RL algorithms make use of deep Neural Networks (NN) as their function approximators, however real-life applications generally involve time series data, as is the case with the control of ABS, hereby adding the complexity of sequence dependency amongst the input variables. Knowing this, it might be beneficial to store this information over a long period of time and learn from it, thus introducing an additional aspect of predictive modelling to the control problem to aid in its performance. Two types of NNs which have shown to be good at handling sequence dependency as well as the modelling of time series data are Long Short Term Memory (LSTM) [14] and Temporal Convolutional Networks (TCN) [15] architectures. It has been proven that the TCN not only exhibits longer memory but consistently outperforms the LSTM on a vast range of tasks [15] and is thus explored further.

In order to model the controller, a state space and action space are to be defined. Several combinations of states can be used as input, such as $[P, v, w, z, \lambda, F_n]$, however these inputs need to be able to be easily accessible in a physical system, and not limit the performance of the model controller by being redundant. Therefore taking this into account, internal investigations show that a simple state space consisting of $[\frac{v}{r}, w]$ yield the best results, with a discrete action setting of [pump, dump, hold].

2.4.1 Reward Function

The reward function determines the behaviour of the brake control, thus in order to ensure the reliability of the brake control, it is crucial to use a properly defined reward function. Different aspects to determining the ideal reward/penalty function exist, such as regulating the control slip. This approach involves regulating the slip by means of predetermined slip thresholds, and is the most commonly used penalty system for ABS systems, however it needs to be adapted according to the type of terrain and restricts the model - no longer representing a model-free approach. An alternative approach would be to monitor and prevent wheel lockup. This approach focuses solely on preventing wheel lockup and penalizes the model when the slip is 100% or the angular velocity is equal to 0. Lastly, the braking distance of the vehicle can be monitored and used as a penalty system. This approach is not ideal as it neglects the importance of lateral control of the vehicle and provides sparse rewards.

When considering the reward function for the ABS model, there is conflict between two intuitive objectives of the ABS: 1) the vehicle is encouraged to brake hard to achieve the shortest possible stopping distance and 2) the wheels are to not slip, which is common during hard braking. If it is unbalanced, the agent becomes either too conservative or reckless, leading to slow braking or wheels locked at 100% respectively. Taking this into consideration, the following reward function is proposed:

$$R = P - 10 - j(\lambda), \quad j(\lambda) = \begin{cases} 10\lambda & \lambda > 20\% \\ 0 & \lambda < 20\% \end{cases} \quad (8)$$

The function serves as a penalty function which always returns a negative reward. The agent tries to maximize the rewards which here it leads to minimizing the penalties received. The penalty function also penalizes the duration of the braking process, the longer the vehicle brakes - the more penalties received. This encourages the agent to additionally slow down the vehicle as quickly as possible. Knowing that the maximum pressure pumped from the brake value is 10 Mpa the negative rewards is achieved by vertically shifting the rewards to below 0 with the constant value -10. To further encourage the model to stop as quickly as possible, the braking force can be taken into consideration via two possible options: linear acceleration or pressure. By increasing either of these, the braking force is increased and the vehicle stops quicker. When travelling over rough terrain, the linear acceleration is sensitive to big changes in the normal force acting at the tyre contact patch, thus it is proposed that the pressure P be used. This ensures that the first intuitive objective of the ABS is met, to meet the second objective, the function $j(\lambda)$ is introduced. This function is dependant on the wheel slip and penalizes the model as the wheels begin to slip. Maximum braking forces tend to occur between the region of 10-30% slip [4], however the peak braking force is closer to the 10% region with slippage occurring almost instantly past 30%, thus slip greater than 20% is penalised. A constant weight of 10 is multiplied by the wheel slip to discourage wheel lockup.

3 Results and Discussion

In order to simulate braking over rough terrain, a varying normal force at the tyre contact patch is required. Previous studies have been conducted at the Gerotek Test Facilities' Belgian paving [3, 10, 11], and this is used as vertical excitation. Due to the difficulty in modelling a vertical tyre that can capture the tyre dynamics over rough terrain, an FFT of the vertical force over the Belgian paving during braking was obtained by simulating a validated FTire model in ADAMS at varying speeds. This model is used to artificially generate a vertical force during simulation, and captures the speed dependency effect of the model, allowing for a random excitation force to be created and better generalization. To further encourage generalization by the model several steps are taken:

1. The starting speeds of the vehicle are randomised between 15 and 19m/s
2. The moment of Inertia I , μ_s and μ_c of the tyre model and mass of the body are randomised for each episode
3. Different normal force profiles are randomly sampled every 5 episodes

It should be noted that all the simulation results are based on the assumption of straight-line braking in the brake system model, and neglect the effects of the transportation delay caused by the brake hydraulic systems and brake pad travel as in [16]. The simulation is terminated below a speed of $v = 2\text{m/s}$ as ABS algorithm often stop working at very low speeds [3].

