

CRR-TMPC: Certificate-Triggered Reset and Reacquisition for Tube MPC under Abrupt Actuation-Effectiveness Changes

Dhayaa Khudher

`dhayaa.khudher@uobasrah.edu.iq`

University of Basrah

Research Article

Keywords: tube model predictive control, set-membership identification, certificate invalidation, regime change, fault-tolerant control, mobile robot

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11171736/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

CRR-TMPC: Certificate-Triggered Reset and Reacquisition for Tube MPC under Abrupt Actuation-Effectiveness Changes

Abstract

Set-membership tube model predictive control (MPC) normally contracts an uncertainty set under a common-parameter premise. An abrupt actuation-effectiveness change can invalidate that premise while the tightened optimization remains feasible. We introduce Certificate-Triggered Reset and Reacquisition Tube MPC (CRR-TMPC), a two-mode protocol that interprets an empty completed-transition consistency intersection as rejection of the active local certificate. It resets independently to a declared global envelope, receives the measured state with a finite-horizon global tube controller, reconstructs a post-reset candidate, and requires three consecutive candidate-specific robust-closure passes before returning to local operation. The closure test exactly maximizes the disturbance P -norm over box vertices and bounds contraction jointly over candidate and reference-rate vertices. A frozen campaign comprises 1,800 controller episodes: 60 paired realizations in five scenario families and six controllers. In the three in-scope step-change families, CRR-TMPC detected 297 of 300 changes, reduced invalid-LOCAL exposure from 14,550 samples for intersection-only adaptation to 865, and reduced covered successor-envelope exceedances from 11,131 to 103. The no-independent-reset ablation produced 3,958 exceedances. Local operation reduced median terminal tube radius by 26.6% and median input tightening by 23.4% relative to the fixed-global construction. All episodes completed without physical-constraint breach. The fixed-global controller produced only 17 covered exceedances and slightly lower mean tracking error, so no safety or performance-superiority claim is made. Gradual drift missed 50 of 60 onset windows. Under the tested abrupt-change regimes, certificate rejection and reacquisition reduced stale-set exposure and recovered local authority; the experiments do not establish nonlinear containment, gradual-drift coverage, or cross-mode guarantees.

Keywords: tube model predictive control, set-membership identification, certificate invalidation, regime change, fault-tolerant control, mobile robot

1 Introduction

Tube model predictive control (MPC) converts bounded model error into tightened nominal constraints. Its interpretation depends on a simple premise: the realized plant must remain inside the uncertainty set used to propagate the tube. Terrain, payload, actuator-loss, or contact-regime changes can violate that premise while leaving the numerical optimization feasible. Numerical feasibility is then not evidence that the active tube still covers the plant.

Fixed-global robust MPC avoids local-set invalidation if its envelope is sound. Adaptive robust MPC instead updates parameter sets and controller ingredients online [1–5]. Most contraction arguments retain the true plant because a

constant or rate-bounded parameter is assumed. For an unannounced jump outside an active local set, monotone intersection cannot restore that lost containment. Recovery needs both a non-monotone reset and a controller able to receive the state reached when the contradiction becomes observable.

We consider the case in which a completed transition falsifies the active local parameter enclosure and develop CRR-TMPC as the response. The contributions are:

1. an auditable command-boundary protocol for abrupt certificate loss, in which contradiction of the active local set triggers rejection of that tube certificate, independent reset to the declared global envelope, and receipt of the

- reached state without allowing the rejected local set to constrain reacquisition;
2. a post-reset reacquisition and re-entry mechanism that constructs a fresh candidate from post-change evidence and permits local tightening only after candidate-specific robust closure; the implemented three-pass persistence rule is treated as a conservative design choice rather than a standalone novelty claim;
 3. conditional finite-horizon and fixed-candidate results that explicitly exclude the first unannounced change interval and arbitrary switching; and
 4. a frozen six-controller, 1,800-episode confirmatory study with paired ablations, detector operating characteristics, stratification, complete failure accounting, and an out-of-scope gradual-drift stress test.

The claim is mechanism-specific. CRR-TMPC is not asserted to outperform a feasible fixed-global tube in tracking or containment.

Set-membership inconsistency tests, feasible-set reinitialization, recovery modes, and uncertainty reduction are established ideas. CRR-TMPC does not claim any of those operations in isolation. Its contribution is their command-boundary integration with tube-certificate semantics: the contradicted local enclosure is rejected, the reset and subsequent candidate are not constrained by that stale enclosure, the reached state is handled under the declared global tube, and local tightening resumes only after the reacquired candidate satisfies the closure gate. B2 removes reset, B4 retains the stale candidate during transition, and B5 removes persistence.

2 Related work and contribution boundary

2.1 Adaptive tubes and set-membership MPC

Robust tubes can be nonlinear, multistage, configuration-constrained, time-varying, or system-level [6–9]. Adaptive variants combine set-membership identification with model, tube, gain, tightening, excitation, or terminal redesign [1–5, 10, 11]. Their central benefit is reduced conservatism while the identification set remains a valid outer approximation. The learning-based

MPC review by Hewing et al. [12] likewise emphasizes that safety statements inherit the assumptions behind learned uncertainty bounds.

An empty consistency set is not itself a new fault or change indicator. In set-membership fault diagnosis, inconsistency between bounded-error data and the assumed model has long been used to invalidate the current model set; Puig [13] also describes resetting the feasible parameter set after fault indication so that post-fault values can be identified. CRR-TMPC therefore does not claim novelty for contradiction detection or feasible-set reset in isolation. The distinction pursued here is how such a contradiction is propagated into the semantics and execution of an adaptive tube controller, including stale-set-independent reacquisition and certified return to local tightening.

