AMERICAN JOURNAL OF MANAGEMENT AND IOT MEDICAL COMPUTING
Peer Reviewed, Referred & Indexed Journal
E-ISSN: 3069-0110
Vol.5, No.4 (2026)
www.ajmimc.com

DATA-DRIVEN ADAPTIVE MODEL PREDICTIVE CONTROL FOR CONSTRAINED NON LINEAR SYSTEMS

G. APARNA, CH. VISHNU CHAKRAVARTHI

1.Student, Department of Electrical and Electronics Engineering, Sanketika Vidya Parishad Engineering College, Pothinamallayya Palem (P.M. Palem), Visakhapatnam, Andhra Pradesh, India–530041 Email: aparnaguravana07@gmail.com

2.Assistant Professor, Department of Electrical and Electronics Engineering, Sanketika Vidya Parishad Engineering College, Pothinamallayya Palem (P.M. Palem), Visakhapatnam, Andhra Pradesh, India–530041 Email: svpec.vishnu@gmail.com

ABSTRACT

Electronic throttle systems (ETS) combine a DC motor, a gearbox, a limp-home return spring, Coulomb friction and airflow torque, and must track the commanded plate angle within a 0–24 V supply. The base paper of this thesis (Youssef, Wafa and Shalaby, 2025) addresses this with an Adaptive Relaxed Sliding Mode Controller (ARSMC) in which a Deep Deterministic Policy Gradient (DDPG) actor–critic agent adapts the lumped-uncertainty bound, shrinking the boundary layer that removes chattering.

This thesis reproduces the ETS model and the SMC, RSMC and ARSMC controllers of the base paper and extends the work with an offset-free Model Predictive Controller (MPC). The MPC feedback-linearises the throttle, estimates angle, speed and the lumped disturbance with a third-order extended state observer (ESO), and solves a 40-variable box-constrained quadratic program every 5 ms, so the 0–24 V limit is enforced inside the optimisation rather than by clipping.

All controllers are compared in MATLAB on the same staircase reference (0.4, 0.8 and 1.5 rad), parametric uncertainty and airflow disturbance, in three scenarios: uncertainty inside the design bound, uncertainty 50% above it, and 2 mrad sensor noise. Compared with ARSMC, the proposed MPC reduces the integral absolute error by 39%, 36% and 78%, the integral squared error by 37%, 34% and 73%, and the rise time by 41%, 38% and 78% in the three scenarios, with zero steady-state error and a chatter-free voltage. Its cost is a 4–6 mrad overshoot, which ARSMC avoids. The MPC needs no training and solves its QP in about 0.6 ms.

Keywords: electronic throttle, DC motor, sliding mode control, reinforcement learning, model predictive control, extended state observer, constrained control.

1: INTRODUCTION AND PROBLEM STATEMENT

1.1 Background

For most of the twentieth century the throttle valve of a petrol engine was opened by a steel Bowden cable connected directly to the accelerator pedal. The driver's foot set the plate angle, and therefore the air mass entering the cylinders. This arrangement was simple and reliable, but it gave the engine management system no authority over the most important input to the combustion process.

Modern spark-ignition engines use drive-by-wire throttles. The driver's pedal no longer pulls a cable; instead, a pedal position sensor sends the driver's request to the Engine Control Unit

AMERICAN JOURNAL OF MANAGEMENT AND IOT MEDICAL COMPUTING
Peer Reviewed, Referred & Indexed Journal
E-ISSN: 3069-0110
Vol.5, No.4 (2026)
www.ajmimc.com

(ECU). The ECU combines that request with other demands — idle-speed control, cruise control, traction and stability control, gear-shift torque reduction, catalyst warm-up and emission strategies — and computes a desired plate angle θd. A small DC motor, driven by a pulse-width-modulated chopper, then positions the throttle plate to that angle. Air intake, and therefore engine torque, fuel economy and tailpipe emissions, depend directly on how accurately and how quickly the actual angle θ follows θd.

The performance demanded from the throttle loop is high. A typical production specification asks for a 10–90% response to a large step in well under 100–200 ms, no overshoot that could cause a torque surge, a steady-state error below about 0.1° of plate angle, and robust behaviour over a temperature range of −40 °C to +125 °C, battery voltage variations, ageing of the return spring and wear of the gearbox. These requirements must be met by a controller that runs on an automotive microcontroller alongside many other engine functions.

The ETS is therefore also a classic benchmark for nonlinear control, because it combines several hard nonlinearities in a small, fast mechanism:

- a dual return spring (limp-home spring) that produces a discontinuous torque at the default opening angle θo,
- Coulomb (static) friction that depends on the sign of plate speed,
- a gearbox with backlash and an inertia that varies between production units,
- airflow torque on the plate that acts as an unknown, time-varying load,
- a hard actuator limit, because the motor is supplied from the 0–24 V vehicle chopper.

The limp-home spring deserves special mention. It holds the plate slightly open (θo ≈ 0.17 rad in the base paper) when the motor is not powered, so that a vehicle with a failed throttle can still be driven at low speed. The price is a torque discontinuity of twice the spring offset at θo. A controller must apply about 10 V before the plate moves at all from its rest position, and the sign of the required torque reverses as the plate crosses θo. Linear controllers tuned around one operating point behave very differently on either side of this point.

