Gibbs sensor fusion: Multi-sensor integration via maximum-entropy recognition-science cost weighting

Dylan Funk[1] and Jonathan Washburn[1, a)]
Recognition Physics Institute, Austin, Texas 78701, USA

Merging redundant channels is routine in plasma diagnostics and control, but noise is often far from Gaussian and channel covariances may be missing, drifting, or meaningless after a fault. **Gibbs Sensor Fusion** (GSF) assigns each channel a weight $w_i \propto \exp(-J(x_i)/T)$ from its ratio $x_i = \hat{\theta}_i/\theta_{\text{ref}}$ to a reference, using the Recognition Science cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ (uniqueness of J is established elsewhere[1]) and a scalar temperature T. No variance matrix is required: Gibbs weights are the maximum-entropy distribution at fixed expected cost[12], and for fixed weights the fused value minimizes the weighted J-loss in closed form. We implement a one-pass $O(N)$ reference-based step and benchmark it against the mean, median, and sliding-window inverse-variance Kalman fusion on multi-sensor runs driven by a gyrokinetic reference trace. On CGYRO-driven tests, fused RMSE is roughly 15–30% below sliding-window inverse-variance Kalman when a channel fails or variance estimates mis-rank outliers; correctly specified inverse-variance weights can still win on the fixed-bias outlier (Sec. 5).

Keywords: sensor fusion; Gibbs weighting; plasma diagnostics; robust estimation; maximum entropy; Kalman filter; recognition science

I. INTRODUCTION

Combining several sensors into one estimate is standard in signal processing and control[2]. In plasmas, interferometry, Thomson scattering, ECE, magnetics, reflectometry, and other diagnostics often return overlapping information on density, temperature, and field. Merging them sensibly matters for real-time control and for analysis after the shot[3,4].

The usual toolkit—Kalman and particle filters[4,5], covariance intersection[6], Dempster–Shafer combination[7]—leans on explicit noise models, covariances, or likelihoods. In experiment those ingredients are often wrong: noise is non-Gaussian, sensors fault or spike, covariances are costly to track, and cross-diagnostic correlations are guesswork[8,9]. Robust M-estimation[9] helps, but someone still picks the loss.

Gibbs Sensor Fusion (GSF) sets weights from the Recognition Science cost functional

$$J(x) = \tfrac{1}{2}\big(x + x^{-1}\big) - 1, \tag{1}$$

the unique cost satisfying coherent composition (d'Alembert functional equation)[10], normalization $J(1) = 0$, and quadratic calibration at unity[1]. Each sensor i providing estimate $\hat{\theta}_i$ of an unknown quantity θ^* receives weight

$$w_i = \frac{\exp(-J(x_i)/T)}{Z}, \tag{2}$$

where the partition function is

$$Z = \sum_{j=1}^{N} \exp(-J(x_j)/T), \tag{3}$$

and $x_i = \hat{\theta}_i/\theta_{\text{ref}}$ is the ratio to a reference value (typically the median), and $T > 0$ is a temperature parameter. This formulation applies to positive quantities ($\hat{\theta}_i$, $\theta_{\text{ref}} > 0$); for signed quantities, a transform to a positive ratio is used (Sec. 6.2). The cost J is fixed by the coherence axioms[1]; convexity and growth of J carry the usual robustness intuition[9]. Propositions III.1 and III.2 give the maximum-entropy weights under a cost constraint and the unique θ that minimizes $\sum_i w_i J(\hat{\theta}_i/\theta)$ when w is fixed. Because $J(x) \to \infty$ as $x \to 0^+$ or $x \to \infty$, extreme ratios receive negligible weight (Proposition III.3). Section 3.4 discusses choices for T.

It is also useful to collect the recognition and entropy terms in a single free-energy objective,

$$\begin{aligned}\mathcal{L}(\theta, w; \lambda) = &\sum_i w_i J(\hat{\theta}_i/\theta) \\ &+ T \sum_i w_i \ln w_i \\ &+ \lambda\Big(\sum_i w_i - 1\Big),\end{aligned} \tag{4}$$

over $\theta > 0$, $w_i \geq 0$, with λ enforcing $\sum_i w_i = 1$. Alternating minimization of $\mathcal{L}$ would alternate the Gibbs and $\theta^\star$ updates in Propositions III.1–III.2; the procedure in Sec. III.1 performs only the first Gibbs evaluation at θ_{ref} and the subsequent $\theta^\star$ step.

II. RECOGNITION SCIENCE COST FUNCTIONAL

The cost J is determined by three axioms on $F : \mathbb{R}_{>0} \to \mathbb{R}$:

(A1) $F(1) = 0$;

(A2) the d'Alembert composition law $F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y)$ for all $x, y > 0$[10];

a)Electronic mail: jonathan@recognitionscience.org

(A3) in log coordinates $t = \ln x$, $\lim_{t\to 0} 2G(t)/t^2 = 1$ for $G(t) = F(e^t)$.

Under A1–A3, F is uniquely $J(x) = \frac{1}{2}(x + x^{-1}) - 1$[1].

Key properties of J:

- $J(x) \geq 0$ with equality iff $x = 1$;
- reciprocity $J(x) = J(x^{-1})$;
- boundary divergence $J(x) \to +\infty$ as $x \to 0^+$ or ∞;
- strict convexity $J''(x) = x^{-3} > 0$;
- quadratic near unity: $J(e^\epsilon) = \cosh(\epsilon) - 1 \approx \frac{1}{2}\epsilon^2$;
- linear asymptotics: $J(x) \approx x/2$ (large x), $J(x) \approx 1/(2x)$ (small x).

The linear growth for large deviations is the structural origin of outlier robustness[9].