Figure 2 plots the mean rolling episode scores against the episode number for the five different RL algorithms using a deep NN as the function approximator. An average score -680 over 1000 episodes was achieved for the Bosch algorithm and is used as a benchmark for the intelligent algorithms.

It can be seen that none of the algorithms are capable of matching or exceeding the performance of the Bosch algorithm, however it is important to note that this is as a result of using the simplified single tyre model. The dynamics that cause the Bosch algorithm to fail namely: unmodeled wheel torsional dynamics as well as noisy wheel speed signal leading to problems in slip calculation and angular acceleration are non existent in these simulations i.e. the Bosch is performing non real world conditions. It is recommended that a more complex model be used to expose the Bosch to these real life conditions, however despite this, it is

good practice to compare the proposed intelligent models to the Bosch algorithm. It can be seen that the DDQN, A3C and A2C algorithms demonstrate successful learning with an increase in the average number of rewards per each episode, and are used in further simulations.

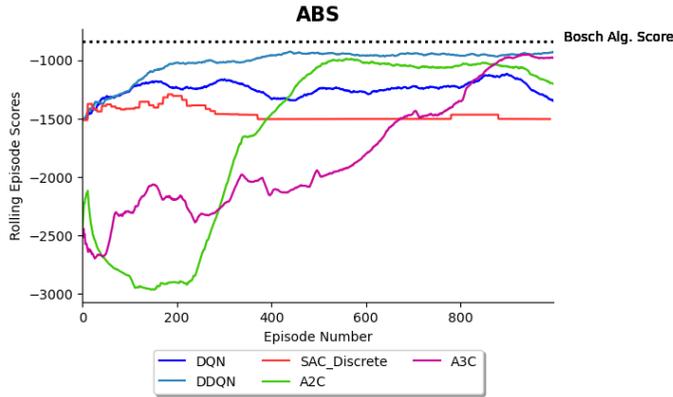


Figure 2: Comparison of rolling episode scores over 1000 episodes for each RL algorithm on rough terrain

As an alternative to the NN function approximator, the use of the intelligent models in conjunction with a TCN is investigated. By storing 100 previous step intervals, the benefit of using a history of data for time series data is seen in Figure 3, where despite the A3C and A2C models unsuccessfully learning, the DDQN algorithm which was able to better the Bosch algorithm over time.

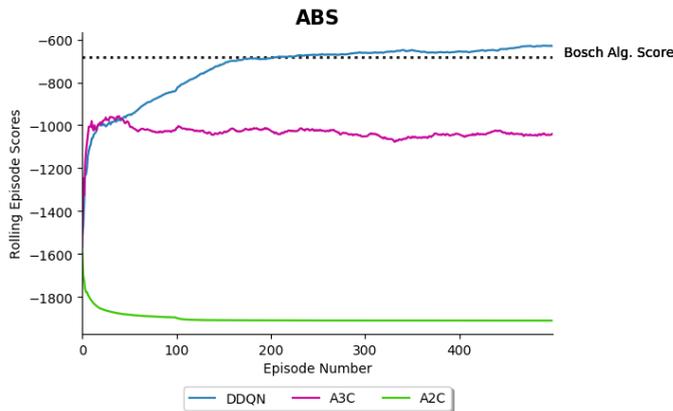


Figure 3: A comparison of the DDQN, A2C and A3C with the use of TCN as a function approximator, over 500 episodes on rough terrain

The DDQN algorithm has shown to be the most consistent and best performing of all the algorithms and is used for post processing. It is important to further validate the performance of this algorithm to ensure it is performing as expected. 10 episodes of the DDQN model are performed against 10 episodes of the Bosch algorithm at the same starting speed of 17m/s. Different performance measures that are conducted based on these runs include: the average stopping distance, the average longitudinal slip and its percentage time spent in certain slip

ranges (this gives a good indication of the how much lockup occurs) as well as the vehicle and wheel speed profile, and the pressure, slip and reward profiles for a single episode. Table 1 summarises the results of the average stopping distances, as well as the different longitudinal slip performance measures of the Bosch, DDQN with and without the TCN (labelled DDQN-NN and DDQN-TCN).