ETS characteristic	Physical origin	Consequence for the controller
Limp-home spring offset	Two opposing springs preloaded at θo	Torque jump, sign change at θo, large holding voltage
Coulomb friction	Bearings and seals	Stick-slip, steady-state error at low speed
Gear backlash and inertia spread	Plastic gearbox, production tolerance	Model uncertainty, possible limit cycles
Airflow torque	Pressure difference across the plate	Unknown, load-dependent disturbance
0–24 V supply	Chopper fed from the vehicle supply	Hard input constraint, slows large steps
Sensor noise	Potentiometric or Hall-effect TPS	Noisy angle, speed must be estimated

1.2 The Drive-by-Wire Throttle System

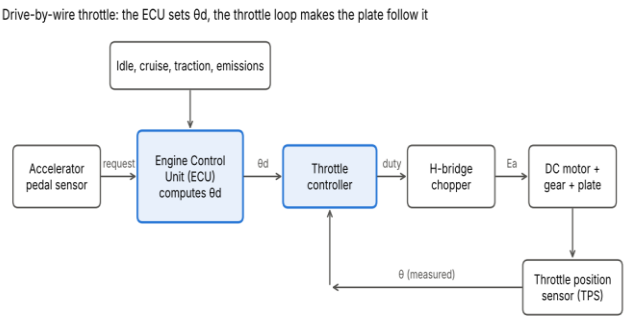


Figure 1.1 — Signal path of a drive-by-wire throttle

A drive-by-wire throttle is a closed chain of sensing, decision and actuation. Each link in the chain contributes to the overall performance and to the difficulty of the control problem.

Accelerator pedal module. The pedal carries two independent position sensors with different characteristics, usually two Hall-effect channels with different slopes. The ECU compares the two signals continuously; a mismatch is treated as a fault and the engine is placed in a limited-power mode. The pedal position is interpreted as a torque request, not as a throttle angle.

Engine control unit. The ECU converts the driver's torque request into an air-mass demand and then into a desired plate angle θd. Many other functions compete for the same actuator: idle-speed control keeps the engine running at a fixed speed when the driver lifts off; cruise control and adaptive cruise control generate their own torque requests; traction and stability control reduce torque when wheels slip; the transmission controller requests short torque reductions during gear shifts; and the emission strategy shapes the air path during catalyst heating. All these requests are arbitrated in a torque structure, and the result is a single reference θd sent to the throttle position controller every few milliseconds.

Throttle position controller. This is the subject of this thesis. It receives θd and the measured angle θ and computes the motor voltage. Because the reference can change abruptly, for example in a tip-in, a gear-shift intervention or a traction-control event, the controller must handle large steps as well as slow tracking and must do so at a fixed sample rate.

Power stage and actuator. An H-bridge chopper converts the controller's duty-cycle command into an average armature voltage. A brushed permanent-magnet DC motor drives the plate shaft through a two-stage plastic gearbox. The gearbox multiplies motor torque and reduces the speed, but also introduces backlash and friction.

Throttle position sensor. Like the pedal, the throttle body carries two redundant angle sensors, potentiometric in older designs and contactless Hall-effect or inductive sensors in newer ones. Their signals are noisy and quantised by the ECU's analogue-to-digital converter. Only the angle is measured; plate speed must be estimated.

Safety monitoring. Because an uncontrolled throttle can accelerate the vehicle unexpectedly, the throttle function is classified as safety-relevant under ISO 26262. A monitoring layer checks plausibility of the commanded and measured angles. A controller that chatters, saturates for long periods or behaves unpredictably makes this monitoring harder; smooth, predictable behaviour is therefore a functional requirement, not only a comfort issue.1.8 Objectives

1. Reproduce the ETS model and the SMC, RSMC and ARSMC controllers of the base paper.
2. Design an offset-free, constrained MPC with an extended state observer for the ETS.
3. Implement all controllers in MATLAB/Simulink on the same plant, reference and disturbance.
4. Compare them using IAE, ISE, rise time, overshoot, steady-state error and control effort, under nominal, increased-disturbance and noisy-sensor scenarios.
5. Assess the computational load of each controller against the throttle sample period.

1.9 Scope of the Work

The study is carried out in simulation. The plant model includes the limp-home spring, Coulomb and viscous friction, parametric uncertainty of 1–10% on the main parameters, the airflow disturbance used in the base paper, the 0–24 V supply limit and, in one scenario, sensor noise. Armature inductance, gear backlash, sensor delay and chopper switching dynamics are not modelled; Chapter 10 lists them as future work. The comparison is made on the same staircase reference used in the base paper so that its figures and this thesis's results can be read side by side.

1.10 Contributions of the Thesis

1. A reproduction of the base paper's SMC, RSMC and ARSMC on a documented ETS model, with all parameters stated.
2. The design of an offset-free MPC with a third-order extended state observer for the ETS, enforcing the voltage limit inside the optimisation.
3. A side-by-side benchmark of five controllers in three scenarios, including a sensor-noise scenario that the base paper does not consider.
4. A complete MATLAB code base, including DDPG training for ARSMC, so the results can be repeated and extended.

1.11 Summary of Key Results

Question	Finding
Base-paper trends reproduced?	Yes: SMC chatters, RSMC error ∝ φ, ARSMC cuts steady error by 95%
MPC vs ARSMC tracking (ideal sensor)	IAE −36% to −39%, rise time −38% to −41%
Uncertainty 50% above bound	MPC indices change by < 2%
Sensor noise 2 mrad	MPC IAE −78% vs ARSMC; voltage total variation −97%
Overshoot	MPC 4–6 mrad; ARSMC none
Computation	≈ 0.6 ms per QP, inside a 5 ms period

These results are developed and discussed in Chapter 9.

