

I Move U Move: V2X-Enabled Wireless Towing

Constantine Ayimba¹, Vincenzo Sciancalepore², Paolo Casari³ and Vincenzo Mancuso¹

¹IMDEA Networks Institute, Avenida del Mar Mediterráneo 22, 28918 Leganés, Madrid, Spain

²NEC Laboratories Europe, Kurfürsten-Anlage 36, 69115 Heidelberg, Germany

³University of Trento, Via Sommarive, 9, 38123 Trento, Italy

Abstract

As smart connected vehicles become increasingly common, their ability to provide enhanced services has improved. One such service is the emergency transport of drivers in medical distress. In this paper, we show how such a service can be run from the network and discuss the importance of having a human in the loop in order to expedite driving. We present a Monte-Carlo-based driver assessment system that the network can use in selecting the most suitable candidate to tow an autonomous vehicle with an incapacitated driver. We show that this mechanism results in a selection policy that ensures sufficiently short spacing between the autonomously driven tail/towed vehicle and the human driven lead/towing vehicle ensuring that no other vehicles get in the way to disrupt the service.

Keywords

Connected vehicles, health, emergency, reinforcement learning, wireless towing

1. Introduction

In cases of emergency, it may not always be possible to avail assistance in a timely fashion to people in need. This is especially true in situations where there is no one around the patient to assist directly or to ask for help. The situation becomes even more alarming if the victim becomes incapacitated while engaged in an activity that requires consistent concentration and focus, such as driving.

To cope with the above issue, a recent work proposes to exploit smart cars to continuously monitor the health of the driver to detect events such as cardiac arrests or strokes [1]. The authors of [2] focus on driver fatigue, given the implications this could have on road safety. Vehicle-to-everything (V2X) communications [3, 4], enhanced by Beyond 5G/6G New Radio (NR), are projected to enable a slew of services requiring low latency and high reliability [5]. The

AI6G'22: First International Workshop on Artificial Intelligence in beyond 5G and 6G Wireless Networks, July 21, 2022, Padua, Italy


✉ constantine.ayimba@imdea.org (C. Ayimba); vincenzo.sciancalepore@neclab.eu (V. Sciancalepore); paolo.casari@unitn.it (P. Casari); vincenzo.mancuso@imdea.org (V. Mancuso)

🌐 <https://networks.imdea.org/team/imdea-networks-team/people/constantine-ayimba/> (C. Ayimba); <http://www.sciancalepore.info/> (V. Sciancalepore); <https://webapps.unitn.it/du/it/Persona/PER0221586/Didattica> (P. Casari); <https://networks.imdea.org/team/imdea-networks-team/people/vincenzo-mancuso/> (V. Mancuso)

📞 0000-0003-4032-9735 (C. Ayimba); 0000-0002-0680-7150 (V. Sciancalepore); 0000-0002-6401-1660 (P. Casari); 0000-0002-4661-381X (V. Mancuso)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

latter, coupled with the rise of connected autonomous vehicles [6], makes it possible to create an emergency towing service that can help an incapacitated driver.

In such cases, it is crucial to have a human in the loop. This is because the vehicle will be required to operate like an ambulance without strictly adhering to traffic rules. Ordinary autonomous driving and strict respect of traffic rules works for commonplace, everyday passenger transportation, but may jeopardize the timely delivery of medical assistance to a driver in need. Recent field studies also report public reluctance to fully trust self driving ambulances [7]. Given the sensitivity of the operation, the choice of the lead human driver becomes key, since their driving behavior has a direct impact on the well-being of the incapacitated driver in the towed vehicle. Thus, the agent that picks the suitable lead vehicle needs to know the driving behaviors. This knowledge should be built up over time so as to make a robust decision of which available driver should respond to the towing request.

The closest analog to our proposal is platooning, which involves the coordinated autonomous driving of connected vehicles at shortest but safest inter-vehicular distances [8]. Platooning is however largely aimed at exploiting driving on highways or other open road situations and as such is ill-suited to towing in much more constrained road environments.

