DASC 45th
Digital Avionics Systems Conference
Orlando, Florida, USA - September 13-17, 2026

Runtime Safety Filtering for Learned Small UAS Separation Policies under GNSS Degradation

Action filtering vs observation filtering

Alex Zongo · Peng Wei

Department of Mechanical and Aerospace Engineering
George Washington University, Washington, DC, USA.

A learned separation policy can resolve conflicts

With nominal GNSS, the policy produces a separation-preserving action.

The policy uses nearby aircraft information to adjust ownship speed in real time.

100 m separation standard
10 km routes

Next: What information does the policy receive, and how can GNSS degradation corrupt it?

SIMULATION · NO GNSS DEGRADATION (NOMINAL GNSS)

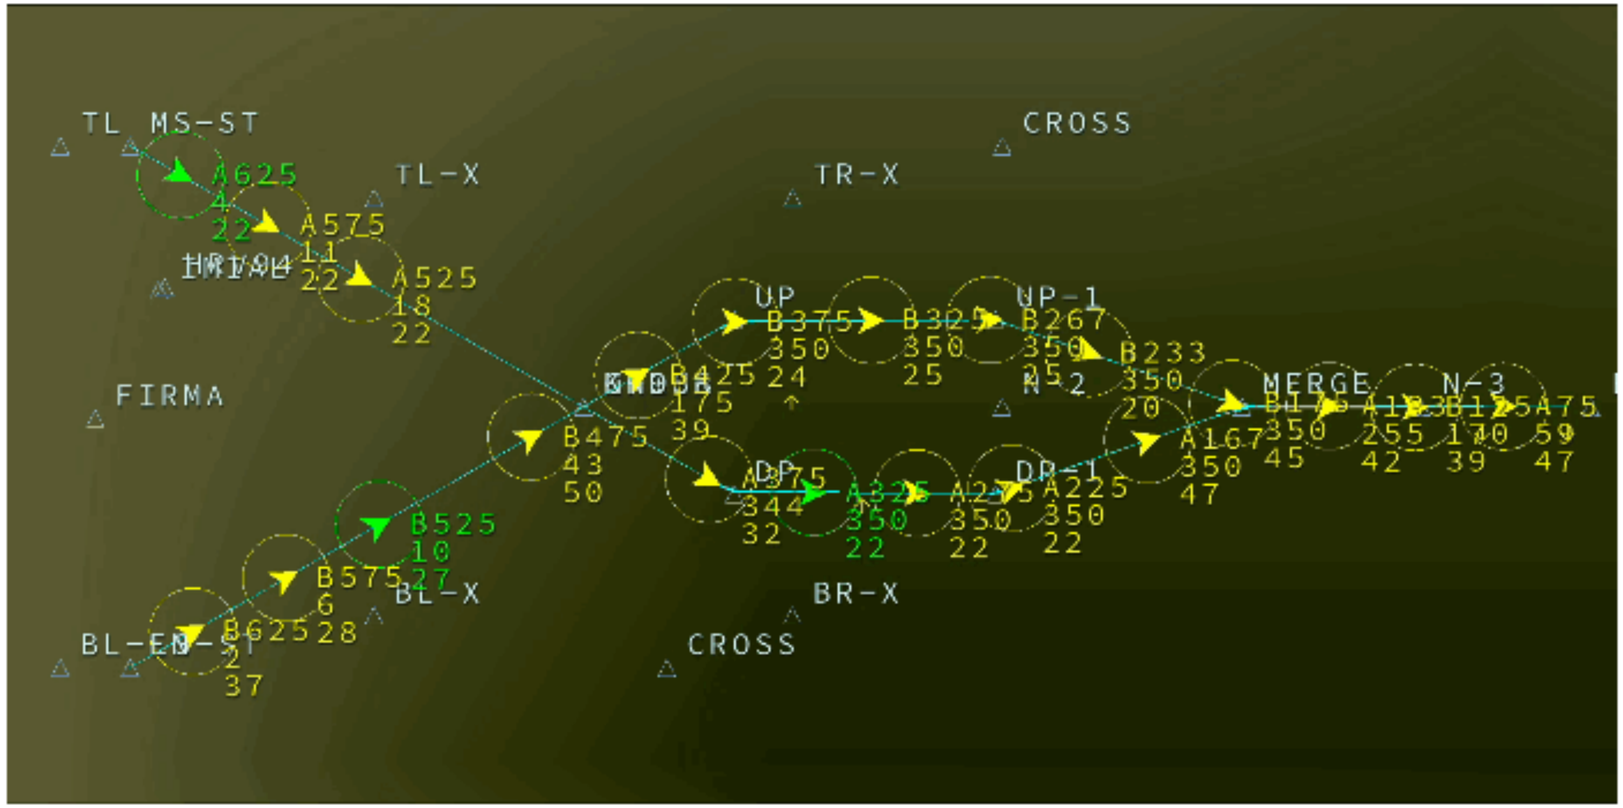

OBSERVE → DECIDE → MOVE

How a decentralized learned separation policy works

1. LOCAL TRAFFIC

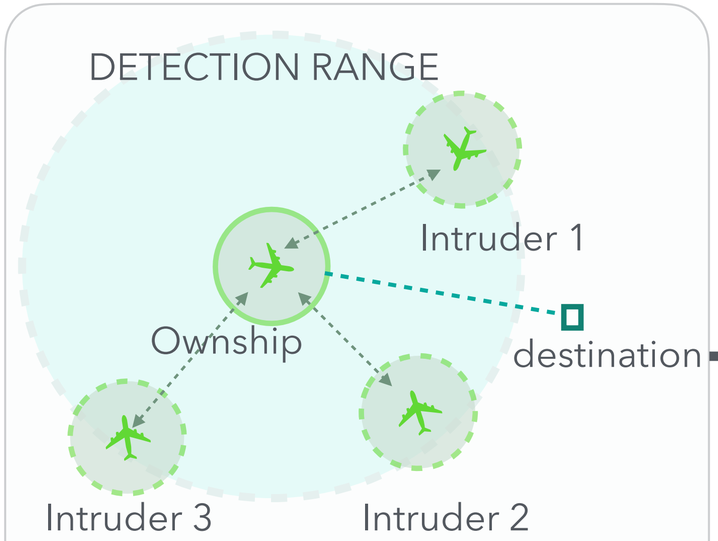

Traffic within sensing range, broadcast over Remote ID

2. LOCAL OBSERVATION

- Ownship: position, speed, heading
- Intruder 1: position, speed, heading
- Intruder 2: position, speed, heading
- Intruder 3: position, speed, heading
- … up to N in range

3. POLICY TO ACTION

$$\pi_\theta(a \mid \text{obs})$$

slow-down	hold speed	speed up
$-\Delta v$	± 0	$+\Delta v$

$\Delta v \approx 2.57$ m/s per decision step

heading follows the structured route

A speed change by one UAS alters what its neighbors observe.

GNSS degradation distorts the state the policy relies on

1. WHY URBAN GNSS DEGRADES

