Personalised Electric Vehicle Acoustics with Generative AI: Dynamic Sonification and User Acceptance Study

WESSEL VERSTELLE, Eindhoven University of Technology, The Netherlands
MD SHADAB ALAM, Eindhoven University of Technology, The Netherlands
PAVLO BAZILINSKYY, Eindhoven University of Technology, The Netherlands

Electric vehicles (EVs) reduce powertrain noise, creating safety challenges and opportunities for sound design. This paper examines whether generative audio supports personalised dynamic acoustics for future EVs. We report a research through design study in which AI-generated sound samples were selected, edited into seamless loops and embedded in a sonification (the process of translating non-auditory data into sound) prototype. The prototype connects a Processing vehicle interface to a Pure Data audio engine using Open Sound Control. Vehicle speed and throttle input modulate pitch, amplitude and load related parameters. A user study with 20 participants combined the Acceptance Scale with open questions. The results show moderately positive acceptance, with an average usefulness of 0.88 and a satisfaction of 0.62 on a scale from -2 to +2. Participants valued personalisation and responsiveness, but requested stronger recognisability and better throttle mapping. Generative AI is useful for early EV acoustic prototyping, while the final design requires expert refinement and safety evaluation.

Additional Key Words and Phrases: Electric vehicles, Sonification, Generative AI, Vehicle acoustics, Sound design

1 Introduction

EVs are changing the acoustic character of road traffic. Compared to vehicles with an internal combustion engine, EVs produce substantially less noise from the powertrain, especially at low speeds [5]. This can reduce environmental noise, but it also creates challenges for pedestrians and other road users who rely on sound to recognise traffic. Regulations therefore require acoustic warning systems for quiet vehicles in several markets [8]. At the same time, sound is no longer only a safety requirement. Manufacturers increasingly use artificial vehicle acoustics to communicate performance, brand identity, and driving feedback [10, 12].

This change creates a design opportunity. Since EV sounds are deliberately authored rather than mechanically produced, they can be adaptive, personalised, and context aware. However, prior research also shows that artificial vehicle sounds are difficult to design well. Sounds that simply imitate combustion engines can be perceived as inauthentic, while abstract futuristic sounds may not communicate that they belong to a vehicle [6, 15]. Effective EV acoustics must therefore balance safety, recognisability, emotional acceptance, and technological authenticity.

Research on EV warning sounds shows that pedestrians and drivers do not evaluate sounds only by loudness. Detectability, reaction time, expectation, and perceived vehicle identity all matter [13, 15]. Engine sound has traditionally contributed to perceived performance, luxury, and brand identity [12, 14]. As EVs remove much of this mechanical cue, manufacturers increasingly design artificial sound as part of the user experience [10]. Continuous auditory feedback is also explored for vehicle automation, where subtle sounds can communicate driving status while raising concerns about cabin comfort if the feedback is too persistent [3]. Research on EV sound quality suggests that drivers often prefer refined, pleasant, and comfortable sound profiles to aggressive or sharp ones [16, 17]. High frequency sharpness can reduce perceived quality, while subtle low frequency components can improve comfort and perceived power.

Authors' Contact Information: Wessel Verstelle, w.i.verstelle@student.tue.nl, Eindhoven University of Technology, Eindhoven, The Netherlands; Md Shadab Alam, m.s.alam@tue.nl, Eindhoven University of Technology, Eindhoven, The Netherlands; Pavlo Bazilinskyy, p.bazilinskyy@tue.nl, Eindhoven University of Technology, Eindhoven, The Netherlands.

Generative audio models offer a new way to explore this design space. Recent work with Veo 3 also shows that deep learning based generative models can produce traffic scenes with synchronised audio from text prompts, suggesting their potential for exploring audio visual traffic environments [1]. Text-to-sound systems can rapidly produce many sound concepts and can be prompted with descriptive characteristics such as calm, powerful, smooth, or futuristic. Yet static generated samples are not sufficient for vehicle interaction: vehicle sound must respond continuously to speed, acceleration, and driver input. Dynamic sonification offers a framework for this task because it translates changing data into changing sound [11]. This paper investigates how AI generated samples can be connected to dynamic sonification for personalised EV sound design.