III. GIBBS SENSOR FUSION FRAMEWORK

A. Algorithm

The objective $\mathcal{L}$ in (4) couples recognition cost to an entropy term over $\theta > 0$ and $w \in \Delta_N = \{w_i \geq 0,\ \sum_i w_i = 1\}$. The multiplier $\lambda(\sum_i w_i - 1)$ pins w to the simplex. *If* one were to minimize $\mathcal{L}$ by alternating in (θ, w), the two blocks would be: (a) for fixed θ, the unique entropy-maximizing weights on Δ_N at fixed expected cost are Gibbs weights $w_i \propto \exp(-J(\hat{\theta}_i/\theta)/T)$ (Proposition III.1); equivalently, constrained stationarity with a multiplier on $\sum_i w_i - 1$ yields the same form. (b) for fixed w, θ updates to $\sqrt{(\sum_i w_i\hat{\theta}_i)/(\sum_i w_i/\hat{\theta}_i)}$ (Proposition III.2). The **algorithm** below stops after one pass: compute Gibbs weights using costs at a fixed θ_{ref}, then apply the $\theta^\star$ map from (b). That pair is motivated by alternating minimization of $\mathcal{L}$ but is not asserted to minimize $\mathcal{L}$ over (θ, w).

Given estimates $\{\hat{\theta}_i\}_{i=1}^N$ and $T > 0$:

1. The reference is defined as follows: for odd N, $\theta_{\text{ref}} = \hat{\theta}_{((N+1)/2)}$ (the median order statistic); for even N, $\theta_{\text{ref}} = \sqrt{\hat{\theta}_{(N/2)}\hat{\theta}_{(N/2+1)}}$ (the geometric mean of the two central order statistics, i.e., the median in log space)[11].
2. $x_i = \hat{\theta}_i/\theta_{\text{ref}}$, $J_i = \frac{1}{2}(x_i + x_i^{-1}) - 1$.
3. Weights are obtained via a log-sum-exp stabilization for numerical stability: $a_i = -J_i/T$, $a_{\max} = \max_i a_i$, and $w_i = \exp(a_i - a_{\max})/\sum_j \exp(a_j - a_{\max})$.
4. $\hat{\theta}_{\text{fused}} = \theta^\star := \sqrt{\left(\sum_i w_i\hat{\theta}_i\right)/\left(\sum_i w_i/\hat{\theta}_i\right)}$ (Proposition III.2).

For the weights w from step 3 (which depend on θ_{ref}, not on $\theta^\star$), $\hat{\theta}_{\text{fused}}$ is the unique minimizer of $\sum_i w_i J(\hat{\theta}_i/\theta)$ (Proposition III.2). No further θ–w iteration is required for the estimate we report. One could iterate—recomputing $w_i \propto \exp(-J(\hat{\theta}_i/\theta^\star)/T)$ and updating $\theta^\star$—as an optional extension; that full loop is not what we implement or evaluate here.

A median reference is one way to resist outliers; in tracking, the reference might instead be a prior or the previous fused value.

B. Reference dependence and breakdown

The one-pass estimator is fixed once we choose a reference rule $\theta_{\text{ref}} = R(\hat{\theta}_1, \ldots, \hat{\theta}_N)$, Gibbs weights $w_i \propto \exp(-J(\hat{\theta}_i/\theta_{\text{ref}})/T)$, and aggregation $\hat{\theta}_{\text{fused}} = A(\hat{\theta}, w) = \sqrt{(\sum_i w_i\hat{\theta}_i)/(\sum_i w_i/\hat{\theta}_i)}$. Throughout, R = median. In log space, $J(\hat{\theta}_i/\theta_{\text{ref}}) = \cosh(\delta_i) - 1$ with $\delta_i = \ln(\hat{\theta}_i/\theta_{\text{ref}})$, so the scheme is a log-space kernel smoother centered at the reference: it rewards agreement with the reference, not correctness. Because $\delta_k = \ln(\hat{\theta}_k/\theta_{\text{ref}})$ satisfies $\partial\delta_k/\partial\theta_{\text{ref}} = -1/\theta_{\text{ref}}$, the log-derivative in θ_{ref} obeys $\partial_{\ln\theta_{\text{ref}}} \ln(w_i/w_j) = (\sinh\delta_i - \sinh\delta_j)/T$, i.e. $\partial_{\theta_{\text{ref}}} \ln(w_i/w_j) = (\sinh\delta_i - \sinh\delta_j)/(T\theta_{\text{ref}})$. If the reference is wrong or the majority is biased, weights shift accordingly. With median reference, the breakdown point is 50%: if more than half of the sensors are coherently biased (e.g. calibration drift), the median moves to the biased cluster and the fused estimate reinforces that bias. Without an external anchor (prior, calibration, or known-good sensor), no purely data-consensus estimator can distinguish "majority wrong" from "minority wrong."[11] Practical mitigations include using a trusted prior in the joint objective (see future work in Sec. VII) or external calibration.

C. Maximum entropy and bounded influence

Proposition III.1 (Gibbs weights maximize entropy under a cost constraint)**.** *Fix $J_1, \ldots, J_N \in \mathbb{R}$ and suppose there exists $w \in \Delta_N$ with $\sum_i w_i J_i = c$. Among all $w \in \Delta_N$ satisfying $\sum_i w_i J_i = c$, the unique maximizer of the Shannon entropy $S[w] = -\sum_i w_i \ln w_i$ is*

$$w_i = \frac{\exp(-J_i/T)}{Z(T)}, \qquad Z(T) = \sum_{j=1}^{N} \exp(-J_j/T), \tag{5}$$

for a unique $T = T(c)$ (equivalently, a unique Lagrange multiplier for the cost constraint).