Urban environments disrupt navigation integrity

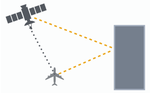

Multipath
Reflections from building faces bias position estimation

Blockage
Urban canyons create gaps and unstable fixes

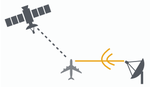

Spoofing / jamming
Delayed signals, false or replayed signals, indistinguishable from the truth

2. CONSEQUENCE

A closing conflict can be reported as clear traffic

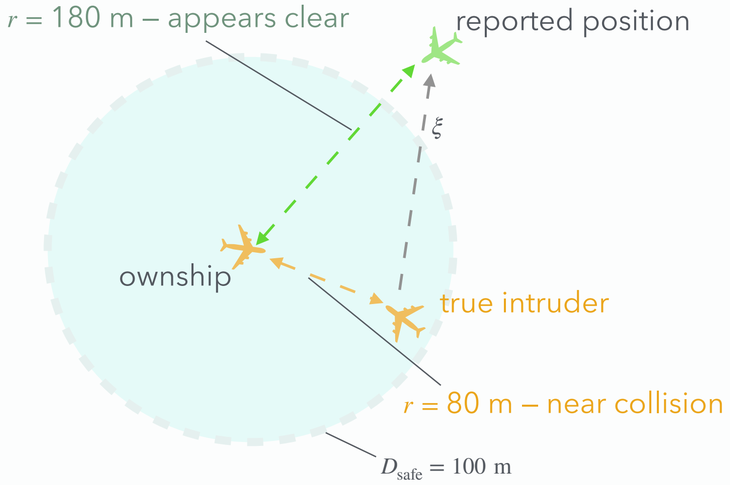

3. DEGRADATION MODEL

A bounded state degradation model

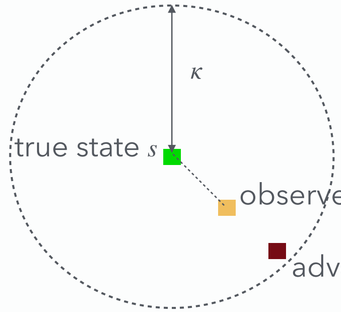

$o = (1 - R)s + R\xi$, where $\|\xi - s\|_\infty \leqslant \kappa$

Perturbation is bounded componentwise

$R = 0$ · nominal — $R \uparrow$ · degraded

The separation policy then acts on observation o, while safety is measured on true state s.

Robustness can be built in — or handled at runtime

ADVERSARIAL / ROBUST TRAINING

- Shapes how the policy responds to corrupted or threatening observations
- Robustness depends on the assumed disturbances and training distribution
- A train-deployment gap can remain when real degradation exceeds those assumptions

RUNTIME FILTERING

- Acts during execution, independent of how the policy was trained
- Requires no retraining of the deployed policy
- Uses the current degradation bound at each decision step
- Has no training-time knowledge and inherits whatever the policy learned

How should a runtime filter attach to a learned policy?

Before the policy sees the state — or after it proposes an action?

