Selection-Induced Contraction of Innovation Statistics in Gated Kalman Filters

Barak Or, *Member, IEEE*

***Abstract*—Validation gating is a fundamental component of classical Kalman-based tracking systems. Only measurements whose normalized innovation squared (NIS) falls below a prescribed threshold are considered for state update. While this procedure is statistically motivated by the chi-square distribution, it implicitly replaces the unconditional innovation process with a conditionally observed one, restricted to the validation event.**

This paper shows that innovation statistics computed after gating converge to gate-conditioned rather than unconditional nominal reference quantities. Under classical linear-Gaussian assumptions, we derive exact expressions for the first- and second-order moments of the innovation conditioned on ellipsoidal gating, and show that gating induces a deterministic, dimension-dependent contraction of the innovation covariance relative to the nominal reference.

The analysis is extended to nearest-neighbor (NN) association, which is shown to act as an additional statistical selection operator. We prove that selecting the minimum norm innovation among multiple in-gate measurements introduces an unavoidable energy contraction, implying that nominal innovation reference statistics cannot be preserved under nontrivial gating and association due to deterministic selection effects, even under perfectly matched linear-Gaussian Kalman filter assumptions. Monte Carlo validation is provided for the analytic truncation and order-statistic results, together with an end-to-end multi-time-step tracking example showing that nominal NIS-based measurement-noise tuning learns a biased measurement-noise scale whereas a gate-aware correction remains close to the true scale.

***Index Terms*—Kalman filtering, validation gating, normalized innovation squared (NIS), nearest-neighbor association, target tracking, innovation statistics, adaptive tuning, normalized estimation error squared (NEES).**

I. Introduction

Kalman-based tracking systems are a cornerstone of modern aerospace, radar, and navigation applications, where reliable state estimation is required under uncertainty. A central role in such systems is played by innovation-based statistics, which are used for measurement validation, data association, consistency monitoring, and adaptive tuning. Among these, the normalized innovation squared (NIS) is particularly attractive due to its simple statistical characterization under nominal linear Gaussian assumptions: in the absence of additional selection mechanisms, the NIS follows a chi-square distribution whose moments provide natural reference values for diagnostic and tuning procedures.

In operational tracking systems, however, innovation statistics are rarely observed in their unconditional form. Prior to data association and state update, measurements are typically subjected to ellipsoidal validation gating, whereby only innovations whose Mahalanobis distance falls below a prescribed threshold are accepted. Validation gating is widely justified on statistical and practical grounds, as it limits computational complexity and suppresses gross outliers. As a result, nearly all innovation-based diagnostics used in practice operate on a post-gate innovation stream rather than on the nominal innovation process assumed in classical Kalman filter theory.

Despite its ubiquity, the statistical consequences of validation gating are rarely treated explicitly. Once gating is applied, the innovations that enter association, filtering, and diagnostic logic are no longer samples from the nominal Gaussian distribution, but from a distribution truncated to the validation region. Consequently, innovation-based statistics computed after gating need not satisfy the classical reference properties associated with the chi-square law. In particular, empirical NIS statistics obtained from accepted measurements may systematically deviate from their nominal expectations even when the underlying system model and noise statistics are perfectly matched.

A closely related issue arises in data association. When multiple measurements fall inside the validation gate, NN association is commonly employed to select a single measurement for the state update. While computationally efficient and widely used, NN association further conditions the observed innovation through order-statistic selection, as it selects the innovation with minimum normalized distance among multiple candidates. The combined effect of validation gating and NN association therefore induces a multi-stage selection process whose impact on innovation statistics is not captured by nominal Kalman filter assumptions.

The objective of this paper is to isolate and characterize the statistical effects induced solely by validation gating and NN association, independent of clutter, false alarms, nonlinear dynamics, or modeling errors. Within the classical linear-Gaussian Kalman filtering framework, we provide an exact characterization of the innovation moments conditioned on ellipsoidal gating and show that validation gating induces a deterministic, dimension-dependent contraction of the innovation covariance relative to the nominal reference. We further show that NN association acts as an additional statistical selection operator and introduces an unavoidable, multiplicity-dependent contraction of innovation energy through order-statistic selection. Together, these results imply that nominal innovation statistics cannot be preserved under nontrivial gating and association, even when all Kalman filter assumptions are satisfied.

The analysis is complemented by a Monte Carlo study with a multi-time-step constant-velocity tracking problem.

B. Or is with metaor artificial intelligence, Haifa, Israel, and with the Google Reichman Tech School, Reichman University, Herzliya, Israel. E-mail: barakorr@gmail.com

This study compares ungated filtering, gated filtering with fixed covariance, nominal NIS-based covariance tuning, and a gate-aware NIS correction, and reports NIS, corrected NIS, NEES, position RMSE, update rate, and final measurement-noise scale.

A. Related Work

Innovation-based statistics are a fundamental component of Kalman filtering theory and practice, where quantities such as the NIS are routinely used for consistency monitoring, measurement validation, and adaptive tuning. Under nominal linear Gaussian assumptions, the statistical properties of the innovation process are well understood and documented in classical estimation references [1]–[5].

In practical tracking systems, innovation statistics are commonly employed in adaptive filtering and noise covariance estimation schemes, often relying on assumed chi-square properties of the NIS [6]–[9]. Such approaches implicitly assume that the observed innovation sequence is representative of the nominal innovation distribution. However, in operational settings, innovation statistics are almost always evaluated after validation gating and data association, conditions under which these assumptions may no longer hold.

Validation gating and NN association are standard components of tracking systems in aerospace, radar, and navigation applications [4], [10]–[12]. Despite their widespread use, the statistical impact of these selection mechanisms on innovation-based diagnostics has received comparatively little explicit theoretical treatment.

While truncation of Gaussian distributions is classical, the implications of mandatory truncation and order-statistic selection on innovation-based consistency diagnostics have not been explicitly characterized in the Kalman tracking literature. In particular, existing analyses typically focus on detection performance or association accuracy, rather than on the induced bias in innovation statistics themselves.