Proof. Consider maximizing $S[w]$ over Δ_N subject to $\sum_i w_i J_i = c$. Introduce Lagrange multipliers λ and μ for the constraints $\sum_i w_i = 1$ and $\sum_i w_i J_i = c$. Stationarity of $\mathcal{F}(w) = -\sum_i w_i \ln w_i - \lambda(\sum_i w_i - 1) - \mu(\sum_i w_i J_i - c)$ with respect to w_i on the interior gives $\ln w_i + 1 + \lambda + \mu J_i =$

0, hence $w_i \propto \exp(-\mu J_i)$. Writing $\mu = 1/T$ yields the Gibbs form. Uniqueness follows because $S[w]$ is strictly concave on Δ_N and the linear constraint defines a convex feasible set, so the maximizer is unique when nonempty. The free-energy form $F = \langle J \rangle - TS$ is the Legendre dual statement[12]. □

Proposition III.2 (Minimizer of weighted cost in θ)**.** *Let $\hat{\theta}_1, \ldots, \hat{\theta}_N > 0$ and $w \in \Delta_N$. The unique $\theta > 0$ minimizing $\sum_{i=1}^N w_i J(\hat{\theta}_i/\theta)$ is*

$$\theta^* = \sqrt{\frac{\sum_i w_i \hat{\theta}_i}{\sum_i w_i/\hat{\theta}_i}}. \tag{6}$$

Proof. Set $f(\theta) = \sum_i w_i J(\hat{\theta}_i/\theta) = \frac{1}{2}\sum_i w_i(\hat{\theta}_i/\theta + \theta/\hat{\theta}_i) - 1$. Then $f'(\theta) = \frac{1}{2}\sum_i w_i(-\hat{\theta}_i/\theta^2 + 1/\hat{\theta}_i)$. Solving $f'(\theta) = 0$ gives $\theta^2 = (\sum_i w_i\hat{\theta}_i)/(\sum_i w_i/\hat{\theta}_i)$. Since $J''(x) > 0$ for $x > 0$, each term $w_i J(\hat{\theta}_i/\theta)$ is strictly convex in $\theta > 0$ when $w_i > 0$, hence f is strictly convex and θ^* is the unique global minimizer. □

Proposition III.3 (Weight decay for large cost)**.** *Let $w_i(T) \propto \exp(-J_i/T)$ on Δ_N. Fix k with $J_k < \infty$. If $J_i \to \infty$ while $\{J_j\}_{j\neq i}$ are fixed, then $w_i/w_k \to 0$.*

Proof. $w_i/w_k = \exp(-(J_i - J_k)/T) \to 0$ as $J_i \to \infty$ for fixed $T > 0$. □

Together, Propositions III.1–III.2 are the two blocks of an alternating minimization of $\mathcal{L}$; Sec. III.1 stops after one pass: weights at θ_{ref}, then $\hat{\theta}_{\text{fused}} = \theta^\star$ for that w.

The weight ratio for any two sensors i, j is

$$\frac{w_i}{w_j} = e^{-(J_i - J_j)/T}. \tag{7}$$

Let $J_{\min} = \min_j J_j$. Then each weight satisfies

$$\frac{1}{N} e^{-(J_i - J_{\min})/T} \le w_i \le e^{-(J_i - J_{\min})/T} \le 1. \tag{8}$$

In particular, relative to any finite-cost sensor k, $w_i/w_k \to 0$ as $J_i \to \infty$ (Proposition III.3): any sensor with arbitrarily large cost receives exponentially vanishing Gibbs weight. As $x_i \to 0^+$ or ∞, $J(x_i) \to \infty$, so $w_i \to 0$ exponentially. This is distinct from the classical influence-function formalism for location estimators[13], though it limits the effect of a single extreme ratio on the *weights*. For example, a sensor with ratio $x = 10$ has $J(10) = 4.05$; at $T = 1$ its weight is $\propto \exp(-4.05) \approx 0.017$ relative to a perfect sensor; at $T = 2$, ≈ 0.143; at $T = 5$, ≈ 0.445. In contrast, quadratic-loss methods have unbounded influence[9].

D. Asymptotic limits and temperature choice

As $T \to 0^+$, weights concentrate on $\arg\min_i J_i$ (hard selection); if a single sensor i^* dominates, $\hat{\theta}_{\text{fused}} = \theta^\star \to \hat{\theta}_{i^*}$. When the reference is the median, the cost-minimizing sensor at $T \to 0$ is the median order statistic for odd N, or one of the two central order statistics for even N, so GSF hard-selection coincides with selecting the median in that case. As $T \to \infty$, $w_i \to 1/N$ and

$$\hat{\theta}_{\text{fused}} \to \sqrt{\frac{N^{-1}\sum_i \hat{\theta}_i}{N^{-1}\sum_i 1/\hat{\theta}_i}}, \tag{9}$$

the square root of the ratio of arithmetic to harmonic sample means (equal to the arithmetic mean only when all $\hat{\theta}_i$ coincide). For intermediate T, the weight distribution interpolates between these extremes.

The temperature T may be selected according to domain-specific criteria. **Option (i):** Let $\kappa > 1$ denote a multiplicative threshold beyond which a sensor is regarded as unreliable, and let $s \in (0, 1)$ denote the desired relative weight for such a sensor; then

$$T = \frac{J(\kappa)}{\ln(1/s)}, \tag{10}$$

so that a sensor with ratio κ receives weight s relative to a sensor with unit ratio ($w(\kappa)/w(1) = s$). **Option (ii):** Specify a target effective number of sensors $N_{\text{eff}}^{\text{target}}$ (e.g., $N/2$) and determine T such that $N_{\text{eff}}(T) = N_{\text{eff}}^{\text{target}}$, where $N_{\text{eff}}(T) = 1/\sum_i w_i^2$ or $\exp(-\sum_i w_i \ln w_i)$[23]; this requires a one-dimensional root-find in T. On synthesized random data, the fused estimate asymptotes to order statistics at low T and to $\sqrt{(\text{AM})/(\text{HM})}$ of the sensors at high T (Sec. 3.4), and can lie outside the interval between the median and the arithmetic mean at intermediate T. In the mild CGYRO scenario used for the figures (Sec. 5), RMSE still exhibits a broad minimum at moderate T.