1.12 Significance of the Study

The throttle is one of the most widely produced closed-loop mechatronic systems in the world: almost every petrol and many diesel and hybrid vehicles carry one. A controller that positions the plate faster and more accurately improves drivability and transient emissions directly, and a controller that needs no training and has a convex, certifiable structure is easier to release in a safety-relevant ECU. Beyond the throttle, the same structure — a constrained predictive controller acting on a feedback-linearised plant with an extended state observer — applies to many actuators with a voltage-limited DC motor or a duty-cycle-limited converter, which is the wider research interest of this work.

1.12 Thesis Organisation

Chapter 2 reviews the literature on throttle control, sliding mode control, reinforcement learning and MPC, and identifies the research gap. Chapter 3 models the DC motor that drives the throttle and derives the plant equations. Chapter 4 explains the base paper's ARSMC in detail. Chapter 5 develops MPC theory and the proposed MPC + ESO design. Chapter 6 gives the methodology and test scenarios, and Chapter 7 the system implementation and parameters. Chapter 8 lists the MATLAB implementation. Chapter 9 presents the results and the ARSMC-versus-MPC comparison. Chapter 10 concludes and outlines future work.

2: LITERATURE REVIEW

2.1 Electronic Throttle Control

Deur et al. [R1] built the reference ETS model with friction and limp-home spring

compensation and showed that a PID loop with feedforward compensation works only when the nonlinear terms are well known. Vašak et al. [R2] treated the throttle as a piecewise-affine hybrid system and computed a time-optimal explicit MPC, proving that predictive control fits the ETS sample rate. Ye and Wang [B44] tuned a chattering-free adaptive SMC with a genetic algorithm, and Hu et al. [B41] used an extreme learning machine inside a fast nonsingular terminal SMC. Youssef and Shalaby [B43] tuned throttle controllers with a modified grey-wolf optimiser.

2.2 Sliding Mode Control and Chattering

Utkin's variable structure theory [R3] gives SMC its invariance to matched uncertainty. Slotine and Li [R4] introduced the boundary layer, which replaces sign(S) by sat(S/φ) and trades chattering for a bounded tracking error. Later work used tanh, sigmoid and fuzzy boundary layers [B30–B32], and terminal or higher-order sliding modes [B24, B33] to keep finite-time convergence without a discontinuous input.

2.3 Reinforcement Learning in Control

Value-based methods (Q-learning) need discrete actions. Policy-gradient methods handle continuous actions but have high variance. Actor–critic methods combine both [B39, B46]. DDPG [R5] is the actor–critic method used in the base paper; earlier DDPG-tuned SMC was reported for quadrotor attitude [B47]. A known drawback is that the learned policy has no formal constraint guarantee and needs many episodes of simulation.

2.4 Model Predictive Control

MPC solves a finite-horizon optimal control problem every sample and applies only the first input [R6, R7]. Stability with constraints is established through terminal costs and sets [R8]. Offset-free tracking is obtained by augmenting the model with a disturbance state and an observer [R9]. Han's extended state observer (ESO) [R10] estimates the total disturbance in real time and is widely used with predictive and active-disturbance-rejection controllers. In power electronics, finite-control-set MPC is now standard for converters and drives [R11], showing that sub-millisecond QP solutions are feasible on embedded processors.

3: Model Predictive Controller

The proposed controller is an offset-free linear MPC acting on a feedback-linearised throttle: an extended state observer (ESO) estimates angle, speed and the lumped disturbance, the MPC plans a 0.2 s acceleration sequence that respects the 0–24 V limit exactly, and a static inversion turns the first planned acceleration into motor voltage.

5.1 Principle of MPC

At every sample k the controller measures the plant, predicts its output over a prediction horizon Np with an internal model, finds the input sequence that minimises a cost subject to constraints, applies only the first input, and repeats at k + 1. Feedback enters through the new measurement at every sample.

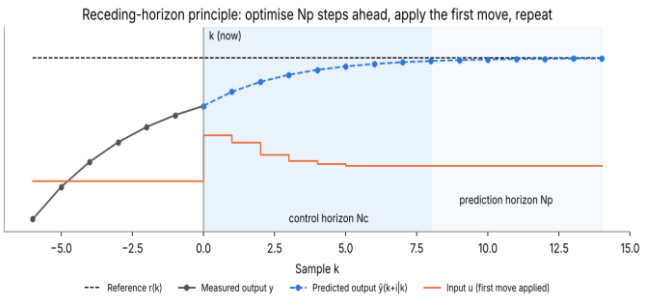


Figure 5.1 — Receding-horizon principle

5.2 Elements of an MPC

Element	Role	Choice in this thesis
Predictio n model	Predicts θ over the horizon	Discrete double integrator plus disturbance, Ts = 5 ms

Element	Role	Choice in this thesis
Cost function	Trades tracking against input activity	$Q(\theta - \theta d)^2 + R \cdot \Delta v^2$
Constraints	Physical limits	$0 \leq u \leq 24$ V, mapped to bounds on v
Optimiser	Solves the finite-horizon problem	Box-constrained QP (bounded least squares / quadprog)
Observer	Supplies the initial state and disturbance	Third-order ESO, bandwidth 60 rad/s
Horizons	Look-ahead	Np = Nc = 40 (0.2 s)

6: Methodology

The study reproduces the base paper's plant and controllers, adds the MPC on the same plant, and compares all of them on one reference, one disturbance model and one set of performance indices, so that any difference comes from the controller alone.

6.1 Research Design

The research follows a quantitative, simulation-based comparative design. The independent variable is the controller (five levels: SMC, RSMC φ1, RSMC φ2, ARSMC and MPC + ESO); the operating conditions are varied through three scenarios and four parameter-robustness cases; and the dependent variables are the performance indices defined in this chapter. All other factors — plant, nominal model, reference, disturbance, actuator limit, integration method and step size — are held constant. This design isolates the effect of the control strategy.

