EXTENDED ABSTRACT OF THE DOCTORAL DISSERTATION

Deep Learning-Based Active Control of Noise Transmission through Encapsulated Structures

Candidate: Alkahf Aboutiman
Supervisors: Prof. Dr. Hamid Reza Karimi (Politecnico di Milano), Prof. Francesco Ripamonti (Politecnico di Milano), Prof. dr hab. inż. Marek Pawełczyk (Silesian University of Technology), Dr inż. Stanisław Wrona (Silesian University of Technology)
Doctoral cycle 38 (2023–2026) — Joint (double-degree) doctorate between Politecnico di Milano and the Silesian University of Technology

---

1. Scope of the Dissertation

This dissertation addresses active noise control (ANC) in complex vibro-acoustic environments, with a particular focus on encapsulated structures — enclosures such as automotive cabins, aircraft fuselages, railway compartments, and industrial housings in which an internal cavity must be acoustically isolated from an external environment. Encapsulated structures are characterised by strong modal behaviour, sensitivity to openings and boundary conditions, and often exhibit non-minimum phase (NMP) dynamics that fundamentally limit the performance of classical adaptive ANC algorithms. In practice, such enclosures frequently include functional openings — for ventilation, cabling, thermal regulation, or airflow — that further modify the acoustic transfer paths and introduce additional nonlinear and diffraction effects that classical narrowband control models do not capture.

The scope of the work spans three complementary layers of investigation:

1. The physical and mathematical modelling of encapsulated vibro-acoustic systems, obtained through finite element analysis and experimental system identification, and the design of model-based control strategies — specifically Kalman filtering — able to cope with non-minimum phase dynamics that classical adaptive filters cannot compensate.
2. Hybrid deep learning architectures that combine lightweight convolutional neural networks with classical adaptive filtering to handle nonlinearities introduced by transducers and enclosure openings, evaluated through both objective metrics (dB attenuation, noise reduction spectra) and subjective, human-centred perceptual metrics (annoyance, affective quality, loudness).

3. Deep transfer reinforcement learning (DTRL) controllers, built on the proposed FxMDNoiseNet architecture, that adapt online to the highly non-stationary, time-varying acoustic conditions found inside a moving vehicle, including secondary-path mismatch caused by varying vehicle occupancy and cabin configuration.

Beyond the control algorithms themselves, the dissertation places explicit emphasis on computational feasibility for embedded, real-time deployment — a requirement that recurs throughout all three research threads and that motivates the systematic use of lightweight network architectures (in the order of a few thousand trainable parameters), fixed-point-compatible operations, and low sampling rates suitable for automotive digital signal processors.

The research was carried out as a double-degree doctorate between Politecnico di Milano (Department of Mechanical Engineering) and the Silesian University of Technology, within the framework of the Horizon Europe Marie Skłodowska-Curie project IN-NOVA (grant agreement no. 101073037). Research secondments were completed at University College London (subjective ANC perception, with Prof. Jian Kang and Dr. Francesco Aletta), at a partner institution in Poland, and at Müller-BBM Active Sound Technology GmbH, Germany, where industrial automotive noise control practices informed the experimental design of the final research thread.

2. Dissertation Objectives and Thesis Statement

The overarching objective of the dissertation is to improve the robustness, adaptability, and perceptual effectiveness of active noise control systems operating in complex, nonlinear, and time-varying vibro-acoustic environments — with vehicle cabin noise control as the principal application scenario, given its combination of strong modal coupling, transducer nonlinearity, and continuously varying operating conditions (engine load, road surface, vehicle speed, occupancy).

The central thesis defended in this work is that classical, purely adaptive ANC algorithms (e.g., FxNLMS) are structurally limited when confronted with non-minimum phase dynamics, transducer nonlinearities, and non-stationary acoustic conditions typical of encapsulated structures, and that these limitations can be systematically overcome by combining physics-based modelling with learning-based control — specifically state estimation via Kalman filtering, hybrid deep-learning filter selection and adaptation (SFANC/GFANC families), and deep transfer reinforcement learning (DDPG/TD3) — while keeping the resulting controllers computationally lightweight enough for real-time, embedded deployment.

