

Acoustic-driven Interior Vehicle Adaptation based on Deep Reinforcement Learning to Improve Driver’s Comfort

Erickson R. Nascimento^{1,2}, Ruzena Bajcsy¹, Gregorij Kurillo¹, and Ismael Villegas¹

Abstract—The safety and comfort of drivers have been improved over the decades through new technologies and driver modeling studies that broadened understanding of predicting driver’s behaviors. Despite remarkable advances in autonomous systems and interactive systems, there is a significantly lack of approaches considering the passengers and the vehicle as components of a dynamical vibro-acoustical system. This paper aims to investigate the interplay between the perceived sounds of a vehicle and the psychoacoustic annoyance metrics. Our goal is to create an intelligent agent that acts to improve the driver’s pleasantness through acoustic-driven learning. To tackle the problem of acting to reduce the annoyance, we present a method based on reinforcement learning that learns from the environment, i.e., the vehicle interior. Our method changes the state inside the vehicle (closing or opening the window and choosing the cruise speed) to avoid annoying sounds in its interior. The results of this work, all performed using the GTA V simulator, showed that the trained agent learned to take actions to avoid creating annoying sounds.

I. INTRODUCTION

A challenge in developing safety systems for driving is predicting and responding to the environment state and the behaviors of human agents. These systems must interact in mixed environments and be aware of different sources of information about the environment state such as acoustical and visual signals. In addition to ensuring safe and effective autonomous systems, it is also essential to take into account the pleasantness and comfort of the driving experience. While there is a great deal of work going into bringing autonomous vehicles into fruition, it is noteworthy that an intelligent assistance system should focus on improving the in-car experience for the driver, which is one of the most critical factors on the complex transportation system.

In this paper, we go beyond safety systems and by learning how to accommodate the driver best. An important factor to consider is the vehicle interior sound, which provides information about the state of the environment and can impact the physiological behavior of drivers. The sound in vehicles is not only informative of the state of the vehicle and the environment, but it can also affect driver’s performance, attention, and comfort of driving. After all, when driving, we are secluded from the surrounding world by the vehicle, which decreases our sensory experience of the world and increases the risk of accidents.

*The authors would like to thank CAPES, CNPq, FAPEMIG, and NSF VeHICaL project (#1545126) for funding this work.

¹The authors are affiliated with the Electrical Engineering and Computer Sciences at the University of California, USA. Emails: {bajcsy, gregorij, ismaelvillegas}@berkeley.edu

²Erickson R. Nascimento is also affiliated to Universidade Federal de Minas Gerais (UFMG), Brazil. Email:erickson@dcc.ufmg.br

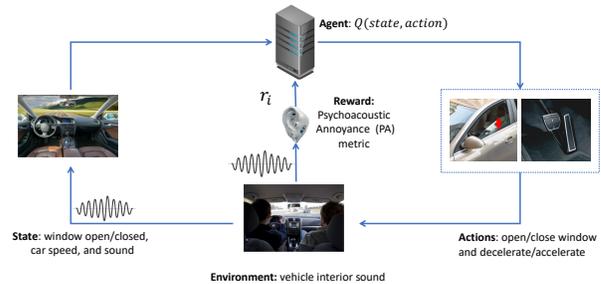


Fig. 1. Our approach learns from the environment how to change the state of the vehicle to avoid an unpleasant environment. After the agent taking an action a_i considering the environment state s_i , our system computes the reward using the psychoacoustic annoyance metric. The reward, the action, and the environment states are used to train a neural network that approximates the Q-function. The network is used by the agent to chose the best action that decreases the annoyance.

While on one hand sound can provide a calming effect on the driver and improve their driving experience, on the other hand, the vehicle interior noise may increase the probability of traffic accidents because of the physiological and psychological effects of sounds on the human. In the context of this paper, *comfortable* and *pleasant* concern a conscious state related to the satisfaction of a person while driving a vehicle. For a complete picture of the neural bases of emotion and mood, we refer the reader to the works of Dalglish [1] and Ruckmick [2].