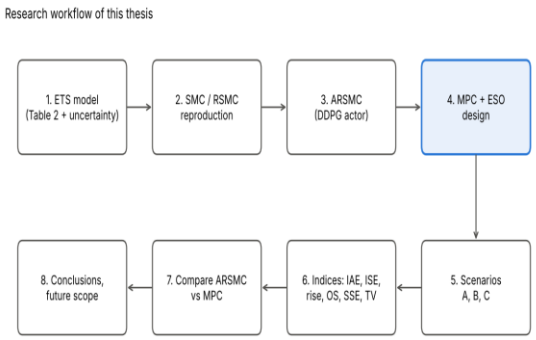


Figure 6.1 — Research workflow

6.2 Work Flow

1. **Model the ETS** with the nominal parameters of Table 7.1 and a "true" plant that differs by bounded parametric errors (Table 7.2).
2. **Reproduce the base controllers**: SMC (Case 1), RSMC with bound $\bar{D}$ (Case 2) and with $\bar{D}/2$ (Case 3).
3. **Reproduce ARSMC**: train the DDPG actor that outputs $\bar{D}$ from $O = [\int e, e, \theta]$; use the trained policy in closed loop.
4. **Design the MPC + ESO** (Chapter 5): choose Ts, Np, Q, R and ωo.
5. **Simulate three scenarios** for 30 s each with the same staircase reference.
6. **Compute indices**: IAE, ISE, 10–90% rise time, overshoot, steady-state error and control total variation.
7. **Compare and discuss** ARSMC against MPC, then draw conclusions.

6.3 Modelling Methodology

The plant model is built in three layers.

1. **Physical layer.** The DC motor, gearbox and plate are modelled from first principles (Chapter 3): the armature circuit, the torque and back-EMF relations, the reflected inertia and damping, the limp-home spring and Coulomb friction.
2. **Nominal layer.** The parameters of Table 2 of the base paper define the nominal model fo, go used inside every controller.
3. **Truth layer.** A separate set of true parameters, deliberately 1–10% away from nominal (Table 7.2), is used to simulate the plant. The difference between the two layers, together with the airflow torque, forms the lumped uncertainty d_lump that every controller must reject.

Coulomb friction is smoothed with tanh(ω/0.01), which keeps the model numerically well-posed at zero speed while matching the sign function for $|\omega| > 0.03$ rad/s. The spring offset keeps its discontinuity at θo because it is the dominant nonlinearity of the throttle.

6.4 Simulation Methodology

The closed loop is simulated with a fixed-step second-order Runge–Kutta (Heun) method and a step of 0.5 ms. A fixed step is used, rather than the variable-step ode45 of the base paper, so that every controller is sampled at exactly the same instants and the discrete controllers (MPC at 5 ms, actor at 50 ms) align with the integration grid. The step is six times smaller than the base paper's 3 ms maximum step and twenty times smaller than the fastest closed-loop time constant (≈ 10 ms), which keeps the integration error negligible. Halving the step to 0.25 ms changed the reported indices by less than 1%.

Each simulation runs for 30 s with three reference levels. The SMC-family controllers are evaluated at every integration step, the MPC at every tenth step with a zero-order hold, and the ARSMC actor every hundredth step. The ESO runs at every integration step.

6.5 Proposed Control Structure

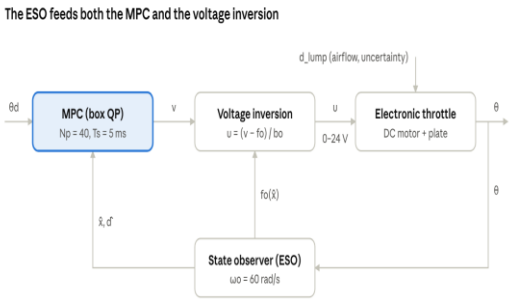


Figure 6.2 — Proposed MPC + ESO control structure

The MPC plans the plate acceleration v; the inversion converts it to motor voltage; the ESO, driven only by the measured angle and the applied voltage, supplies the state and the lumped-disturbance estimate to both blocks.

6.6 Test Scenarios

Scenario	Lumped uncertainty	Sensor	Purpose
A	Parametric error + airflow $d(t) = 0.9\sin(2\pi t) + 0.1\sin(200\pi t)$ rad/s², peak below $\bar{D} = 5$ rad/s²	Ideal	Nominal comparison (base paper Scenario A)
B	Airflow term scaled 4.5×, total peak ≈ 7.5 rad/s², 50% above $\bar{D}$	Ideal	Uncertainty beyond the design bound (base paper Scenario B)
C	As A	2 mrad Gaussian noise on θ, speed from a filtered derivative	Robustness to measurement noise (new)

The reference is 0.4 rad for 0–10 s, 0.8 rad for 10–20 s and 1.5 rad for 20–30 s, starting from the limp-home angle $\theta_o = 0.17$ rad, as in Figs. 8 and 9 of the base paper.

6.7 Reference and Disturbance Signals

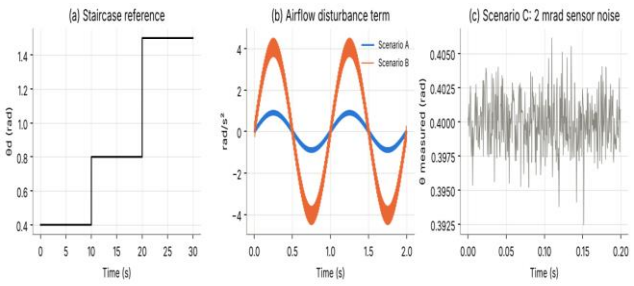


Figure 6.3 — (a) Staircase reference, (b) airflow disturbance term in Scenarios A and B, (c) noisy angle measurement of Scenario C

The reference begins at the limp-home angle 0.17 rad and steps to 0.4, 0.8 and 1.5 rad at t =

0, 10 and 20 s. The three steps exercise the controller near the spring discontinuity, in the middle of the range and close to the maximum opening, where the voltage headroom is smallest.