A secondary but methodologically important claim of the thesis is that objective attenuation metrics (expressed in dB) are necessary but not sufficient to assess the practical value of an ANC system, and that human perceptual response — annoyance, affective quality, and perceived loudness — must be evaluated directly through controlled listening experiments in order to validate that algorithmic improvements translate into genuine gains in occupant comfort.

The specific objectives derived from this thesis are:

- To characterise, through finite-element modelling and system identification, the non-minimum phase behaviour of an encapsulated vibro-acoustic structure and to quantify its impact on classical adaptive ANC performance.
- To design and validate a Kalman filter-based ANC strategy that improves robustness and attenuation under NMP dynamics relative to FxNLMS, while quantifying the associated computational overhead.
- To develop hybrid deep learning ANC frameworks (SFANC-NLMS, GFANC-THFxNLMS) that remain computationally efficient while improving robustness to loudspeaker nonlinearities and enclosure openings.
- To assess, through controlled listening experiments with human participants, how algorithmic ANC performance translates into perceived annoyance, affective quality, and perceived loudness, using validated psychoacoustic protocols (ISO 15666) and appropriate multivariate statistical analysis.
- To design a lightweight neural ANC controller (FxMDNoiseNet) and to integrate it with deep reinforcement learning (DDPG, TD3) and transfer learning to obtain a deep transfer reinforcement learning controller (FxMDNoiseNet-DTRL), validated under realistic, experimentally measured automotive conditions, including secondary-path mismatch and linear time-varying (LTV) dynamics.
- To evaluate, throughout all three research threads, the trade-off between control performance and computational cost, expressed in terms of trainable parameters, multiply-accumulate operations (MACs), and compatibility with embedded digital signal processors used in automotive applications.

3. Description of the Research Subject

3.1 Active Noise Control in Non-Minimum Phase Encapsulated Systems

The first research thread investigates ANC in an encapsulated vibro-acoustic structure — a laboratory "noise box" representative of an automotive cabin — modelled by the finite

element method (FEM) to capture the coupled structural-acoustic dynamics. The structural domain is discretised using shell elements based on Reissner–Mindlin plate theory with MITC (Mixed Interpolation of Tensorial Components) formulation, while the acoustic cavity is modelled within a unified vibro-acoustic finite element framework. Structural damping is introduced through an isotropic loss factor of $\eta = 0.003$, consistent with typical values reported for the plate material considered. The resulting frequency response functions (FRFs) between the primary source, the secondary (control) source, and the error microphone are computed and subsequently validated against experimental measurements obtained using calibrated instrumentation, including a microphone with a sensitivity of 37.37 mV/Pa calibrated with a Brüel & Kjær 4231 acoustic calibrator at 1 kHz and 94 dB.

Pole-zero analysis of the identified primary, secondary, and reference paths reveals zeros located outside the unit circle in the discrete-time domain — i.e., non-minimum phase behaviour. This NMP characteristic introduces phase delays and causality constraints that are shown, both analytically and experimentally, to fundamentally cap the achievable attenuation of the classical Filtered-x Normalised Least Mean Squares (FxNLMS) algorithm, whose optimal control filter formally requires an unstable inversion of the secondary path when the latter is non-minimum phase.

To overcome this limitation, a discrete-time state-space identification of the acoustic paths is performed, and a Kalman filter-based ANC strategy is proposed, in which the filter's improved state estimation compensates for the NMP-induced uncertainty without requiring explicit inversion of the secondary path. All simulations for this study were performed at a sampling frequency of 8 kHz, consistent with the bandwidth of interest for automotive low-frequency noise. The computational complexity of both approaches is analysed in terms of multiplication and addition floating-point operations (FLOPs) per sample, confirming that the Kalman filter-based approach, while more demanding than FxNLMS, remains compatible with real-time implementation on embedded automotive hardware, at the cost of a moderate but justified increase in computational load.