Two ways to add a runtime safety filter

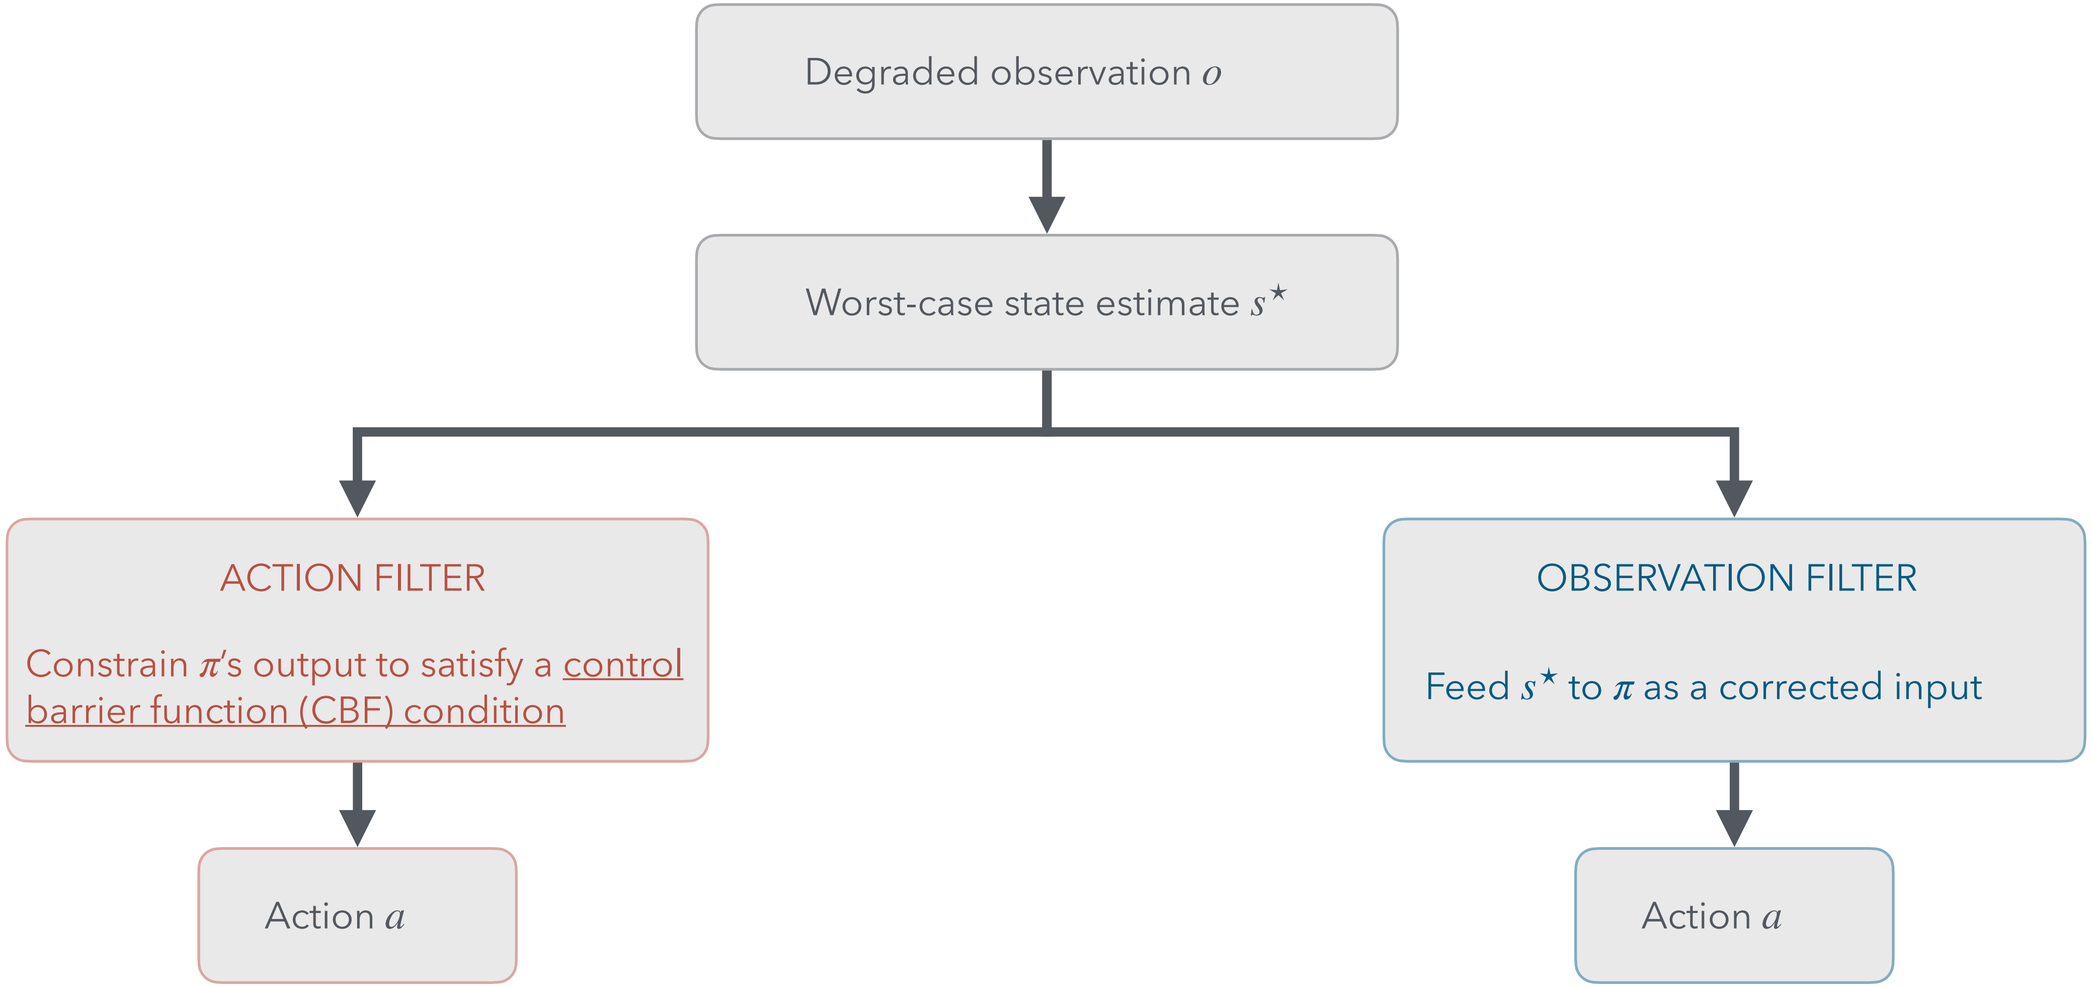

Both filters start from the same worst-case state estimation

1. SAFETY MARGIN

Barrier function as safety margin

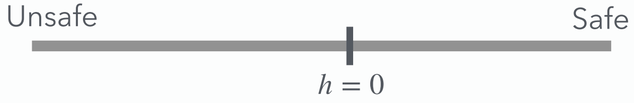

$$h(s) = (\|r\|^2 - D_{\text{safe}}^2) + \alpha \cdot r^T v_{\text{rel}}$$

$\|r\|^2 - D_{\text{safe}}^2$

How much room there is between two UAS

$\alpha \cdot r^T v_{\text{rel}}$