E. Variance reduction and bounded influence

Conditional on fixed weights w, $\hat{\theta}_{\text{fused}} = \theta^\star$ is a nonlinear function of $\{\hat{\theta}_i\}$; for independent additive noise a simple variance formula is not as transparent as for a weighted arithmetic mean. The operational algorithm uses data-dependent w_i, so unconditional variance is not given by a closed form here. Differentiating the implicit equation $f'(\theta^\star) = 0$ from the proof of Proposition III.2 with respect to $\hat{\theta}_k$ (holding θ_{ref} fixed so w is fixed) yields bounded partial sensitivities when $w_k \to 0$ as $\hat{\theta}_k$ becomes inconsistent with the reference, so gross outliers do not shift $\theta^\star$ without bound through the weight channel[13].

IV. COMPARISON WITH CLASSICAL APPROACHES

Inverse-variance and Kalman-style rules need variances or covariances; linear Gaussian Kalman updates are not robust to gross outliers in the usual sense. For a single-time batch of N scalar measurements, both GSF and standard

recursive Kalman measurement updates scale linearly in N, so comparing only leading-order cost is uninformative—what matters for GSF is the J-based Gibbs map without an explicit noise model. Particle methods address nonlinear state evolution with explicit dynamics and likelihoods, which is a different target than the static or windowed scalar fusion in Sec. V. Where M-estimation leaves the loss ρ as a modeling choice, here J is fixed by the axioms and inherits the maximum-entropy interpretation of Proposition III.1. Near equilibrium, Gibbs weights in log-residual approximate a Gaussian kernel, connecting GSF to log-normal multiplicative-noise models[14]. For large x_i, $J(x_i) \sim x_i/2$, and for $x_i \to 0^+$, $J(x_i) \sim 1/(2x_i)$; in both limits $J \to \infty$, driving Gibbs weights for extreme ratios toward zero (Proposition III.3). We compare only to the mean, median, and sliding-window inverse-variance Kalman; trimmed means, Huber-type M-estimators, innovation gating, and heavy-tailed likelihoods are left aside.

A. Relation to filtering practice in physics and plasmas

State-space filtering—Kalman, extended/unscented Kalman, and particle filters—is the standard tool when a dynamical model, measurement model, and noise statistics are available: the update combines predicted state uncertainty with measurement innovation using gains derived from covariances[3–5]. That line of work targets *full* estimation of a latent state trajectory. Here the question is smaller: at each time (or each control cycle), form a single representative value from several *parallel* scalar estimates of the same positive quantity when variance and cross-channel correlations are poorly known, wrong after a fault, or non-Gaussian. It is a *fusion map* on diagnostics, not a substitute for equilibrium reconstruction, profile inversion, or MHD forward modeling.

Robust statistics and influence-function analysis treat contamination, outliers, and model misspecification via chosen losses, redescending scores, or high-breakdown location estimators[9,11,13]. Robust Kalman-type filters extend that idea by modifying innovations or covariances under heavy-tailed or faulty data; those methods still embed explicit stochastic structure. GSF occupies a middle ground: it inherits Gibbs weighting and a convex cost J with fixed asymptotics (Sec. II), avoiding ad hoc selection of ρ, yet it does not propagate a full state covariance as in classical filtering.

In fusion and laboratory plasma systems, multi-sensor integration appears at control and data-analysis layers: combining interferometry, Thomson, ECE, magnetics, and radiometry for density, temperature, shape, and stability metrics[2,8,17]. Operational practice often relies on inverse-variance or hand-tuned weights, median-like robustness, or Kalman filters with noise matrices adjusted by experience[17]. GSF is meant as a lightweight option when many channels must be merged on millisecond clocks and covariance tuning is fragile; caveats (common-mode bias, consensus failure) appear in Sec. VI. How one pipes GSF into a larger estimator or controller is left to the installation.

V. SIMULATION RESULTS

Simulations compare simple averaging, the median, a sliding-window inverse-variance Kalman filter, and GSF at $T = 0.5$ (aggressive down-weighting of channels far from the reference). The main run drives synthetic sensors from a CGYRO trace instead of a hand-picked scalar "truth."

CGYRO-based test. CGYRO is a first-principles gyrokinetic code used for turbulence and stability in magnetized plasmas[24]. We run a linear electrostatic case (Miller geometry[25], single k_y, ion and electron species) with gentler gradients and higher collisionality ("mild" scenario) to produce a smooth time series. The trace is plotted over $t \in [0, 17.5]$ s (875 points), analogous in character to a line-averaged density or fluctuation amplitude. From this reference series we construct synthetic multi-sensor signals: several channels are the reference value plus Gaussian noise; one channel is an outlier (fixed bias); and in a second scenario one sensor fails midway (constant wrong value). A third test uses four unbiased channels with a common mean but unequal Gaussian noise, one of them substantially noisier than the other three. Inverse-variance Kalman—with variances either specified from the simulation or estimated in a sliding window—systematically down-weights that high-noise channel, whereas temporal GSF does not model per-sensor variances and forms weights only from ratios to the reference, so GSF fused RMSE can exceed Kalman's or the median's in this geometry even though no channel is biased. Temporal GSF uses $\theta_{\text{ref}}^{(t)} = \hat{\theta}_{\text{fused}}^{(t-1)}$ and is compared to averaging, median, and Kalman on fused RMSE. The trace comes from a plasma-physics code; injected noise, bias, and mid-run failure resemble what diagnostics see; GSF assumes no tuned covariance; and the lag-one reference update is the kind of rule used in simple real-time loops. All figures in this section use that mild CGYRO setup except Fig. 4, where $\theta_{\text{ref}}^{(t)}$ is the cross-diagnostic median each time (the static rule of Sec. 3.1) so a stuck bad channel is compared to the bulk rather than to the lagged fused estimate.