Recent virtual reality work further shows why generated EV sound should be evaluated as an interaction cue rather than only as an audio artefact. Pedestrian crossing studies with EVs emitting synthetic sounds demonstrate that external vehicle sound can be investigated through controlled immersive scenarios in which participants judge an approaching vehicle from the perspective of a vulnerable road user [4]. Complementary psychoacoustic work on synthetic EV sounds shows that exterior sound design involves a trade off between noticeability, informativeness, and annoyance, and that psychoacoustic indicators can provide more meaningful evaluation criteria than sound level alone [2]. These studies motivate a broader view of AI-generated vehicle sound: a useful sound profile should be recognisable and informative for pedestrians, acceptable for occupants, and still open to identity driven customisation.

1.1 Aim of Study

This study examines whether contemporary generative audio models can be used to develop dynamic, personalised, and identity driven sound profiles for future EVs. Rather than treating AI generated audio as a finished vehicle sound, the study investigates how static AI generated samples can act as source material for responsive sonification. Specifically, it explores how such samples can be transformed into dynamic sound that reacts to simulated vehicle speed and throttle input, and how users perceive and accept personalised AI-generated vehicle acoustics. The work contributes by comparing contemporary generative media tools for EV sound prototyping, presenting a functional prototype that combines AI generated loops with a vehicle control interface and real time audio engine, and evaluating user acceptance of personalised AI generated dynamic EV sounds with 20 participants.

2 Method

The study followed the research through design (RtD) process (Figure 1). Instead of starting from a fixed final concept, the project progressed through exploration, prototyping, evaluation, and reflection. The process had four phases: selection of a suitable generative audio model, an informal exploratory demonstration, development of a dynamic sonification prototype, and a structured user acceptance study. The user study was approved by the Eindhoven University of Technology ethical review board, and all participants signed an informed consent form before participating. The study focused on the evaluation of the early stage concept rather than the implementation of production ready vehicles.

Several generative media tools for the generation of EV sounds were considered. Early exploration used video and audio generation systems, including Kling (https://kling.ai), Grok Imagine (https://grok.com/imagine) and Veo (https://deepmind.google/models/veo), accessed through Artlist. These systems produced outputs with some creative potential but were not suitable for the project because the audio was coupled with the generation of videos. This increased cost and generation time while producing visual output that was irrelevant to the intended application. ElevenLabs Text to Sound (https://elevenlabs.io/sound-effects) was selected as the main tool because it generated audio directly, allowed control over prompt influence and duration, exported usable audio files and produced sufficiently

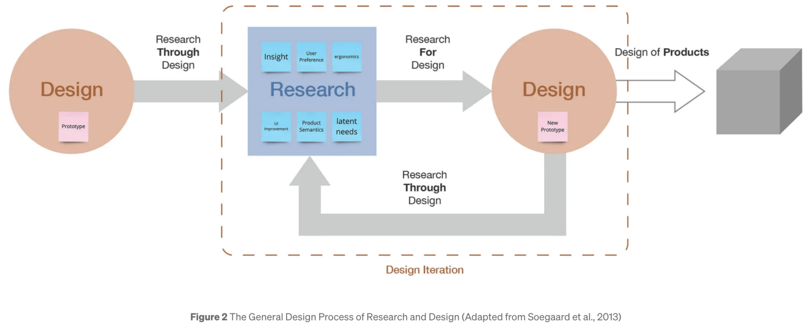

Fig. 1. Research through design workflow used in the study. The figure shows how the project moved from an initial design challenge through iterative research, prototyping, evaluation and reflection before converging on the dynamic sonification demonstrator.

Table 1. Exploratory comparison of generative media tools for EV acoustic prototyping.

Model	Main observation
Kling	Audio matched the generated video, but video was unnecessary and costly for audio prototyping.
Grok Imagine	Output had distorted or mismatched audio and weak control over vehicle sound qualities.
Veo	Generated audio was tied to video and did not provide reliable stand alone acoustic samples.
ElevenLabs Text to Sound	Generated audio directly, supported fast iteration and exported samples that could be edited into loops.

consistent results for iteration. Generation was fast, typically under 10 seconds, which made the model suitable for rapid sound exploration. A limitation was that repeated prompts did not reproduce identical sounds, which complicates exact replication. Table 1 summarises the comparison of the four models.