A closely related recent direction is the use of matrix-valued tracking statistics and Wishart-based consistency measures. Forsling *et al.* [13] emphasize that scalar consistency statistics compress vector-valued tracking information, and propose matrix-valued measures for offline and online assessment. The present paper is complementary: it studies how the commonly used scalar NIS itself is altered by mandatory gating and association before any consistency test or adaptive tuning rule sees the data.

Recent work on adaptive and learning-enhanced Kalman filtering further highlights the reliance on innovation statistics for online consistency assessment and tuning [14]–[16]. These methods benefit from a precise understanding of how structural elements of the tracking pipeline, such as gating and association, affect the observed innovation process.

The present work complements the existing literature by providing an explicit statistical characterization of gate-conditioned innovation moments and by showing that NN association introduces an unavoidable order-statistic bias. Unlike prior adaptive or robust filtering approaches, the results here isolate selection-induced effects under ideal modeling assumptions, thereby clarifying fundamental limitations in the interpretation of innovation-based diagnostics.

The main contributions of this paper are summarized as follows:

- Validation gating is interpreted as a statistical conditioning operation, and exact expressions are derived for the first- and second-order moments of the gate-conditioned innovation under linear-Gaussian assumptions.
- It is shown that ellipsoidal gating induces a deterministic contraction of the innovation covariance that depends only on the gate threshold and the measurement dimension.
- NN association is characterized as an order-statistic selection mechanism, and it is proven that this selection introduces an additional, unavoidable contraction of innovation energy, leading to an impossibility result: nominal innovation statistics cannot be preserved under nontrivial gating and association.
- Monte Carlo simulations validate the analytic gate-conditioned and NN order-statistic expressions, including confidence intervals and direct theory-versus-simulation comparisons.
- An end-to-end tracking study shows how post-gate NIS bias propagates into adaptive measurement-noise tuning and affects RMSE and NEES.

The remainder of the paper is organized as follows. Section II introduces the problem formulation and the innovation model. Section III analyzes validation gating as a statistical selection mechanism and derives gate-conditioned innovation moments. Section IV discusses the implications of these results for innovation-based consistency diagnostics. Section V analyzes the additional bias induced by NN association. Section VI presents closed-form two-dimensional results, Monte Carlo validation, and an end-to-end tracking study. Section VII concludes the paper and discusses implications and limitations.

II. Problem Formulation

A. Kalman Filter Innovation Model

The analysis is conducted within the classical linear Gaussian Kalman filtering framework. Throughout the paper, the nominal assumptions are the following: the dynamic and measurement models used by the filter are correctly specified; the process and measurement noises are zero-mean, white, Gaussian, mutually independent, and independent of the initial estimation error; and the covariance matrices used by the filter match the true covariances. These assumptions are stated explicitly to distinguish the structural selection effects studied here from model mismatch, non-Gaussian disturbances, or suboptimal covariance tuning.

Consider the discrete-time linear measurement model

$$z_k = Hx_k + v_k, \tag{1}$$

where $x_k \in \mathbb{R}^n$ denotes the system state, $z_k \in \mathbb{R}^m$ the measurement vector, $H \in \mathbb{R}^{m\times n}$ the measurement matrix, and the measurement noise is zero-mean, white, and Gaussian with covariance R.

Let $\hat{x}_k^-$ and P_k^- denote the predicted state estimate and error covariance produced by a Kalman filter. The innovation is defined as

$$\nu_k \triangleq z_k - H\hat{x}_k^-, \tag{2}$$

with associated innovation covariance

$$S_k \triangleq HP_k^- H^\top + R. \tag{3}$$

Under the nominal assumptions stated above, the innovation sequence $\{\nu_k\}$ is zero-mean, Gaussian, and white. In the sequel, we analyze a single time index and omit the subscript k for notational clarity. Accordingly, we model the innovation as the random vector

$$\nu : \Omega \to \mathbb{R}^m, \qquad \nu \sim \mathcal{N}(0, S), \tag{4}$$

where $S \in \mathbb{S}_{++}^m$ denotes the symmetric positive-definite innovation covariance matrix.

B. Post-Gate Innovation Distribution

Under the nominal linear Gaussian Kalman filter assumptions, the normalized innovation squared (NIS) is defined as

$$Z \triangleq \nu^\top S^{-1} \nu, \tag{5}$$

and follows a chi-square distribution with m degrees of freedom.

Let

$$\mathcal{A} \triangleq \{Z \le \tau\} \tag{6}$$

denote the validation event induced by ellipsoidal gating. All innovation samples that pass the gate are therefore distributed according to the conditional distribution of ν given $\mathcal{A}$, denoted by $\nu \mid \mathcal{A}$. This conditional distribution corresponds to a Gaussian law truncated to the ellipsoidal region

$$\mathcal{E}_\tau \triangleq \{\nu \in \mathbb{R}^m : \nu^\top S^{-1} \nu \le \tau\}$$

rather than the unconditional Gaussian distribution in (4). The statistical properties of this gate-conditioned innovation distribution form the basis of the analysis in the remainder of this paper.

III. Innovation Gating as a Statistical Selection Mechanism

We interpret ellipsoidal gating as a statistical selection mechanism acting on the innovation process. Specifically, gating restricts the observed innovation stream to realizations satisfying the acceptance event $\mathcal{A}$ defined in (6). Consequently, all innovations that enter data association, state update, and diagnostic logic are drawn from the conditional distribution $\nu \mid \mathcal{A}$, rather than from the nominal Gaussian distribution $\mathcal{N}(0, S)$.

A. Distribution of the NIS and Ellipsoidal Gating

Proposition 1 (Distribution of the NIS)**.** *If $\nu \sim \mathcal{N}(0, S)$ with $S \in \mathbb{S}_{++}^m$, then*

$$Z \sim \chi_m^2. \tag{7}$$

This classical result provides the probabilistic basis for ellipsoidal validation gating: selecting the threshold τ as a chi-square quantile ensures

$$\mathbb{P}\{\mathcal{A}\} = P_g \tag{8}$$

under the nominal innovation model.

