Safe Control Allocation of Articulated Heavy Vehicles Using Machine Learning

(article starts on next page)

Safe Control Allocation of Articulated Heavy Vehicles Using Machine Learning

Sander van Dam[1,2], Lukas Wisell[1,2], Kartik Shingade[1,2], Mikael Kieu[1], Umur Erdinc[1,2](✉), Maliheh Sadeghi Kati[1], Esteban Gelso[1], and Dhasarathy Parthasarathy[1]

[1] Volvo Group Trucks Technology, Gothenburg, Sweden
umur.erdinc@volvo.com

[2] Chalmers University of Technology, Gothenburg, Sweden

Abstract. As articulated heavy vehicles are over-actuated, achieving a safe control allocation is crucial to ensure stability. This study introduces a machine learning model developed to identify unsafe behaviours and modes, such as jack-knifing and trailer swing, enabling the control scheme to prioritize stability. High-fidelity simulations, focusing on high-risk scenarios, generate data for training the machine learning model. This model is integrated into the control scheme to predict safe braking allocations and prevent unsafe vehicle modes during real-time driving scenarios. Initial tests showed promising results regarding prediction accuracy and a safety margin that can be implemented to further ensure that safe vehicle motion is achieved.

Keywords: Control allocation · articulated heavy vehicles · machine learning · yaw instability

1 Introduction

Articulated heavy vehicles (AHVs) have proven to be excellent candidates for road transportation due to their significant positive impact on environmental and economic efficiency [1,2]. However, concerns remain regarding their stability, including risks like jackknifing and trailer swing. The electrification of trucks introduced the concept of distributed propulsion across different vehicle units. To maximize power regeneration, braking with the propelled vehicle units and axles is preferred, which might cause instabilities such as jackknifing [3]. Distributed propulsion in AHVs adds another layer of complexity, impacting these vehicles which are already dynamically complex due to over-actuation and their articulation joints. Therefore, achieving a safe control allocation (CA) is crucial to ensure stability while maximizing the regenerated energy [4].

Using machine learning (ML) in the automotive industry is growing as a strong technique, bringing several benefits such as improved safety measures, enhanced vehicle efficiency, and better performance (e.g., [5–7]). To the best of our knowledge, there has been no work done to use ML in the control allocation

G. Mastinu et al. (Eds.): AVEC 2024, LNME, pp. 1–7, 2024.
https://doi.org/10.1007/978-3-031-70392-8_1

of forces in heavy vehicle applications. In this study, ML is applied to heavy vehicle applications to identify unsafe behaviours for a tractor-semitrailer vehicle combination, specifically yaw instabilities. This enables the control scheme to maintain the vehicle within safe operating conditions by requesting safe braking allocation, ensuring no unsafe modes will appear.

2 Machine Learning Model Development

To predict safe allocations, a high-fidelity vehicle model, referred to as Volvo Transport Model (VTM) [8], was used to perform simulations, collecting data on factors contributing to the unsafe behaviour of the tractor-semitrailer vehicle combination. Multiple simulations were performed focusing particularly on high-risk environmental and operational conditions such as low friction and high lateral acceleration on a circular track for different radii.

Approximately 25,000 uniformly random distributed simulations were performed for a range of brake-in-turn manoeuvres, varying parameters such as the initial velocity at which braking begins, friction, radius of the turn, load on the trailer, and braking allocations. The simulation was initialised during steady-state cornering of the vehicle, and data samples were collected until the vehicle either reached a standstill or an unsafe mode occurred. To differentiate between safe and unsafe modes, criteria were checked during the entire simulation on important vehicle states to monitor the state of the vehicle and classify the simulation as safe or unsafe. The criteria used were based on the value of the side-slip angle rate of the vehicle combination, as discussed in [9]. This metric captures whether the truck or trailer exhibits a fast change in yaw angle within a short amount of time, indicating unsafe behaviour in the form of either jack-knifing or trailer swing. The chosen thresholds to limit the side-slip angle rates of the tractor and trailer were set at 6 deg/s. This threshold value was validated with VTM and found to be effective in identifying yaw instabilities.