Although the human perception of sound as pleasant or annoying is very subjective, some general findings in human subject studies have resulted in the determination of several psychoacoustic properties. Due to individual differences and context of sound exposure, different people may have a different perception of annoyance, which may depend on and consequently affect their emotional state. In this paper, we aim at investigating the interplay between the perceived sounds inside a vehicle and the driver’s attention/enjoyment that can be improved through individualized acoustic-driven preference learning. The sound in the interior of a vehicle is influenced by several uncorrelated and dynamics factors such as road type, road quality, radio, and the number of passengers [3], to name a few. To tackle this problem, we can learn from the environment how to change the state inside the car to avoid an annoyance in the interior of the vehicle and keep driver’s attention, closing the vehicle’s windows or suggesting a different cruise speed.

Traditional learning techniques use an existing dataset and associated features to find trends in the data. Typically, this means the features need to be defined beforehand and once the input-output mapping is learned, and unless adapted to an online framework, remains stagnant. More recent advances in deep learning eliminate the need for identifying features beforehand but requires incredibly large amounts of data and computation time to learn the functional relationship between the data and output. Similar to the more traditional methods, these new artificial intelligence (AI) methods do not have an efficient updating scheme. We focus on developing/designing the learning framework for active vehicle adaptation based on psychoacoustic metric, derived from previous studies. The main advantage of our approach is in its capability of learning from a realistic simulation and not requiring annotated data. We evaluate the performance of the developed agent in a simulated environment without a user input.

II. RELATED WORK

The safety of drivers has been improved over the decades through driver modeling studies [4]–[6] that broadened understanding of predicting driver’s behaviors. Despite remarkable advances in visual perception, mostly applied in autonomous systems, and driver modeling, there is an important gap when considering the passengers, the driver and the vehicle as components of a vibro-acoustical system. This work builds upon findings from the areas of sound recognition and psychoacoustic sound evaluations.

Virtually all works on noise treatment in the vehicle have focused on noise detection or reduction for adjustments during design time of a vehicle. Several psychoacoustic indices have been introduced in sound quality evaluation engineering [7]–[9]. Although several psychoacoustic indices have been proposed in the past years, when driving, this information alone is not useful for identification of the sound source of annoyance. A straightforward approach is to separate the sound events and compute the annoyance of each one. Unfortunately, the detection and segmentation of general sound events is still a not well-solved problem.

In the last few years, thanks to the advances in transfer learning techniques and cross-modal learning, deep learning method attain the state-of-the-art results for most sound recognition problems. However, one particular problem of learning from sound samples is the lack of massive annotated data. This absence is particularly accentuated for sounds from the vehicle interior as most of the available sound datasets are composed of ambient, event or mixed sounds. Besides being influenced by several uncorrelated and dynamics factors, the sound in the interior of a vehicle also affects people (i.e., drivers) differently. Thus, in addition to the absence of data, a challenging task is how to learn to adapt from the environment and also from the drivers preferences to keep a pleasant environment inside the car for a particular driver.

Due to this scarcity of labeled data in sound classification tasks, several deep learning approaches build upon some transfer learning technique. Han et al. [10] tackle the absence

of labeled data by combining active learning and a self-training approach to minimize the human annotation. Their methodology was able to provide a reduction of 52.2% in human labeled instances on the FindSounds dataset¹. This dataset has nine categories and more than 16,930 sound instances with durations ranging from 1 to 10 seconds. SoundNet [11], for instance, provides a large-scale and rich representation of natural sound by using frames from labeled videos for learning a representation for raw audio waveforms.

Çakir et al. [12] presented a convolution recurrent neural network (CRNN) for polyphonic sound event detection, where multiple overlapping sound events must be detected. After convoluting spectrograms with multiple convolutional layers and max-pooling, they fed a recurrent neural network (RNN), which has the output layers binarized for the event classification. Their experiments showed their CRNN outperformed traditional Gaussian Mixture Models-Hidden Markov Models (GMM-HMM) classifiers and methods using convolutional neural network or RNN.

Despite the impressive results on the supervised approaches, its full integration on real-world scenarios faces several challenges. Foremost, is the requirement of a large annotated dataset, which can be unfeasible to be created in some cases, where the set of labels is ill-defined or too large.

III. METHODOLOGY

We detail our exploration-exploitation approach based on a deep reinforcement learning, where an agent learns actions to perform in order to decrease the annoyance. Since the sound in the interior of a vehicle is influenced by several uncorrelated and dynamics factors, our agent can learn from the environment how to change the state inside the car to reduce the sound annoyance in the interior of the vehicle. Figure 1 depicts the main steps of the learning method.