An informal mid project demonstration was used to gather early reactions to static AI generated samples. A set of 12 ten second samples was generated, including calm atmospheric sounds, synthetic futuristic sounds, and engine like sounds. Students, staff, and visitors listened using Beats Studio3 headphones and verbally shared their first impressions, which were noted by the researcher. No participant count or demographic information was recorded, so the demonstration was treated as design exploration rather than formal evaluation. The feedback showed that static sounds could be emotionally engaging but often lacked a clear connection to vehicle movement. Several attendees stated that the sound should react to speed and driver input, which led to a pivot from static sample evaluation toward dynamic sonification.

The final prototype consists of two connected software components. A Processing interface acts as a simplified vehicle simulator in which users control speed and throttle. A Pure Data patch, edited via plugdata, acts as the audio engine (https://msp.ucsd.edu/software.html,https://processing.org). Communication between the components is implemented with Open Sound Control (OSC), which is commonly used for low latency communication between interactive systems and audio software (https://ccrma.stanford.edu/groups/osc). Within the Pure Data patch, AI generated audio samples are loaded as seamless loops. The speed input changes the sample rate and pitch to create the impression of increasing or decreasing the speed of the vehicle. Throttle input changes load rate, load depth and amplitude to communicate driver intent. The interaction was inspired by the continuous behaviour of a continuously variable transmission, so the sound responds smoothly rather than through discrete gear changes.

The mapping was intentionally simple because the goal was to evaluate the plausibility of the interaction rather than to optimise an automotive sound algorithm. Speed and throttle were treated as independent control dimensions. Speed provided a continuous kinematic cue, while throttle provided a temporary effort cue that could rise even when speed was already high. This distinction was important because participants in the exploratory phase asked for sound that communicated both motion and the intention of the driver. The sound engine therefore did not only make the sample louder at high speed; it also changed pitch, load, and depth so that the vehicle appeared to work harder during acceleration. This made the prototype closer to active sound design than to a fixed warning sound.

Before user testing, the generated samples were edited to reduce discontinuities at the loop point. Samples that contained abrupt transients, strong rhythmic events, or non vehicle like artefacts were avoided because they became distracting when repeated continuously. The selected samples were normalised to a comparable listening level before being loaded into Pure Data. This preparation step was necessary because the generative tools produced sounds with different loudness, spectral balance, and temporal structure. The resulting workflow combined automated generation with manual curation: AI produced a broad sound palette, while the designer selected, cleaned, and mapped the material to vehicle variables. Generative AI was also used as a co-development tool during prototyping. ChatGPT supported understanding Pure Data object logic and troubleshooting, while Claude was used to iterate Processing layouts and Pure Data patch structures. This role is consistent with recent descriptions of vibe coding, where natural language prompts lower the barrier to software prototyping while still requiring human judgement to evaluate and integrate generated output [9]. Figure 2 presents the two main parts of the demonstrator: the Processing based vehicle controller and the Pure Data audio patch.

The final study involved 20 participants. The median age was 24.50 years (SD = 10.46). Fifteen participants identified as male and five identified as female. Nineteen participants had a driver's licence, providing familiarity with conventional vehicle sound and driving behaviour of conventional vehicles. The participants first selected descriptive keywords from a prompt framework. The words were grouped into character, texture, motion, and emotion. Examples included refined, sporty, synthetic, smooth, pulsing, comfortable, and powerful. These keywords were inserted into a predefined Text to Sound prompt to generate four five second sound samples. The participants listened to the four samples using Beats Studio3 headphones and selected the one they preferred. The selected sample was then loaded into the dynamic sonification prototype.

Participants interacted with the virtual vehicle interface by adjusting speed, throttle and sound toggles. This allowed them to experience how the selected AI generated sound responded to simulated vehicle behaviour. After the interaction, the participants completed the Acceptance Scale and answered five open questions. The Acceptance scale [18] evaluates user attitude toward transport technology using nine bipolar items scored from -2 to $+2$. Following the original procedure, items 3, 6, and 8 were mirrored. In the analysis file, the item values were checked against the existing scale columns. The already favourable oriented scoring matched the existing columns exactly (difference = 0.00), whereas applying an additional mirroring step produced a larger difference (difference = 1.22). Therefore, the already favourable oriented item values were used for the reported dimensions. Usefulness was calculated from items 1, 3, 5, 7, and 9. Satisfaction was calculated from items 2, 4, 6, and 8. The open questions addressed desired changes in sound behaviour, comparison with combustion engine sound, interest in customisation, whether vehicles should make sound at all, and acceptance of AI generated sound as part of vehicle identity.