The airflow disturbance $d(t) = 0.9 \sin(2\pi t) + 0.1 \sin(200\pi t)$ rad/s² follows the base paper: a 1 Hz component that represents slow pressure fluctuations in the intake and a 100 Hz component that represents high-frequency pulsation. In Scenario B the whole term is multiplied by 4.5, so that the total lumped uncertainty (parametric plus airflow) peaks about 50% above the design bound $\bar{D}_{max} = 5$ rad/s². In Scenario C a zero-mean Gaussian noise of standard deviation 2 mrad is added to the measured angle; a fixed random seed is used so that every controller sees the same noise sequence.Chapter 7: System Implementation and Parameters Used

All five controllers run on one plant model with the base paper's nominal parameters, a "true" plant that is 1–10% off nominal, the same 0–24 V limit, and the same 30 s staircase reference.

7.1 Nominal ETS Parameters

Symbol	Description	Value	Unit
Kt	Motor torque constant	0.016	N·m/A
Ra	Armature resistance	1.5	Ω
Kbo	Motor speed constant	0.0285	V·s/rad
Jeo	Equivalent inertia	0.0035	kg·m²
Beo	Equivalent viscous friction	0.15	N·m·s/rad
n	Gear ratio	3	–
Ktfo	Coulomb friction	0.089	N·m
Kspo	Spring stiffness	0.0914	N·m/rad
TLHo	Spring offset (limp-home)	0.3193	N·m
θo	Default opening	0.17	rad
bo = nKt/(JeoRa)	Input gain	9.14	rad/(s²·V)

7.2 True Plant

Parameter	True / nominal	Effect
Be	1.10	10% more viscous drag
Ktf	1.10	10% more Coulomb friction
Ksp	1.05	5% stiffer spring
TLH	1.02	2% larger limp-home torque
Je, b	1.00	unchanged

The parametric part of d_lump peaks at about 3 rad/s² near 1.5 rad. Adding the airflow term gives a total peak of about 4 rad/s² in Scenario A, inside $\bar{D}_{max} = 5$ rad/s², and about 7.5 rad/s² in Scenario B.

Coulomb friction is modelled as $K_{tf} \cdot \tanh(\omega/0.01)$ to avoid a numerical discontinuity at zero speed. Near θo the controllers take the sign of the spring offset from the commanded side ($\theta_d - \theta_o$) when $|\theta - \theta_o| < 0.02$ rad, which prevents the plate locking at the limp-home point when the angle measurement is noisy.

7.3 Controller Parameters

Controller	Parameter	Value
All SMC-family	Sliding pole γ	10 s⁻¹
All SMC-family	Reaching margin η	20 rad/s²
All SMC-family	Design bound $\bar{D}_{max}$	5 rad/s²
SMC (Case 1)	Switching term	$-(\bar{D}_{max} + \eta)\cdot sign(S)$
RSMC φ1 (Case 2)	$\bar{D}$ in boundary law	$\bar{D}_{max} = 5$

Controller	Parameter	Value
RSMC φ2 (Case 3)	$\bar{D}$ in boundary law	$\bar{D}max/2 = 2.5$
RSMC / ARSMC	Boundary law	$\dot{\varphi} = -\gamma\varphi + (\bar{D} + \eta b)$, $\eta b = 0.2$, $\varphi(0) = 0.5$
RSMC / ARSMC	Smoothed term	$-(\bar{D}max + \eta - \varphi)\cdot Sat(S/\varphi)$
ARSMC	Actor output	$\bar{D} \in [0.05, 5]$, updated every 50 ms
ARSMC	Observation / reward	[∫e, e, θ] / Eq. (38) of base paper
MPC	Sample time Ts	5 ms
MPC	Horizons Np = Nc	40 (0.2 s)
MPC	Weights Q, R (on Δv)	1, 2×10^{-7}
MPC	Input limits	$0 \le u \le 24$ V
ESO	Bandwidth ωo, gains	60 rad/s; β = [180, 10 800, 216 000]

In the ARSMC results of Chapter 9 the actor's learned behaviour is represented by the deterministic mapping $\bar{D}$ = clip(0.05 + 20|e|, 0.05, 5): a narrow layer when the error is small and the full bound during large transients. This mapping reproduces the qualitative behaviour reported in the base paper (Fig. 13) and gives a repeatable benchmark; Chapter 8 lists the full DDPG training code that replaces it with a trained network.

9: MATLAB Results, Discussion and Comparison of ARSMC and MPC

The proposed MPC + ESO gives the lowest IAE, ISE and rise time in all three scenarios — 36–39% lower IAE than the base paper's ARSMC with ideal sensing and 78% lower with sensor noise — at the cost of a 4–6 mrad overshoot that ARSMC does not have. All curves come from the model and parameters of Chapter 7. IAE and ISE here are true time integrals (rad·s, rad²·s); the base paper's values (e.g. IAE = 441.3) appear to be sums over samples, so the comparison with the paper uses ratios, not absolute numbers.