2.2 Time-varying uncertainty and recovery

Tanaskovic et al. [14] construct feasible parameter sets for linear time-varying MIMO systems under stated variation assumptions. Bujarbaruah et al. [15] treat time-varying additive uncertainty with a known domain and maximum rate of change. LPV scheduling [16], stochastic switching MPC [17], and data-driven chance-constrained tubes [18] provide alternative structures. Peschke and Görges [19] are the closest recovery-oriented adaptive-tube reference: their controller accommodates varying quantities under its declared bounds and updates tracking ingredients. CRR-TMPC does not supersede those rate-bounded formulations. It targets an abrupt jump model, explicitly logs rejection of the current local certificate, and prevents the rejected set from constraining reacquisition.

Fault-tolerant tube MPC also separates detection and accommodation, often using an observer or fault estimator. Qin et al. [20], for example, combine set-membership estimation and tube MPC for robust fault-tolerant control and report hardware-in-the-loop validation. Daniels et al. [21] combine active set-membership fault identification, informative excitation, tube-based robust control, and a transition back toward nominal control as parameter confidence improves. These works make clear that neither set-membership fault identification nor return from a robust recovery regime is new by itself. CRR-TMPC

addresses a different failure event: contradiction of an already active local tube certificate while the optimization may remain numerically feasible. The response rejects that certificate, receives the reached state under the global construction, reacquires a candidate without inheriting the stale set, and gates local re-entry on candidate-specific closure.

2.3 Mobile-robot tracking and exact gap

Wheel slip and wheel–terrain variation motivate the validation plant; adaptive robust and predictive tracking approaches are reported in [22–27]. Those works primarily improve tracking under slip or interaction models. Here the coefficients are bounded actuation-effectiveness losses rather than an identified tire law. The gap addressed here is therefore not “MPC for slippery robots,” nor the generic use of set-membership inconsistency for fault detection. It is the command-boundary integration needed when an adaptive tube’s locally contracted certificate is contradicted by an abrupt admissible regime change: reject the stale certificate, recover under a broader receiving construction, reacquire independently, and re-enter local tightening only after candidate-specific certification.

Table 1 makes this boundary operational. CRR-TMPC is complementary to constant-parameter adaptation, rate-bounded time variation, and fault-identification-based tube control; it sacrifices a continuous-variation claim to make abrupt certificate invalidation, independent reacquisition, and certified local return explicit.

3 Plant, surrogate, and assumptions

3.1 Nonlinear validation plant

Let $e = (e_x, e_y, e_\psi)^\top$ be path-frame tracking error, $u = (u_v, u_\omega)^\top$ the correction to reference commands (v_r, ω_r) , and $s_i = 1 - \theta_i$. The validation dynamics are

$$\dot{e}_x = s_v(v_r + u_v) \cos e_\psi - v_r + \omega_r e_y, \quad (1)$$

$$\dot{e}_y = s_v(v_r + u_v) \sin e_\psi - \omega_r e_x, \quad (2)$$

$$\dot{e}_\psi = s_\omega(\omega_r + u_\omega) - \omega_r. \quad (3)$$

The uncertain coefficient is $\theta = (\theta_v, \theta_\omega)^\top \in \Theta_G = [0, 0.25]^2$. A clipped additive increment d_k is applied after numerical integration, with $|d_k| \leq (0.004, 0.006, 0.004)^\top$.

The plant is integrated by fourth-order Runge–Kutta at 0.02 s. Control is updated at $h = 0.1$ s. The circle reference uses $v_r = 0.8$ m/s and $\omega_r = 0.18$ rad/s; the S-curve uses $\omega_r = 0.26 \sin(0.55t)$ rad/s. Angle error is wrapped to $(-\pi, \pi]$.

3.2 Sampled affine surrogate

The controller uses

$$e_{k+1} = A_k e_k + B(\theta_k) u_k + c(\theta_k, \omega_{r,k}) + w_k, \quad (4)$$

where

$$A_k = \begin{bmatrix} 1 & h\omega_{r,k} & 0 \\ -h\omega_{r,k} & 1 & hv_r \\ 0 & 0 & 1 \end{bmatrix}, \quad (5)$$

$$B(\theta) = \text{diag}_{3 \times 2}(h(1 - \theta_v), 0, h(1 - \theta_\omega)), \quad (6)$$

$$c(\theta, \omega_r) = (-h\theta_v v_r, 0, -h\theta_\omega \omega_r)^\top. \quad (7)$$

The notation $\text{diag}_{3 \times 2}(a, 0, b)$ denotes nonzero entries $(1, 1) = a$ and $(3, 2) = b$. This first-order surrogate deliberately differs from the RK4 plant, so the operational containment audit tests the combined discretization, nonlinear, and disturbance residual rather than replaying the prediction model.

The completed-transition consistency set is evaluated on a 0.01 grid:

$$\mathcal{C}_k = \{\theta \in \Theta_G : |e_k - \hat{e}_k(\theta)| \leq (0.008, 0.020, 0.005)^\top\}. \quad (8)$$

The box hull of feasible grid points is retained. Empty feasibility invalidates the active local certificate. Quantization is part of the declared implementation and not presented as an exact continuous set-membership solution.

4 CRR-TMPC method

4.1 Tube sections and MPC

A common ancillary pair (K, P) was designed against the joint vertex set $\theta_v, \theta_\omega \in [0, 0.10]$ and

Table 1 Positioning against the closest uncertainty-update and fault-accommodation structures. “Contradiction” means that new bounded-error data have empty intersection with the active local parameter enclosure.