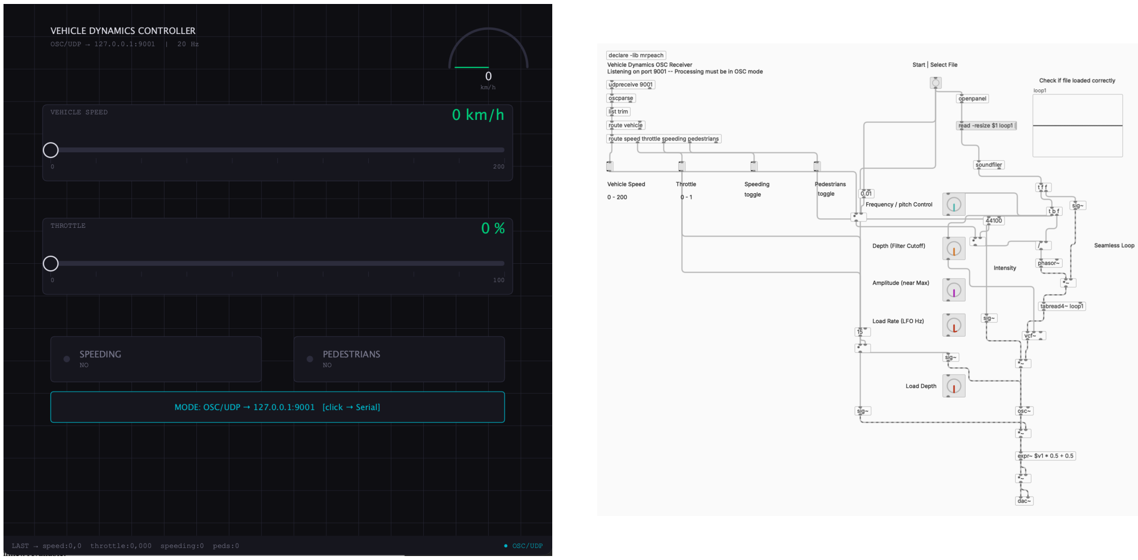

Fig. 2. Dynamic sonification prototype used in the participant evaluation. Left: the Processing based vehicle controller, where participants adjusted speed, throttle input and sound state. Right: the Pure Data audio patch, which received OSC messages and mapped vehicle parameters to pitch, amplitude and load related audio modulation.

Table 2. Summary of user acceptance results and qualitative implications.

Measure or theme	Result	Interpretation
Usefulness	+0.88	Participants saw value in responsive and personalised EV sound.
Satisfaction	+0.62	The experience was moderately pleasant but still required refinement.
Responsiveness	Recurring theme	Throttle mapping, bass response and volume scaling should be improved.
Recognisability	Recurring theme	Some AI sounds were not immediately understood as vehicle sounds.
Personalisation	Strong interest	Several participants liked the idea of choosing a vehicle sound identity.

3 Results

The usefulness score calculated was +0.88 and the satisfaction score calculated was +0.62. Because the scale ranges from −2 to +2, both values indicate a moderately positive attitude. The participants did not treat the prototype as a finished automotive sound system, but generally considered the concept useful and acceptable for further development. Figure 3 shows the distributions of the nine original Acceptance Scale item scores alongside the two computed dimensions, usefulness and satisfaction. Items' scores are shown separately from the computed dimensions to make it clear that usefulness and satisfaction are derived summary measures rather than additional questionnaire items.