3.2 Hybrid Deep Learning ANC: Objective and Subjective Evaluation

The second research thread introduces hybrid architectures that combine a lightweight one-dimensional convolutional neural network (1D CNN) with classical adaptive filtering. In the SFANC-NLMS framework, the CNN selects, frame by frame, the most suitable pre-trained control filter from a fixed-filter bank inspired by Selective Fixed-filter Active Noise Control (SFANC), while an NLMS stage refines the selected filter's coefficients online. This hybrid design combines the low latency of filter selection — which allows the controller to

react instantaneously to changes in the noise spectrum — with the stability and fine-tuning capability of adaptive refinement.

The CNN classifier was trained on a synthetic dataset comprising 84,000 broadband random noise samples, split into 80,000 training samples, 2,000 validation samples, and 2,000 test samples. Training was performed using the Adam optimizer with a learning rate of 10^{-4}, a batch size of 250, over 30 epochs, with L2 regularisation applied to all trainable parameters to reduce overfitting and preserve a lightweight architecture suitable for embedded implementation. The resulting model achieves high classification accuracy while significantly reducing the number of trainable parameters compared to larger reference architectures, confirming its suitability for real-time ANC applications on embedded hardware. During operation, incoming noise is sampled at 16 kHz, with the CNN analysing one-second frames and relaying the selected filter index to the real-time controller, which then continuously refines the filter's parameters at the sampling rate; validation against recordings from the SONYC urban-sound dataset confirmed generalisation beyond the synthetic training distribution.

Beyond objective attenuation, this thread investigates how ANC performance is perceived by human listeners — a dimension often neglected in purely algorithmic ANC research. Using a noise-box mock-up of a vehicle interior equipped with a Head and Torso Simulator (HATS) and calibrated headphone playback, a total of 35 participants (16 men and 19 women, aged 23 to 62, mean age approximately 31) were recruited and exposed to combinations of noise type (motorcycle, street, train), sound pressure level (55, 65, 72 dB(A)), and control condition (no control, NLMS, SFANC-NLMS). Two participants were excluded from the final analysis due to unreliable response patterns (Spearman correlation coefficient of 0.8 combined with a p-value exceeding 20% on repeated-stimulus consistency checks). Prior to the main experiment, participants completed demographic questionnaires and a pre-test assessment including the Weinstein Individual Noise Sensitivity (INS) scale, to account for inter-individual variability and potential confounding factors such as stress and mood.

Participants rated perceived annoyance (PAY) using the standardised ISO 15666 single-item scale, affective quality (PAQ) using a two-dimensional affective grid (arousal/valence), and perceived loudness (PLN) via magnitude estimation. Statistical analysis was performed using a three-way repeated-measures ANOVA (3WR-ANOVA) for PAY and PLN, with noise type, noise level, and ANC condition as within-subject factors; for PAQ, a three-way repeated-measures permutational multivariate ANOVA (3WR-PERMANOVA) was used instead, due to violations of multivariate normality confirmed by Mardia's test. Post-hoc pairwise comparisons were applied where significant main effects or interactions were

detected. SFANC-NLMS was found to consistently outperform NLMS on all three perceptual dimensions across noise types and levels, and the relationship between objective attenuation and perceived comfort was shown to be non-trivial — in particular, low-frequency content (20–500 Hz) emerged as a dominant factor in perceived discomfort even when overall broadband attenuation was similar between conditions — motivating human-centred evaluation as a necessary complement to dB-based metrics in ANC design.