In this paper, we specifically look into how to select the most suitable driver based on his driving behavior collected over time. To build this knowledge-based evaluation system we use the Monte-Carlo reinforcement learning approach. We make use of the open-source CARLA platform [9] developed to test self-driving algorithms in realistic scenarios. Concretely, our main contributions are as follows:

1. we implement a follower navigation agent in CARLA that tracks the motion of a lead vehicle;
2. we design a driving behavior profiling service leveraging Monte Carlo Reinforcement Learning;
3. we extend CARLA to carry out driver profiling and selection using ancillary python scripts running in parallel to the simulator;
4. we evaluate the effectiveness of the selection algorithm in choosing the most suitable driving profile to tow a target tail vehicle.

To the best of our knowledge, ours is the first approach to leverage autonomous driving and connected cars to deliver an ad hoc emergency towing service with a human in the loop. The latter element is crucial with regards to gaining public confidence, given the reticence we have mentioned regarding fully autonomous ambulances.

2. Wireless towing scenario and service

2.1. Wireless towing definition

We define wireless towing as the autonomous movement of a vehicle in sympathy with the movement of another vehicle that precedes it, based on directives it receives from the lead vehicle itself, the network or both. While this is similar to platooning, the difference here is that the spacing between the lead and tail vehicle is not maintained relatively constant, and in the same vein the relative speed of the two vehicles is not zero. This fact is of great importance,

because the emergency towing service should be versatile enough to work in crowded urban environments, with intervening traffic as well as stops due to traffic lights and pedestrians.

We assume that there exists an on-board unit on the tail vehicle that can respond to driving directives from the network prescribing when and by how much to accelerate or brake. Recent works such as [10, 11] point to the practicality of this assumption. Furthermore, the ability of an autonomous vehicle to drive itself solely based on received instructions has been demonstrated in field trials [12].

2.2. Scenario description

In the scenario we envisage, the vehicle's sensor signal that the driver is unresponsive as has been described in, e.g., [1, 2]. Upon such a detection, the vehicle sends out an alert to the network which reacts by running the emergency service. In the alert, the vehicle specifies its geographic information, model and engine data. The network then polls available drivers in the vicinity of the target vehicle. The candidate lead vehicles that respond to the poll start moving towards the target. The network continuously monitors the motion of these vehicles. After a short duration (in the order of 30 s) the vehicle evaluates the driving behavior of each responding vehicle. The network agent then selects a candidate vehicle from the ones available to carry out the task.

When the chosen lead vehicle reaches the target tail vehicle, it overtakes it and stops slightly ahead of it after confirmation from the network. The network then alerts the tail vehicle to activate its tow mode so as to follow the lead vehicle ahead of it. The tail and lead vehicle maintain constant communication with each other and the network. Both vehicles constantly communicate their geographic positions and speed among others to the network.

When the vehicles reach their intended destination, the network agent evaluates the driving experience based on the data it received over the course of towing. The agent calculates the reward based on the evolution of the spacing between the lead and tail vehicle. The rationale behind this is that, for an effective service, the space between the vehicles needs to be maintained small enough not to allow other vehicles to get in between, as this would result in delays depending on how the preceding vehicle is being driven. Moreover, a conscious tail vehicle passenger should not get the feeling that the wireless towing service is leading him or her astray. Both events would negatively affect the effectiveness of the service.

The network agent will then update the corresponding driving behavior profile with the calculated reward in its reinforcement learning data base. As the service is increasingly used, more of these experiences will be similarly evaluated and a selection policy progressively developed. Following the policy, the selection of the most suitable lead vehicle will continuously improve. In due course, the network agent will consistently select the candidate lead vehicle with the driving behavior that exhibits the best spacing discipline, hence the best service quality.

