Control-Oriented Input Design with Simplified Data-Informed Models: Applications to Solid-State Batteries and Human–Autonomy Driving

by

Patrick Alan Linford

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in The University of Michigan
2026

Doctoral Committee:

Professor Kira Barton, Chair
Associate Professor Tulga Ersal
Assistant Professor Maani Ghaffari
Associate Professor Chris Vermillion

Patrick Alan Linford
plinford@umich.edu
ORCID iD: 0009-0001-7255-0293

© Patrick Alan Linford 2026

DEDICATION

This dissertation is dedicated to the people, mentors, colleagues, friends, and family members who supported the long arc of this work. Especially my wife, Melanie, who shouldered the burden of the family during this late adventure, and my children, Frieda and Peter, who had to deal with a father always gone and sleepy. Thank you; I love you all!

ACKNOWLEDGEMENTS

This work was funded by the Advanced Civil Schooling fellowship in the US Army, as well as funding from the Ground Vehicle Systems Center, US Army Development Command. Formulation and concurrent work from my collaborators: Kieth Ng, Aleksandra Dudek, and Zihan Yu, as well as Ian Holbrook who was instrumental in setting up the simulator.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF FIGURES viii

LIST OF TABLES xiv

LIST OF ACRONYMS xvi

LIST OF SYMBOLS xvii

ABSTRACT xix

CHAPTER

1 Introduction 1

1.1 Control-Oriented Modeling 1
1.2 Application Context 2
1.3 Dissertation Problem Statement 2
1.4 Dissertation Thesis 3
1.5 Research Questions 3
1.5.1 Chapter-level contributions 4
1.6 Claim Boundaries 4
1.7 Dissertation Organization 5

2 Motivation and Background 6

2.1 Why These Problems Matter 6
2.2 Purpose-Built Models in Control Engineering 7
2.2.1 Model fidelity is relative to intended use 7
2.2.2 Model reduction, system identification, and data-driven representations 7
2.3 Position of This Dissertation 10
2.4 Primer A: Control-Oriented Modeling of Solid-State Battery Interface Evolution 10
2.4.1 Promise and limitation of solid-state batteries 10
2.4.2 Established low-order battery models and their control uses 11

2.4.3 Why existing simplified strategies are insufficient for this study . . . 11
2.4.4 Collaborative data and dissertation contribution 12
2.5 Primer B: Human–Autonomy Intervention and Control Adaptation 12
2.5.1 Conditional automation and intervention 12
2.5.2 Existing control approaches that respond to the human 13
2.5.3 Known-route learning as event-based disturbance mitigation 13
2.5.4 Transfer beyond the learned route 14
2.6 Comparative Structure of the Three Case Studies 14
2.7 Application-Specific Research Gaps 16

3 Reduced-Order Modeling of Solid-State Battery Interface Evolution for Feedforward Discharge Control 18

3.1 Introduction 18
3.2 Background and Related Work 20
3.2.1 Interface Evolution in Solid-State Batteries 20
3.2.2 Control-Oriented Battery Models and Their Limits 20
3.2.3 From Stability Characterization to Transition-Region Design 21
3.2.4 Discharge Input Shaping 22
3.3 Methodology 22
3.3.1 Data Sources and Scope 22
3.3.2 Areal Capacity, Overpotential, and Interface-Health Metrics 23
3.3.3 Reduced-Order Response Families 24
3.3.4 Transition-Region Identification 25
3.3.5 Reference Library and Transition Interpolation 26
3.3.6 Simulation-Based Evaluation of Interface Health 28
3.3.7 Feedforward Discharge Profile Design 28
3.4 Results 34
3.4.1 Aggressive and Conservative Profile Families 34
3.4.2 Two-Hour Discharge Performance 35
3.4.3 Interpretation of Capacity and Interface-Health Tradeoffs 36
3.5 Discussion 37
3.5.1 What the Reduced-Order Model Captures 37
3.5.2 Role of Physical and Simulation Data 38
3.5.3 Limitations 38
3.5.4 Implications for Control Design 39
3.6 Chapter Summary 40

4 Learning-Based Disturbance Rejection for Human Takeover on a Known Route 41

4.1 Chapter Overview 41
4.2 Related Work 42
4.2.1 Human Intervention in Conditionally Automated Driving 42
4.2.2 From Explanatory Human Models to Control-Oriented Disturbance Rejection 43
4.2.3 Iterative Learning Control for Repeated Tasks 43

4.2.4 Route-Local Speed Adaptation as a Simplified Model 44
4.3 Methods 45
4.3.1 Problem Formulation 45
4.3.2 Backward-Reachable Speed Formulation 47
4.3.3 Learning Update for the Assumed Friction Map 49
4.3.4 Simulation and Human-Subject Testing Protocol 49
4.3.5 Analysis Plan and Sensitivity Checks 51
4.4 Results 52
4.4.1 Handling-Course Segmentation 52
4.4.2 Participant Groups and Sensitivity Analysis 52
4.4.3 Lap-Time and Continuous-Task Performance 56
4.4.4 Learned Coefficient-of-Friction Profile and Route Location of Takeover 56
4.4.5 Takeover Evolution Across Iterations 68
4.4.6 Lateral Error and Vehicle-Performance Context 68
4.4.7 Summary of Known-Route Learning Results 75
4.5 Discussion and Chapter Conclusion 78
4.5.1 Interpretation of the Learning Result 78
4.5.2 Participant Heterogeneity and Sensitivity of the Cohort Trend . . . 78
4.5.3 Role and Limits of the Assumed Friction Profile 79
4.5.4 Performance Tradeoffs 80
4.5.5 Implications for Experimental Design 80
4.5.6 Limitations 81
4.5.7 Chapter Conclusion 81

5 Transfer of Learned Human-Intervention Profiles to Novel Routes 83

5.1 Introduction 83
5.2 Related Work and Motivation 84
5.3 Methods 86
5.3.1 Problem Formulation 86
5.3.2 Route-Context Feature Representation 86
5.3.3 Normalized Feature Space 87
5.3.4 Weighted Three-Nearest-Neighbor Transfer 89
5.3.5 Single-Nearest-Neighbor Transfer 89
5.3.6 Mapping the Transferred Profile onto the Target Course 90
5.3.7 Evaluation Conditions and Metrics 91
5.4 Results 91
5.4.1 Baseline and Transferred-Profile Summary 93
5.4.2 Transferred Assumed-Friction Profiles 96
5.4.3 Intervention Changes Across the Aggressive Course 97
5.4.4 Driving Performance and Lateral Error 99
5.4.5 Transfer Diagnostics and Expectation-Mismatch Locations 101
5.5 Discussion and Limitations 102
5.6 Chapter Conclusion 105

6 Conclusion and Future Work 107

6.1 Dissertation Summary . . . 107
6.2 Research Contributions . . . 107
6.2.1 Solid-state battery discharge input design . . . 107
6.2.2 Known-route human-intervention disturbance rejection . . . 108
6.2.3 Transfer of learned profiles to a novel route . . . 109
6.3 Claim–Evidence Boundaries . . . 110
6.4 Cross-Domain Lessons . . . 110
6.5 Limitations . . . 111
6.6 Future Work . . . 112
6.6.1 Battery-control extensions . . . 112
6.6.2 Known-route disturbance-rejection extensions . . . 113
6.6.3 Transfer-framework extensions . . . 113
6.7 Closing Remarks . . . 114

APPENDIX . . . 115

BIBLIOGRAPHY . . . 117

LIST OF FIGURES

FIGURE

2.1 Dissertation-level modeling approach linking the three research domains. The common workflow is to select a domain, identify a dominant behavior that limits performance, build a simplified representation around that behavior, and use the resulting model to modify an input or planning signal. 16

3.1 Mixed-source workflow used to construct the feedforward discharge-design library. Physical-cell data ground the reduced-order transition characterization in measured voltage response, while high-fidelity simulation data provide additional transition estimates and internal interface-health metrics for evaluating shaped profiles. 23
3.2 Transition interpolation across current-density conditions for the lithium–sodium anode case. Interpolation provides transition-capacity estimates between sampled constant-current simulations and enables a denser set of candidate step-down points for feedforward discharge design. 27
3.3 High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.005 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements. 29
3.4 High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.8 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements. 30
3.5 High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 2.5 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements. 31
3.6 High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.005 MPa for hte 10at% sodium anode. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements. 32
3.7 Candidate feedforward discharge profiles derived from transition-region behavior. The profiles range from aggressive shapes that remain at high current density until a later transition point to conservative shapes that step down earlier to preserve interface health. 33

3.8 Aggressive feedforward discharge profiles for the lithium–sodium anode simulation at 0.8 MPa. Aggressive designs maintain high current density longer and therefore provide higher early discharge rate, but they preserve too little margin to maintain interface health. The yellow region indicates additional delivered capacity relative to the paired constant-current reference over the displayed portion of the discharge. 34
3.9 Conservative feedforward discharge profiles for the lithium–sodium anode simulation at 0.8 MPa. Conservative designs transition earlier within the identified transition region, which reduces exposure to high-current operation and improves robustness to interface degradation. 35
3.10 Input-shape simulation results compared with their respective constant-current references. From left to right, the panels show the most aggressive, aggressive, conservative, and most conservative shaped profiles. The most aggressive and aggressive profiles retain too little contact area fraction, whereas the most conservative profile achieves the largest capacity improvement relative to its matched constant-current reference while preserving comparable final contact area fraction. 37

4.1 Experimental setup used for the experiment. Participants sit in a chair with steering wheel and pedal controls, while a tablet sits to the side. One the tablet is a picture and participants are required to count the number of objects in the pictures to earn points. A zoomed in picture of a typical SPOT task is presented on the top right. 50
4.2 Handling-course segmentation used for known-route learning. Eighteen turns are marked from T1 to T18, with bold hash lines marking turn boundaries and thin hash lines marking internal turn partitions. The start/end point is marked by a circle, and curvature magnitude is shown in color. 53
4.3 Lap-time evolution for the full cohort across handling-course iterations. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The mean trend is shown with a one-standard-deviation band, with participant-level variability retained to show the full experimental spread. Only the high-intervention participant intervened more than 10 times. 57
4.4 Lap-time evolution for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The sensitivity view shows that lap time does not materially change for the remaining cohort. This indicates the lap time cost for the adaptable profiles is low. 58
4.5 Continuous-task score evolution for the full cohort across handling-course iterations. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Scores increase across repeated laps while retaining substantial participant-level variation. Only the high-intervention participant intervened more than 10 times. . . 59

4.6 Continuous-task score evolution for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The trend remains upward after the high-intervention participant is removed. Nearly half of participants completed only 5 iterations, which explains the sharp decrease in standard deviation at the sixth iteration. 60
4.7 Lap-time evolution for the high-intervention participant across handling-course iterations. This exemplar shows a participant whose interventions persisted beyond the initial updates. This participant was more than 25 seconds slower than the average without them. 61
4.8 Lap-time evolution for the non-intervention subgroup. Because these participants did not take over, the route-local learning rule received no event signal for further profile adaptation. 62
4.9 Continuous-task score evolution for the high-intervention participant across handling-course iterations. This single-participant exemplar is interpreted alongside the takeover plots because persistent intervention does not prevent improvement in the secondary-task metric. The large score increase at the end is partially explained by the longer lap duration. 63
4.10 Continuous-task score evolution for the non-intervention subgroup. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Score improvement in this subgroup is a boundary case for the method because it occurs without takeover-based route-local updates. Despite not intervening in takeover tasks for all iterations, their later iteration task score increase suggests more attention on the continuous task. 64
4.11 Route-local learned assumed-friction profile across the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Lower values indicate route bins where the speed planner became more conservative over the course of learning. The lowest μ values correspond to areas before and during sharp turns as well as before long straight sections. 65
4.12 Velocity profile of the initial and final averages for the total cohort in the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Despite the varied mu profiles, standard deviations are fairly tight and mean differences appear mostly at the tightest curves. 66
4.13 Total number of takeover events by section of the handling course. Takeover events cluster in selected route regions rather than being distributed uniformly across the course. Takeover events are clustered around the sharpest turn, T5, but also before the long straightaway, T12. The high-intervetion driver represents large proportions of takeover events at the end of T10 and early T11, as well as late T12 through T13. 67

4.14 Evolution of takeover count across learning iterations for the full cohort. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The mean trend is shown with a one-standard-deviation band; participant-level variation is retained to show the full spread. 69
4.15 Evolution of takeover count across learning iterations for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The remaining decrease indicates that the cohort-level takeover reduction is not explained solely by the high-intervention case. 70
4.16 Evolution of total takeover time across learning iterations for the full cohort. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The plot complements the takeover-count result by showing whether interventions became shorter across learning iterations. 71
4.17 Evolution of total takeover time across learning iterations for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The trend is used as a sensitivity check on the total intervention-duration result. 72
4.18 Evolution of takeover count for the high-intervention participant. The participant continued to produce takeover events after multiple local speed reductions, illustrating a boundary case for the proposed disturbance representation. 73
4.19 Evolution of total takeover time for the high-intervention participant. The trend is interpreted as partial mitigation rather than complete convergence. 74
4.20 Signed lateral error for the initial measured lap compared with the final learned iteration across the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The restricted y-axis in Figure 4.20b shows the mean and one-standard-deviation behavior near the path centerline while the full-scale excursion information is preserved in Figure 4.20a. 75
4.21 Signed lateral-error behavior for the non-intervention subgroup. These participants did not intervene despite path-tracking deviations at some route locations, illustrating that lateral error and takeover are related but not equivalent measures. 76
4.22 Initial and final signed lateral-error behavior for the high-intervention participant. The participant continued to intervene even when lateral error alone did not provide a complete explanation of the intervention pattern, supporting the decision to treat takeover as the measured disturbance signal. 77

5.1 Aggressive-course layout and curvature map used for the transfer evaluation. The circle marks the start/finish location, and turns T1–T11 identify the route bins used for transfer and results interpretation. The color scale indicates curvature magnitude, which is the primary geometric quantity used to construct the route-context feature vector. 92
5.2 Velocity profile across the transfer methods for the total cohort in the aggressive course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The baseline standard deviation denotes the many times drivers lost control or had to slow down to regain control. The much narrower band for the 3-NN method shows consitent performance, despite the varied mu profiles for the total cohort. 96
5.3 Transferred assumed-friction profiles across aggressive-course progress for all available transfer runs. Dark solid lines denote mean velocity while shaded regions show +/- standard deviations. The fixed baseline remains at $\mu = 0.7$; the weighted 3-NN and single-nearest-neighbor profiles are transferred from the handling-course learned library using normalized curvature-context features. The curvature overlay provides the geometric context used by the transfer feature vector. 97
5.4 Transferred assumed-friction profiles across aggressive-course progress for the sensitivity cohort excluding participant 112. Dark solid lines denote mean velocity while shaded regions show +/- standard deviations. The fixed baseline, weighted 3-NN transfer, and single-nearest-neighbor transfer are shown using the same route-progress coordinate as Fig. 5.3. The similarity to the full-cohort profile indicates that the transferred profile structure is not driven solely by the high-intervention participant. 98
5.5 Manual takeover events across aggressive-course progress for all available transfer runs. Bars show the number of takeover events occurring at each course-progress location for the baseline $\mu = 0.7$ profile, the weighted 3-NN transferred profile, and the single-nearest-neighbor transferred profile. The curvature trace is included to relate intervention locations to route geometry; the aggregate event counts are reported in Table 5.2. 99
5.6 Manual takeover events across aggressive-course progress for the sensitivity cohort excluding participant 112. Bars show event counts for the fixed baseline, weighted 3-NN transfer, and single-nearest-neighbor transfer using the same route-progress coordinate as Figure 5.5. The corresponding paired statistical comparison is reported in Table 5.3 Changes noted in turns T2, T4, T6, T7, and T11. 100
5.7 Signed lateral error across aggressive-course progress for all available transfer runs. Bold traces show condition means, shaded bands show approximately +/- standard deviation, and thin outer traces show the observed range across runs. The full-scale y-axis preserves large lateral-error excursions; Fig. 5.7b provides the corresponding zoomed view near the path centerline. 101

5.8 Signed lateral error across aggressive-course progress for the sensitivity cohort excluding participant 112. Bold traces show condition means, shaded bands show approximately ±1 standard deviation, and thin outer traces show the observed range across runs. The full-scale y-axis preserves large lateral-error excursions; Fig. 5.8b provides the corresponding zoomed view. 102

A.1 Example physical cells used for anode-stripping experiments. The images show cells with a lithium anode and a lithium–sodium alloy anode, illustrating the experimental configurations that support the battery-interface modeling chapter. 116

LIST OF TABLES

TABLE

2.1 Illustrative foundational and recent examples of purpose-built modeling approaches used for control and prediction. The examples are not an exhaustive taxonomy; they show that compact, application-oriented models remain active across control engineering and adjacent application areas. 9

2.2 Comparative structure of the dissertation's three application-specific case studies. Each study uses a compact data-informed representation, but the representation, input interface, evidence, and limitations are domain specific. 15

3.1 Reference data used to generate feedforward discharge input shapes. Dashed experimental entries indicate cases where the feedforward design used simulation-derived transition information without a corresponding experimental value in the reference table. Even with increasing stack pressure, the transition capacity decreases with current density. The 10at% sodium stabilized the interface, as shown by the higher discharge capacities for the same pressure and current densities. High current densities showed the greatest disagreement between experimental and simulation data. 26

3.2 Input-shaping simulation metrics after a two-hour discharge. Larger capacities and contact area fractions are favorable. Conservative and Most Conservative shaped inputs show similar interface health but higher areal capacity. The Aggressive and Most Aggressive inputs do not show improvement. The lower overpotential represents the lower current density at the end of the shaped input but simulation demonstrates the interface health is as damaged as the $1.00\,\mathrm{mA\,cm^{-2}}$ 36

4.1 Vehicle and planner parameters used in the known-route human-intervention disturbance-rejection study. 46

4.2 Participant groups used to interpret the known-route learning results. 54

4.3 Paired baseline-to-final outcome statistics for the handling-course learning experiment. Differences are computed as final minus baseline. Negative differences indicate reductions in lap time, takeover count, takeover time, and lateral error; positive differences indicate improvement in continuous-task score. 55

5.1 Route-context features and transfer hyperparameters used for aggressive-course profile transfer. 88

5.2 Aggressive-course performance summary for the fixed baseline assumed-friction profile and two transferred driver-profile methods. Values are mean ± standard deviation. Changes are computed relative to the baseline $\mu = 0.7$ condition. Negative changes indicate reductions in lap time, takeover outcomes, and lateral error; positive changes indicate higher continuous-task score. Table supports improvement using the 3-NN method due to reduced lap time, takeover events/time, peak and mean lateral error, and improved continuous task score. The sensitivity group also shows improvement, but not as much due to poor scores in driver 112. 93
5.3 Participant-level paired statistics for the aggressive-course transfer experiment using the sensitivity cohort excluding participant 112. Differences are computed as comparison condition minus reference condition. Negative differences indicate reductions in lap time, takeover outcomes, and lateral error, while positive differences indicate higher continuous-task score. 95
5.4 Diagnostic aggressive-course exemplars illustrating limitations of curvature-feature transfer. The table includes locations where the weighted 3-NN and single-nearest-neighbor (Nearest) transfer operators produce different assumed-friction profiles (greater than 0.05), and locations where the adaptive profile improves the response but does not fully eliminate the local lateral error. Bolded entries highlight major differences. 103

6.1 Claim–evidence boundaries for the research chapters. 110

LIST OF ACRONYMS

3-NN three-nearest neighbor

BMS battery management system

CI confidence interval

EPA Environmental Protection Agency

GHG greenhouse gas

HAT human–autonomy team

ILC iterative learning control

IRB Institutional Review Board

K-NN K-nearest neighbor

NN nearest neighbor

NHTSA National Highway Traffic Safety Administration

ODD operational design domain

SAE Society of Automotive Engineers

SSB solid-state battery

SOC state of charge

LIST OF SYMBOLS

a	Lumped coefficient in the first-order battery response model.
a_x	Longitudinal acceleration used by the backward-reachable speed planner.
a_y	Lateral acceleration demand, approximated by $v^2\kappa$ for path following.
b	Lumped input coefficient in the first-order battery response model.
c_A	Contact area fraction used as a simulation-based interface-health metric.
$d_{m,i}$	Euclidean distance between target bin m and source bin i in normalized route-context feature space.
e_y	Lateral path-following error.
g	Gravitational acceleration.
g_z	Small-sample corrected paired standardized effect size.
$h^{(k)}(t)$	Manual-takeover indicator during learning iteration k.
I_z	Vehicle yaw moment of inertia.
j, p	Source-route turn and part indices.
ℓ, q	Target-route turn and part indices.
m	Vehicle mass or, when used as an index in Chapter 5, a target-route bin.
N_{TO}	Number of distinct manual-takeover events.
p_1, p_2	Fitted lumped parameters in the reduced-order battery response models.
Q	Areal capacity.
Q_1	Transition areal capacity used in the early battery stability-characterization work.
Q_l	Lower transition-point areal capacity used in the battery transition-region method.
Q_h	Upper transition-point areal capacity used in the battery transition-region method.
$r(n)$	Event-dependent reduction factor for the learned assumed-friction update.

s	Path station or progress coordinate along a route.
$T_{j,p}^{(k)}$	Takeover time assigned to route bin (j, p) during learning iteration k.
$u(t)$	Applied input in the reduced-order battery model.
v	Vehicle speed.
v_{lat}	Curvature-limited speed cap from the assumed-friction envelope.
v_{ref}	Reference speed generated by the backward-reachable speed planner.
v_{straight}	Straight-line reference-speed cap used when curvature is near zero.
V_{ss}	Nominal working overpotential/voltage level used in transition-region identification.
$w_{m,i}$	Inverse-distance weight assigned to source bin i when transferring to target bin m.
$\boldsymbol{M}^{\mathrm{src}}$	Learned source-route assumed-friction matrix.
$\widehat{\boldsymbol{M}}^{\mathrm{tgt}}$	Transferred target-route assumed-friction matrix.
$\boldsymbol{x}_r$	Route-context feature vector for bin r.
$\boldsymbol{z}_r$	Normalized route-context feature vector for bin r.
Δ	Difference between comparison and reference condition.
ϵ	Small positive constant used to avoid singular inverse-distance weights.
ϵ_κ	Curvature threshold below which a segment is treated as effectively straight.
η	Battery overpotential.
κ	Local path curvature.
μ	Assumed coefficient of friction used as a speed-planner parameter and learned disturbance-rejection proxy.
μ_0	Initial or baseline assumed coefficient of friction.
$\mu_{\max}$	Upper saturation value for the learned assumed-friction profile.
μ_{3NN}	Assumed-friction value transferred using weighted three-nearest-neighbor averaging.
μ_{NN}	Assumed-friction value transferred from the single nearest source bin.
ρ_r	Normalized part location within a turn for route-context feature vector $\boldsymbol{x}_r$.
τ	Battery response time-scale parameter.
ψ	Path or vehicle heading angle, depending on context.

ABSTRACT

Complex engineered systems often require control-oriented representations tailored to a specific decision. This dissertation evaluates three application-specific uses of established control-oriented modeling principles: feedforward discharge input design for solid-state batteries with evolving interfaces, route-local speed adaptation for a human–autonomy driving team, and transfer of a learned route-local profile to a new course. In each study, available data are organized into a compact, purpose-built representation that is used to modify an input or reference policy. The studies are comparative case studies rather than instances of a common cross-domain model-reduction algorithm.

The first research contribution develops a reduced-order representation of solid-state battery interface evolution for feedforward discharge control. Physical-cell data and high-fidelity simulation data from collaborators are used to identify transition behavior in the voltage response during discharge. This transition-region model is then used to design shaped discharge-current profiles. Simulation results show that conservative shaped profiles can increase delivered areal capacity relative to a constant-current baseline while preserving comparable interface health. This chapter demonstrates that a simple transition feature can be sufficient for a useful discharge input design, even when the underlying interface evolution remains complex.

The second research contribution treats human takeover in a simulated SAE Level-3 driving task as a repetitive disturbance on a known route. Participants drive a 'handling' course in TruckSim while completing a continuous secondary task. After each measured lap, an iterative learning controller is used to update the local route coefficients of friction that are then applied to the backward-reachable reference speed calculation. Through this iterative process, the controller updates a route-indexed profile of assumed-friction values to reduce future speed demands in regions associated with takeover events. The results show that the takeover count, the takeover time, and the average lateral error decrease for the tested participants, while the continuous-task score increases.

The third research contribution evaluates whether a learned driving profile for a human-autonomy driving team can be transferred to a new course. Route segments are represented using a compact curvature-feature vector, and learned friction values from the handling course are transferred to the target route using weighted three-nearest-neighbor and single-

nearest-neighbor methods to develop new parameters in the control-oriented point-object speed model. These transferred profiles are compared against a fixed baseline friction value. Results show that the weighted three-nearest-neighbor method provides the strongest improvement for the tested drivers and target course, reducing takeover outcomes, increasing continuous-task score, and reducing mean lateral error relative to the fixed baseline.

Together, the studies provide bounded evidence that purpose-built, data-informed representations can support useful input or reference redesign in these three settings. The results do not establish a universal simplified-modeling method or generalization beyond the evaluated cells, participants, routes, and transfer rules.

CHAPTER 1

Introduction

1.1 Control-Oriented Modeling

Engineering models are constructed for particular purposes. Models used to explain a physical mechanism, reproduce a high-fidelity simulation, estimate an internal state, or select a control input require different levels of detail and different error structures. Model reduction, system identification, surrogate modeling, and control-relevant experiment design all seek representations that retain the input–output behavior needed for analysis or decision making while omitting details that are unnecessary for the intended task [Moore, 1981, Gevers and Ljung, 1986, Ljung, 1999]. Data-driven methods extend this practice by learning compact dynamics or predictors from measurements when a complete first-principles model is unavailable or impractical [Brunton et al., 2016, Korda and Mezić, 2018].

This dissertation presents three application-specific case studies that use this established control-oriented modeling philosophy in settings where the relevant simplified representation is not already available for the control objective. The novelty lies in the construction, use, and evaluation of the three domain-specific representations: a transition-region model for feedforward solid-state battery discharge, a route-indexed intervention profile for known-route speed adaptation, and a curvature-context library for transferring that profile to a new route.

The three studies share a practical question. When an existing model or controller does not represent a behavior that limits measured performance, can available data be organized into a compact representation that supports a useful change to the input or reference policy? The answer is necessarily application dependent. The battery chapter uses voltage and simulation data to identify a transition in the interface degradation in the cell. The known-route driving chapter uses repeated takeover events to update a local speed-planning parameter. The transfer chapter uses route geometry to map a previously learned profile to a new course. These are not instances of a formal cross-domain algorithm; they are separate control designs that apply a common, established modeling philosophy.

1.2 Application Context

The first application concerns morphogenic interfaces in solid-state batteries. Solid-state batteries can support lithium-metal-based architectures and offer potential safety and energy-density advantages, but their performance is strongly dependent on electrochemical and mechanical processes at solid–solid interfaces [Janek and Zeier, 2016, Albertus et al., 2018, Pang et al., 2021]. Recent work demonstrates that a deformable secondary phase can reorganize at the lithium–solid-electrolyte interface in response to local electro-chemo-mechanical conditions, producing a morphogenic interface whose contact state evolves during operation [Yoon et al., 2025]. The physical-cell measurements and high-fidelity simulation outputs used in this dissertation were generated through collaboration with battery-materials and battery-modeling researchers [Yoon et al., 2025]. The dissertation contribution is the control-oriented use of those data: Chapter 3 identifies a transition-region representation in the discharge response and uses it to construct feedforward discharge-current profiles.

The second application concerns human intervention in conditionally automated driving. A vehicle-dynamics model may indicate that a planned speed profile is feasible, but a human operator may still intervene. The substantial literature studies takeover time, driver readiness, attention, workload, emotion, steering behavior, and other variables that help explain or predict intervention [Eriksson and Stanton, 2017, Zhang et al., 2019, Gaspar and Carney, 2019, Du et al., 2020, Wang et al., 2018, Sekadakis and Yannis, 2025]. The vehicle chapters pursue a different objective. They do not infer why a participant intervenes or interpret reduced takeover as increased trust. Chapter 4 treats takeover as an observed event that can be used to modify a route-local speed policy on a repeated course. Chapter 5 then evaluates two simple methods to transfer the learned policy to a new course using a curvature-derived route context.

1.3 Dissertation Problem Statement

This dissertation addresses three related, application-specific control problems:

> *How can established control-oriented modeling ideas be instantiated in settings where the available low-order model does not represent the behavior relevant to the input-design decision, and do those application-specific representations improve selected measured outcomes under the tested conditions?*

For the solid-state battery study, the missing representation is a control-oriented indicator of when the evolving interface behavior should trigger a change in the discharge current.

For the known-route human–autonomy study, the missing representation is a route-local record of where the nominal speed policy repeatedly produces human intervention. For the transfer study, the missing representation is a compact description that connects the learned source-route adjustments to geometrically similar locations on a target route.

The dissertation does not derive common applicability conditions, convergence guarantees, or a domain-independent procedure for selecting the correct reduced state. Instead, it evaluates whether purpose-built, data-informed representations can be useful in three specific applications and identifies the evidence and limitations associated with each case.

1.4 Dissertation Thesis

The thesis of this dissertation is that established control-oriented modeling principles can be productively applied to three application-specific problems: feedforward discharge shaping for solid-state batteries with evolving interfaces, route-local reduction of repeated human takeover, and transfer of a learned route-local profile to a new course. In each case, a compact representation constructed from available data is used to modify an input or reference policy, and the resulting performance is evaluated against an application-appropriate baseline.

The thesis is supported only within the boundaries of the studies. The battery result is a simulation-supported feedforward discharge design based on collaborative physical-cell and high-fidelity simulation data. The known-route result is an empirical human-in-the-loop demonstration for the tested participants and course. The transfer result is a bounded comparison of a fixed baseline, inverse-distance weighted three-nearest-neighbor transfer, and single-nearest-neighbor transfer on one target course. Together, the chapters provide comparative evidence that the established modeling philosophy is useful in these settings; they do not establish a universal theory of simplified modeling.

1.5 Research Questions

The dissertation is organized around three research questions.

RQ1: In solid-state batteries with evolving interface behavior, can a transition-region representation derived from discharge data be used to design feedforward discharge-current profiles, and what performance–interface-health tradeoffs result?

RQ2: In a human–autonomy driving team on a repeated route, how does using recurring takeover as a route-local learning signal to modify the backward-

reachable speed profile affect takeover, secondary-task performance, lap time, and lateral error?

RQ3: On a novel route, can inverse-distance weighted three-nearest-neighbor and single-nearest-neighbor transfer of a learned route-local profile out-perform a fixed friction baseline?

The questions are linked by applying the same broad modeling philosophy, but each requires a distinct representation, controller interface, experimental design, and validation criterion. The first question concerns a transition in a physical-system response. The second considers repeated event-based adaptation on a known route. The third investigates representation and transfer rather than continued learning on the same route.

1.5.1 Chapter-level contributions

- Chapter 3 develops a transition-region representation of solid-state battery discharge response from collaborative physical-cell and high-fidelity simulation data. This representation is used to construct piecewise feedforward current-density profiles to evaluate the tradeoff between delivered areal capacity and simulated interface-health metrics.
- Chapter **??** develops and evaluates a route-local learning law that uses measured takeover events to update the assumed-friction profile used by a backward-reachable speed planner. The contribution is a disturbance-mitigation strategy that changes future autonomy behavior without estimating a causal human-state model.
- Chapter 5 develops and compares two interpretable transfer rules for applying a driver-specific learned profile to a novel course: inverse-distance weighted three-nearest-neighbor averaging and single-nearest-neighbor assignment in a normalized curvature-feature space.

1.6 Claim Boundaries

Because the dissertation crosses application domains, the limits of each claim are explicit. Chapter 3 is not a general theory of solid-state battery degradation and does not experimentally validate the shaped profiles. This chapter presents a control-oriented study of feedforward discharge shaping based on a transition-region representation and simulation-based interface-health evaluation. Chapter 4 does not model trust, comfort, attention, or intent. This chapter

defines intervention as a measured event and asks whether localized speed-policy adaptation reduces its recurrence under the tested protocol. Chapter 5 does not establish generalization across drivers, roads, or operational design domains. This chapter evaluates two transfer rules for the tested drivers, one learned source course, and one target course.

The distinction between Chapters 4 and 5 is particularly important. Chapter 4 concerns repeated-route adaptation: information from one lap changes the next lap on the same course. Chapter 5 concerns route-context representation: a completed learned profile is mapped to a different course without repeating the known-route learning process on the new target course.

1.7 Dissertation Organization

Chapter 2 positions the dissertation within the established literature on control-oriented, reduced-order, and data-driven modeling; provides examples from other engineering domains; and then develops the battery and human–autonomy application gaps. Chapter 3 presents the solid-state battery transition-region and feedforward discharge-control study. Chapter 4 presents the known-route human-intervention disturbance-mitigation study. Chapter 5 presents transfer of the learned profile to a novel route. Chapter 6 synthesizes the evidence as three bounded case studies and identifies the additional validation required across battery conditions, participants, routes, and transfer representations.

CHAPTER 2

Motivation and Background

2.1 Why These Problems Matter

The dissertation's application domains are tied to significant engineering needs. Transportation remains a major source of U.S. greenhouse-gas emissions, and transportation electrification depends on energy-storage systems that can provide high energy, sufficient power, long life, and acceptable safety. Solid-state batteries are attractive because they can support lithium-metal-based architectures and eliminate flammable liquid electrolyte, but practical deployment remains constrained by coupled electrochemical and mechanical behavior at internal interfaces [Janek and Zeier, 2016, Albertus et al., 2018, Hatzell et al., 2020b, Pang et al., 2021].

Road safety and the changing allocation of control between humans and automation motivate the driving chapters. In conditionally automated driving, the automation performs the dynamic driving task within an operational design domain, while the human remains part of the broader fallback and supervisory system [SAE International, 2021]. The resulting engineering problem is not only whether the vehicle can follow a path, but whether the coupled human–autonomy system can operate without repeated human intervention. This dissertation studies one narrow control response to that problem: using intervention history to modify the route-local speed policy rather than attempting to infer the latent psychological cause of intervention.

The battery and driving studies are not linked by common physics. They are linked by a common use of established control-oriented modeling practice: construct a representation that is no more detailed than necessary for a defined decision, then evaluate the resulting input or reference change. In the battery study, the input is discharge current density. In the driving studies, the relevant input to the existing controller is a route-dependent reference-speed profile generated from an assumed-friction parameter. The chapters evaluate whether those purpose-built representations improve selected outcomes under bounded experimental or

simulation conditions.

2.2 Purpose-Built Models in Control Engineering

2.2.1 Model fidelity is relative to intended use

Control engineering has a long history of using models that are deliberately less detailed than the underlying system. Balanced realization and truncation retain states that are important to the input–output map while discarding states that are weakly controllable or observable [Moore, 1981]. System identification similarly treats model structure, experiment design, estimation, and validation as choices that must be aligned with the intended application [Gevers and Ljung, 1986, Ljung, 1999]. These traditions establish the principle that a model should be judged by its suitability for a particular analysis or control decision, not by its ability to reproduce every internal mechanism. Recent work continues to refine both the guarantees and the range of systems covered by these methods. Reiter et al. characterize conditions under which the classical balanced-truncation error bound is attained for structured linear systems, while Benner and Goyal extend balancing-based reduction to quadratic-bilinear control systems [Reiter et al., 2024, Benner and Goyal, 2024].

This principle does not imply that simpler is always better. A low-order model can fail if it removes the dynamics that determine constraint activation, stability, safety, or performance. Conversely, a high-fidelity model can be difficult to identify, too expensive for repeated optimization, or poorly matched to the available measurements. The engineering task is therefore to select or construct a representation whose error is acceptable for the intended use.

2.2.2 Model reduction, system identification, and data-driven representations

The terms used in this dissertation describe overlapping but distinct practices. A low-order model has relatively few states or parameters. A reduced-order model is obtained by simplifying or projecting a more detailed description. A data-driven model obtains its structure or parameters primarily from measured or simulated data. A data-informed model, as used here, combines prior knowledge about the application with features or parameters identified from data. None of these labels guarantees control usefulness; usefulness must be evaluated against the intended decision and operating conditions.

Recent work illustrates how these established ideas continue to evolve. In fluid mechanics,

data-driven manifold coordinates paired with state-space or Koopman evolution models have reduced high-dimensional direct numerical simulations to two- or three-dimensional representations while retaining coherent-state dynamics [Constante-Amores and Graham, 2024]. For nonlinear system identification, physics-enhanced sparse identification incorporates known discontinuous or frictional structure into equation discovery from noisy measurements [Lathourakis and Cicirello, 2024]. In predictive control, machine-learned input-augmented Koopman models preserve a tractable lifted representation while accounting explicitly for nonlinear input effects [Li et al., 2024]. These recent studies complement the foundational balanced-POD, SINDy, and Koopman-MPC formulations [Rowley, 2005, Brunton et al., 2016, Korda and Mezić, 2018]; they do not eliminate the need to align the learned or reduced representation with the intended control decision.

Table 2.1 establishes the appropriate context for the dissertation. The general idea of using a compact model for control is not new. The post-2023 studies are included as contemporary extensions of the classical methods rather than replacements for the foundational citations. The battery-health example likewise illustrates a hybrid data-informed modeling principle in an adjacent energy-storage application; it is not a model of the morphogenic solid-state interface studied in Chapter 3. The dissertation's research questions concern whether useful compact representations can be constructed for three specific application objectives that are not directly addressed by the standard model forms already used in those domains.

Table 2.1: Illustrative foundational and recent examples of purpose-built modeling approaches used for control and prediction. The examples are not an exhaustive taxonomy; they show that compact, application-oriented models remain active across control engineering and adjacent application areas.

Application context	Modeling approach	Information retained	Intended use
Linear dynamical systems	Balanced realization and truncation [Moore, 1981, Reiter et al., 2024]	States important to controllability and observability, with error characterized through Hankel singular values	Reduced input–output models for analysis and controller design
Quadratic-bilinear control systems	Balancing-based nonlinear model reduction [Benner and Goyal, 2024]	Reachability and observability information represented through generalized Gramians	Reduced models that retain control-relevant nonlinear input–state interactions
Fluid dynamics	Balanced proper orthogonal decomposition and data-driven manifold models [Rowley, 2005, Constante-Amores and Graham, 2024]	Dominant input–output flow structures or low-dimensional invariant-manifold coordinates	Tractable models for flow prediction and feedback-control development
Nonlinear dynamical systems and system identification	Sparse and physics-enhanced equation identification [Brunton et al., 2016, Lathourakis and Cicirello, 2024]	A small set of active governing terms, augmented when appropriate by known physical structure	Interpretable low-order dynamics identified from measured or simulated data
Nonlinear predictive control	Koopman predictors with learned state and input liftings [Korda and Mezić, 2018, Li et al., 2024]	Approximately linear evolution in a lifted coordinate space with explicit input dependence	Computationally tractable model predictive control of nonlinear systems
Battery health estimation	Physics-informed neural networks [Wang et al., 2024]	Empirical and state-space degradation structure combined with learned dynamics	State-of-health estimation and prognosis across varying cell conditions

2.3 Position of This Dissertation

The dissertation applies the established purpose-built modeling philosophy through three case studies. It does not present a formal model-reduction algorithm shared across the chapters. The source data, representation, input interface, and evidence differ in each study:

- The battery study uses physical-cell and high-fidelity simulation data to identify a transition region in discharge response and to schedule changes in current density.
- The known-route driving study uses repeated takeover events to update a route-indexed assumed-friction profile that changes a backward-reachable reference speed.
- The transfer study uses normalized curvature features to assign previously learned assumed-friction values to a different route through weighted three-nearest-neighbor and single-nearest-neighbor rules.

The value of placing these studies in one dissertation is therefore comparative rather than methodological. Together, they show how the same established modeling philosophy leads to different domain-specific choices, what evidence is needed to defend those choices, and what information is lost by each simplification.