Whether two UAS are closing in on each other

α sets how far ahead the margin looks

2. WORST CASE

What looked clear now reads as a conflict

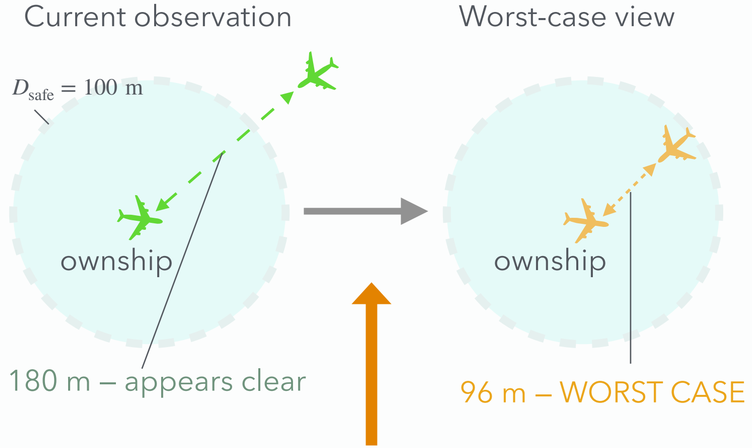

3. IN CLOSED FORM

A closed-form displacement

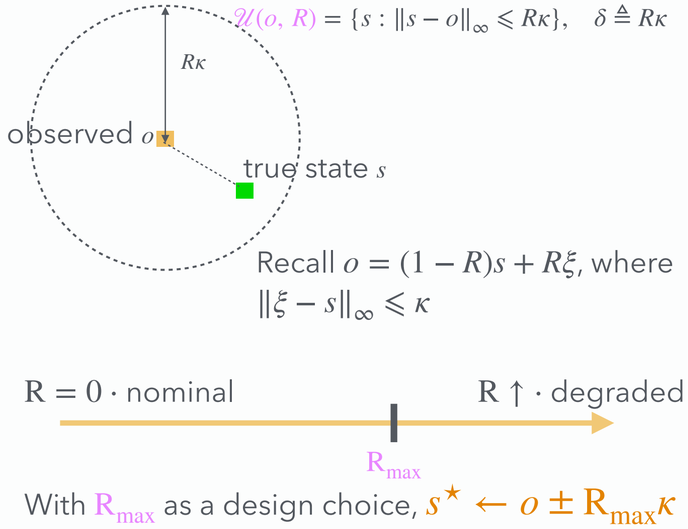

Both filters compute $s^\star$ from the same safety margin h. However, only the action filter lets h veto the policy.