Once validation gating is applied, however, the innovation process is no longer observed unconditionally, but only through realizations satisfying the acceptance event $\mathcal{A} = \{Z \le \tau\}$. As a result, all innovation-based statistics computed after gating are statistics of a conditionally observed random variable. In the innovation space $\mathbb{R}^m$, the acceptance event $\mathcal{A}$ corresponds to the ellipsoidal region $\mathcal{E}_\tau$. Validation gating therefore implements a deterministic truncation of the innovation distribution to $\mathcal{E}_\tau$, retaining only realizations within a fixed Mahalanobis radius.

B. Gate-Conditioned Innovation Moments

Proposition 2 (Gate-Conditioned Innovation Moments)**.** *Let $\nu \sim \mathcal{N}(0, S)$ with $S \in \mathbb{S}_{++}^m$, and let $\mathcal{A} = \{\nu^\top S^{-1}\nu \le \tau\}$ denote the validation event. Then the gate-conditioned innovation satisfies*

$$\mathbb{E}[\nu \mid \mathcal{A}] = 0, \tag{9a}$$

$$\mathbb{E}[\nu\nu^\top \mid \mathcal{A}] = \gamma(\tau, m)\, S, \tag{9b}$$

where the scalar contraction factor $\gamma(\tau, m)$ is given by

$$\gamma(\tau, m) \triangleq \frac{1}{m}\, \mathbb{E}[Z \mid Z \le \tau]\,, \qquad Z \sim \chi_m^2, \tag{10}$$

and satisfies $0 < \gamma(\tau, m) < 1$.

Proposition 2 shows that ellipsoidal gating preserves the zero mean of the innovation while deterministically shrinking its second-order moment by a scalar factor. The factor depends only on the gate threshold and the measurement dimension, not on the particular innovation covariance matrix.

C. Gate-Conditioned NIS Statistics

An immediate consequence of (9b) is an explicit expression for the mean NIS after gating.

Corollary 1 (Gate-Conditioned Mean NIS)**.** *Under the assumptions of Proposition 2,*

$$\mathbb{E}[Z \mid \mathcal{A}] = m\, \gamma(\tau, m), \tag{11}$$

where $Z = \nu^\top S^{-1} \nu$.

In contrast to the nominal reference value $\mathbb{E}[Z] = m$, the expected NIS computed from accepted measurements reflects the conditioning induced by the validation gate.

IV. Gate-Aware Consistency and Interpretation of Innovation Statistics

The results of Section III establish that ellipsoidal gating alters the statistical properties of the innovation process. In this section, we examine the implications of these results for innovation-based consistency diagnostics, with particular emphasis on NIS.

A. Classical Use of NIS for Consistency Validation

In Kalman-based tracking systems, the normalized innovation Z_k is commonly employed as a diagnostic quantity for filter consistency. Under the nominal Kalman filter assumptions recalled in Section II, the NIS follows a chi-square distribution with m degrees of freedom. As a result, standard practice is to compare instantaneous NIS values to chi-square confidence bounds, or to monitor the empirical mean of $\{Z_k\}$ over time and compare it to the nominal reference value m, as commonly done in online consistency testing [17].

If the empirical mean is significantly below m, the filter is often declared conservative in the innovation space; if it is significantly above m, it is declared inconsistent or underestimating innovation uncertainty. This interpretation is valid for the unconditional innovation sequence, but it is not valid without modification for the accepted post-gate sequence.

B. Post-Gate NIS Is Not Chi-Square

As a direct consequence of the conditioning induced by ellipsoidal gating, the normalized innovation squared no longer follows its nominal distribution. Specifically,

$$Z \mid \mathcal{A} \not\sim \chi^2_m, \tag{12}$$

but instead follows a truncated chi-square law obtained by renormalizing χ^2_m on the interval $[0, \tau]$ [18].

Accordingly, as established in Proposition 2,

$$\mathbb{E}[Z \mid \mathcal{A}] = m\,\gamma(\tau, m), \tag{13}$$

with $\gamma(\tau, m) \in (0, 1)$. The nominal reference value $\mathbb{E}[Z] = m$ therefore ceases to be valid whenever gating is applied.

C. Gate-Aware Interpretation of NIS Statistics

The results above imply that standard NIS diagnostics must be interpreted relative to gate-conditioned reference values. Specifically, when NIS statistics are computed from accepted innovations only, the correct reference mean is given by Corollary 1. Equivalently, a gate-aware normalized statistic is defined as

$$Z_k^{\text{corr}} \triangleq \frac{1}{\gamma(\tau, m)}\, Z_k, \tag{14}$$

which satisfies

$$\mathbb{E}[Z_k^{\text{corr}} \mid \mathcal{A}] = m. \tag{15}$$

This normalization allows the empirical mean of accepted NIS samples to be compared to the classical reference mean, while preserving the existing Kalman filter recursion and gating logic. The correction is a diagnostic and tuning-reference correction; it does not change the innovation used in the state update.

D. Implications for Adaptive Tuning and Diagnostics

Many adaptive noise-tuning and consistency-monitoring schemes rely, either explicitly or implicitly, on innovation covariance estimates or NIS-based statistics [19], and often enforce nominal chi-square innovation behavior as a tuning objective [20].

If such schemes operate on post-gate data without accounting for the conditioning induced by gating, they will systematically underestimate the innovation covariance, may reduce the estimated measurement-noise level, and may tighten the effective validation logic by making valid measurements more likely to be rejected. The resulting behavior is better described as an increased probability of missed updates or rejection of valid measurements, rather than as an increase in false detections. This feedback mechanism arises even under ideal modeling assumptions and is a direct consequence of ignoring the gate-conditioned nature of the observed innovation process.

The contraction factor $\gamma(\tau, m)$ admits intuitive limiting behavior. As $\tau \to \infty$, validation gating becomes inactive and $\gamma(\tau, m) \to 1$, recovering the nominal innovation statistics. Conversely, as $\tau \to 0$, only vanishingly small innovations are accepted and $\gamma(\tau, m) \to 0$, driving the post-gate innovation energy to zero.

V. Bias Induced by Nearest-Neighbor Association

The previous sections characterized the statistical effect of ellipsoidal gating on innovation moments. In practical tracking systems, however, gating is only the first stage of a selection process. When multiple measurements fall inside the validation region, a data association rule is applied to select a single measurement for the state update. The most widely used rule in classical tracking is NN association [21], [22].