Approach	Declared change model	Response to local-set contradiction	Condition for returning to a local tube
Set-membership fault diagnosis [13]	Bounded-error model becomes inconsistent with measurements	Empty feasible set invalidates the assumed model; feasible parameter set may be reset for post-fault identification	Fault isolation/estimation from the reinitialized set; no tube-specific local re-entry rule
Set-membership adaptive MPC [1, 2]	Unknown parameter retained by a valid outer set	Not the principal modeled event; monotone updates preserve containment under assumptions	Updated valid set and standard adaptive-MPC feasibility ingredients
Time-varying adaptive MPC [14, 15]	Known domain with stated temporal-variation structure	Variation is propagated within the declared model	Feasible parameter set remains valid under rate/domain assumptions
Recovery-oriented adaptive tracking MPC [19]	Varying quantities under specified relative bounds	Recovery logic uses the method’s declared changing-plant assumptions	Method-specific recovery and updated tracking ingredients
Active fault identification + tube MPC [20, 21]	Unknown bounded faults estimated online	Fault/parameter information drives robust accommodation; active identification can accelerate uncertainty reduction	Return toward nominal/less conservative control as fault information improves
CRR-TMPC (this work)	Abrupt, piecewise-constant admissible effectiveness changes	Reject contradicted local tube certificate; reset independently to Θ_G and receive the reached state	Fresh post-reset candidate must satisfy width and candidate-specific closure; three-pass persistence is an implementation choice

$\omega_r \in [-0.26, 0.26]$. The implemented gain is

$$K = \begin{bmatrix} -2.02680 & 0.07314 & 0.03568 \\ -0.09077 & -5.08087 & -3.19124 \end{bmatrix}. \quad (9)$$

The associated positive-definite metric is

$$P = \begin{bmatrix} 53.56631 & 0.32756 & 0.02683 \\ 0.32756 & 228.68080 & 73.48377 \\ 0.02683 & 73.48377 & 42.38163 \end{bmatrix}. \quad (10)$$

Full-precision values are stored in the frozen implementation. Positive scaling of P leaves λ_Θ and the sign of the unbuffered closure margin invariant; the reported normalization of P is fixed when applying the 10^{-4} numerical buffer. The stage weights remain $Q = \text{diag}(3, 18, 7)$ and $R = \text{diag}(2, 1.5)$. For a parameter box $\Theta = [\underline{\theta}, \bar{\theta}]$, let its centre be $\bar{\theta}$, half-width be $\Delta\theta$, and $F = A + B(\bar{\theta})K$. Axis-aligned error radii are initialized by $r_0 = (0.003, 0.004, 0.003)^\top$ and propagated for $i = 0, \dots, N-1$ as

$$r_{i+1} = |F|r_i + q(\Theta, \omega_r), \quad (11)$$

with

$$q = \begin{bmatrix} h\Delta\theta_v(v_r + 0.9) + 0.006 \\ 0.012 \\ h\Delta\theta_\omega(|\omega_r| + 1.6) + 0.005 \end{bmatrix}. \quad (12)$$

For the ancillary tube analysis, define the deviation from the centred nominal trajectory as $\xi_i = e_i - z_i$ and the ancillary correction as $\delta u_i = u_i - v_i = K\xi_i$. Relative to the centred nominal surrogate, the deviation dynamics take the form

$$\xi_{i+1} = [A(\omega_r) + B(\theta)K]\xi_i + d_i, \quad (13)$$

where d_i collects the bounded parameter-mismatch, discretization, and process terms represented by the componentwise radius in $q(\Theta, \omega_r)$. At each receding-horizon command boundary $z_0 = e_k$, hence $\xi_0 = 0$, while the propagated tube sections bound future deviations. The nominal QP, horizon $N = 8$, minimizes

$$\sum_{i=0}^{N-1} (z_i^\top Q z_i + v_i^\top R v_i) + z_N^\top P z_N \quad (14)$$

subject to $z_0 = e_k$, the centred affine surrogate, and

$$|z_i| \leq x_{\max} - r_i, \quad |v_i| \leq u_{\max} - |K|r_i, \quad (15)$$

where $x_{\max} = (0.75, 0.60, 0.75)^\top$ and $u_{\max} = (0.90, 1.60)^\top$. Since $z_0 = e_k$, the first receding-horizon correction is $u_k = v_{0|k}$. A solve that is infeasible, erroneous, or later than 100 ms invokes $\text{clip}(Ke_k, \pm 0.65u_{\max})$ and is logged.

4.2 Invalidation, reset, and re-entry

LOCAL starts from $[0, 0.10]^2$ and replaces its box by $\Theta_k^L \cap \mathcal{C}_k$. If the intersection is empty, CRR-TMPC declares certificate invalidation, sets the outer controller box to Θ_G , enters TRANSITION, and initializes a candidate to Θ_G . Subsequent nonempty consistency boxes contract the candidate. During TRANSITION a fixed zero-mean identification signal $(0.14(-1)^{\lfloor k/2 \rfloor}, 0.36(-1)^k)$ is added after the QP and clipped to the physical correction limits.

For candidate Θ , define $\Omega = [-0.26, 0.26]$ and the full-reference disturbance bound

$$\begin{aligned} \bar{q}_\Theta &= \max_{\omega_r \in \Omega} q(\Theta, \omega_r) \\ &= \begin{bmatrix} h\Delta\theta_v(v_r + 0.9) + 0.006 \\ 0.012 \\ h\Delta\theta_\omega(0.26 + 1.6) + 0.005 \end{bmatrix}. \end{aligned} \quad (16)$$

Here and below, the maximum defining $\bar{q}_\Theta$ is taken componentwise. The robust-closure disturbance box is $\mathcal{D}_\Theta = \{d : |d| \leq \bar{q}_\Theta\}$. The closure audit computes

$$\lambda_\Theta = \max_{\substack{\theta \in \text{vert}(\Theta) \\ \omega_r \in \text{vert}(\Omega)}} \left\| P^{1/2} [A(\omega_r) + B(\theta)K] P^{-1/2} \right\|_2, \quad (17)$$

$$\omega_\Theta = \max_{d \in \text{vert}(\mathcal{D}_\Theta)} (d^\top P d)^{1/2}, \quad (18)$$

$$r_{\text{adm}} = \min\{r_x, r_u\}, \quad (19)$$

$$r_x = \min_i \frac{x_{\max, i}}{\sqrt{(P^{-1})_{ii}}}, \quad (20)$$

$$r_u = \min_j \frac{u_{\max, j}}{\|(KP^{-1/2})_{j,:}\|_2}. \quad (21)$$

The spectral norm of an affine matrix is convex, so the maximum over the joint parameter/reference box is attained at a vertex. Likewise, the convex quadratic disturbance norm attains its maximum over the box at a sign vertex; all eight disturbance vertices are enumerated. The gate passes when candidate widths are at most 0.10 and

$$(1 - \lambda_\Theta)r_{\text{adm}} - \omega_\Theta \geq 10^{-4}. \quad (22)$$

Three consecutive passes are required. Any failure resets the counter; the third pass copies the candidate to LOCAL. Figure 1 summarizes the command-boundary logic.