The filtering approaches differ in what happens next

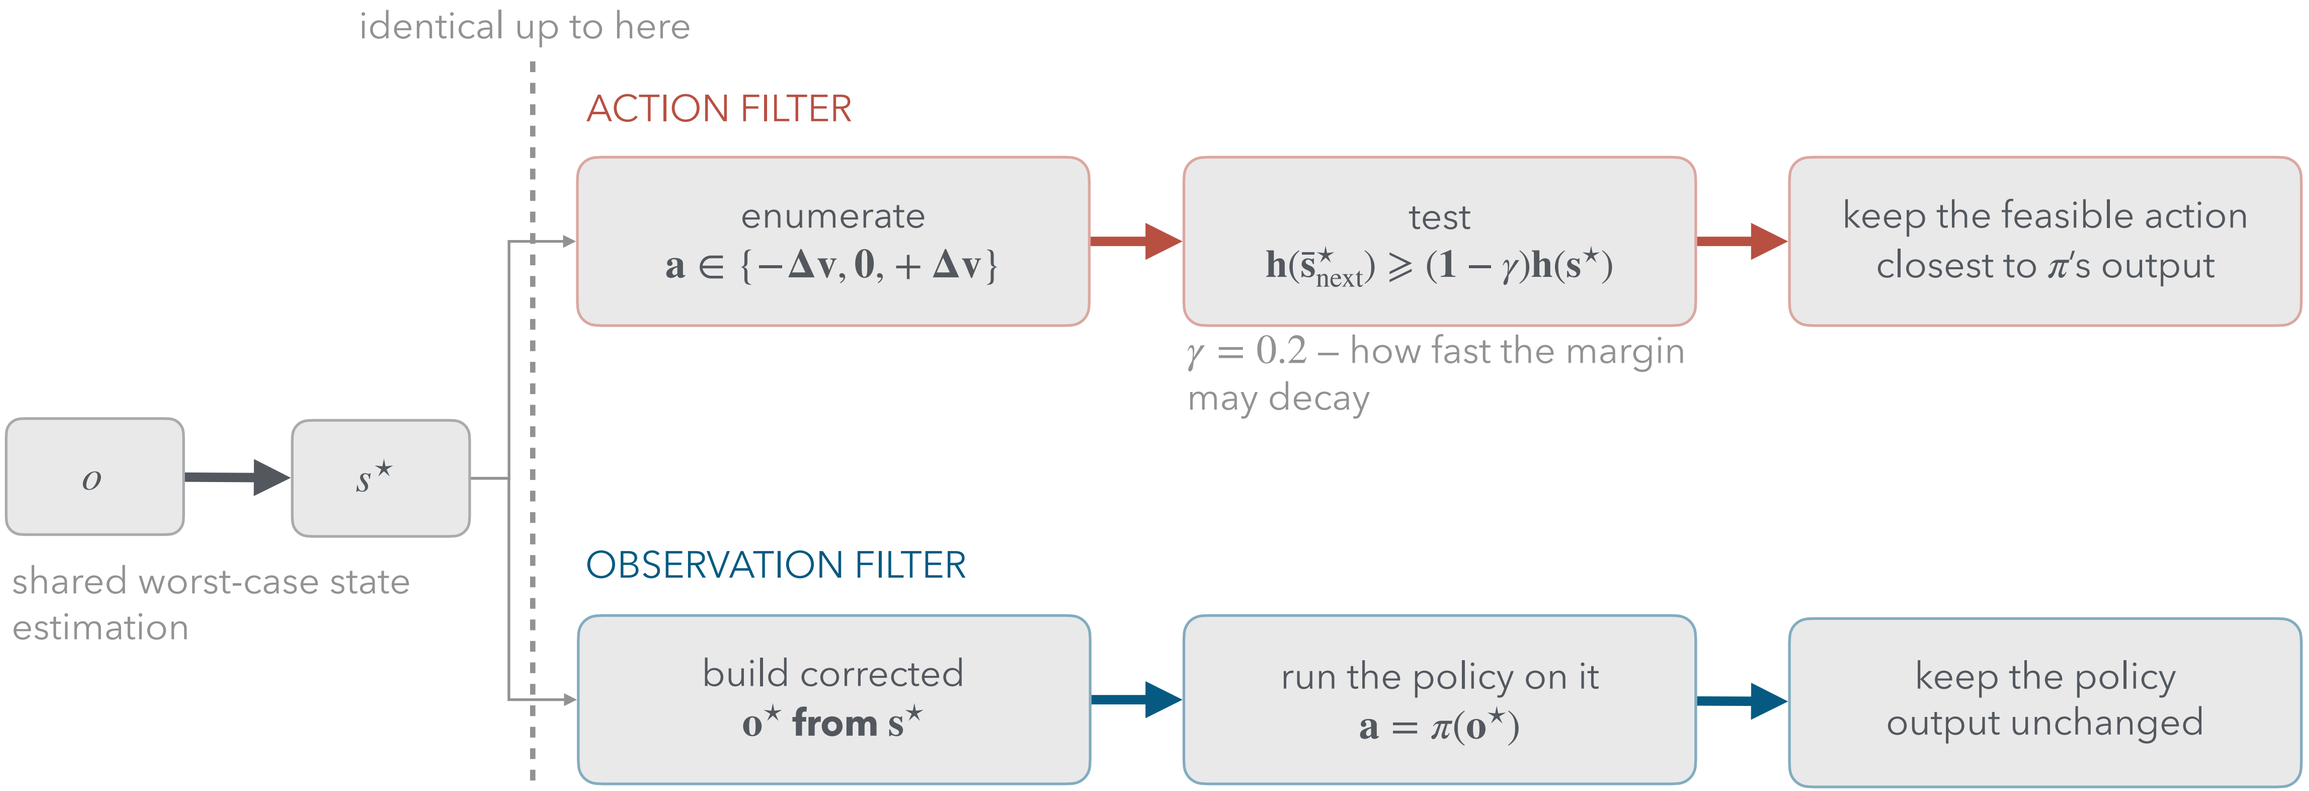

Same estimate, different use. The discrete-time CBF condition applies only to the action filtering.

Action filtering barely moves the safety needle

NMAC

Near mid-air collision when two aircraft pass within $D_{\text{safe}} = 100$ m; counted per episode.

R

Contamination rate: how degraded GNSS actually is. 0 is nominal; 1 is fully adversarial.

R_{max}

What the filter assumes: a design choice fixed before deployment.