This section shows that NN association introduces an additional and unavoidable statistical bias that compounds the gate-induced effects. Unlike modeling errors or tuning artifacts, this bias arises solely from order-statistic selection and persists even under ideal nominal assumptions.

A. Nearest-Neighbor Association as Statistical Selection

Consider a time step at which $M \geq 1$ measurements pass the validation gate $\mathcal{A}$. Let $\{\nu^{(i)}\}_{i=1}^{M}$ denote the corresponding innovation vectors, each generated according to the same post-gate distribution $(\nu \mid \mathcal{A})$. NN association selects the innovation with minimum NIS,

$$i^\star \triangleq \arg\min_{1 \leq i \leq M} \|\nu^{(i)}\|^2_{S^{-1}}, \tag{16}$$

and uses $\nu^{(i^\star)}$ for the update step.

While (16) is often motivated algorithmically as a simple and efficient approximation to more complex association schemes, it constitutes a nonlinear statistical selection operator acting directly on the innovation.

B. Energy Contraction Under NN Association

The central effect of NN association is most clearly exposed at the level of second-order energy. The following proposition formalizes this effect without invoking Gaussianity or any specific parametric form.

Proposition 3 (Post-Gate NN Energy Contraction)**.** *Let $\nu \in \mathbb{R}^d$ be an innovation vector and let $\mathcal{A}$ denote the gating acceptance event. Let $\{\nu^{(i)}\}_{i=1}^{M}$ be independent samples from the conditional distribution $(\nu \mid \mathcal{A})$, and let $i^\star = \arg\min_i \|\nu^{(i)}\|$. Then*

$$\mathbb{E}\Big[\|\nu^{(i^\star)}\|^2 \mid \mathcal{A}\Big] \;\le\; \mathbb{E}\big[\|\nu\|^2 \mid \mathcal{A}\big]\,, \tag{17}$$

with strict inequality for $M > 1$ whenever the conditional distribution $(\nu \mid \mathcal{A})$ is non-degenerate.

This bound shows that NN association induces a systematic contraction of innovation energy beyond that caused by gating alone. The effect is purely statistical and follows from order-statistic selection.

Corollary 2 (Impossibility of Preserving Nominal Innovation Energy Under Selection)**.** *Let $\nu \sim \mathcal{N}(0, S)$ with $S \in \mathbb{S}_{++}^m$ and let $\mathcal{A} = \{\nu^\top S^{-1}\nu \le \tau\}$ be a nontrivial ellipsoidal gate with $\Pr(\mathcal{A}) \in (0,1)$. Assume that, whenever the gate admits $M \ge 1$ candidate measurements, the association rule selects the minimum-norm innovation among M independent post-gate candidates, yielding $\nu^{(i^\star)}$. Then, for any $M > 1$ for which $(\nu \mid \mathcal{A})$ is non-degenerate,*

$$\mathbb{E}\Big[\|\nu^{(i^\star)}\|^2 \mid \mathcal{A}\Big] \;<\; \mathrm{tr}(S). \tag{18}$$

Proof. Proposition 3 gives, for $M > 1$ and non-degenerate $(\nu \mid \mathcal{A})$,

$$\mathbb{E}\Big[\|\nu^{(i^\star)}\|^2 \mid \mathcal{A}\Big] < \mathbb{E}\big[\|\nu\|^2 \mid \mathcal{A}\big]\,.$$

Moreover, Proposition 2 implies

$$\mathbb{E}\big[\|\nu\|^2 \mid \mathcal{A}\big] = \mathrm{tr}(\gamma(\tau,m)S) = \gamma(\tau,m)\,\mathrm{tr}(S),$$

with $\gamma(\tau, m) \in (0,1)$ for any nontrivial gate. Therefore,

$$\mathbb{E}\Big[\|\nu^{(i^\star)}\|^2 \mid \mathcal{A}\Big] < \gamma(\tau,m)\,\mathrm{tr}(S) < \mathrm{tr}(S),$$

which proves the claim. □

C. Interpretation via Order Statistics

The purpose of this subsection is to provide intuition for Proposition 3 by isolating the purely statistical mechanism underlying NN association. NN association selects the minimum-norm element from a finite set of independent post-gate innovations. As a result, the selected innovation corresponds to the first-order statistic of the squared norms.

For any nonnegative, non-degenerate random variable X and $M \ge 2$ independent copies $\{X^{(i)}\}$,

$$\mathbb{E}\Big[\min_i X^{(i)}\Big] < \mathbb{E}[X]. \tag{19}$$

Applying this property to $X = \|\nu\|^2 \mid \mathcal{A}$ yields the strict inequality in (17). Notably, this argument relies only on elementary properties of order statistics and does not invoke Gaussianity, dimensionality, or the specific shape of the validation gate.

D. Other Data Association Methods

The analysis focuses on NN association because it yields a clean order-statistic statement. Other association methods do not remove the selection issue, but they expose it in different mathematical forms. In probabilistic data association (PDA), the update uses likelihood-weighted mixtures over gated measurements; the observed update is therefore affected by both truncation and likelihood weighting rather than by a hard minimum alone. In multiple-hypothesis tracking (MHT), multiple hypotheses are propagated, but gating, likelihood scoring, pruning, and hypothesis management still impose selection on the set of innovations that influence the maintained tracks. Deriving exact gate-conditioned reference statistics for PDA and MHT requires explicit assumptions on clutter, detection probability, and hypothesis management, and is left as a natural extension.

VI. Monte Carlo Validation and Tracking Case Study

This section provides explicit illustrations, Monte Carlo validation, and a multi-time-step tracking example. All simulations use white Gaussian process and measurement noises, correctly specified covariances, and the same ellipsoidal validation rule used in the analysis.

