Accepted for publication in the 2026 IEEE Cyber Awareness and Research Symposium (CARS).
© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Task-State Integrity Attacks and Defenses for Fixed-Route Autonomous Ice Resurfacing

Michael Hajostek
School of Electrical Engineering & Computer Science, University of North Dakota, Grand Forks, ND, USA
michael.hajostek@und.edu orcid.org/0009-0002-8445-0801

Tingjun Lei
tingjun.lei@und.edu orcid.org/0000-0002-9043-5600

***Abstract*—Fixed-route robotic service systems separate navigation from service application. A platform follows a repeated coverage route while adapting a scalar service command from a spatial task-state map, and when that map is corrupted, service quality can degrade even though the route stays safe, deterministic, and collision-free. Task-state integrity is thus a cyber-physical concern distinct from navigation safety. The concern is studied through a reproducible evaluation framework instantiated for autonomous ice resurfacing, in which a resurfacer computes water commands from an observed surface-demand map that an attacker may corrupt. False-positive, false-negative, combined, and kernel-aware map corruptions are evaluated against a naive controller, an uncertainty-aware Bayesian controller, and several low-cost filters. The naive attacks produce measurable excess service, missed service, and readiness delay. Bayesian smoothing mitigates the naive attacks. Kernel-aware false positives yield greater absolute regional excess for Bayesian than naive allocation under budget-matched controls, but smaller attack-induced increases, exposing a distinction between service quality and attack sensitivity. The low-cost filters expose complementary tradeoffs rather than a dominant defense. These results indicate that adaptive robotic service systems must protect task-state integrity in addition to safe coverage and navigation.**

***Index Terms*—Autonomous ice resurfacing, cyber-physical systems, false-data injection, task-state integrity**

I. Introduction

Fixed-route robotic service systems often separate *where* a robot travels from *what* service it applies. A platform follows a repeated coverage route while adapting a scalar service command from a spatial task-state map that encodes how much service each location needs. Prescription-map variable-rate agriculture provides a concrete instance: management zones, position, and sensed field state determine spatially varying application rates [1]–[3]. In such systems it is the task-state map, not the route, that sets the service intensity, so a corrupted map can degrade service quality even when navigation remains safe, deterministic, and collision-free. Task-state integrity is therefore a cyber-physical concern distinct from navigation safety, since the robot may reach the right places while applying the wrong service.

Autonomous and instrumented ice resurfacing is an underserved candidate instance of this problem class. Public project and commercial sources show relevant building blocks—autonomous or patterned coverage, connected machine data, and position-linked ice-state sensing—but do not establish that any deployed system implements the exact demand-map water controller studied here [4], [5]. This paper therefore uses ice resurfacing as a benchmark instantiation rather than claiming a vulnerability in a specific product. When a controller trusts a corrupted surface-condition map, the two failure modes are concrete and opposite. A false positive marks a low-demand region as needing service, causing unnecessary water application, local over-service, and readiness delay while the surface re-freezes. A false negative masks a genuinely high-demand region, so the controller under-waters it and leaves a degraded patch behind. Neither failure requires tampering with the route, the vehicle dynamics, or the actuator, since corrupting the map alone is sufficient.

This map-to-service abstraction and its integrity risk connect several lines of work. Fixed-route service is naturally modeled as coverage path planning [6]–[9], and the map-driven command follows precision-agriculture variable-rate application, in which zone maps and prescription files set spatially varying application rates [1]. The integrity threat follows cyber-physical false-data injection, where corrupted observations change downstream control actions without compromising actuator logic [10]–[12], while precision-agriculture security work identifies analogous threats to connected farm systems and application equipment [13], [14]. Its robotic form includes injected or removed perceived obstacles, spoofed sensor inputs, and secure mapping under manipulated localization or ranging data [15]–[20]. A recurring lesson in secure control is that a defense must be evaluated against an attacker who understands the estimator or filter, not only against random noise [21], which motivates the kernel-aware stress test for task-state maps. Finally, ice-rink studies show that resurfacing water volume and transient heat transfer govern recovery time and ice quality [22], [23], motivating the service and readiness proxies without requiring a full thermodynamic model.

To isolate task-state integrity from navigation, the coverage route is held fixed and deterministic, leaving the observed surface-demand map as the only adaptive input to the water command. This setting is realized as a reproducible ice-resurfacing framework and used to evaluate four map-corruption families, namely false positive (FP), false negative (FN), combined, and a smoother-aware FP, against a naive controller, an uncertainty-aware Bayesian controller, and a set of low-cost filters. Bayesian smoothing reduces the tested attack-induced changes, but kernel-aware inputs yield greater absolute regional excess than naive allocation. Clean regional baselines explain why these controller orderings differ, even under matched perturbation budgets. The main contributions are summarized as follows.