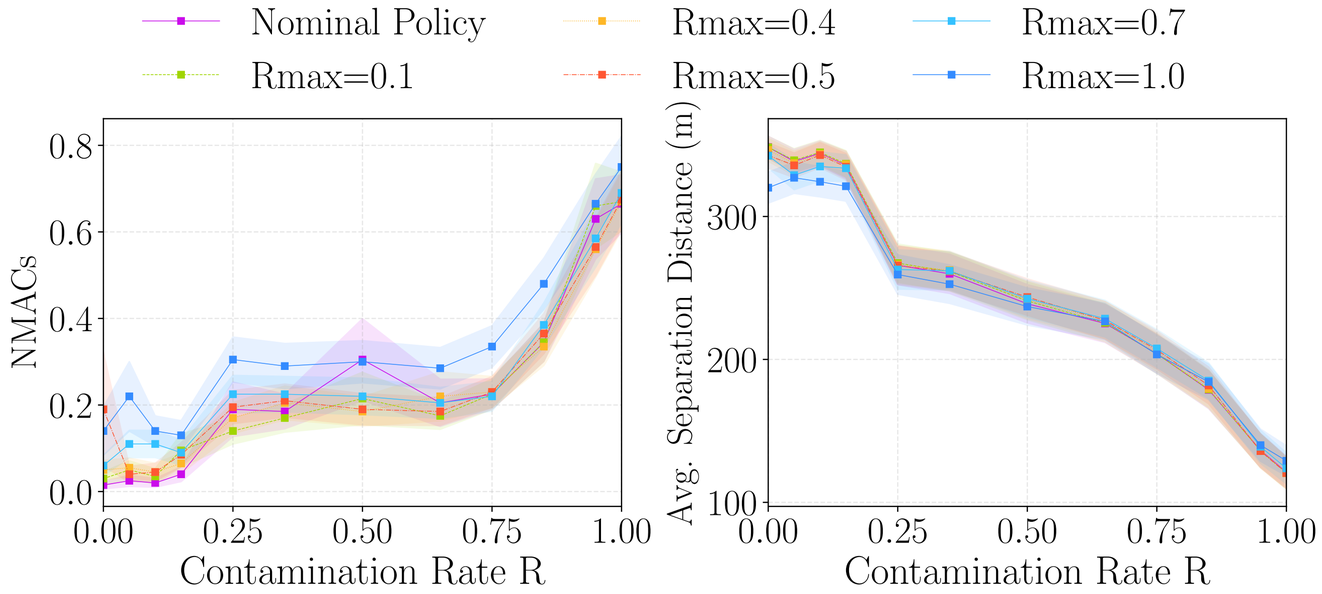

Shaded bands are ± standard error.

Every setting tracks the nominal policy. Despite changing R_{max}, action filtering does not separate from the unfiltered baseline.

Observation filtering removes nine near collisions in ten

90%

fewer near mid-air collisions

0.24 → 0.02

NMACs per episode

250 m → 375 m

average separation

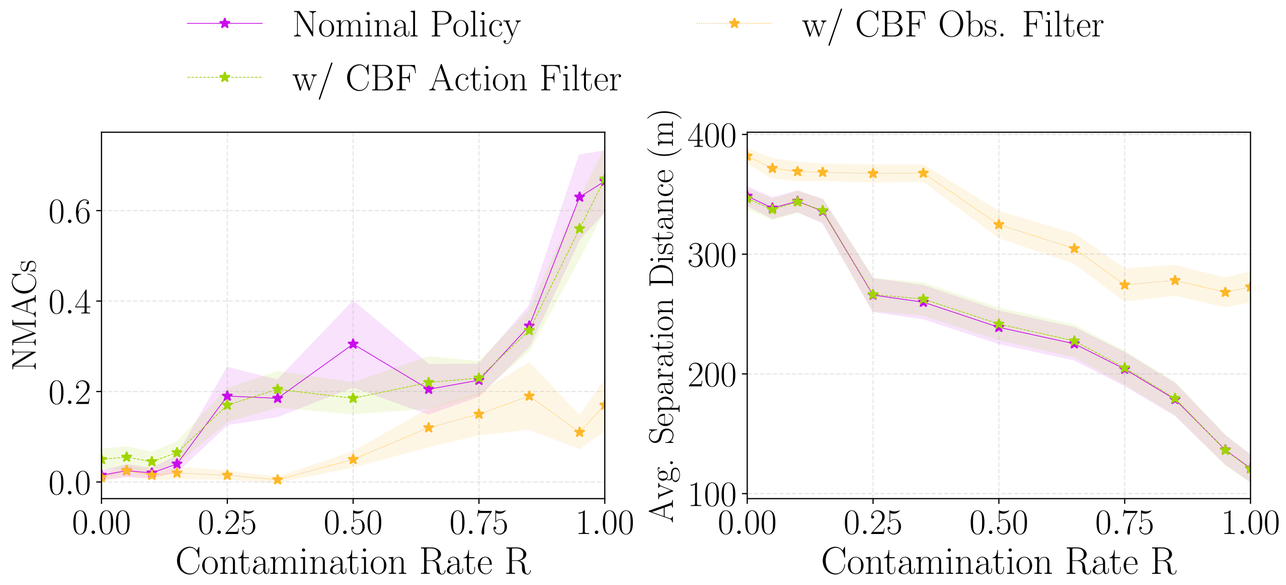

With the same worst-case estimate, the observation filter performs better on its own.

The advantage grows with what the filter assumes

$R_{max} = 0.1 - 0.5$

Little benefit. The assumed bound is too small to surface the danger.

$R_{max} = 0.7 - 1.0$

NMACs are near zero; separation holds around 380 m as nominal policy collapses to 130 m.

No tuning

The filter never observes the true R; R_{max} is fixed before deployment.