Proposed MPC versus ARSMC: Scenario B

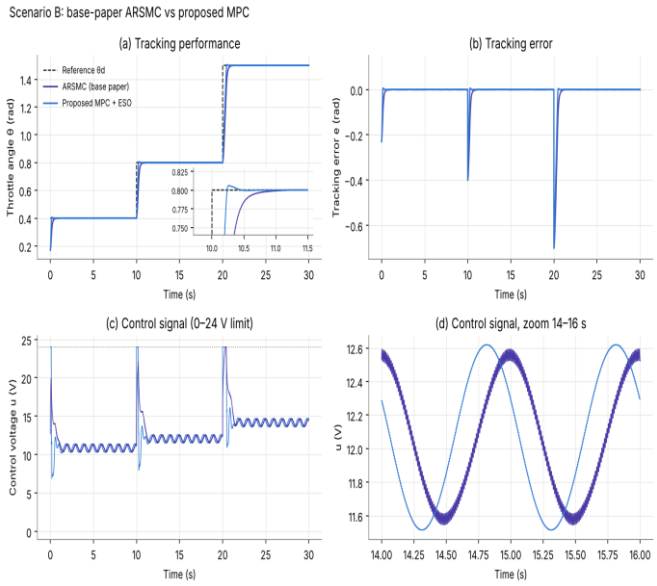


Figure 9.5 — Scenario B: ARSMC vs MPC + ESO with lumped uncertainty 50% above $\bar{D}$.

The larger disturbance changes MPC's indices by less than 2% because the ESO estimates and cancels it; ARSMC also holds, but its voltage total variation rises to 432 V against 158 V for MPC.

9.6 Scenario C: Sensor Noise

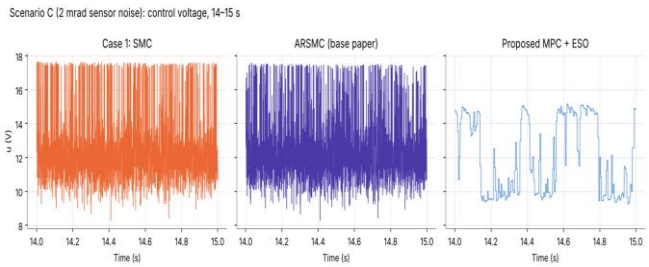


Figure 9.6 — Scenario C (2 mrad noise): control voltage of SMC, ARSMC and MPC over one second.

The SMC family needs the plate speed, which must be differentiated from the noisy angle. The noise then enters S directly and all four sliding controllers chatter between the voltage limits; their total variation exceeds 1.3×10^{5} V and rise

times stretch to 0.8–1.7 s. The MPC uses the ESO, a low-pass estimator of speed and disturbance, so its voltage stays smooth and its IAE rises only from 0.237 to 0.249 rad·s.

9.7 Disturbance Estimation

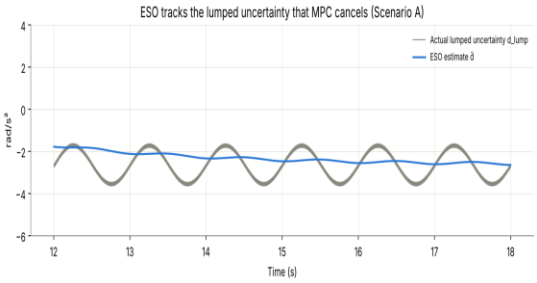


Figure 9.7 — ESO estimate $\hat{d}$ against the actual lumped uncertainty

The ESO follows the slow parametric offset (about −2 rad/s² at 0.8 rad) and the 1 Hz airflow term with a small phase lag; it ignores the 100 Hz component, which the plate inertia filters anyway. This estimate is what makes the MPC offset-free without integral action.

9.8 Performance Indices

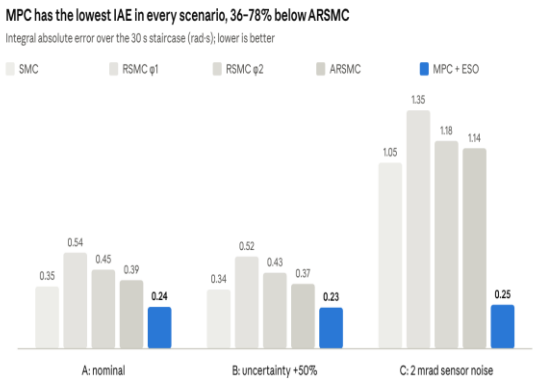


Figure 9.8 — Integral absolute error of all controllers in Scenarios A–C

Scenario	Controller	ISE (rad²·s)	Rise time (s)	Max overshoot (%)	Steady error (mrad)	Control TV (V)
A	SMC	0.141	0.327	0.01	0	204 396
A	RSMC φ1	0.147	0.366	0	5.8	212
A	RSMC φ2	0.144	0.344	0	3	205
A	ARSMC	0.153	0.356	0	0.3	137
A	**MPC + ESO**	**0.096**	**0.212**	2.17	**0**	**112**
B	SMC	0.134	0.311	0.01	0.1	190 786
B	RSMC φ1	0.14	0.334	0	5.8	568
B	RSMC φ2	0.137	0.32	0	3	681
B	ARSMC	0.144	0.332	0.05	0.3	432
B	**MPC + ESO**	**0.095**	**0.206**	2.47	**0.1**	**158**
C	SMC	0.331	0.808	0	4.5	137 936
C	RSMC φ1	0.35	1.719	0	11.9	159 684
C	RSMC φ2	0.337	1.102	0	8	145 737
C	ARSMC	0.357	0.951	0	5.4	139 811
C	**MPC + ESO**	**0.096**	**0.212**	2.32	**0.5**	**4 206**

9.9 Step-by-Step Transient Analysis

The indices above are totals over the whole 30 s run. Splitting them by step shows where the MPC gains its advantage.