A. Two-Dimensional Closed-Form Gate Contraction

For $m = 2$, the NIS $Z = \nu^\top S^{-1}\nu$ follows a chi-square distribution with two degrees of freedom under nominal linear Gaussian assumptions. Equivalently, Z is exponentially distributed with density

$$f_Z(z) = \frac{1}{2}e^{-z/2}, \qquad z \ge 0, \tag{20}$$

and cumulative distribution function

$$F_Z(z) = 1 - e^{-z/2}. \tag{21}$$

Given a validation gate with acceptance probability P_g, the gate threshold is

$$\tau = -2\ln(1 - P_g). \tag{22}$$

The gate-conditioned mean NIS is

$$\mathbb{E}[Z \mid Z \le \tau] = \frac{2 - (\tau+2)e^{-\tau/2}}{1 - e^{-\tau/2}}, \tag{23}$$

or equivalently,

$$\mathbb{E}[Z \mid Z \le \tau] = 2\left(1 + \frac{(1-P_g)\ln(1-P_g)}{P_g}\right). \tag{24}$$

Thus

$$\gamma(P_g, 2) = 1 + \frac{(1-P_g)\ln(1-P_g)}{P_g}. \tag{25}$$

Table I and Fig. 1 show close agreement between the analytic gate-conditioned mean and Monte Carlo estimates. For the common 0.95 gate probability, the expected accepted NIS is approximately 1.685 rather than the nominal value 2.

TABLE I
GATE-INDUCED CONTRACTION IN TWO DIMENSIONS. MONTE CARLO VALUES ARE REPORTED WITH 95% CONFIDENCE HALF-WIDTHS.

P_g	τ	γ	Theory	MC
0.90	4.605	0.744	1.488	1.487 ± 0.004
0.95	5.991	0.842	1.685	1.684 ± 0.005
0.99	9.210	0.953	1.907	1.907 ± 0.006

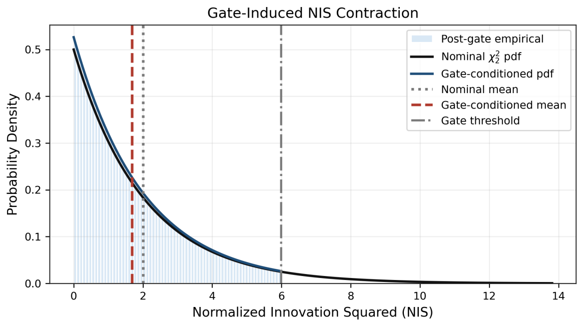

Fig. 1. Monte Carlo validation of the gate-conditioned NIS distribution for the two-dimensional case with gate probability 0.95. The accepted NIS samples follow the truncated chi-square density, and their mean is lower than the nominal reference.

TABLE II
NN ORDER-STATISTIC CONTRACTION FOR $m = 2$ AND $P_g = 0.95$.

M	Theory	MC	MC / nominal
1	1.685	1.689 ± 0.005	0.845
2	0.911	0.912 ± 0.003	0.456
3	0.619	0.616 ± 0.002	0.308
5	0.375	0.375 ± 0.001	0.187
10	0.189	0.189 ± 0.001	0.094

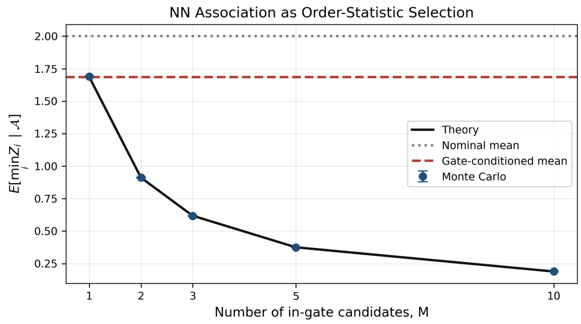

Fig. 2. Expected selected NIS under NN association as a function of the number of in-gate candidates. Monte Carlo results match the order-statistic prediction and show rapid contraction as the number of candidates increases.

B. NN Order-Statistic Validation

For the NN experiment, independent samples are drawn from the gate-conditioned NIS distribution and the minimum is selected among a fixed number of in-gate candidates. For $Z_i \sim (Z \mid Z \le \tau)$ independent,

$$\mathbb{E}\left[\min_{1\le i\le M} Z_i\right] = \int_0^\tau \left(\frac{F_Z(\tau) - F_Z(z)}{F_Z(\tau)}\right)^M dz. \tag{26}$$

The results in Table II and Fig. 2 validate the order-statistic analysis. Even with two in-gate candidates, the expected selected NIS is reduced to approximately 0.912, less than half of the nominal two-dimensional reference.

C. End-to-End Tracking and Tuning Study

To illustrate the practical implication for adaptive tuning, a multi-time-step constant-velocity tracking example is considered. The state is

$$x_k = \begin{bmatrix} p_{x,k} & v_{x,k} & p_{y,k} & v_{y,k} \end{bmatrix}^\top,$$

with sampling period $\Delta t = 1$. The process and measurement models are

$$x_k = Fx_{k-1} + w_{k-1}, \tag{27}$$
$$z_k = Hx_k + v_k, \tag{28}$$

where w_k and v_k are zero-mean white Gaussian noises with correctly specified covariances. The measurement dimension is two, the gate probability is 0.95, and the Monte Carlo study uses 1200 trials of 120 time steps. The compared cases are: no gate with fixed measurement covariance, gate with fixed measurement covariance, gate with nominal NIS-based covariance tuning, and gate with gate-aware NIS-based covariance tuning.

The end-to-end study confirms that the analytic gate-conditioned correction has practical consequences for tuning. The ungated fixed-covariance filter has the lowest RMSE because it uses all measurements and is included only as an ideal baseline without validation rejection. The fixed gated filter has larger RMSE because valid measurements are occasionally rejected by the gate, which is the expected cost of applying validation. With nominal NIS tuning, the final measurement-noise scale is biased downward to approximately 0.797, because the tuning rule attempts to force a gate-contracted statistic toward an unconditional reference. With the gate-aware statistic, the final scale remains close to the true value, approximately 1.042. Therefore, the performance-relevant comparison is within the gated adaptive cases: gate-aware tuning yields lower mean RMSE and lower NEES than nominal post-gate NIS tuning in this experiment. The NEES remains above the nominal value, emphasizing that NIS consistency and NEES consistency are distinct properties.