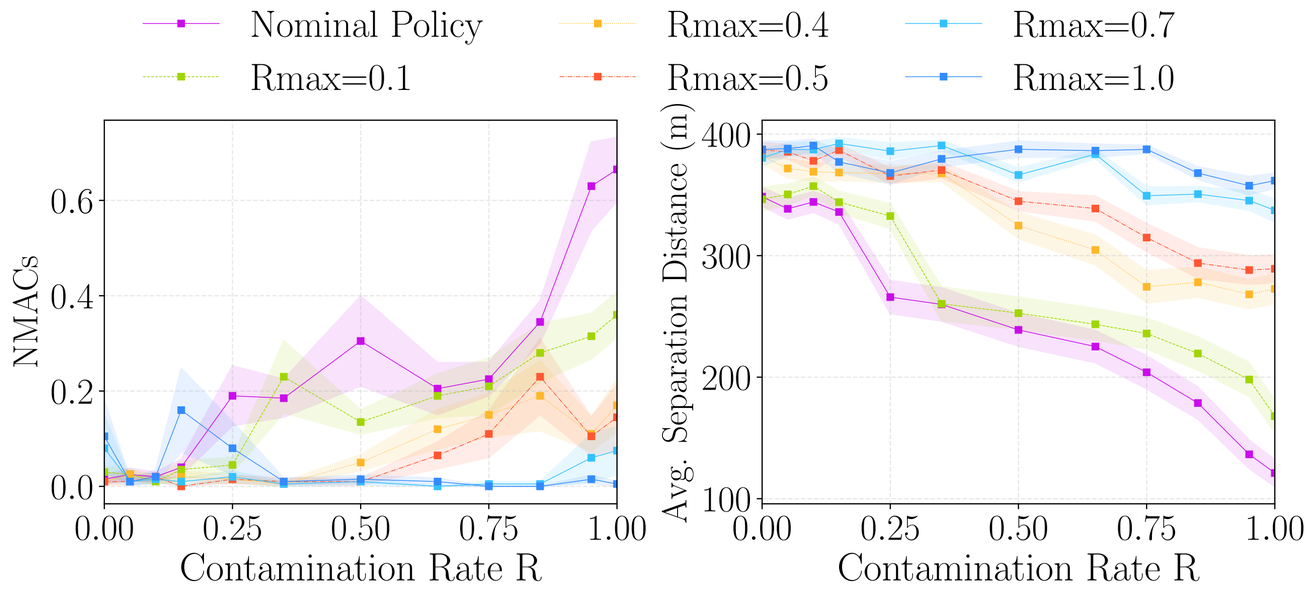

The filter is never told how bad the degradation really is.

Action filtering breaks as α grows; observation filtering does not

ACTION FILTER

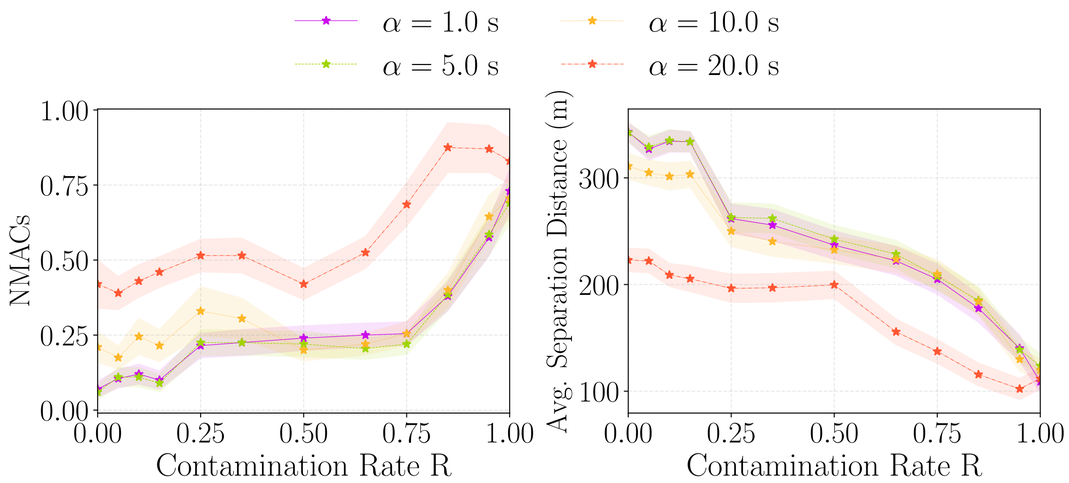

$\alpha = 20\ s$ more than doubles the NMACs.

OBSERVATION FILTER

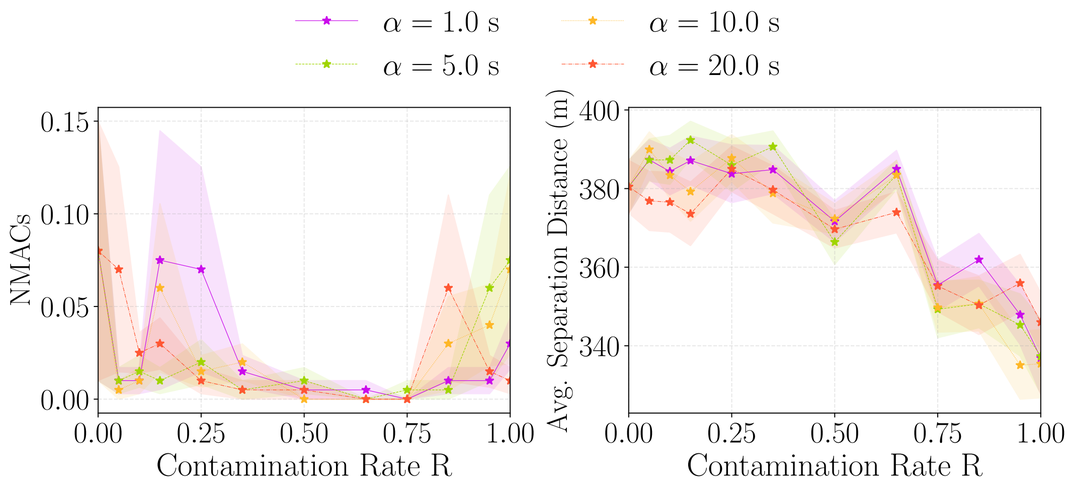

All α values' effects seem to coincide.