(1) A reproducible evaluation framework for task-state integrity is proposed, in which an ice resurfacer follows a fixed coverage route and computes water commands from an observed surface-demand map, and the impact of a map attack is measured against ground-truth demand through a matched no-attack baseline.
(2) An uncertainty-aware Bayesian water-allocation controller is developed that maintains a per-cell demand belief along the route footprint and forms an upper-confidence acquisition surface, so that isolated map spikes are damped before they reach the water command.
(3) A kernel-aware false-positive attack is proposed that shapes the injected demand to the controller's smoothing scale, with peak-matched and L_1-matched budget controls that compare input shape without increasing the corresponding perturbation budget.

II. Problem Scope and Threat Model

The framework isolates fixed-route variable-rate service: the route is deterministic, and only the scalar service command adapts. In the ice-resurfacing instantiation this command represents water allocation along a repeated coverage backbone. Holding the route fixed separates task-state integrity from navigation safety and makes map-induced service errors directly observable, to test whether corrupted task-state maps degrade service despite following the intended route.

A. Formal Problem Definition

Let G denote the set of in-rink grid cells, each with area ΔA. A scenario contains a true normalized surface-demand map $d_i \in [0, d_{\max}]$ for each cell $i \in G$, along with a high-demand mask $H \subseteq G$ used for under-service evaluation. The resurfacer follows a fixed route $p_{1:T}$, where $B_t \subseteq G$ is the brush or water-footprint support around waypoint p_t. A water-allocation policy f maps an observed task-state map $\tilde{d}$ to a water-command sequence $w_{1:T} = f(\tilde{d}, p_{1:T})$, where each command is bounded by actuator limits and by a mean-water budget. The resulting applied service at cell i is modeled as

$$s_i(w) = \sum_{t=1}^{T} \eta_{it} w_t,$$

where $\eta_{it} = 0$ when $i \notin B_t$ and otherwise encodes the local footprint contribution. Evaluation is always against the true demand d, not the observed map $\tilde{d}$.

An attack is a map transform $a(\cdot)$ applied before water allocation, so that the controller receives $\tilde{d} = a(d + \epsilon)$, where ϵ is observation noise. The headline experiments set $\epsilon = 0$ to isolate attack effects. Let $\tilde{d}_0 = d + \epsilon$ denote the unattacked observation and $\tilde{d}_a = a(d + \epsilon)$ the attacked observation. For a metric m on the same spatial support in both runs, paired attack impact is

$$\Delta m_{a,f} = m(d, \tilde{d}_a, f(\tilde{d}_a, p_{1:T})) - m(d, \tilde{d}_0, f(\tilde{d}_0, p_{1:T})),$$

so that each attacked run is compared against a no-attack run with the same seed, demand field, route, and policy. Reported values are seed-level means with 95% confidence intervals (CIs); absolute FP outcomes are distinguished below.

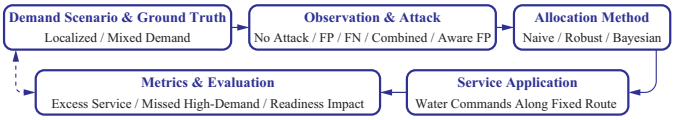


Fig. 1. Overview of the task-state integrity evaluation pipeline. The selected allocator maps a clean or attacked surface-demand observation to water commands along the fixed route; resulting service is evaluated against ground truth via the dashed path.

Three metrics instantiate m. The primary false-positive metric is integrated over-service on the fake region,

$$E_F(w) = \Delta A \sum_{i \in F} [s_i(w) - d_i]_+,$$

which, because d_i and s_i are normalized service intensities, is a square-meter-equivalent excess-service quantity rather than a calibrated physical water volume. FP tables and figures report absolute attacked-run $E_F(w_a)$; absolute recovery is the relative reduction from naive attacked-run E_F. One m^2-eq. represents one unit of excess normalized service integrated over one square meter. Attack-induced FP change instead uses $\Delta E_F = E_F(w_a) - E_F(w_0)$, applying the attack-defined region F to both service outcomes after allocation. A clean run's empty fake-region mask cannot provide that regional baseline. Paired FP recovery is the relative reduction in mean ΔE_F from naive allocation. FN and readiness results report paired changes. The primary false-negative metric is thresholded missed high-demand area,

$$M_H(w) = \Delta A \sum_{i \in H} \mathbf{1}\{d_i - s_i(w) > \tau\},$$

and the readiness proxy is

$$R(w) = \max\left(0, \frac{L(p_{1:T})}{v} + \lambda \max_{i \in G} [s_i(w) - d_i]_+ - T_{\text{win}}\right),$$

where $L(p_{1:T})/v$ is route time, λ is the drying-delay weight, and T_{win} is the operational readiness window.