Scenario 1: Five sensors, one outlier. Figs. 1–3 use the mild CGYRO trace on $t \in [0, 17.5]$ s; $\langle\theta\rangle$, θ_{sc}, and noise $0.03\,|\langle\theta\rangle|$ are computed from $t \geq 1$ s on the trace. One sensor is fixed at $1.5\,\theta_{\text{sc}}$. *Kalman (specified)* assigns variances σ^2 on the good channels and $(0.5\,\langle\theta\rangle)^2$ on the outlier draw (a finite, conventional covariance slot for the faulty channel rather than the generative fixed bias). RMSE over $t \geq 1$ s ranks as in Fig. 2 (left): *Kalman (specified)* and the *median* are lowest and close; *GSF* at $T = 0.5$ is intermediate; *average* and *Kalman (estimated)* are highest and close—the sliding window often fails to mark the biased channel as high-variance, so inverse-variance fusion resembles the mean. Absolute

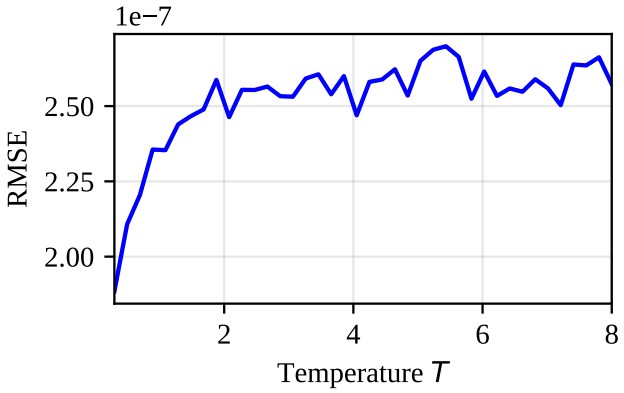

FIG. 1. RMSE of GSF versus temperature T for Scenario 1 (five sensors, one outlier; Monte Carlo over sensor draws; $\langle\theta\rangle$ sets the noise scale on the CGYRO trace). Small T favors a single channel; large T washes out toward equal weights and weaker suppression of bad channels.

RMSEs are orders of magnitude below unity because the underlying trace is small in the mean. These panels do *not* reproduce an older standalone scalar experiment with $\theta^* = 100$, $\sigma = 5$, and multi-unit RMSEs from 1000 trials.

Scenario 2: Sensor failure over time. Four sensors observe the CGYRO-derived reference series; at $t \approx 9.25\,\mathrm{s}$—the temporal midpoint of $[1, 17.5]\,\mathrm{s}$—one sensor fails (constant wrong value). Before failure all methods perform comparably. After failure, the simple average and Kalman filter (which trusts the failed sensor's low apparent variance) degrade sharply, while GSF degrades only slightly because the failed sensor incurs higher J-cost and receives lower weight; the fused estimate tracks the reference.

RMSE versus T for Scenario 1 has a broad minimum at moderate T; performance degrades at very low T (over-selection) and very high T (insufficient outlier suppression). Figure 1 shows RMSE vs T; Fig. 2 gives RMSE bar comparisons for Scenarios 1–3. Figure 3 plots absolute error versus time; after the failure, GSF stays closer to zero than the average or Kalman in the panels shown. Figure 4 uses a different reference choice and compares GSF to the average (left, biased channel) and to Kalman (right, failed channel). In both panels GSF uses the per-time diagnostic median as θ_{ref}; GSF (blue) tracks lower absolute error than the average (left) and than sliding-window inverse-variance Kalman (right), which over-weights the stuck channel once its empirical variance collapses.

VI. RELEVANCE TO PLASMA PHYSICS

Experiments stack redundant views of the same quantities: interferometry and polarimetry for line density; Thomson, ECE, and soft x-ray for T_e; coils and loops for field and flux; magnetic and optical blends for shape[15,16]. Controllers still need one number per quantity on millisecond clocks, often with poorly known noise and channels that drop or lie[17]. Those are the conditions the present paper targets.

A. Diagnostic heterogeneity and failure

Plasma diagnostics differ in principle, calibration, and noise. Interferometer phases can be corrupted by vibration, refraction, or fringe jumps[8]; ECE channels can saturate or be affected by non-thermal radiation[18]; Thomson scattering is typically sparse in time and space[15]; magnetic signals are subject to integrator drift and electromagnetic interference[19]. Noise is often non-Gaussian (transients, digitization, MHD-induced excursions)[20]. Sensors can fail temporarily (fiber break, detector fault, lost line of sight) or produce outliers during edge-localized modes (ELMs), disruptions, or start-up transients[21]. Classical Kalman fusion assumes known or estimated covariances and Gaussian noise[4]; when covariances are wrong or a diagnostic fails, the fused estimate can be severely biased[17]. GSF requires no covariance model: each diagnostic is weighted by its consistency with the reference (e.g., the median over diagnostics). Large ratio errors yield large J_i and, by Proposition III.3, vanishingly small weight relative to a finite-cost channel. Yet at temperatures and errors typical of a running system, the map can be gentle: with $J(0.6) \approx 0.13$, $J(1.5) \approx 0.08$, and $T = 0.5$, the Gibbs factors $e^{-J/T}$ differ by roughly ten percent—soft steering rather than hard dropout. The weights therefore reward proximity to whatever reference rule was fixed, not an independent certificate of which hardware is "good"; a shared systematic error that pulls most channels together can move the reference with it, as discussed below.

Like any consensus rule, GSF has a breakdown point: with a median reference, more than half bad channels can drag the reference[11]. Common-mode drift or a shared systematic can push most interferometers or magnetics together[8,22]. In that situation the median lies within the biased cluster and GSF reinforces the bias, down-weighting the minority of correct sensors. In the absence of an external anchor (e.g., a trusted prior, independent calibration, or a known-good diagnostic), no purely data-driven fusion method can distinguish a biased majority from a biased minority; proposed mitigations include the use of a prior in the joint objective (Sec. VII) or external calibration[17].