α is a barrier-function parameter. It matters only where the barrier can veto the policy's action.

05 WHY

The barrier looks one step ahead; the policy looks further

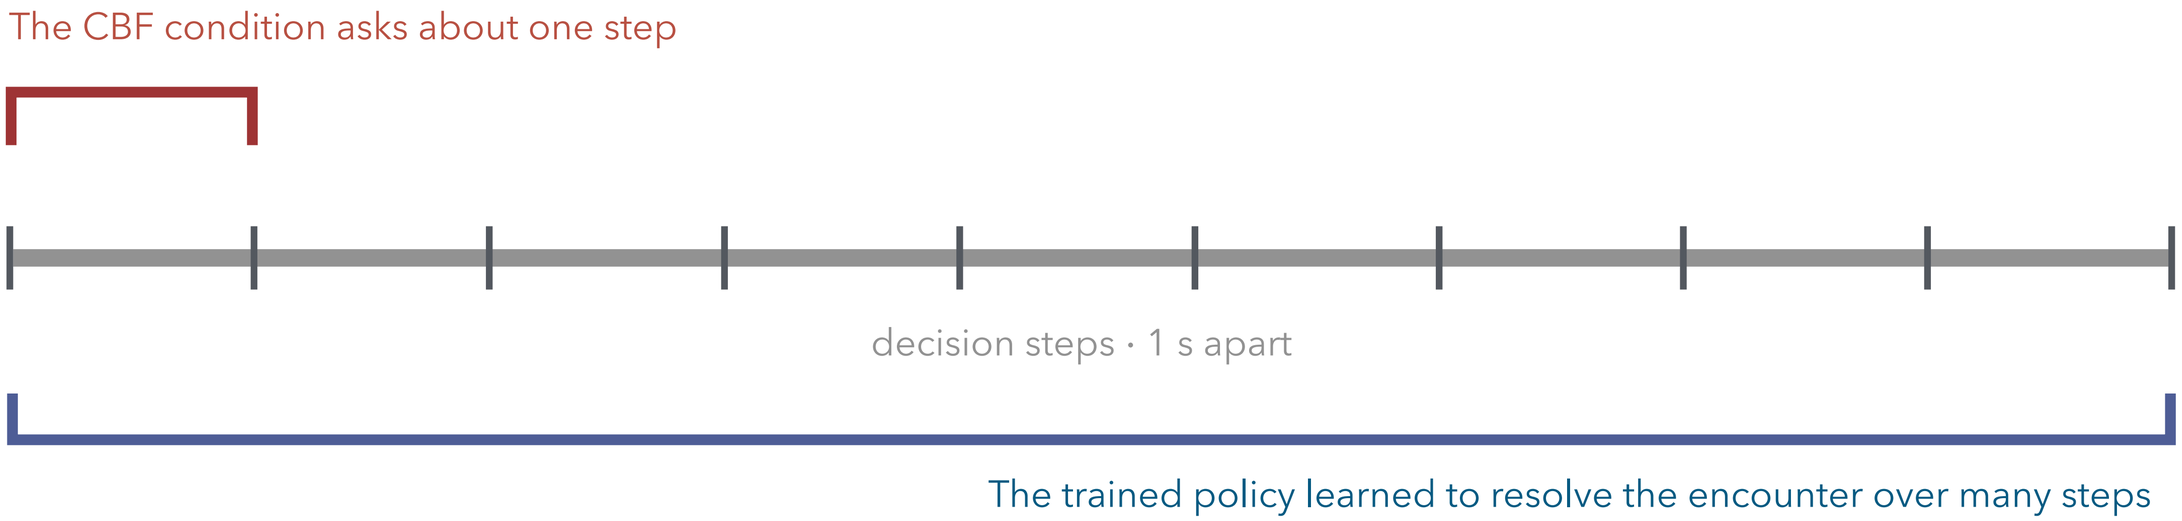

Action filtering replaces multi-step reasoning with a single-step test.

Preserve the trained policy's decision authority

When a learned controller has internalized safety behavior,
informing it outweighs overriding its decision.

ACTION FILTERING

No measurable safety gain; highly sensitive to the filtering assumptions.

OBSERVATION FILTERING

Could be more effective depending on how the policy was trained.

THE COST

Conservatism, e.g., earlier and more frequent slowing down.

DASC 45th
Digital Avionics Systems Conference
Orlando, Florida, USA - September 13-17, 2026

Thank you!

Questions?

Runtime Safety Filtering for Learned Small UAS Separation Policies under GNSS Degradation

Alex Zongo · Peng Wei

Department of Mechanical and Aerospace Engineering
George Washington University, Washington, DC, USA.

Contact: a.zongo@gwu.edu

Preprint: arXiv:2607.10014