2.4 Primer A: Control-Oriented Modeling of Solid-State Battery Interface Evolution

2.4.1 Promise and limitation of solid-state batteries

Solid-state batteries replace the liquid electrolyte in conventional lithium-ion batteries with a solid electrolyte. This change can enable lithium-metal negative electrodes and improve safety and energy density, but it also makes solid–solid contact and coupled electro-chemo-mechanical behavior central to performance [Janek and Zeier, 2016, Albertus et al., 2018, Hatzell et al., 2020b]. The lithium–solid-electrolyte interface can lose contact, form voids, support nonuniform reaction, or experience pressure-dependent evolution during plating and stripping [Luntz et al., 2015, Zhang et al., 2020, Vishnugopi et al., 2023, Singh et al., 2023].

Recent work on morphogenic interfaces further shows that the interface need not behave as a fixed boundary. Yoon et al. demonstrated that a deformable secondary phase can reorganize in response to local electro-chemo-mechanical stimuli and alter contact behavior during operation [Yoon et al., 2025]. This evolving morphology motivates a control problem in which the relationship between commanded current density, voltage response, and interface health changes with operating history.

2.4.2 Established low-order battery models and their control uses

Battery management commonly relies on a hierarchy of models. Equivalent-circuit models represent terminal-voltage behavior using lumped electrical elements and are attractive for real-time estimation because they are inexpensive to simulate and straightforward to parameterize [Hu et al., 2012, Plett, 2015]. Electrochemical models retain more internal structure. The Doyle–Fuller–Newman model resolves transport and reaction through porous electrodes and electrolyte, while single-particle models and their electrolyte-corrected extensions reduce that description to obtain tractable predictors for estimation and control [Speltino et al., 2009, Marquis et al., 2019, Liu et al., 2022].

Recent battery-health research also illustrates a hybrid, data-informed strategy. Wang et al. combine empirical degradation and state-space structure with neural-network components for lithium-ion state-of-health estimation across multiple battery types and operating protocols [Wang et al., 2024]. That objective differs from morphogenic-interface control, but the modeling principle is relevant: preserve available physical structure and learn the unresolved behavior needed for a defined estimation or control task.

These approaches demonstrate that low-order and reduced-order battery models are already standard tools. Their success, however, is tied to the phenomena they retain and the applications for which they were developed. Equivalent-circuit models can reproduce terminal-voltage dynamics without uniquely resolving the physical origin of a changing impedance. Single-particle-type models are primarily derived for intercalation-based lithium-ion cells and retain diffusion and electrolyte dynamics selected for those systems. More detailed solid-state battery models can resolve coupled mechanics, transport, reaction, and contact evolution, but their internal states and parameters are difficult to identify online and their computational burden is poorly matched to simple feedforward scheduling [Bates et al., 2015, Bistri et al., 2021, Pang et al., 2021].

2.4.3 Why existing simplified strategies are insufficient for this study

The limitation relevant to Chapter 3 is specific. The desired control variable is discharge current density, and the relevant performance limit is the onset of a degradation-dominated response associated with evolving solid–solid contact. A conventional equivalent-circuit fit may approximate interface, but some morphogenic SSB's have interfaces that evolve and repair as long as their interface does not degrade too far. A conventional single-particle model does not contain the contact-area and morphology dynamics central to the collaborating solid-state cell and simulation system. Conversely, the available high-fidelity model contains

the desired interface information but is not the representation used to schedule the feedforward input.

Chapter 3 therefore does not argue that existing low-order battery models are generally inadequate. It addresses a narrower gap: the absence of a simple, experimentally accessible representation designed specifically to schedule discharge-current changes around the observed transition in morphogenic-interface behavior. The chapter uses measured voltage and simulated contact information to identify a transition region, then evaluates whether that representation can improve delivered areal capacity relative to a constant-current baseline at comparable simulated interface health.

2.4.4 Collaborative data and dissertation contribution

The physical-cell measurements and high-fidelity simulation outputs used in Chapter 3 were generated by collaborators in a multi-university battery project [Yoon et al., 2025]. Those contributions establish the material system, experiments, and detailed simulation environment. The dissertation contribution is the control-oriented reduction and input design: selecting a transition-region representation from those data, interpolating it across operating conditions, constructing shaped discharge profiles, and evaluating the resulting performance–interface-health tradeoff.

2.5 Primer B: Human–Autonomy Intervention and Control Adaptation

2.5.1 Conditional automation and intervention

The SAE J3016 taxonomy distinguishes levels of driving automation by allocation of the dynamic driving task and fallback responsibility. In conditional automation, the vehicle performs the driving task within a defined operational design domain, while the human remains part of the broader system and may resume or assert control [SAE International, 2021]. Human intervention therefore reflects the behavior of a coupled human–automation system rather than the vehicle controller alone.

Empirical studies have examined how takeover timing and quality vary with urgency, non-driving-related tasks, driver state, interface design, roadway context, and prior experience [Eriksson and Stanton, 2017, Zhang et al., 2019, Gaspar and Carney, 2019, Du et al., 2020, Sekadakis and Yannis, 2025]. Other work estimates intervention intent from measurable signals such as steering torque or characterizes voluntary intervention separately from system-requested takeover [Wang et al., 2018, Zhong et al., 2025]. A recent systematic review of 38

peer-reviewed studies groups the relevant mental-model factors around drivers' understanding of automation, attentional processing, and attitudes toward the automated system [Nasser et al., 2025]. The breadth of those factors reinforces the dissertation's decision not to infer a single latent cause from a voluntary takeover event. This literature is primarily explanatory or predictive: it seeks variables and relationships that account for when intervention occurs or how well the transition is performed.

2.5.2 Existing control approaches that respond to the human

Human-aware control is not limited to explanation. Shared-control and adaptive-automation studies change control authority or haptic assistance using measurements such as workload, gaze, input torque, or muscle activity. For example, workload-adaptive haptic shared control has been evaluated in human-in-the-loop driving experiments, and adaptive steering assistance has been modulated using forearm electromyography [Luo et al., 2020, Wang et al., 2020]. These approaches demonstrate that human measurements can be used to change automation behavior in real time.

The driving contribution in this dissertation is distinct from both the explanatory literature and these immediate shared-control approaches. It does not estimate a latent human state and does not continuously reallocate steering authority. Instead, it uses the occurrence and route location of voluntary takeover as a repeated-task performance signal. The actionable response is also different: the autonomy modifies the future reference-speed profile locally, before the vehicle returns to the same route context.

2.5.3 Known-route learning as event-based disturbance mitigation

A repeated route supplies the structure needed to associate intervention with route location. The vehicle encounters similar path geometry on each lap, the nominal planner produces a similar reference profile, and takeover events can be assigned to turns or route parts. Information from one lap can therefore modify the next lap.

Iterative learning control provides the conceptual basis for using previous-trial information to improve repeated-task performance [Arimoto et al., 1984, Bristow et al., 2006]. Chapter 4 adapts that logic to an event-based outcome. The learned quantity is a route-indexed assumed-friction profile used by a backward-reachable speed planner. When intervention occurs, the assumed-friction value is reduced in the associated route bin, which reduces the planned speed on a subsequent lap. Limited recovery toward the nominal ceiling is allowed after a previously reduced bin is traversed without intervention.

The chapter does not interpret the learned assumed-friction value as physical road friction

or as a direct measure of human preference. It is a planner parameter that records where a more conservative speed reference has been associated with reduced intervention. The research gap is therefore specific: using repeated voluntary intervention directly as a route-indexed learning signal for reference-speed redesign without first constructing a causal model of the human state.

2.5.4 Transfer beyond the learned route

A profile indexed only by distance along one course cannot be applied directly to a different course. Chapter 5 therefore represents each source and target route bin using five normalized features: current-bin mean absolute curvature, current-bin peak absolute curvature, preceding-bin mean absolute curvature, preceding-bin peak absolute curvature, and normalized bin location within the turn. The features encode a minimal geometric context related to lateral acceleration demand and turn progression.

Two transfer operators are evaluated. Single-nearest-neighbor transfer assigns the learned value from the closest source bin in normalized feature space. Inverse-distance weighted three-nearest-neighbor transfer blends the values from the three closest source bins. Both are compared with the fixed baseline $\mu = 0.7$. This transfer formulation is related to instance-based learning and personalized autonomy [Cover and Hart, 1967, Dudani, 1976, Karnan et al., 2024b, Cui et al., 2024], but its purpose is narrower: reuse an intervention-mitigation profile rather than infer a general model of driver preference.

Curvature is intentionally a minimal feature set. It does not encode road width, grade, banking, visual occlusion, obstacle placement, traffic, shoulder proximity, or all transient vehicle states. The diagnostic examples in Chapter 5 are therefore important evidence, because they identify target locations where the two transfer operators disagree or where lateral error remains elevated despite adaptation.

2.6 Comparative Structure of the Three Case Studies

The three research chapters apply an established modeling philosophy through separate domain-specific constructions. Table 2.2 summarizes the comparison. It is not a prescriptive workflow: the table does not imply that the transition-region, event-based learning, and nearest-neighbor transfer procedures are interchangeable or that they provide a general recipe for a fourth domain.

The comparison clarifies what is and is not unified. All three studies use available data to construct a control-oriented representation and evaluate an altered input or reference. The

Table 2.2: Comparative structure of the dissertation's three application-specific case studies. Each study uses a compact data-informed representation, but the representation, input interface, evidence, and limitations are domain specific.

Case study	Compact representation and source data	Control use	Bounded evidence and limitation
Solid-state battery discharge	Transition-region model identified from physical-cell voltage and high-fidelity simulation data	Schedule feedforward discharge-current reductions	Simulated capacity and interface-health comparison for the evaluated cells and profiles; no experimental validation of the shaped inputs
Known-route human–autonomy driving	Turn/part-indexed assumed-friction profile updated from takeover events over repeated laps	Recompute the backward-reachable speed profile on the same route	Human-in-the-loop takeover, task-score, lap-time, and lateral-error outcomes for the tested participants and course
Transfer to a novel route	Weighted 3-NN or single-NN mapping of a learned source profile using normalized curvature features	Construct a target-route assumed-friction and speed profile	Paired comparison with a fixed baseline on one target course; local diagnostic cases expose curvature-only limitations

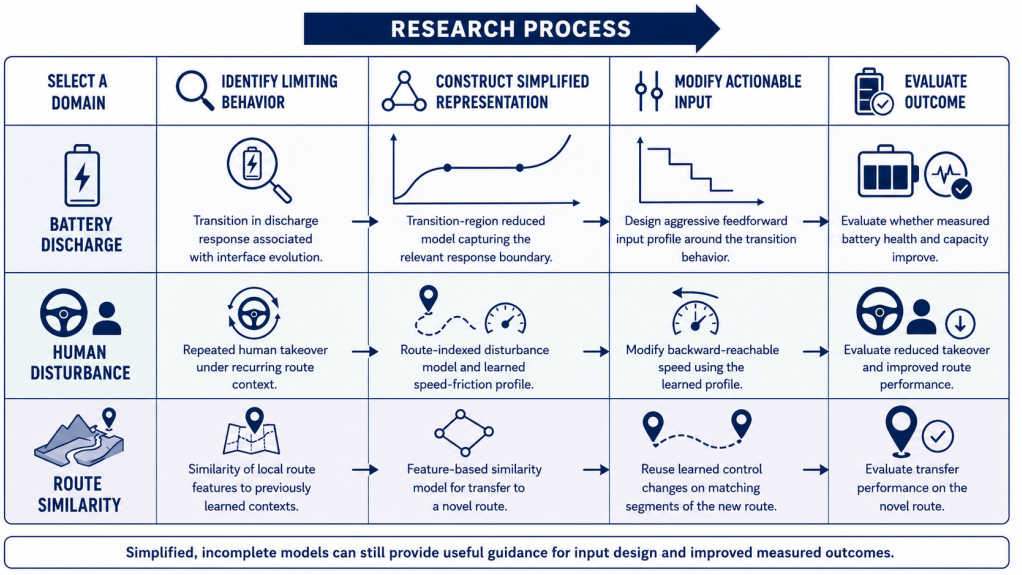

Figure 2.1: Dissertation-level modeling approach linking the three research domains. The common workflow is to select a domain, identify a dominant behavior that limits performance, build a simplified representation around that behavior, and use the resulting model to modify an input or planning signal.

mathematical models, identification steps, and validation procedures are not shared. The dissertation's synthesis is therefore evidence that the established philosophy can be useful in these three settings, together with an explicit account of the limitations of each application. But there is a common workflow used for the dissertation, as shown in 2.1.

2.7 Application-Specific Research Gaps

The literature supports three application-specific gaps addressed by the dissertation.

First, low-order battery models are widely used for estimation and control, but the standard equivalent-circuit and single-particle model forms are not designed to represent the current-driven transition associated with evolving solid–solid contact in the collaborative morphogenic-interface system. Recent physics-informed lithium-ion health models likewise address estimation and prognosis rather than current scheduling around solid–solid contact evolution [Wang et al., 2024]. Detailed solid-state battery models can represent that evolution, but they are not the simple input-scheduling representation evaluated in Chapter 3. The battery gap is therefore a transition-oriented feedforward discharge model that connects accessible response data to current-profile design.

Second, the automated-driving literature provides rich explanatory and predictive models of takeover and also includes human-aware shared control. Recent systematic evidence

further shows that takeover performance depends on several cognitive and interface-related factors, supporting the decision to treat intervention as an observed signal rather than a direct measurement of a single human state [Nasser et al., 2025]. The known-route gap is narrower: route-indexed use of repeated voluntary takeover as a direct learning signal for future speed-reference modification, without estimating the causal human state and without changing the underlying path-following controller.

Third, personalized and example-based autonomy methods motivate reuse of prior experience, but the transfer of a driver-specific intervention-mitigation profile through a compact five-feature curvature representation remains an open application question. Chapter 5 evaluates two transparent transfer rules and reports both aggregate improvements and local failure cases.

These gaps define the dissertation's contributions more precisely than a general claim about simplified modeling. The dissertation applies an established engineering practice to three specific control objectives, develops the required domain-specific representations, and evaluates their usefulness and limitations under the tested conditions.

CHAPTER 3

Reduced-Order Modeling of Solid-State Battery Interface Evolution for Feedforward Discharge Control

3.1 Introduction

Solid-state batteries (SSBs) are attractive for next-generation electrified systems because they offer a plausible path to improved safety, higher energy density, and compatibility with lithium-metal or lithium-alloy negative electrodes [Janek and Zeier, 2016, Albertus et al., 2018, Hatzell et al., 2020a]. These advantages depend strongly on the behavior of the solid electrolyte–electrode interface. Unlike an idealized fixed boundary, this interface evolves as current is applied. Contact can be lost, pores and voids can form, lithium morphology can change, and the resulting transport and mechanical conditions can alter usable capacity, overpotential, and long-term interface health [Pang et al., 2021, Luntz et al., 2015, Singh et al., 2023, Yoon et al., 2025]. For a control system, this is a difficult setting: the dominant limiting behavior is internal and history dependent, while the routinely available measurements are external electrical quantities such as voltage and applied current density.

This chapter addresses this problem from a control-oriented perspective. Rather than attempting to reconstruct every electrochemical, mechanical, and morphological state at the solid electrolyte–electrode interface, the objective is to identify a reduced-order representation that preserves the behavior needed to design a useful discharge input. The central premise is that the measured voltage response during constant-current discharge contains a transition-like signature. At lower stress or earlier in discharge, the response may be approximately steady, slowly varying, or only weakly increasing. As interface degradation accumulates, the response can transition into a regime in which overpotential grows more rapidly. If that transition can be identified consistently, then it can be used as a design feature for feedforward discharge shaping.

The data used in this chapter were generated through collaboration. Physical-cell data provide experimentally grounded voltage responses under specified conditions of chemistry,

pressure, and current-density. High-fidelity simulation data, developed by collaborators, supplement those measurements by providing internal interface-health quantities, especially contact area fraction, which are generally inaccessible in routine physical testing. This mixed-source structure is essential to this chapter. The physical cells ground the reduced-order model in measured behavior, while the simulation provides a way to evaluate whether shaped discharge inputs preserve interface health while improving practical discharge performance.

The contribution of this chapter is a reduced-order transition-region framework for feedforward discharge input design in SSBs with evolving interfaces. The method identifies a voltage-based transition region from constant-current data, fits simple local response models to characterize the measured behavior, interpolates transition behavior across current-density conditions, and designs non-constant discharge profiles that begin at higher current density before stepping down to lower current density. The resulting input shapes test whether a cell can provide a higher current density for part of the discharge while maintaining interface health comparable to a lower constant-current baseline.

The principal result is that a conservative shaped discharge profile increased the two-hour areal capacity from $1.00\,\mathrm{mA\,h\,cm^{-2}}$ to approximately $1.10\,\mathrm{mA\,h\,cm^{-2}}$, a 9.5 % improvement, while ending with a comparable overpotential and contact area fraction to the $0.50\,\mathrm{mA\,cm^{-2}}$ constant-current baseline. More aggressive profiles provided higher early current density but produced poorer interface health, illustrating the performance–robustness tradeoff associated with transition-region selection. These results support a core tenant of the thesis that simplified, data-informed models can preserve the dominant behavior needed for input design, even when they do not fully resolve the underlying physics.

The remainder of the chapter is organized as follows. Section 3.2 summarizes the relevant background on SSB interface evolution, control-oriented battery modeling, and the progression from early stability characterization to transition-informed input shaping. Section 3.3 presents the data sources, variables, reduced-order response models, transition-region identification method, and feedforward discharge-profile design. Section 3.4 presents simulation-based discharge-performance results. Section 3.5 discusses what the reduced-order model captures, the role of physical and simulation data, limitations, and implications for future feedback control.

3.2 Background and Related Work

3.2.1 Interface Evolution in Solid-State Batteries

A central challenge in SSB development is that the interface between the solid electrolyte and the negative electrode is not static. During stripping and plating, the interface can lose contact, form voids, support dendrite growth, and undergo mechanical or morphological changes that alter the effective impedance of the cell [Luntz et al., 2015, Pang et al., 2021, Singh et al., 2023]. These effects matter because electrochemical performance depends not only on the bulk material properties, but also on local transport, pressure, and contact conditions at the interface. A cell can therefore exhibit unacceptable voltage growth or interface degradation even when the nominal material system is attractive.

Recent work on morphogenic interfaces emphasizes that these interfacial structures can evolve dynamically rather than acting as fixed boundaries [Yoon et al., 2025]. In the collaborative material systems motivating this chapter, lithium and lithium–sodium alloy negative electrodes are used in SSB configurations, and the interface can change with cycle history, applied current density, and stack pressure. From a controls perspective, this creates a familiar but difficult problem: the dominant limiting behavior is internal, history dependent, and only partially observable through external electrical measurements.

Physically significant interface-health variables are difficult to measure directly during ordinary operation. Contact area fraction is an especially important example because loss of contact can increase overpotential and reduce accessible capacity. However, direct contact-area characterization typically requires specialized characterization, destructive evaluation, or high-fidelity simulation. A more practical control-oriented model should translate accessible electrical measurements—specifically voltage response—into a model that drives input design, without claiming to fully reconstruct the underlying interface state

3.2.2 Control-Oriented Battery Models and Their Limits

Battery models span a wide range of fidelity. Molecular and particle-scale models can provide detailed insight into mechanisms, while continuum and cell-scale models can support simulation and control design. Conventional lithium-ion battery control often relies on reduced-order electrochemical models, single-particle models, pseudo-two-dimensional models, equivalent-circuit models, or data-driven state estimators [Plett, 2015, Speltino et al., 2009, Liu et al., 2022, Wang et al., 2023]. These approaches remain valuable, but their assumptions do not directly eliminate the interface-evolution challenge in SSBs.

The difficulty is not simply that detailed SSB models do not exist. High-fidelity SSB

models can capture coupled electrochemical, mechanical, and morphological processes and are essential for understanding mechanisms and guiding material design [Bates et al., 2015, Vishnugopi et al., 2023, Zhang et al., 2020, Masias et al., 2019]. The difficulty is that such models are often computationally expensive, parameter intensive, and difficult to deploy directly as real-time control models across many cells. Conversely, very simple empirical models may be fast but fail to preserve the behavior needed for input design, especially when the cell response changes qualitatively as interface degradation accumulates.

This chapter uses an intermediate strategy. The model is not a full physics reconstruction of the solid electrolyte–electrode interface. It is also not a purely black-box regression. Instead, it uses a low-dimensional, data-informed representation of the voltage response and identifies a transition region that acts as a control-relevant feature. This representation is useful because the input-design task does not require a complete reconstruction of the interface. It requires a way to decide when high discharge current density becomes undesirable and when a decrease in current density is likely to preserve interface health.

3.2.3 From Stability Characterization to Transition-Region Design

The earliest stage of this work treated the voltage response of experimental cells as a linear-systems characterization problem. Constant-current data were filtered and fit to simple response families. The associated time constants and pole-location estimates were then interpreted as indicators of whether the measured response was stable, marginal, or unstable over the discharge window. In that early framing, lower current densities produced stable or slowly varying behavior, while higher current densities produced faster voltage growth. The preliminary study reported stable behavior at lower current densities, unstable behavior at higher current density, and a transitional or stabilizable region between them.

That preliminary stability framework was useful for two reasons. First, it recasts a difficult interface-degradation problem as a control problem: how does the measured output respond as a function of input current density and what input regimes appear to move the system toward undesirable voltage growth? Second, it suggests that the most useful control feature may not be a single binary stable/unstable label. Instead, a useful feature is a transition region that identifies when the cell begins to move from a relatively benign response to a degradation-dominated response. The present chapter builds on that insight by moving from pole-location characterization to feedforward discharge input design.

3.2.4 Discharge Input Shaping

Non-constant current profiles are common in battery operation, but shaped discharge control for individual SSB cells remains much less developed than standard charging and battery-management methods. Intermittent or shaped current profiles can matter because the interface can tolerate a high current density for short periods of time even when constant operation at that current density would lead to excessive degradation [Lee and Sakamoto, 2024]. This observation is especially relevant for high-power applications. For example, a vehicle battery pack may need a short-duration power burst during acceleration, even if maintaining that same current demand over a long window would be harmful.

The discharge input-shaping problem can therefore be stated as a tradeoff. A higher current density provides higher instantaneous power and greater early discharge rate, but may accelerate interface degradation. Conservative current densities protect interface health but fail to exploit the high-power capabilities of the system. The transition-region model developed in this chapter provides a way to navigate this tradeoff: remain at higher current density while the measured and modeled response remains within a conservative transition window, then step down before the interface response enters a more rapidly degrading regime.

3.3 Methodology

3.3.1 Data Sources and Scope

In this chapter, two complementary data sources are used. The first source is the data of physical-cell discharge collected by the collaborators [Yoon et al., 2025]. These data provide a measured voltage response for specified cell configurations, including negative-electrode chemistry, stack pressure, and applied current density. In the principal discharge dataset, lithium composite working electrodes and solid electrolyte layers were tested against lithium counter electrodes. Physical cells are essential because they anchor the reduced-order model in experimentally observed behavior. However, physical-cell measurements primarily provide accessible external quantities, especially the measured voltage and applied current density. They do not directly provide internal interface-health variables such as contact area fraction during ordinary operation.

The second source is high-fidelity simulation data generated by collaborators [Vishnugopi et al., 2023] for related solid electrolyte–electrode interface systems. These simulations provide the voltage response and internal health quantities of the interface under constant-current discharge. In particular, contact area fraction is used as a simulation-derived metric for

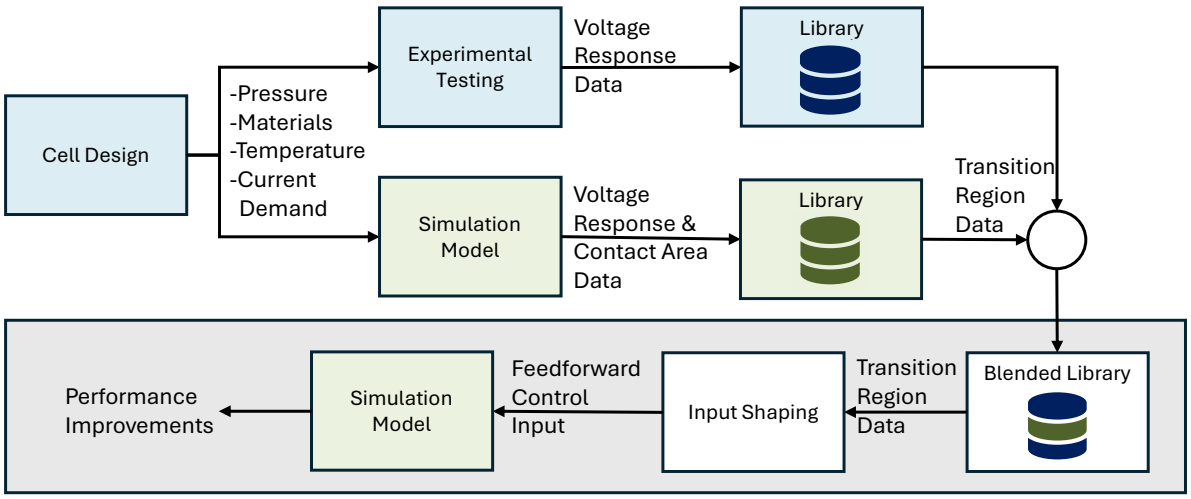

Figure 3.1: Mixed-source workflow used to construct the feedforward discharge-design library. Physical-cell data ground the reduced-order transition characterization in measured voltage response, while high-fidelity simulation data provide additional transition estimates and internal interface-health metrics for evaluating shaped profiles.

interface health. Simulation models include coupled interfacial mechanisms such as contact mechanics, lithium creep, lithium self-diffusion, and electrochemical stripping behavior. The simulation data are not treated as a substitute for experiment; rather, they serve as a complementary source that allows the reduced-order transition model to be evaluated against internal degradation metrics that are otherwise difficult to measure directly.

Figure 3.1 summarizes the mixed-source workflow used in the chapter. Physical-cell data are used to identify the voltage-current density response structure and transition behavior. Simulation data are then used to expand the available operating-condition library and to evaluate whether transition-informed feedforward discharge inputs can improve accessible capacity without producing unacceptable interface degradation. This division of labor is important for the dissertation argument: the reduced model is intentionally simple, but it is tied to both measured electrical behavior and simulation-derived internal health metrics.

3.3.2 Areal Capacity, Overpotential, and Interface-Health Metrics

Let $j(t)$ denote the applied discharge current density, with units of $\mathrm{mA\,cm^{-2}}$. Positive $j(t)$ denotes stripping/discharge in this chapter. The areal capacity discharged up to time t is

$$Q(t) = \int_0^t j(\xi)\, d\xi, \tag{3.1}$$

with units of $\mathrm{mA\,h\,cm^{-2}}$ when time is expressed in hours. For constant-current tests, $Q(t) = jt$, so time and areal capacity provide equivalent independent variables up to a known scaling. Because shaped inputs change the current density over time, areal capacity provides a convenient coordinate for comparing transition behavior across current density segments.

The measured voltage response is converted to an overpotential-like signal, denoted η. In the reduced model, η is not claimed to be a complete interface-health state. It is an accessible electrical response that reflects contributions from transport, polarization, and interface degradation. The usefulness of η comes from the fact that increased overpotential during stripping is associated with a deteriorating interface response, including contact loss and increased effective resistance [Pang et al., 2021, Singh et al., 2023]. The model therefore uses η as a practical proxy to identify when the cell leaves a relatively benign response regime.

The simulation-derived interface-health metric used for the final evaluation is the contact area fraction, denoted c_A. This quantity is not assumed to be measured directly in a deployed controller. Instead, it is used in this chapter to evaluate whether a feedforward profile that was designed using transition information also preserves the interface health predicted by the high-fidelity simulation.

3.3.3 Reduced-Order Response Families

The reduced-order model begins with a first-order linear response form,

$$\tau \dot{y}(t) + a y(t) = b u(t), \tag{3.2}$$

where $y(t)$ represents the fitted output response, $u(t)$ represents the applied input, τ is a time-scale parameter, and a and b are lumped coefficients. Under a step input, the solution is exponential when $a \neq 0$ and linear when $a = 0$. This motivates two simple response families for fitting the measured voltage response over a local region:

$$\eta_{\mathrm{exp}}(Q; j) = -p_1(j) \exp\left(-\frac{Q}{\tau_Q(j)}\right) + p_2(j), \tag{3.3}$$

$$\eta_{\mathrm{lin}}(Q; j) = p_1(j) Q + p_2(j). \tag{3.4}$$

Here the response is written as a function of areal capacity Q, rather than time, because constant-current characterization traces satisfy $Q = jt$. The fitted parameter τ_Q therefore has units of areal capacity when the response is expressed in the Q-domain. The parameters p_1, p_2, and τ_Q are identified from data for each current-density and cell-configuration condition. The fit family is selected by comparing the residual error over the corresponding regime.

This simple family is not intended to reproduce every physical mechanism. Its role is to capture whether the observed response is slowly varying, approximately linear, exponentially decaying, or exponentially growing over the region relevant to transition detection and input design. The preliminary work used the same logic to compute effective pole-location metrics and interpret cell behavior as stable, marginal, or unstable. In the present chapter, the pole-location interpretation is retained as intuition, but the primary design object is the transition in areal capacity rather than the pole estimate itself. That shift makes the framework better suited for feedforward discharge design, because the input must decide when to step from one current-density value to another.

3.3.4 Transition-Region Identification

For each constant-current discharge response, the transition-region algorithm identifies a nominal voltage regime and then determines when the voltage response leaves that regime. Let Q_f denote the final areal capacity of the constant-current trace used for characterization. The nominal response is defined using the average voltage or overpotential between 20% and 50% of the cycle capacity,

$$\bar{\eta}_{20-50} = \frac{1}{0.3Q_f} \int_{0.2Q_f}^{0.5Q_f} \eta(Q)\, dQ. \tag{3.5}$$

In implementation, this integral is evaluated by averaging the sampled data over the corresponding areal-capacity window. A tolerance band is then formed around the nominal response,

$$\eta^- = (1-\epsilon)\bar{\eta}_{20-50}, \qquad \eta^+ = (1+\epsilon)\bar{\eta}_{20-50}, \tag{3.6}$$

where ϵ is selected from a small range such as $0.025 \leq \epsilon \leq 0.05$. The lower and upper transition capacities, Q_l and Q_h, are then defined from the areal-capacity values at which the response enters and exits this nominal band. Q_l is the conservative transition boundary and Q_h is the aggressive transition boundary.

This transition interval encodes a performance–robustness tradeoff. If the shaped input remains at high current density until Q_h, the profile is aggressive because it extracts more capacity at high current density before stepping down. If the profile steps down near Q_l, it is conservative because it sacrifices some high-current operation to reduce the risk of entering a rapidly degrading response regime.

Table 3.1: Reference data used to generate feedforward discharge input shapes. Dashed experimental entries indicate cases where the feedforward design used simulation-derived transition information without a corresponding experimental value in the reference table. Even with increasing stack pressure, the transition capacity decreases with current density. The 10at% sodium stabilized the interface, as shown by the higher discharge capacities for the same pressure and current densities. High current densities showed the greatest disagreement between experimental and simulation data.

Na (at.%)	Stack Pressure (MPa)	Applied Current Density (mA/cm^2)	Experiment			Simulation		
			Q (mAh/cm^2)	Pole-Location Metric	Areal Capacity Window (mAh/cm^2)	Q_1 (mAh/cm^2)	Pole-Location Metric	Areal Capacity Window (mAh/cm^2)
0	0.005	0.25	—	—	—	0.649	42.282	0.75
0	0.005	0.50	—	—	—	0.118	52.900	0.21
0	0.800	0.75	0.291	18.20	0.45	0.141	10.700	0.40
0	0.800	1.00	0.155	4.32	0.85	0.134	17.800	0.25
0	2.500	1.00	0.253	16.18	1.01	0.161	16.600	0.50
0	2.500	1.50	0.140	5.38	1.36	0.143	10.500	0.50
10	0.005	0.25	—	—	—	0.741	37.700	0.93
10	0.005	0.50	—	—	—	0.143	15.500	0.35
10	0.005	1.50	—	—	—	0.330	70.000	0.40
10	2.500	1.00	1.673	0.78	2.96	0.185	0.000	0.80

3.3.5 Reference Library and Transition Interpolation

Table 3.1 summarizes the reference transition data used to construct the feedforward discharge profiles. The table combines experimental and simulation-derived quantities across sodium composition, stack pressure, and applied current density. The experimental and simulation transition-capacity columns retain the names used in the source data, Q and Q_1, respectively. These values identify the transition capacities used to select current-density step-down points. The pole-location values are retained as response-characterization metrics from the reduced-order fits; they are used comparatively to describe response growth behavior rather than as complete physical descriptors of the interface.

Because the library contains only a finite set of current-density conditions, interpolation is used to estimate transition capacities between sampled points. Within the reduced-order design model, transition capacity is treated as a smooth function of applied current density for a fixed material and pressure condition. A modified Akima interpolation is used to generate intermediate transition estimates without requiring additional high-fidelity simulations at every candidate current density. Figure 3.2 illustrates this interpolation for the lithium–sodium anode case.

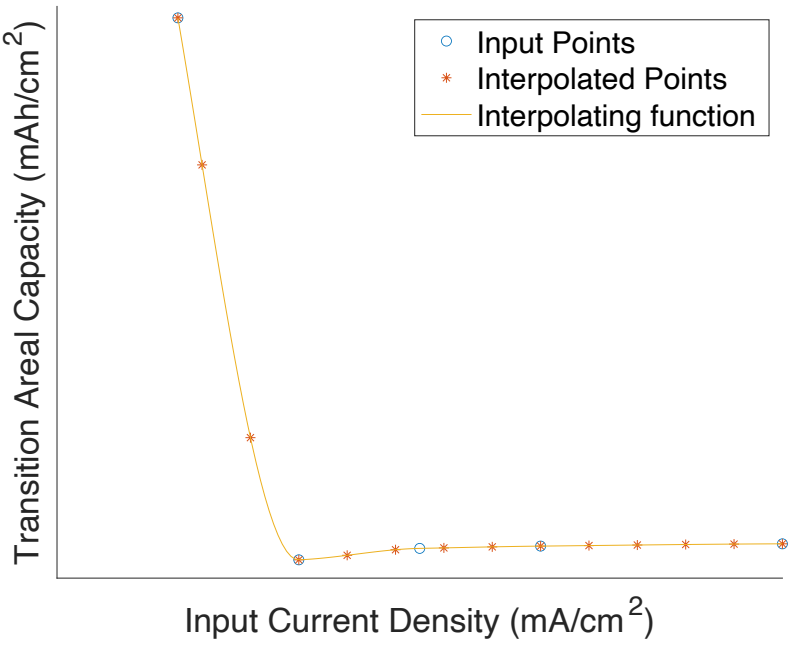

Figure 3.2: Transition interpolation across current-density conditions for the lithium–sodium anode case. Interpolation provides transition-capacity estimates between sampled constant-current simulations and enables a denser set of candidate step-down points for feedforward discharge design.

3.3.6 Simulation-Based Evaluation of Interface Health

The transition-region model uses measured or simulated voltage response as the design signal, but the final evaluation also requires an interface-health metric. For this purpose, high-fidelity simulation provides contact area fraction during discharge. The simulation environment resolves interface behavior that the reduced model intentionally omits, including contact-loss dynamics associated with the solid electrolyte–electrode interface. Figure **??** shows representative simulation-derived responses used to evaluate the proposed input-shaping strategy.

3.3.7 Feedforward Discharge Profile Design

The feedforward discharge profiles are designed as piecewise-constant current-density signals. Each profile begins with a high current density, $j = 1.5\,\mathrm{mA\,cm^{-2}}$, and then steps down when the discharged areal capacity reaches a transition threshold. The general decision rule is

$$j(Q) = j_k, \qquad Q_k^\star \le Q < Q_{k+1}^\star, \tag{3.7}$$

where j_k is the current density applied over segment k, and $Q_k^\star$ is the transition capacity used to trigger the next step-down. The thresholds are selected from the transition region,

$$Q^\star(j_k;\gamma) = Q_l(j_k) + \gamma\left[Q_h(j_k) - Q_l(j_k)\right], \tag{3.8}$$

where $\gamma \in [0,1]$ is an aggressiveness parameter. Smaller values of γ produce conservative profiles that step down earlier, while larger values produce aggressive profiles that remain at high current density longer.

Four representative profiles are evaluated, as shown in Figure 3.7. The most aggressive profile remains at high current density until the late transition boundary and then steps down gradually. The aggressive profile also uses gradual step-downs but begins the transition earlier. The conservative profile steps from the initial high current density to an intermediate current density. The most conservative profile steps from the initial high current density to an intermediate current density and then to a lower current density based on conservative transition boundaries. These profiles are not claimed to be globally optimal. They are deliberately simple test cases designed to determine whether transition-informed feedforward shaping can improve discharge capacity while preserving comparable interface health.

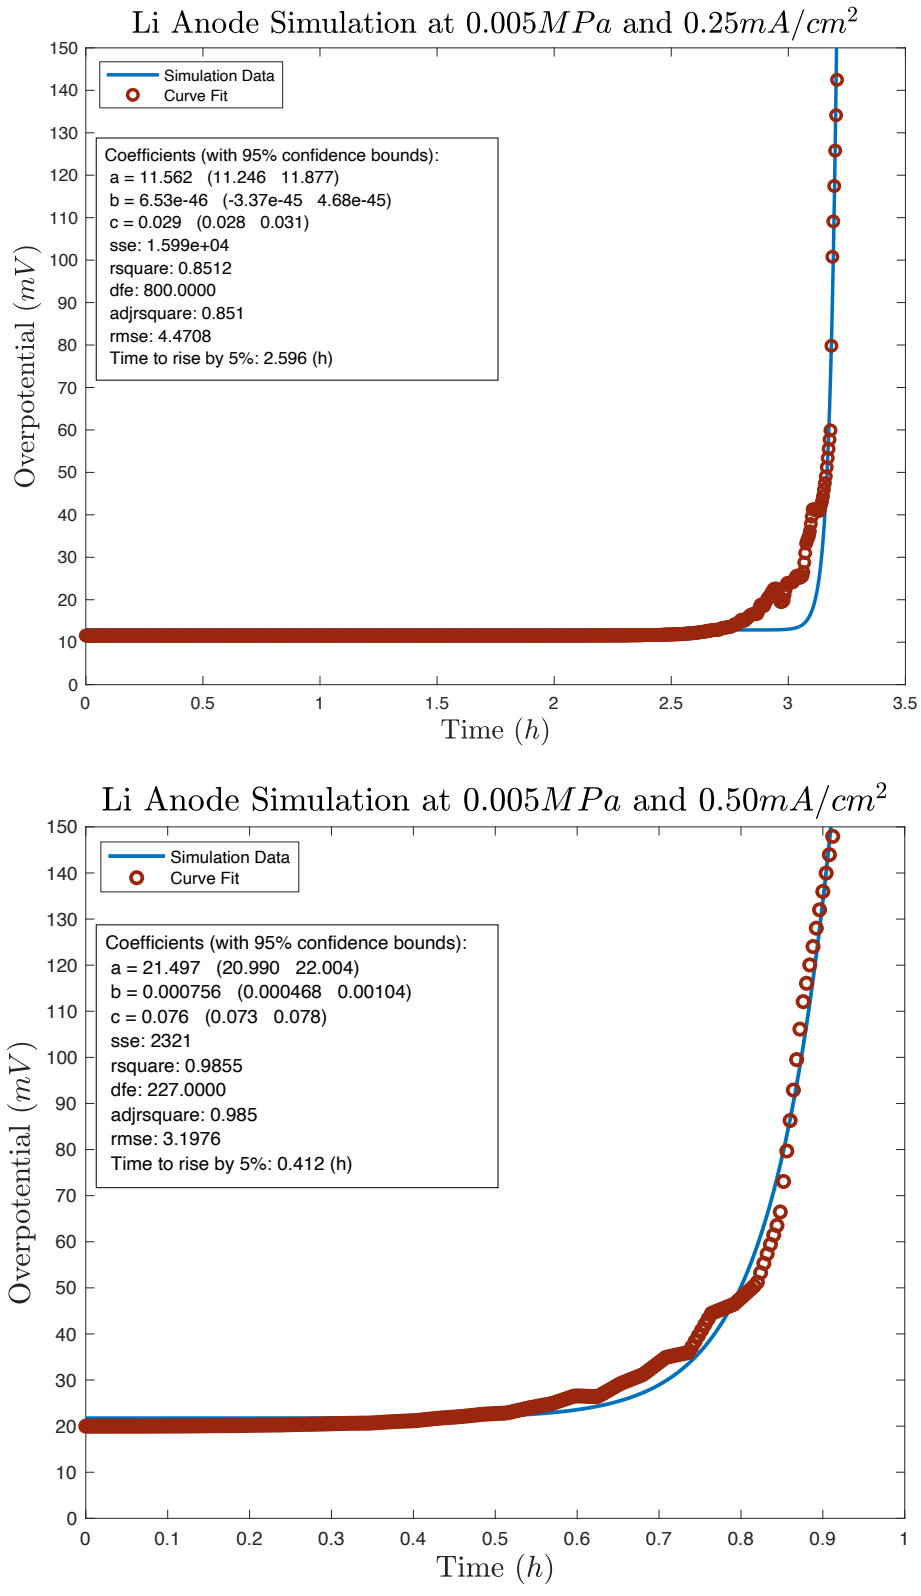

Figure 3.3: High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.005 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements.