B. Map-Integrity Threat Model

The attacker corrupts the task-state map consumed by the water-allocation controller. The corruption could be realized by spoofed sensors, a compromised sensor module, modified map-generation software, or a data-pipeline compromise. Following prior false-data-injection and robotics perception-attack work [10], [15], [19], these mechanisms are abstracted as corruption of the surface-demand map before water allocation, and the resulting threat model is summarized in Table I.

Four attack families are considered. *False positive (FP)* injects a fake high-demand region, causing unnecessary water. *False negative (FN)* suppresses a true high-demand region, causing under-service. *Combined* adds a fake region and masks a true region at once. *Aware FP* tunes a false-positive pattern to the Bayesian smoother kernel scale. The aware-FP case tests inputs shaped to the known spatial response of the smoother. This mirrors broader secure-control and adversarial-planning

TABLE I
MAP-INTEGRITY THREAT MODEL USED IN THE FRAMEWORK.

Dimension	Assumption
Attacker knowledge	Knows the rink grid, fixed route, true demand for target selection, and map-to-water control class; the aware-FP case also knows the Bayesian smoother scale.
Attacker capability	Can modify the observed surface-demand map before water allocation.
Fixed by assumption	Route, vehicle dynamics, physical rink geometry, mission budget, actuator limits, and evaluation masks.
Attack targeting	FP targets are sampled from low-demand regions and FN targets from high-demand regions in the generated scenario. This is a targeted stress test, not a blind attacker model.
Non-goal	The framework evaluates map-integrity risk in a fixed-route service abstraction; it does not attribute a vulnerability to any specific deployed resurfacer.

settings in which a defense must be evaluated against attackers who understand the estimator, detector, or planner structure [24], [25].

III. FRAMEWORK AND EXPERIMENTAL DESIGN

A. Rink, Demand, and Observation Maps

Fig. 1 summarizes the end-to-end evaluation flow. Each run generates a synthetic rink-scale demand field on a grid mask with rounded rink geometry, where the true demand map represents where surface service is needed. Two regimes are used, namely *localized*, with compact demand patches, and *mixed*, which combines local and broader structure. The observed demand map is produced by adding noise and then applying the selected attack transform, with additive noise set to zero in the headline runs to isolate attack effects while the operators still support noise for sensitivity checks. The water-allocation controller sees only the observed map, and evaluation is always against the true map. Fig. 2 shows a representative attacked map together with the resulting excess and missed service.

B. Water-Allocation Methods

All water-allocation methods share the same fixed route, water-on mask, actuator limits, and mean-water budget, and differ mainly in the map or acquisition surface used to form the scalar command. For a map-based score q_i, the score is sampled along the route, normalized by a demand scale S, and mapped to a bounded command through a clipped affine rule with limits $w_{\min}$ and $w_{\max}$. Commands are zeroed where the path has water off, and active commands are scaled down uniformly if their mean exceeds the budget $\bar{w}_{\max}$. Controller and budget parameters are listed in Table II.

The naive controller uses the corrupted map directly, $q_i = \tilde{d}_i$. The fixed-constant baseline ignores the map and applies a constant command on water-on points. The `robust_water` reference multiplies observed demand by a local spatial-consistency and confidence cue, which downweights isolated spikes, and is retained as a lightweight baseline rather than a tuned defense.

The Bayesian controller maintains a per-cell Gaussian demand belief with mean μ_i and variance σ_i^2, initialized from a

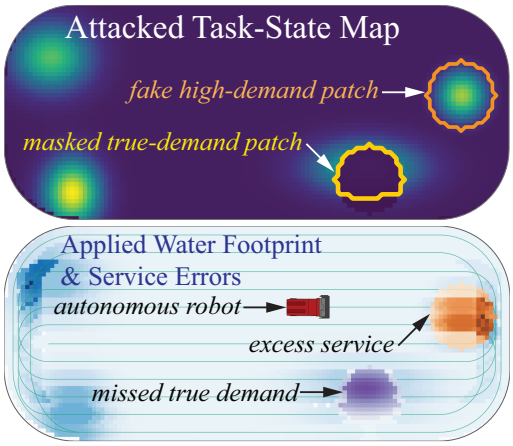


Fig. 2. Qualitative example of a task-state attack on fixed-route water allocation. The route is fixed, and the attack changes the observed surface-demand map used to compute water commands. A false positive (orange) induces excess service on a fake patch, while a masked true-demand (yellow) region produces missed service. The red rectangle in the applied water footprint is the ice resurfacer, shown for visualization.

fixed prior and updated along the route footprint by a standard scalar Kalman step for cells inside the brush footprint B_t. A separable Gaussian smoothing prior is then applied within the rink mask and the variance is inflated to avoid overconfidence. The allocator forms an upper-confidence acquisition surface

$$A_i = \mu_i + \kappa\sigma_i,$$

and maps its route-footprint average to a water command under the shared limits and budget. This design damps isolated map spikes, since a lone high observation raises μ_i only in proportion to the Kalman gain and is then spread by the smoothing prior, whereas a genuine high-demand region supported by its neighbors survives both steps.