4.3 Executable protocol

At each boundary k , CRR-TMPC executes the following order: (1) form $\mathcal{C}_k$ from (e_{k-1}, u_{k-1}, e_k) ; (2) update LOCAL or trigger reset; (3) in TRANSITION, update the candidate and closure counter; (4) select the local or global box; (5) build radii and solve the QP; (6) apply the deadline rule and identification signal; and (7) log mode, sets, predicate, margins, command, residual, solver status, and latency. Thus the new mode affects u_k , not the already completed transition.

5 Conditional properties and claim boundary

Proposition 1 (Finite-horizon containment implication). *If the realized uncertainty and residual belong to the sets used in (11), the tube deviation belongs to the initial section, all propagated section inclusions hold, and (15) is feasible, then the executed state and correction satisfy their physical constraints over the covered prediction horizon.*

Proof. Induction in (11) yields componentwise tube-deviation containment. Minkowski addition with each tightened nominal constraint gives the original state and correction constraints. $\square$

Proposition 2 (Reset restores parameter outer containment). *If the post-change parameter belongs to Θ_G , resetting to Θ_G restores parameter containment. Intersecting only with the invalidated local box does not.*




Fig. 1 Implemented command-boundary CRR-TMPC logic. A completed transition first forms C_k . A nonempty local intersection updates Θ_k^L ; otherwise the local certificate is rejected, mode enters TRANSITION, and the reacquisition candidate is reset independently to Θ_G before contraction with C_k . Candidate width and the robust closure margin $m(\Theta) = (1 - \lambda_\Theta)r_{\text{adm}} - \omega_\Theta$ determine each closure pass; a failure resets the persistence counter, while the third consecutive pass promotes the candidate to LOCAL. The local or global tube controller is solved only after this set/mode update, so the new mode determines u_k . The change-containing interval precedes detection and is outside the coverage claim.

Proof. The first claim is membership. For the second, take any feasible post-change parameter in $C_k \setminus \Theta_{k-1}^L$. $\square$

Theorem 3 (Fixed-candidate terminal closure). *For a fixed candidate satisfying (22), the ellipsoid $\mathcal{E}_\Theta = \{\xi : \|\xi\|_P \leq r_{\text{adm}}\}$ is robust positively invariant for the tube-deviation dynamics (13) under the ancillary law $\delta u = K\xi$, for every $\theta \in \Theta$, every $\omega_r \in [-0.26, 0.26]$, and every $d \in \mathcal{D}_\Theta$, while the candidate and declared bounds remain valid.*

Proof. Joint vertex maximization bounds the induced norm for the complete affine parameter/reference box, and disturbance-vertex maximization bounds every $d \in \mathcal{D}_\Theta$. Hence $\|[A(\omega_r) + B(\theta)K]\xi + d\|_P \leq \lambda_\Theta r_{\text{adm}} + \omega_\Theta \leq r_{\text{adm}} - 10^{-4}$. The definition of r_{adm} provides admissibility of the ancillary state/input deviation; induction proves invariance of $\mathcal{E}_\Theta$ while the fixed candidate and declared bounds remain valid. $\square$

These statements do not prove recursive feasibility across arbitrary switches. The first interval containing an unannounced jump is observable

only through its successor residual and is excluded from the coverage claim. A global cross-mode theorem would additionally require an invariant receiving region and certified treatment of every timeout and repeated switch.

6 Experimental design

6.1 Frozen configuration

Table 2 reports the complete numerical configuration. After invalidating an earlier development certificate, the corrected controller was tested only on development seeds 1011, 1023, 1037, and 1041 in three families. Those results were not pooled with confirmation. Before opening the confirmatory record, code, endpoints, and analysis roles were frozen under SHA-256 configuration hash prefix `d1acedc196dec258d`; implementation and audit hashes are recorded in the freeze file. Untouched confirmatory seeds were 4000–4059 per family. No post-outcome controller, gate, threshold, scenario, or endpoint change was made.

Table 2 Frozen CRR-TMPC confirmatory-study configuration.

Quantity	Value
Control / integration step	0.10 / 0.02 s
Episode / horizon	12 s / 8 steps
Global / initial-local box	$[0, 0.25]^2$ / $[0, 0.10]^2$
Grid / candidate max width	0.01 / 0.10
Observation residual bound	(.008, .020, .005)
Process increment bound	(.004, .006, .004)
State limits	(.75, .60, .75)
Correction limits	(.90, 1.60)
Re-entry / numeric buffer	3 passes / 10^{-4}
QP solver	CVXPY 1.9.3, OSQP
Tolerance / iteration cap	10^{-5} / 4000
Deadline	100 ms
Confirmatory design	5 families $\times$ 60 seeds $\times$ 6 methods
Bootstrap replicates	5,000

For each seed, geometry alternates between circle and S-curve and disturbance scale alternates between 0.35 and 0.85 of the bound. Initial error is uniform on $[-.08, .08] \times [-.10, .10] \times [-.07, .07]$. Low coefficients are independently drawn from $[.025, .075] \times [.030, .085]$, and high coefficients from $[.165, .225] \times [.175, .235]$. Abrupt episodes switch at 4 s; repeated episodes switch at 3, 6, and 9 s; near-edge episodes switch at 4 s to $[.225, .245] \times [.230, .248]$. These three families define the in-scope step-change population. No-change tests false triggering. Gradual episodes interpolate from low to high during 3–8 s and are an explicitly out-of-scope diagnostic stress test. A fixed subset of abrupt and repeated episodes receives one declared isolated residual outlier.