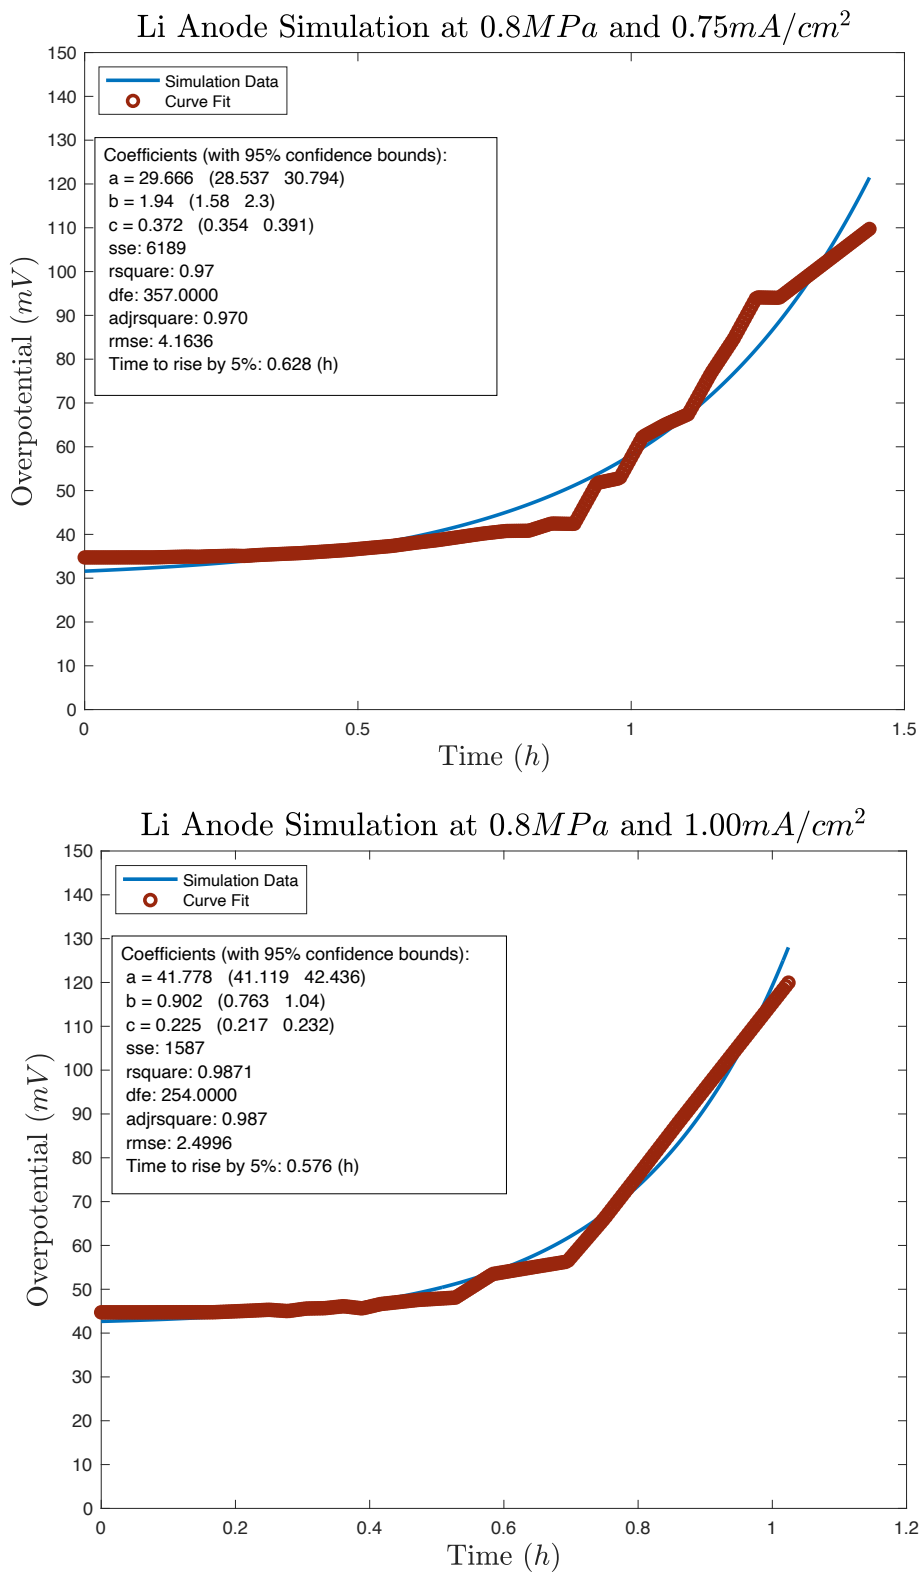

Figure 3.4: High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.8 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements.

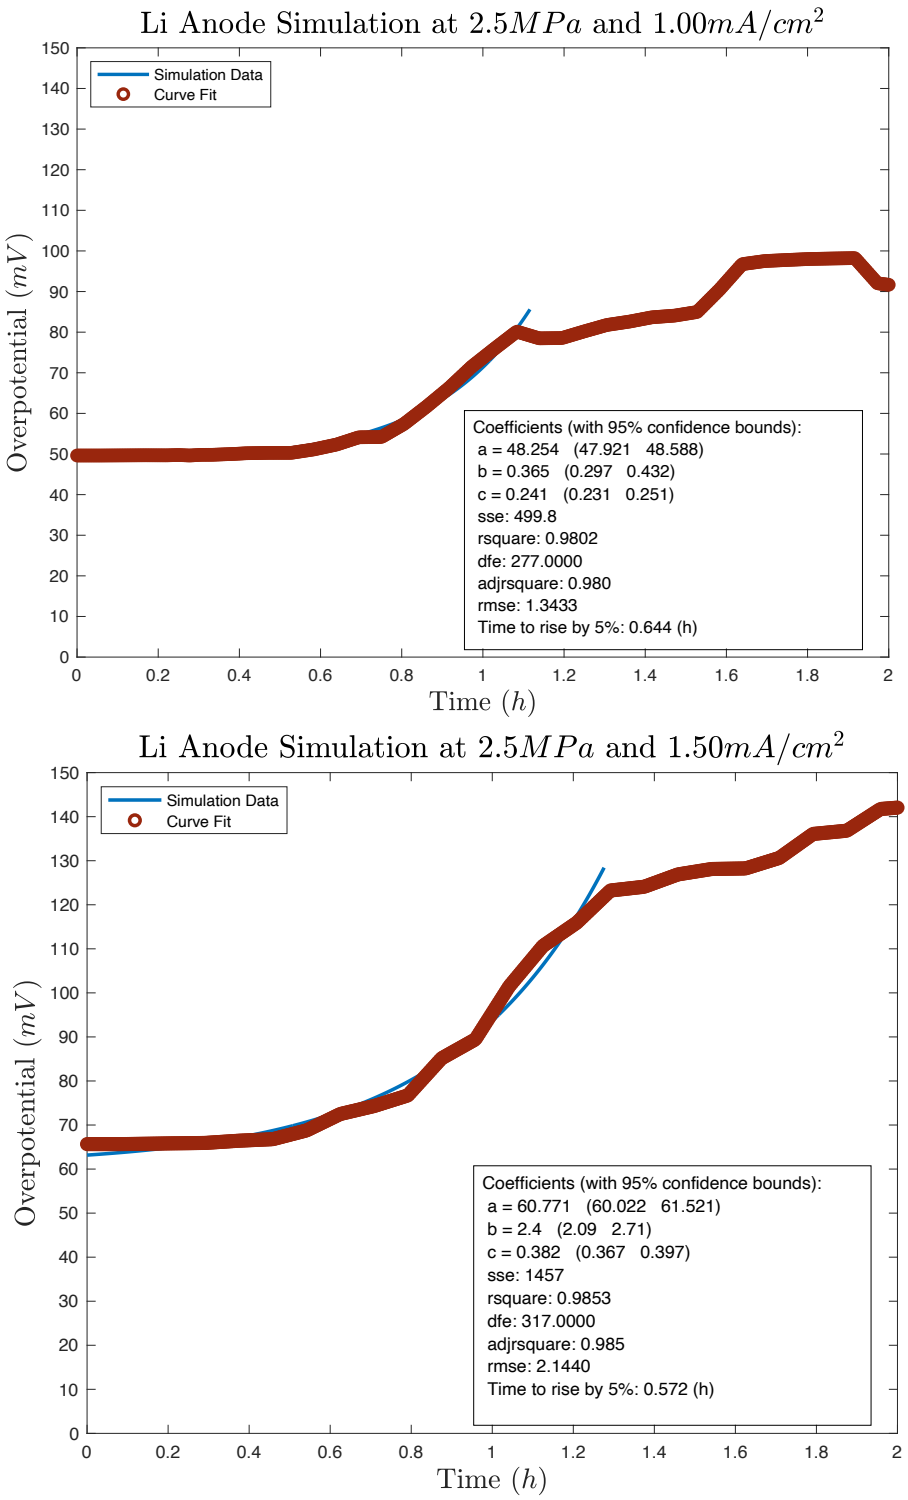

Figure 3.5: High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 2.5 MPa. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements.

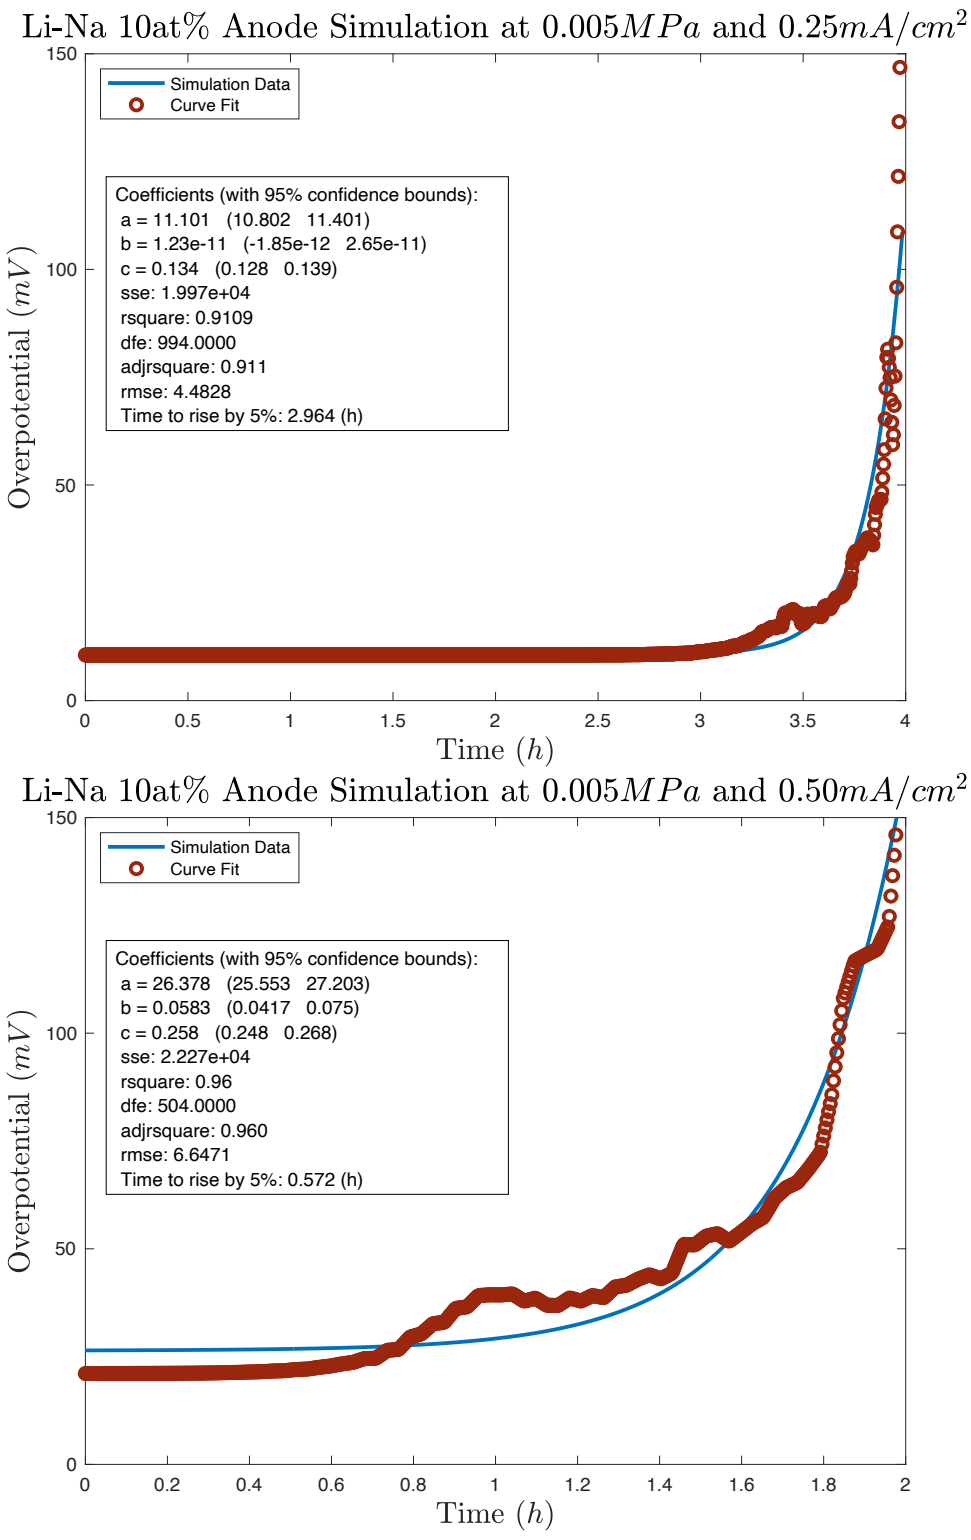

Figure 3.6: High-fidelity simulation responses used to evaluate the feedforward discharge profiles at 0.005 MPa for hte 10at% sodium anode. The simulation data supplement physical-cell voltage measurements by providing contact area fraction and other internal responses that are not directly available from routine physical-cell measurements.

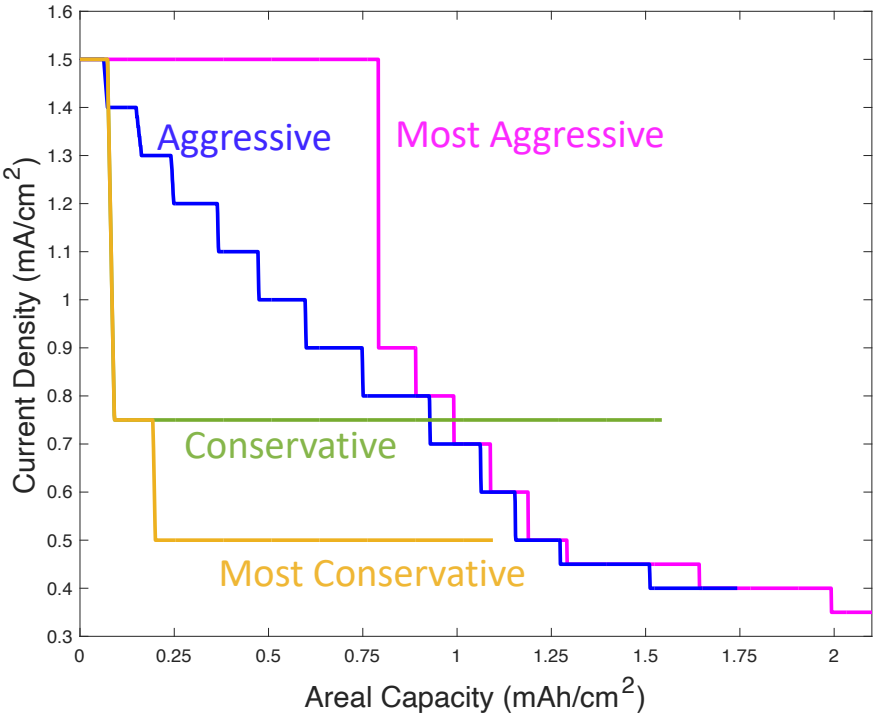

Figure 3.7: Candidate feedforward discharge profiles derived from transition-region behavior. The profiles range from aggressive shapes that remain at high current density until a later transition point to conservative shapes that step down earlier to preserve interface health.

3.4 Results

3.4.1 Aggressive and Conservative Profile Families

The transition-region method produces a family of shaped discharge profiles rather than a single design. This is useful because the transition point is not a perfectly known physical boundary. It is a data-informed design feature extracted from the voltage response, and different placements within the transition interval correspond to different assumptions about robustness. Profiles that step down late seek higher performance but risk greater interface degradation. Profiles that step down early sacrifice some performance to preserve interface health.

Figure 3.8 shows representative aggressive input cases. These profiles keep the discharge current density high for longer portions of the discharge. As expected, aggressive profiles can increase early delivered power, but also move the cell closer to degradation-dominated behavior. In the simulation results, aggressive designs that transition later produce poor interface-health outcomes, with contact area fraction approaching the degraded levels associated with a high constant-current discharge. This result is important because it shows that transition-informed shaping is not automatically beneficial; the transition-region location matters.

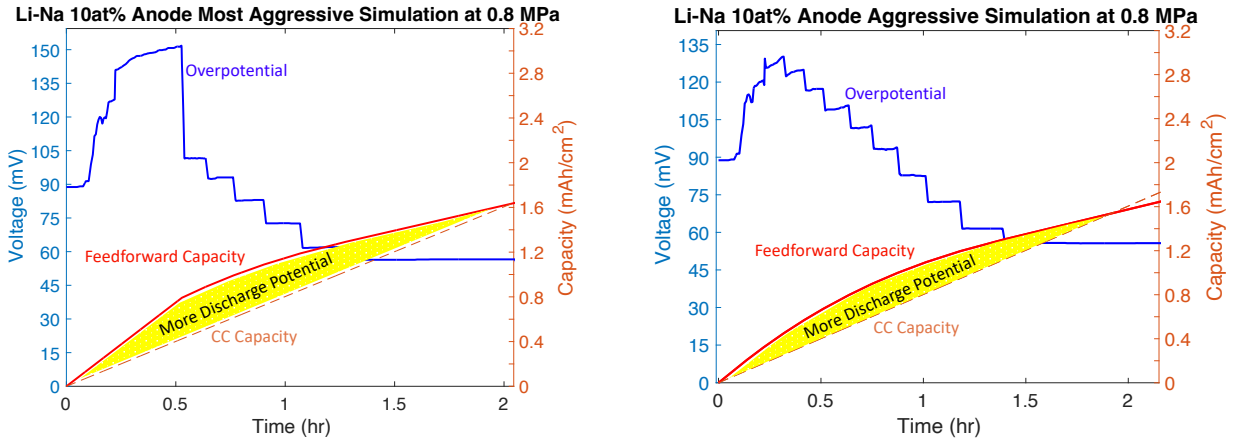

(a) Most aggressive profile.

(b) Aggressive profile.

Figure 3.8: Aggressive feedforward discharge profiles for the lithium–sodium anode simulation at 0.8 MPa. Aggressive designs maintain high current density longer and therefore provide higher early discharge rate, but they preserve too little margin to maintain interface health. The yellow region indicates additional delivered capacity relative to the paired constant-current reference over the displayed portion of the discharge.

Figure 3.9 shows representative conservative input cases. These profiles keep the discharge current density high for a short portion of the discharge, then steps down. As expected,

conservative profiles can increase early delivered power, but unlike the aggressive, it does so without degrading the cell. In the simulation results, conservative feedforward profiles produce the same interface-health outcomes, with contact area fraction approaching the same levels associated with a safe constant-current discharge. Since the constant current reference for the most conservative is only $0.50\,\mathrm{mA\,cm^{-2}}$ The conservative inputs demonstrate that the performance of the battery can improve without degrading the interface health.

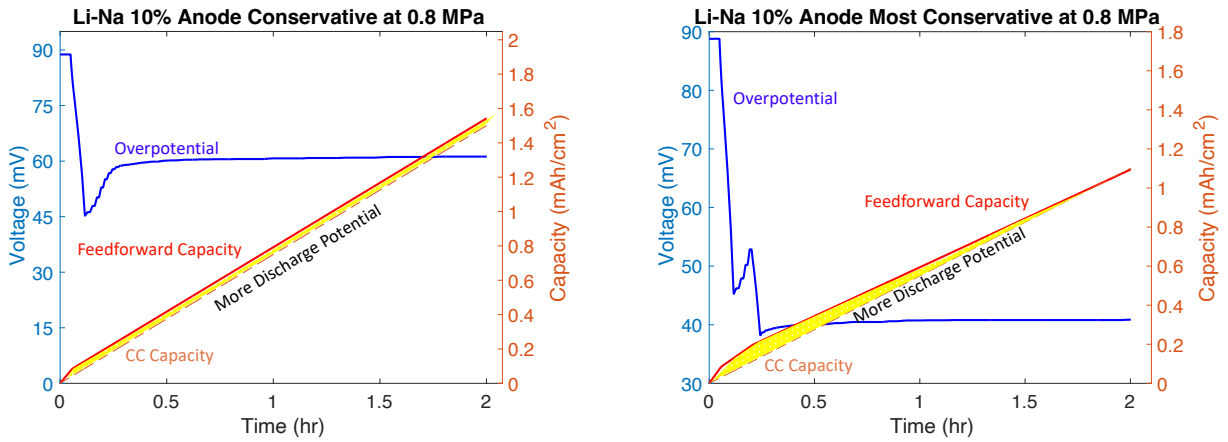

(a) Conservative profile.

(b) Most conservative profile.

Figure 3.9: Conservative feedforward discharge profiles for the lithium–sodium anode simulation at 0.8 MPa. Conservative designs transition earlier within the identified transition region, which reduces exposure to high-current operation and improves robustness to interface degradation.

3.4.2 Two-Hour Discharge Performance

The primary comparison evaluates each shaped profile after a two-hour discharge window. The shaped inputs are compared against constant-current baselines selected to provide relevant interface-health comparisons. The simulation outputs include final areal capacity, overpotential, and contact area fraction. Table 3.2 summarizes the results.

The most aggressive and aggressive shaped profiles did not provide the desired tradeoff. Although their final overpotentials were lower than the $1.00\,\mathrm{mA\,cm^{-2}}$ constant-current baseline because the current density was stepped down later in the discharge, their final contact area fractions were poor. The most aggressive profile ended with a contact area fraction of 19 %, and the aggressive profile ended with 15 %. These values are close to the degraded interface state of the $1.00\,\mathrm{mA\,cm^{-2}}$ constant-current baseline, which ended at 14 %. Thus, late transition points provided too little robustness to preserve interface health.

Table 3.2: Input-shaping simulation metrics after a two-hour discharge. Larger capacities and contact area fractions are favorable. Conservative and Most Conservative shaped inputs show similar interface health but higher areal capacity. The Aggressive and Most Aggressive inputs do not show improvement. The lower overpotential represents the lower current density at the end of the shaped input but simulation demonstrates the interface health is as damaged as the $1.00\,\mathrm{mA\,cm^{-2}}$

Profile	Current label	Areal capacity ($\mathrm{mA\,h\,cm^{-2}}$)	Overpotential (mV)	Contact area fraction (%)
Constant current	$1.00\,\mathrm{mA\,cm^{-2}}$	2.00	127	14
Constant current	$0.75\,\mathrm{mA\,cm^{-2}}$	1.50	62	60
Constant current	$0.50\,\mathrm{mA\,cm^{-2}}$	1.00	40	60
Most aggressive	shaped	1.62	57	19
Aggressive	shaped	1.575	56	15
Conservative	shaped	1.54	60	60
Most conservative	shaped	1.10	40	60

The conservative profile produced a more favorable result. It ended with an areal capacity of $1.54\,\mathrm{mA\,h\,cm^{-2}}$, compared with $1.50\,\mathrm{mA\,h\,cm^{-2}}$ for the $0.75\,\mathrm{mA\,cm^{-2}}$ constant-current baseline. It also ended with the same contact area fraction, 60 %, and a similar overpotential. This corresponds to a modest capacity improvement while maintaining comparable simulated interface health.

The most conservative profile provided the strongest result relative to its matched baseline. It ended with an areal capacity of $1.10\,\mathrm{mA\,h\,cm^{-2}}$, compared with $1.00\,\mathrm{mA\,h\,cm^{-2}}$ for the $0.50\,\mathrm{mA\,cm^{-2}}$ constant-current baseline. The final overpotential was 40 mV for both cases, and the final contact area fraction was 60 % for both cases. Thus, the most conservative shaped input produced a 9.5 % increase in areal capacity while preserving comparable end-of-discharge interface health. Figure 3.10 provides the side-by-side simulation comparison that supports this tradeoff interpretation across the full profile family.

3.4.3 Interpretation of Capacity and Interface-Health Tradeoffs

The results support three conclusions. First, the cell response is sensitive to when the discharge current density is reduced. Aggressive profiles that delay the step-down can preserve high early current density, but they do not preserve interface health. Second, earlier step-down profiles can improve practical discharge capacity relative to conservative constant-current baselines. Third, the transition-region model provides a useful design coordinate because the selected transition point directly determines the performance–robustness tradeoff.

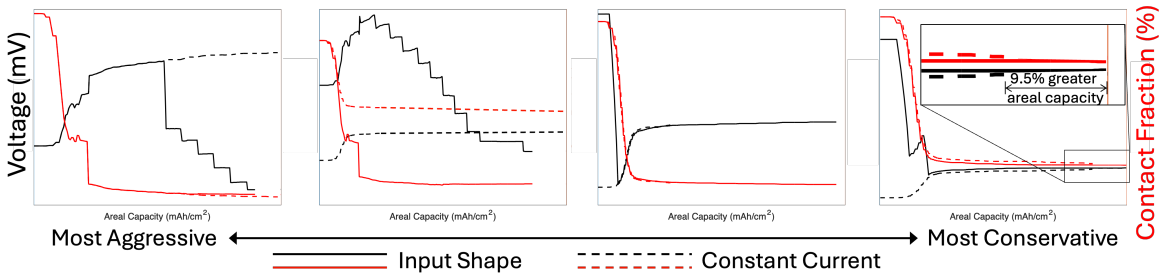

Figure 3.10: Input-shape simulation results compared with their respective constant-current references. From left to right, the panels show the most aggressive, aggressive, conservative, and most conservative shaped profiles. The most aggressive and aggressive profiles retain too little contact area fraction, whereas the most conservative profile achieves the largest capacity improvement relative to its matched constant-current reference while preserving comparable final contact area fraction.

The most conservative profile is especially important because it demonstrates the value of a simplified model. The feedforward design did not require a full online reconstruction of interface morphology. It used a transition feature inferred from accessible voltage-response behavior, then evaluated the resulting input in a high-fidelity simulation environment. The result is not a proof of global optimality, but it is evidence that a low-dimensional representation of dominant interface behavior can support useful input redesign.

3.5 Discussion

3.5.1 What the Reduced-Order Model Captures

The reduced-order transition model captures the feature of the discharge response that is most important for the feedforward design problem: the approximate areal-capacity location at which a high-current discharge begins to lose robustness. This is not equivalent to identifying all interface mechanisms. The model does not separately estimate void geometry, local pressure, lithium creep, local current constriction, or spatially resolved contact morphology. Instead, it treats the voltage response as a measurable signature of an evolving interface and extracts the transition interval that determines when a current-density step-down should occur.

This abstraction is appropriate for the design question studied here. A feedforward profile does not need to know the exact interface morphology if it can identify when continued high-current operation is likely to degrade interface health. In that sense, the transition

region is analogous to a control boundary. It is not a fundamental material constant, but it is a useful design feature derived from the measured and simulated response behavior.

The results also show why a single binary stable/unstable classification is insufficient. The most useful design information is not that high current density can be harmful. That point is already expected from interface-degradation physics. The useful information is where the response begins to transition and how conservative the input should be relative to that transition. The aggressive and most conservative cases make this distinction clear. Both use the same initial high current density, but their different step-down timing leads to substantially different interface-health outcomes.

3.5.2 Role of Physical and Simulation Data

The chapter relies on the combination of physical-cell and simulation data. The physical-cell data are necessary because they demonstrate that the reduced-order transition behavior is grounded in the measured voltage response. The simulation data are necessary because they provide internal interface-health metrics that are not directly available from the physical-cell voltage trace. This mixed-source structure is a strength of the chapter, but it must be described with caution.

The reduced-order model should not be presented as a replacement for high-fidelity simulation. High-fidelity models remain necessary for interpreting mechanisms and evaluating internal states such as contact area fraction. The contribution of the reduced-order model is different: it provides a computationally simple, control-oriented representation that can generate candidate inputs and guide design exploration. The collaborative origin of the experimental and simulation data is therefore part of the technical interpretation of the chapter: the control-design contribution is the reduced-order characterization and feedforward input-design framework, while the physical-cell measurements and high-fidelity simulation outputs provide the experimental and mechanistic evidence used to evaluate that framework.

3.5.3 Limitations

Several limitations should temper the claims of the chapter. First, the shaped inputs are not globally optimized. They are structured profiles selected to test whether the concept of the transition-region can improve discharge performance. A formal optimization of current-density trajectories, transition thresholds, and health constraints could produce better profiles and should be considered in future work.

Second, the strongest interface-health evaluation is simulation based. The shaped profiles should eventually be experimentally validated on physical cells, especially because manufactur-

ing variability and experimental noise may shift the transition region. The chapter therefore supports the claim that transition-informed feedforward discharge shaping is promising and simulation-validated, but it should not claim complete experimental validation of the shaped profiles unless those tests are added.

Third, the model parameters are configuration dependent. A transition region identified for one chemistry, pressure, and current-density range should not be assumed to transfer unchanged to another cell design. This does not undermine the framework, but it means that the reduced-order model must be re-identified or adapted when the cell configuration changes. In a practical battery-management setting, this motivates future work on online identification, uncertainty bounds, and cell-to-cell variability.

Fourth, overpotential is an imperfect proxy for interface health. The most conservative shaped input shows why this matters: the overpotential response and contact area fraction can tell related but non-identical stories during parts of the discharge. The final overpotential and contact area fraction match the conservative baseline, but intermediate behavior may still differ. A final controller should therefore treat voltage-based transition features as indicators rather than direct health measurements.

3.5.4 Implications for Control Design

Despite these limitations, the results provide useful insight into the design of the controller. SSB interface evolution does not necessarily require a fully resolved online interface model before input design can begin. A reduced-order transition feature can be extracted from constant-current data, interpolated across current-density conditions, and used to produce feedforward profiles that improve performance relative to conservative constant-current operation. This is the same modeling philosophy used throughout the dissertation: identify the dominant unmodeled behavior, represent it at the level required for control, and use that representation to redesign the input.

The next technical step is to convert the feedforward design into an optimization and eventually into feedback. A formal optimization could define an objective such as

$$\max_{j(t)} \; Q(T) - \lambda_\eta \eta(T) - \lambda_c \left[c_{\min} - c_A(T)\right]_+ , \tag{3.9}$$

where $Q(T)$ is final areal capacity, $\eta(T)$ is final overpotential, $c_A(T)$ is contact area fraction, $c_{\min}$ is a minimum acceptable contact-area threshold, and $[\cdot]_+$ denotes the positive part. In physical implementation, $c_A(T)$ would not be directly measured, so feedback designs would require either an observer, a learned health proxy, or conservative constraints based on the transition-region uncertainty. The present chapter does not solve that full problem, but it

provides the reduced-order feature around which such a controller could be built.

3.6 Chapter Summary

This chapter developed a reduced-order transition-region framework for feedforward discharge input design in SSBs with evolving interfaces. The method used constant-current voltage responses to identify transition intervals, interpolated transition behavior across current-density conditions, and designed piecewise-constant shaped discharge profiles ranging from aggressive to conservative. High-fidelity simulation was then used to evaluate whether these profiles improved discharge capacity while preserving contact-area-based interface health.

The key result is that the most conservative discharge profile increased the two-hour areal capacity by 9.5 % compared to a conservative constant-current baseline, while ending with a comparable overpotential and contact area fraction. More aggressive profiles demonstrated why transition-region selection matters: delaying the step-down can preserve high early current density, but it can also produce poor final interface health. These results support the claim that simplified, data-informed models can be useful when they preserve the dominant behavior needed for input design, even when they do not fully resolve the underlying system physics.

CHAPTER 4

Learning-Based Disturbance Rejection for Human Takeover on a Known Route

4.1 Chapter Overview

This chapter develops a learning-based disturbance-rejection framework for a human–autonomy team operating repeatedly over a known route. The measured disturbance is a human takeover event: an interval in which the participant interrupts automated operation and assumes manual control. The chapter does not infer whether a takeover was caused by trust, distraction, discomfort, situation awareness, perceived risk, simulator familiarity, or another latent human state. It asks a narrower control question: when takeover events recur at similar route locations, can the autonomy reduce later takeover by changing the local speed profile?

The distinction is central to the contribution. Much of the takeover literature is explanatory: it studies takeover time, driver state, intervention quality, and post-intervention behavior in conditionally automated driving [Zhang et al., 2019, Sekadakis and Yannis, 2025, Sekadakis et al., 2025, Yu et al., 2025a]. Those models are valuable when the objective is to explain or predict the human driver. The objective here is different. This chapter uses takeover as a measured output of the coupled human–autonomy system and uses that output to revise the autonomy's future reference behavior.

The known-route setting makes the learning problem well posed. Chapter 5 considers transfer to a new route. In contrast, this chapter assumes that the same handling course is traversed repeatedly and that route-local information from one iteration can be used to modify the next iteration. This repeated-task structure is similar in spirit to iterative learning control, in which previous-trial information improves future-trial performance [Arimoto et al., 1984, Bristow et al., 2006]. The quantity learned in this chapter is not a direct steering or acceleration command. It is a turn/part-indexed assumed coefficient-of-friction profile used by a backward-reachable speed planner. Reducing the assumed coefficient of friction in a

route bin reduces the local feasible speed and generates a more conservative reference profile for the next lap.

The chapter contribution is therefore a simplified, data-informed representation of human intervention as a repetitive, route-local disturbance. The representation is simplified because it does not estimate the internal human states. It is data-informed because measured takeover events update the route-local profile. It is control-oriented because the learned profile enters the speed-planning calculation that determines the autonomy reference behavior. The hypothesis is that strategic, localized speed reduction can reduce takeover events for participants who intervene while preserving useful route performance and improving measured continuous-task performance.

The remainder of the chapter is organized as follows. Section 4.2 reviews related work on takeover modeling, intervention sensing, and learning for repetitive systems. Section 4.3 formulates the known-route learning problem, defines the simulation and human-subject protocol, and derives the backward-reachable speed calculation used to construct the learned reference profile. Section 4.4 presents the results of the handling-course, including cohort-level trends, sensitivity analysis of participant outliers, assumed-friction profiles that are learned, evolution of the takeover, and path-following behavior. Section 4.5 discusses the implications and limits of treating takeover as a disturbance rather than as a detailed human-state model.

4.2 Related Work

4.2.1 Human Intervention in Conditionally Automated Driving

Human intervention remains a central challenge for conditionally automated driving. In SAE Level 3 operation, the automated system performs the dynamic driving task within a defined operational design domain, but the human occupant may still need to resume or assert control under some conditions. Therefore, a large body of literature studies the takeover time, the takeover quality, the driver's state, non-driving-related tasks and intervention behavior [Eriksson and Stanton, 2017, Zhang et al., 2019, Sekadakis and Yannis, 2025, Sekadakis et al., 2025, Gaspar and Carney, 2019, Yu et al., 2025a]. These studies show that the human response to automation is not governed by a single scalar variable. It depends on the driving scenario, task demand, environmental context, driver attention, interface design, and interaction history with automation.

That complexity motivates two complementary research directions. One direction seeks richer human models: driver attention, gaze, hand position, torque input, intent, trust, or

takeover readiness can be estimated or inferred using sensors and behavioral data [Wang et al., 2018, Zhong et al., 2025]. The second direction, pursued here, asks whether a lower-dimensional control proxy can reduce repeated intervention when the human model is not known. This chapter follows the second direction. It does not claim that intervention is fully explained by speed, curvature, friction margin, lateral acceleration, or any other single variable. It treats intervention as an output of the coupled human–autonomy system and tests whether local speed reduction can reduce repeated takeover events for the tested participants and route.

4.2.2 From Explanatory Human Models to Control-Oriented Disturbance Rejection

An explanatory model of human intervention asks why the driver intervened. The control problem in this chapter is narrower: given that an intervention occurred at a route location during a previous iteration, how should the autonomy modify the next reference behavior at that location? A detailed explanation may require variables that are difficult to measure, unavailable in real time, or poorly identified for a specific driver. A control-oriented disturbance representation can instead use the intervention event itself as the measured signal.

In this interpretation, takeover is analogous to a repetitive disturbance observed during a repeated task. The disturbance is not additive in the vehicle equations of motion in the conventional sense. It is a discrete event that interrupts or modifies the autonomous control loop. If the event recurs under similar context, the autonomy can modify the feedforward reference behavior before the disturbance reappears. This approach does not remove the need for human modeling in general, and it does not assert that reduced speed is the correct response to every takeover event. It tests whether a route-local speed-policy adaptation is sufficient to reduce repeated intervention for the studied drivers and route.

4.2.3 Iterative Learning Control for Repeated Tasks

iterative learning control uses data from previous repetitions of a task to improve the input applied on later repetitions. Early work established the idea that previous-trial error can be used to improve the next operation of a robotic system [Arimoto et al., 1984]. Later surveys formalized iterative learning control as a learning-based method for high-performance tracking in repeated tasks, including convergence, robustness, and learning transients [Bristow et al., 2006]. The setting in this chapter differs from classical trajectory-tracking iterative learning control in two ways. First, the learned quantity is not the low-level vehicle command. Second,

the error signal is not a continuous tracking error, but an event-based signal indicating that the human took over control.

Point-to-point tracking [Freeman and Tan, 2013] can simplify the trajectory tracking problem by selecting only certain points to track. This has further been expanded for systems where the critical point initial conditions are trial-varying yet can still guarantee robustness [Tao et al., 2020]. In the case of more complex trajectories, Pareto optimal methods can be used to determine optimal locations, optimal parameters, and thus optimal inputs for subsequent iterations [Lim and Barton, 2014]. This has been extended to ensure that conflicts in inputs can be resolved as multi-objective optimizaiton problems in complex systems [Zhang et al., 2025]. When applying economic optimizations to shared driving tasks, identifying parameters for route-segments tailored to the individual has also been shown to demonstrate better performance Yu et al. [2025b].

Despite these differences, the repeated-route structure preserves the central idea of iterative learning control: previous-iteration information changes future-iteration behavior. The learned input in this chapter is a map of assumed coefficient-of-friction values indexed by route turn and route part. Reducing the assumed coefficient of friction lowers the backward-reachable speed in that local region and creates a more conservative reference profile in subsequent laps. This preserves the existing autonomy architecture while providing a route-local learning mechanism to reduce repeated human takeover.