Building on these results, the Generative Fixed-filter ANC combined with a Transducer/Harmonic-aware FxNLMS stage (GFANC-THFxNLMS) was developed to address nonlinear distortions introduced by the secondary loudspeaker, particularly relevant for encapsulated structures with openings. The experimental configuration used for this study consisted of an enclosure with a circular opening of 0.05 m radius, three microphones positioned inside the enclosure, and a secondary loudspeaker positioned centrally behind the opening. System identification confirmed that the primary path exhibits strong resonant behaviour below 350 Hz due to coupled dynamics between the vibrating plate and acoustic leakage through the opening, while the secondary path, driven by the internal loudspeaker, shows a smoother, more attenuated response consistent with more direct acoustic propagation. A 1D CNN dynamically selects and linearly combines sub-filters decomposed from a broadband control filter, while the THFxNLMS stage compensates for loudspeaker nonlinearities in real time using a gradient-scaling mechanism that adapts the effective step size according to the instantaneous noise level, improving stability and convergence under non-stationary conditions. The final trained CNN model for this study contains 6.74K parameters and requires approximately 19.33 million multiply-accumulate operations (MACs) per one-second frame at a 3 kHz sampling rate — a computational footprint low enough to allow the model to be fully stored in on-chip memory on embedded digital signal processors such as the Analog Devices ADSP-SC598, minimising inference latency.

3.3 Deep Transfer Reinforcement Learning ANC for Vehicle Interiors

The third and most advanced research thread addresses the highly non-stationary, nonlinear acoustic conditions encountered in a moving vehicle, where engine load, road excitation, vehicle speed, and cabin occupancy continuously alter the acoustic transfer paths. A lightweight neural network controller, FxMDNoiseNet, is proposed within a filtered-x architecture. Its core consists of stacked Atrous Scaled Convolution (ASC) modules with progressively increasing dilation factors, combined with causal and pointwise one-dimensional convolutions and residual/skip connections, enabling efficient multi-scale extraction of time-domain features from raw audio while keeping the total

parameter count extremely low — a design choice directly motivated by the real-time and embedded-hardware constraints identified throughout the thesis. Training data for this stage were derived from impulse responses recorded in a real vehicle equipped with door-mounted woofers and a factory-installed subwoofer located in the trunk, with signals originally recorded at 48 kHz and subsequently downsampled to 3 kHz to reduce computational complexity, consistent with the low-frequency content of interest for cabin noise control. Optimisation was performed using the Adam optimizer with a learning rate of 2×10^{-4} and a weight decay of 10^{-4}, combined with a step learning-rate scheduler (decay factor 0.5 every 5 epochs) and Xavier weight initialisation to ensure stable convergence, using the same 80,000/2,000/2,000 training/validation/test split methodology applied in the preceding research thread.

To improve adaptability under time-varying conditions, a deep reinforcement learning (RL) layer is integrated on top of FxMDNoiseNet. The ANC problem is formulated as a sequential decision-making (Markov decision) task, in which at each discrete time step n an agent observes a state s_n derived from the estimated disturbance and outputs a continuous control action $a_n = \pi\phi(s_n)$, applied through the secondary path to generate the anti-noise signal; the resulting residual error e(n) is used to compute the reward $r(n) = -|e(n)|$, so that maximising cumulative reward corresponds to minimising residual noise energy over time. Two continuous-control, off-policy actor–critic RL algorithms are investigated: Deep Deterministic Policy Gradient (DDPG) and its more stable successor, Twin Delayed DDPG (TD3). TD3 addresses two well-known limitations of DDPG — overestimation bias in the value function and error accumulation from bootstrapped critic targets — through three complementary mechanisms: (i) Clipped Double Q-learning, in which two independent critics Qθ1 and Qθ2 are trained and the minimum of their estimates is used to form the Bellman target, mitigating value overestimation; (ii) target policy smoothing, in which clipped Gaussian noise is added to the target action to regularise the value estimate against sharp, error-prone peaks in the Q-function; and (iii) delayed policy updates, in which the critics are updated at every environment step while the actor (and target networks) are updated only every d iterations, allowing the value estimate to stabilise before it is used to guide policy improvement.

