Sequential Filtering in the Presence of Uniform Measurement Errors

James S. McCabe

NASA Johnson Space Center, Houston, TX

2023 AAS/AIAA Astrodynamics Specialist Conference, Big Sky, MT

August 13th – 17th, 2023

Outline

Where Are Uniform Measurement Errors Present?

- Many of our estimation and simulation efforts focus on Gaussian measurement noises
- Indeed, many (most) systems are well-represented using Gaussian statistics!
- However, some prevalent systems are corrupted by **uniform** measurement errors
  - Quantization, thermal noise, clock errors, emerging quantum devices, PFA/PFR
  - Subjective observations, polling, broad input measurements
- In this case, the classic Kalman filter can be used safely (and it often is)
  - A common misconception is that the Kalman filter *requires* Gaussian noises
  - It is the optimal MMSE estimator so long as noise statistics are known sufficiently
- However, does the application need stronger optimality (better estimation performance)?
  - Other methods exist to obtain stronger (usually Bayesian) optimality, i.e., PF, GMF, etc.
  - Generally, however, these methods consume much, much more computational throughput
  - **This paper derives a new, efficient (approximately) Bayes-optimal sequential filter.**

But What About The Central Limit Theorem?

- It is common to lean upon the **Central Limit Theorem** to justify Gaussian methods
  - There is nothing wrong with this! But is it the best for the application?
  - The moments may well be matched by a Gaussian, but what about their frequency content?
  - (The CLT is be nuanced and complex, but here seek to emphasize it's not a catch-all)
- Consider the 10,000,000 sample demo below, sampled from randomly selected uniforms

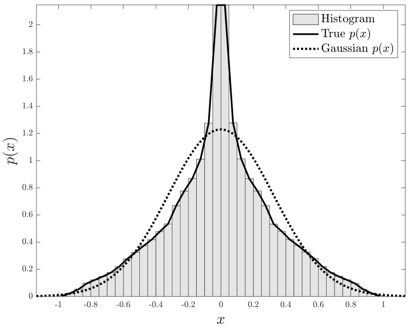

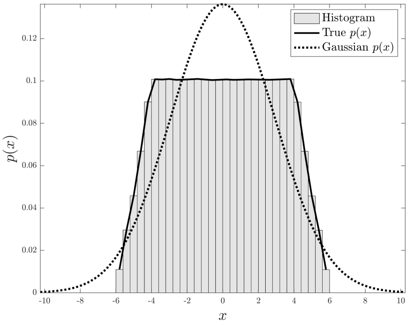

System Design and Requirements: Intent v. Realization

- Another key observation (though anecdotal) is in the community interpretation of "**3σ**"
  - There is a tendency in system and requirements design to interpret 3σ as a "maximum error"
  - Designers often intend it to be used as a "bound", often with little other specification
  - Sometimes it may be safer to interpret as uniform rather than matching a Gaussian 3σ
- There's a disconnect between the **design intent** and what gets **designed/flown**

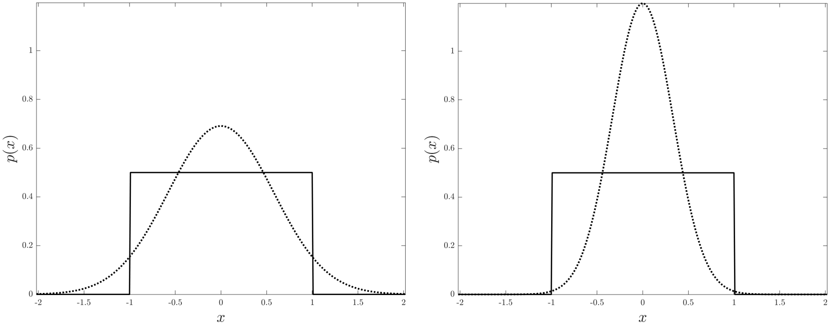

The Masreliez Filter (1)

- In an astounding though seemingly forgotten work in 1975, Masreliez derived an extremely general solution to (approximate) Bayesian estimation for arbitrary noise densities

$$\boldsymbol{m}^{+}_{x,k} = \boldsymbol{m}^{-}_{x,k} + \boldsymbol{P}^{-}_{xx,k}\boldsymbol{H}^{T}_{x,k}\boldsymbol{g}_k(\boldsymbol{z}_k)$$
$$\boldsymbol{P}^{+}_{xx,k} = \boldsymbol{P}^{-}_{xx,k} - \boldsymbol{P}^{-}_{xx,k}\boldsymbol{H}^{T}_{x,k}\boldsymbol{G}_k(\boldsymbol{z}_k)\boldsymbol{H}_{x,k}\boldsymbol{P}^{-}_{xx,k}$$