Table 1: Summary of performance measures of the Bosch, DDQN-NN and DDQN-TCN

Stopping Distance (m)	BOSCH	DDQN-NN	DDQN-TCN
Mean	24.37	24.5	25.21
Std. dev	1.19	1.21	1.98
Best	23.45	23.74	23.17
Longitudinal slip (%)			
Mean	9.53	3.84	5.56
Std. dev	19.64	10.87	15.97
$\lambda < 10\%$	76.55	89.96	85.98
$10\% < \lambda < 20\%$	6.90	3.95	6.45
$\lambda > 20\%$	16.55	6.08	7.57

Analysing this summary, it can be seen that the DDQN models (with and without TCN) have a lower average slip percentage and produce a lower percentage of slip values greater than 20%, highlighting their ability to prevent wheel lockup on rough terrain. Along with ensuring the stability and lateral control of the model, the stopping distances are within 1m of the stopping distance achieved by the Bosch algorithm, and thus acceptable. Comparisons between the vehicle and wheel speed for the DDQN-TCN algorithm against the Bosch algorithm are presented in Figure 4.

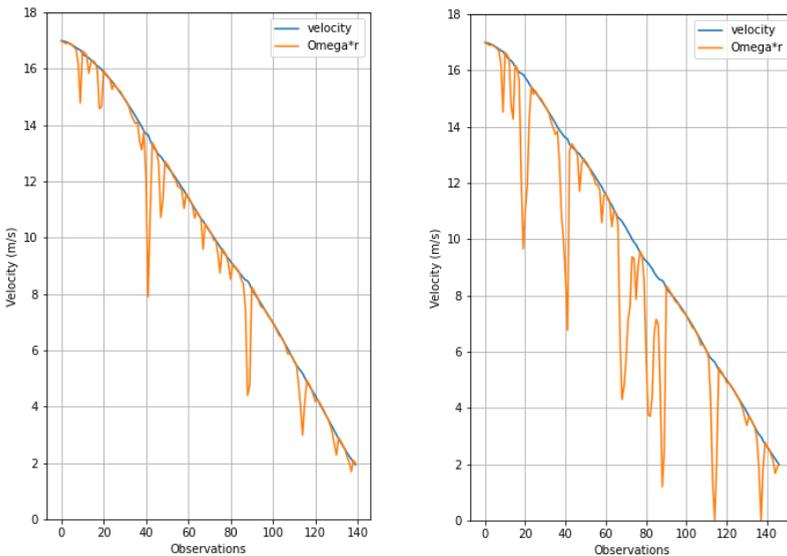


Figure 4: A comparison of the vehicle and wheel speed of the DDQN-TCN Algorithm (left) vs. the Bosch Algorithm (right)

It can be seen that the DDQN-TCN algorithm cycles the wheel speed better than the Bosch algorithm, within the same amount of step intervals (observations), hereby achieving small slippage as confirmed in Figure 5, where a maximum slip of 50% is achieved compared to 100% for the Bosch algorithm. This confirms the results from Table 1, where low percent of slippage is achieved by the DDQN-TCN model, resulting in better lateral control, as well preventing wheel lockup over rough terrain.

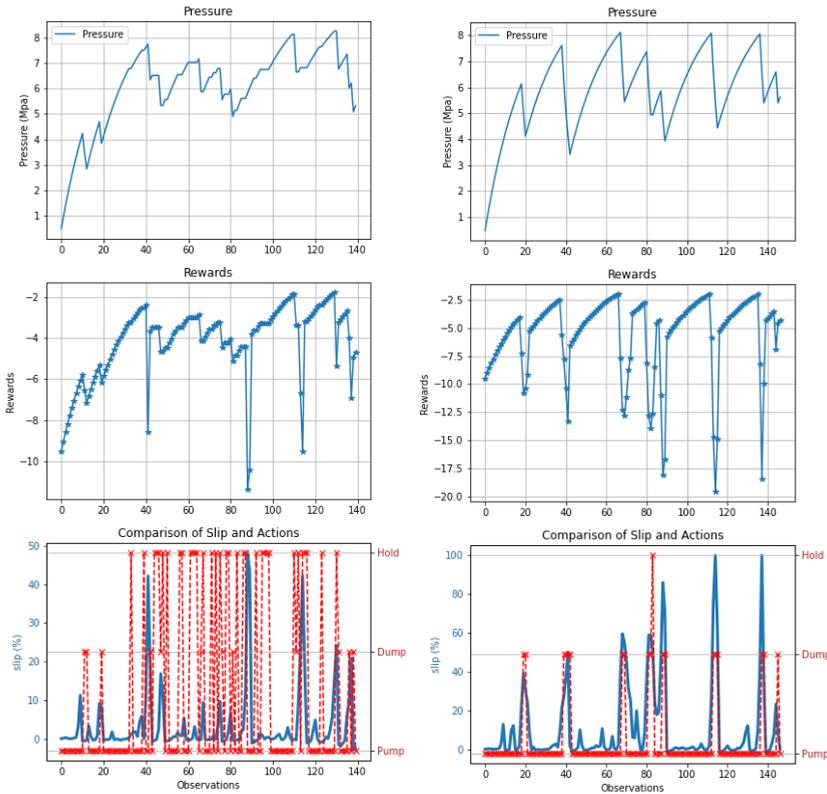


Figure 5: A comparison of the pressure, rewards and the slip and actions chosen of the DDQN-TCN Algorithm (left) vs. the Bosch Algorithm (right)