Simple-defense variants are label-free transforms applied before allocation. Percentile clipping caps extreme observed values, median- and morphology-style filters suppress small unsupported regions, local-consistency weighting downweights cells unsupported by their neighborhood, and a trust-region variant caps the naive schedule relative to a nominal constant schedule. These are tradeoff probes, not optimized defenses.

C. Attacks, Protocol, and Defense Calibration

The core factorial design varies demand regime, attack condition, and water-allocation method. The demand regimes are localized and mixed, and the attack conditions are no attack, FP, FN, combined, and aware FP. All headline comparisons use paired random seeds. For each seed and regime, clean and attacked runs share the same true demand field, route, and allocation policy. Attack-induced changes are computed within each seed before averaging, with FP changes evaluated on the same attack-defined region in both runs. FP tables and figures instead report absolute attacked-region excess, unless explicitly identified as paired changes. Each headline cell uses 90 seeds and reports means with 95% confidence intervals computed from the corresponding seed-level quantities.

TABLE II
SIMULATION, ATTACK, AND CONTROLLER PARAMETERS.

Component	Parameter	Value used in experiments
Grid	Rink/resolution	Rounded National Hockey League (NHL)-style rink mask with $0.5\ m$ cells.
Localized demand	Compact patches	3 Gaussian patches; background demand 0.05; amplitude $U(0.8, 1.5)$; $\sigma_x = U(2.0, 5.0)\ m$; $\sigma_y = U(1.5, 3.5)\ m$; clipped to 1.8.
Mixed demand	Broad patches plus skate proxy	2 broad Gaussian patches; background demand 0.03; broad amplitude $U(0.3, 0.9)$; $\sigma_x = U(4.0, 8.0)\ m$; $\sigma_y = U(2.0, 5.0)\ m$; normalized skate-cut component from 14 simulated skaters, 130 steps, and cut depths from -3.0 to $-0.4\ mm$, weighted by 0.8; clipped to 1.8.
Observation	Additive noise	0.0 in headline runs to isolate attack effects.
Default attack	FP/FN budget	Circular support radius $4.0\ m$; FP amplitude 1.2; FN multiplier 0.05 inside the targeted patch; combined attack applies both FP and FN transforms.
Aware attack	Kernel match and budget modes	Gaussian FP matched to the Bayesian smoother scale, $\sigma = 0.75\ m$, with the same $4.0\ m$ support radius. The stress case uses a $2\times$ peak factor; budget controls include peak matching at amplitude 1.2 and L_1 matching to the normal FP perturbation.
Dose response	Sweep grid	FP radius $\{2.0, 3.0, 4.0, 5.5, 7.0\}\ m$ and amplitude $\{0.6, 0.9, 1.2, 1.5, 1.8\}$; the default attack budget is the $4.0\ m$, 1.2-amplitude cell.
Allocator	Command mapping	Demand scale $S = 1.2$; $w_{\min} = 0.10$; $w_{\max} = 1.0$; mean-water budget $\bar{w}_{\max} = 0.45$; fallback demand scale is the in-rink 95th percentile.
Bayesian controller	Belief and acquisition	Prior mean 0.30; prior variance 0.10; observation variance 10^{-3}; smoothing scale $0.75\ m$; variance inflation 0.01; acquisition weight $\kappa = 1.0$.

The FP dose-response experiment varies attack amplitude and patch radius over the grid in Table II, testing whether the qualitative conclusion holds across weaker and stronger fake patches and exposing differences in absolute under-attack excess. The aware-FP experiment replaces the generic fake patch with a Gaussian profile matched to the Bayesian smoother scale. Besides the original $2\times$ stress condition, two budget-matched controls are reported, namely a peak-matched aware attack with the same peak amplitude as the normal FP attack, and an L_1-matched aware attack whose positive perturbation mass matches the normal FP attack for the same paired seed, regime, observed map, center, radius, and amplitude. These controls compare input shapes at matched perturbation budgets.

The simple-defense experiment fixes one parameterization per defense family on calibration seeds 90-149 and reports held-out impacts on seeds 0-89, matching the 90-seed pool used elsewhere. It evaluates cheap map transforms such as percentile clipping, median-style filtering, and trust-region caps against the naive, robust, and Bayesian methods. Intervals use seed-level absolute FP outcomes or paired changes, as specified.