- This relies "solely" upon computing very tricky likelihood derivatives (often impossible)

$$g_{i,k}(\boldsymbol{z}_k) = -\left[\frac{\partial p(\boldsymbol{z}_k|\boldsymbol{Z}_{1:k-1})}{\partial z_{i,k}}\right]\frac{1}{p(\boldsymbol{z}_k|\boldsymbol{Z}_{1:k-1})}, \qquad [\boldsymbol{G}_k(\boldsymbol{z}_k)]_{i,j} = \frac{\partial g_{i,k}(\boldsymbol{z}_k)}{\partial z_{j,k}}$$

- Approximation to Bayes rule that yields very similar results to the true Bayes posterior[1]
- Despite its appearance, it's a nonlinear estimator (in contrast to the KF/EKF/UKF/etc.)
  - However, for Gaussian noises... *These derivatives just give us the Kalman filter!*

[1] M. Brunot recently used this to produce a recursion for Gauss-uniform noises, but the result is unwieldy for our purposes (especially for onboard applications).

The Masreliez Filter (2)

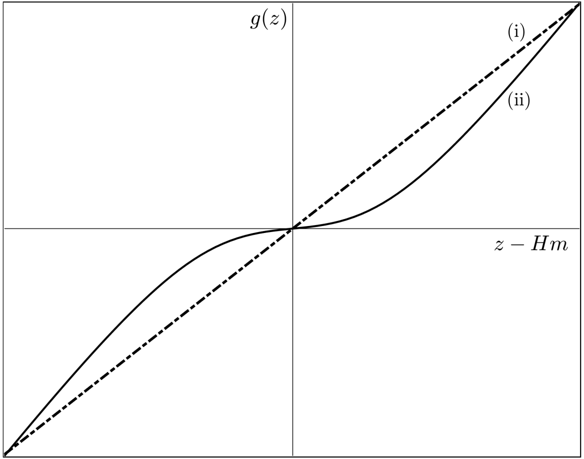

A New Filter for Uniform Measurement Noises

- This paper takes all the ugly derivatives to compute $\boldsymbol{g}_k(\boldsymbol{z}_k)$ and $\boldsymbol{G}_k(\boldsymbol{z}_k)$ for uniform noises
- The result winds up being surprisingly tidy and is of the form

$$\boldsymbol{g}_k(\boldsymbol{z}_k) = \boldsymbol{S}_k^{-T}\boldsymbol{\xi}_k$$
$$\boldsymbol{G}_k(\boldsymbol{z}_k) = \boldsymbol{S}_k^{-T}\boldsymbol{\Xi}_k\boldsymbol{S}_k^{-1}$$

- Above, $\boldsymbol{S}_k$ is the square-root factor of $\boldsymbol{H}_{x,k}\boldsymbol{P}_{xx,k}^{-}\boldsymbol{H}_{x,k}^{T}$ and

$$\boldsymbol{\xi}_k \triangleq \boldsymbol{\psi}_k \sinh\{\boldsymbol{b}_k^T\boldsymbol{\beta}_k\}$$
$$\boldsymbol{\Xi}_k \triangleq \mathrm{diag}\{\boldsymbol{\psi}_k \odot \boldsymbol{\zeta}_k + \boldsymbol{\xi}\}$$

- This only requires a few more operations than the trusty Kalman filter but provides approximately Bayes-optimal estimation performance
- Simply plug $\boldsymbol{g}_k(\boldsymbol{z}_k)$ and $\boldsymbol{G}_k(\boldsymbol{z}_k)$ into the previous updates for $\boldsymbol{m}_{x,k}^{+}/\boldsymbol{P}_{xx,k}^{+}$ and you're done

Implementation Considerations

1. What was odd to me was that hyperbolic trigonometry pops out of the derivatives
   - This can actually be tricky and/or expensive for spaceflight-grade CPUs...
   - However, lookup tables are handily implemented with great accuracy
   - A total non-issue for modern computing architectures

2. Computing $\boldsymbol{\xi}_k$ requires evaluating a function containing **error functions**

$$\phi_{i,k} \triangleq \operatorname{erf}\left\{\frac{b_{i,k}-\beta_{i,k}}{\sqrt{2}}\right\}+\operatorname{erf}\left\{\frac{b_{i,k}+\beta_{i,k}}{\sqrt{2}}\right\}$$

   - Anyone who has messed with error functions knows that they can be troublesome numerically
   - However, this is an age-old problem with a myriad of attractive solutions, even for lean CPUs