4.2.4 Route-Local Speed Adaptation as a Simplified Model

The proposed method uses route-local assumed friction as a proxy variable. The real road friction in the simulation does not need to change. Instead, the assumed friction used by the speed planner is reduced in selected route bins. The learned friction profile is therefore not an estimate of physical tire–road friction. It is a compact representation of how conservative the autonomy should be at a given route location for the tested participant.

This is a simplified, data-informed model. The route-local assumed friction profile is incomplete by design. It does not explain human behavior and does not capture every factor that can cause intervention. It is useful if it changes the speed profile in the route bins where the intervention history indicates that the nominal profile was not accepted by a participant. To evaluate the proposed method, this chapter must demonstrate that the route-local update accomplishes three goals: it modifies the speed profile, reduces measured takeovers for intervention-prone participants, and avoids unacceptable degradation in overall route performance.

4.3 Methods

4.3.1 Problem Formulation

Consider a closed reference route parameterized by station $s \in [0, S)$, where S is the route length. The route geometry is represented by a centerline position $\boldsymbol{p}(s) = [x(s), y(s)]^\top$, heading $\psi(s)$, and signed curvature $\kappa(s)$. The route is partitioned into labeled bins,

$$\mathcal{B} = \{(j, p) : j = 1, \ldots, N_T,\ p = 1, \ldots, N_P(j)\}, \tag{4.1}$$

where j indexes the turn and p indexes the part of that turn. Each path point is assigned to a route bin through

$$b(s) = (j(s), p(s)) \,. \tag{4.2}$$

The segmentation is fixed throughout the known-route experiment. Figure 4.2 shows the handling-course route, start/end location, turn labels, and part labels used for the learning update.

The autonomy follows a reference path and speed profile in TruckSim. TruckSim 2024.2 x64 Revision 235317, dated July 16, 2024, running the "Off-Road Pickup: Ind_SA Throttle Lag {Pickup Trucks}" vehicle model with the 265/75 R16 Off-Road tire model on a paved surface with true road friction $\mu = 1.0$, was used throughout the study. Hardware input was provided through a Logitech G29 steering wheel and pedal set. User input, TruckSim interfacing, and data logging were implemented in MATLAB/Simulink. Vehicle and planner parameters used by the simplified controller and backward-reachable speed planner are summarized in Table 4.1. Parameters implemented only inside the high-fidelity TruckSim model are distinguished from quantities used directly by the simplified planner when needed.

For iteration k, the route-local learned quantity is an assumed coefficient-of-friction map

$$\boldsymbol{\mu}^{(k)} = \{\mu_{j,p}^{(k)} : (j, p) \in \mathcal{B}\}. \tag{4.3}$$

This map is converted to a path-indexed friction profile by assigning

$$\mu^{(k)}(s_i) = \mu_{j(s_i),p(s_i)}^{(k)} \tag{4.4}$$

at each discrete route point s_i. The physical road friction in TruckSim is not changed by this update. The learned quantity is the assumed friction used by the planner. A lower value reduces the planner's assumed acceleration envelope and produces a more conservative speed reference.

Table 4.1: Vehicle and planner parameters used in the known-route human-intervention disturbance-rejection study.

Symbol	Description	Value	Notes
m	Vehicle mass	1306 kg	Total mass of the TruckSim vehicle and the mass used in planner calculations where needed.
I_z	Yaw moment of inertia	1816 kg m^2	TruckSim vehicle parameter.
a	Distance from CG to front axle	1.35 m	TruckSim vehicle parameter.
b	Distance from CG to rear axle	1.43 m	TruckSim vehicle parameter.
L	Wheelbase, $L = a + b$	2.78 m	TruckSim vehicle parameter.
g	Gravitational acceleration	9.81 m/s^2	Used in the friction-limited speed calculation and TruckSim simulation.
μ_0	Initial assumed coefficient of friction	0.7	Nominal planner value before learning.
$\mu_{\max}$	Upper bound on learned assumed friction	0.7	Selected as the study's pre-test safe operating ceiling for the handling course.
$\mu_{\min}$	Lower bound on learned assumed friction	None imposed	The implemented update used multiplicative reductions without an explicit lower saturation bound; repeated reductions remain positive but can become very conservative.

Human takeover is represented by a route-local event signal. Let $h^{(k)}(t) \in \{0, 1\}$ denote whether the participant has taken over during iteration k at time t. Through the vehicle position and route projection, this event is mapped to route bins. For bin (j, p), define

$$D_{j,p}^{(k)} = \begin{cases} 1, & \text{if a takeover event intersects bin } (j,p) \text{ during iteration } k, \\ 0, & \text{otherwise,} \end{cases} \tag{4.5}$$

and let $T_{j,p}^{(k)}$ denote the total takeover duration assigned to that bin. If a takeover event spans multiple parts or turns, each intersected bin is marked as having an event for route-local learning and heatmap visualization. For iteration-level event counts, the same physical takeover is counted once even if it spans multiple bins. This convention separates the route-local disturbance map from the aggregate event-count metric.

4.3.2 Backward-Reachable Speed Formulation

The speed planner computes a friction-limited route speed using local curvature and the assumed coefficient-of-friction profile. For a vehicle traveling at speed v through a path point with curvature κ, the lateral acceleration magnitude is approximated by

$$a_y(s) = v(s)^2 |\kappa(s)|. \tag{4.6}$$

The local lateral acceleration constraint is

$$v(s)^2 |\kappa(s)| \leq \mu(s) g. \tag{4.7}$$

Thus, the local curvature-limited speed cap is

$$v_{\text{lat}}(s_i) = \begin{cases} \sqrt{\dfrac{\mu(s_i) g}{|\kappa(s_i)|}}, & |\kappa(s_i)| \geq \epsilon_\kappa, \\ v_{\text{straight}}, & |\kappa(s_i)| < \epsilon_\kappa, \end{cases} \tag{4.8}$$

where ϵ_κ prevents division by zero and v_{straight} is a straight-section speed cap.

The backward-reachable pass accounts for the longitudinal deceleration required to slow down before future high-curvature regions. At each point, the total acceleration is limited by a friction-circle approximation,

$$a_x(s)^2 + a_y(s)^2 \leq \left(\mu(s) g\right)^2. \tag{4.9}$$

For discrete path points s_i with segment length $\Delta s_i = s_{i+1} - s_i$, the available longitudinal deceleration at point $i+1$ is computed from the remaining acceleration envelope,

$$\bar{a}_{x,\mathrm{br}}(s_{i+1}) = \sqrt{\max\left((\mu(s_{i+1})g)^2 - \left(\left(v_{i+1}^{\mathrm{back}}\right)^2 |\kappa(s_{i+1})|\right)^2,\ 0\right)}. \tag{4.10}$$

The backward pass is then

$$v_i^{\mathrm{back}} = \min\left(v_{\mathrm{lat}}(s_i), \sqrt{\left(v_{i+1}^{\mathrm{back}}\right)^2 + 2\bar{a}_{x,\mathrm{br}}(s_{i+1})\Delta s_i}\right). \tag{4.11}$$

This equation enforces that the vehicle can brake from point i to satisfy the reachable speed at point $i+1$ under the assumed friction envelope.

A forward pass is applied to enforce longitudinal acceleration feasibility. The available drive acceleration at point $i-1$ is

$$\bar{a}_{x,\mathrm{drv}}(s_{i-1}) = \sqrt{\max\left((\mu(s_{i-1})g)^2 - \left(\left(v_{i-1}^{\mathrm{ref}}\right)^2 |\kappa(s_{i-1})|\right)^2,\ 0\right)}. \tag{4.12}$$

The forward-pass speed is

$$v_i^{\mathrm{ref}} = \min\left(v_i^{\mathrm{back}}, \sqrt{\left(v_{i-1}^{\mathrm{ref}}\right)^2 + 2\bar{a}_{x,\mathrm{drv}}(s_{i-1})\Delta s_{i-1}}\right). \tag{4.13}$$

For the closed handling course, the path mesh is preserved and the closed-route boundary is handled consistently across the start/end point. In implementation, the route arrays can be duplicated during the backward and forward passes and then reduced back to the original mesh after the speed profile is computed. This preserves the original path points, route length, curvature values, and turn/part labels while replacing the speed row used by the autonomy.

The choice $\mu_{\max} = 0.7$ defines the upper bound of the learned profile. It should not be interpreted as a claim that every participant preferred that exact aggressiveness level. Instead, $\mu = 0.7$ was selected through pre-study engineering evaluation as the highest assumed-friction value at which the autonomous vehicle remained on the handling course under the implemented TruckSim vehicle, controller, and route conditions. The value is therefore a planner safety ceiling rather than an estimate of either the true road friction or a participant's desired aggressiveness. It was intentionally aggressive enough to create a mismatch between the nominal autonomy behavior and some participants' perceived safety or comfort. In practice, a manufacturer would likely impose a more conservative operating envelope than one allowing lateral path errors on the order of several meters. The role of this ceiling in the present study is to provide a repeatable upper-bound baseline from which route-local

conservative adjustments can be learned.

4.3.3 Learning Update for the Assumed Friction Map

The takeover signal is converted into an update to the assumed friction map. The implemented update is event based: a bin is modified if a takeover intersects that bin during the previous iteration. Let $n_{j,p}^{(k)}$ denote the cumulative number of reductions applied to bin (j,p) up to iteration k. The event-dependent reduction factor is

$$r\left(n_{j,p}^{(k)}\right) = \begin{cases} 0.20\mu_0, & n_{j,p}^{(k)} = 0, \\ 0.10\mu_0, & n_{j,p}^{(k)} = 1, \\ 0.05\mu_0, & n_{j,p}^{(k)} \geq 2. \end{cases} \tag{4.14}$$

The implemented update also allows recovery toward the nominal ceiling when a previously reduced bin does not contain a takeover event. Recovery is represented as an additive coefficient-of-friction increment $\Delta_r = 0.01\mu_0$ applied once per iteration to eligible bins. The upper-saturation operator is

$$\mathrm{sat}_{\max}(x) = \min(x, \mu_{\max}). \tag{4.15}$$

The route-bin update law is

$$\mu_{j,p}^{(k+1)} = \begin{cases} \left[1 - r\left(n_{j,p}^{(k)}\right)\right] \mu_{j,p}^{(k)}, & D_{j,p}^{(k)} = 1, \\ \mathrm{sat}_{\max}\left(\mu_{j,p}^{(k)} + 0.01\mu_0\right), & D_{j,p}^{(k)} = 0 \text{ and } \mu_{j,p}^{(k)} < \mu_{\max}, \\ \mu_{j,p}^{(k)}, & D_{j,p}^{(k)} = 0 \text{ and } \mu_{j,p}^{(k)} = \mu_{\max}. \end{cases} \tag{4.16}$$

The first intervention in a bin reduces the assumed friction by 20% of the initial assumed coefficient of friction, $\mu_0 = 0.7$, the second event reduces μ by 10%, and all subsequent interventions result in a reduction of μ by 5%. This update law reduces the assumed friction in bins associated with takeover events and thus lowers the next-iteration speed reference through Eqs. (4.8)–(4.13).

4.3.4 Simulation and Human-Subject Testing Protocol

The study was conducted on a laboratory simulator using TruckSim as a vehicle-dynamics environment. The participants first received the study introduction and an introduction to the simulator controls, as shown in 4.1. They then completed a familiarization lap with access to both manual and automated modes. This lap was used for practice and was not



Figure 4.1: Experimental setup used for the experiment. Participants sit in a chair with steering wheel and pedal controls, while a tablet sits to the side. One the tablet is a picture and participants are required to count the number of objects in the pictures to earn points. A zoomed in picture of a typical SPOT task is presented on the top right.

included in the learning analysis.

Following the simulator familiarization, they completed a familiarization of the continuous task as presented by [Ali et al., 2024]. The familiarization provides participants with five images they must count objects in, including a picture without objects. Five types of objects are presented: people, vehicles, bags, barrels, and antennas. A participant scores a point for every correctly identified object in the picture, and a bonus point for correctly identifying all objects in the image.

After familiarization, each participant completed one measured baseline lap using the nominal assumed friction profile $\boldsymbol{\mu}^{(0)} = \mu_0 \mathbf{1}$. The reference-speed profile for the vehicle in TruckSim was then updated based on the measured takeover events, and the participant drove another measured lap with the updated profile. This iterative sequence continued until a stopping condition was met. In the current study, participants were required to complete at least five measured iterations and no more than fourteen. The loop stopped early if the convergence criteria indicated that the participant had reached a stable learned profile, including cases in which consecutive measured laps produced no takeover events. Convergence was reached if two consecutive laps showed no intervention or if 14 measured

iterations were completed.

The participant instructions deliberately framed the takeover as a safety- and comfort-based decision rather than a performance-optimization problem. Participants were instructed to imagine riding in a large truck operating as an SAE Level 3 vehicle. They were told to take control whenever they felt the vehicle's behavior was unsafe or uncomfortable, and to maximize their performance on the continuous tablet task while remaining safe. Additionally, participants were informed that the vehicle's behavior might vary from lap to lap. This instruction set defines the measured intervention signal: takeover indicates that the participant was unwilling to continue with the current autonomy behavior under simultaneous route-following and secondary-task demand.

The primary intervention metrics are the number of distinct takeover events and the total takeover time per measured iteration. The primary task-performance metrics are lap time and continuous-task score. Lateral error is reported as a supporting vehicle-performance metric. All human-subject tests were conducted under the appropriate institutional review or exemption, and all tests were performed in a manner that was safe for participants. The study was granted an exception to full IRB review by the U-M IRB under Study ID HUM00282808, Adaptive Variable Intervention Driver Training (AVID Training), Amendment ID Ame00174441.

4.3.5 Analysis Plan and Sensitivity Checks

Data are reported using aggregate views and participants' perspectives. The entire cohort is retained to show the full range of participant behavior. A sensitivity analysis excludes the especially cautious participant whose takeover frequency and total takeover time dominate the full-cohort mean and variance. This exclusion is not used to remove inconvenient data; it is used to show whether the central learning trend persists when a single high-leverage participant is removed. The other comparatively cautious participants are retained in the sensitivity cohort because they represent a plausible variation in participant tolerance rather than a qualitatively distinct use of the simulator.

A second subgroup contains participants who never took over during the handling-course experiment. For these participants, the event-based learning rule does not receive a route-local correction signal. Their continuous-task scores are still reported, but any score improvement should be interpreted as familiarization with the task, simulator, or automation, rather than as evidence that takeover-based disturbance rejection occurred.

The primary quantitative analysis compares each participant's measured baseline lap with that participant's final measured lap. Lap time, continuous-task score, peak lateral error, and

mean lateral error are evaluated using paired baseline-to-final differences with paired t-tests, bootstrap confidence intervals for the paired mean difference, and small-sample corrected paired standardized effect size g_z. The takeover count and total takeover time are discrete or zero-inflated results, so the primary p-value for these results is from the Wilcoxon signed-rank test. The same paired summaries are reported for the full cohort, the sensitivity cohort excluding the high-intervention participant, and the non-intervention subgroup where the metric is meaningful. Iteration-wise figures provide the longitudinal context for the paired summaries.

4.4 Results

This section reports on the learning results of the handling-course. The analysis tests whether route-local speed profile adaptation reduced takeover for participants who intervened, whether continuous-task performance improved across the repeated laps, and whether the resulting speed profile preserved acceptable vehicle performance. The results are intentionally separated into full-cohort, sensitivity-cohort, high-intervention, and non-intervention views because the participant data contain meaningful heterogeneity rather than a single representative response.

4.4.1 Handling-Course Segmentation

The handling course was divided into route sections that support both event attribution and local updates to the assumed coefficient of friction. Figure 4.2 shows the start/end location, turn boundaries, part boundaries, and route direction. This segmentation is used to assign takeover events to route bins and to apply local coefficient-of-friction updates.

4.4.2 Participant Groups and Sensitivity Analysis

The participant data contain two types of behavior that strongly affect interpretation. First, three participants were comparatively cautious, and one of these participants produced substantially more takeover events than the rest of the cohort. That participant strongly affects full-cohort means and standard deviations for lap time, takeover count, and total takeover time. For that reason, the results report both the full cohort and a sensitivity cohort that excludes only this one high-intervention participant. The other two comparatively cautious participants remain in the sensitivity cohort because they appear to represent plausible participant variation rather than a qualitatively different experimental condition.

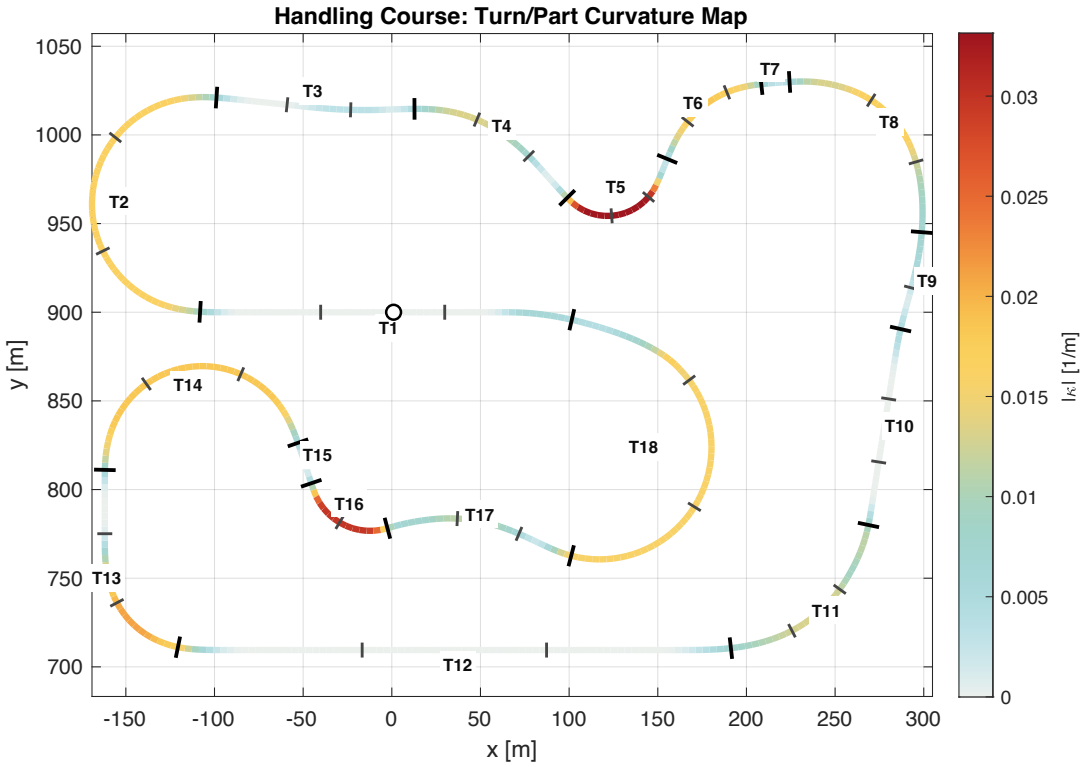

Figure 4.2: Handling-course segmentation used for known-route learning. Eighteen turns are marked from T1 to T18, with bold hash lines marking turn boundaries and thin hash lines marking internal turn partitions. The start/end point is marked by a circle, and curvature magnitude is shown in color.

Second, five participants never took over during the handling-course experiment. For these participants, the learning algorithm received no takeover-based update signal. This does not imply that the nominal profile fully represented their preferences. This means only that the measured disturbance signal used in this chapter did not reveal a route-local incompatibility between the nominal profile and their takeover threshold. Their continuous-task scores remain informative, but the score improvement for this subgroup should be interpreted separately from the takeover-based disturbance rejection. Table 4.2 summarizes the groups used for the interpretation of the results.

Table 4.2: Participant groups used to interpret the known-route learning results.

Group	Size	Used for	Interpretation
Full cohort	18	Primary transparency view	Includes all participant behavior and shows the full experimental spread.
Sensitivity cohort	17	Main cohort trend	Excludes only the one high-intervention participant to reduce leverage from a single participant.
Comparatively cautious participants retained in sensitivity cohort	2	Variability check	Participants with elevated takeover relative to the cohort, but not treated as separate exclusions.
High-intervention participant	1	Exemplar and sensitivity case	Used to show a case where speed reduction alone did not fully remove intervention.
Non-intervention subgroup	5	Boundary case	Participants with no takeover signal; route-local disturbance rejection cannot infer additional profile information from takeover.

The paired baseline-to-final summary statistics are given in Table 4.3. These values provide the primary quantitative comparison, while the following figures show the iteration-wise evolution of the same outcomes.

Table 4.3 supports three direct observations. First, the continuous-task score increased for the full cohort and the sensitivity cohort, and also increased for the non-intervention subgroup. Second, the takeover count and total takeover time decreased for the full and sensitivity cohorts. Third, lap time did not improve; the sensitivity cohort showed essentially no baseline-to-final change in lap time. The mean lateral error decreased, while the peak lateral error decreased numerically but should be interpreted cautiously because the paired-test p-values

Table 4.3: Paired baseline-to-final outcome statistics for the handling-course learning experiment. Differences are computed as final minus baseline. Negative differences indicate reductions in lap time, takeover count, takeover time, and lateral error; positive differences indicate improvement in continuous-task score.

Outcome	Cohort	N	Baseline	Final	Difference	95% Diff CI	p	g_z
Lap time (s)	Full	18	119.1 ± 11.8	121.3 ± 11.3	2.2 ± 15.0	$[-4.2,\ 9.2]$	0.534	0.143
	Sens.	17	119.4 ± 12.1	119.3 ± 7.2	-0.1 ± 11.5	$[-5.6,\ 5.0]$	0.966	−0.010
	Non-int.	5	111.2 ± 0.0	111.2 ± 0.0	0.0 ± 0.0	$[0.0,\ 0.0]$	1.000	–
Continuous-task score	Full	18	87.1 ± 33.8	151.2 ± 32.7	64.1 ± 39.6	$[46.3,\ 82.2]$	< 0.001	1.547
	Sens.	17	88.9 ± 33.9	150.4 ± 33.5	61.4 ± 39.1	$[44.1,\ 79.5]$	< 0.001	1.498
	Non-int.	5	96.4 ± 22.4	127.2 ± 24.0	30.8 ± 20.9	$[14.6,\ 47.0]$	0.030	1.180
Takeover count	Full	18	1.39 ± 2.00	0.11 ± 0.32	-1.28 ± 2.05	$[-2.28,\ -0.44]$	0.016	−0.595
	Sens.	17	1.41 ± 2.06	0.06 ± 0.24	-1.35 ± 2.09	$[-2.35,\ -0.47]$	0.016	−0.617
	Non-int.	5	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	$[0.00,\ 0.00]$	1.000	–
Total takeover time (s)	Full	18	14.45 ± 21.78	0.85 ± 2.93	-13.60 ± 22.17	$[-24.39,\ -4.61]$	0.008	−0.586
	Sens.	17	14.66 ± 22.43	0.72 ± 2.96	-13.94 ± 22.81	$[-25.25,\ -4.47]$	0.016	−0.582
	Non-int.	5	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	$[0.00,\ 0.00]$	1.000	–
Peak lateral error (m)	Full	18	11.81 ± 17.51	3.33 ± 3.59	-8.48 ± 18.96	$[-17.82,\ -1.26]$	0.075	−0.427
	Sens.	17	12.16 ± 17.98	3.49 ± 3.62	-8.67 ± 19.53	$[-18.65,\ -1.01]$	0.086	−0.423
	Non-int.	5	4.57 ± 0.00	4.57 ± 0.00	0.00 ± 0.00	$[0.00,\ 0.00]$	1.000	–
Mean lateral error (m)	Full	18	1.10 ± 1.44	0.28 ± 0.17	-0.82 ± 1.51	$[-1.58,\ -0.23]$	0.035	−0.517
	Sens.	17	1.13 ± 1.48	0.29 ± 0.17	-0.84 ± 1.55	$[-1.63,\ -0.22]$	0.041	−0.512
	Non-int.	5	0.40 ± 0.00	0.40 ± 0.00	0.00 ± 0.00	$[0.00,\ 0.00]$	1.000	–

Notes: Full = all participants; Sens. = sensitivity cohort excluding the high-intervention participant; Non-int. = participants who did not intervene during the handling-course experiment. Values are mean ± standard deviation. Confidence intervals are participant-level bootstrap 95% confidence intervals for the paired mean difference. The primary p-value is from a paired t-test for lap time, continuous-task score, and lateral-error outcomes, and from the Wilcoxon signed-rank test for takeover count and total takeover time. g_z is the small-sample corrected paired standardized effect size; it is omitted when the paired-difference standard deviation is zero.

are above 0.05.

4.4.3 Lap-Time and Continuous-Task Performance

Lap time and continuous-task score provide complementary performance metrics. The lap time measures the cost of conservative speed-profile adaptation, while the continuous-task score measures the participant's ability to remain engaged with the secondary task while monitoring the automation. The desired outcome is not simply a slower route. The desired outcome is a route-local profile that reduces takeover while maintaining a reasonable lap-time penalty and improving secondary-task performance.

Figures 4.3 and 4.4 show lap times for the full and sensitivity cohorts, respectively. Figures 4.5 and 4.6 show the corresponding continuous-task scores. The full-cohort figures show the complete distribution of the participants. The sensitivity-cohort figures remove only the high-intervention participant and, therefore, provide a clearer view of the cohort-level trend for the remaining participants.

The paired statistics clarify the visual trends. The final lap time of the full cohort was 121.3 ± 11.3 s, with a baseline-to-final difference of 2.2 ± 15.0 s and a 95% confidence interval of $[-4.2,\ 9.2]$ s. The sensitivity cohort final lap time was 119.3 ± 7.2 s, with a baseline-to-final difference of -0.1 ± 11.5 s and a 95% confidence interval of $[-5.6,\ 5.0]$ s. Thus, the learned profile did not produce a meaningful lap-time reduction. In contrast, the continuous-task score increased from 87.1 ± 33.8 to 151.2 ± 32.7 for the full cohort and from 88.9 ± 33.9 to 150.4 ± 33.5 for the sensitivity cohort. The corresponding paired differences were 64.1 ± 39.6 and 61.4 ± 39.1, with confidence intervals that did not include zero.

The high-intervention participant and the non-intervention subgroup are useful exemplars because they show the limits of using takeover as the only learning signal. Figures 4.7 and 4.9 isolate the high-intervention participant. Figures 4.8 and 4.10 summarize the subgroup that did not intervene. The high-intervention participant shows that a participant can continue to intervene even after local speed reductions. The non-intervention subgroup shows the opposite boundary case: the autonomy receives no route-local correction signal, even though the participant may still improve at the continuous task over repeated laps.

4.4.4 Learned Coefficient-of-Friction Profile and Route Location of Takeover

The learned coefficient-of-friction profile identifies where the speed planner became more conservative. Figure 4.11 summarizes the route-local learned profile across the handling course. The figure should be interpreted as a speed-policy map. Locations with repeated

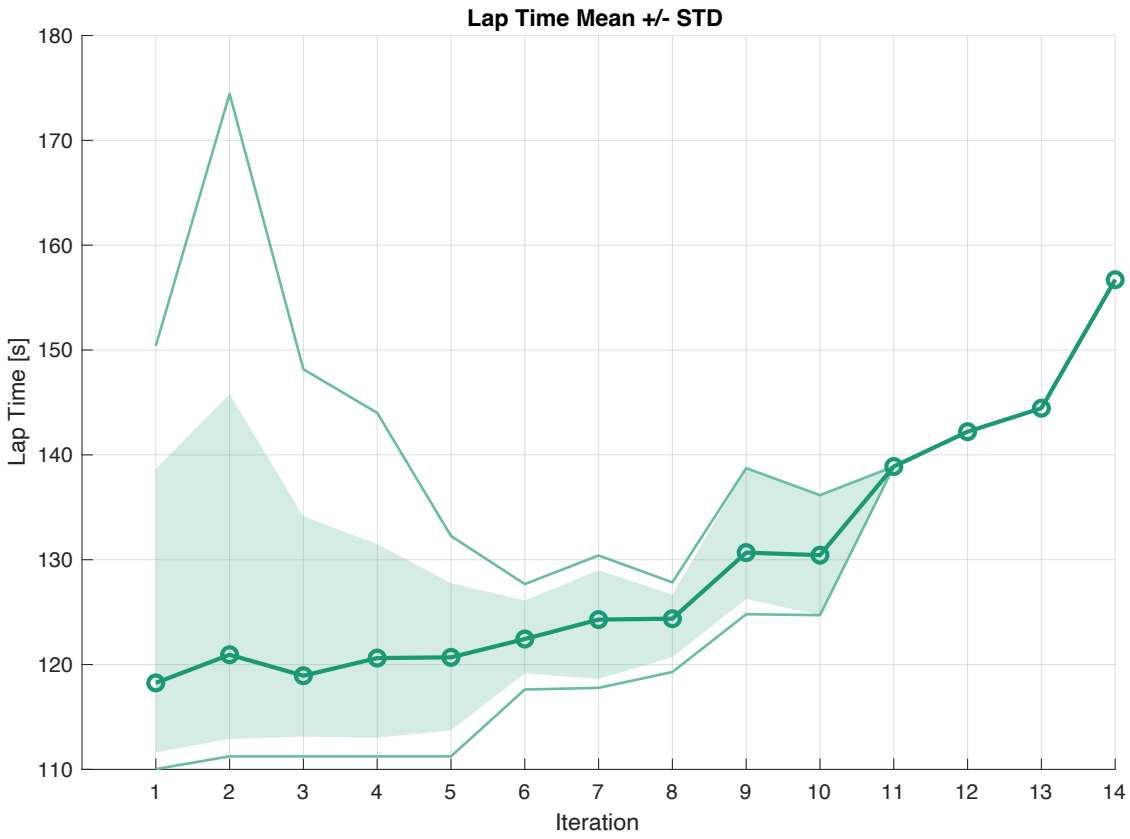

Figure 4.3: Lap-time evolution for the full cohort across handling-course iterations. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The mean trend is shown with a one-standard-deviation band, with participant-level variability retained to show the full experimental spread. Only the high-intervention participant intervened more than 10 times.

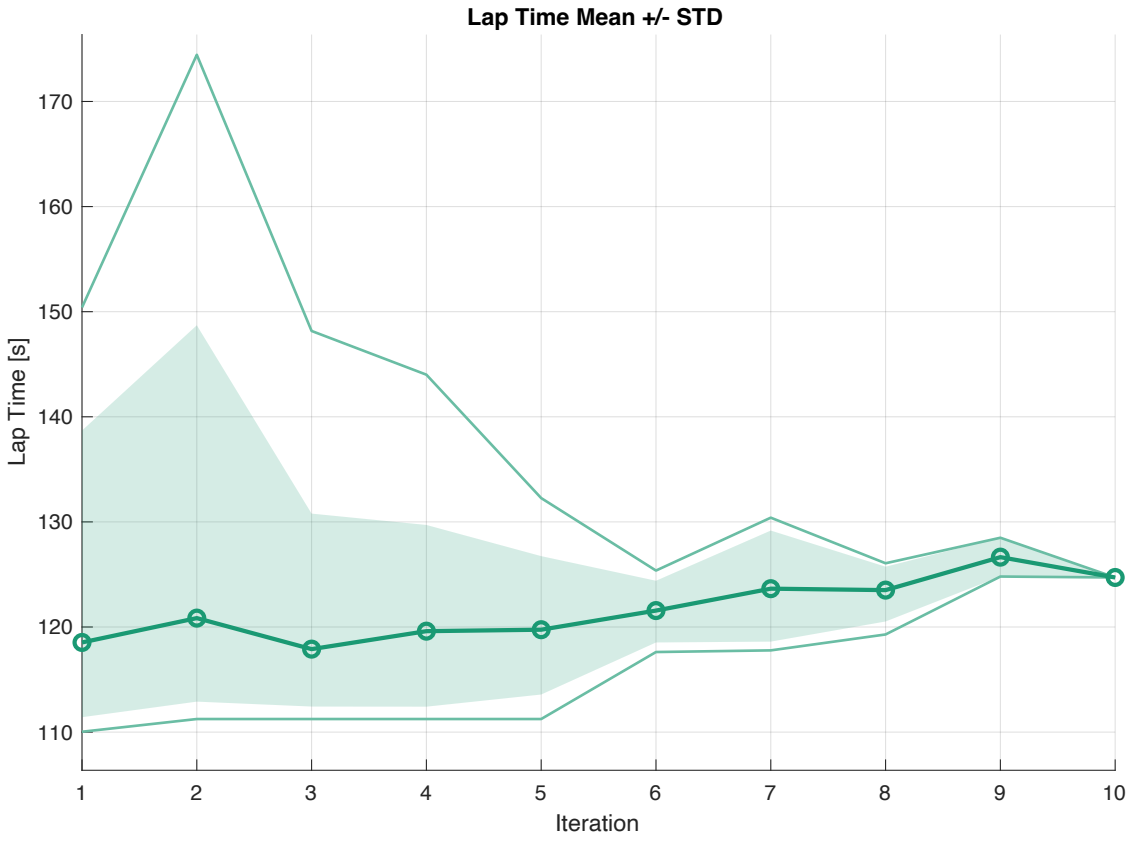

Figure 4.4: Lap-time evolution for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The sensitivity view shows that lap time does not materially change for the remaining cohort. This indicates the lap time cost for the adaptable profiles is low.

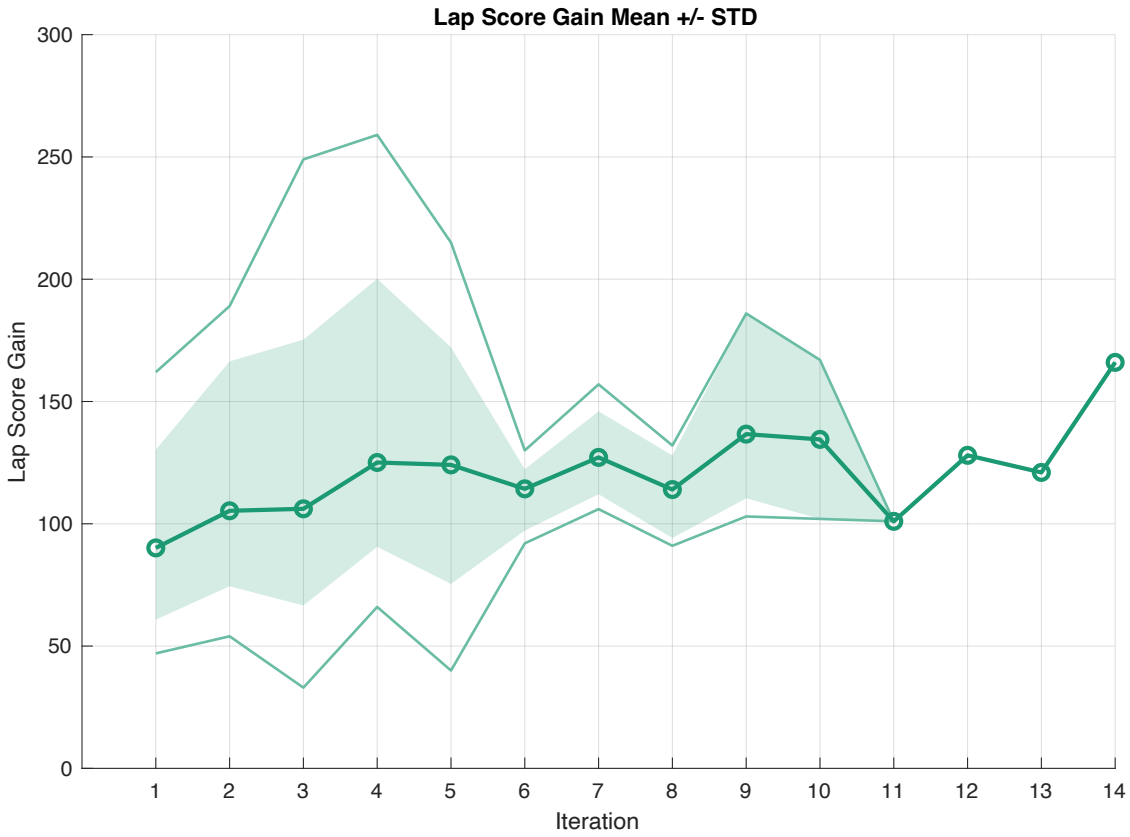

Figure 4.5: Continuous-task score evolution for the full cohort across handling-course iterations. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Scores increase across repeated laps while retaining substantial participant-level variation. Only the high-intervention participant intervened more than 10 times.

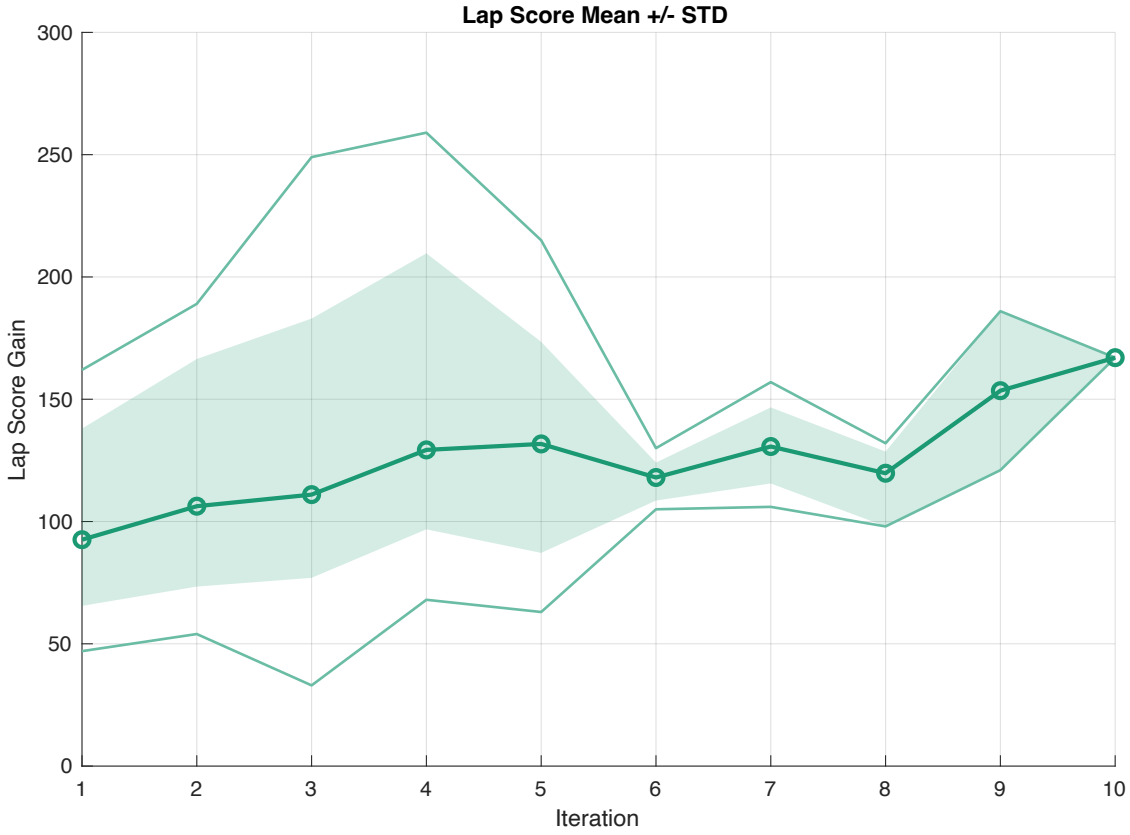

Figure 4.6: Continuous-task score evolution for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The trend remains upward after the high-intervention participant is removed. Nearly half of participants completed only 5 iterations, which explains the sharp decrease in standard deviation at the sixth iteration.