3. Design of driving behavior selector

As stated earlier, we leverage Monte-Carlo reinforcement learning [13] to create our driving behavior evaluation and selection algorithm. In particular, we characterize the environment




Figure 1: Migrator State and Action Spaces. The driving behavior indicator (χ) and vehicle power ratio, (η), constitute the state space. The action space is represented by the red and green colored cells. The green cells represent candidate lead vehicles (labelled \downarrow) which are less powerful than the vehicle to be towed. The red cells are candidate lead vehicles (labelled \uparrow) with the same or higher engine power as the vehicle to be towed.

as having a driving behavior indicator, χ , of the candidate lead driver and a ratio between the engine power of the lead vehicle and tail vehicle, η . State and action spaces are shown in Fig. 1.

When the alert is received, a candidate lead vehicle sets up a path to the position of the target vehicle. The path is constituted by many waypoints at regular intervals, which dictate the movement of the vehicle to the meeting point. If \dot{x}_j is the sample speed at waypoint j , $\sigma_{\dot{x}}^2$ the empirical variance of the speeds across all waypoints, \bar{x} the mean speed between waypoints and W is the number of waypoints until the meeting point then, χ is given as:

$$\sigma_{\dot{x}}^2 = \frac{1}{W} \sum_{j \in W} (\dot{x}_j - \bar{x})^2; \quad (1)$$

$$\chi = \frac{\sum_{j \in W} \dot{x}_j}{W \sum_{j \in W} \dot{x}_j^2} \times \frac{\sqrt{\sigma_{\dot{x}}^2}}{3}. \quad (2)$$

The first term of Eq. (2) is the fairness index [14] of speed between waypoints and the second term is points to the standard deviation in speeds between waypoints normalised according to the 3σ empirical rule [15]. To obtain η , we use the ratio between the maximum revolutions per minute (RPM) that can be reached by candidate lead vehicle and target vehicle to be towed, i.e.:

$$R_{\text{rpm}} = \frac{\text{MaxRPM}_{\text{candidate}}}{\text{MaxRPM}_{\text{target}}}; \quad (3)$$

$$\eta = \begin{cases} R_{\text{rpm}}, & \text{if } R_{\text{rpm}} < 1.0; \\ \frac{1}{R_{\text{rpm}}}, & \text{otherwise.} \end{cases} \quad (4)$$

In order to keep the state space small, we quantize χ to 10 levels $0, \dots, 9$ using a uniform quantizer. We similarly quantize η to 5 levels $0, \dots, 4$. We update the state values using the following modified update:

$$V(S) \leftarrow \frac{1}{k} [\Delta + (k-1)V(S)] , \quad (5)$$

where $\Delta = R + (\gamma - 1)V(S)$, $\gamma \leq 1.0$ is a discount parameter, $k \geq 1$ is the number of times the agent was in state S , and R is the reward at the end of the episode. R is calculated as

$$R = -\frac{1}{N} \sum_{i \leq N} \sqrt{(x_i - x_{\text{ideal}})^2} , \quad (6)$$

where x_i is the spacing between the lead car and target tail vehicle taken at regular intervals, N is the number of intervals until the end of the episode and x_{ideal} is the desired spacing that should be maintained between the lead and tail vehicle.¹ The formula in Eq. (5) is modified because it does not consider the simple reward R as the update is done at the end of the episode and the lead driver/vehicle selected in a previous episode have no bearing on the current episode.

At the start of each episode, we use ϵ -greedy selection to choose the lead car and driving behavior option. We also employ the weighted fair exploration mechanism presented in our previous work [16] to improve the trade-off between exploration and exploitation.

4. Performance evaluation

4.1. System setup using the CARLA simulator

The CARLA simulator is a distributed system, where the server spawns the map and ancillary infrastructure, i.e., roads, traffic lights, and buildings. The client runs python scripts which we use to introduce the lead and tail vehicle in the simulation. Multiple clients can run on the same server, and we exploit this feature to spawn other vehicles and pedestrians into the environment to mirror real-world traffic.