6.2 Comparators and outcomes

All methods receive identical paired plant realizations. B0 is nominal MPC at $\theta = (.05, .05)$ without tube tightening. B1 is a fixed-global tube MPC and is the conservative robustness reference. B2 intersects the local set but never resets. CRR (internally B3) is the complete method. B4 detects inconsistency but removes independent non-monotone reset by retaining the stale candidate in TRANSITION. B5 is CRR with one-pass rather than three-pass re-entry. All other costs,

limits, horizon, solver settings, excitation, and backup rules are identical where applicable.

The primary mechanism endpoint is invalid-LOCAL exposure after annotated step changes. The co-primary operational endpoint is the covered successor-envelope exceedance rate. Detector outcomes are sensitivity within 3 s, unmatched events, and delay. Lateral RMSE is a performance endpoint, not a superiority endpoint. Prespecified conservatism endpoints compare local and global state-tube radii and input tightening over 105 reference-rate samples. Additional outcomes are heading RMSE, maximum error, physical-constraint breach, completion, effort, dwell, optimization failure, backup, and latency. A jump-containing interval is reported separately; all other intervals are “covered” for the operational audit. Change events are greedily matched within 3 s and remain clustered by episode. Means and paired CRR-minus-comparator effects use deterministic 5,000-replicate percentile-bootstrap 95% intervals. The episode is the independent unit.

7 Results

7.1 Certificate audit

The corrected full-range audit gave $\lambda = 0.872082$, $\omega = 0.279649$, $r_{\text{adm}} = 3.249189$, $r_{\text{req}} = 2.186155$, and buffered margin 0.135882 for the initial local box. The critical model vertex was $(\theta_v, \theta_\omega, \omega_r) = (0.10, 0.10, 0.26)$. The global box failed: $\lambda = 0.917120$, $\omega = 0.395736$, $r_{\text{req}} = 4.774790 > r_{\text{adm}} = 3.249189$, and margin -0.126542 (Fig. 2). Of 2,304 gridded candidate boxes in the sensitivity audit, 2,286 passed; the minimum positive buffered margin was 2.05×10^{-4} . A near-edge width-0.10 candidate failed, whereas width 0.08 passed. Consequently B1 and TRANSITION remain finite-horizon tube controllers; no global terminal-invariance claim is made.

7.2 Primary step-change evidence

All 1,800 controller episodes completed, and none contained a physical state or correction-limit breach. Table 3 aggregates the 180 plant episodes in the three in-scope step-change families. CRR detected 297 of 300 annotated changes (99.0% sensitivity within 3 s). It reduced invalid-LOCAL exposure from 14,550 samples for stale




Fig. 2 Corrected robust-closure audit over the complete reference-rate interval and exact disturbance-box sign vertices. Lines show the buffered margin $m(r) = (1 - \lambda_\Theta)r - \omega_\Theta - 10^{-4}$ for the initial local and global parameter boxes; the horizontal dashed line is the closure boundary, the vertical dotted line marks r_{adm} , circles show the margin at r_{adm} , and diamonds show the zero-margin radius r_{req} . The initial local box passes whereas the global envelope fails the fixed-candidate terminal-closure condition.

intersection-only B2 to 865, a 94.1% reduction. Covered successor-envelope exceedances fell from 11,131 for B2 to 103 for CRR (99.1% reduction) and from 3,958 for no-independent-reset B4 to 103 (97.4% reduction). Figure 3 shows episode-bootstrap uncertainty for these mechanisms.

The fixed-global B1 controller produced 17 covered exceedances, compared with 103 for CRR. Its lower count is consistent with continuous use of the broad envelope and does not support a claim that adaptation is safer than a feasible global tube. The 103 nonzero CRR exceedances also show that the implemented axis-aligned radius is an operational approximation rather than a verified nonlinear enclosure over every sampled interval.

7.3 Detector operating characteristics and ablations

CRR detected all 60 abrupt changes and all 60 near-edge changes. It detected 177 of 180 repeated changes. Mean delay was 0.205 s for abrupt changes, 0.206 s for repeated changes, and 0.103 s near the envelope edge. Thirteen of the 16 unmatched CRR events occurred in abrupt episodes containing the declared residual challenge; three occurred in repeated episodes. No-change episodes produced zero CRR triggers. Figure 4 separates sensitivity from delay and unmatched-event burden.

B4’s 6,610 invalid-LOCAL samples and 3,958 covered exceedances isolate the importance of reacquiring independently of the rejected set. B5 nearly matched CRR (109 versus 103 exceedances), with the same number of detected changes and a shorter transition dwell. The experiment therefore does not establish an empirical advantage for three-pass over one-pass return. The persistence rule remains a conservative design choice, not a validated contribution. The representative abrupt trace in Fig. 5 shows reset, global receipt, candidate contraction, and subsequent LOCAL re-entry.

7.4 Recovered authority relative to B1

Across 105 samples of $\omega_r \in [-0.26, 0.26]$, returning from the global to the initial local construction reduced median terminal state-radius sum by 26.6% and median input tightening at prediction step seven by 23.4% (Fig. 6). These quantities measure recovered state and command authority directly; they do not imply better nonlinear containment or tracking in every episode. B1 retains the broad envelope continuously, whereas CRR accepts transition overhead to recover a smaller tube when candidate-specific evidence and closure permit it.