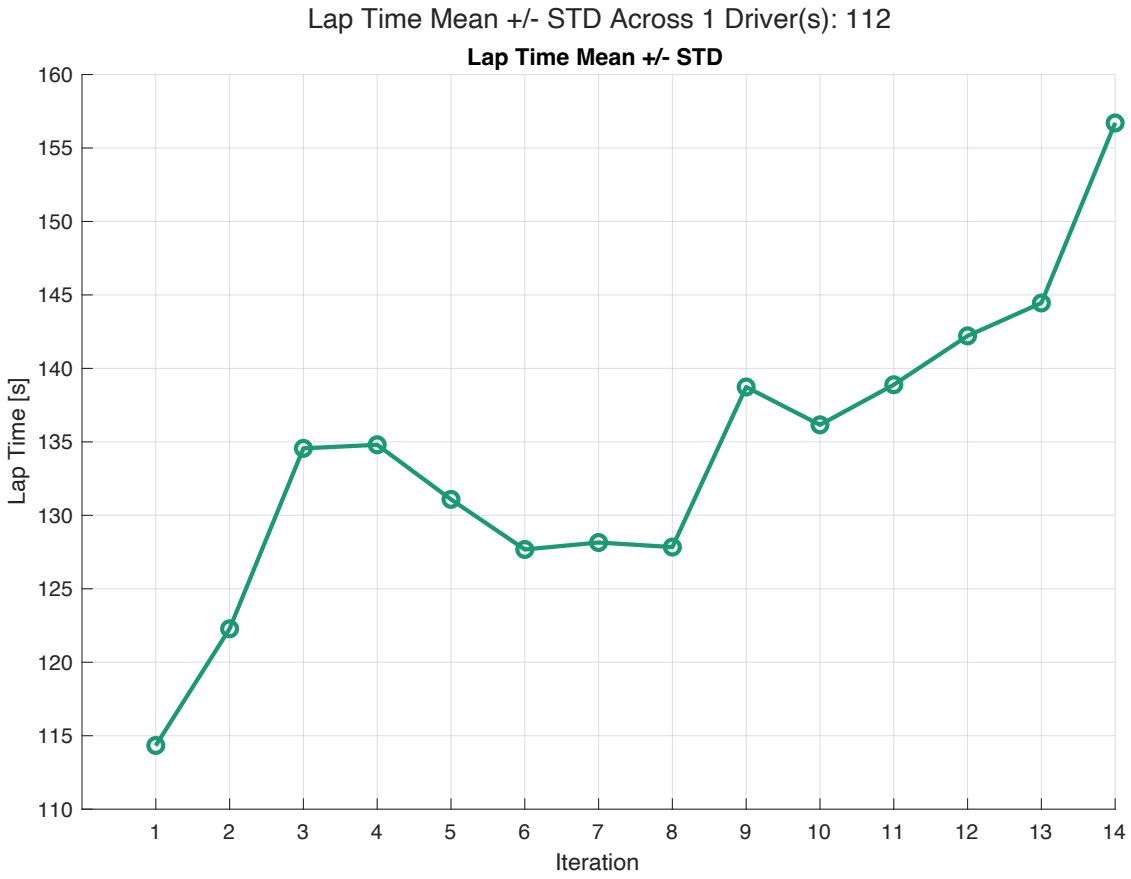

Figure 4.7: Lap-time evolution for the high-intervention participant across handling-course iterations. This exemplar shows a participant whose interventions persisted beyond the initial updates. This participant was more than 25 seconds slower than the average without them.

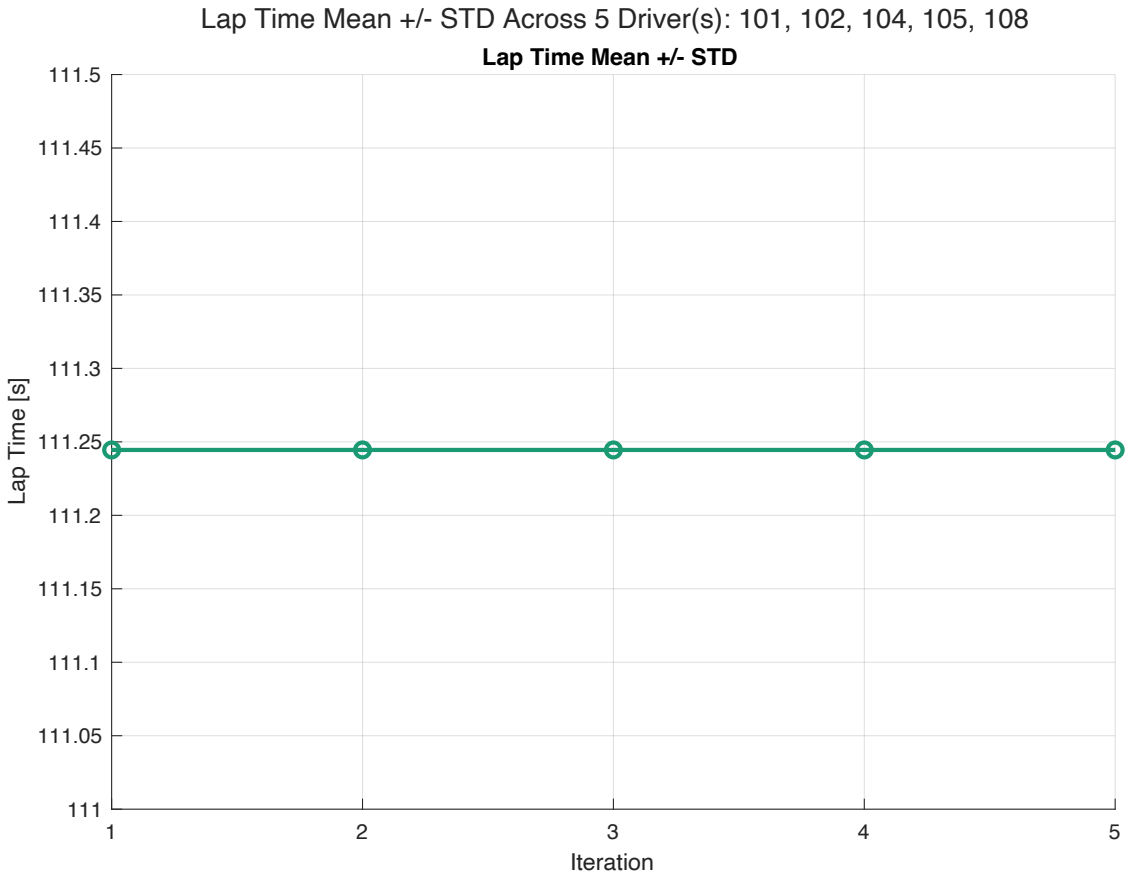

Figure 4.8: Lap-time evolution for the non-intervention subgroup. Because these participants did not take over, the route-local learning rule received no event signal for further profile adaptation.

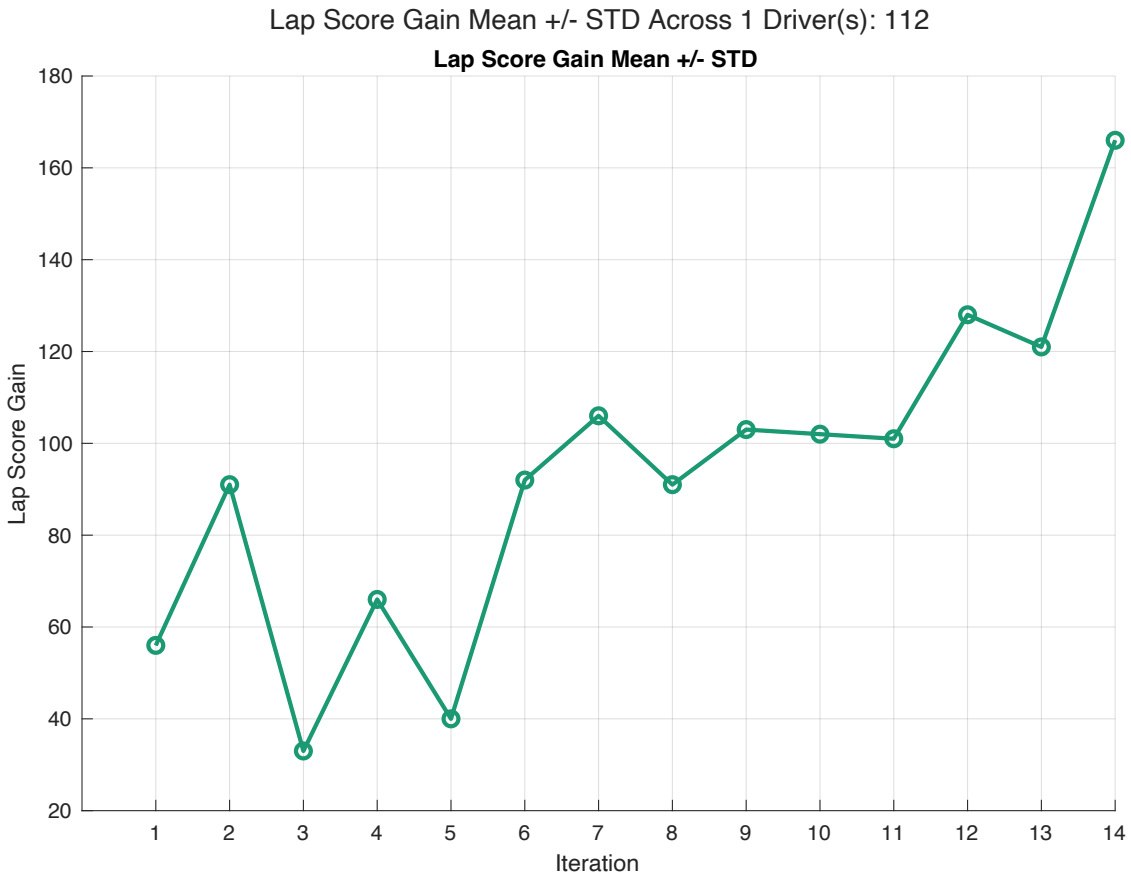

Figure 4.9: Continuous-task score evolution for the high-intervention participant across handling-course iterations. This single-participant exemplar is interpreted alongside the takeover plots because persistent intervention does not prevent improvement in the secondary-task metric. The large score increase at the end is partially explained by the longer lap duration.

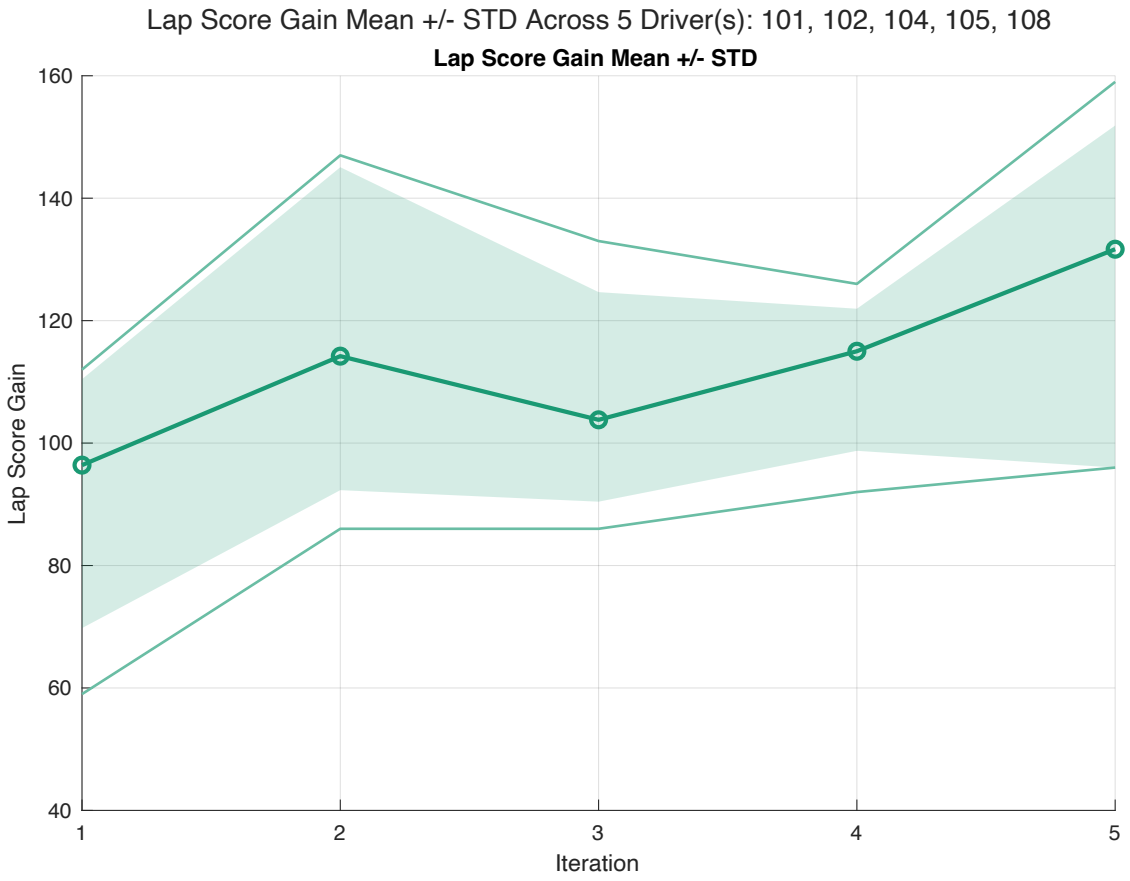

Figure 4.10: Continuous-task score evolution for the non-intervention subgroup. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Score improvement in this subgroup is a boundary case for the method because it occurs without takeover-based route-local updates. Despite not intervening in takeover tasks for all iterations, their later iteration task score increase suggests more attention on the continuous task.

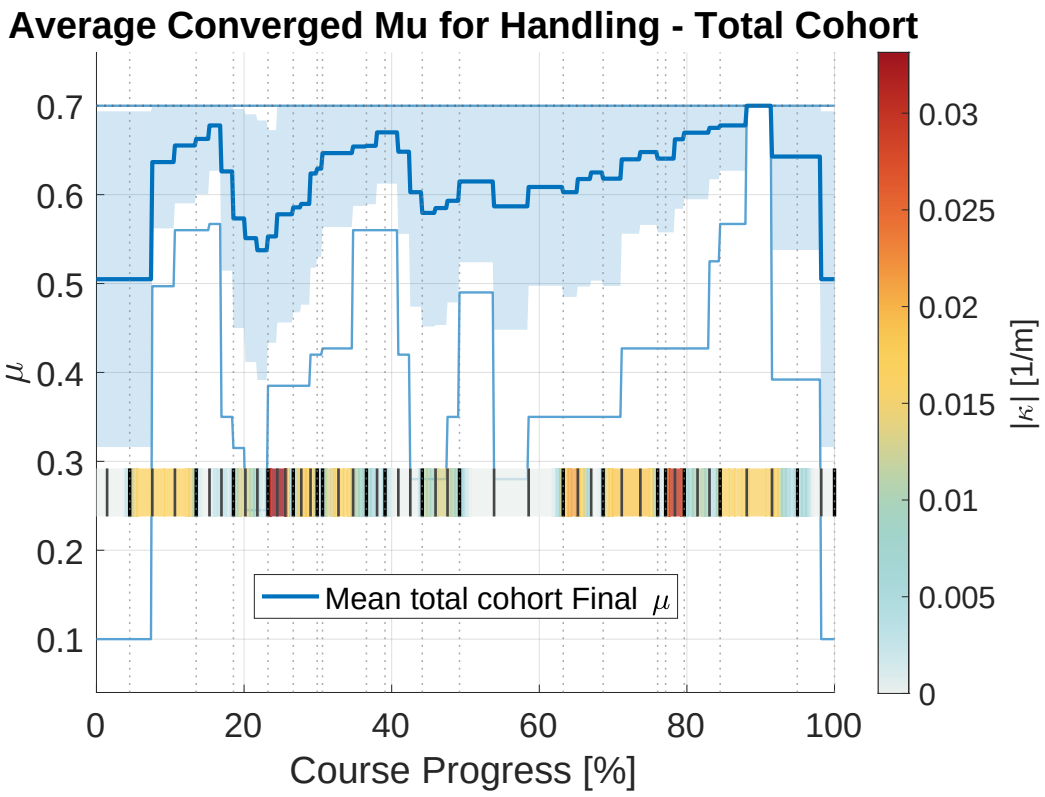

Figure 4.11: Route-local learned assumed-friction profile across the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Lower values indicate route bins where the speed planner became more conservative over the course of learning. The lowest μ values correspond to areas before and during sharp turns as well as before long straight sections.

takeover events tend to receive larger reductions in the assumed coefficient of friction, while locations with little or no takeover remain near the nominal safe-operating value.

Velocity averages are shown in Figure 4.12. The large negative standard deviations during the initial run show the many locations participants lost control and had to come to a stop or near stop to regain control. Despite the varied profiles, the mean velocities across the total cohort are fairly close for the initial and final runs, and the standard deviation for the final run is much tighter.

Figure 4.13 summarizes the spatial distribution of takeover events by course section. The distribution is not uniform, and takeover events cluster in specific route regions. Several high-count regions occur after high-speed approach sections, suggesting that curvature alone is not a sufficient descriptor of intervention likelihood. This observation is consistent with the route-local learning design: the update is tied to the measured takeover location rather than to a single global curvature threshold.

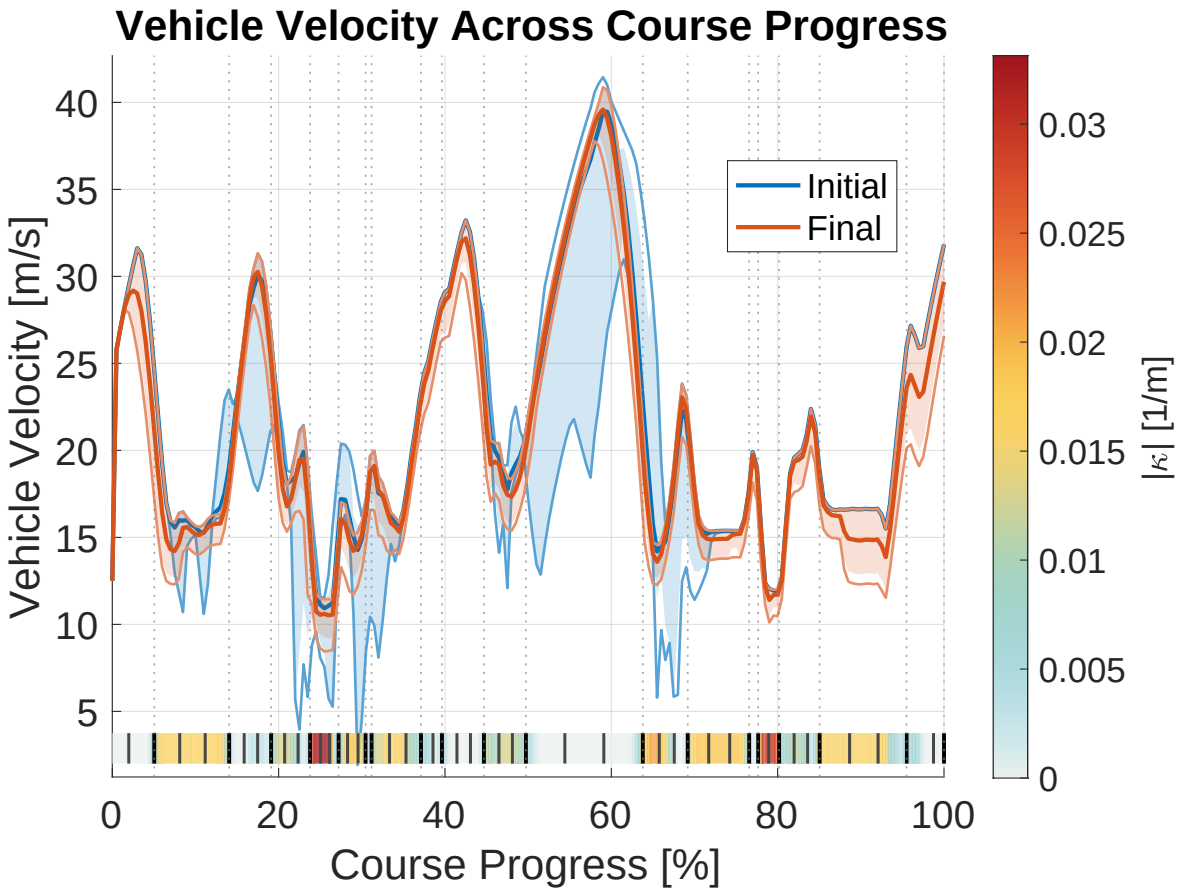

Figure 4.12: Velocity profile of the initial and final averages for the total cohort in the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Despite the varied mu profiles, standard deviations are fairly tight and mean differences appear mostly at the tightest curves.

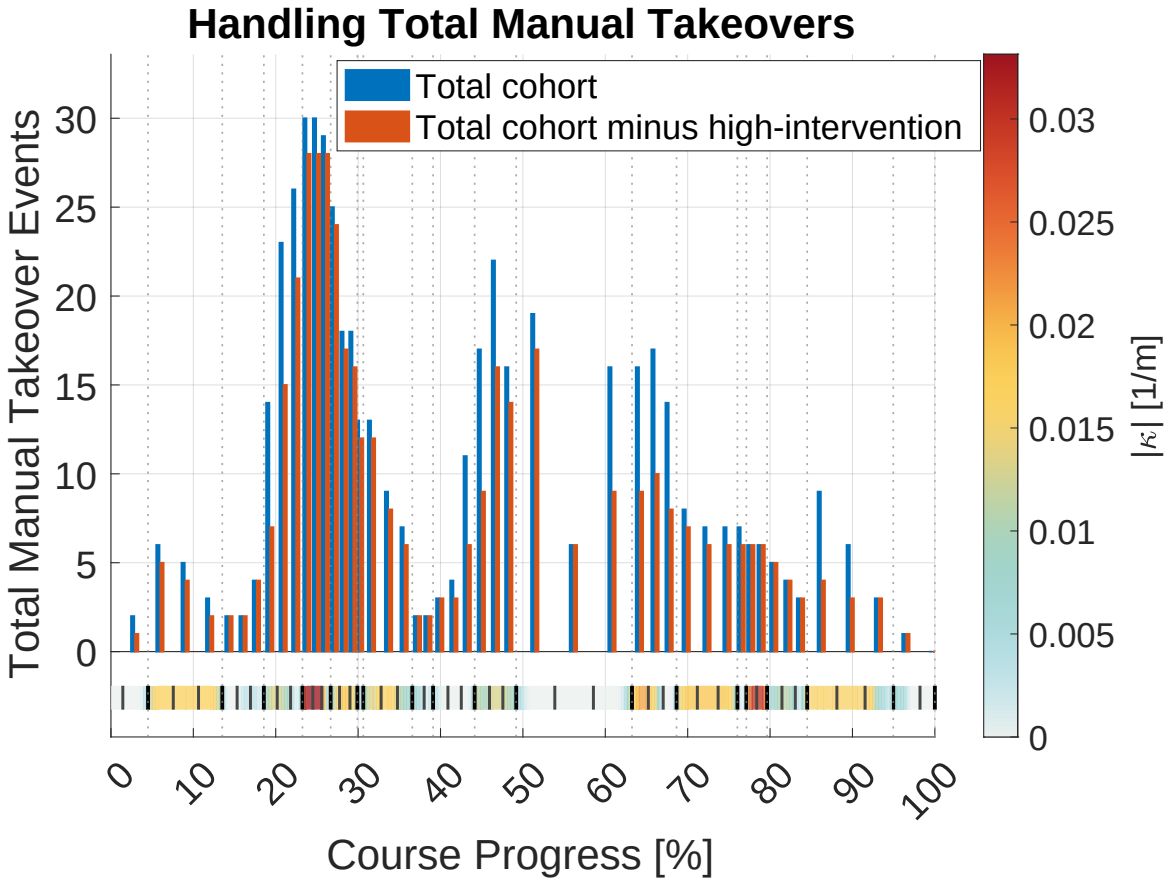

Figure 4.13: Total number of takeover events by section of the handling course. Takeover events cluster in selected route regions rather than being distributed uniformly across the course. Takeover events are clustered around the sharpest turn, T5, but also before the long straightaway, T12. The high-intervetion driver represents large proportions of takeover events at the end of T10 and early T11, as well as late T12 through T13.

4.4.5 Takeover Evolution Across Iterations

The main disturbance-rejection result is the evolution of the takeover count and the total takeover time between iterations. Let $N_{\mathrm{TO}}^{(k)}$ denote the number of distinct takeover events during iteration k, and let

$$T_{\mathrm{TO}}^{(k)} = \sum_{r=1}^{N_{\mathrm{TO}}^{(k)}} \left(t_{r,\mathrm{end}}^{(k)} - t_{r,\mathrm{start}}^{(k)} \right) \tag{4.17}$$

be the total takeover time during that iteration.

Figures 4.14 and 4.15 show the takeover count for the full and sensitivity cohorts, respectively. Figures 4.16 and 4.17 show the corresponding total takeover time. The full-cohort view is necessary for transparency, while the sensitivity view tests whether the decreasing trend remains after removing the highest-average participant.

The paired statistics support the chapter hypothesis for the intervention-prone portion of the cohort. The full cohort change in takeover count was -1.28 ± 2.05 events with a 95% confidence interval of $[-2.28,\ -0.44]$, and the full cohort change in total takeover time was -13.60 ± 22.17 s with a confidence interval of $[-24.39,\ -4.61]$ s. The corresponding sensitivity-cohort changes were -1.35 ± 2.09 events and -13.94 ± 22.81 s. The similarity between the full and sensitivity results indicates that the reduction in takeover is not explained solely by removing the high-intervention participant.

Figures 4.18–4.19 show the high-intervention participant separately. This participant is important because the case demonstrates a limit of the current disturbance representation. The participant's total takeover time decreases, but takeover does not disappear. This suggests that speed-policy reduction captured part of the participant's intervention behavior, but not all of it.

4.4.6 Lateral Error and Vehicle-Performance Context

Lateral error is reported as a supporting metric rather than the primary outcome. The chapter's main question is whether the learned speed profile reduces takeover. Lateral error is still important because it shows whether the learned reference profile preserves acceptable path-following behavior and because it helps interpret cases in which takeover does not align with path-tracking error. Figures 4.20a and 4.20b compare the initial measured lap with the final learned iteration, while Figures 4.21 and 4.22 show subgroup boundary cases.

Mean lateral error decreased from 1.10 ± 1.44 m to 0.28 ± 0.17 m for the full cohort and from 1.13 ± 1.48 m to 0.29 ± 0.17 m for the sensitivity cohort. Peak lateral error also decreased numerically, but the paired-test p-values exceeded 0.05. The lateral-error evidence is therefore best interpreted as supporting context: the learned profiles did not produce a path-following

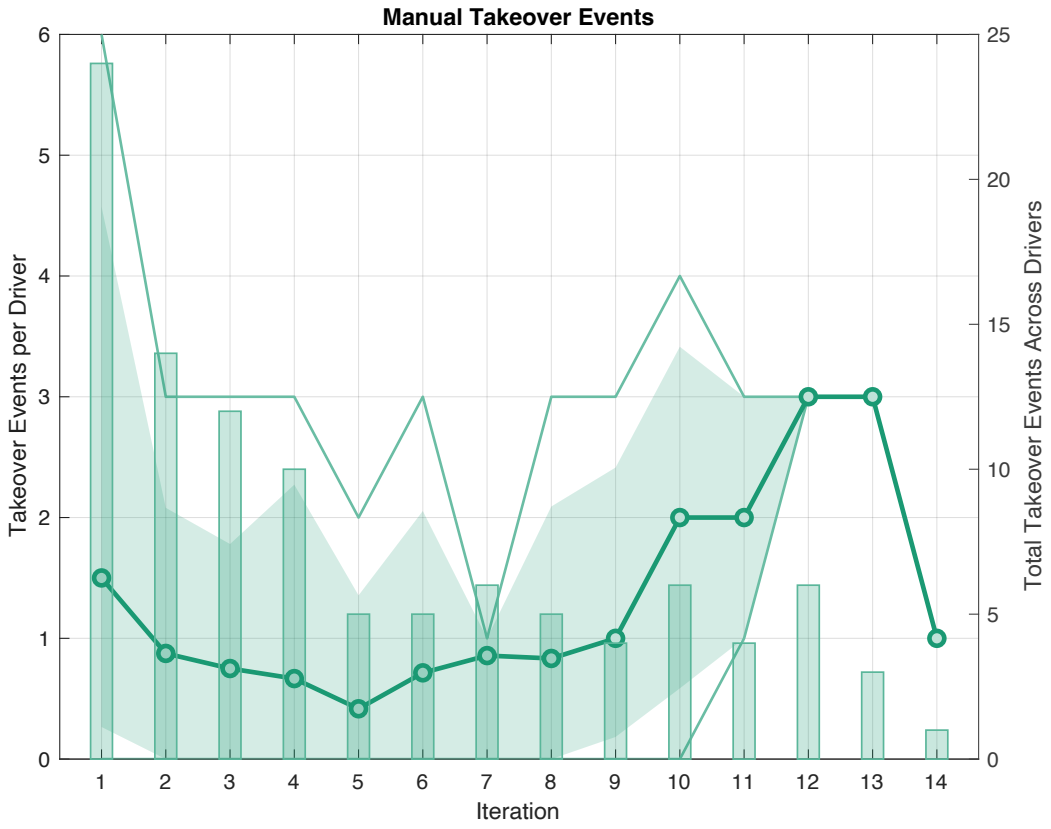

Figure 4.14: Evolution of takeover count across learning iterations for the full cohort. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The mean trend is shown with a one-standard-deviation band; participant-level variation is retained to show the full spread.

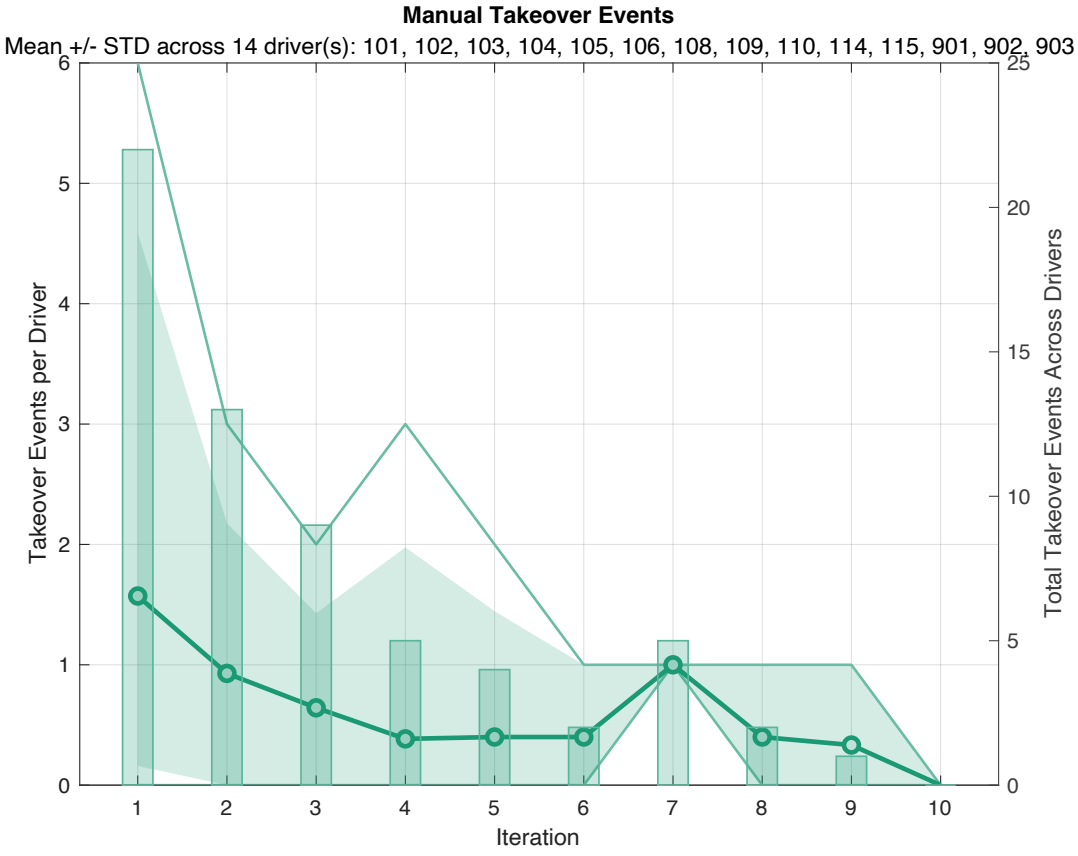

Figure 4.15: Evolution of takeover count across learning iterations for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The remaining decrease indicates that the cohort-level takeover reduction is not explained solely by the high-intervention case.

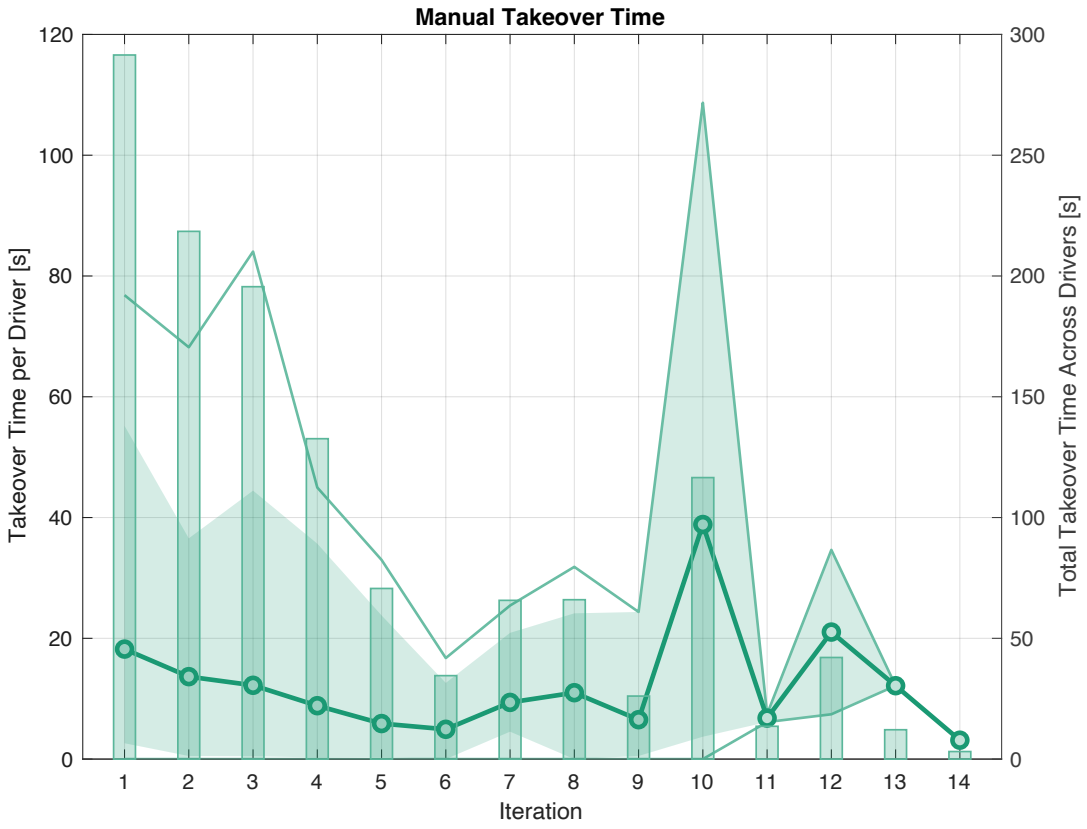

Figure 4.16: Evolution of total takeover time across learning iterations for the full cohort. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The plot complements the takeover-count result by showing whether interventions became shorter across learning iterations.

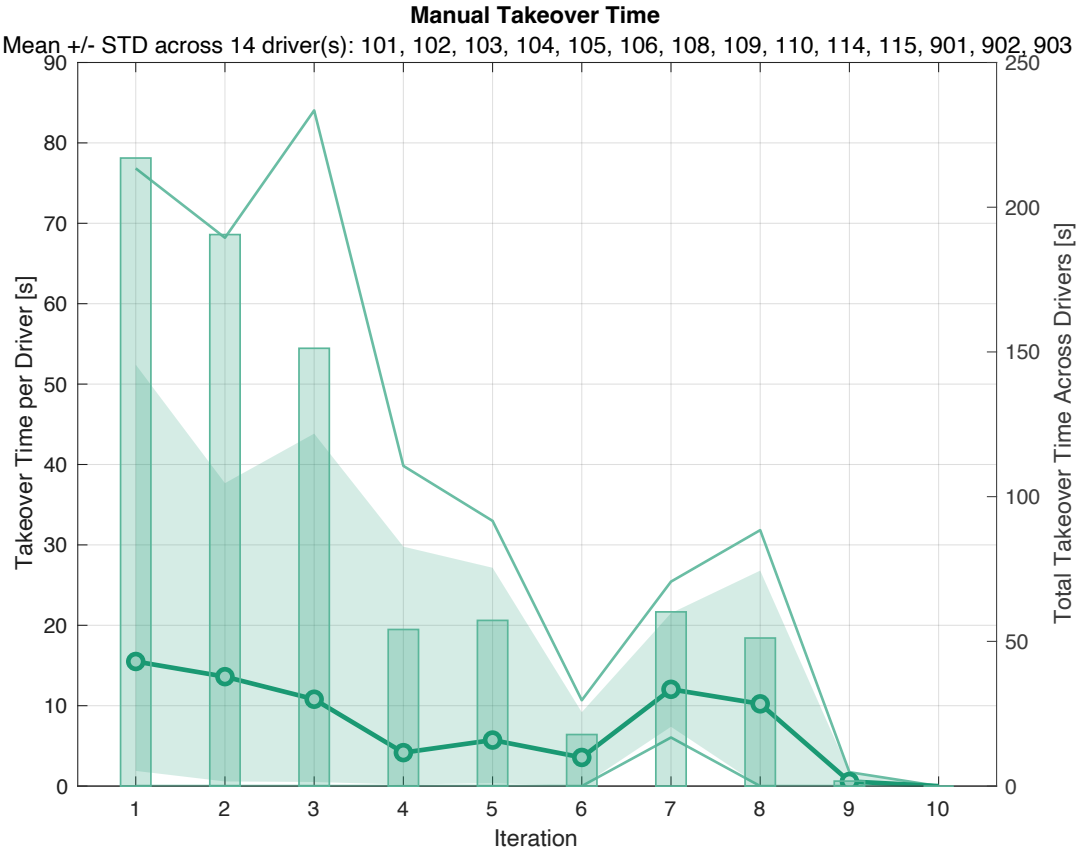

Figure 4.17: Evolution of total takeover time across learning iterations for the sensitivity cohort excluding only the high-intervention participant. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. Bars indicate total number of takeover events across the cohort. The trend is used as a sensitivity check on the total intervention-duration result.

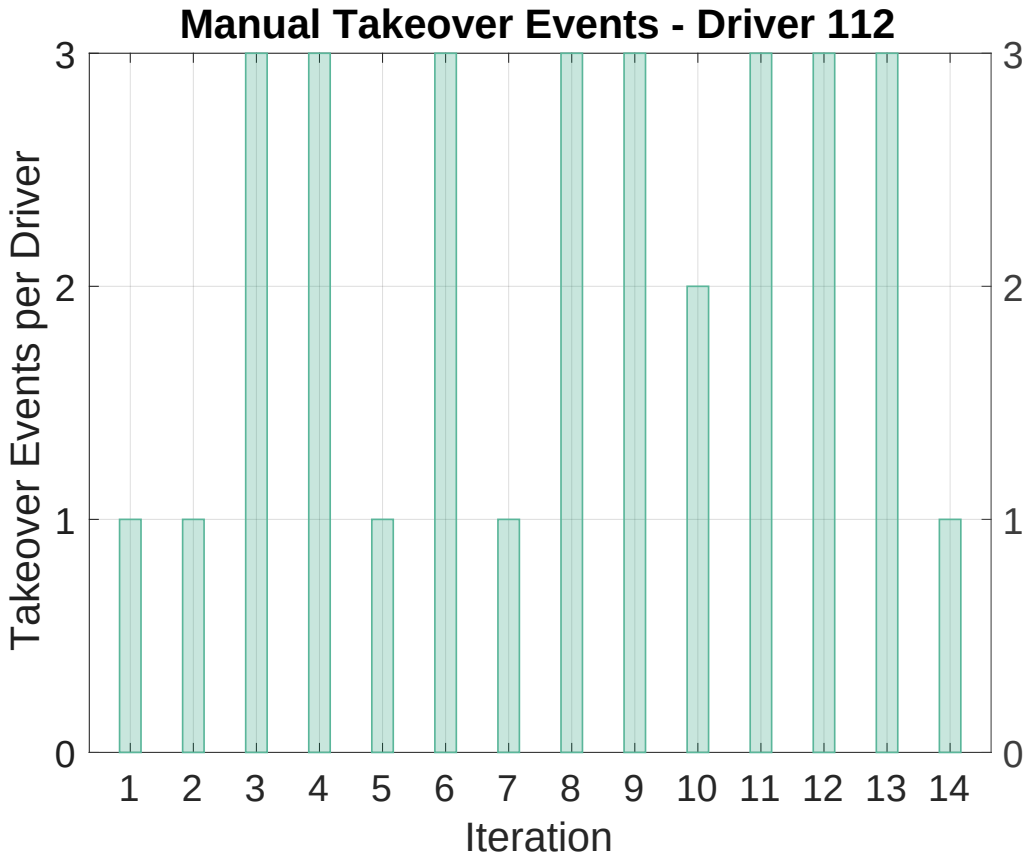

Figure 4.18: Evolution of takeover count for the high-intervention participant. The participant continued to produce takeover events after multiple local speed reductions, illustrating a boundary case for the proposed disturbance representation.