Five recurring themes emerged from the open responses. First, participants wanted a stronger coupling between vehicle behaviour and sound output. Suggestions such as more bass, better throttle response, and clearer acceleration

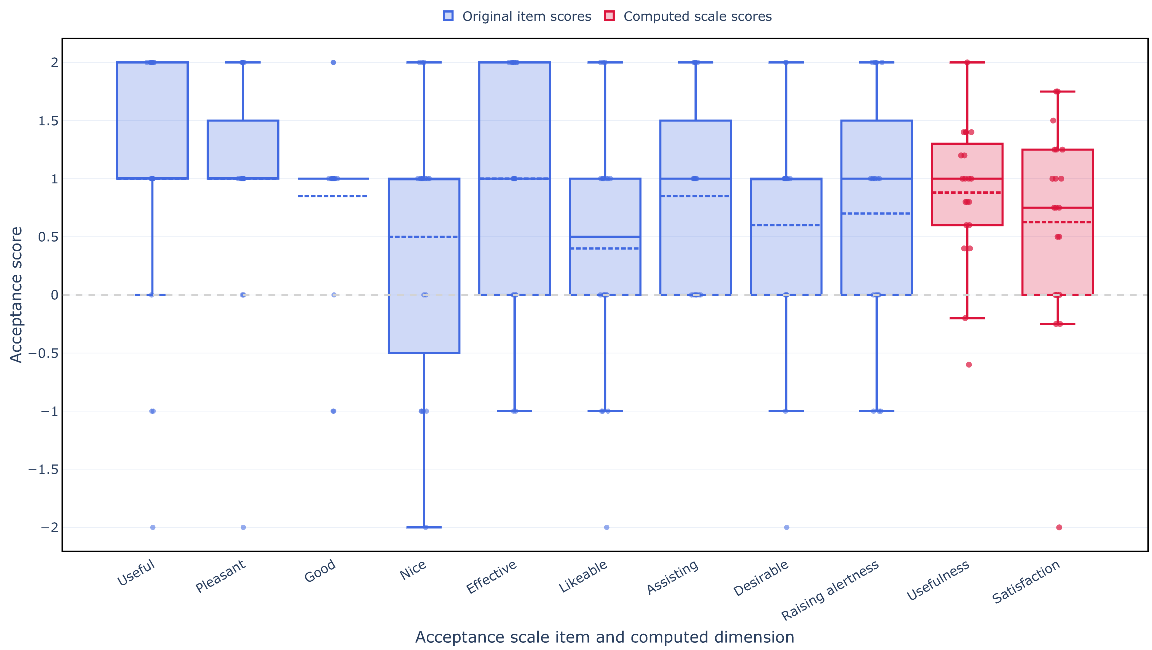

Fig. 3. Acceptance scale results for the dynamic sonification prototype. Blue boxplots show the nine original item scores. Red boxplots show the two acceptance scale dimensions computed from those items: usefulness and satisfaction. Scores are shown on the original scale from −2 to +2, with higher values indicating more favourable judgements.

cues indicate that the dynamic mapping was as important as the generated sample itself. Second, the sound of traditional combustion engines remained a strong reference point. Many participants described combustion engine sounds as more organic, engaging, and connected. However, this did not mean they rejected AI generated EV sound. Rather, they judged the prototype as promising but not yet equivalent to established vehicle acoustics.

Third, personalisation was attractive. Participants often expressed interest in being able to customise the sound of their own vehicle, especially if the sound reflected their identity while remaining safe and recognisable. Fourth, most of the participants believed that vehicles should have some sound for safety, awareness, or driver feedback. Finally, resistance to AI generated sound was low when participants retained meaningful control over the result. Acceptance therefore depended less on whether AI was used and more on whether the output felt useful, intentional, and appropriate.

4 Discussion

The findings suggest that generative AI can support the acoustic design of EVs, but primarily as an exploratory and prototyping tool. The model made it possible to generate diverse sound directions quickly and translate identity words into audio candidates. This is valuable in the early design phases, where designers need to compare many possible directions before committing to detailed sound engineering.

However, the study also shows that AI generated sound alone is insufficient. Dynamic behaviour, recognisability, and contextual appropriateness determine whether the sound is perceived as a vehicle sound. Some samples were pleasant but too abstract, while others were recognisable but closer to combustion imitation. This confirms the tension identified in previous research: future electric vehicle sound should not simply copy combustion engines, but it must still communicate movement, size, and intent [6, 15].