The following results prove that the proposed intelligent model-free ABS control algorithm is able to prevent wheel lockup over rough terrain, without deteriorating the stopping distance of the system. Future work includes increasing the complexity of the model from a single tyre to a validated full vehicle simulation model and training the intelligent algorithms on this vehicle model using FTire simulations in ADAMS. Additional tests to determine the success of the model are also to be performed according to the ABS performance metric proposed in [3]. These results should provide sufficient evidence on the application of a model-free approach as an ABS control approach.

4 Conclusion

Straight line braking simulations on rough terrain were performed using a single tyre system. The control of the ABS was modelled by multiple model-free RL techniques, and compared to the performance of a baseline Bosch algorithm. The DDQN algorithm with the use of TCN (labelled DDQN-TCN) demonstrated the most consistency of all the different algorithms considered, and was able to exceed the average rolling episode scores benchmark set by the Bosch algorithm. Through a series of ten straight line braking simulations, the DDQN-TCN model was able to achieve a lower average slip percentage as well as a lower percentage of slip values greater than 20%, highlighting its ability to prevent wheel lockup on rough terrain, whilst maintaining an average stopping distance within 1m of the Bosch algorithm. A more comprehensive model, along with the performance metrics defined in [3] are recommended for future work, in order to fully validate the performance of the model-free algorithm.

Bibliography

- [1] SAE, Antilock Brake System Review, SAE International (2014)
- [2] Garrott, W. R. & Mazzae, E. N., An overview of the national highway traffic safety administration's light vehicle antilock brake systems research program, Technical report SAE Technical Paper (1999)
- [3] Hamersma, H. A, ABS braking on rough terrain, PhD thesis, University of Pretoria (2017)
- [4] Aly, A. A., Zeidan, E.-S., Hamed, A., Salem, F. et al., 'An antilock-braking systems (abs) control: A technical review', *Intelligent Control and Automation* **2**, 186–195 (2011)
- [5] Radac, M.-B. & Precup, R.-E., 'Data-driven model-free slip control of anti-lock braking systems using reinforcement q-learning', *Neurocomputing* **275**, 317–329 (2018)
- [6] Zhao, D., Wang, B. & Liu, D., 'A supervised actor-critic approach for adaptive cruise control', *Soft Computing* **17**(11), 2089–2099, (2013)
- [7] Sardarmehni, T. & Heydari, A, Optimal switching in anti-lock brake systems of ground vehicles based on approximate dynamic programming, in 'Dynamic Systems and Control Conference', American Society of Mechanical Engineers **57267**, p. V003T50A010, (2015)
- [8] Bosch Gmbh, R. & Girling, P, Bosch Driving-safety systems, Society Of Automotive Engineers (1999)
- [9] Canudas-de Wit, C., Tsiotras, P., Velenis, E., Basset, M. & Gissinger, G, 'Dynamic friction models for road/tire longitudinal interaction', *Vehicle System Dynamics* **39**(3), 189–226, (2003)
- [10] Penny, W. C. W, 'Anti-lock braking system performance on rough terrain', Master's thesis, University of Pretoria, (2016)
- [11] van der Merwe, N. A., Els, P. S. & Žuraulis, V, 'Abs braking on rough terrain', *Journal of Terramechanics* **80**, 49–57, (2018)
- [12] Sutton, R. S. & Barto, A. G, Reinforcement learning: An introduction, MIT press, (2018)
- [13] Christodoulou, P. Deep-reinforcement-learning-algorithms-with-pytorch, (2017). URL: Available at: <https://github.com/p-christ/Deep-Reinforcement-Learning-Algorithms-with-PyTorch>
- [14] Hochreiter, S. & Schmidhuber, J, 'Long short-term memory', *Neural computation* **9**(8), 1735–1780, (1997)

- [15] Bai, S., Kolter, J. Z. & Koltun, V, 'An empirical evaluation of generic convolutional and recurrent networks for sequence modeling', arXiv:1803.01271 (2018)
- [16] Keshmiri, R. & Shahri, A. M, 'Intelligent abs fuzzy controller for diverse road surfaces', World Academy of Science, Engineering and Technology **2(2)**, 62–67, (2007)
- [17] Stallmann, M. J. & Els, P. S, 'Parameterization and modelling of large off-road tyres for ride analyses: Part 2–parameterization and validation of tyre models', Journal of Terramechanics **55**, 85–94, (2014)