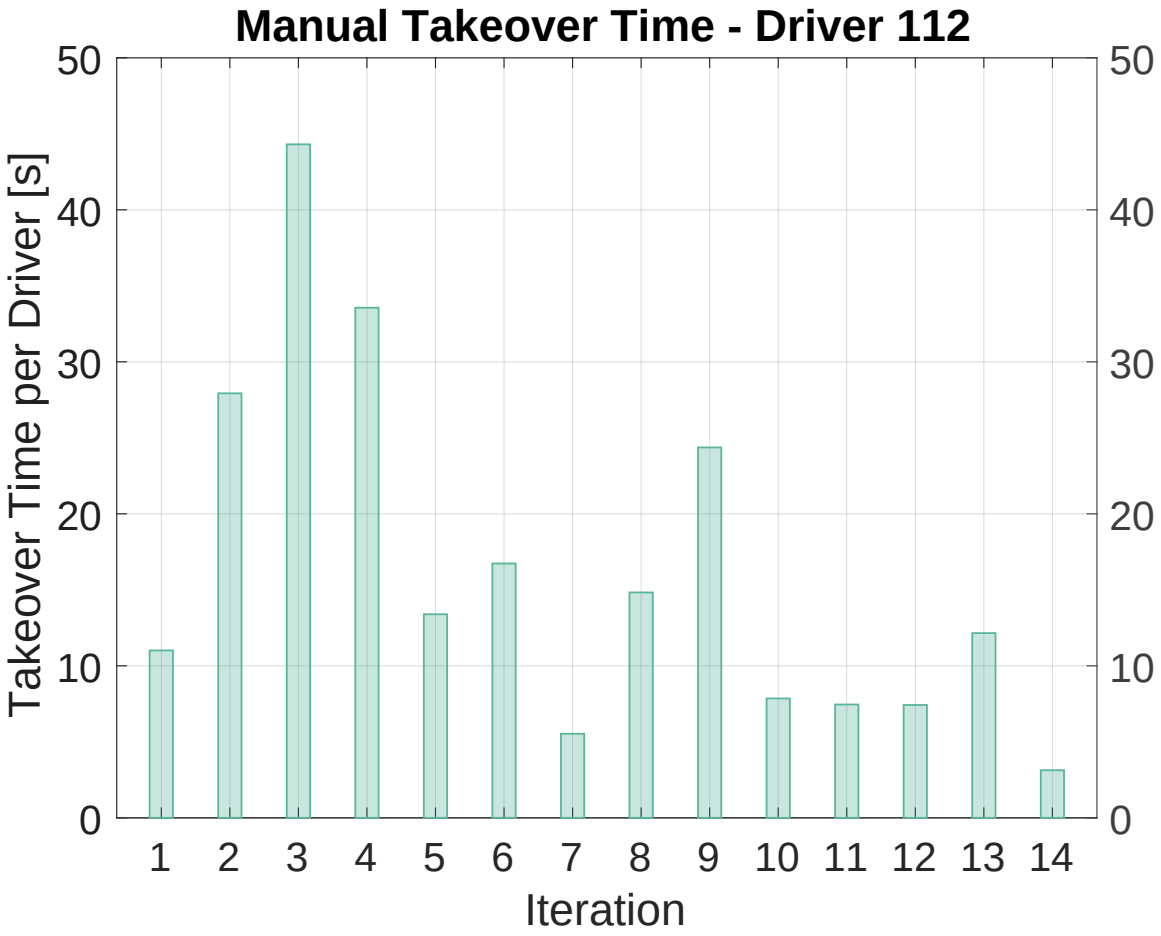

Figure 4.19: Evolution of total takeover time for the high-intervention participant. The trend is interpreted as partial mitigation rather than complete convergence.

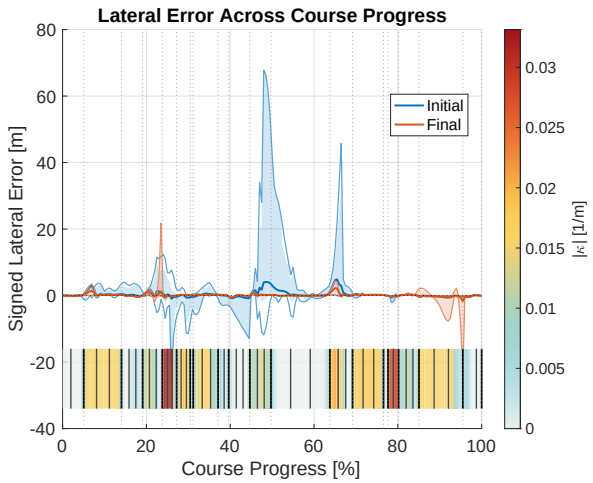

(a) Signed lateral error for the initial measured lap compared with the final learned iteration across the handling course.

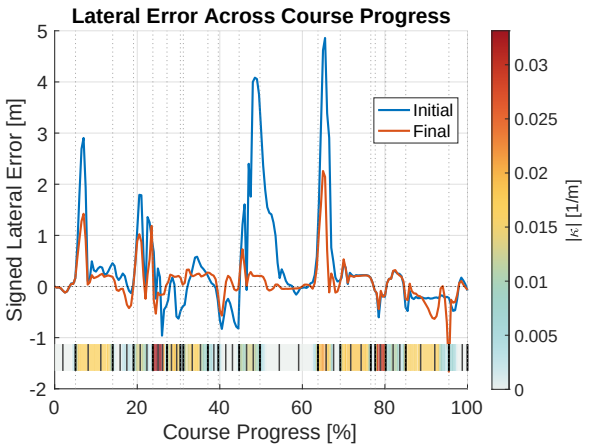

(b) Zoomed signed lateral-error comparison between the initial measured lap and the final learned iteration.

Figure 4.20: Signed lateral error for the initial measured lap compared with the final learned iteration across the handling course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The restricted y-axis in Figure 4.20b shows the mean and one-standard-deviation behavior near the path centerline while the full-scale excursion information is preserved in Figure 4.20a.

degradation in the reported metrics, and mean lateral error improved, but the chapter's primary claim remains takeover reduction rather than tracking-control improvement.

The subgroup figures show why lateral error should not be treated as a complete explanation of takeover. The non-intervention subgroup did not intervene despite path-tracking deviations at some route locations. In contrast, the high-intervention participant continued to intervene even when lateral error was comparatively small. This contrast supports the chapter's methodological choice: the proposed method does not try to explain the participant's latent reason for takeover. It uses the observed takeover signal to modify later behavior.

4.4.7 Summary of Known-Route Learning Results

The results support four bounded conclusions. First, route-local coefficient-of-friction reduction provides a practical mechanism for altering the handling-course speed profile without changing the physical TruckSim road friction. Second, the takeover count and total takeover time decreased for the full and sensitivity cohorts. Third, the continuous-task score increased throughout the cohort, including for the non-intervention subgroup. Fourth, lateral-error results provide supporting evidence that learned profiles did not degrade reported vehicle

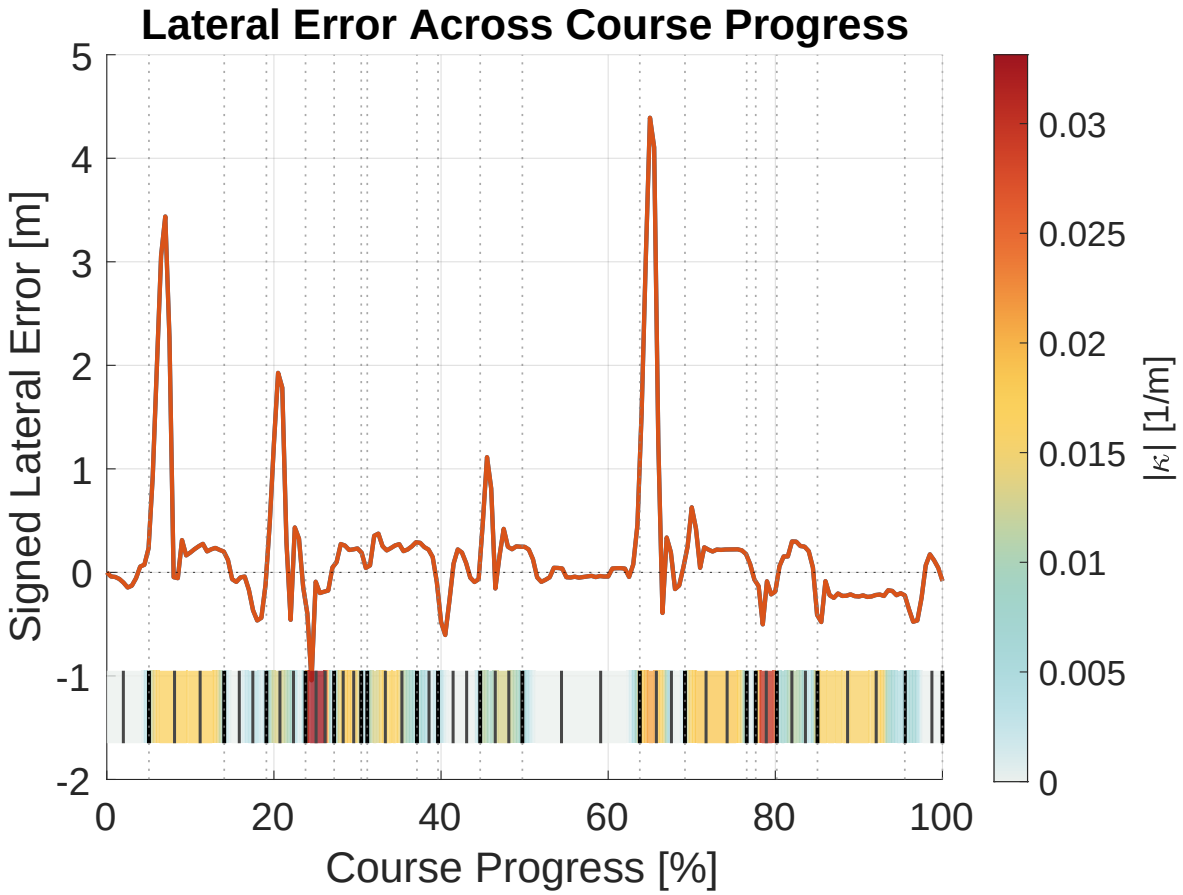

Figure 4.21: Signed lateral-error behavior for the non-intervention subgroup. These participants did not intervene despite path-tracking deviations at some route locations, illustrating that lateral error and takeover are related but not equivalent measures.

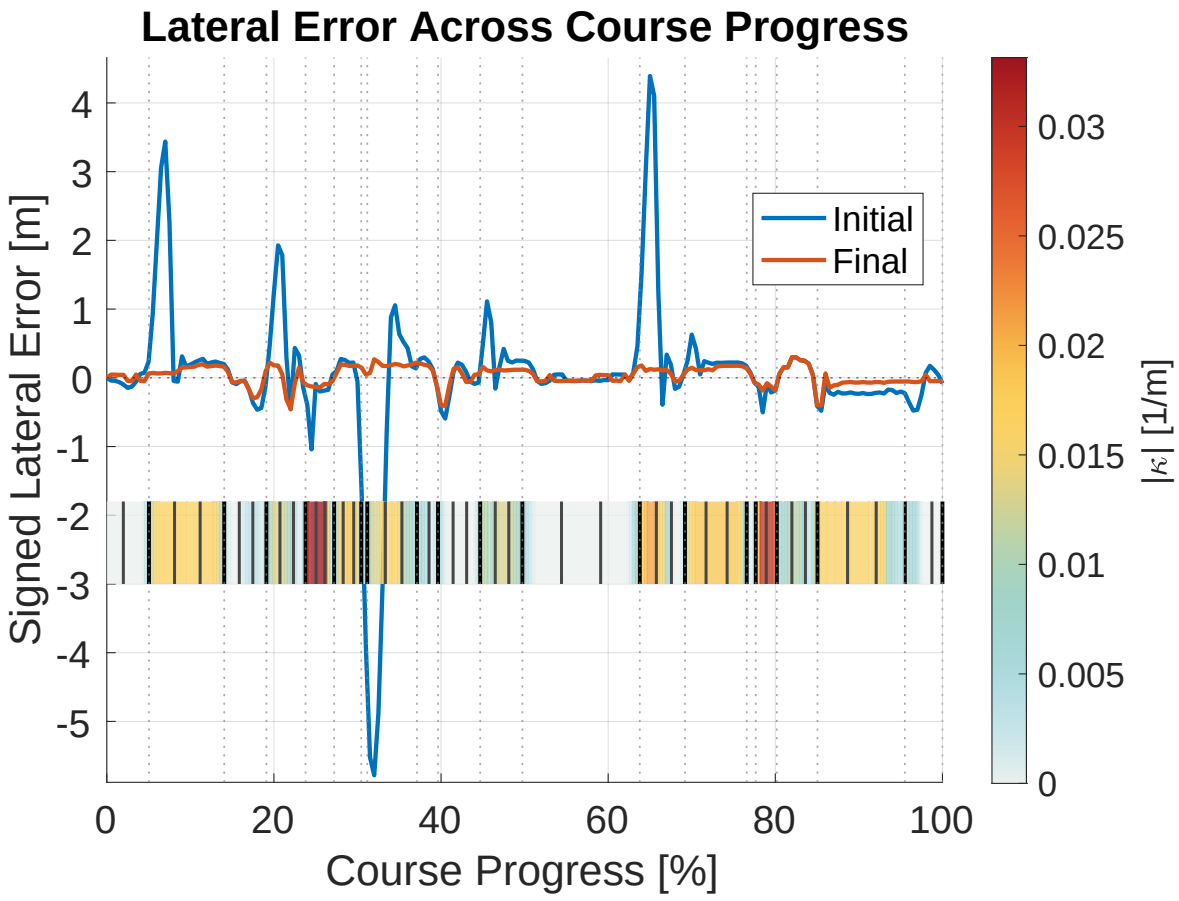

Figure 4.22: Initial and final signed lateral-error behavior for the high-intervention participant. The participant continued to intervene even when lateral error alone did not provide a complete explanation of the intervention pattern, supporting the decision to treat takeover as the measured disturbance signal.

path-following performance, although peak lateral-error improvements should be described cautiously.

These observations are consistent with the chapter hypothesis, but the final claim remains narrow. The method does not explain why the participants intervened. It shows that for many participants on a repeated known route, the adaptation of the local speed-policy can reduce the measured takeover while preserving route performance.

4.5 Discussion and Chapter Conclusion

4.5.1 Interpretation of the Learning Result

The central result of this chapter is not that all participants converged to a common preferred speed profile, and it is not that human intervention was fully explained. The narrower result is that repeated human takeover on the handling course can be treated as a control-relevant disturbance for participants who produce takeover events. When takeover occurred in a route bin, reducing the planner's assumed coefficient of friction in that bin produced a more conservative speed profile for the next iteration and was associated with reduced future intervention in the cohort-level paired results.

This interpretation should remain tied to the measured signal. The method does not infer whether a participant intervened because of trust, discomfort, attention, perceived instability, steering motion, lateral deviation, proximity to a road edge, or other latent state. It treats takeover as an observed indication that the current autonomy behavior was not accepted by that participant at that route location. This framing is less ambitious than a full human-behavior model, but it is well matched to the control objective in this chapter: reducing repeated intervention through reference-behavior adaptation.

4.5.2 Participant Heterogeneity and Sensitivity of the Cohort Trend

Two participant-level patterns are essential to interpret the results. The first is the high-intervention participant. Two additional participants were also comparatively cautious, but one participant exerted a much larger influence on full-cohort means and variances for lap time, takeover count, and total takeover time. Reporting a sensitivity cohort that excludes only this participant is therefore appropriate, but it should be read as a sensitivity analysis rather than as the primary dataset. The full cohort shows the complete study outcome. The sensitivity cohort shows that the dominant learning trend is not solely an artifact of one

high-leverage participant.

The high-intervention participant is not a failure case in a simple sense. Takeover time decreased, but intervention persisted. This pattern suggests that route-local speed reduction captured part of the participant's response to the autonomy but did not capture all relevant factors. One plausible explanation is that the participant responded to a cue not represented by the learned friction profile, such as steering-rate behavior, yaw motion, proximity to a road edge, or a local geometry feature not captured by the turn/part bin. This remains an interpretation rather than a measured mechanism.

The second pattern is the non-intervention subgroup. Five participants never intervened during the handling-course experiment. For these participants, the disturbance-rejection mechanism received no event signal and therefore could not infer a more specific route-local profile. This does not prove that the nominal profile was ideal for them. It only means that the measured signal used by this chapter did not indicate a need for speed-profile revision. Their continuous-task scores still increased, but that improvement should be interpreted conservatively. It may reflect repeated exposure to the continuous task, greater familiarity with the simulator, reduced monitoring effort, or some combination of these factors.

4.5.3 Role and Limits of the Assumed Friction Profile

The learned coefficient-of-friction map is a control proxy, not a road-friction estimate. In TruckSim, the true tire–road interaction is governed by the simulation configuration. The learned μ values only modify the planner's assumed acceleration envelope. Lower assumed friction reduces curvature-limited and braking-feasible speeds, producing a more conservative reference through the affected route bins.

This proxy is useful because it is compact, physically interpretable, and compatible with the existing speed planner. It also creates a clear limitation. Since $\mu_{\max} = 0.7$ was selected as the safe operating ceiling for the handling course, the learning law could not create profiles that were more aggressive than the baseline. The value was determined from pre-study engineering evaluation as the highest assumed-friction value at which the vehicle remained on the course under the implemented controller and TruckSim configuration. It is best interpreted as a manufacturer-like planner limit: a design team may choose a maximum operating profile that keeps the vehicle within a predefined envelope, even if some individual occupants would tolerate or prefer more speed. In this study, that ceiling was intentionally aggressive enough to expose mismatch between nominal autonomy behavior and some participants' experience. For example, allowing lateral path errors on the order of several meters would not normally be an acceptable production comfort or safety target. The absence of upward adaptation

should therefore not be interpreted as evidence that all participants preferred $\mu = 0.7$; it reflects the imposed upper bound of the study design.

4.5.4 Performance Tradeoffs

Route-local speed reduction can reduce takeover events, but it can also change route time and path-following behavior. Table 4.3 shows that takeover count and total takeover time decreased for the full and sensitivity cohorts, while continuous-task score increased. Lap time did not improve, and the sensitivity cohort showed essentially no baseline-to-final lap-time change. This combination is important: the evidence does not support a claim that the learned profile made the route faster. It supports the narrower claim that takeover decreased and continuous-task performance improved without a meaningful cohort-level lap-time penalty in the sensitivity analysis.

The lateral-error results should also be interpreted carefully. Mean lateral error decreased, which is consistent with the more conservative route-local speed profiles. Peak lateral error decreased numerically, but the primary paired-test p-values exceeded 0.05. Thus, lateral error should be used as supporting vehicle-performance context rather than as the primary evidence for the chapter claim. The chapter's core result remains takeover reduction through route-local speed-policy adaptation.

4.5.5 Implications for Experimental Design

The experiment demonstrates one workable design for driver-specific profile adaptation, but it also exposes the controls that would strengthen the next study. The most important addition would be a global-speed-reduction baseline. That condition would test whether route-local learning is more effective than simply lowering speed everywhere. Without that comparison, the current study can show that route-local adaptation is feasible and associated with reduced takeover, but it cannot fully isolate the benefit of route locality relative to a global conservative profile.

A second useful addition would be a no-update or sham-update condition. That condition would separate participant familiarization from improvement caused by the learned speed profile. This is especially important for the non-intervention subgroup, whose continuous-task scores improved even without takeover-based updates.

A third extension would be to collect a lightweight comfort or perceived-safety rating after each lap or after selected turns. The method in this chapter intentionally avoids explaining why intervention occurred, but a simple subjective rating would help distinguish between cases where takeover is plausibly safety-driven, comfort-driven, or associated with simulator/task

familiarity. That information would not replace the event-based disturbance signal; it would help interpret where the simplified model succeeds and where it is incomplete.

Finally, repeating the handling-course experiment on a second day or after a delay would help separate true profile learning from within-session adaptation to the simulator and continuous task.

4.5.6 Limitations

Several limitations remain explicit. First, this chapter studies a known route. The same route segmentation is used across iterations, and the learning update relies on repeated exposure to the same turn/part bins. This chapter therefore does not establish transfer to a new course. That question is reserved for Chapter 5.

Second, the method is evaluated only for the participants and the simulator conditions tested. It should not be described as a universal model of all drivers or all handling-course preferences. The appropriate claim is that the study provides evidence that route-local speed-policy learning can reduce takeover events for many participants under the tested conditions.

Third, the assumed coefficient-of-friction update is a proxy for speed-policy conservatism. It is not a physical friction estimator and should not be interpreted as identifying the true tire–road coefficient of friction.

Fourth, takeover alone is not an exhaustive descriptor of human preference. The non-intervention subgroup is the clearest example. A participant can produce no takeover events while still changing secondary-task performance over time. This is not a failure of the experiment, but it sets a limit on what this specific disturbance signal can identify.

Fifth, because this experiment did not include a no-update, sham-update, or global-speed-reduction control condition, it is difficult to completely isolate the effects of profile learning from participant familiarization. While the current results effectively demonstrate the feasibility of the adaptive framework, definitive causal claims regarding the drivers of this performance change should be deferred until these controls can be rigorously evaluated.

4.5.7 Chapter Conclusion

This chapter treated human takeover on a repeated handling course as an observed disturbance to the autonomy system. Takeover events were mapped to route bins, used to update a route-local assumed coefficient-of-friction profile, and converted through a backward-reachable speed planner into a revised reference-speed profile. The results show that the takeover count and the total takeover time decreased for the full and sensitivity cohorts, while the

continuous-task score increased and the lap time remained approximately unchanged in the sensitivity cohort.

These findings support the bounded claim that measured takeover can serve as a useful event signal for route-local speed-profile adaptation on a known, repeated route. While this method neither explains the underlying reasons for participant intervention nor identifies a universally preferred driving profile, it demonstrates that a simplified, data-informed disturbance representation can successfully modify reference behavior and reduce interventions for many tested participants. Chapter 5 builds on this known-route foundation to investigate whether a learned profile can be transferred to a novel course using route-geometry features.

CHAPTER 5

Transfer of Learned Human-Intervention Profiles to Novel Routes

5.1 Introduction

Chapter 4 treated human takeover as an observed disturbance on a repeated, known route. Rather than modeling the driver's internal state, intention, or trust, the framework used repeated intervention as evidence that the autonomy's nominal route-local speed policy failed to avoid takeover for that specific driver and context. In response, the known-route controller reduced the assumed coefficient of friction within the corresponding route bins and recomputed the backward-reachable speed profile. The resulting learned profile served as a route-indexed control adjustment that recorded where prior takeover events caused the speed planner to become more conservative for the tested participant.

This chapter asks whether elements of the learned profile can be transferred to a different course without repeated iterations of the route to learn a new profile. The objective is not to prove universal generalization across drivers or courses. The objective is to evaluate whether a profile learned in the handling course can be converted into a feedforward profile for a novel course using a compact route-context representation. The transfer problem is framed as a library-matching problem. Each source-route and target-route bin is represented by curvature-derived features, the target bins are matched to similar source bins in normalized feature space, and the learned assumed-friction values from the source route are mapped onto the target route. The transferred profile is then used to recompute the backward-reachable speed reference for the target course.

Two transfer operators are evaluated. The first uses a weighted three-nearest-neighbor average of learned source-route values. The second assigns the value from the single nearest source bin in the normalized feature space. Both are compared against the same fixed baseline used in Chapter 4: a uniform assumed coefficient of friction of $\mu = 0.7$. The feature vector used for transfer contains the mean absolute curvature of the current bin, peak absolute

curvature of the current bin, mean absolute curvature of the preceding bin, peak absolute curvature of the preceding bin, and normalized location of the bin within the turn. This feature set is intentionally small. It tests whether the curvature context alone contains enough information to reuse a learned driving profile on a new route.

The central hypothesis is that a driver-specific profile learned in the handling course can reduce interventions in a novel course if the target route contains segments with curvature features similar to those of the source route. To maintain a tight scope, this hypothesis is deliberately limited. The underlying method does not attempt to explain why a driver intervenes, estimate trust, or claim that curvature captures all factors that affect human behavior. Rather, it evaluates whether a simple and interpretable route-context transfer rule can outperform a uniform baseline profile.

The chapter contributes a transfer framework that maps learned route-local coefficient-of-friction values from a source course to a target course using nearest-neighbor methods in a normalized curvature-feature vector space. The framework is a concrete example of the theme: a simplified, data-informed model can be useful for control design when it captures the structure needed for the decision. In this case, the simplified model is a route-context library that links local curvature features to driver-specific reductions in the assumed friction used by the speed planner.

The remainder of the chapter is organized as follows. Section 5.2 positions the chapter relative to personalization, route-context transfer, and example-based control adaptation. Section 5.3 defines the learned library, target-route feature extraction, weighted three-nearest-neighbor transfer rule, single-nearest-neighbor transfer rule, and backward-reachable profile synthesis. Section 5.4 compares the uniform $\mu = 0.7$ baseline with the two transferred driver-specific profiles on the aggressive target course. Section 5.5 discusses what the results support and where curvature-only transfer remains insufficient.

5.2 Related Work and Motivation

Automated-driving systems are often designed using a nominal route representation, vehicle model, and controller objective. However, the desired behavior of an automated vehicle may depend on the driver, the driving context, and the driver's tolerance for the autonomy's speed and path behavior. Prior work studies takeover behavior, driver state, trust, distraction, and human intention in automated driving. That work helps explain why interventions occur, but explanation is not the modeling objective here. This chapter treats takeover as a measured disturbance event indicating a mismatch between the autonomy's behavior and the tested driver's intervention boundary in that route context.

Recent work in personalized autonomous driving and preference-aligned planning shows that driver-specific information can modify automated-driving behavior [Cui et al., 2024, Karnan et al., 2024a]. Those approaches motivate the broader idea that a control policy can adapt to a human user. The present chapter differs in two ways. First, the transferred profile is not generated from natural-language preference statements or from a learned high-dimensional policy. It is a route-local assumed-friction profile learned from repeated interventions on a source course. Second, the transfer rule is intentionally simple and interpretable: target-route parts are matched to source-route parts using Euclidean distance in a normalized vector space of curvature-derived features.

This design choice reflects the emphasis on simplified, control-oriented models. A richer transfer formulation could include road width, elevation, banking, surface condition, vehicle speed history, driver state, visual preview, or secondary-task loading. Those features may improve prediction in future work, but they increase the data requirement and can make the transfer mechanism harder to diagnose. This chapter begins with a minimal structure that tests whether learned route-local speed conservatism can be reused on a novel course.

The transfer problem is framed as instance-based library matching. The source library is the learned handling-course coefficient-of-friction profile from Chapter 4, indexed by turn and part. Each source turn/part bin is described by curvature-derived features. A target route is partitioned into analogous turn/part bins, and each target bin receives either a weighted average of the learned values from the three nearest source bins or the learned value from the single nearest source bin. The transferred profile is then used by the same backward-reachable speed planner introduced in Chapter 4. This creates a direct comparison among a uniform baseline speed policy, a smoothed weighted-neighbor transfer policy, and a single-neighbor transfer policy.

The main technical risk is that curvature similarity may not be sufficient to predict intervention behavior. A target segment can have similar curvature features to a source segment but differ in visual appearance, entry speed, preceding maneuver, route width, or timing within the drive. For that reason, the results are interpreted as evidence for one transfer mechanism, not as evidence that curvature alone is an exhaustive model of driver preference or takeover behavior.

5.3 Methods

5.3.1 Problem Formulation

Let the source route be the handling course used in Chapter 4. After repeated iterations with a tested driver, the known-route learning procedure produces a learned assumed coefficient-of-friction matrix

$$\boldsymbol{M}^{\mathrm{src}} = \left[\mu_{j,p}^{\mathrm{src}}\right], \tag{5.1}$$

where j indexes the source-course turn and p indexes the part within that turn. Each entry $\mu_{j,p}^{\mathrm{src}}$ represents the coefficient of friction used by the backward-reachable speed planner for that source-route bin after learning. The physical road friction in TruckSim is not changed. Instead, $\mu_{j,p}^{\mathrm{src}}$ is an assumed-friction design parameter that makes the planned speed more conservative in contexts where prior interventions occurred.

The target route is a new course that the driver has not used to learn a route-specific coefficient-of-friction profile. The target route is partitioned into turns and parts, producing bins indexed by (ℓ, q), where ℓ is the target-course turn and q is the target-course part. The transfer problem is to predict a target-course assumed-friction matrix

$$\widehat{\boldsymbol{M}}^{\mathrm{tgt}} = \left[\widehat{\mu}_{\ell,q}^{\mathrm{tgt}}\right] \tag{5.2}$$

from the source-course learned matrix $\boldsymbol{M}^{\mathrm{src}}$ and the geometric features of the two courses.

The baseline controller uses a uniform assumed coefficient of friction,

$$\mu^{\mathrm{base}}(s) = 0.7, \tag{5.3}$$

for every point on the target route. The transfer controllers use target-course matrices $\widehat{\boldsymbol{M}}^{\mathrm{tgt}}$, map those values onto the target-course path mesh, and recompute the backward-reachable speed profile using the equations from Chapter 4. The primary comparison is therefore among the uniform $\mu = 0.7$ baseline, the weighted three-nearest-neighbor transferred profile, and the single-nearest-neighbor transferred profile.

5.3.2 Route-Context Feature Representation

Each source-course and target-course bin is represented by a vector of curvature-derived features. The implemented feature vector is

$$\boldsymbol{x}_r = \left[\bar{\kappa}_r \quad \kappa_r^{\max} \quad \bar{\kappa}_r^{\mathrm{prev}} \quad \kappa_r^{\max,\mathrm{prev}} \quad \rho_r\right]^{\top}, \tag{5.4}$$

where r indexes a turn/part bin. The first two entries are the mean absolute curvature and peak absolute curvature within the current bin. The next two entries are the mean absolute curvature and peak absolute curvature over the preceding bin. The final entry, ρ_r, is the normalized progress of the part within the turn. These features correspond to the implemented variables:

$$[\texttt{MeanAbsKappa}, \texttt{PeakAbsKappa}, \texttt{PrevMeanAbsKappa}, \texttt{PrevPeakAbsKappa}, \texttt{PartProgressMid}]\,. \tag{5.5}$$

The preceding-bin terms are included because a driver's intervention in a given part may depend not only on the local curvature but also on the curve entry condition or immediately preceding curvature context. The part-progress term distinguishes early, middle, and late portions of turns with otherwise similar curvature magnitudes. Table 5.1 summarizes the implemented feature representation and transfer hyperparameters used in the aggressive-course evaluation.

Let $\mathcal{R}_{\text{src}}$ denote the set of valid source-course bins and $\mathcal{R}_{\text{tgt}}$ denote the set of valid target-course bins. For each source bin $i \in \mathcal{R}_{\text{src}}$, let $\boldsymbol{x}_i^{\text{src}}$ be the feature vector and let μ_i^{src} be the learned coefficient of friction for that source bin. For each target bin $m \in \mathcal{R}_{\text{tgt}}$, let $\boldsymbol{x}_m^{\text{tgt}}$ be the corresponding feature vector.

5.3.3 Normalized Feature Space

Because the feature entries have different units and numerical scales, the distance calculation is performed in normalized source-feature coordinates. The source-feature mean and scale are

$$\boldsymbol{c} = \frac{1}{N_{\text{src}}} \sum_{i \in \mathcal{R}_{\text{src}}} \boldsymbol{x}_i^{\text{src}}, \tag{5.6}$$

$$\boldsymbol{\sigma} = \text{std}\left(\{\boldsymbol{x}_i^{\text{src}}\}_{i \in \mathcal{R}_{\text{src}}}\right), \tag{5.7}$$

where entries of $\boldsymbol{\sigma}$ that are zero or numerically small are replaced with one to avoid division by zero. The normalized source and target feature vectors are

$$\boldsymbol{z}_i^{\text{src}} = (\boldsymbol{x}_i^{\text{src}} - \boldsymbol{c}) \oslash \boldsymbol{\sigma}, \tag{5.8}$$

$$\boldsymbol{z}_m^{\text{tgt}} = \left(\boldsymbol{x}_m^{\text{tgt}} - \boldsymbol{c}\right) \oslash \boldsymbol{\sigma}, \tag{5.9}$$

where $\oslash$ denotes element-wise division. Normalizing the target-course features with the source-course center and scale makes the transfer rule depend on the learned library distribution. Target segments far outside the source-feature range are therefore identifiable as extrapolation cases.

Table 5.1: Route-context features and transfer hyperparameters used for aggressive-course profile transfer.

Quantity	Implementation	Purpose
Current-bin mean absolute curvature	$\bar{\kappa}_r$	Represents average lateral-acceleration demand within the target or source bin.
Current-bin peak absolute curvature	$\kappa_r^{\max}$	Represents the most demanding local curvature within the bin.
Preceding-bin mean absolute curvature	$\bar{\kappa}_r^{\text{prev}}$	Captures entry context from the route segment immediately before the current bin.
Preceding-bin peak absolute curvature	$\kappa_r^{\max,\text{prev}}$	Captures sharp preceding curvature that may influence the vehicle state entering the current bin.
Normalized part location within turn	ρ_r	Distinguishes early, middle, and late portions of turns with similar curvature values.
Feature normalization	Source-course mean and standard deviation	Places all features on comparable scales and allows target bins far from the source distribution to be identified.
Distance metric	Euclidean distance in normalized feature space	Defines source–target route-context similarity.
Weighted transfer	$K = 3$, inverse-distance weights $1/(d+\epsilon)$	Smooths the transferred assumed-friction value across the three closest source bins.
Direct transfer	Single nearest source bin	Provides an interpretable direct source-to-target assignment for comparison.
Candidate restriction	Same turn sign when available	Prevents left-turn target bins from being matched to right-turn source bins when sign-specific context exists.

Note: The transferred value is an assumed-friction parameter used by the speed planner, not a measurement of physical road friction.

For each target bin m, the distance to a candidate source bin i is computed as

$$d_{m,i} = \left\| \boldsymbol{z}_m^{\mathrm{tgt}} - \boldsymbol{z}_i^{\mathrm{src}} \right\|_2 . \tag{5.10}$$

The candidate source set is denoted $\mathcal{C}_m$. In the implementation, $\mathcal{C}_m$ can be restricted to source bins with the same turn sign as the target bin when source bins of that sign exist. This optional sign constraint prevents a left-turn target segment from being matched to a right-turn source segment when sign-specific route context is relevant. When the sign-constrained set is empty, the method falls back to all valid source bins.

5.3.4 Weighted Three-Nearest-Neighbor Transfer

The first transfer method uses the three nearest source bins in the normalized feature space. Let $\mathcal{N}_3(m)$ be the set of the three source bins in $\mathcal{C}_m$ with the smallest distances $d_{m,i}$. The weighted transferred value is

$$\widehat{\mu}_m^{\mathrm{tgt,W3NN}} = \frac{\sum_{i \in \mathcal{N}_3(m)} w_{m,i} \mu_i^{\mathrm{src}}}{\sum_{i \in \mathcal{N}_3(m)} w_{m,i}}, \tag{5.11}$$

where $w_{m,i}$ is the distance-based weight assigned to source bin i for target bin m. The implemented transfer rule uses inverse-distance weighting,

$$w_{m,i} = \frac{1}{d_{m,i} + \epsilon}, \tag{5.12}$$

where $\epsilon > 0$ prevents a singular weight when a target bin exactly matches a source bin in normalized feature space. This weighting gives larger influence to source bins that are closer to the target bin while still smoothing the transferred profile across multiple similar source contexts.

In the figures and results tables, this method is labeled "weighted 3-NN" because it uses a weighted K-nearest-neighbor average with $K = 3$.

5.3.5 Single-Nearest-Neighbor Transfer

The second transfer method assigns each target bin the learned value from the single nearest source bin. The selected source bin is

$$i^\star(m) = \arg\min_{i \in \mathcal{C}_m} d_{m,i}, \tag{5.13}$$

and the transferred target-bin coefficient of friction is

$$\widehat{\mu}_m^{\mathrm{tgt,NN}} = \mu_{i^\star(m)}^{\mathrm{src}}. \tag{5.14}$$

The nearest-neighbor mapping table records the selected source turn, selected source part, nearest-neighbor distance, selected source μ, target turn sign, source turn sign, and whether the target feature lies outside the observed source-feature range. These diagnostics are useful for interpreting unexpected transfer behavior and for identifying target-course locations where the simplified feature representation may be insufficient.

The weighted three-nearest-neighbor method and the single-nearest-neighbor method represent different transfer assumptions. The weighted method smooths across multiple similar source contexts, which can reduce sensitivity to a single source-course bin. The single-nearest-neighbor method preserves interpretability because every target-course adjustment can be traced to one learned source-course context. Comparing both methods against the fixed $\mu = 0.7$ baseline tests whether smoothing improves transfer or whether direct nearest-bin assignment is sufficient.

5.3.6 Mapping the Transferred Profile onto the Target Course

The target-course transfer produces a coefficient-of-friction value for each target turn/part bin. The value is then assigned to every path point belonging to that bin:

$$\widehat{\mu}^{\mathrm{tgt}}(s_n) = \widehat{\mu}^{\mathrm{tgt}}_{\ell(s_n),q(s_n)}, \tag{5.15}$$

where s_n is the station of the nth target-route path point and $\ell(s_n), q(s_n)$ are the target turn and part associated with that path point. The original target path mesh is preserved. The transfer process changes the assumed-friction row used by the speed planner but does not resample the path geometry.

After the path-indexed profile $\widehat{\mu}^{\mathrm{tgt}}(s_n)$ is assigned, the backward-reachable speed planner is recomputed. As in Chapter 4, the curvature-limited speed cap is

$$v_{\mathrm{lat}}(s_n) = \begin{cases} \sqrt{\dfrac{\widehat{\mu}^{\mathrm{tgt}}(s_n) g}{|\kappa(s_n)|}}, & |\kappa(s_n)| \geq \epsilon_\kappa, \\ v_{\mathrm{straight}}, & |\kappa(s_n)| < \epsilon_\kappa, \end{cases} \tag{5.16}$$

with longitudinal feasibility enforced through the backward and forward passes previously defined in Chapter 4. Thus, the transfer profile changes the target route only through the assumed-friction input to the same speed-generation procedure. This preserves a direct

experimental comparison between the baseline and transferred conditions.

5.3.7 Evaluation Conditions and Metrics

The transfer method is evaluated by comparing three target-course conditions:

1. a baseline condition using the uniform assumed-friction profile $\mu = 0.7$;
2. a weighted three-nearest-neighbor transfer condition using Eq. (5.11); and
3. a single-nearest-neighbor transfer condition using Eqs. (5.13)–(5.14).

All conditions are evaluated in TruckSim using the same target route, vehicle model, controller structure, and human-intervention logging procedure. The TruckSim 'Self-Contained Handling Course' scene was used as the target course (noted as 'aggressive' to distinguish it from the source course) for this study. The transfer evaluation used drivers from the human intervention ILC experiment, with each driver assigned one lap under each profile condition: the fixed baseline $\mu = 0.7$ profile, the weighted 3-NN transferred profile, and the single-nearest-neighbor transferred profile. Due to data-transfer corruption, some baseline and weighted 3-NN runs were unavailable in the aggregate data set; therefore, the results report both condition-level summaries and matched participant-level paired statistics for the sensitivity cohort.