In both DDPG and TD3 configurations, the RL policy network is initialised with the pre-trained FxMDNoiseNet weights rather than random initialisation — a transfer-learning strategy that markedly accelerates convergence, improves training stability, and reduces the amount of interaction data required with the (costly to simulate) real acoustic environment, yielding the proposed FxMDNoiseNet-DTRL framework. Exploration during training is guided by an adaptive, envelope-scaled Gaussian noise process, whose standard deviation is modulated by the instantaneous envelope of the disturbance signal,

allowing more aggressive exploration during high-energy transients and more conservative behaviour during quiet periods. The critic networks are trained to minimise the Bellman error with $\ell 2$ regularisation, drawing mini-batches from a replay buffer populated during interaction with the (simulated and, subsequently, experimentally measured) acoustic environment.

The framework was evaluated using experimentally measured impulse responses recorded in a real vehicle, under both linear time-invariant (LTI) and linear time-varying (LTV) conditions, and including deliberate secondary-path mismatch — training the controller using the impulse response of the empty vehicle while evaluating it using the impulse response measured with two occupants present — to explicitly test robustness to modelling uncertainty. These conditions closely approximate in-service automotive operation, where the true secondary path is rarely known exactly and varies continuously with vehicle load and cabin configuration.

4. Main Results and Conclusions

Across the three research threads, the results consistently show that combining physics-based modelling with learning-based control yields substantial improvements over classical adaptive ANC, while keeping the resulting controllers compact enough for real-time use on embedded automotive hardware.

Summary of principal quantitative results:

Study	**Compared methods**	**Key quantitative outcome**
Paper A — NMP encapsulated structure	Kalman filter-based ANC vs. FxNLMS	16.92 dB vs. 2.72 dB attenuation (0–450 Hz band)
Paper B — subjective perception	SFANC-NLMS vs. NLMS (N = 33 valid participants)	Significant improvement in PAY, PAQ, and PLN across all tested conditions (3WR-ANOVA / 3WR-PERMANOVA, $p < 0.05$)
Paper C — GFANC-THFxNLMS	GFANC-THFxNLMS vs. GFANC vs. THFxNLMS	9.49 dB vs. 8.01 dB vs. 6.74 dB attenuation
Papers D–E — FxMDNoiseNet-DTRL	FxMDNoiseNet-DTRL (DDPG/TD3) vs. FxMDNoiseNet vs. FxNLMS	Improved NMSE and PSD attenuation under both LTI and LTV conditions with secondary-path mismatch

For the non-minimum phase encapsulated structure (Paper A), the Kalman filter-based ANC strategy achieved up to 16.92 dB of attenuation in the 0–450 Hz band, compared to only 2.72 dB for FxNLMS. Time-domain residual error signals and spectrogram comparisons confirm that FxNLMS leaves substantial residual energy concentrated around the resonance and anti-resonance frequencies of the coupled structural-acoustic system, whereas the Kalman filter-based approach substantially suppresses this residual energy across the full analysed bandwidth, at the cost of a moderate increase in computational load that remains compatible with real-time implementation. These results confirm that accurate state estimation is essential when the acoustic paths exhibit non-minimum phase dynamics that classical adaptive filters, by construction, cannot compensate.

The subjective study (Paper B) demonstrated that SFANC-NLMS is consistently preferred over NLMS in terms of perceived annoyance and perceived loudness, with statistically significant effects confirmed by three-way repeated-measures ANOVA (for PAY and PLN) and permutational multivariate ANOVA (for PAQ). While higher sound pressure levels increase perceived annoyance overall, as expected, they do not reverse the relative ranking of the ANC algorithms — SFANC-NLMS remains the preferred condition across all tested noise types and levels. Low-frequency energy content, rather than broadband attenuation alone, emerged as a dominant predictor of perceived discomfort, reinforcing the value of the hybrid, learning-based approach beyond purely objective, dB-based performance metrics, and validating the methodological choice to combine algorithmic and psychoacoustic evaluation throughout the dissertation.

The hybrid GFANC-THFxNLMS framework (Paper C) achieved 9.49 dB of noise reduction, outperforming both GFANC alone (8.01 dB) and THFxNLMS alone (6.74 dB), in an encapsulated structure with an acoustic opening — a configuration in which loudspeaker nonlinearities and diffraction effects are particularly pronounced. These results demonstrate that combining CNN-based filter selection with nonlinearity-aware adaptive filtering improves robustness to loudspeaker nonlinearities without sacrificing computational efficiency, confirming the hypothesis that hybrid architectures can outperform either pure fixed-filter selection or pure adaptive filtering taken in isolation.