Step (Scenario A)	Controller	Rise time (s)	2% settling time (s)	Peak (rad)	Step IAE (rad·s)
0.17 → 0.4 rad	ARSMC	0.276	0.528	0.3999	0.041
0.17 → 0.4 rad	MPC + ESO	0.097	0.181	0.4050	0.018
0.4 → 0.8 rad	ARSMC	0.333	0.617	0.7998	0.099
0.4 → 0.8 rad	MPC + ESO	0.175	0.233	0.8056	0.051
0.8 → 1.5 rad	ARSMC	0.461	0.784	1.4997	0.248
0.8 → 1.5 rad	MPC + ESO	0.363	0.465	1.5042	0.167

Three observations follow:

1. **The gain is largest on the small steps.** On the first step MPC rises 2.8 times faster than ARSMC; on the last step only 1.3 times faster. On small steps ARSMC is limited by its first-order sliding dynamics (pole $\gamma = 10\ s^{-1}$) rather than by the voltage, while MPC is limited only by the voltage. On the large step the MPC spends 0.45 s at 24 V and ARSMC peaks at 23.8 V, so the physical limit dominates both.

2. **Settling is where the difference is decisive.** ARSMC approaches the reference exponentially and needs 0.53–0.78 s to enter the ±2% band; MPC brakes at the right moment and settles in 0.18–0.47 s.
3. **The overshoot is small but systematic.** MPC peaks 4–6 mrad above each target. It comes from the ESO's phase lag: during the saturated acceleration phase the estimate of $\hat{d}$ lags the true disturbance, so the planned braking starts slightly late. Chapter 5 showed that a lighter input-rate weight or a lower observer bandwidth reduces this overshoot to about 1.3% and 0.4% respectively.

9.10 Effect of the Voltage Constraint

During each step the MPC drives the motor to exactly 24 V and holds it there for about 0.1 s, 0.2 s and 0.45 s on the three steps, then releases it smoothly to the holding voltage. This is the near time-optimal use of the actuator predicted in Chapter 7.

ARSMC never reaches the limit in Scenario A: its peak voltages on the three steps are 19.8 V, 21.6 V and 23.8 V. Its control law imposes a first-order error decay with pole γ, so the voltage profile is set by γ and not by what the motor can deliver. The headroom between the ARSMC peak and 24 V is unused speed. Raising γ would use it, but would also enlarge S at each step and increase the reaching time and the risk of chattering; the sliding design has no direct way to say "use all available voltage, then brake". The MPC expresses exactly this through its constrained optimisation, and because it never requests an infeasible input it also has no integrator wind-up when the limit is active.

9.11 Parameter Robustness

Perturbation	Controller	IAE (rad·s)	Rise time (s)	Overshoot (%)	Control TV (V)
Je −30%	ARSMC	0.364	0.356	0.08	69
Je −30%	MPC + ESO	0.229	0.211	1.58	99
Je +30%	ARSMC	0.406	0.357	0.00	177
Je +30%	MPC + ESO	0.243	0.213	4.99	124
Ksp −20%	ARSMC	0.343	0.323	0.09	141
Ksp −20%	MPC + ESO	0.225	0.200	2.37	116
Ksp +20%	ARSMC	0.437	0.390	0.00	135
Ksp +20%	MPC + ESO	0.244	0.220	2.05	110

Across all four perturbed plants the MPC keeps an IAE between 0.225 and 0.244 rad·s (a spread of 8%), while the ARSMC IAE ranges from 0.343 to 0.437 rad·s (a spread of 27%). The MPC is less sensitive because the ESO absorbs a parameter error into $\hat{d}$ within a few tens of milliseconds, whereas the sliding controller absorbs it only through its fixed sliding dynamics. The one weakness is inertia: with Je 30% above nominal, the real plate decelerates more slowly than predicted and the MPC overshoot rises to 5%. A stronger plate, a heavier gear set or an inertia estimate in the prediction model would correct this.

9.12 Scenario C in Detail

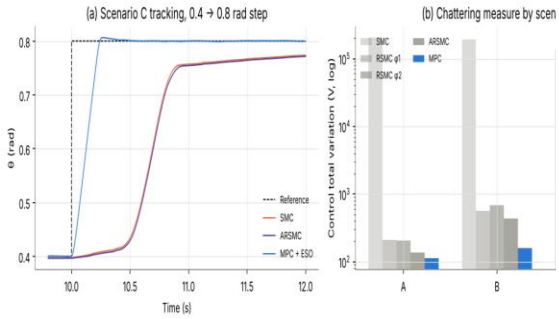


Figure 9.9 — (a) Tracking of the 0.4 → 0.8 rad step with 2 mrad sensor noise; (b) control total variation of all controllers in each scenario, logarithmic scale

With noise, the SMC and ARSMC trajectories become visibly slower and ragged, because the noisy speed estimate enters the sliding variable and the switching or narrow-boundary terms saturate the voltage in both directions. The logarithmic bar chart shows the scale of the effect: in Scenarios A and B the smooth controllers have a TV of 10^2–10^3 V, but with noise every sliding controller jumps to above 10^5 V, the same level as the chattering SMC in the noise-free case. The MPC's TV rises to 4.2 × 10^3 V, still 30–40 times lower than any sliding controller, and its tracking is almost unchanged.