3. Covariance matrices are inherently plagued by floating point arithmetic
   - Loss of symmetric positive semidefiniteness (PSD) makes any covariance-based filter collapse
   - **Square-root** and **UDU** factorized representation are derived in the paper
   - These all but solve any symmetry or PSD concerns for the covariance matrix

Numerical Studies

- 10,000 trial Monte Carlos
- These runs compare perf. of
  1. Kalman filter with different assumptions
  2. optimal Gaussian mixture filter
  3. this paper's uniform filter
- **Case 1:** True noise is uniform, KF with true uniform stats.
- **Case 2:** True noise is uniform, KF with 3σ matched to uniform
- **Case 3:** Notched Gaussian noise
- **Case 4:** Heavy-tailed noise
- Note that the KF has all the standard "bells and whistles"

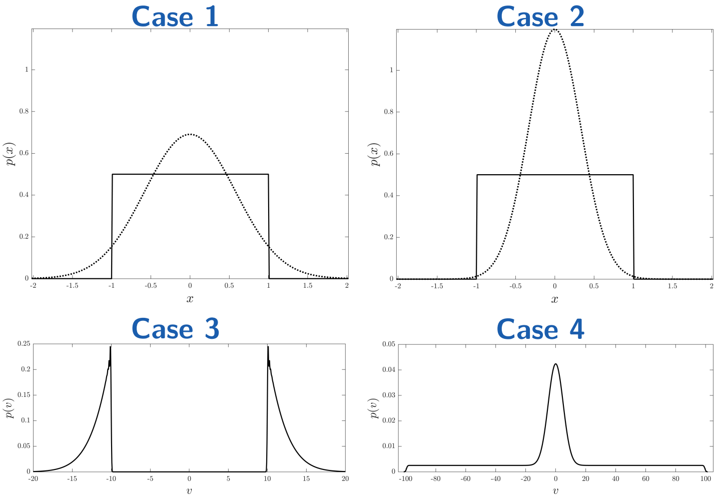

Case 1: Uniform Noise, Ideal Kalman Filter

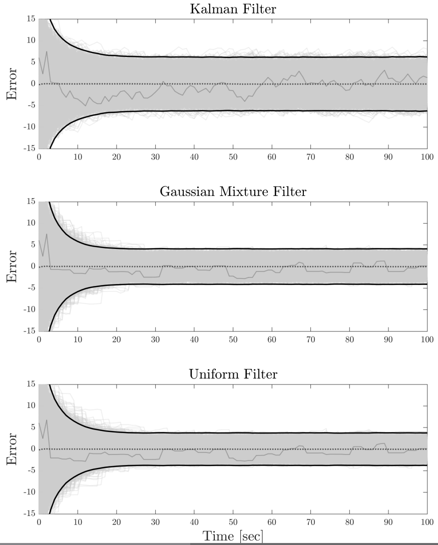

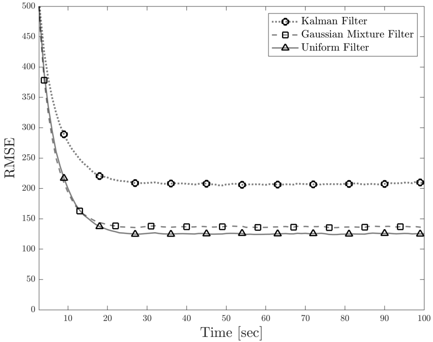

Case 1

Filter	Total RMSE	Avg. Norm. Time
Kalman	3165.8	1.0
Gaussian Mixture	2725.2	15704.4
Uniform	2690.0	16.0

Case 2: Uniform Noise, Matched Kalman Filter

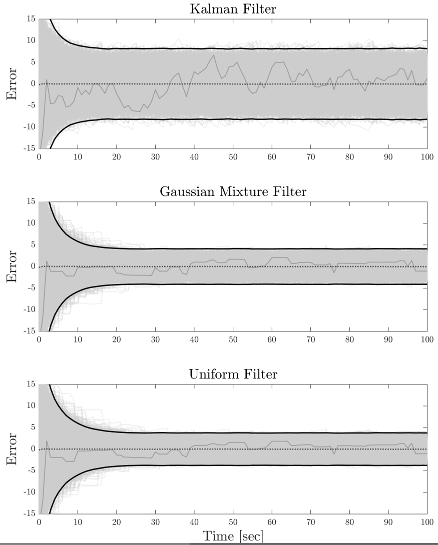

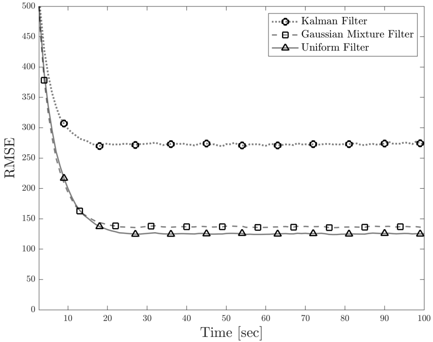

Case 2

Filter	Total RMSE	Avg. Norm. Time
Kalman	3584.6	1.0
Gaussian Mixture	2724.1	15188.3
Uniform	2688.9	15.9