From these simulations, 20% were used for test data and 80% for training data. Each simulation yielded about 25 data points by sampling the vehicle state throughout the manoeuvre, resulting in around 625,000 data points overall. This extensive dataset was obtained with a total simulation time of about 275 h.

During each simulation, data was collected about the most informative states of the vehicle from a vehicle safety perspective. The selected states were used as features for the ML model, with labels classifying the state as feasible or infeasible. Using this labelled dataset, an ML model could be trained offline to be used as an online prediction tool within the control scheme. The chosen state vector as the features of the ML model included road friction, steering wheel angle, brake force of the tractor, brake force of the trailer, longitudinal velocity, side slip angle of each unit, side slip angle rate of each unit, articulation angle, articulation angle rate and axle loads of each axle, i.e. a total of 17 features. Considering the high dimensional inputs of the problem with these many features, it is too intricate to determine the safe allocation with simple boundaries and any analytical solution. Therefore, ML is a good tool to address this complexity.

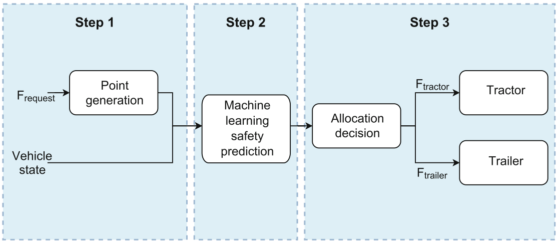

Fig. 1. Control scheme overview

Several ML models were evaluated on the dataset, including decision trees, random forest and neural networks. Among these, neural networks proved to be the most accurate with a training and validation accuracy of around 97%. The models were trained using the MATLAB classification learner app with default values for the hyperparameters not explicitly mentioned in this paper. The chosen model was a neural network consisting of three layers, each with 12 nodes. A higher cost was assigned to false positives since it is crucial for a safety critical system to accurately predict all unsafe points, even if it allows some error in predicting safe points.

3 Machine Learning Based Control Allocation Strategy

The control system is divided into a three-part process, as illustrated in Fig. 1. It begins by reading the vehicle state and brake request from the vehicle. This information is then used to calculate a set of potential braking allocations. The braking request, $F_{request}$, represents the total force request and is divided between the tractor and the trailer as $F_{request} = F_{tractor} + F_{trailer}$. From this set, 100 allocations are uniformly sampled that fulfil the total braking request, as illustrated by the points in Fig. 2a. All forces are normalised and therefore varying between 0 and -1, where -1 corresponds to the maximum possible braking that each unit can achieve given the current friction, μ. This normalization is expressed as $F^* = F/(F_z \cdot \mu)$ for each unit, where F_z is the unit normal force. These points together with the vehicle state are then inputted to the ML model, which predicts the safety of each allocation, distinguishing between safe (green) and unsafe (red) points as in Fig. 2a. The predictions are performed assuming the driver input to the vehicle is kept constant until standstill. Lastly, a decision is made to select one of these safe allocations, as indicated by an orange circle in the last plot of Fig. 2a.

The decision of which allocation to choose provides considerable flexibility. As the ML model assigns labels to each point, this part becomes modular, allowing for an easily interchangeable strategy without affecting the overall design. This paper will not present any optimal solution to this, instead showing multiple viable options are possible, and that the safety prediction is reliable. For a given state, the predictions are similar to what can be seen in Fig. 2a (steps 2 and

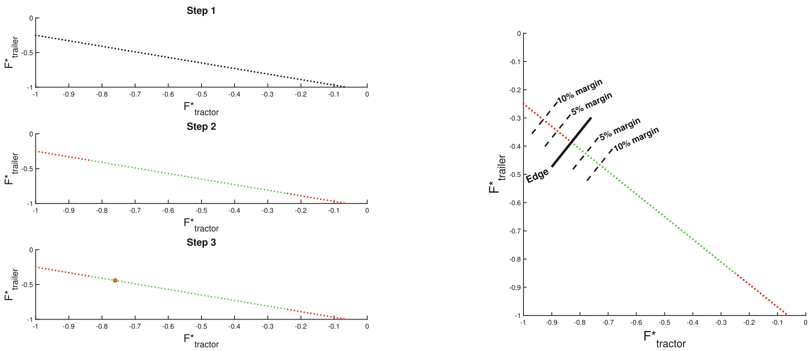

(a) Steps for selection of safe CA

(b) Safety margin selection in step 3

Fig. 2. Illustrations of the control allocation and the implemented safety margins.