In order to solely test the performance of the selection agent against driving behaviors, we spawn the same car model for both the lead and tail vehicle. For the lead vehicle, we implement behavior agents which build on the basic agent provided with the simulator. This allows us to emulate different driving styles for simulated human drivers. We supply the destination coordinates to a method of the behavior agent, which leverages the CARLA world object with its global view of the map in order to generate waypoints to the specified location. The lead vehicle records its current position and bearing at regular intervals, thereby generating waypoints which the tail vehicle uses to follow the lead. For each of the four candidate vehicles shown in Fig. 2, labeled as A, B, C and D, we specify a behavior agent exhibiting a given driving behavior: Extra-cautious (EC), Cautious (CS), Normal (NL) and Aggressive (AG). The characterisation of each is given by the parameters in Table 1. We remark that these are base behaviors from which a driving behavior indicator χ can be derived.

¹We set $x_{\text{ideal}} = 10$ m.




Figure 2: Simulation scenario. After the emergency signal is received, each candidate lead vehicle (colored cars marked as A, B, C, D) starts driving towards the meeting point MP, from where the target vehicle will be towed. The other vehicles represent prevailing traffic present on the road, while the dots represent pedestrians.

Table 1
Driving behaviours of lead vehicle

	EC	CS	NL	AG
Max. Speed [km/h]	30	40	50	70
Deviation from Speed limit [km/h]	10	6	3	1
Braking level [km/h]	15	12	10	8
Sudden braking time to collision [s]	5	3	3	3
Distance to object before avoidance [m]	15	12	10	8
Emergency braking distance [m]	10	6	5	4

After a 30 s behavior evaluation period, χ is determined for each candidate vehicle using Eq. (2) and one of the cars/driving behavior is selected and proceeds to an agreed Meeting point (MP) shown in Fig. 2. Given that the seed for each simulation is generated randomly the behaviors are not completely deterministic, in each simulation round. Therefore, one behavior may map into a different indicator in a subsequent simulation. This reflects the fact that human drivers do not drive exactly the same way each time. We prefer to use the indicator χ to cover this spectrum of behaviors.

In order to implement the towing behavior, we modify the basic agent provided with the simulator so that its array of waypoints is not generated by the world object but is instead supplied by the lead vehicle. The autonomous driving behavior of the tail vehicle is thus greatly influenced by the behavior of the lead vehicle for the path between the meeting point and the destination (cf. Fig. 2). Owing to this influence, the question of which vehicle to follow becomes an important decision.

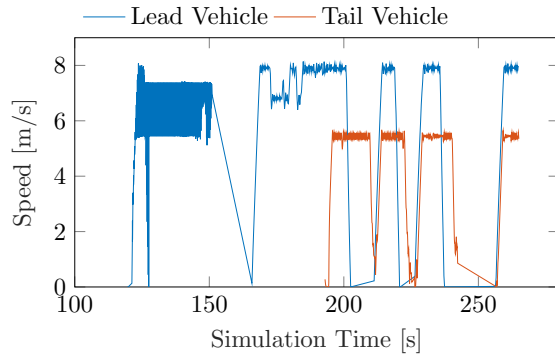


Figure 3: Instantaneous speeds of the lead and tail vehicles. The tail vehicle moves in sympathy with the lead vehicle but always has to catch up to it.

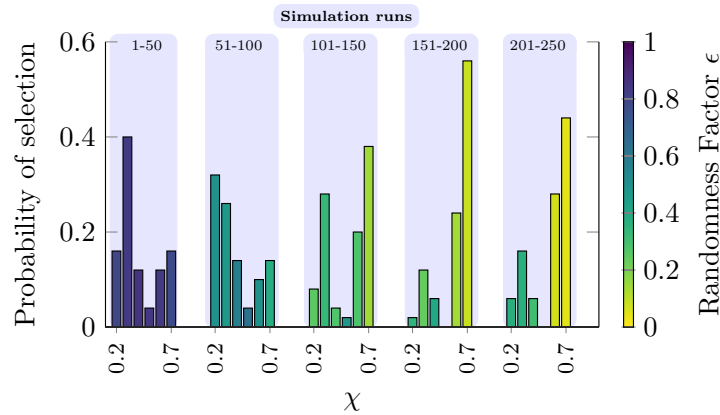


Figure 4: Lead Car selection learning. The bars represent behavior indicators from $\chi = 0.2$ to $\chi = 0.7$. The color scale indicates the probability that the selection was a random decision not a learned policy.