The prototype also illustrates the importance of human control. AI was used to generate samples and accelerate development, but all samples required selection, editing and mapping. The sound became meaningful only when it was connected to the vehicle parameters through designed sonification. For real vehicles, additional constraints would be necessary, including legal AVAS requirements, exterior propagation, brand consistency, pedestrian testing and long term driver comfort.

Overall, the study indicates that AI powered dynamic sonification is a promising direction for personalised EV acoustics. The approach can support early exploration of sound identity, but its value depends on careful mapping between vehicle behaviour and sound behaviour. In this sense, generative AI should be treated as a material for sound designers rather than a replacement for sound design expertise.

A key implication is that personalised EV sound cannot be designed only for the driver. The same sound may be heard by pedestrians, cyclists, and nearby residents, so the design space is shared. Personalisation should therefore operate within boundaries that preserve a common acoustic language for road users. This also highlights a tension between individual customisation and manufacturer brand identity: open customisation can strengthen ownership, but it may conflict with the intended vehicle character or recognisability. The requested identity words could control timbre, texture or emotional colour, while safety critical cues such as approach, acceleration and vehicle presence should remain consistent and testable. Recent VR and psychoacoustic studies of synthetic EV sounds provide a useful methodological direction for this next step because they link sound design to pedestrian perception, crossing behaviour and annoyance [2, 4]. For AI generated sonification, this means that success should be evaluated through both occupant experience and external road user response.

5 Limitations and Future Work

The study has several limitations. The sample size was modest and the group of participants was not balanced in gender or driving experience. Participants interacted with a simulated interface rather than an instrumented vehicle, so the findings should be interpreted as early stage concept feedback rather than evidence of real world deployment readiness. The exploratory model comparison was qualitative and based on available platform access rather than controlled benchmarking. The prototype used pre generated loops, not fully real time generation inside a vehicle. Pedestrian safety was discussed but not directly evaluated, and the study did not measure long term listening comfort, exterior sound propagation or compliance with AVAS regulations.

Future work should test the system in immersive driving simulation and eventually in controlled vehicle trials. Further studies should compare different sonification mappings for speed, throttle, and load, and should evaluate pedestrian detectability alongside driver acceptance. It would also be valuable to study how much personalisation can be allowed before vehicle sounds become confusing, unsafe or inconsistent with brand identity. Finally, future prototypes should investigate real time generation or adaptive recombination of sound layers so that AI generated acoustics can respond more flexibly to driving context while remaining safe, recognisable, and comfortable.

Supplementary Material

In line with current open science practices and recommendations for transparency in automotive user research [7], the authors openly provide these research artefacts to support reproducibility, collaboration and further advancements in the field. The materials used in the study, analysis code and anonymised responses of the participants are available at https://www.dropbox.com/scl/fo/xkycbumjwh0zgyh38kx79/AK4jn5LCIYmV1XoLblyJLQ4?rlkey=47412fjfh845oywzfb2mdcr2q&st=ijm8wbsl. A maintained version of the code is available is https://github.com/Shaadalam9/llm-ev-sound

References

[1] Md Shadab Alam, Zi Wang, Linghan Zhang, and Pavlo Bazilinskyy. 2025. *Exploring Veo 3's capabilities for generating urban traffic scenes in 76 cities worldwide.* Technical Report.

[2] Pavlo Bazilinskyy, Md Shadab Alam, and Roberto Merino-Martínez. 2025. Psychoacoustic assessment of synthetic sounds for electric vehicles in a virtual reality experiment. In *Proceedings of the 11th Convention of the European Acoustics Association, Forum Acusticum / Euronoise 2025.* European Acoustics Association (EAA), Málaga, Spain, 5635–5642. doi:10.61782/fa.2025.0710

[3] Pavlo Bazilinskyy, Pontus Larsson, Emma Johansson, and Joost CF de Winter. 2019. Continuous auditory feedback on the status of adaptive cruise control, lane deviation, and time headway: an acceptable support for truck drivers? *Acoustical science and technology* 40, 6 (2019), 382–390. doi:10.1250/ast.40.382