7.5 Tracking and out-of-scope gradual drift

CRR was not tuned or hypothesized to dominate B1 in tracking. Its mean lateral RMSE over the step-change families was 0.01790 m, versus 0.01770 m for B1. Paired CRR-minus-B1 intervals excluded zero for abrupt and repeated changes, while the near-edge interval included zero. Figure 7 gives all paired effects, including the diagnostic gradual family. The excitation and TRANSITION periods improve reacquisition evidence at a small tracking cost in some families.

The gradual-drift stress test met the 3 s onset-detection window in only 10 of 60 episodes; 50 windows were missed, and 54 later events were unmatched. CRR accumulated 2,768 invalid-LOCAL samples and 64 covered exceedances in this family. The demonstrated scope is therefore abrupt regime changes rather than slowly drifting

Table 3 Confirmatory totals over the abrupt, repeated, and near-edge step-change families (180 paired plant episodes per controller; 21,600 executed intervals per controller). Dashes indicate that B0/B1 do not run the local-set detector. Exceedances exclude the explicitly uncovered jump-containing intervals.

Controller	Detected/true	Missed	Unmatched events	Invalid LOCAL	Covered exceed.	Breach episodes	Mean RMSE (m)
B0 nominal	—	—	—	—	10,548	0	.01802
B1 fixed global	—	—	—	0	17	0	.01770
B2 intersection only	262/300	38	1,924	14,550	11,131	0	.01789
CRR-TMPC	297/300	3	16	865	103	0	.01790
B4 no independent reset	261/300	39	1,893	6,610	3,958	0	.01788
B5 one-pass return	297/300	3	17	854	109	0	.01787




Fig. 3 Primary mechanism evidence across the abrupt, repeated, and near-edge step-change families. Panel (a) reports invalid-LOCAL exposure for the adaptive-set methods; B0 and B1 do not run LOCAL-set adaptation and are therefore not applicable. Panel (b) reports covered successor-envelope exceedances for all six controllers. Bars are episode-mean counts per 1,000 intervals classified as covered, labels give aggregate totals, and whiskers are deterministic 5,000-replicate episode-bootstrap 95% percentile intervals. Jump-containing intervals are excluded from the covered-exceedance endpoint.

parameters. A windowed or rate-aware estimator is required for that extension.

7.6 Computation and failure accounting

The campaign ran on a Linux 6.18.44 x86-64 virtual host with Python 3.12.14, NumPy 2.3.5, pandas 2.2.3, and CVXPY 1.9.3. Table 4 and Fig. 8 summarize measured wall-clock latency over 300 episodes per controller. No solve was declared infeasible. Eight deadline misses across all controllers invoked the declared backup; CRR accounted for three. Since timing was measured on a shared virtual host, these counts characterize this execution rather than a real-time hardware guarantee.

8 Discussion

Across the tested abrupt regimes, independent reset and post-reset reacquisition removed most

Table 4 Runtime and operational counts over 300 episodes per controller. Times are milliseconds.

Method	Median	Mean ep. p95	Max	Late/backup
B0	1.711	2.378	106.3	1/1
B1	1.713	2.368	104.3	2/2
B2	1.710	2.350	94.3	0/0
CRR	1.730	2.389	110.0	3/3
B4	1.749	2.386	109.8	1/1
B5	1.735	2.420	100.9	1/1

of the stale-set exposure and operational containment failures seen in B2 and B4. The corrected full-reference-range certificate, 297/300 event sensitivity, zero no-change triggers, paired ablations, and command-boundary trace are consistent with the intended certificate-rejection and reacquisition mechanism.

Tracking and containment results also constrain the interpretation. CRR does not generally improve tracking over B1; its identification signal and TRANSITION dwell increase RMSE slightly in abrupt and repeated families. B1 has fewer containment-audit failures because its




Fig. 4 Confirmatory CRR-TMPC detector outcomes for the three in-scope step-change families. Panel (a) shows sensitivity using annotated change events as the denominator, with labels giving detected/annotated totals. Panel (b) shows mean matched-detection delay, and panel (c) shows unmatched detector events per episode, with aggregate unmatched-event totals in parentheses. Whiskers are deterministic 5,000-replicate episode-bootstrap 95% percentile intervals over the 60 plant episodes in each family; unmatched events are detections not paired to an annotation within the 3 s matching window.

broad envelope remains active continuously. Reacquisition instead recovers local authority, with a 26.6% median terminal-radius reduction and 23.4% median input-tightening reduction once local certification is restored. Gradual drift defeats the onset-window detector in 50 of 60 episodes, so rate-aware set propagation, a sliding window, or a sequential detector is needed before gradual-change coverage can be claimed.

The 103 in-scope CRR covered-envelope exceedances expose another limitation: the axis-aligned first-order radius formula is an operational approximation, not a verified enclosure of the RK4 nonlinear plant over the entire domain. The corrected ellipsoidal closure theorem is valid for the affine surrogate and declared box disturbance, but the nonlinear RK4 successor audit shows that these surrogate premises do not hold on every sampled interval. Future work should compute a certified nonlinear remainder, replace the grid hull with verified set propagation, and construct a cross-mode invariant receiving region. Hardware validation should follow, not substitute for, those proof obligations.

Threats to validity include 60 plant seeds per family, deterministic balancing rather than random assignment of geometry and disturbance level, one mobile-robot model, one tuning budget, a gridded consistency approximation, and wall-clock measurements on a shared virtual host. The campaign has no hardware or outside- Θ_G claim. Bootstrap intervals quantify seed variability, not model-form uncertainty. The confirmatory seeds

are distinct from the earlier pilot and reference-study seeds.