3), demonstrating that extreme choices, braking only with the tractor or trailer, are the most unsafe, while intermediate options are safe. This is consistent with findings from previous research [3].

The proposed strategy was tested by picking an allocation that brakes as hard as predicatively possible with the tractor and with some safety margin. This corresponds to, for example, the orange circled point in "step 3" of Fig. 2a. However, when implementing this strategy, the edge can change rapidly due to the changing state of the vehicle. This is unwanted and could lead to unsafe behaviour of the vehicle combination. To accommodate this issue, a low-pass filter was added to the output of the CA. In Fig. 3a, the CA for the tractor is shown in one of these cases. The important observation here is that the CA varies rapidly and discretely. Figure 3a illustrates how the CA appears without a filter, while Fig. 3b shows the same scenario with the filter activated. It can be seen that the filter improves the allocation by eliminating these very fast changes in the CA.

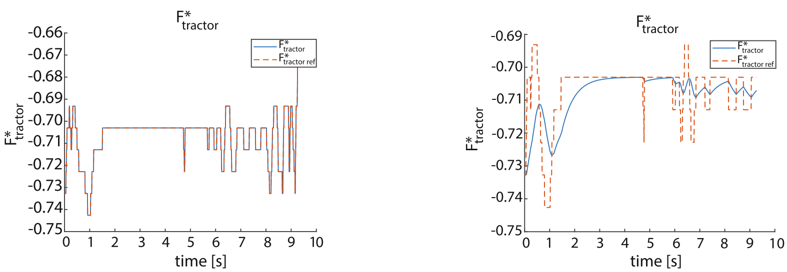

(a) Without low-pass filtering.

(b) with low-pass filtering.

Fig. 3. Illustrations of force allocation for the tractor in the scenario with parameters: $v_x = 43\ kph$, $\mu = 0.3$, $R = 72\ m$, $F^*_{request} = -0.6$.

4 Validation of Proposed Control Allocation Strategy

Simulations were conducted to validate these predictions and ensure the controller functions as intended. During these simulations, a safety margin of 5 and 10% of the total amount of sampled allocations was applied around the prediction border, with simulations running on points from both safe and unsafe margins (see Fig. 2b). The safety margin addresses uncertainties and inaccuracies in vehicle parameters, modelling simplifications, etc. Ideally, selecting a safe point should result in a safe simulation, and choosing an allocation predicted as unsafe should lead to an unsafe outcome. Doing this verifies that the estimation of the prediction border is accurate. This was done for a number of test scenarios that were divided into edge and normal cases. Edge cases are defined as scenarios where the vehicle is in an initial state close to where the vehicle can no longer operate without showing unsafe behaviour. Normal cases are further from an unsafe initial state than edge cases which still contain control inputs that can result in unsafe modes.

The system was tested by running 26 cases, 13 normal and 13 edge cases, of the above mentioned cases. Table 1 illustrates the success rate (whether the prediction matches the simulated outcome) for different margins. Throughout the simulations, it was assumed that the tractor is electric with a conventional trailer, prioritizing tractor braking for enhancing energy regeneration. Besides testing with constant steering and perfect knowledge about the state (as during training), the system was also tested with disturbances. An active path-follower, emulating a driver trying to stay on the road (marked by "steering" in Table 1), was introduced. The system was also tested under "wrong friction" conditions, where friction measurement was intentionally increased to simulate a misestimation. During these tests, the steering input was kept constant. Additionally, all other parameters remained the same as the initial test case.