Case 3: Notched Measurement Noise

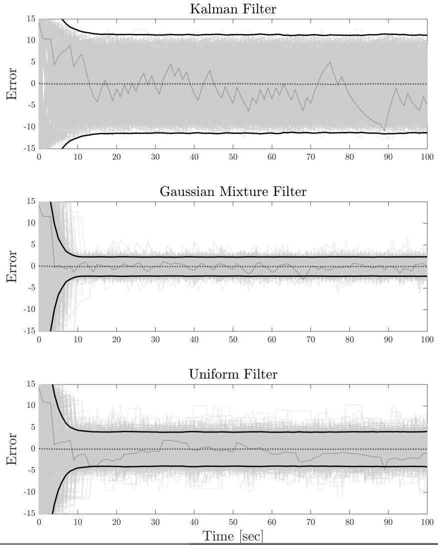

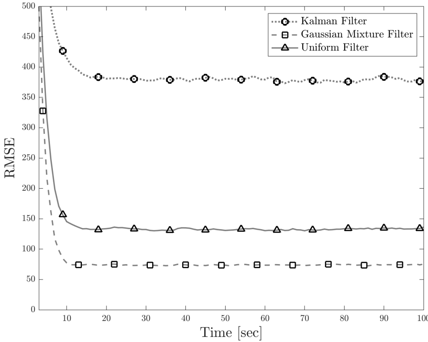

Case 3

Filter	Total RMSE	Avg. Norm. Time
Kalman	4599.3	1.0
Gaussian Mixture	2571.7	2851.6
Uniform	2851.6	16.1

Case 4: Heavy-Tailed Noise

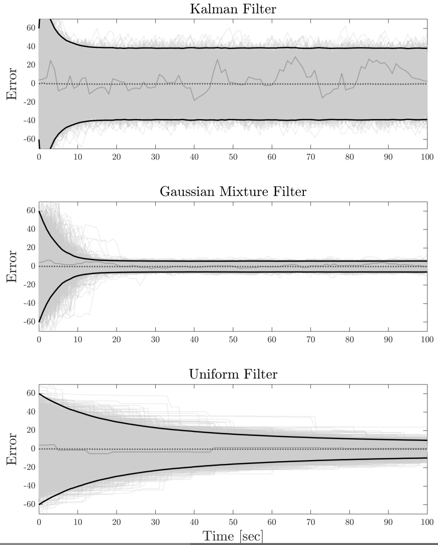

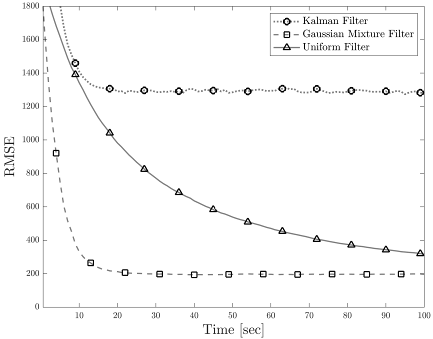

Case 4

Filter	Total RMSE	Avg. Norm. Time
Kalman	14207.8	1.0
Gaussian Mixture	3993.2	2236.1
Uniform	8170.1	5.2

Conclusions

- This paper presents **a new sequential estimator for uniform measurement noise**
  - Produces approximately Bayes-optimal estimation performance
  - Very efficient, slightly more computationally burdensome than the Kalman filter
  - Numerical stability adjustments developed, including factorized formulations
- Numerical studies are compelling and illustrate its advantages
  - Demonstrates ideal performance when estimating with uniform measurement noises
  - Also demonstrates improvements upon KF to accomodate robust estimation strategies
- All proofs provided explicitly in the paper (including the key likelihood function)
- These results build upon Masreliez's key discovery and Brunot's recent findings
- A similar strategy could be used to develop a filter for **uniform process noise**
- Seems to be room to discover ideal strategies for accommodating those error functions

Any Questions?