A. Learning

We formulate the problem of choosing an action that avoids annoying sounds as a Markov Decision Process (MDP). Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a discrete action set and \mathcal{S} the state set. Thus, the agent takes action a_i considering the environment state s_i represented at the i -step of the episode. An action a_i leads to a state transition from the state s_i to s_{i+1} and a immediate reward r_i . In order to maximize the accumulated reward R defined as $R = \sum_i \gamma^{k-1} r_i$, where $\gamma \in [0, 1]$ is the discount factor for future rewards, the learning process tries to minimize the loss function:

$$\mathcal{L} = \frac{1}{2} \left(r + \max_{a'} Q(s', a') - Q(s, a) \right)^2, \quad (1)$$

where $Q(s, a)$ is the function that gives the best score after performing the action a in the state s .

Our approach uses a ϵ -greedy strategy to explore the action space and learn the result of each action on the annoyance metric. Thus, probability $1 - \epsilon$ the action is determined by the policy

$$\pi(s_i) = \underset{a}{\operatorname{argmax}} Q(s_i, a). \quad (2)$$

¹<http://www.findsounds.com>

The Q -function is modeled by a multilayer perceptron network, which receives as input a state vector and returns a vector containing the Q -value for each possible action. This network is initialized with random parameters.

B. Immediate reward

The sound impression for a human listener can be estimated by the following psychoacoustic properties [13]:

- *Fluctuation and Roughness*: Multiple signals with different frequencies might interfere constructively and destructively with each other creating modulation. Fluctuation and roughness measure the modulation of a signal over the time. Fluctuation was designed to work with up to 20 modulations per second and can be given by: $F \approx \Delta L / (4Hz / f_{mod} + f_{mod} / 4Hz)$, where ΔL is the modulation depth and f_{mod} the modulation frequency. The roughness R describes sounds with modulations range from 20 to 300 times per second, and can be computed as being the product $R \approx \Delta L \times f_{mod}$. A modulated signal is considerably more unpleasant when having a higher roughness and fluctuation;
- *Loudness*: The loudness is a psychological phenomenon based on the perceived loudness. Unlike the sound level that is a physical measurement, the loudness metric is based on subject studies using a group of people with normal hearing. Formally, let E_{TQ} be the excitation at threshold in quiet, and E_0 be the excitation of the reference intensity, the specific loudness of a sound with excitation E is given by

$$N' = 0.08 \left(\frac{E_{TQ}}{E_0} \right) \left[\left(0.5 + 0.5 \frac{E}{E_{TQ}} \right) - 1 \right]. \quad (3)$$

The total loudness is the result of integrating the specific loudness over critical-band rates ², i.e.,

$$N = \int_0^{24} N' dz, \quad (4)$$

where z is the critical-band in Bark.

- *Sharpness*: It is a function of the spectral composition. It is estimated by a weighted sum of specific loudness levels in different bands. The total sharpness is given by $S = \int_0^{24} S' dz$, where $S' = \frac{0.11}{N} \int_0^{24} N' g(z) z dz$ is the specific sharpness and $g(\cdot)$ is a critical-band-rate dependent weighting function. The sound with higher sharpness is more unpleasant and more annoying.

The sound annoyance is closely related to the aforementioned psychoacoustic indices. Zwicker proposes to compute the psychoacoustic annoyance (PA) [13] as a function of sharpness, loudness, fluctuation, and roughness as follows:

$$PA = N_5 \left(1 + \sqrt{\omega_S^2 + \omega_{FS}^2} \right), \quad (5)$$

²A critical band is the smallest band of frequencies that activates the same part of the basilar membrane in the human hearing system.