D. NEES and Nonlinear Filters

NIS consistency does not imply NEES consistency. Gating conditions not only the innovation statistic but also the event under which a measurement update occurs. Consequently, the posterior state error after a gated update is also conditioned on a selection event, and the resulting NEES need not follow the unconditional chi-square reference even when the innovation correction restores the mean accepted NIS. The Monte Carlo NEES results in Table III and Fig. 5 illustrate this distinction.

For nonlinear filters such as EKFs and UKFs, the exact linear-Gaussian derivation does not apply globally, but the same selection mechanism operates after local linearization or sigma-point prediction whenever an ellipsoidal innovation gate

TABLE III

END-TO-END TRACKING RESULTS. VALUES ARE MONTE CARLO MEANS WITH 95% CONFIDENCE HALF-WIDTHS.

Case	Used NIS	Corrected NIS	Final R scale	Update rate	RMSE	NEES
No gate, fixed R	1.986 ± 0.011	1.986 ± 0.011	1.000 ± 0.000	1.000 ± 0.000	3.284 ± 0.019	3.970 ± 0.036
Gate, fixed R	1.717 ± 0.008	2.039 ± 0.010	1.000 ± 0.000	0.936 ± 0.004	5.444 ± 1.138	4.477 ± 0.073
Gate, nominal NIS tuning	1.838 ± 0.006	2.182 ± 0.007	0.797 ± 0.006	0.910 ± 0.004	4.890 ± 0.776	4.962 ± 0.076
Gate, gate-aware NIS tuning	1.706 ± 0.005	2.025 ± 0.006	1.042 ± 0.008	0.939 ± 0.003	4.142 ± 0.410	4.398 ± 0.060

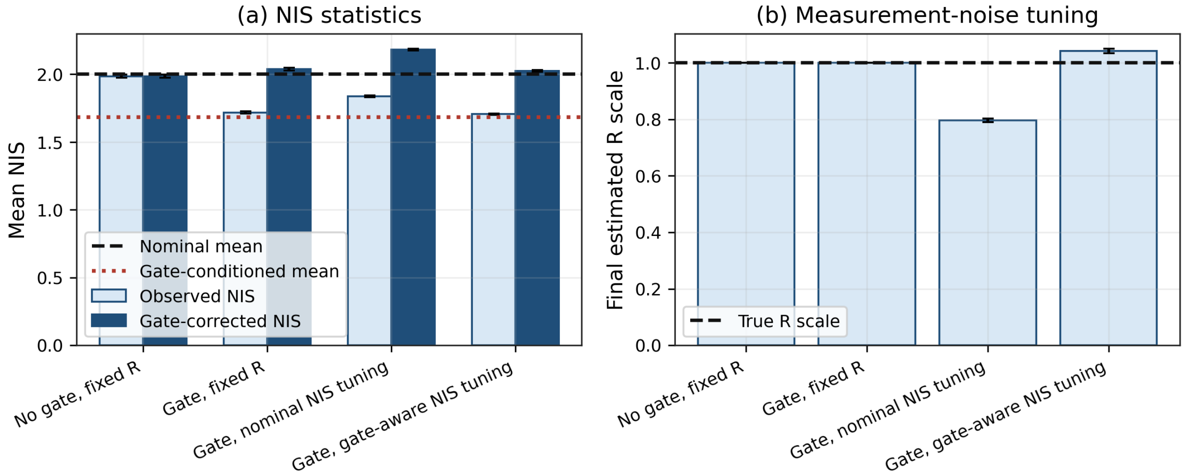

Fig. 3. End-to-end tracking results for NIS statistics and adaptive measurement-noise tuning. Nominal post-gate tuning learns a biased final measurement-noise scale, whereas the gate-aware statistic remains close to the true scale.

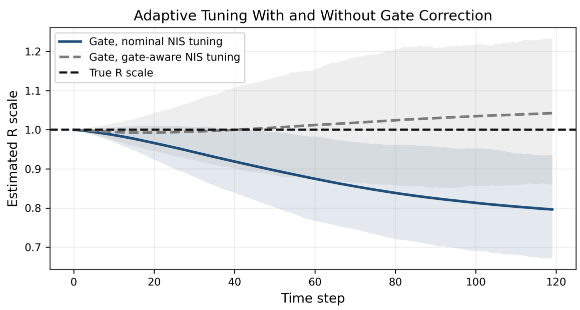

Fig. 4. Time evolution of the estimated measurement-noise scale. Nominal NIS tuning reacts to gate-contracted innovations by reducing the estimated scale, while gate-aware tuning remains near the true reference.

is used. Problems such as bearings-only tracking can already exhibit fragile chi-square approximations due to geometry and linearization error; validation gating then further conditions the observed innovation sequence. A complete nonlinear theory is outside the scope of this paper, but the linear result provides the baseline selection effect that any nonlinear extension must account for.

VII. DISCUSSION

The analysis developed in this paper provides a principled explanation for a widely observed phenomenon in Kalman-based tracking systems: the tendency of innovation-based diagnostics evaluated after gating and association to exhibit systematically reduced innovation energy relative to nominal reference values, even in well-functioning trackers.

Sections III–VI show that this behavior arises from two structural selection mechanisms inherent to practical tracking pipelines:

- ellipsoidal gating, which deterministically truncates the innovation distribution and contracts its covariance, and
- NN association, which introduces an additional order-statistic bias by selecting the minimum-norm innovation among multiple candidates.

Both mechanisms operate even under ideal modeling assumptions. They do not rely on clutter, false alarms, nonlinear dynamics, or model mismatch. As a result, tuning-based explanations alone are insufficient to account for systematic bias observed in post-gate innovation statistics.

The Monte Carlo results show that the effect is not merely an asymptotic or algebraic observation. In a multi-time-step tracking loop, using the unconditional NIS reference on accepted innovations biases adaptive measurement-noise tuning, whereas using the gate-conditioned correction substantially reduces that bias.