B. Positive physical quantities and ratio structure

Many plasma quantities are positive (densities, temperatures, intensities, line integrals). The GSF cost $J(x)$ is defined for ratios $x = \hat{\theta}/\theta_{\mathrm{ref}} > 0$, so it applies directly to such quantities. The ratio formulation is scale-free: the same logic applies whether densities are in $10^{19}\,\mathrm{m}^{-3}$ or arbitrary units, and T is dimensionless. For quantities

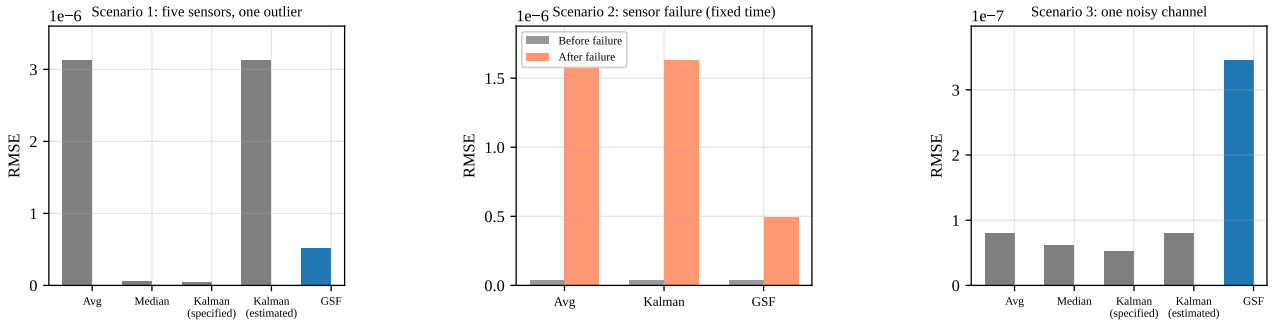

FIG. 2. Scenario 1: RMSE by method. *Kalman (specified)* uses inverse-variance fusion with weights fixed from per-sensor variances set to match the generative model (an idealized baseline when the variance specification is correct). *Kalman (estimated)* uses the same fusion rule but with weights from variances estimated in a backward sliding window over past samples, as when channel noise levels are not known a priori. Scenario 2: RMSE before and after sensor failure (average, *Kalman (estimated)*, GSF); GSF exhibits only modest degradation. Scenario 3: four sensors, three with common noise and one with larger Gaussian noise (unbiased); *Kalman (specified)* uses the four variances of the generative model.

that can be negative (e.g., radial position error ξ), one may work with a positive quantity such as $\theta = \xi - \xi_{\min} + c$ for a constant $c > 0$ (or another one-to-one map to $\mathbb{R}_{>0}$), then apply GSF to the transformed values; the fused result is transformed back to the original scale.

C. Real-time and temporal fusion

Control systems need low-latency, low-cost updates. GSF is $O(N)$ per fusion step with no matrix inversions or iterative solvers, making it suitable for real-time loops where many diagnostics are combined every few milliseconds. The temporal extension $J_i^{(t)} = J(\hat{\theta}_i^{(t)}/\hat{\theta}_{\text{fused}}^{(t-1)})$ incorporates the previous fused estimate as reference, so diagnostics that suddenly jump (e.g., after a fault) are down-weighted relative to the recent consensus. This provides a simple, robust filter without propagating full covariance matrices as in a Kalman filter.

D. Concrete applications

Line-integrated density. Multiple interferometer or polarimeter chords give $\int n_e\, dl$ along different paths. GSF can fuse these into a single representative line-averaged density (or per-chord estimates), assigning lower weight to channels whose ratios to the reference are spoiled by fringe jumps or noise spikes.

Electron temperature. ECE radiometry, Thomson scattering, and soft x-ray diagnostics provide T_e at various locations. Combining them with GSF yields a robust profile or core value without assuming Gaussian errors or known variances.

Magnetic reconstruction. Many magnetic probes and flux loops supply redundant information for equilibrium reconstruction. GSF can fuse scalar quantities (e.g., estimated stored energy or normalized beta) from different algorithms or subsets of sensors, or weight individual sensor contributions in a preprocessing step, with failed or aberrant coils tending to receive lower weight.

Disruption and stability indicators. Multiple indicators (e.g., from temperature, density, MHD, and shape) can be fused into a single robust metric for disruption prediction or stability assessment[17,21], where outlier sensors (e.g., a single bad channel) do not dominate the decision.

E. Summary for the plasma community

For plasma experiments the advantage is straightforward: no full covariance matrix, no Gaussian assumption, and channels whose ratios to the reference are absurd end up with almost no weight (subject to the consensus caveats above). T can be pinned from rough engineering targets—how strongly to down-weight a multiplicative error κ, or roughly how many channels should contribute—without sweeping a fine grid. Similar sensor-fusion problems show up outside plasmas; the combination of messy diagnostics and tight control deadlines is simply very visible on tokamaks.

VII. DISCUSSION AND CONCLUSION

Once J is fixed by the axioms, Gibbs weights and the $\theta^\star$ map follow. The only control left on the bench is T, which fixes how sharply weights penalize disagreement with the reference.