Finally, the FxMDNoiseNet-DTRL framework (Papers D and E) outperformed both the standalone deep-learning controller (FxMDNoiseNet) and classical FxNLMS under realistic automotive conditions, including secondary-path mismatch and LTV dynamics — the most demanding test scenario considered in the thesis, combining engine and road noise components. In the frequency domain, FxMDNoiseNet-DTRL consistently achieved the strongest attenuation in the low-frequency band (50–250 Hz), where engine noise energy is dominant, confirming its robustness in targeting the most energetic acoustic components

despite the modelling errors deliberately introduced through LTV conditions and secondary-path mismatch. The transfer-learning initialisation strategy — seeding the RL policy with pre-trained FxMDNoiseNet weights rather than random weights — was found to be a key factor in achieving stable convergence within a practical number of training interactions, a result with direct implications for the feasibility of deploying RL-based ANC controllers in real automotive development cycles, where interaction with the true acoustic environment is costly.

Overall, the dissertation demonstrates a coherent progression — from physics-based modelling of non-minimum phase dynamics, through hybrid deep-learning filtering validated with human-perception experiments, to deep transfer reinforcement learning — that systematically extends the operating envelope of ANC systems toward realistic, nonlinear, time-varying automotive environments, while respecting the computational constraints of embedded, real-time deployment at every stage. Across all three threads, the number of trainable parameters was kept in the range of a few thousand, and the computational cost was consistently expressed in terms compatible with existing embedded automotive digital signal processors, directly supporting the practical relevance of the proposed methods beyond simulation studies.

The dissertation identifies three main directions for future work. First, moving from a sample-by-sample control formulation to a block-based one would allow the controller to better exploit longer temporal dependencies in the disturbance signal and to reduce computational overhead through batched processing, at the cost of introducing additional algorithmic latency that would need to be carefully characterised. Second, exploring more advanced continuous-control reinforcement learning algorithms — such as Distributed Distributional Deep Deterministic Policy Gradient (D4PG) or other recent off-policy methods incorporating distributional value estimation — could further improve convergence speed and robustness to modelling uncertainty beyond what is achieved with TD3. Third, implementing the proposed FxMDNoiseNet-DTRL framework on embedded real-time hardware, rather than validating it offline using pre-recorded impulse responses, would allow the approach to be assessed under genuine latency, sensor-noise, and hardware-nonlinearity constraints, closing the gap between the present simulation-based validation and full in-vehicle deployment.

5. References to the Author's Main Publications

The dissertation is presented as a collection of five appended, thematically-related papers (Papers A–E), listed below. No patents resulted from this research; all outcomes were disseminated through peer-reviewed publications and conference proceedings.

Appended papers forming the basis of the dissertation:

Paper A. A. Aboutiman, R. Shams, H.R. Karimi, F. Ripamonti, M. Pawełczyk, "Active noise control in encapsulated structures with non-minimum phase characteristics using a Kalman filter approach," *Journal of Sound and Vibration*, 615 (2025) 119187.

Paper B. A. Aboutiman, Z. Rachman, T. Oberman, F. Aletta, J. Kang, H.R. Karimi, F. Ripamonti, "Subjective perception analysis of active noise control algorithms in an encapsulated structure: An experimental study," *Applied Acoustics*, 239 (2025) 110823.

Paper C. A. Aboutiman, K.S.A. Maamoun, H.R. Karimi, F. Ripamonti, "Hybrid deep learning–based active noise control for encapsulated structures with openings," *Expert Systems With Applications*, 309 (2026) 131247.

Paper D. A. Aboutiman, J. Kirchhof, F. Ripamonti, "Transfer Learning-Enhanced DDPG for Adaptive Active Noise Control in Vehicle Interiors," 39th Conference on Neural Information Processing Systems (NeurIPS 2026), under review.