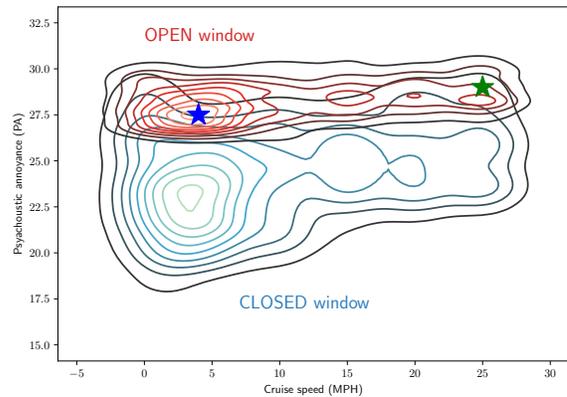


Fig. 2. Psychoacoustic annoyance metric measured in different configurations of the window status and cruise speed. The red curve depicts when the window is open and the blue curve when window is closed. All measures include additional noise. Blue star represents a state where the agent can decrease the annoyance by simply closing the window. Green star is a more challenging state where the agent has to combine the actions of selecting a lower speed and closing the window to decrease the annoyance.

where N_5 is the 95th percentile of loudness and

$$\omega_S = \begin{cases} -(S - 1.75) \log(N_5 + 10), & \text{if } S > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

$$\omega_{FS} = \frac{2.78}{N_5^{0.4}} (0.4F + 0.6R). \quad (7)$$

We define the immediate reward as being a function of the psychoacoustic annoyance metric, i.e., $r_i = f(PA)$ is given by the PA metric, where $f(\cdot)$ is the shape function $f(x) = 1 - (x/MAX_{PA})^{0.4}$ and MAX_{PA} is the maximum acceptable value for PA. In our experiments we used $MAX_{PA} = 27$.

IV. EXPERIMENTS

In this section, we present the experiments and the analysis of the effectiveness of our proposed framework.

A. Implementation details

We used a 5-layers neural network to represent the Q -function. The first layer is a fully connected layer of 128 neurons followed by a rectifier linear unit (ReLU) activation function. The second, third and fourth layers have 64, 32 and 24 neurons, respectively. A ReLU follows all of them. The last layer is composed of $|\mathcal{A}|$ neurons, producing the outputs as a weighted sum of the 4th layer followed by a sigmoid function. The action set \mathcal{A} is composed of 7 actions: changing the window state (open or closed) and six different cruise speed: 2, 4, 8, 10, 15 and 20 miles per hour. The learning rate used was equal to 0.001. The exploration rate from 1 and to 0.1 with a decay rate of 0.95 and discount factor $\gamma = 0.8$. We set the memory size to 256 transitions, batch size 32 and 300 epochs. A window of size equal to 3 seconds is used to compute the psychoacoustic annoyance metric after performing the selected action.

We used the Grand Theft Auto (GTA) V (Rockstar Games, NY) game to generate different interactive scenarios. During

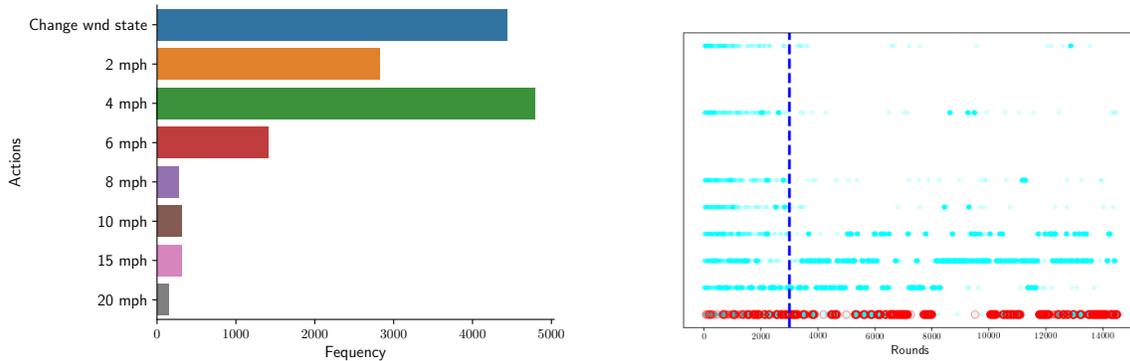


Fig. 3. On the left, we can see the number of times the agent chose each action. These results show that, in general, agent learned to take actions that prevent large PA values. On the right, the plot depicts the rounds where there was noise, and the window’s state was open (red circles). In several rounds, the agent prefers closing the window when is perceived noise that is coming from outside of the vehicle.

the task, the agent was exposed to various sounds due to environmental factors, and other external auditory stimuli, such as pedestrians, and traffic. Additionally to the sound environment of the simulator, we also included sounds of bells and beeps from outside in half of the time. Figure 2 shows the values of psychoacoustic annoyance metric when setting different cruise speeds and changing the state of the vehicle’s window. Whenever the window state is closed, the bells and beeps noise is not perceived by the agent.