9 Conclusion

CRR-TMPC addresses a specific failure mode in adaptive tube MPC: an abrupt admissible regime change that falsifies the local uncertainty certificate. Across 180 confirmatory step-change plant episodes, it detected 297 of 300 changes, reduced invalid-local exposure by 94.1% relative to intersection-only adaptation, and reduced covered envelope exceedances by 97.4% relative to the no-independent-reset ablation. Re-entry to the certified local construction reduced median terminal tube radius by 26.6% and input tightening by 23.4% relative to continuous global operation. The demonstrated contribution is certificate-triggered reset and reacquisition for recovering local control authority after abrupt admissible changes. The study does not establish safety or tracking superiority over a fixed-global tube, gradual-drift coverage, a universally valid nonlinear enclosure, arbitrary-switching recursive feasibility, or hardware safety.

Data availability

All raw episode metrics, traces used in figures, summary tables, configuration manifests, certificate-audit records, and frozen analysis outputs are supplied in Online Resource 1 (CRR-TMPC V3 reproducibility archive).




Fig. 5 Prespecified representative confirmatory abrupt episode (seed 4100). Panel (a) shows the true actuation-effectiveness losses together with the active local or TRANSITION candidate bounds; panel (b) shows lateral tracking error and its physical limits; panel (c) shows protocol mode, the persistence counter normalized by three, and successful closure-gate evaluations; and panel (d) shows the buffered closure margin relative to the zero closure boundary. The shaded interval is TRANSITION, the red dotted line marks local-certificate invalidation, and the green dash-dot line marks LOCAL re-entry after three consecutive closure passes.

Code availability

The complete Python reference implementation, certificate-audit script, and deterministic manuscript-figure regeneration script are supplied in Online Resource 1. The confirmatory configuration hash and untouched seed list are reported in the manuscript and machine-readable manifest.

Appendix A Reproducibility map

Online Resource 1 contains the frozen V3 implementation and analysis record. Running `python crr.tmpc_study_v3.py heldout` reproduces the 1,800 controller-episode record, summaries,

paired effects, traces, and manifest. Running `python audit_v3_certificate.py` reproduces the 2,304-case certificate sensitivity grid and local-global conservatism records. The supplied `generate_submission_figures_v31.py` script regenerates the manuscript figures from the frozen CSV/JSON/trace records without changing the controller or endpoints. CSV columns preserve controller, family, geometry, disturbance, seed, outcomes, detector counts, margins, solver counts, and latency. The V3 freeze file states hashes, analysis unit, scope, endpoints, untouched seed range, uncovered-interval rule, and prohibition on post-opening changes.




Fig. 6 Direct local-versus-global conservatism comparison over 105 reference yaw-rate samples in $[-0.26, 0.26]$ rad/s. Panel (a) compares the prediction-horizon terminal ℓ_1 state-tube radius for the initial local set and global envelope; panel (b) gives the corresponding terminal-radius reduction, whose dashed line marks the 26.6% median. Panel (c) compares retained input authority at prediction step seven, and panel (d) gives the corresponding input-tightening reduction, whose dashed line marks the 23.4% median. Vertical dotted lines mark zero reference yaw rate.

Appendix B Interpretation of containment counts

For each executed interval, the audit forms the centre-model successor and compares its residual componentwise with the observation bound plus the active parameter half-width contribution. If a declared abrupt jump occurs at boundary k , interval $[k, k + 1]$ is labelled uncovered because the controller can observe the new coefficient only in e_{k+1} . Exceedances on those intervals are reported as change-interval counts and excluded from the covered count. This accounting prevents retroactive certification while retaining all failures in the archive.

References

- [1] Lu, X., Cannon, M. & Koksal-Rivet, D. Robust adaptive model predictive control: Performance and parameter estimation.

International Journal of Robust and Nonlinear Control **31**, 8703–8724 (2021). <https://doi.org/10.1002/rnc.5175>.

- [2] Köhler, J., Kötting, P., Soloperto, R., Allgöwer, F. & Müller, M. A. Robust adaptive model predictive control framework for nonlinear uncertain systems. *International Journal of Robust and Nonlinear Control* **31**, 8725–8749 (2021). <https://doi.org/10.1002/rnc.5147>.
- [3] Sasfi, A., Zeilinger, M. N. & Köhler, J. Robust adaptive mpc using control contraction metrics. *Automatica* **155**, 111169 (2023). <https://doi.org/10.1016/j.automatica.2023.111169>.
- [4] Aboudonia, A. & Lygeros, J. Adaptive learning-based model predictive control for uncertain interconnected systems: A set membership identification approach. *Automatica* **171**, 111943 (2025). <https://doi.org/10.1016/j.automatica.2024.111943>.




Fig. 7 Paired lateral-RMSE effects, reported as CRR-TMPC minus comparator, for no-change, abrupt, repeated, near-edge, and gradual-drift families. Points are paired episode means and whiskers are deterministic 5,000-replicate episode-bootstrap 95% percentile intervals. Negative values favor CRR-TMPC and positive values favor the comparator; filled markers denote intervals excluding zero and open markers denote intervals including zero. Gradual drift is shown only as an out-of-scope diagnostic.




Fig. 8 Runtime and deadline accounting over 300 episodes per controller. Panel (a) shows the distribution of each episode's 95th-percentile solve latency, with diamonds marking the mean episode p95. Panel (b) shows the maximum observed solve latency for each controller relative to the 100 ms deadline; annotations report the maximum latency and the corresponding late/backup counts. Measurements were obtained on the shared virtual host used for the confirmatory campaign and are not a real-time hardware guarantee.

- [5] Dhar, A. & Bhasin, S. Homothetic tube-based adaptive mpc. *International Journal of Robust and Nonlinear Control* **35**, 5705–5716 (2025). <https://doi.org/10.1002/rnc.8010>.