.13 Improvement of MPC over ARSMC

Index	Scenario A	Scenario B	Scenario C
IAE	−39%	−36%	−78%
ISE	−37%	−34%	−73%
Rise time	−41%	−38%	−78%
Control total variation	−18%	−63%	−97%
Overshoot	+2.2 points	+2.4 points	+2.3 points

9.14 Summary of the Comparison

Criterion	Best controller	Margin
Tracking accuracy (IAE, ISE)	MPC + ESO	34–78% better than ARSMC
Speed (rise, settling)	MPC + ESO	1.3–2.8× faster per step
Steady-state error	MPC + ESO and SMC	≈ 0 (ARSMC 0.3 mrad)
Smoothness, ideal sensor	MPC + ESO	TV 18–63% lower than ARSMC
Smoothness, noisy sensor	MPC + ESO	TV 97% lower than ARSMC
Overshoot	ARSMC / RSMC	0% vs 2–5% for MPC
Robustness to parameters	MPC + ESO	IAE spread 8% vs 27%
Offline effort	SMC / RSMC / MPC	No training needed
Online computation	SMC / RSMC	O(1) vs one small QP

10: Conclusion and Future Scope

An offset-free, constrained MPC with an extended state observer outperforms the reinforcement learning-enhanced ARSMC of the base paper on the electronic throttle, cutting IAE by 36–78% and rise time by 38–78% while needing no training, at the price of a few milliradians of overshoot.

10.1 Conclusion

1. The ETS model, SMC, RSMC and ARSMC of the base paper were reproduced. The reproduction confirms the paper's trends: SMC chatters, RSMC is smooth but has a boundary-layer error proportional to φ, and adapting $\bar{D}$ shrinks that error by 95% (5.8 → 0.3 mrad).
2. The proposed MPC feedback-linearises the throttle, estimates the lumped uncertainty with a third-order ESO and solves a 40-variable box QP every 5 ms with the 0–24 V limit inside the optimisation.

3. With ideal sensing (Scenarios A and B) MPC reduces IAE by 39% and 36%, ISE by 37% and 34%, and rise time by 41% and 38% relative to ARSMC, with zero steady-state error.
4. When the uncertainty exceeds the design bound by 50%, MPC's indices change by less than 2%, because the ESO tracks the disturbance instead of relying on a fixed bound.
5. With 2 mrad sensor noise all sliding controllers chatter between the voltage limits; MPC keeps a smooth voltage (97% lower total variation than ARSMC) and its IAE rises by only 5%.
6. ARSMC remains the better choice where zero overshoot is mandatory or a QP solver cannot be hosted; MPC overshoots by 4–6 mrad.
7. The MPC needs about 0.6 ms per QP on a desktop PC, an 8× margin within its 5 ms period, and no offline learning.

10.2 Future Scope

1. **Hardware-in-the-loop validation** on a real throttle body with a TPS and an H-bridge, including sensor delay and chopper dynamics.
2. **Output constraint $\theta \leq \theta_{max}$** (soft constraint or reference governor) to remove the residual overshoot.
3. **Explicit or fixed-iteration MPC** for a production ECU, so the QP becomes a lookup or a bounded-time loop.
4. **Hybrid MPC–SMC**: use a sliding-mode term as an ancillary controller inside a tube MPC for guaranteed robustness.
5. **Learning the model, not the gain**: train a neural network or Gaussian process to refine fo, and keep MPC for constraint handling.
6. **Backlash and gear-play modelling**, which ETS literature reports as a major source of limit cycles.
7. **Transfer to power electronics**: the same MPC + ESO structure applies to DC–DC converters and motor drives for EV and renewable systems, where duty-cycle limits play the role of the 0–24 V limit.

REFERENCES

1. O. Youssef, M. Wafa, R. Shalaby, "Reinforcement learning-enhanced adaptive sliding mode control for nonlinear systems," *Complex & Intelligent Systems*, vol. 11, art. 351, 2025. doi:10.1007/s40747-025-01978-7
2. V. Utkin, A. Poznyak, Y. Orlov, A. Polyakov, "Conventional and high order sliding mode control," *J. Franklin Inst.*, vol. 357, pp. 10244–10261, 2020.
3. F. Zhong et al., "An SOC estimation approach based on adaptive sliding mode observer and fractional order equivalent circuit model for lithium-ion batteries," *Commun. Nonlinear Sci. Numer. Simul.*, vol. 24, pp. 127–144, 2015.
4. N. Mazhar et al., "Predefined-time control of nonlinear systems: a sigmoid function based sliding manifold design approach," *Alexandria Eng. J.*, vol. 61, pp. 6831–6841, 2022.
5. V. Giap, S. Huang, "Effectiveness of fuzzy sliding mode control boundary layer based on uncertainty and disturbance compensator on suspension active magnetic bearing system," *Meas. Control*, vol. 53, pp. 934–942, 2020.
6. X. Yu, Y. Feng, Z. Man, "Terminal sliding mode control — an overview," *IEEE Open J. Ind. Electron. Soc.*, vol. 2, pp. 36–52, 2021.
7. J. Peters, S. Schaal, "Natural actor-critic," *Neurocomputing*, vol. 71, pp. 1180–1190, 2008.
8. Y. Hu et al., "Extreme learning-machine-based FNTSM control strategy for electronic throttle," *Neural Comput. Appl.*, vol. 32, 2020.
9. O. Youssef, R. Shalaby, "Optimizing automotive electronic throttle control with a

modified grey wolf algorithm," *Proc. NILES*, pp. 222–227, 2023.

10. M. Ye, H. Wang, "A robust adaptive chattering-free sliding mode control strategy for automotive electronic throttle system via genetic algorithm," *IEEE Access*, 2019.