IV. EXPERIMENT RESULTS

A. Naive Attacks Have Operational Effects

Table III summarizes the primary single-mode 90-seed results. Under FP attack, naive allocation produces $17.58\ m^2$-eq. of absolute regional excess in localized maps and $18.76\ m^2$-eq. in mixed maps. The matched-region increases are 15.19 ± 0.62 and $14.35 \pm 0.71\ m^2$-eq.; readiness impacts are 6.15 and $11.41\ s$, respectively. FN attacks add $6.74\ m^2$ of missed high-demand area in localized maps and $5.38\ m^2$ in mixed maps. Combined FP+FN corruption preserves both harms: Bayesian allocation reduces localized absolute FP excess from 17.58 ± 0.75 to $12.55 \pm 0.60\ m^2$-eq. and paired missed area from 6.97 ± 1.18 to $2.06 \pm 0.52\ m^2$. In the mixed regime, it reduces the corresponding quantities from 18.72 ± 0.90 to $13.96 \pm 0.77\ m^2$-eq. and from 5.63 ± 1.15 to $1.84 \pm 0.68\ m^2$, respectively.

TABLE III
NINETY-SEED MEANS ± 95% CIS. FP EXCESS/RECOVERY ARE ABSOLUTE; FN MISSED AREA/RECOVERY AND READINESS CHANGES ARE PAIRED. RECOVERY IS RELATIVE TO NAIVE ALLOCATION.

Regime	Metric	Naive	Bayes	Rec.
Localized	FP excess (m^2-eq.)	17.58 ± 0.75	12.42 ± 0.59	**29.3%**
	FP ready (s)	6.15 ± 4.41	-0.19 ± 0.20	–
	FN missed (m^2)	6.74 ± 1.03	1.58 ± 0.42	**76.5%**
Mixed	FP excess (m^2-eq.)	18.76 ± 0.90	13.94 ± 0.75	**25.7%**
	FP ready (s)	11.41 ± 5.72	0.02 ± 0.08	–
	FN missed (m^2)	5.38 ± 1.09	2.49 ± 0.56	**53.8%**

B. Bayesian Smoothing Provides Partial Mitigation

Bayesian water allocation reduces absolute FP excess by 29.3% in localized maps and 25.7% in mixed maps. Matched-region paired FP recovery is 95.8% and 92.5%, respectively. It recovers 76.5% of localized and 53.8% of mixed FN missed area and nearly eliminates the naive FP readiness penalty. These results support mitigation, but attack sensitivity alone does not establish service quality.

A paired 90-seed observation-noise sweep over $\sigma \in \{0.00, 0.02, 0.05, 0.10\}$ preserves positive Bayesian absolute-FP reduction (25.4%–29.3%) and paired FN recovery (55.5%–84.8%) across nonzero noise levels in both regimes. These comparisons are therefore not restricted to noise-free observations. Fig. 3 shows greater absolute Bayesian excess at some low FP amplitudes, with more consistent absolute reduction at higher amplitudes.

TABLE IV
L_1-MATCHED KERNEL-AWARE FP: 90-SEED MEAN EXCESS (m^2-EQ.) ON THE SAME REGION F, WITH 95% CIS FOR PAIRED INCREASES.

Regime	Method	Clean	Attacked	Paired increase
Localized	Naive	2.39	11.43	9.04 ± 0.52
	Bayesian	11.79	12.54	0.75 ± 0.44
Mixed	Naive	4.41	13.08	8.67 ± 0.58
	Bayesian	12.87	13.96	1.09 ± 0.37

C. Kernel-Aware FP Changes Under-Attack Rankings

The kernel-aware FP experiment tests inputs shaped to the smoother scale. Fig. 4 sweeps the 2× stress condition; Table V adds peak- and L_1-matched controls. Absolute Bayesian excess exceeds naive excess even with matched budgets (L_1-matched Bayes-minus-naive: localized $+1.11 \pm 0.64$, mixed $+0.88 \pm 0.53$). This absolute ordering does not establish greater attack-induced harm.

Table IV separates clean regional excess from the attack-induced increase at the headline 4.0 m, 1.2-amplitude setting. Across stress and matched-budget aware variants at this setting, Bayesian paired recovery is 87.5%–96.7%. Bayesian therefore has smaller attack-induced increases despite greater absolute excess.

D. Simple Defenses Expose Tradeoffs

No single low-cost defense dominates across attack modes. Fig. 5 and Table VI compare defenses calibrated on seeds 90-149 and evaluated on held-out seeds 0-89.