A. Implications for Classical Tracking Theory

The apparent discrepancy between nominal and observed innovation statistics arises because classical theory characterizes the unconditional innovation process, whereas practical

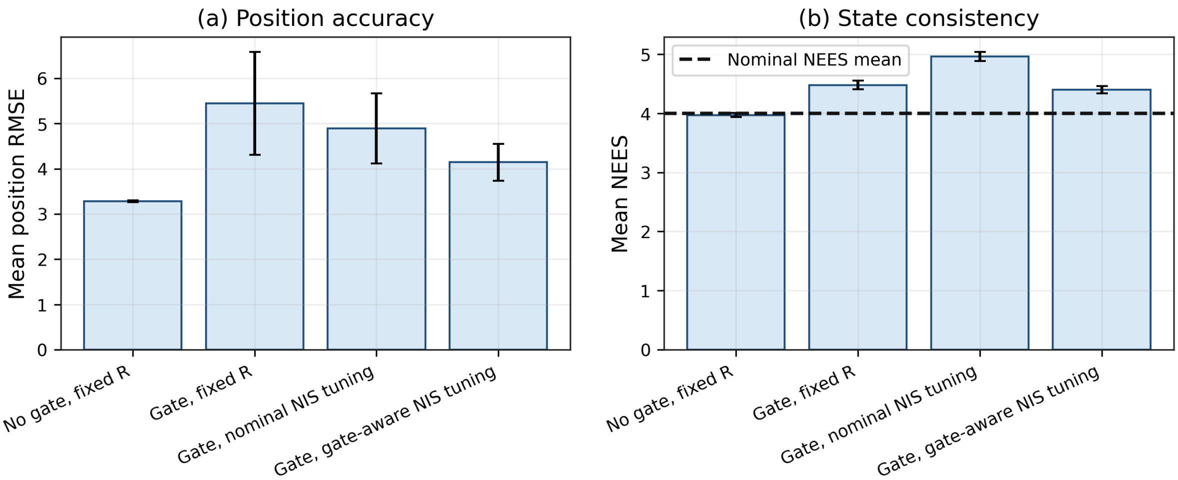

Fig. 5. Position RMSE and NEES in the end-to-end tracking experiment. The ungated case is a baseline that uses every measurement; it is not the target that gated tuning is expected to outperform. The relevant tuning comparison is between the two gated adaptive cases, where the gate-aware reference reduces the bias induced by nominal post-gate NIS tuning. The NEES panel also illustrates that NIS correction does not by itself guarantee exact state-space consistency.

tracking systems operate on a conditionally observed innovation stream. The present work does not propose an alternative to adaptive or robust Kalman filtering. Rather, it provides a missing statistical layer that clarifies how innovation-based diagnostics should be interpreted before invoking adaptation mechanisms. In this sense, the analysis complements existing tuning and robustness methods by separating structural selection effects from genuine model mismatch.

The connection to the matrix-valued measures of Forsling et al. [13] suggests an additional forward path. Scalar NIS gating compresses vector-valued innovation information into a single quadratic form; matrix-valued consistency measures may retain richer directional information and could be studied under analogous gate-conditioned selection laws.

B. Implications for Implementation and Evaluation

From an implementation perspective, the results suggest several practical guidelines:

- Innovation covariance and NIS statistics computed after gating should be interpreted using gate-conditioned reference values.
- Systematically reduced innovation energy observed in post-gate diagnostics does not necessarily indicate filter mistuning.
- Adaptive tuning schemes based on innovation statistics should explicitly account for selection-induced bias to avoid self-reinforcing measurement-noise underestimation and valid-measurement rejection.
- When ground truth is available, NEES should be evaluated separately from NIS because correcting the accepted NIS reference does not guarantee state-space consistency.

These guidelines can be incorporated into existing tracking systems without modifying the Kalman recursion, validation logic, or association rules themselves.

C. Limitations and Scope

The analysis in this paper is deliberately restricted to linear measurement models, Gaussian noise, ellipsoidal validation gates, and NN association. These restrictions ensure that the identified effects are not artifacts of nonlinearity or non-Gaussianity.

The Monte Carlo tracking experiment is also intentionally restricted to a single-target, correctly specified, linear-Gaussian setting so that the observed bias can be attributed to selection rather than to clutter modeling or association failures. Extensions to nonlinear filters, PDA, MHT, and explicit clutter models are important directions for future work, but require additional modeling choices that would obscure the isolated selection mechanism studied here.

VIII. Conclusion

Validation gating and NN association introduce predictable and quantifiable biases in innovation statistics, even under ideal Kalman filter assumptions. Ellipsoidal gating deterministically contracts the innovation covariance, while NN association introduces an additional multiplicity-dependent bias through order-statistic selection.

These structural effects explain why innovation-based diagnostics evaluated after gating and association systematically deviate from nominal chi-square references in practical tracking systems. Monte Carlo validation confirms the analytic gate-conditioned and order-statistic predictions, and an end-to-end tracking example demonstrates that nominal NIS-based tuning can learn a biased measurement-noise scale when applied to accepted post-gate innovations. Gate-aware interpretation restores the appropriate innovation reference, although state-space consistency as measured by NEES remains a separate issue.

Acknowledgment

The author gratefully acknowledges Prof. Yaakov Bar-Shalom for his insightful comments that helped clarify the interpretation and terminology of the paper.

Appendix

We prove the statements in Proposition 2 by transforming the innovation into whitened coordinates and exploiting rotational invariance of the Gaussian distribution.

Since $S \in \mathbb{S}^m_{++}$, there exists a symmetric matrix $S^{1/2}$ such that $S^{1/2}S^{1/2} = S$. Define the whitened innovation

$$u \triangleq S^{-1/2}\nu. \tag{29}$$

By construction,

$$u \sim \mathcal{N}(0, I_m), \tag{30}$$

and the NIS becomes

$$Z = \nu^\top S^{-1}\nu = u^\top u = \|u\|^2. \tag{31}$$

Accordingly, the gating event can be written as

$$\mathcal{A} = \{\|u\|^2 \le \tau\}. \tag{32}$$

The distribution of u is symmetric about the origin, and the event $\|u\|^2 \le \tau$ is invariant under $u \mapsto -u$. Therefore,

$$\mathbb{E}[u \mid \mathcal{A}] = 0,$$

and transforming back gives $\mathbb{E}[\nu \mid \mathcal{A}] = 0$.