Limitations: the ratio x_i requires positive quantities (signed quantities require a transform); the presentation assumes independent sensors (correlations would require a joint cost structure); the one-pass estimator does not jointly minimize the postulated free energy (4) over (θ, w); and temperature adaptation, extension to full state-space models (e.g., position and velocity), and finite-sample or

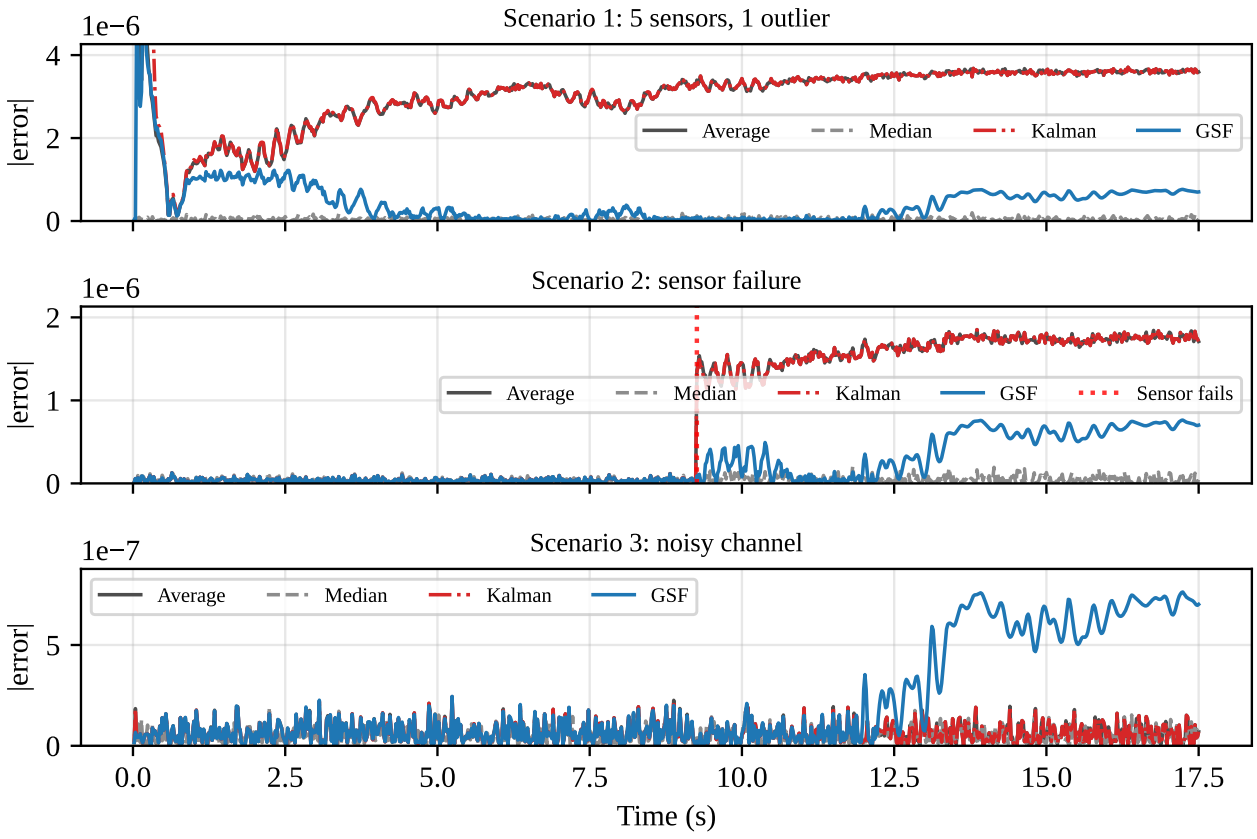

FIG. 3. Absolute error vs time for average, median, *Kalman (estimated)*, and GSF in each scenario (the same sliding-window inverse-variance rule as in Fig. 2). In the second panel (sensor failure), error for average and *Kalman (estimated)* rises after the failure while GSF stays comparatively stable. In the third panel (Scenario 3: one channel with larger Gaussian noise but unbiased), *Kalman (estimated)* persistently down-weights that high-variance channel; GSF does not estimate channel variances and only penalizes ratios to the temporal reference, so its absolute error can exceed *Kalman (estimated)* or the median even though no channel is biased. Time range $t \in [0, 17.5]$ s.

minimax guarantees under specific noise families remain open. In addition, (i) the reference-based estimator exhibits a 50% breakdown point when the reference is the median[11]: if more than half of the sensors are arbitrarily corrupted, the reference is displaced and the fused estimate may degrade. (ii) When a majority of sensors are coherently biased (e.g., common-mode calibration drift), the fused estimate reinforces that bias; no purely data-consensus method can distinguish "majority wrong" from "minority wrong" without an external anchor[17]. (iii) Use of the previous fused value as reference in temporal fusion may introduce lag or preferential weighting of a sensor that remains consistent with the past. Proposed mitigations include external calibration or trusted priors and, as outlined below, a prior-term extension of the joint objective.

Future work. Costs built from known variances or per-sensor T_i (near equilibrium, weights resemble Gaussian kernels in log-residual and connect to log-normal noise when T tracks log-variance[14]), vector costs, temporal schemes with a free reference plus a prior toward $\hat{\theta}^{(t-1)}_{\text{fused}}$, adaptive T from the spread of $\{J_i\}$ or from a target effective channel count, hierarchical clustering of sensors, and soft priors inside joint Gibbs objectives[12] are not implemented here.

In summary, multi-sensor fusion is written in terms of the Recognition Science cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$; uniqueness of J under the coherence axioms is in Ref.[1]. Proposition III.1 recovers maximum-entropy weights at fixed expected cost; Proposition III.2 gives the closed θ-update for fixed weights. The estimator in Sec. 3.1 forms weights at θ_{ref} and sets $\hat{\theta}_{\text{fused}} = \theta^\star$, which minimizes weighted J-loss for that w without claiming joint stationarity of (4) in (θ, w). Proposition III.3 and Sec. 3.2 spell out sensitivity of weights to the reference. Section 3.4 covers T; the fusion step is $O(N)$. On the CGYRO-driven tests, fused RMSE is roughly 15–30% lower than sliding-window inverse-variance Kalman in the mid-run failure case and when variance estimates mis-rank biased or stuck channels; inverse-variance fusion with correctly specified variances can still beat GSF on the Scenario 1 fixed-bias outlier (Sec. 5). Sec. IV A places the method next to stan-

When GSF wins: robust to outlier (left) and sensor failure (right)