Clipping observed demand at the 90th percentile reduces absolute FP and aware-FP excess but increases paired FN misses, because it suppresses the high-demand evidence needed to recover masked true demand. A trust-region cap with $\Delta = 0.05$ improves FN behavior and is competitive with Bayesian water

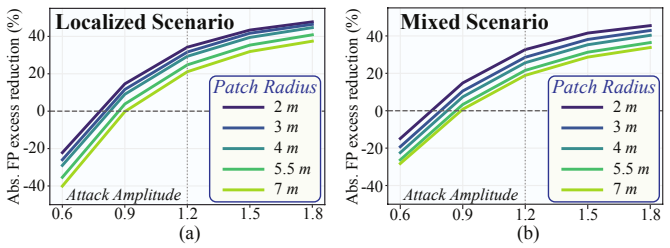


Fig. 3. False-positive dose response over 90 seeds. Markers show attack budgets; lines connect equal-radius amplitude sweeps. Absolute recovery compares attacked-run E_F, not attack-induced change; the dashed line marks zero.

TABLE V
KERNEL-AWARE FP: ABSOLUTE ATTACKED-REGION EXCESS (MEAN ± 95% CI; m^2-EQ.).

Regime	Variant	Naive	Bayes	Abs. rec.
Localized	Normal FP	17.58 ± 0.75	12.42 ± 0.59	29.3%
	Aware stress 2×	8.35 ± 0.45	12.05 ± 0.53	**-44.4%**
	Aware peak-matched	6.02 ± 0.32	11.91 ± 0.50	**-97.7%**
	Aware L_1-matched	11.43 ± 0.63	12.54 ± 0.69	**-9.7%**
Mixed	Normal FP	18.76 ± 0.90	13.94 ± 0.75	25.7%
	Aware stress 2×	9.92 ± 0.59	13.22 ± 0.65	**-33.2%**
	Aware peak-matched	7.75 ± 0.48	13.04 ± 0.62	**-68.2%**
	Aware L_1-matched	13.08 ± 0.78	13.96 ± 0.78	**-6.7%**

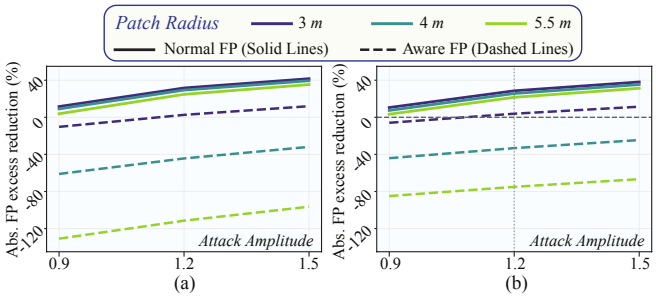


Fig. 4. Kernel-aware FP sensitivity over 90 seeds. Solid/dashed lines show normal/aware FP. Absolute recovery compares attacked-run E_F; negative values mean greater Bayesian excess, not a paired-effect reversal.

TABLE VI
HELD-OUT SIMPLE-DEFENSE COMPARISON (MEAN ± 95% CI). FP EXCESS IS ABSOLUTE; FN IS PAIRED. TRUST CAP $\Delta = 0.05$.

Regime	Method	FP ex. (m^2-eq.)	FP Abs. rec.	Aware-FP ex. (m^2-eq.)	FN miss (m^2)
Localized	naive	17.58 ± 0.75	0.0%	8.35 ± 0.45	6.74 ± 1.03
	Bayesian	$\mathbf{12.42 \pm 0.59}$	**29.3%**	12.05 ± 0.53	1.58 ± 0.42
	clipped p90	12.52 ± 0.69	28.8%	$\mathbf{5.50 \pm 0.37}$	11.52 ± 1.30
	trust cap	13.43 ± 0.59	23.6%	11.81 ± 0.53	$\mathbf{1.47 \pm 0.22}$
Mixed	naive	18.76 ± 0.90	0.0%	9.92 ± 0.59	5.38 ± 1.09
	Bayesian	13.94 ± 0.75	25.7%	13.22 ± 0.65	2.49 ± 0.56
	clipped p90	$\mathbf{12.47 \pm 0.72}$	**33.5%**	$\mathbf{6.56 \pm 0.41}$	6.37 ± 1.31
	trust cap	14.57 ± 0.69	22.3%	13.07 ± 0.64	$\mathbf{0.63 \pm 0.18}$

for FN, though it is not the strongest FP defense. Robust water is often strong in these absolute-FP summaries but remains a local-filtering reference baseline, not a validated defense.

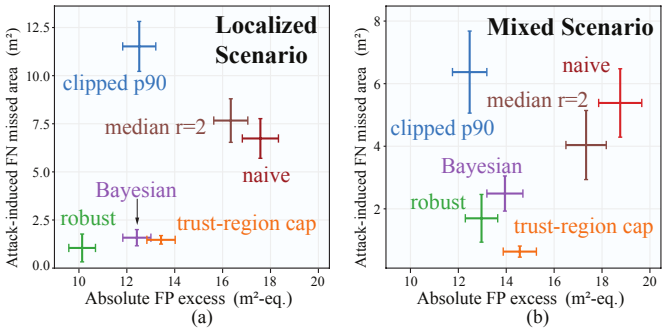


Fig. 5. Simple-defense tradeoff: calibration seeds 90-149; held-out seeds 0-89. Axes show absolute attacked-region FP excess and paired FN missed-area change. Lower-left is better for these quantities.