The primary comparison metrics are lap time, number of takeover events, total takeover time, continuous secondary-task score, peak lateral error, and mean lateral error. The expected result is not that every metric improves in every condition. A transferred profile may reduce intervention at the cost of a longer lap time, and a curvature-only transfer can fail at specific route locations where unmodeled context affects takeover behavior. The results are therefore interpreted as a tradeoff among intervention reduction, task performance, route time, and tracking behavior.

5.4 Results

The transferred profiles were evaluated on the aggressive course, a target route that is narrower and more demanding than the handling course used for learning. The course contains eleven labeled turns and a track width less than 5 m. Figure 5.1 shows the course layout, start/end location, turn labels, and the curvature distribution. The route includes several regions where curvature and path geometry create higher lateral-acceleration demand than a straight or mildly curved segment, making it a suitable test case for a curvature-based transfer rule. However, there is essentially no elevation change in the course.

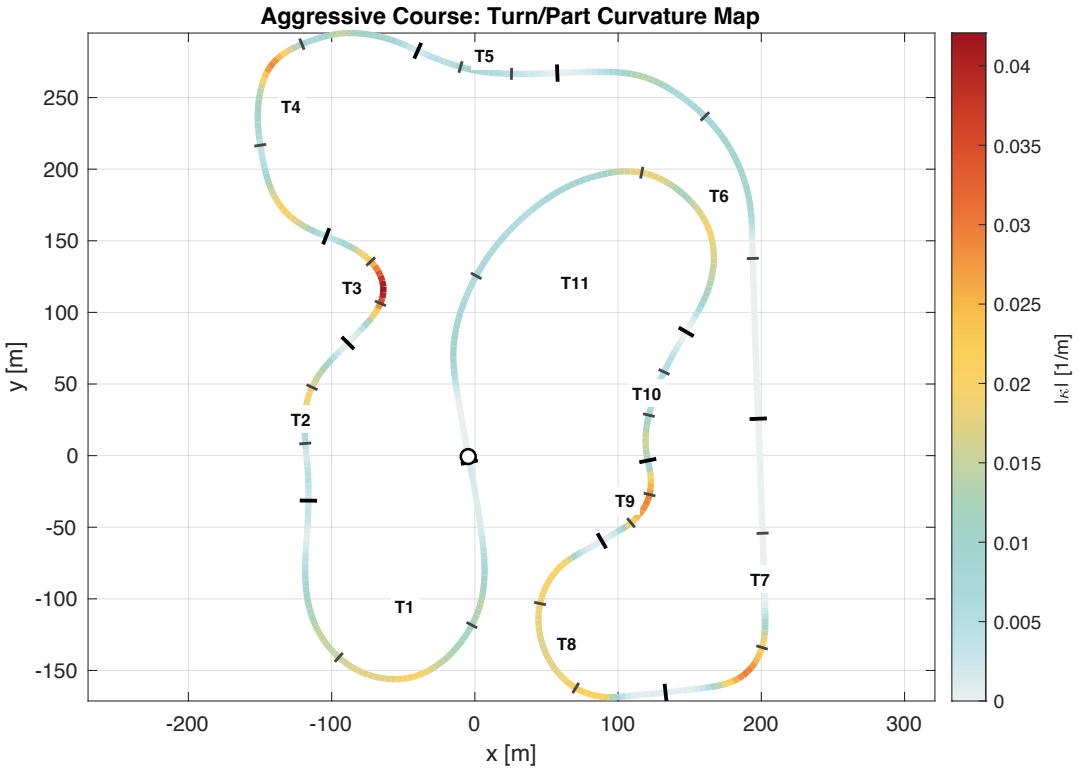

Figure 5.1: Aggressive-course layout and curvature map used for the transfer evaluation. The circle marks the start/finish location, and turns T1–T11 identify the route bins used for transfer and results interpretation. The color scale indicates curvature magnitude, which is the primary geometric quantity used to construct the route-context feature vector.

Table 5.2: Aggressive-course performance summary for the fixed baseline assumed-friction profile and two transferred driver-profile methods. Values are mean ± standard deviation. Changes are computed relative to the baseline $\mu = 0.7$ condition. Negative changes indicate reductions in lap time, takeover outcomes, and lateral error; positive changes indicate higher continuous-task score. Table supports improvement using the 3-NN method due to reduced lap time, takeover events/time, peak and mean lateral error, and improved continuous task score. The sensitivity group also shows improvement, but not as much due to poor scores in driver 112.

Cohort	Outcome	Baseline $\mu = 0.7$	Weighted 3-NN	Nearest	Δ Weighted 3-NN	Δ Nearest
Full	Lap time (s)	123.6 ± 16.1	119.5 ± 6.8	126.8 ± 14.9	-4.1	$+3.3$
	Takeover events	1.40 ± 1.64	0.27 ± 0.80	0.69 ± 0.95	-1.13	-0.71
	Total takeover time (s)	19.9 ± 28.8	0.9 ± 2.5	12.8 ± 19.5	-19.0	-7.1
	Continuous-task score	98.1 ± 47.7	134.0 ± 35.6	126.6 ± 34.6	$+35.9$	$+28.4$
	Peak lateral error (m)	9.32 ± 7.15	4.74 ± 0.00	8.47 ± 6.13	-4.58	-0.85
	Mean lateral error (m)	1.03 ± 0.94	0.33 ± 0.04	0.72 ± 0.52	-0.70	-0.32
Sens.	Lap time (s)	121.1 ± 13.5	118.4 ± 5.5	124.9 ± 13.1	-2.7	$+3.8$
	Takeover events	1.14 ± 1.35	0.07 ± 0.27	0.53 ± 0.74	-1.07	-0.61
	Total takeover time (s)	16.8 ± 27.1	0.5 ± 1.7	11.8 ± 19.8	-16.4	-5.0
	Continuous-task score	103.4 ± 44.7	136.5 ± 35.6	128.5 ± 34.9	$+33.1$	$+25.1$
	Peak lateral error (m)	8.47 ± 6.61	4.74 ± 0.00	8.41 ± 6.34	-3.74	-0.06
	Mean lateral error (m)	0.88 ± 0.77	0.33 ± 0.04	0.70 ± 0.53	-0.56	-0.19

Notes: Full = all evaluated runs included in the aggressive-course transfer summary; Sens. = sensitivity cohort excluding participant 112. Weighted 3-NN denotes the weighted three-nearest-neighbor transfer method. Nearest denotes the single-nearest-neighbor transfer method. The data set contains $N = 15$ runs for the full-cohort baseline and weighted 3-NN conditions, $N = 16$ runs for the full-cohort nearest-neighbor condition, $N = 14$ runs for the sensitivity-cohort baseline and weighted 3-NN conditions, and $N = 15$ runs for the sensitivity-cohort nearest-neighbor condition because several baseline and weighted-transfer runs were corrupted during data transfer. This table reports condition-level summaries; matched participant-level paired statistics for the sensitivity cohort are reported in Table 5.3.

5.4.1 Baseline and Transferred-Profile Summary

Table 5.2 summarizes aggressive-course outcomes for the fixed baseline method $\mu = 0.7$, the weighted three-nearest-neighbor transfer method, and the single-nearest-neighbor transfer method. The full-cohort summary includes all evaluated runs in the current data set. The sensitivity cohort excludes participant 112, the high-intervention participant identified in Chapter 4. This sensitivity comparison is retained for consistency with the known-route analysis, although participant 112 has less influence on the transfer-course summary than on the handling-course learning results.

The weighted three-nearest-neighbor method shows the strongest aggregate improvement relative to the $\mu = 0.7$ baseline. For the full cohort, it reduces takeover events from 1.40 ± 1.64 to 0.27 ± 0.80, reduces total takeover time from 19.9 ± 28.8 s to 0.9 ± 2.5 s, and increases the continuous-task score from 98.1 ± 47.7 to 134.0 ± 35.6. The same pattern remains after excluding participant 112: takeover events decrease from 1.14 ± 1.35 to 0.07 ± 0.27, total takeover time decreases from 16.8 ± 27.1 s to 0.5 ± 1.7 s, and continuous-task score increases from 103.4 ± 44.7 to 136.5 ± 35.6. The lap time also decreases modestly under the weighted

method in both cohorts.

The single-nearest-neighbor method also reduces the intervention relative to baseline, but the improvement is less consistent between outcomes. In the full cohort, takeover events decrease by 0.71 events and the continuous-task score increases by 28.4 points relative to baseline. However, lap time increases by 3.3 s, and lateral-error reductions are smaller than those produced by the weighted three-nearest-neighbor method. The sensitivity cohort shows the same pattern. These results indicate that preserving a single nearest source-bin value can transfer useful route-local conservatism, but smoothing across three nearby source contexts produces a more favorable aggregate tradeoff for this course and data set.

The zero standard deviation reported for the weighted 3-NN peak lateral error is retained because the weighted transferred profile produced nearly identical closed-loop trajectories across participants at the peak-error location. Under this condition, the takeover was nearly eliminated, so the vehicle followed the same planned path rather than diverging through participant intervention. This value is therefore interpreted as a consequence of repeatable closed-loop behavior under the transferred profile rather than as evidence of a broader lack of variability across the entire course.

Because each participant completed the fixed baseline condition and the transferred-profile conditions, the aggregate comparison in Table 5.2 was supplemented with participant-level paired statistics. Participant-level differences were computed after aggregating each participant within each profile condition. For continuous results, the primary test was a paired t-test; for zero-heavy takeover results, the primary test was the Wilcoxon signed-rank test. The confidence intervals were computed using bootstrap resampling at the participant-level. Table 5.3 reports the sensitivity-cohort results excluding participant 112, the high-intervention participant. This cohort is used as the primary statistical summary because it evaluates whether the transfer effect persists without being driven by a single high-leverage participant.

The paired analysis supports the same interpretation as the aggregate summary. Compared to the fixed $\mu = 0.7$ baseline, the weighted 3-NN profile increased the continuous-task score, reduced the takeover count, reduced the total takeover time, and reduced the average lateral error in the sensitivity cohort. Lap time did not show a statistically supported reduction, which indicates that the primary benefit of the transferred profile was not faster route completion. Peak lateral error decreased numerically, with a bootstrap interval below zero, but the paired t-test was slightly above the conventional $p = 0.05$ threshold; therefore, peak lateral error is interpreted as supportive but not primary evidence. The single-nearest-neighbor transfer increased continuous-task score, but it produced weaker takeover and lateral-error improvements and increased lap time relative to baseline. A direct comparison between the two transfer methods shows that the weighted 3-NN profile achieved lower lap times, shorter

Table 5.3: Participant-level paired statistics for the aggressive-course transfer experiment using the sensitivity cohort excluding participant 112. Differences are computed as comparison condition minus reference condition. Negative differences indicate reductions in lap time, takeover outcomes, and lateral error, while positive differences indicate higher continuous-task score.

Comparison	Outcome	N	Difference	95% CI	p	g_z
Weighted 3-NN − Baseline	Lap time (s)	14	-2.68 ± 10.13	$[-7.97,\ 2.13]$	0.340	−0.249
	Continuous-task score	14	33.07 ± 44.09	$[12.21,\ 56.43]$	0.015	0.706
	Takeover count	14	-1.07 ± 1.39	$[-1.79,\ -0.43]$	0.031	−0.728
	Total takeover time (s)	14	-16.38 ± 27.46	$[-31.21,\ -3.82]$	0.031	−0.561
	Peak lateral error (m)	14	-3.74 ± 6.61	$[-7.50,\ -0.99]$	0.054	−0.532
	Average lateral error (m)	14	-0.56 ± 0.79	$[-0.99,\ -0.19]$	0.021	−0.663
Nearest − Baseline	Lap time (s)	14	4.17 ± 6.46	$[1.00,\ 7.45]$	0.031	0.608
	Continuous-task score	14	24.43 ± 30.61	$[9.86,\ 40.71]$	0.011	0.751
	Takeover count	14	-0.57 ± 1.16	$[-1.14,\ 0.00]$	0.125	−0.464
	Total takeover time (s)	14	-4.17 ± 13.13	$[-10.68,\ 2.63]$	0.301	−0.298
	Peak lateral error (m)	14	0.20 ± 9.31	$[-4.69,\ 4.95]$	0.938	0.020
	Average lateral error (m)	14	-0.16 ± 0.77	$[-0.56,\ 0.22]$	0.451	−0.195
Weighted 3-NN − Nearest	Lap time (s)	14	-6.85 ± 9.99	$[-12.33,\ -2.11]$	0.023	−0.646
	Continuous-task score	14	8.64 ± 34.72	$[-8.07,\ 26.43]$	0.369	0.234
	Takeover count	14	-0.50 ± 0.85	$[-0.93,\ -0.07]$	0.094	−0.550
	Total takeover time (s)	14	-12.22 ± 20.63	$[-23.56,\ -3.22]$	0.031	−0.557
	Peak lateral error (m)	14	-3.94 ± 6.49	$[-7.64,\ -1.03]$	0.041	−0.570
	Average lateral error (m)	14	-0.40 ± 0.56	$[-0.69,\ -0.14]$	0.019	−0.674

Notes: Baseline refers to the fixed $\mu = 0.7$ profile. Weighted 3-NN uses inverse-distance weighting over the three nearest handling-course bins in normalized curvature-feature space. Nearest uses the single closest handling-course bin. Values are paired mean difference ± standard deviation. Confidence intervals are participant-level bootstrap 95% confidence intervals for the paired mean difference. The primary p-value is from a paired t-test for lap time, continuous-task score, and lateral-error outcomes, and from a Wilcoxon signed-rank test for takeover count and total takeover time. g_z is the Hedges-corrected paired standardized effect size.

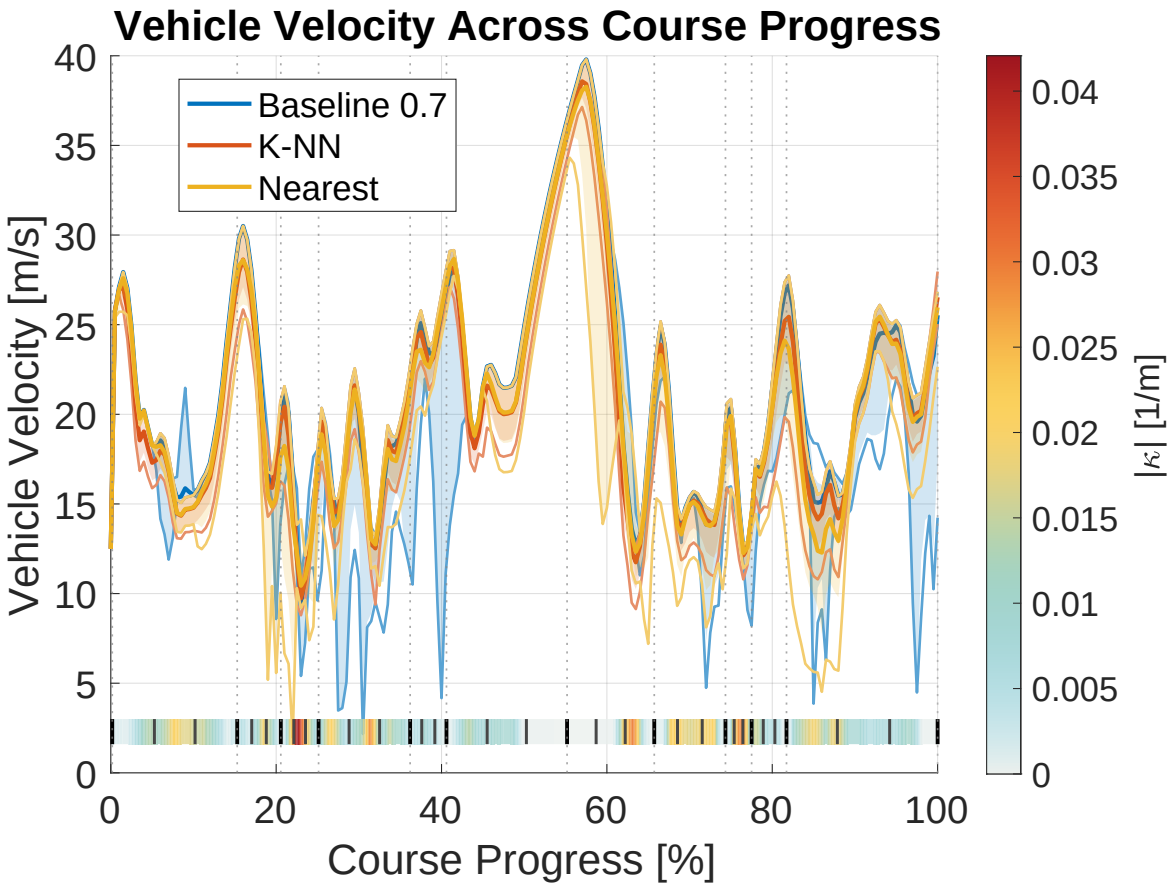

Figure 5.2: Velocity profile across the transfer methods for the total cohort in the aggressive course. Dark solid lines denote mean velocity, while shaded regions show +/- standard deviations, and light lines represent the max and mean values. The baseline standard deviation denotes the many times drivers lost control or had to slow down to regain control. The much narrower band for the 3-NN method shows consitent performance, despite the varied mu profiles for the total cohort.

total takeover times, and reduced lateral error compared to the single-nearest-neighbor profile.

Velocity averages are shown in Figure 5.2. The large negative standard deviations during the baseline run show the many locations in which participants lost control and had to come to a stop or near stop to regain control. The much narrower band for the 3-NN method shows consistent performance, despite the varied mu profiles for the total cohort.

5.4.2 Transferred Assumed-Friction Profiles

Figures 5.3 and 5.4 show how the assumed-friction profiles of the baseline and transferred methods vary over the aggressive course. These figures connect the route-context matching procedure to the observed driving outcomes. The baseline remains fixed at $\mu = 0.7$ throughout

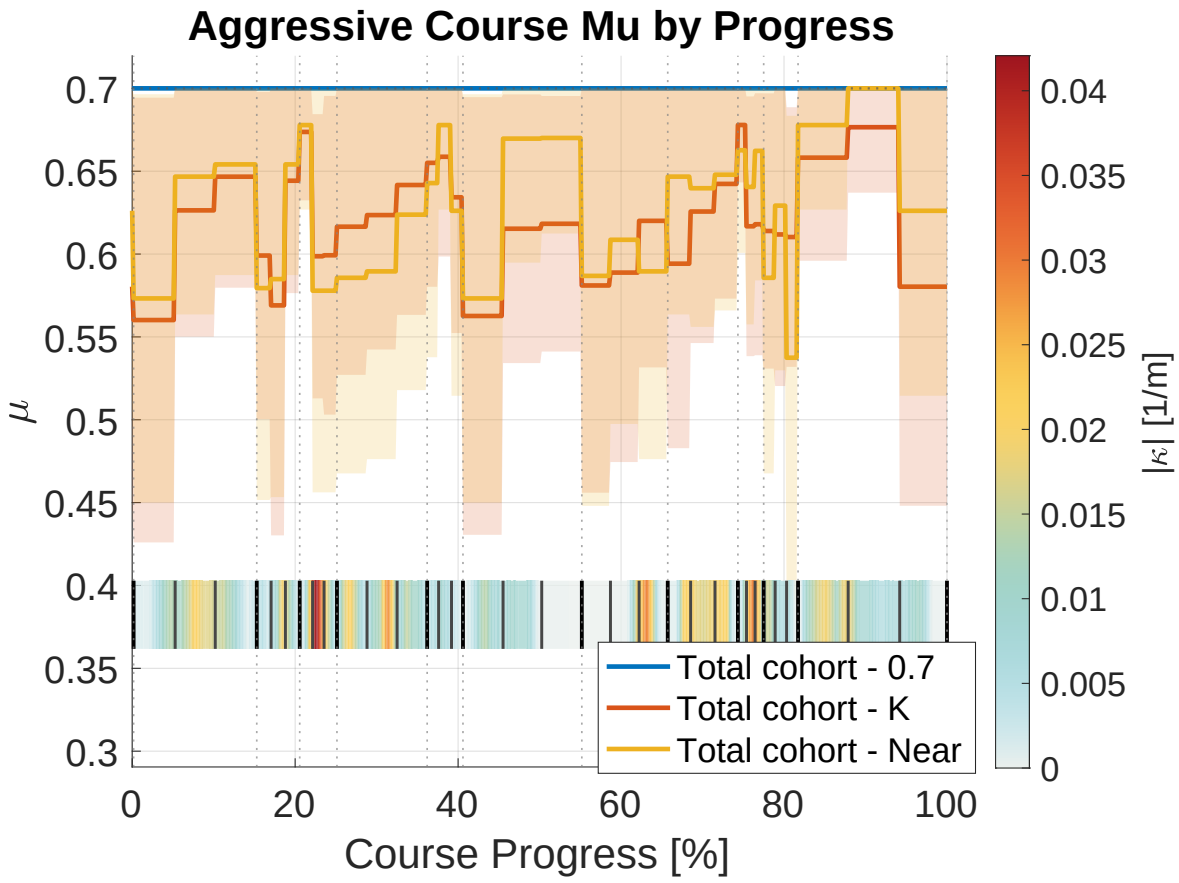

Figure 5.3: Transferred assumed-friction profiles across aggressive-course progress for all available transfer runs. Dark solid lines denote mean velocity while shaded regions show +/- standard deviations. The fixed baseline remains at $\mu = 0.7$; the weighted 3-NN and single-nearest-neighbor profiles are transferred from the handling-course learned library using normalized curvature-context features. The curvature overlay provides the geometric context used by the transfer feature vector.

the route, while the two transferred methods reduce the assumed friction in selected regions of the route. Lower assumed friction causes the backward-reachable planner to generate a more conservative reference-speed profile at those locations. The weighted three-nearest-neighbor profile varies more smoothly than the single-nearest-neighbor profile because it blends several nearby source contexts instead of transferring one source-bin value directly.

5.4.3 Intervention Changes Across the Aggressive Course

Figures 5.5 and 5.6 show takeover events throughout the aggressive-course progress. These plots are important because the claim of the chapter is spatial: the transferred profile should reduce interventions in target-route regions where the baseline speed policy remains too

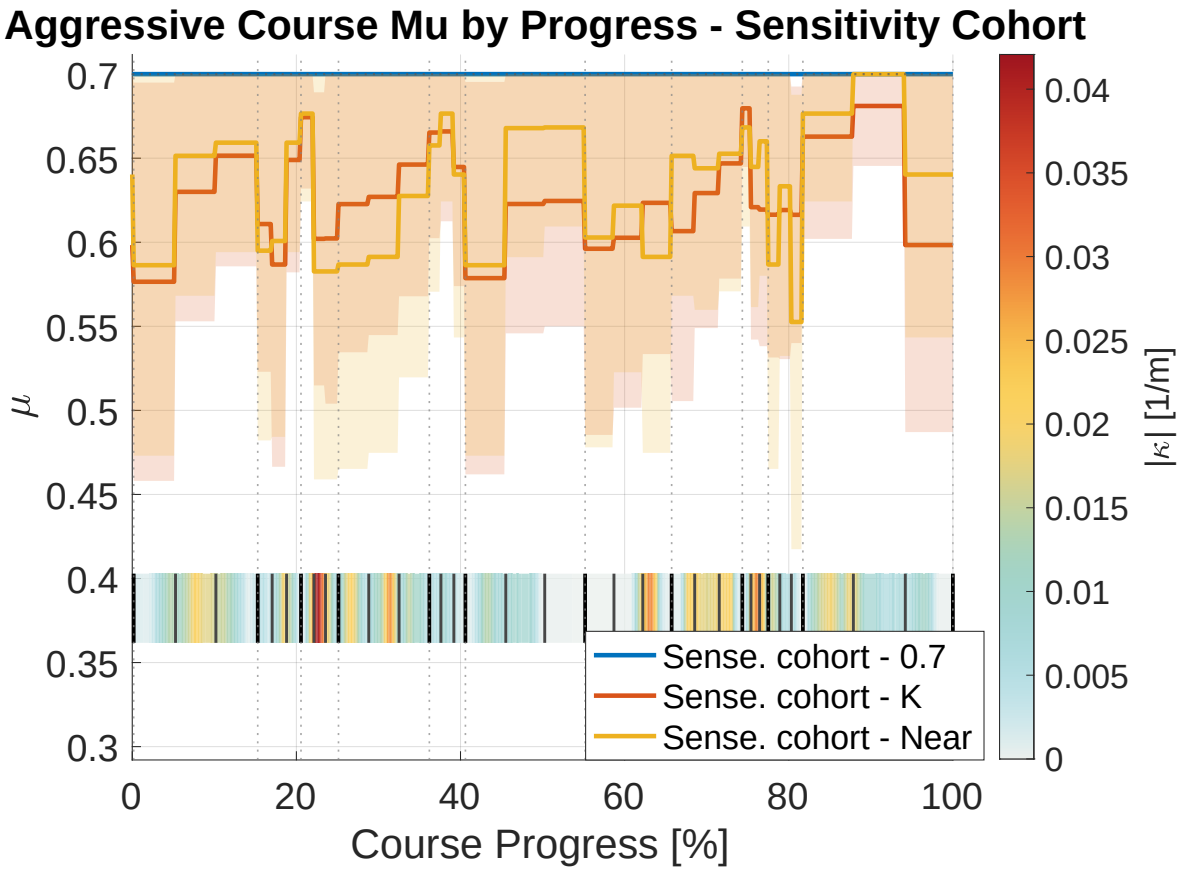

Figure 5.4: Transferred assumed-friction profiles across aggressive-course progress for the sensitivity cohort excluding participant 112. Dark solid lines denote mean velocity while shaded regions show +/- standard deviations. The fixed baseline, weighted 3-NN transfer, and single-nearest-neighbor transfer are shown using the same route-progress coordinate as Fig. 5.3. The similarity to the full-cohort profile indicates that the transferred profile structure is not driven solely by the high-intervention participant.

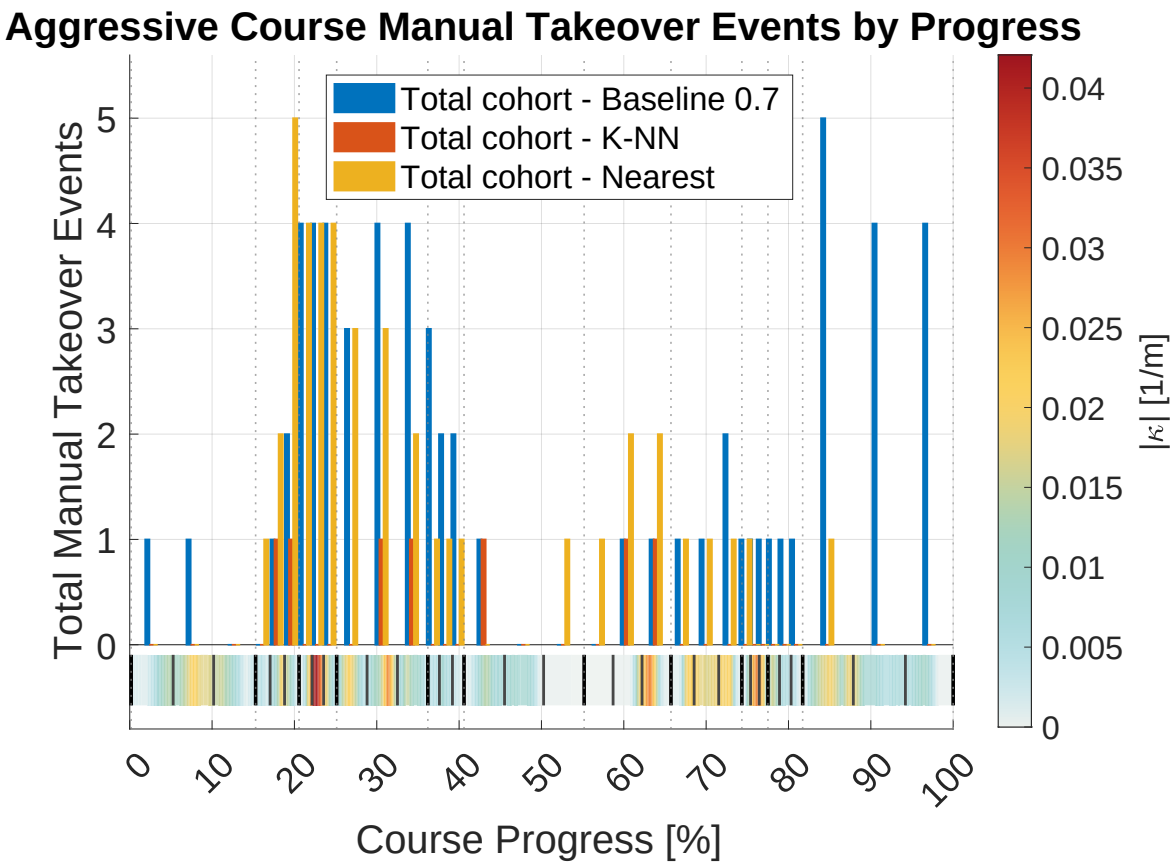

Figure 5.5: Manual takeover events across aggressive-course progress for all available transfer runs. Bars show the number of takeover events occurring at each course-progress location for the baseline $\mu = 0.7$ profile, the weighted 3-NN transferred profile, and the single-nearest-neighbor transferred profile. The curvature trace is included to relate intervention locations to route geometry; the aggregate event counts are reported in Table 5.2.

aggressive for the tested drivers. The weighted three-nearest-neighbor method produces the lowest total number of takeover events throughout the course in both the full and the sensitivity cohorts, consistent with Table 5.2. The single-nearest-neighbor method reduces takeovers relative to the baseline but retains more intervention events than the weighted method.

5.4.4 Driving Performance and Lateral Error

The transfer methods also change the vehicle's tracking behavior on the aggressive course. Table 5.2 shows that the weighted three-nearest-neighbor method produces the lowest mean lateral error for both cohorts. In the full cohort, the mean lateral error decreases from 1.03 ± 0.94 m below the baseline to 0.33 ± 0.04 m below the weighted transfer. In the

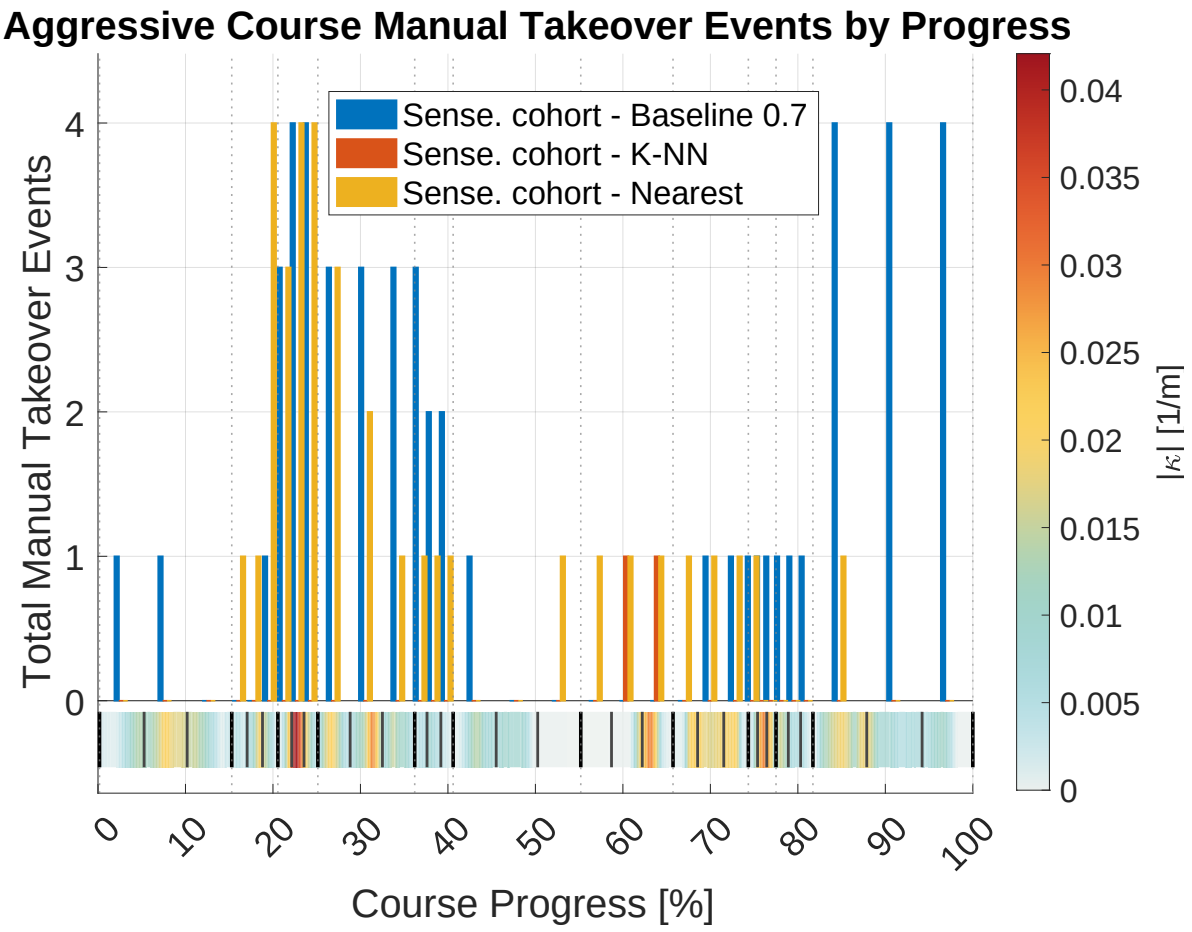

Figure 5.6: Manual takeover events across aggressive-course progress for the sensitivity cohort excluding participant 112. Bars show event counts for the fixed baseline, weighted 3-NN transfer, and single-nearest-neighbor transfer using the same route-progress coordinate as Figure 5.5. The corresponding paired statistical comparison is reported in Table 5.3 Changes noted in turns T2, T4, T6, T7, and T11.

sensitivity cohort, the mean lateral error decreases from 0.88 ± 0.77 m to 0.33 ± 0.04 m. The single-nearest-neighbor method also reduces mean lateral error relative to baseline, but less strongly.

Figures 5.7–5.8 show the spatial distribution of the signed lateral error along the course. The full-scale plots retain the large-error envelope and show where individual runs produce high lateral-error excursions. Zoomed plots make it easier to compare the mean behavior and standard-deviation bands in the range where most of the trajectories lie. These figures support the table-level result that the weighted transfer produces the most consistent reduction in lateral-error magnitude, whereas the nearest-neighbor transfer remains closer to the baseline in several route regions.

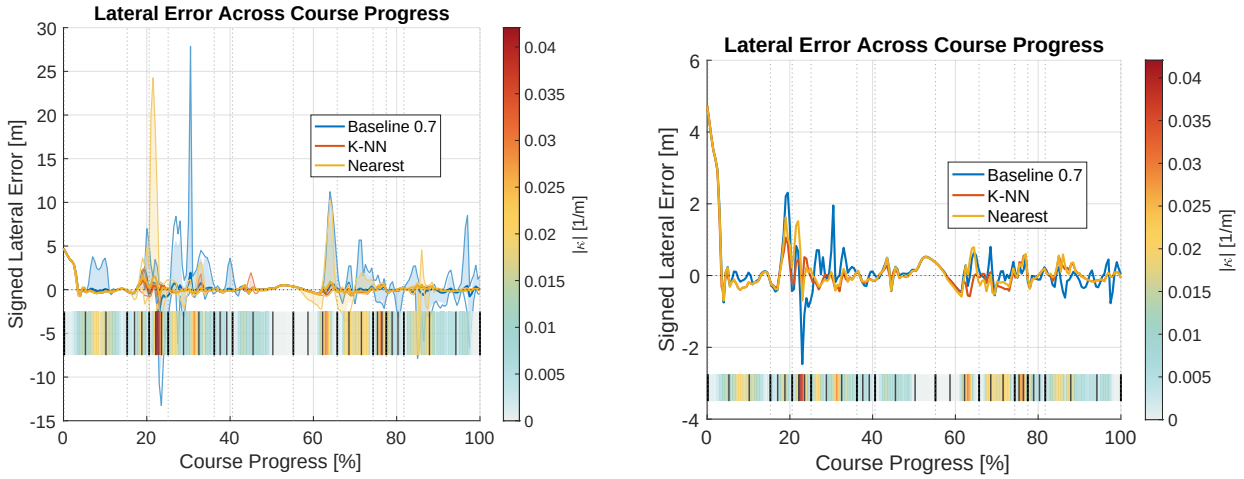

(a) Signed lateral error across aggressive-course progress for all available transfer runs.

(b) Zoomed signed lateral error across aggressive-course progress for all available transfer runs.

Figure 5.7: Signed lateral error across aggressive-course progress for all available transfer runs. Bold traces show condition means, shaded bands show approximately +/- standard deviation, and thin outer traces show the observed range across runs. The full-scale y-axis preserves large lateral-error excursions; Fig. 5.7b provides the corresponding zoomed view near the path centerline.

5.4.5 Transfer Diagnostics and Expectation-Mismatch Locations

The aggregate results support the transfer concept most strongly for the weighted three-nearest-neighbor method, but local diagnostic cases are needed to characterize the limits of the simplified route-context representation. A curvature-only feature vector can assign a target segment to a source segment that is similar in curvature but different in visual context, route width, entry speed, preceding maneuver, or vehicle transient response. Such cases

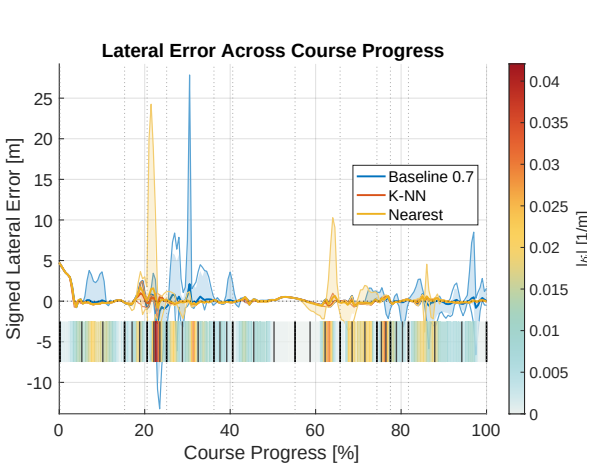

(a) Signed lateral error across aggressive-course progress for the sensitivity cohort excluding participant 112.

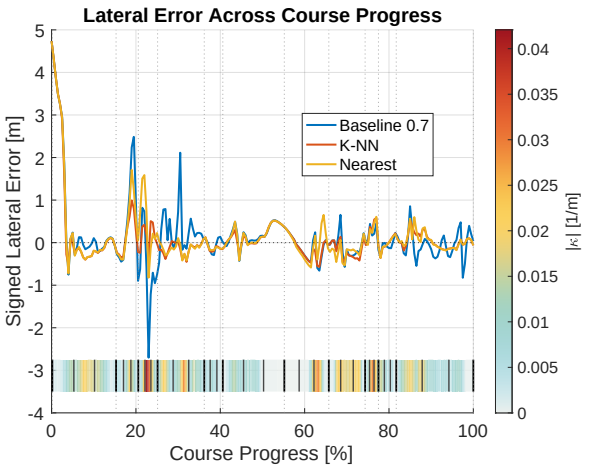

(b) Zoomed signed lateral error across aggressive-course progress for the sensitivity cohort excluding participant 112.

Figure 5.8: Signed lateral error across aggressive-course progress for the sensitivity cohort excluding participant 112. Bold traces show condition means, shaded bands show approximately ±1 standard deviation, and thin outer traces show the observed range across runs. The full-scale y-axis preserves large lateral-error excursions; Fig. 5.8b provides the corresponding zoomed view.

identify where the selected feature set is insufficient, not where the transfer algorithm is necessarily implemented incorrectly.

Table 5.4 summarizes local diagnostic examples from the aggressive course. The first group shows locations where the weighted three-nearest-neighbor and single-nearest-neighbor operators produced meaningfully different transferred assumed-friction values. The second group shows locations where the adaptive profile improved the local response but did not fully remove takeover or lateral-error limitations. These examples strengthen the bounded interpretation of the results: curvature-context transfer is useful for the evaluated course, but curvature alone does not explain every local outcome.