Define

$$C \triangleq \mathbb{E}[uu^\top \mid \mathcal{A}].$$

Since u is rotationally invariant and the event $\|u\|^2 \le \tau$ depends only on the norm of u, the conditional distribution $u \mid \mathcal{A}$ is also rotationally invariant. Consequently, C commutes with all orthogonal transformations and must be of the form $C = \gamma(\tau, m)I_m$. Taking traces gives

$$m\gamma(\tau, m) = \mathbb{E}[\mathrm{tr}(uu^\top) \mid \mathcal{A}] = \mathbb{E}[Z \mid Z \le \tau],$$

which establishes (10). Transforming back to the original coordinates yields

$$\mathbb{E}[\nu\nu^\top \mid \mathcal{A}] = S^{1/2}CS^{1/2} = \gamma(\tau, m)S.$$

Finally, because Z is nonnegative, nondegenerate, and has mean m, truncating to the nontrivial event $Z \le \tau$ gives $0 < \mathbb{E}[Z \mid Z \le \tau] < m$, hence $0 < \gamma(\tau, m) < 1$.

Condition on the gating acceptance event $\mathcal{A}$. Let $\nu^{(1)}, \ldots, \nu^{(M)}$ be independent samples from $(\nu \mid \mathcal{A})$, and define

$$X_i \triangleq \|\nu^{(i)}\|^2, \qquad i = 1, \ldots, M.$$

Then the X_i are i.i.d. nonnegative random variables under the conditional law. NN association selects

$$X_{(1)} \triangleq \|\nu^{(i^\star)}\|^2 = \min_{1 \le i \le M} X_i.$$

Since $X_{(1)} \le X_1$ almost surely,

$$\mathbb{E}[X_{(1)} \mid \mathcal{A}] \le \mathbb{E}[X_1 \mid \mathcal{A}] = \mathbb{E}[\|\nu\|^2 \mid \mathcal{A}].$$

If $M > 1$ and the conditional distribution is nondegenerate, then $\Pr(X_{(1)} < X_1 \mid \mathcal{A}) > 0$, which yields strict inequality.

References

[1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
[2] R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Journal of Basic Engineering*, vol. 83, no. 1, pp. 95–108, 1961.
[3] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. Courier Corporation, 2007.
[4] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*. John Wiley & Sons, 2001.
[5] R. G. Brown and P. Y. C. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*, 3rd ed. Wiley, 1992.
[6] R. Mehra, "On the identification of variances and adaptive kalman filtering," *IEEE Transactions on Automatic Control*, vol. 15, no. 2, pp. 175–184, 1970.
[7] S. Akhlaghi, N. Zhou, and Z. Huang, "Adaptive adjustment of noise covariance in kalman filter for dynamic state estimation," in *2017 IEEE Power & Energy Society General Meeting*, 2017, pp. 1–5.
[8] J. Wang, T. Zhang, X. Xu, and Y. Li, "A variational bayesian based strong tracking interpolatory cubature kalman filter for maneuvering target tracking," *IEEE Access*, vol. 6, pp. 52 544–52 560, 2018.
[9] W. Youn, N. Y. Ko, S. Gadsden, and H. Myung, "A novel multiple-model adaptive kalman filter for an unknown measurement loss probability," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 10, pp. 7953–7966, 2020.
[10] K. V. Ramachandra, *Kalman Filtering Techniques for Radar Tracking*. CRC Press, 2018.
[11] S. Shats, B. Z. Bobrovsky, and U. Shaked, "Discrete-time state estimation of analog double integrators," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, no. 6, pp. 670–677, 1988.
[12] R. Karlsson and F. Gustafsson, "The future of automotive localization algorithms: Available, reliable, and scalable localization: Anywhere and anytime," *IEEE Signal Processing Magazine*, vol. 34, no. 2, pp. 60–69, 2017.
[13] R. Forsling, S. J. Julier, and G. Hendeby, "Matrix-valued measures and wishart statistics for target tracking applications," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 61, no. 5, pp. 12 234–12 244, 2025. [Online]. Available: https://doi.org/10.1109/TAES.2025.3571685
[14] B. Or, B.-Z. Bobrovsky, and I. Klein, "Kalman filtering with adaptive step size using a covariance-based criterion," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–10, 2021.
[15] B. Or and I. Klein, "A hybrid model and learning-based adaptive navigation filter," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–11, 2022.
[16] ——, "Learning vehicle trajectory uncertainty," *Engineering Applications of Artificial Intelligence*, vol. 122, p. 106101, 2023.
[17] R. Piché, "Online tests of kalman filter consistency," *International Journal of Adaptive Control and Signal Processing*, vol. 30, no. 1, pp. 115–124, 2016.
[18] J. Duník, O. Kost, O. Straka, and E. Blasch, "Covariance estimation and gaussianity assessment for state and measurement noise," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 1, pp. 132–139, 2020.
[19] J. Duník, O. Straka, O. Kost, and J. Havlík, "Noise covariance matrices in state-space models: A survey and comparison of estimation methods–part i," *International Journal of Adaptive Control and Signal Processing*, vol. 31, no. 11, pp. 1505–1543, 2017.
[20] Z. Chen, H. Biggie, N. Ahmed, S. Julier, and C. Heckman, "Kalman filter auto-tuning through enforcing chi-squared normalized error distributions with bayesian optimization," *arXiv preprint arXiv:2306.07225*, 2023.
[21] D. Reid, "An algorithm for tracking multiple targets," *IEEE transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
[22] Y. Bar-Shalom, T. E. Fortmann, and P. G. Cable, "Tracking and data association," 1990.

Barak Or (Member, IEEE) received a B.Sc. degree in aerospace engineering (2016), a B.A. degree (cum laude) in economics and management (2016), and an M.Sc. degree in aerospace engineering (2018) from the Technion–Israel Institute of Technology. He graduated with a Ph.D. degree from the University of Haifa, Haifa (2022). His research interests include navigation, deep learning, sensor fusion, and estimation theory.