V. DISCUSSION

The results show that task-state integrity matters even with fixed navigation: a corrupted map changes service along an unchanged route. Bayesian smoothing reduces the tested paired effects, while clean regional baselines explain the apparent kernel-aware defense reversal. These results do not establish exploitation of the smoother. The fixed-constant baseline illustrates why sensitivity is not service quality: its clean and attacked excess are identical (localized 12.79 ± 0.57, mixed 13.95 ± 0.68 m^2-eq. at the headline FP support), giving zero attack-induced change despite nonzero excess. Simple filters still expose mode-specific tradeoffs: clipping reduces absolute FP excess but worsens FN misses, while trust caps better protect masked demand. These findings motivate evaluating

both service quality and corruption sensitivity, alongside secure telemetry and operator visibility.

These findings are simulation-based. Bayesian hyperparameters are locked; component ablations are needed to distinguish smoothing, belief updates, uncertainty, and water-budget effects. The simple defenses are label-free probes, not tuned detectors, and targeted white-box attacks are stress tests rather than deployed-system guarantees. Regional FP excess is not a complete utility measure. Service and readiness are operational proxies: real freezing depends on water volume, heat transfer, operating history, and facility conditions [22], [23], [26], [27], so magnitudes are comparative.

VI. Conclusion

This paper studied the task-state integrity risk that arises when water allocation depends on a surface-demand map. In a fixed-route framework, false-positive and false-negative map attacks produced measurable excess service, missed service, and readiness impact, and Bayesian smoothing mitigated the tested paired effects. Kernel-aware false positives produced greater absolute Bayesian excess than naive allocation under matched budgets, but smaller attack-induced increases; simple defenses revealed mode-specific tradeoffs. The practical lesson is that adaptive resurfacing systems should treat surface-condition maps as security-critical task state even when the route is fixed and deterministic. Future work should calibrate the service and readiness proxies against real resurfacer telemetry, study adaptive or randomized smoother scales and detector-aware attackers, and extend the analysis from the fixed-route setting toward joint route-and-service planning.

Acknowledgment

This material is based upon work supported by the National Science Foundation CISE Graduate Fellowship under Grant No. 2313998. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation. The authors would like to thank Minara Sooriyaarachchi, Dr. Sicong Shao, Dr. Jun Chen, and Dr. Chaomin Luo for providing valuable feedback. OpenAI Codex helped edit and review throughout the paper; the authors verified all outputs and remain responsible for the work.

References

[1] J. Campos, M. Gallart, J. Llop, P. Ortega, R. Salcedo, and E. Gil, "On-farm evaluation of prescription map-based variable rate application of pesticides in vineyards," *Agronomy*, vol. 10, no. 1, p. 102, Jan. 2020.

[2] C. Song *et al.*, "Variable-rate control system for UAV-based granular fertilizer spreader," *Comput. Electron. Agric.*, vol. 180, p. 105832, Jan. 2021.

[3] S. A. O'Shaughnessy, S. R. Evett, and P. D. Colaizzi, "Dynamic prescription maps for site-specific variable rate irrigation of cotton," *Agric. Water Manag.*, vol. 159, pp. 123–138, Sep. 2015.

[4] R. Shah, N. Carcione, Y. Cai, J. Qiu, and K. Shen, "Autonomous Zamboni convoy: final report," Carnegie Mellon University, Pittsburgh, PA, MRSD Capstone Project Final Report, Dec. 2022. [Online]. Available: https://mrsdprojects.ri.cmu.edu/2022teami/documents/

[5] WM technics, "The Ice Master innovation: catalogue Mammoth," WM technics Srl, Product Catalogue, 2026. [Online]. Available: https://www.wm-technics.com/wp-content/uploads/2024/10/Catalogue-Mammoth-2026-en-low-1.pdf

[6] M. Höffmann, S. Patel, and C. Büskens, "Optimal guidance track generation for precision agriculture: a review of coverage path planning techniques," *J. Field Robot.*, vol. 41, no. 3, pp. 823–844, May 2024.

[7] T. Lei, P. Chintam, C. Luo, and S. Rahimi, "Multi-robot directed coverage path planning in row-based environments," in *Proc. IEEE 5th Int. Conf. Artif. Intell. Knowl. Eng. (AIKE)*, Sep. 2022, pp. 114–121.