Paper E. A. Aboutiman, J. Kirchhof, H.R. Karimi, F. Ripamonti, "Deep transfer reinforcement learning for active noise control in vehicle interiors," *Nature Machine Intelligence*, submitted, 2026.

Additional publications by the candidate:

- K.S.A. Maamoun, A. Aboutiman, A.M. Ali, "Synergistic use of inertial actuators for advanced noise control in transparent patched vibrating plates," *Measurement*, 2025.
- C. Liang, A. Aboutiman, J. Kirchhof, F. Ripamonti, H.R. Karimi, "Experimental Evaluation of Simultaneous Perturbation Stochastic Approximation Algorithms for Active Engine Noise Control in Automotive Applications," *Applied Acoustics*.
- A. Aboutiman, H.R. Karimi, F. Ripamonti, "A hybrid learning-based active noise control algorithm for encapsulated structures with linear time-varying characteristics," 31st International Congress on Sound and Vibration (ICSV31), 2025 — **Best Paper Award (1st Classified)**.
- A. Aboutiman, H.R. Karimi, F. Ripamonti, "Hybrid Learning-based Active Noise Control in Encapsulated Structures," International Conference on Machine Learning and Artificial Intelligence for Audio, 2025.
- A. Aboutiman, H.R. Karimi, F. Ripamonti, "Frequency-Based Modeling and Active Noise Control Design for Encapsulated Structures," ICSV30, 2024.

6. Summary of the Candidate's Own Contribution

The candidate, Alkahf Aboutiman, is the first and corresponding author of all five appended papers (A–E) and conceived, in collaboration with the supervisors, the overall research programme linking physics-based modelling, hybrid deep learning, and deep transfer reinforcement learning for active noise control.

Specifically, the candidate's own contribution comprises:

- Development of the finite element and state-space models of the encapsulated vibro-acoustic structure, identification of its non-minimum phase characteristics through pole-zero analysis of experimentally measured frequency response functions, and design, implementation, and experimental validation of the Kalman filter-based ANC strategy, including the computational complexity analysis supporting its real-time feasibility (Paper A).
- Design and execution of the subjective listening experiments — including participant recruitment, experimental protocol design following ISO 15666 and validated affective-grid methodology, calibrated stimulus preparation using a Head and Torso Simulator, and the application of three-way repeated-measures ANOVA and PERMANOVA statistical analysis — carried out in collaboration with University College London during the research secondment, and analysis of the resulting perceptual data on annoyance, affective quality, and loudness (Paper B).
- Design, implementation, and experimental validation of the hybrid SFANC-NLMS and GFANC-THFxNLMS algorithms, including the 1D CNN filter-selection and sub-filter combination module, the nonlinearity-aware THFxNLMS adaptive stage with adaptive gradient scaling, and the corresponding training pipeline and embedded-feasibility analysis (Paper C).
- Design of the lightweight FxMDNoiseNet architecture, based on causal, pointwise, and atrous scaled convolutions, and its integration with DDPG- and TD3-based deep reinforcement learning and transfer learning, together with the design and execution of the experimental campaign in a real vehicle — including impulse-response measurement under LTI and LTV conditions with secondary-path mismatch — used to validate the resulting FxMDNoiseNet-DTRL framework (Papers D and E).
- Co-supervision of two master's students (Maxime Zaman and Reza Shams) contributing to related aspects of the research, and collaborative work with Dr. Chao Liang, Dr. Zulfi Rachman, and Dr. Khaled Said Ahmed Maamoun on connected

publications extending the core methods of the dissertation to related vibro-acoustic control problems.

Where publications involved multiple authors, the individual percentage and substantive contribution of each co-author is provided separately, as required by §2(4) of the Regulations, in the corresponding co-authorship statements attached to the doctoral application. In accordance with the same provision, where it was not possible to obtain individual statements from all co-authors, a statement was instead provided by the corresponding author (the candidate) or the relevant project/team leader, as applicable.