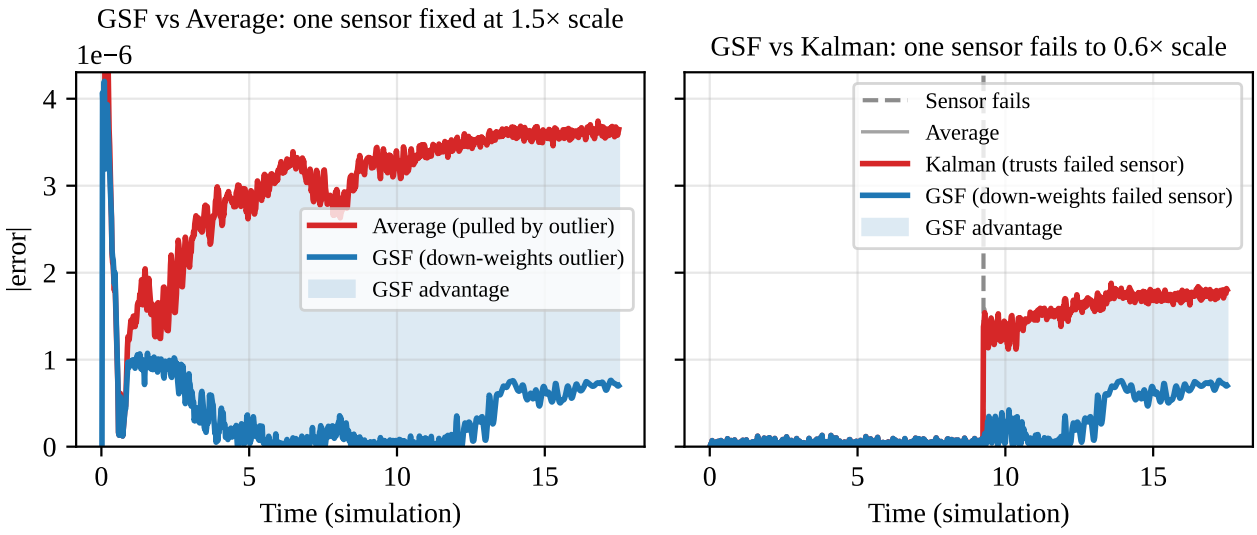

FIG. 4. GSF versus baselines under the reference choice described in the text. Outlier and failure levels use $\theta_{\mathrm{sc}} = \max(\langle|\theta|\rangle, \sigma_\theta, \max_t|\theta|)$ on the CGYRO trace. Left: one sensor fixed at $1.5\,\theta_{\mathrm{sc}}$; GSF (blue) vs average (red). Right: one sensor fails to $0.6\,\theta_{\mathrm{sc}}$; GSF (blue) vs *Kalman (estimated)* (red), i.e. sliding-window inverse-variance fusion as elsewhere in the simulation section. GSF uses the diagnostic median each time as θ_{ref} (see text). Time range $t \in [0, 17.5]$ s.

dard filtering; Sec. VI sketches plasma use cases. What is still missing are finite-sample and minimax guarantees, correlated channels, and runs on hardware.

Data from the simulations reported in this work are available from the authors upon request.

ACKNOWLEDGMENTS

The authors thank collaborators at the Recognition Science Institute and Recognition Physics Institute for discussions on cost-functional applications. We thank Anil Thapa and Philip Beltracchi for detailed comments that helped strengthen the manuscript. No external funding was received. The authors declare no competing interests.

[1] J. Washburn and A. Rahnamai Barghi, Axioms **15**, 151 (2026); https://doi.org/10.3390/axioms15020151.
[2] D. L. Hall and J. Llinas, Proc. IEEE **85**, 6 (1997).
[3] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation* (Wiley, New York, 2001).
[4] R. E. Kalman, J. Basic Eng. **82**, 35 (1960).
[5] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics* (MIT Press, Cambridge, MA, 2005).
[6] S. J. Julier and J. K. Uhlmann, Proc. SPIE **3068**, 182 (1997).
[7] G. Shafer, *A Mathematical Theory of Evidence* (Princeton University Press, Princeton, NJ, 1976).
[8] A. J. H. Donné *et al.*, Nucl. Fusion **48**, 075002 (2008).
[9] P. J. Huber, Ann. Math. Stat. **35**, 73 (1964).
[10] J. Aczél, *Lectures on Functional Equations and Their Applications* (Academic Press, New York, 1966).
[11] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection* (Wiley, New York, 1987).
[12] E. T. Jaynes, Phys. Rev. **106**, 620 (1957); **108**, 171 (1957).
[13] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel, *Robust Statistics: The Approach Based on Influence Functions* (Wiley, New York, 1986).
[14] E. Limpert, W. A. Stahel, and M. Abbt, BioScience **51**, 341 (2001).
[15] I. H. Hutchinson, *Principles of Plasma Diagnostics* (Cambridge University Press, Cambridge, 1987).
[16] T. C. Luce, Phys. Plasmas **19**, 055501 (2012).
[17] J. R. Ferron *et al.*, Fusion Sci. Technol. **59**, 61 (2011).
[18] M. E. Austin, Rev. Sci. Instrum. **78**, 033501 (2007).
[19] J. B. Lister *et al.*, Fusion Eng. Des. **36**, 327 (1997).
[20] S. J. Zweben *et al.*, J. Plasma Phys. **82**, 655820301 (2016).
[21] T. C. Hender *et al.*, Nucl. Fusion **47**, S128 (2007).
[22] P. Blanchard *et al.*, Rev. Sci. Instrum. **87**, 11E512 (2016).
[23] L. Kish, *Survey Sampling* (Wiley, New York, 1965).
[24] E. A. Belli and J. Candy, J. Comput. Phys. **324**, 290 (2016).
[25] R. L. Miller *et al.*, Phys. Plasmas **5**, 973 (1998).