5. Results

We first look at the towing performance overall. From Fig. 3, we see that the tail vehicle starts to move after the lead vehicle gets to the meeting point and tries to keep up with it. Given that the lead vehicle precedes the tail, it bears the inconvenience of waiting for longer at traffic lights, so that by the time the tail vehicle has caught up it does not have to wait as long, or at all.

In Fig. 4 we plot the probability of a lead car behavior being selected in every set of 50 simulations. In the initial sets, {1..50} and {51..100}, when the selection algorithm is still developing a policy, $\chi = 0.3$ and $\chi = 0.2$ are selected most often. The spacing discipline between the tail and lead vehicle progressively improves as can be seen by comparing Fig. 5a for $\chi = 0.3$ and Fig. 5b for $\chi = 0.2$.

As the policy is further developed, beyond 100 simulations, the selection leans towards driving behaviors $\chi = 0.6$ and $\chi = 0.7$. The spacing maintained between the lead and tail vehicle also further improves as can be seen from Fig. 5c. The eventual policy that is settled on, choosing $\chi = 0.7$ results in the most favourable spacing between the tail and lead vehicle as depicted in Fig. 5d.

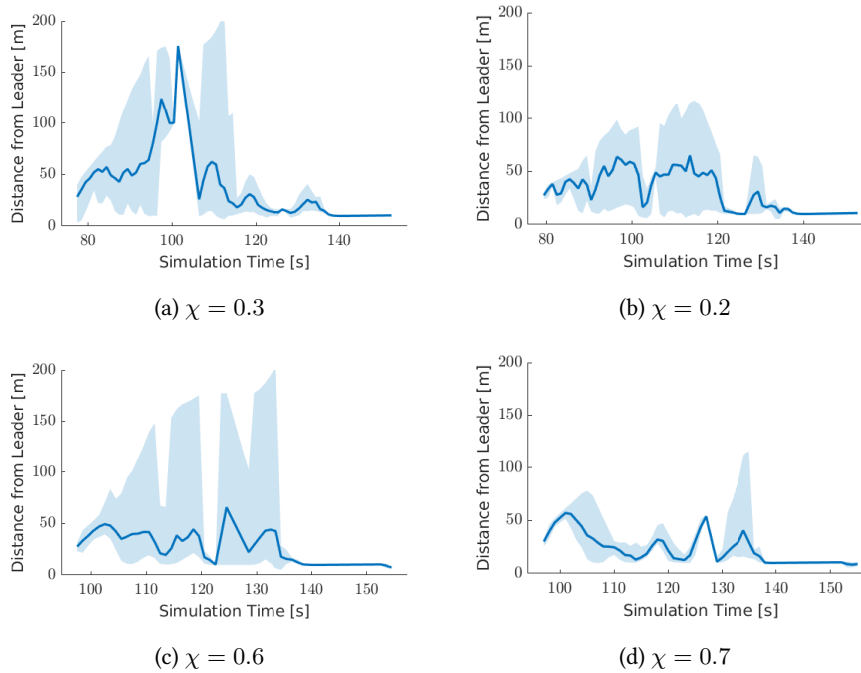


Figure 5: Spacing evolution for popular driving indicators. The main curve is the median value of all the simulations with particular χ values grouped in 2 s intervals. The ribbon around it represents the 95% confidence interval.

6. Conclusions

We have presented a novel network service that leverages connected vehicles and their self-driving capabilities to deliver emergency towing. We have also implemented a selection algorithm for the service that uses Monte-Carlo reinforcement learning to improve the system’s capability to select the most appropriate driver for the towing service, based on driving behavior assessment. We have shown that the algorithm converges to selecting the driving behavior that maintains the best spacing discipline throughout the towing process, thereby ensuring that no intervening vehicles get in between to disrupt the service.

As part of future work, we intend to investigate more complex scenarios and the cases in which a lead vehicle can only tow the tail vehicle for part of the way. In this case the network will need to engage multiple candidate vehicles and will consequently require to plan ahead to minimize delays by aligning the handover process between lead vehicles.