[8] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robot. Auton. Syst.*, vol. 61, no. 12, pp. 1258–1276, Dec. 2013.

[9] T. Lei, C. Luo, G. Jan, and Z. Bi, "Deep Learning-Based Complete Coverage Path Planning With Re-Joint and Obstacle Fusion Paradigm," *Front. Robot. AI*, vol. 9, pp. 1–14, Mar. 2022.

[10] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Trans. Inf. Syst. Secur.*, vol. 14, no. 1, pp. 1–33, May 2011.

[11] A. Teixeira, K. C. Sou, H. Sandberg, and K. H. Johansson, "Secure control systems: a quantitative risk management approach," *IEEE Control Syst. Mag.*, vol. 35, no. 1, pp. 24–45, Feb. 2015.

[12] M. Ahmed and A.-S. K. Pathan, "False data injection attack (FDIA): an overview and new metrics for fair evaluation of its countermeasure," *Complex Adapt. Syst. Model.*, vol. 8, no. 1, p. 4, Apr. 2020.

[13] A. Yazdinejad *et al.*, "A review on security of smart farming and precision agriculture: security aspects, attacks, threats and countermeasures," *Appl. Sci.*, vol. 11, no. 16, p. 7518, Jan. 2021.

[14] M. Freyhof, G. Grispos, S. Pitla, and W. Mahoney, "Investigating the implications of cyberattacks against precision agricultural equipment," *Int. Conf. Cyber Warfare Secur.*, vol. 20, no. 1, pp. 93–104, Mar. 2025.

[15] J. Banfi, Y. Zhang, G. E. Suh, A. C. Myers, and M. Campbell, "Path planning under malicious injections and removals of perceived obstacles: a probabilistic programming approach," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 6884–6891, Oct. 2020.

[16] A. Szvoren, J. Liu, D. Kanoulas, and N. Tuptuk, "Exploring adversarial obstacle attacks in search-based path planning for autonomous mobile robots," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2025, pp. 14 843–14 849.

[17] D. Davidson, H. Wu, R. Jellinek, V. Singh, and T. Ristenpart, "Controlling UAVs with sensor input spoofing attacks," in *10th USENIX Workshop on Offensive Technologies (WOOT 16)*, 2016.

[18] J. Wang, F. Li, X. Zhang, and H. Sun, "Adversarial obstacle generation against LiDAR-based 3D object detection," *IEEE Trans. Multimedia*, vol. 26, pp. 2686–2699, 2024.

[19] A. Devanga, T. Lei, J. Hu, J. Chen, and C. Luo, "Trustworthy cyber-resilient reinforcement learning for secure navigation under adversarial attack," in *Proc. 2025 Cyber Awareness Res. Symp. (CARS)*, Oct. 2025, pp. 1–6.

[20] S. Steen, T. Lei, C. Luo, G. Li, and S. Shao, "Fail-safe multi-robot awareness during GPS outages using CNN-based LiDAR," in *Proc. 2025 Cyber Awareness Res. Symp. (CARS)*, Oct. 2025, pp. 1–7.

[21] G. Chen, Y. Zhang, S. Gu, and W. Hu, "Resilient state estimation and control of cyber-physical systems against false data injection attacks on both actuator and sensors," *IEEE Trans. Control Netw. Syst.*, vol. 9, no. 1, pp. 500–510, Mar. 2022.

[22] L. Li, X. Liu, and T. Zhang, "Investigation of heat and mass transfer characteristics in the ice rink: ice making, maintaining and resurfacing processes," *Build. Environ.*, vol. 196, p. 107779, Jun. 2021.

[23] Z. Zhang *et al.*, "Heat and mass transfer during ice resurfacing process in ice rink," *Build. Environ.*, vol. 252, p. 111270, Mar. 2024.

[24] Q. Dinh Vu, R. Tan, and D. K. Y. Yau, "On applying fault detectors against false data injection attacks in cyber-physical control systems," in *Proc. IEEE INFOCOM*, Apr. 2016, pp. 1–9.

[25] G. Franzè, D. Famularo, W. Lucia, and F. Tedesco, "Cyber–physical systems subject to false data injections: a model predictive control framework for resilience operations," *Automatica*, vol. 152, p. 110957, Jun. 2023.

[26] Z. Zhang, S. Wang, Y. Wang, R. Liu, L. Chang, and H. Wang, "Transient heat and mass transfer characteristics during multiple resurfacing operations in ice rink," *Appl. Therm. Eng.*, vol. 274, p. 126627, Sep. 2025.

[27] R. H. S. Hutchins, Y. Li, G. Taylor, D. Loverock, and S. Impellizzeri, "Arena ice quality and perspectives on optimizing performance and addressing emerging challenges," *Sci. Rep.*, vol. 15, no. 1, p. 13600, Apr. 2025.