- [6] Köhler, J., Soloperto, R., Müller, M. A. & Allgöwer, F. A computationally efficient robust model predictive control framework for uncertain nonlinear systems. *IEEE Transactions on Automatic Control* **66**, 794–801 (2021). <https://doi.org/10.1109/TAC.2020.2982585>.
- [7] Subramanian, S., Lucia, S. & Paulen, R. Tube-enhanced multi-stage model predictive control for flexible robust control of constrained linear systems with additive and parametric uncertainties. *International Journal of Robust and Nonlinear Control* **31**, 4458–4487 (2021). <https://doi.org/10.1002/rnc.5486>.
- [8] Villanueva, M. E., Müller, M. A. & Houska, B. Configuration-constrained tube mpc. *Automatica* **163**, 111543 (2024). <https://doi.org/10.1016/j.automatica.2024.111543>.

- [9] Sieber, J., Didier, A. & Zeilinger, M. N. Computationally efficient system level tube-mpc for uncertain systems. *Automatica* **180**, 112466 (2025). <https://doi.org/10.1016/j.automatica.2025.112466>.
- [10] Fukushima, H., Kim, T.-H. & Sugie, T. Adaptive model predictive control for a class of constrained linear systems based on the comparison model. *Automatica* **43**, 301–308 (2007). <https://doi.org/10.1016/j.automatica.2006.08.026>.
- [11] Lu, X. & Cannon, M. Robust adaptive model predictive control with persistent excitation conditions. *Automatica* **152**, 110959 (2023). <https://doi.org/10.1016/j.automatica.2023.110959>.
- [12] Hewing, L., Wabersich, K. P., Menner, M. & Zeilinger, M. N. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems* **3**, 269–296 (2020). <https://doi.org/10.1146/annurev-control-090419-075625>.
- [13] Puig, V. Fault diagnosis and fault tolerant control using set-membership approaches: Application to real case studies. *International Journal of Applied Mathematics and Computer Science* **20**, 619–635 (2010). <https://doi.org/10.2478/v10006-010-0046-y>.
- [14] Tanaskovic, M., Fagiano, L. & Gligorovski, V. Adaptive model predictive control for linear time varying mimo systems. *Automatica* **105**, 237–245 (2019). <https://doi.org/10.1016/j.automatica.2019.03.030>.
- [15] Bujarbaruah, M., Zhang, X., Tanaskovic, M. & Borrelli, F. Adaptive stochastic mpc under time-varying uncertainty. *IEEE Transactions on Automatic Control* **66**, 2842–2849 (2021). <https://doi.org/10.1109/TAC.2020.3009362>.
- [16] Fleming, J. & Hawari, Q. Robust tube mpc using gain-scheduled policies for a class of lpv systems. *IEEE Control Systems Letters* **8**, 1589–1594 (2024). <https://doi.org/10.1109/LCSYS.2024.3412652>.
- [17] Wu, T., Zhang, L., Lu, S., Che, W. & Xu, K. Efficient and smooth robust model predictive control for stochastic switching systems. *Automatica* **163**, 111572 (2024). <https://doi.org/10.1016/j.automatica.2024.111572>.
- [18] Zhang, S., Jia, R., He, D. & Chu, F. A chance-constrained tube-based model predictive control for tracking linear systems using data-driven uncertainty sets. *International Journal of Robust and Nonlinear Control* **34**, 969–995 (2024). <https://doi.org/10.1002/rnc.7010>.
- [19] Peschke, T. & Görges, D. Robust adaptive tube tracking model predictive control for piece-wise constant reference signals. *International Journal of Robust and Nonlinear Control* **33**, 8158–8182 (2023). <https://doi.org/10.1002/rnc.6814>.
- [20] Qin, W., Tang, W. & Xia, W. Robust fault-tolerant control based on set-membership estimation and tube mpc. *IEEE Transactions on Automation Science and Engineering* **22**, 19499–19510 (2025). <https://doi.org/10.1109/TASE.2025.3594583>.
- [21] Daniels, A., Teutsch, J., Kleindienst, F., Leibold, M. & Wollherr, D. Active fault identification and robust control for unknown bounded faults via volume-based costs. 2025 6th International Conference on Control and Fault-Tolerant Systems (SysTol), pp. 126–131 (2025). URL <https://doi.org/10.1109/SysTol66549.2025.11267332>. IEEE, DOI: 10.1109/SysTol66549.2025.11267332.
- [22] Wang, Y., Liu, Y., Liu, Y.-H. & Su, C.-Y. Differential flatness-based adaptive robust tracking control for wheeled mobile robots with slippage disturbances. *ISA Transactions* **144**, 482–489 (2024). <https://doi.org/10.1016/j.isatra.2023.11.008>.
- [23] Korayem, M. H., Safarbali, M. & Yousefi Lademakhi, N. Adaptive robust control with slipping parameters estimation based on intelligent learning for wheeled mobile robot. *ISA Transactions* **147**, 577–589 (2024). <https://doi.org/10.1016/j.isatra.2024.02.008>.

- [24] Korayem, M. H., Namdarpour, F. & Lademakhi, N. Y. Superiority of model predictive control with robust and stable approach for sliding wheeled mobile systems in the presence of obstacles. *Optimal Control Applications and Methods* **45**, 1503–1523 (2024). <https://doi.org/10.1002/oca.3107>.
- [25] Li, P., Yang, H. & Wang, S. Model predictive tracking control with disturbance compensation for wheeled mobile robots in an environment with obstacles. *Journal of the Franklin Institute* **360**, 6669–6692 (2023). <https://doi.org/10.1016/j.jfranklin.2023.05.004>.
- [26] Li, L., Cao, W., Yang, H. & Geng, Q. Trajectory tracking control for a wheel mobile robot on rough and uneven ground. *Mechatronics* **83**, 102741 (2022). <https://doi.org/10.1016/j.mechatronics.2022.102741>.
- [27] Aro, K., Guevara, L., Torres-Torriti, M., Torres, F. & Prado, A. Robust nonlinear model predictive control for the trajectory tracking of skid-steer mobile manipulators with wheel-ground interactions. *Robotics* **13**, 171 (2024). <https://doi.org/10.3390/robotics13120171>.