Table 1. Prediction accuracies for different cases (%)

Case	Unsafe 10%	Unsafe 5%	Safe 5%	Safe 10%	Safe 15%	Safe 20%
Normal	100	69.23	92.31	100	100	100
Edge	84.62	61.54	61.54	92.31	100	100
Steering normal	76.92	46.15	84.62	100	100	100
Steering edge	92.31	92.31	92.31	100	100	100
Wrong friction normal	100	100	7.69	53.85	84.62	100
Wrong friction edge	100	100	7.69	15.38	69.23	76.92

The controller accurately predicted safety and responded well, with high accuracy rates, particularly for the normal cases and a slight decrease for edge cases. A margin of 10% from the predicted edge gave results close to 100%, while a 5% margin resulted in accuracy ranging from 60–95%, depending on the cases. The controller and ML model performed well when a path controller was active during braking, despite all training data being conducted under constant steering conditions.

5 Conclusion

This paper introduces a novel method of predicting the yaw instability in a tractor-semitrailer vehicle combination using machine learning. Additionally, it allocates the driver's brake request to individual vehicle units to avoid such yaw instability. To enhance the overall performance of the controller including ML, different ML models were tested and strategies were developed for allocation. The neural network classification model was found to outperform other models, even when integrated with the controller during random test manoeuvres in real time. The controller proved to be accurate at predicting the safety and responding accordingly. Accuracies were high, especially for the normal cases, but dropped somewhat for edge cases. The controller and ML model worked well when a path controller was active during braking, even though all training data was obtained with constant steering. This paper has shown that it is possible to predict safe allocations using a data-driven ML model, specifically a neural network model.

References

1. Bienkowski, B.N., Walton, C.M.: The economic efficiency of allowing longer combination vehicles in Texas. University of Texas at Austin. Center for Transportation Research, SWUTC/11/476660-00077-1 (2011)
2. Woodrooffe, J., Ash, L.: Economic efficiency of long combination transport vehicles in Alberta: final report. Technical report, Woodrooffe & Associates (2001)
3. Erdinc, U., Jonasson, M., Kati, M.S., Jacobson, B., Fredriksson, J., Laine, L.: Safe operating envelope based on a single-track model for yaw instability avoidance of articulated heavy vehicles. Veh. Syst. Dyn. **62**(8), 2138–2161 (2023)
4. Hansson, A., Andersson, E., Laine, L., Kati, M.S., Erdinc, U., Jonasson, M.: Safe operating envelope for limiting actuation of electric trailer in tractor-semitrailer combination. In: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), Macau, China, pp. 3886–3893 (2022)
5. Leoni, L., BahooToroody, A., Abaei, M.M., Cantini, A., BahooToroody, F., De Carlo, F.: Machine learning and deep learning for safety applications: investigating the intellectual structure and the temporal evolution. Saf. Sci. **170**, 106363 (2024). ISSN 0925-7535
6. Lattanzi, E., Freschi, V.: Machine learning techniques to identify unsafe driving behavior by means of in-vehicle sensor data. Expert Syst. Appl. **176**, 114818 (2021). ISSN 0957-4174
7. Norouzi, A., Heidarifar, H., Borhan, H., Shahbakhti, M., Koch, C.R.: Integrating machine learning and model predictive control for automotive applications: a review and future directions. Eng. Appl. Artif. Intell. **120**, 105878 (2023). ISSN 0952-1976
8. Fröjd, N.: Handling Analysis and Control Development of Commercial Trucks with Volvo Transport Models. MATLAB EXPO (2021). https://www.mathworks.com/videos/handling-analysis-and-control-development-of-commercial-trucks-with-volvo-transport-models-1622035211192.html

9. Fu, R., Zhang, H., Guo, Y., Yang, F., Lu, Y.: Real-time estimation and prediction of lateral stability of coaches: a hybrid approach based on EKF, BPNN, and online autoregressive integrated moving average algorithm. IET Intell. Transp. Syst. **14**, 1892–1902 (2020). https://doi.org/10.1049/iet-its.2020.0385