5.5 Discussion and Limitations

The aggressive-course results support the central claim of this chapter: a learned handling-course profile can be transferred to a novel course using curvature-context matching, and the transferred profile can reduce human takeover relative to a uniform assumed-friction baseline for the tested drivers. The strongest support comes from the weighted three-nearest-neighbor method. Relative to the $\mu = 0.7$ baseline, the weighted method reduces takeover events,

Table 5.4: Diagnostic aggressive-course exemplars illustrating limitations of curvature-feature transfer. The table includes locations where the weighted 3-NN and single-nearest-neighbor (Nearest) transfer operators produce different assumed-friction profiles (greater than 0.05), and locations where the adaptive profile improves the response but does not fully eliminate the local lateral error. Bolded entries highlight major differences.

Diagnostic type	Location	μ_{3NN}	μ_{NN}	$\lvert\Delta\mu\rvert$	v_{ref} B/3NN/NN	TO B/3NN/NN	$\bar{e}_y$ B/3NN/NN	Interpretation
Method divergence	T10P3	0.610	0.537	**0.073**	28.8/27.1/25.5	1/0/0	0.33/0.14/0.13	Largest disagreement between transfer operators; Nearest is substantially more conservative than weighted 3-NN.
	T6P2	0.615	0.670	**0.061**	22.7/21.1/22.0	0/0/0	0.09/0.08/0.09	Transfer operators disagree in the opposite direction; weighted 3-NN is more conservative than Nearest.
	T8P1	0.594	0.647	**0.053**	20.0/18.3/19.0	1/0/1	0.25/0.14/0.37	High feature-space distance and method disagreement indicate a less direct source–target match.
	T9P3	0.618	0.662	**0.052**	14.5/13.6/14.1	1/0/0	0.58/0.46/0.46	High feature-space distance with modest operator disagreement; useful as a transfer-ambiguity case.
	T11P3	0.580	0.626	**0.050**	23.2/21.0/21.9	4/**0**/**0**	0.87/**0.18**/**0.18**	Transfer operators differ, but both eliminate the local takeover events and reduce mean lateral error.
Outcome limitation	T2P3	0.644	0.654	0.010	19.5/18.6/18.7	2/1/**5**	1.99/0.84/**1.92**	Weighted 3-NN reduces takeovers and error, while Nearest increases takeovers relative to baseline.
	T3P1	0.674	0.678	0.013	18.1/17.4/17.3	4/0/**4**	0.94/0.35/**1.78**	Weighted 3-NN removes local takeovers; Nearest leaves takeovers unchanged and increases mean error.
	T3P2	0.599	0.578	0.038	11.1/10.3/10.1	4/0/4	1.87/**0.88**/1.29	Weighted 3-NN removes local takeovers, but residual mean error remains elevated.
	T4P2	0.624	0.590	0.043	17.4/16.4/15.9	4/**1**/3	1.56/0.24/0.69	Weighted 3-NN improves lateral error but does not eliminate intervention at this location.

Notes: B = baseline $\mu = 0.7$ profile; 3NN = weighted three-nearest-neighbor transfer; NN = single-nearest-neighbor transfer; TO = total local takeover events; $\bar{e}_y$ = mean lateral error in meters; v_{ref} = mean reference speed in m/s. Dashes indicate that the workbook did not report the corresponding local outcome or transferred-μ summary for that exemplar. Locations are labeled by target-course turn and part.

reduces total takeover time, increases continuous-task score, reduces mean lateral error, and does not produce a statistically supported increase in lap time. The same trend remains in the sensitivity cohort excluding participant 112. These results are consistent with the broader thesis that a simplified, data-informed model can be useful for control design when it captures the structure needed for the decision.

The single-nearest-neighbor method provides a more qualified result. It reduces takeover events and increases continuous-task score relative to the baseline, but it increases lap time and provides smaller improvements in lateral error. This comparison is useful because it separates two aspects of the transfer problem. Matching a target bin to one nearest source context can transfer useful driver-specific conservatism, but using only one source context is more sensitive to local mismatch in the feature space. The weighted three-nearest-neighbor method smooths across nearby learned contexts and gives a more robust profile in the evaluated aggressive-course condition.

The diagnostic examples in Table 5.4 clarify why the weighted method produced the stronger result. Several target bins produce meaningfully different transferred values under the two operators, even though both methods use the same curvature-feature vector. These cases show that the transfer outcome depends not only on the nearest source bin but also on the local density and consistency of nearby source-route examples. When the closest source bin is not representative of the neighboring examples, the single-nearest-neighbor method can transfer a value that is either too conservative or not conservative enough. The weighted method reduces this sensitivity by blending across the three closest source contexts.

The same diagnostic examples also show that curvature-based transfer is not an exhaustive explanation of local intervention behavior. In some bins, weighted three-nearest-neighbor transfer removes takeover events and reduces lateral error. In other bins, it improves the local response but does not fully eliminate intervention or residual error. These locations are important because they identify where the simplified representation is insufficient. Curvature is physically meaningful because it is directly related to lateral acceleration demand, but the route-context vector does not encode road width, visual preview, banking, elevation, entry speed, obstacle proximity, path boundary perception, or transient vehicle response. A participant may intervene because of any of these omitted features. Therefore, a successful reduction in takeover indicates that curvature-context transfer was sufficient for useful control adaptation in the tested condition, not that curvature is a complete model of intervention behavior.

The results should not be interpreted as proof of universal driver-profile transfer. The study evaluates one target course, one source course, and the tested drivers for whom source and target data were available. The chapter reports both condition-level summaries for all

available runs and participant-level paired statistics for the sensitivity cohort with matched data. The paired analysis strengthens the design-oriented conclusion that weighted three-nearest-neighbor transfer produced measurable improvement relative to the fixed baseline for the tested drivers, while the unequal full-cohort run counts and single-course evaluation still limit broader population-level claims. Additioanlly, the aggressive course contains esentially no elevation change, and the source (handling) course contains only minor changes. Further exploration is needed before claims about transfer applicability when major elevation changes are present.

Finally, the method transfers a speed-planning proxy rather than a psychological model. The learned μ profile does not represent physical road friction, trust, comfort, or intent. It is a route-local control parameter that changes the backward-reachable speed reference. A reduction in takeover under the transferred profile therefore supports a control-design conclusion: strategic speed reduction derived from a learned route-context library can reduce intervention for the tested drivers on the evaluated course. It does not show why the driver intervened and does not establish that the driver would behave the same way on a different course, in a different vehicle, or under different task demands.

5.6 Chapter Conclusion

This chapter presented a method for transferring a learned human-intervention disturbance-rejection profile from a source course to a novel route. The method represents each part of the route using curvature-derived features, normalizes those features using the source-course distribution, and maps the assumed restrictions of the learned source-route onto target-route bins using nearest-neighbor methods in feature space. Two transfer operators were evaluated: a weighted three-nearest-neighbor average and a single-nearest-neighbor assignment. The resulting target-course μ profiles were used to recompute the backward-reachable speed profile and were compared against a fixed baseline $\mu = 0.7$.

The weighted three-nearest-neighbor transfer produced the strongest result in the aggressive-course evaluation. This method reduced takeover events and total takeover time, increased continuous-task score, and reduced mean lateral error relative to the baseline. Lap time decreased numerically in the aggregate table, but the paired analysis did not support lap-time reduction as a primary effect; the main benefit was reduced intervention and improved secondary-task performance rather than faster course completion. The single-nearest-neighbor transfer also reduced takeovers and increased task score, but with a less favorable route-time and lateral-error tradeoff. These results demonstrate one practical way to construct a driver-specific feedforward profile for a novel route using a compact learned

library and transparent route-context matching.

The claims in this chapter are intentionally constrained. The findings provide bounded evidence that curvature-context matching can be used to transfer a learned assumed-friction profile for the tested drivers and the evaluated source–target course pair. Broader generalization requires richer route features, additional source and target courses, a larger driver population, and systematic analysis of locations where curvature-only matching is insufficient.

CHAPTER 6

Conclusion and Future Work

6.1 Dissertation Summary

This dissertation evaluated three application-specific uses of established control-oriented modeling principles. In each study, available physical, simulation, or human-in-the-loop data were organized into a compact representation of a behavior relevant to a specific input-design decision. The studies do not share a formal model-reduction algorithm, common state definition, or domain-independent identification procedure. Their connection is comparative: each asks whether a purpose-built representation can support a useful change to an input or reference policy under bounded test conditions. Chapter **??** used transition behavior in solid-state battery discharge to design feedforward current-density profiles. Chapter **??** used repeated human takeover on a known route as an event-based signal for route-local speed-policy adaptation. Chapter **??** used curvature-derived route-context features to transfer the resulting learned profile to one target course. Together, the studies provide comparative evidence that the established modeling philosophy was useful in the three evaluated settings; they do not establish a common architecture or general theory.

6.2 Research Contributions

6.2.1 Solid-state battery discharge input design

The first research contribution was a reduced-order transition region for feedforward discharge input design in solid-state batteries with evolving interfaces. The chapter used physical-cell data and high-fidelity simulation data generated through collaboration with battery-materials and battery-modeling researchers. The physical-cell data grounded the reduced-order characterization in the measured voltage response. The simulation data provided access to interface-health quantities, especially the contact area fraction, which are difficult to observe

directly during routine cell operation.

The battery results showed that transition-region discharge shaping can improve practical discharge performance relative to a matched constant-current baseline while preserving comparable final interface-health metrics in simulation. The most conservative profile increased the two-hour areal capacity by approximately 9.5 % relative to the matched case of lower-current constant-current while ending with a comparable overpotential and contact area fraction. That comparison is the appropriate basis for the capacity-improvement claim because the shaped profile is evaluated against a constant-current reference with comparable final interface health, not against a higher-current reference that produces greater degradation. More aggressive profiles showed the other side of the tradeoff: delaying the current density reduction can preserve a high early current density, but it can also degrade final interface health. The chapter therefore supports a bounded claim: a voltage-derived transition-region model can provide a useful design coordinate for feedforward discharge control, even though it does not fully resolve the interface physics.

6.2.2 Known-route human-intervention disturbance rejection

The second research contribution was a learning-based disturbance-rejection framework to reduce human takeover on a repeated known route.

This chapter focused on identifying takeover events as observed disturbances to the human–autonomy team. If takeover occurred repeatedly at similar route locations or under similar route conditions, the event provided evidence that the nominal autonomy behavior was not producing the desired closed-loop outcome for the tested driver.

The simplified representation in this chapter was a route-indexed assumed coefficient-of-friction profile for a point mass. The autonomy used this profile in a backward-reachable speed calculation. When takeover occurred in a route part or turn, the learning rule reduced the assumed coefficient of friction in the corresponding local region. Reducing the assumed friction caused the backward-reachable planner to generate a more conservative speed reference before the relevant route section. The learned coefficient-of-friction profile therefore acted as a control-oriented disturbance-rejection signal, not as a detailed behavioral model of the human.

The handling-course results support the hypothesis that route-local speed reduction can reduce repeated takeover events under the tested conditions. The paired analysis showed reduced takeover count and total takeover time for the full and sensitivity cohorts, an increase in continuous-task score, and a reduction in mean lateral error. Lap time did not provide evidence of faster course completion, and peak lateral error was weaker evidence than mean

lateral error. These results support the chapter's feasibility claim: repeated intervention can be used as a route-local control signal for adapting a speed profile, but the experiment does not isolate the learning rule from all possible familiarization or task-learning effects. The five non-intervention participants also define an important boundary case. For these participants, the takeover-based learning rule did not receive a disturbance signal, although the continuous-task scores still improved.

6.2.3 Transfer of learned profiles to a novel route

The third research contribution was a transfer framework for applying a learned human-intervention profile to a new route. The chapter used a compact route-context feature vector consisting of current-bin mean absolute curvature, current-bin peak absolute curvature, preceding-bin mean absolute curvature, preceding-bin peak absolute curvature, and normalized bin location within the turn. Two transfer operators were evaluated. The first used inverse-distance weighted three-nearest-neighbor averaging over the learned handling-course library. The second transferred the learned value from the single nearest source bin in the normalized feature space. Both transferred profiles were compared with a uniform baseline profile, $\mu = 0.7$.

The aggressive-course results support the transfer claim most strongly for the weighted three-nearest-neighbor method. Relative to the fixed baseline, the weighted method reduced takeover events and total takeover time, increased continuous-task score, and reduced average lateral error for the tested drivers. The single-nearest-neighbor method improved some outcomes relative to the baseline but produced a weaker tradeoff, particularly in lap time and lateral-error behavior. The paired participant-level analysis strengthened the aggregate result by showing that the weighted transfer effect was not only a run-level average. Diagnostic examples further showed that the simplified route-context representation has limits: locations with large differences between the weighted and single-nearest-neighbor transferred values, or locations with persistent error despite adaptation, identify where curvature-derived features do not encode all factors that influence intervention and vehicle response.

The transfer chapter therefore supports a limited but useful conclusion. Curvature-feature matching is one feasible way to construct a driver-specific feedforward profile for a new route, but it does not establish exhaustive transfer, universal driver generalization, or complete route understanding. Broader transfer claims require additional source courses, additional target courses, more participants, and systematic analysis of when a target route lies inside or outside the learned source-feature distribution.

Table 6.1: Claim–evidence boundaries for the research chapters.

Chapter	Supported claim	Primary evidence	Evidence boundary
Solid-state battery discharge input design	A voltage-derived transition-region representation can guide feedforward discharge shaping.	Simulation evaluation of shaped discharge profiles showed increased two-hour areal capacity for conservative profiles while preserving comparable overpotential and contact area fraction relative to matched constant-current references.	The shaped profiles are not globally optimized, and the strongest interface-health evidence is simulation based rather than physical-cell validation of the shaped inputs.
Known-route human-intervention learning	Repeated takeover events can be used as a route-local disturbance signal for speed-profile adaptation.	Paired handling-course statistics showed reduced takeover count, reduced takeover time, improved continuous-task score, and reduced mean lateral error for the full and sensitivity cohorts.	The experiment does not fully separate route-local learning from familiarization, secondary-task learning, or repeated exposure; participants with no intervention provide no takeover-based learning signal.
Novel-route profile transfer	A learned profile can be transferred to a new course using curvature-context matching, especially with weighted 3-NN averaging.	Aggressive-course aggregate and paired statistics showed that weighted 3-NN reduced takeover outcomes, improved continuous-task score, and reduced average lateral error relative to the fixed $\mu = 0.7$ baseline.	The result is limited to the tested drivers and evaluated source/target routes; curvature-only features do not capture all route, vehicle, or human factors.

6.3 Claim–Evidence Boundaries

Table 6.1 summarizes the main claims of the dissertation, the primary evidence supporting each claim, and the corresponding evidence boundary. The table is included to make the distinction between demonstrated feasibility and broader generalization.

6.4 Cross-Domain Lessons

Several lessons emerge in the three research chapters. First, a control model does not need to be a complete model of the underlying system. The battery chapter did not identify the full

morphology of the interface. The known-route driving chapter did not identify the cause of human intervention. The transfer chapter did not model all the route and driver features. In each case, the useful model was the one that preserved the structure needed to change the input.

Second, the quality of a simplified model depends on where the simplification is placed. The battery chapter simplified around a transition in discharge response. The known-route driving chapter simplified around repeated intervention events. The transfer chapter simplified around curvature-feature similarity. These are not arbitrary reductions. Each was chosen because it connects directly to an input-design decision.

Third, simplified models require disciplined validation. Because each representation is incomplete by design, validation must assess whether it is adequate for the intended decision under the evaluated conditions and must identify cases in which omitted variables limit its usefulness. The battery chapter requires interface-health metrics to confirm that increased capacity did not come at an unacceptable degradation cost. The chapter on known-route driving requires takeover, task-performance, lap-time, and path-following metrics to confirm that speed adaptation improved the coupled task rather than simply slowing the system. The transfer chapter requires expectation-mismatch analysis to show where the feature representation succeeds and where it fails.

Finally, interpretability is a practical advantage. The transition region, assumed-friction heatmap, and curvature-feature nearest-neighbor assignment can each be inspected. This matters because the dissertation's goal is not only to report performance improvements, but also to show how a simplified model changes the input.

6.5 Limitations

The dissertation has several limitations. In the battery chapter, the feedforward profiles were not globally optimized, and the strongest interface-health assessment was simulation based. Physical-cell validation of the shaped profiles is still needed before the method can be claimed as an experimentally validated discharge control. The transition-region model is also configuration dependent. An identified profile for one chemistry, pressure, and current-density range should not be assumed to apply unchanged to another cell design.

In the known-route driving chapter, the disturbance-rejection model treats takeover as an event signal. That framing is useful for control, but it intentionally omits the reason for the intervention. If a driver takes over because of a factor unrelated to speed or route context, reducing the assumed coefficient of friction may not address the underlying cause. The method also depends on repeated exposure to the same route. Without repetition, the

learning update cannot identify route-local intervention patterns.

In the transfer chapter, the feature representation is compact. Curvature-derived features are interpretable and route-based, but do not capture all factors that can affect intervention. Visual preview, road width, turn sequence, banking, grade, surface condition, speed history, traffic, and driver workload may all matter. The nearest-neighbor and weighted-nearest-neighbor transfer rules are transparent and easy to diagnose, but they represent only two possible transfer strategies. Richer feature sets and alternative weighting methods may be more effective when more data are available.

The human–autonomy results should be interpreted as applying to the tested drivers, tested routes, tested simulator configuration, and tested protocol. The dissertation does not claim that the learned profiles generalize to all drivers or to all courses. It evaluates whether the proposed simplifications produce useful control modifications in the studied cases and identifies the validation required for broader claims.

6.6 Future Work

6.6.1 Battery-control extensions

The battery work has three potential extensions that will be highlighted here. First, the shaped discharge profiles should be experimentally validated on physical cells. The current results use physical-cell data to inform the transition model and simulation data to evaluate interface health under shaped inputs. Experimental validation of the shaped profiles would test whether the simulated capacity and interface-health trends persist under real manufacturing variability, measurement noise, and cell-to-cell differences.

Second, the feedforward design should be replaced or supplemented with formal optimization. The profiles studied in Chapter 3 demonstrate the value of transition-informed shaping, but are not guaranteed to be optimal. Future work should optimize current-density trajectories subject to constraints on overpotential, contact-area proxy variables, uncertainty in transition timing, thermal limits, and application-specific power demands.

Third, the transition representation should be incorporated into feedback or adaptive control. A practical battery-management system would need to update its estimate of transition behavior as the cell ages and as operating conditions change. This motivates future work on online identification, health observers, uncertainty-aware constraints, and pack-level coordination across many cells.

6.6.2 Known-route disturbance-rejection extensions

The known-route driving framework should be tested in a larger group of participants and more route conditions. The present framework is intentionally driver-specific. Broader conclusions require analysis of whether different drivers produce similar learned friction maps, whether driver groups can be clustered, and whether learning gains persist when the task, route, or controller changes.

The learning update should also be compared against simpler baselines. A global speed reduction may reduce the intervention, but at the cost of unnecessary conservatism. The value of the proposed method depends on showing that route-local updates reduce takeover more efficiently than broad speed reductions. Future studies should therefore compare route-local learning against fixed conservative speed caps, manually tuned profiles, sham-update conditions, and alternative event-triggered update laws.

Future work should also investigate how the disturbance-rejection framework can coexist with richer human-state models. This dissertation intentionally avoided explaining the cause of the intervention. Once the disturbance-rejection baseline is established, additional human-state signals could be layered on top of it to decide when speed reduction is appropriate and when another type of autonomy response is needed.

6.6.3 Transfer-framework extensions

The transfer framework should be evaluated on more source and target routes. The present transfer study used the handling course as the learned source route and the aggressive course as the target route. That source–target pair is sufficient for a feasibility demonstration, but it is not enough to determine when a learned profile transfers reliably. Future work should include multiple handling-style source courses, multiple target courses with different curvature distributions, and repeated tests in which the same learned profile is transferred to several new routes. This would make it possible to distinguish a generally reusable route-context representation from a profile that happens to work for one source–target pair.

Future work should also compare additional transfer operators. The weighted three-nearest-neighbor method used inverse-distance weighting and performed better than the single-nearest-neighbor assignment in the present study. Other weighting strategies may provide different smoothness and locality tradeoffs, including inverse-squared-distance weighting, Gaussian kernel weighting, distance clipping, or adaptive weighting based on source-feature density. The number of neighbors should also be varied systematically. Using more than three neighbors may smooth noise and reduce sensitivity to one source bin, but it may also average over route contexts that are not truly similar. This tradeoff should be evaluated with matched

participant-level statistics and local diagnostic examples.

The route-context feature vector should be expanded carefully. Curvature is a useful first representation because it directly affects lateral-acceleration demand, but it omits many factors that can affect intervention. Future feature sets should test turn sequence, entry speed, road width, lateral-error history, preview distance, elevation, banking, surface condition, and visual boundary information. These features should be added only when they improve the transfer enough to justify the added dimensionality and data requirement.

Finally, future transfer studies should define explicit interpolation and extrapolation diagnostics. A target bin that lies close to several consistent source bins is different from a target bin that lies far from the learned library or near conflicting source examples. Future methods should report feature-space distance, source-neighbor agreement, and out-of-distribution indicators alongside the transferred profile. These diagnostics would make it possible to decide when to apply a transferred profile, when to return to a conservative baseline, and when additional driver-specific learning is required.

6.7 Closing Remarks

The central contribution of this dissertation is the construction, use, and evaluation of three application-specific control-oriented representations. The battery study uses a transition-region representation to shape discharge current; the known-route study uses repeated takeover events to adapt a route-local speed-planning parameter; and the transfer study uses curvature-context matching to map a learned profile to one target route. These representations are linked by an established purpose-built modeling philosophy, not by a common cross-domain algorithm or universal reduction procedure. Across the three studies, usefulness depends on alignment among the representation, the control decision, and the available evidence. The results show that compact data-informed representations can support beneficial input or reference changes under the tested conditions, while the reported limitations identify the additional experiments, routes, participants, and validation required for broader claims.

APPENDIX A

Supplementary Battery Modeling Details

The battery testing and system identification methods were conducted in conjunction with colleagues from both Purdue and The Georgia Institute of Technology. Physical cell testing was conducted by a research team including Sun Geun Yoon. The data was then sent to Keith Ng who was the primary researcher responsible for fitting the discharge data. Some of the system identification data is presented below to help demonstrate the origin of the physical cell data. However, the production, testing, and analysis of the physical cell data was not conducted by the author, so the data is presented here in the appendix to clearly distinguish it as not the author's work.

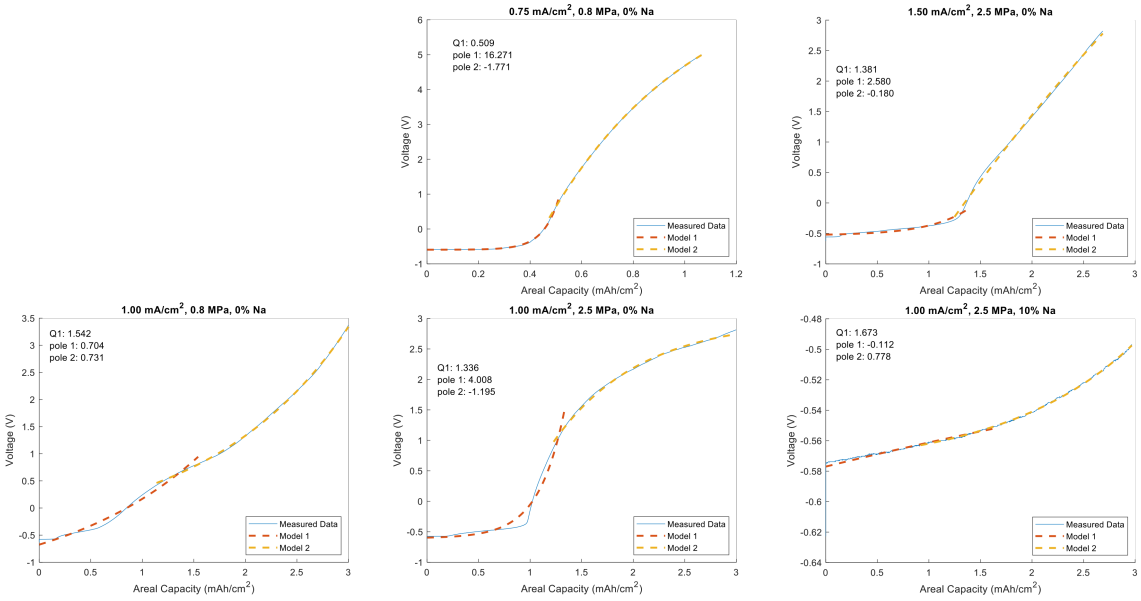

Figure A.1: Example physical cells used for anode-stripping experiments. The images show cells with a lithium anode and a lithium–sodium alloy anode, illustrating the experimental configurations that support the battery-interface modeling chapter.

BIBLIOGRAPHY

Paul Albertus, Susan Babinec, Scott Litzelman, and Aron Newman. Status and challenges in enabling the lithium metal electrode for high-energy and low-cost rechargeable batteries. Nature Energy, 3(1):16–21, 2018. doi: 10.1038/s41560-017-0047-2.

Arsha Ali, Rohit Banerjee, Wonse Jo, Samantha Dubrow, Kayla Riegner, Jonathon M. Smereka, Lionel P. Robert Jr., and Dawn M. Tilbury. Spot report: An open-source and real-time secondary task for human-robot interaction user experiments. In ACM/IEEE International Conference on Human-Robot Interaction, 2024. ISBN 9798400703232.

Suguru Arimoto, Sadao Kawamura, and Fumio Miyazaki. Bettering operation of robots by learning. Journal of Robotic Systems, 1(2):123–140, 1984. doi: 10.1002/rob.4620010203.

A. Bates, S. Mukherjee, N. Schuppert, B. Son, J. Kim, and S. Park. Modeling and simulation of 2d lithium-ion solid state battery. International Journal of Energy Research, 39(11): 1505–1518, 2015. doi: 10.1002/er.3352.

Peter Benner and Pawan Goyal. Balanced truncation for quadratic-bilinear control systems. Advances in Computational Mathematics, 50:88, 2024. doi: 10.1007/s10444-024-10186-9. Article 88.

Davide Bistri, Amir Afshar, and Claudio V. Di Leo. Modeling the chemo-mechanical behavior of all-solid-state batteries: A review. Meccanica, 56:1523–1554, 2021. doi: 10.1007/s11012-021-01341-z.

Douglas A. Bristow, Marina Tharayil, and Andrew G. Alleyne. A survey of iterative learning control. IEEE Control Systems Magazine, 26(3):96–114, 2006. doi: 10.1109/MCS.2006.1636313.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the National Academy of Sciences, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.

C. Ricardo Constante-Amores and Michael D. Graham. Data-driven state-space and Koopman operator models of coherent state dynamics on invariant manifolds. Journal of Fluid Mechanics, 984:R9, 2024. doi: 10.1017/jfm.2024.284.

Thomas M. Cover and Peter E. Hart. Nearest neighbor pattern classification. IEEE Transactions on Information Theory, 13(1):21–27, 1967. doi: 10.1109/TIT.1967.1053964.

Can Cui, Zichong Yang, Yupeng Zhou, Yunsheng Ma, Juanwu Lu, Lingxi Li, Yaobin Chen, Jitesh Panchal, and Ziran Wang. Personalized autonomous driving with large language models: Field experiments. In 2024 IEEE 27th International Conference on Intelligent Transportation Systems, pages 20–27, 2024. doi: 10.1109/ITSC58415.2024.10919978.

Na Du, Feng Zhou, Elizabeth M. Pulver, Dawn M. Tilbury, Lionel P. Robert, Anuj K. Pradhan, and X. Jessie Yang. Examining the effects of emotional valence and arousal on takeover performance in conditionally automated driving. Transportation Research Part C: Emerging Technologies, 112:78–87, 2020.

Sahibsingh A. Dudani. The distance-weighted k-nearest-neighbor rule. IEEE Transactions on Systems, Man, and Cybernetics, SMC-6(4):325–327, 1976. doi: 10.1109/TSMC.1976.5408784.

Alexander Eriksson and Neville A. Stanton. Takeover time in highly automated vehicles: Noncritical transitions to and from manual control. Human Factors, 59(4):689–705, 2017. doi: 10.1177/0018720816685832.

Chris T. Freeman and Ying Tan. Iterative learning control with mixed constraints for point-to-point tracking. IEEE Transactions on Control Systems Technology, 21(3):604–616, 2013. doi: 10.1109/TCST.2012.2187787.

John Gaspar and Cher Carney. The effect of partial automation on driver attention: A naturalistic driving study. Human Factors, 61, 2019. doi: 10.1177/0018720819836310.

Michel Gevers and Lennart Ljung. Optimal experiment designs with respect to the intended model application. Automatica, 22(5):543–554, 1986. doi: 10.1016/0005-1098(86)90046-2.

Kelsey B. Hatzell, Xi Chelsea Chen, Corie L. Cobb, Neil P. Dasgupta, Marm B. Dixit, Lauren E. Marbella, Matthew T. McDowell, Partha P. Mukherjee, Ankit Verma, Venkatasubramanian Viswanathan, Andrew S. Westover, and Wolfgang G. Zeier. Challenges in lithium metal anodes for solid-state batteries. ACS Energy Letters, 5(3):922–934, 2020a. doi: 10.1021/acsenergylett.9b02668.

Kelsey B. Hatzell, Xi Chelsea Chen, Corie L. Cobb, Neil P. Dasgupta, Marm B. Dixit, Lauren E. Marbella, Matthew T. McDowell, Partha P. Mukherjee, Ankit Verma, Venkatasubramanian Viswanathan, Andrew S. Westover, and Wolfgang G. Zeier. Challenges in lithium metal anodes for solid-state batteries. ACS Energy Letters, 5(3):922–934, 2020b. doi: 10.1021/acsenergylett.9b02668.

Xiaosong Hu, Shengbo Li, and Huei Peng. A comparative study of equivalent circuit models for li-ion batteries. Journal of Power Sources, 198:359–367, 2012. doi: 10.1016/j.jpowsour.2011.10.013.

J"urgen Janek and Wolfgang G. Zeier. A solid future for battery development. Nature Energy, 1(9):1–4, 2016. doi: 10.1038/nenergy.2016.141.

Haresh Karnan, Elvin Yang, Garrett Warnell, Joydeep Biswas, and Peter Stone. Wait, that feels familiar: Learning to extrapolate human preferences for preference-aligned path planning. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pages 13008–13014, 2024a. doi: 10.1109/ICRA57147.2024.10611475.

Haresh Karnan, Elvin Yang, Garrett Warnell, Joydeep Biswas, and Peter Stone. Wait, that feels familiar: Learning to extrapolate human preferences for preference-aligned path planning. In Proceedings of the IEEE International Conference on Robotics and Automation, pages 13008–13014, 2024b. doi: 10.1109/ICRA57147.2024.10611475.

Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. Automatica, 93:149–160, 2018. doi: 10.1016/j.automatica.2018.07.004.

Christos Lathourakis and Alice Cicirello. Physics enhanced sparse identification of dynamical systems with discontinuous nonlinearities. Nonlinear Dynamics, 112(13):11237–11264, 2024. doi: 10.1007/s11071-024-09652-2.

K. Lee and J. Sakamoto. Li stripping behavior of anode-free solid-state batteries under intermittent-current discharge conditions. Advanced Energy Materials, 2024. Citation details to verify before final submission.

Zhaoyang Li, Minghao Han, Dat-Nguyen Vo, and Xunyuan Yin. Machine learning-based input-augmented Koopman modeling and predictive control of nonlinear processes. Computers & Chemical Engineering, 191:108854, 2024. doi: 10.1016/j.compchemeng.2024.108854.

Ingyu Lim and Kira L. Barton. Pareto iterative learning control: Optimized control for multiple performance objectives. Control Engineering Practice, 26:125–135, 2014. ISSN 0967-0661. doi: https://doi.org/10.1016/j.conengprac.2014.01.011. URL https://www.sciencedirect.com/science/article/pii/S0967066114000410.

Kailong Liu, Yizhao Gao, Chong Zhu, Kang Li, Minrui Fei, Chen Peng, Xi Zhang, and Qing-Long Han. Electrochemical modeling and parameterization towards control-oriented management of lithium-ion batteries. Control Engineering Practice, 124:105176, 2022. doi: 10.1016/j.conengprac.2022.105176.

Lennart Ljung. System Identification: Theory for the User. Prentice Hall, 2 edition, 1999.

Alan C. Luntz, Johannes Voss, and Karsten Reuter. Interfacial challenges in solid-state li ion batteries. The Journal of Physical Chemistry Letters, 6(22):4599–4604, 2015. doi: 10.1021/acs.jpclett.5b02352.

Ruikun Luo, Yifan Weng, Yifan Wang, Paramsothy Jayakumar, Mark J. Brudnak, Victor Paul, Vishnu R. Desaraju, Jeffrey L. Stein, Tulga Ersal, and X. Jessie Yang. A workload adaptive haptic shared control scheme for semi-autonomous driving, 2020.

Scott G. Marquis, Valentin Sulzer, Robert Timms, Colin P. Please, and S. Jon Chapman. An asymptotic derivation of a single particle model with electrolyte. Journal of The Electrochemical Society, 166(15):A3693–A3706, 2019. doi: 10.1149/2.0341915jes.

Andrew Masias, Nathaniel Felten, Regina Garcia-Mendez, Jeff Wolfenstine, and Jeff Sakamoto. Elastic, plastic, and creep mechanical properties of lithium metal. Journal of Materials Science, 54:2585–2600, 2019. doi: 10.1007/s10853-018-2971-3.

Bruce C. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. IEEE Transactions on Automatic Control, 26(1):17–32, 1981. doi: 10.1109/TAC.1981.1102568.

George Nasser, Ben W. Morrison, and Mark W. Wiggins. Mental models and driver takeover in automated vehicles: A systematic review. Ergonomics, pages 1–14, 2025. doi: 10.1080/00140139.2025.2514599. Advance online publication.

Mei-Chin Pang, Kai Yang, Rowena Brugge, Teng Zhang, Xinhua Liu, Feng Pan, Shichun Yang, Ainara Aguadero, Billy Wu, Monica Marinescu, Huizhi Wang, and Gregory J. Offer. Interactions are important: Linking multi-physics mechanisms to the performance and degradation of solid-state batteries. Materials Today, 49:145–183, 2021. doi: 10.1016/j.mattod.2021.03.013.

Gregory L. Plett. Battery Management Systems, Volume I: Battery Modeling. Artech House, 2015.

Sean Reiter, Tobias Damm, Mark Embree, and Serkan Gugercin. On the balanced truncation error bound and sign parameters from arrowhead realizations. Advances in Computational Mathematics, 50:10, 2024. doi: 10.1007/s10444-024-10105-y. Article 10.

Clarence W. Rowley. Model reduction for fluids, using balanced proper orthogonal decomposition. Journal of Fluid Mechanics, 534:115–173, 2005. doi: 10.1017/S0022112005004479.

SAE International. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. SAE Standard J3016$_2$02104, $SAEInternational$, 2021.

Marios Sekadakis and George Yannis. Systematic review and meta-analysis of take-over time from automated driving at SAE levels 2 and 3 to manual control. Transportation Research Part F: Traffic Psychology and Behaviour, 113:263–306, 2025. doi: 10.1016/j.trf.2025.04.003.

Marios Sekadakis, Stefan Tr"osterer, Philipp Moertl, and George Yannis. Identifying driving profiles after take over request in automated vehicles at SAE levels 2 and 3. Transportation Research Part F: Traffic Psychology and Behaviour, 111:250–263, 2025. doi: 10.1016/j.trf.2025.03.007.

Dharmendra K. Singh, Tim Fuchs, Christoph Krempaszky, Philipp Schweitzer, Christoph Lerch, Felix H. Richter, and J"urgen Janek. Origin of the lithium metal anode instability in solid-state batteries during discharge. Matter, 6(5):1463–1483, 2023. doi: 10.1016/j.matt.2023.03.009.

C. Speltino, D. Di Domenico, G. Fiengo, and A. Stefanopoulou. Comparison of reduced order lithium-ion battery models for control applications. In Proceedings of the 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, pages 3276–3281, 2009. doi: 10.1109/CDC.2009.5400618.

Hongfeng Tao, Jian Li, Yiyang Chen, Vladimir Stojanovic, and Huizhong Yang. Robust point-to-point iterative learning control with trial-varying initial conditions. IET Control Theory & Applications, 14(19):3344–3350, 2020. doi: https://doi.org/10.1049/iet-cta.2020.0557. URL https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-cta.2020.0557.

Bairav S. Vishnugopi, K. G. Naik, H. Kawakami, N. Ikeda, Y. Mizuno, R. Iwamura, T. Kotaka, K. Aotani, Y. Tabuchi, and Partha P. Mukherjee. Asymmetric contact loss dynamics during plating and stripping in solid-state batteries. Advanced Energy Materials, 13(8):2203671, 2023. doi: 10.1002/aenm.202203671.

Fujin Wang, Zhi Zhai, Zhibin Zhao, Yi Di, and Xuefeng Chen. Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis. Nature Communications, 15:4332, 2024. doi: 10.1038/s41467-024-48779-z. Article 4332.

Xin Wang, Long Guo, and Yuanqing Jia. Online sensing of human steering intervention torque for autonomous driving actuation systems. IEEE Sensors Journal, 18(8):3444–3453, 2018. doi: 10.1109/JSEN.2018.2805381.

Yujie Wang, Xuebing Zhang, Kang Li, Guang Zhao, and Zonghai Chen. Perspectives and challenges for future lithium-ion battery control and management. eTransportation, 18: 100260, 2023. doi: 10.1016/j.etran.2023.100260.

Zheng Wang, Satoshi Suga, Edric John Cruz Nacpil, Zhanhong Yan, and Kimihiko Nakano. Adaptive driver-automation shared steering control via forearm surface electromyography measurement, 2020.

Sun Geun Yoon, Adrian Xiao Bin Yong, Congcheng Wang, Cheng-Han Tsai, Abhinand Ayyaswamy, Bairav S. Vishnugopi, Xin Wang, Seungwoo Park, Ming Gong, Gajanand Pawar, Partha P. Mukherjee, Neil P. Dasgupta, and Matthew T. McDowell. Interface morphogenesis with a deformable secondary phase in solid-state lithium batteries. Science, 388: 1062–1068, 2025. doi: 10.1126/science.adt5229.

Zhenhua Yu, Xin Zhao, Kang Jiang, et al. Analyzing the effects of human error on automated driving takeovers. Cognition, Technology & Work, 27:215–239, 2025a. doi: 10.1007/s10111-025-00794-w.

Zihan Yu, Aleksandra Dudek, Patrick Linford, Ian Holbrook, Scott James, Matthew Castanier, Kira Barton, and Chris Vermillion. Iterative learning-based arbitration in shared vehicle control: Methodological formulation and driver-in-the-loop results. In IEEE Conference on Control Technology and Applications, 2025b.

Bo Zhang, Joost C. F. de Winter, Silvia F. Varotto, Riender Happee, and Brady M. Kuhl. Determinants of take-over time from automated driving: A meta-analysis of 129 studies. Transportation Research Part F: Traffic Psychology and Behaviour, 64:285–307, 2019. doi: 10.1016/j.trf.2019.04.020.

Xiaoming Zhang, Q. J. Wang, Katharine L. Harrison, Scott A. Roberts, and Stephen J. Harris. Pressure-driven interface evolution in solid-state lithium metal batteries. Cell Reports Physical Science, 1(2):100012, 2020. doi: 10.1016/j.xcrp.2019.100012.

Zhenfa Zhang, Dong Shen, and Xinghuo Yu. Iterative learning control for pareto optimal tracking in incompatible multisensor systems. IEEE Transactions on Cybernetics, 55(3): 1093–1106, 2025. doi: 10.1109/TCYB.2024.3514688.

Xiaolong Zhong et al. Human-automated vehicle interactions: Voluntary driver intervention in car-following. Transportation Research Part C: Emerging Technologies, 171:104969, 2025.