B. Discussion

In Figure 3, the plot on the right shows the behavior of the agent selecting actions during the learning process. We can see that at the beginning of the training, the agent explores the action state by selecting different actions uniformly. After 3,000 rounds (vertical blue line), the agent starts to avoid cruise speeds larger than 8 mph, since these speeds return large PA values. More interesting, we can see that the agent not only learned to select low speeds but also learned to change the window state when it is necessary. We show red circles where there was noise, and the window’s state was open. The first row on the plot represents the agent choosing to close the vehicle window to decrease the annoyance. It is worth noting that the agent correctly took the action that closes the window when there is noise from outside. Also, this behavior became predominant after 3,000 rounds.

The bars on the left in Figure 3 show the number of times the agent chose each action. As expected, most of the time the agent took actions to prevent large PA values, i.e., change window state and low cruise speed.

V. CONCLUSIONS

In this work, we present an acoustic-driven method based on deep reinforcement learning to adjust the interior vehicle state to improve driver’s pleasantness. Our method feeds a Q-network with psychoacoustic indices, which measure the annoyance coming from the sound using. The network encodes the best action to be taken to avoid annoyance as far as sound is concerned. In our experiments, we showed the relation between the psychoacoustic indices and the cruise

speed as well as the noise coming from outside of the vehicle. The results presented in this paper, all performed using the GTA V simulator, showed that the agent learned to choose actions to avoid creating annoying sounds.

As future works, we will further investigate other approaches, including agent as a recommendation system, where every time the driver follows a suggestion, the system is rewarded. Such reward can also be computed by measuring the behavior of the driver such as engaging in speeding.

REFERENCES

- [1] T. Dalgleish, “The emotional brain,” *Nature Reviews Neuroscience*, vol. 5, no. 7, p. 583, 2004.
- [2] C. A. Ruckmick, “The psychology of pleasantness,” *Psychological Review*, vol. 32, no. 5, p. 362, 1925.
- [3] P. Davies, “An overview of methods to quantify annoyance due to noise with application to tire-road noise,” Citeseer, Tech. Rep., 2008.
- [4] C. Little, “The intelligent vehicle initiative: advancing” human-centered” smart vehicles,” *Public Roads*, vol. 61, no. 2, pp. 18–25, 1997.
- [5] E. Wahlstrom, O. Masoud, and N. Papanikolopoulos, “Vision-based methods for driver monitoring,” in *Intelligent Transportation Systems, 2003. Proceedings. 2003 IEEE*, vol. 2. IEEE, 2003, pp. 903–908.
- [6] K. Driggs-Campbell, V. Govindarajan, and R. Bajcsy, “Integrating intuitive driver models in autonomous planning for interactive maneuvers,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 12, pp. 3461–3472, 2017.
- [7] H. Fastl, “The psychoacoustics of sound-quality evaluation,” *Acta Acustica united with Acustica*, vol. 83, no. 5, pp. 754–764, 1997.
- [8] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [9] M. J. M. Nor, M. H. Fouladi, H. Nahvi, and A. K. Ariffin, “Index for vehicle acoustical comfort inside a passenger car,” *Applied Acoustics*, vol. 69, no. 4, pp. 343–353, 2008.
- [10] W. Han, E. Coutinho, H. Ruan, H. Li, B. Schuller, X. Yu, and X. Zhu, “Semi-supervised active learning for sound classification in hybrid learning environments,” *PLoS one*, vol. 11, no. 9, p. e0162075, 2016.
- [11] D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds., *SoundNet: Learning Sound Representations from Unlabeled Video*. Curran Associates, Inc., 2016.
- [12] E. Çakir, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, “Convolutional recurrent neural networks for polyphonic sound event detection,” *IEEE/ACM Trans. Audio, Speech & Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [13] E. Zwicker and H. Fastl, *Psychoacoustics: Facts and models*. Springer Science & Business Media, 2013, vol. 22.