[4] Pavlo Bazilinskyy, Md Shadab Alam, and Roberto Merino-Martinez. 2025. Pedestrian crossing behaviour in front of electric vehicles emitting synthetic sounds: A virtual reality experiment. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, Vol. 272. Institute of Noise Control Engineering, 2221–2232. doi:10.3397/IN_2025_1076086

[5] Hector Campello-Vicente, Ramon Peral-Orts, Nuria Campillo-Davo, and Emilio Velasco-Sanchez. 2017. The effect of electric vehicles on urban noise maps. *Applied Acoustics* 116 (2017), 59–64. doi:10.1016/j.apacoust.2016.09.018

[6] Elizabeth A Clendinning. 2018. Driving future sounds: imagination, identity and safety in electric vehicle noise design. *Sound Studies* 4, 1 (2018), 61–76. doi:10.1080/20551940.2018.1467664

[7] Patrick Ebel, Pavlo Bazilinskyy, Mark Colley, Courtney Michael Goodridge, Philipp Hock, Christian P Janssen, Hauke Sandhaus, Aravinda Ramakrishnan Srinivasan, and Philipp Wintersberger. 2024. Changing lanes toward open science: Openness and transparency in automotive user research. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications.* 94–105. doi:10.1145/3640792.3675730

[8] European Commission. 2019. Electric and hybrid cars: new rules on noise emitting to protect vulnerable road users. https://single-market-economy.ec.europa.eu/news/electric-and-hybrid-cars-new-rules-noise-emitting-protect-vulnerable-road-users-2019-07-03_en

[9] Margarida Fortes-Ferreira, Md Shadab Alam, and Pavlo Bazilinskyy. 2025. Vibe Coding in Practice: Building a Driving Simulator Without Expert Programming Skills. In *Adjunct Proceedings of the 17th International Conference on Automotive User Interfaces and Interactive Vehicular Applications.* 60–66. doi:10.1145/3744335.3758482

[10] Kiran Govindswamy and Georg Eisele. 2011. Sound character of electric vehicles. In *SAE 2011 Noise and Vibration Conference and Exhibition.* SAE Technical Paper. doi:10.4271/2011-01-1728

[11] Thomas Hermann. 2008. Taxonomy and Definitions for Sonification and Auditory Display. In *Proceedings of the 14th International Conference on Auditory Display (ICAD 2008).* Paris, France, 1–8. https://www.icad.org/Proceedings/2008/Hermann2008.pdf

[12] Louise Humphreys, Sebastiano Giudice, Paul Jennings, Rebecca Cain, Wookeun Song, and Garry Dunne. 2011. The influence of company identity on the perception of vehicle sounds. *Ergonomics* 54, 4 (2011), 338–346. doi:10.1080/00140139.2011.558637

[13] Ming Une Jen and Ming-Hung Lu. 2011. Warning sound of electric vehicle for pedestrian safety. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, Vol. 2011. Institute of Noise Control Engineering, 4646–4653.

[14] Donggun Park, Sunghwan Park, Wonjoon Kim, Ilsun Rhiu, and Myung Hwan Yun. 2019. A comparative study on subjective feeling of engine acceleration sound by automobile types. *International Journal of Industrial Ergonomics* 74 (2019), 102843. doi:10.1016/j.ergon.2019.102843

[15] Jean-François Petiot, Bjørn G Kristensen, and Anja M Maier. 2013. How should an electric vehicle sound? User and expert perception. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 55928. American Society of Mechanical Engineers, V005T06A028. doi:10.1115/DETC2013-12535

[16] DJ Swart and A Bekker. 2019. The relationship between consumer satisfaction and psychoacoustics of electric vehicle signature sound. *Applied Acoustics* 145 (2019), 167–175. doi:10.1016/j.apacoust.2018.09.019

[17] Daniël Johannes Swart, Anriëtte Bekker, and Jörg Bienert. 2018. The subjective dimensions of sound quality of standard production electric vehicles. *Applied Acoustics* 129 (2018), 354–364. doi:10.1016/j.apacoust.2017.08.012

[18] Jinke D Van Der Laan, Adriaan Heino, and Dick De Waard. 1997. A simple procedure for the assessment of acceptance of advanced transport telematics. *Transportation Research Part C: Emerging Technologies* 5, 1 (1997), 1–10. doi:10.1016/S0968-090X(96)00025-3