Acknowledgments

The work was supported by the ECID project (grant PID2019-109805RB-I00 funded by MCIN/AEI/10.13039/501100011033) and by the Italian Ministry for University and Research (MIUR) under the initiative “Departments of Excellence” (Law 232/2016).

References

- [1] S. J. Park, S. Hong, D. Kim, I. Hussain, Y. Seo, Intelligent in-car health monitoring system for elderly drivers in connected car, in: Proc. 20th Congress of the International Ergonomics Association (IEA 2018), Springer International Publishing, Cham, 2019, pp. 40–44.
- [2] G. Sikander, S. Anwar, Driver fatigue detection systems: A review, *IEEE Transactions on Intelligent Transportation Systems* 20 (2019) 2339–2352. doi:10.1109/TITS.2018.2868499.
- [3] K. Abboud, H. A. Omar, W. Zhuang, Interworking of dsrc and cellular network technologies for v2x communications: A survey, *IEEE Transactions on Vehicular Technology* 65 (2016) 9457–9470. doi:10.1109/TVT.2016.2591558.
- [4] Z. MacHardy, A. Khan, K. Obana, S. Iwashina, V2x access technologies: Regulation, research, and remaining challenges, *IEEE Communications Surveys Tutorials* 20 (2018) 1858–1877. doi:10.1109/COMST.2018.2808444.
- [5] X. Cheng, Z. Huang, S. Chen, Vehicular communication channel measurement, modelling, and application for beyond 5g and 6g, *IET Communications* 14 (2020) 3303–3311.
- [6] R. Coppola, M. Morisio, Connected car: Technologies, issues, future trends, *ACM Comput. Surv.* 49 (2016). URL: <https://doi.org/10.1145/2971482>. doi:10.1145/2971482.
- [7] A. Zarkeshev, C. Csiszár, Are people ready to entrust their safety to an autonomous ambulance as an alternative and more sustainable transportation mode?, *Sustainability* 11 (2019). URL: <https://www.mdpi.com/2071-1050/11/20/5595>. doi:10.3390/su11205595.
- [8] F. Dressler, F. Klingler, M. Segata, R. Lo Cigno, Cooperative driving and the tactile internet, *Proceedings of the IEEE* 107 (2019) 436–446. doi:10.1109/JPROC.2018.2863026.
- [9] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, V. Koltun, CARLA: An open urban driving simulator, in: Proc. 1st Annual Conference on Robot Learning, volume 78 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 1–16.
- [10] C. Quadri, V. Mancuso, M. Ajmone Marsan, G. P. Rossi, Platooning on the edge, in: Proceedings of the 23rd ACM MSWiM, MSWiM'20, 2020, p. 1–10.
- [11] C. Ayimba, M. Segata, P. Casari, V. Mancuso, Closer than close: Mec-assisted platooning with intelligent controller migration, in: Proceedings of the 24th ACM MSWiM, MSWiM'21, 2021, pp. 23–32.
- [12] S. Tsugawa, S. Jeschke, S. E. Shladover, A review of truck platooning projects for energy savings, *IEEE Transactions on Intelligent Vehicles* 1 (2016) 68–77. doi:10.1109/TIV.2016.2577499.
- [13] R. S. Sutton, A. G. Barto, Introduction to Reinforcement Learning, 2nd ed., MIT Press, Cambridge, MA, USA, 2020. URL: <http://incompleteideas.net/book/RLbook2020.pdf>.
- [14] R. K. Jain, D.-M. W. Chiu, W. R. Hawe, et al., A quantitative measure of fairness and discrimination, Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA (1984).
- [15] D. Vysochanskij, Y. I. Petunin, Justification of the 3σ rule for unimodal distributions, *Theory of Probability and Mathematical Statistics* 21 (1980).
- [16] C. Ayimba, P. Casari, V. Mancuso, SQLR: Short-Term Memory Q-Learning for Elastic Provisioning, *IEEE Transactions on Network and Service Management* 18 (2021) 1850–1869. doi:10.1109/TNSM.2021.3075619.