Measurement Extraction for a Point Target from an Optical Sensor's Focal Plane Array: Accuracy and ROC

Qin Lu
Yaakov Bar-Shalom, IEEE AESS Distinguished Lecturer
Peter Willett and Bala Balasingam

University of Connecticut, Storrs, CT

Introduction (1/2)

- Optical sensors are widely used in target tracking systems. To date, there seems to be no existing model for measurement extraction from optical sensors that blends physics and statistics: the optics' point spread function (PSF), pixel size, dead zone and Poisson / Gaussian observed intensity levels.
- A typical measurement extractor for a point target subject to a PSF consists of two parts:
  - The first part is the target location estimator, which provides the estimate of the target location in the focal plane array (FPA) given the measured intensities from the pixels in the FPA together with the accuracy of the estimate (measurement noise variance for a tracking system).
  - The second part is the target detector, which determines whether the estimate comes from a target — this is done based on a statistical test that yields the receiver operating characteristics (ROC).

Introduction (2/2)

- In this paper, we use a realistic model in which neighboring pixels are separated by a dead zone in the FPA of an optical sensor.
  - Assuming that the spatial energy density of the target deposited in the FPA conforms to a Gaussian PSF and that the noise intensity in each pixel is zero-mean Gaussian distributed (after demeaning) with moments given by a Poisson process,
    - ⋆ we derive the Cramér Rao lower bound (CRLB) for the covariance of the estimated target location in the FPA.
    - ⋆ the optimal pixel size for minimum CRLB for target localization variance is also obtained.
  - Target detection is performed using the generalized likelihood ratio test (GLRT) to determine whether a target exists at the estimated location. The test statistic is shown to be the output of a matched filter centered at the estimated location, which is the maximum of the matched filter outputs in the FPA.

The Model (1/2)

- The spatial density of the intensity at location (x, y) from a target at $\boldsymbol{\theta} \triangleq (\xi, \eta)'$ in the FPA is given by the (circular) Gaussian PSF

$$f(x, y; \boldsymbol{\theta}) = \frac{1}{2\pi\sigma^2} \exp\left\{ -\frac{(x-\xi)^2 + (y-\eta)^2}{2\sigma^2} \right\} \tag{1}$$

  where σ is the spread of the PSF.
- Assume each pixel has a square shape with side a, and any two neighboring pixels are separated by the dead zone with width μ.

- The portion of energy deposited in the (i, j)th pixel is then

$$\begin{aligned} p_{i,j}(\boldsymbol{\theta}) &= \int_{x_i - a/2}^{x_i + a/2} \int_{y_j - a/2}^{y_j + a/2} f(x, y; \boldsymbol{\theta}) dx dy \\ &= \left[\Phi\left(\frac{x_i + a/2 - \xi}{\sigma} \right) - \Phi\left(\frac{x_i - a/2 - \xi}{\sigma} \right) \right] \\ &\times \left[\Phi\left(\frac{y_j + a/2 - \eta}{\sigma} \right) - \Phi\left(\frac{y_j - a/2 - \eta}{\sigma} \right) \right] \end{aligned} \tag{2}$$

  where $\Phi(\cdot)$ is the standard 1-D Gaussian cumulative distribution function (CDF).

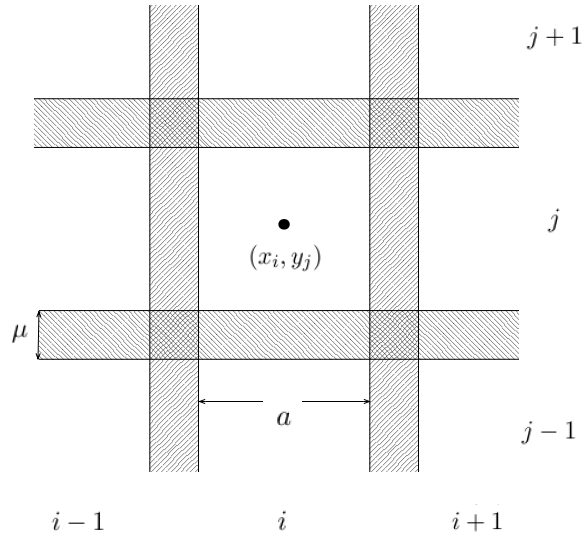

Fig. 1. Topology of the FPA (the shaded area represents the dead zone).

The Model (2/2)

- The FPA fill factor is defined as the ratio of pixel areas to the total area of the FPA and yields (approximately) the total portion of target energy deposited in the FPA

$$\psi \triangleq \frac{a^2}{(a+\mu)^2} \approx \sum_i \sum_j p_{i,j}(\boldsymbol{\theta}) \tag{3}$$

- The measured signal intensity in pixel (i, j) is

$$z_{i,j} = \mathcal{I} p_{i,j}(\boldsymbol{\theta}) + w_{i,j} \tag{4}$$

where $\mathcal{I}$ is the total energy from the target (number of photons), the noise intensity $w_{i,j}$ is assumed to be white and Gaussian with zero mean and variance proportional to the area of the pixel (this approximates a demeaned Poisson noise)[1]

$$\sigma_w^2 = s a^2 \tag{5}$$

where s is the spatial density of the noise intensity and

$$\mathbf{R} = \mathrm{diag}[\sigma_w^2] \tag{6}$$

[1] The number of photons is assumed Poisson distributed, i.e., its mean and variance are the same. The mean is subtracted via the demeaning.

Target Location Estimation

- Define the stacked vectors: $\mathbf{Z} \triangleq \{z_{i,j}\}$, $\mathbf{P}(\boldsymbol{\theta}) \triangleq \{p_{i,j}(\boldsymbol{\theta})\}$ and $\mathbf{W} \triangleq \{w_{i,j}\}$. The log-likelihood function of $\boldsymbol{\theta}$ is then

$$\lambda\left(\boldsymbol{\theta};\mathbf{Z}\right) \triangleq \ln p(\mathbf{Z}|\boldsymbol{\theta}) = -\frac{1}{2}[\mathbf{Z}-\mathcal{I}\mathbf{P}(\boldsymbol{\theta})]^{'}\mathbf{R}^{-1}[\mathbf{Z}-\mathcal{I}\mathbf{P}(\boldsymbol{\theta})]+C \tag{7}$$

  where C is an irrelevant constant.

- The maximum likelihood estimate (MLE) of $\boldsymbol{\theta}$ is obtained as

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \lambda\left(\boldsymbol{\theta};\mathbf{Z}\right) \tag{8}$$

- The Fisher information matrix (FIM) for $\boldsymbol{\theta}$ is obtained as

$$\mathbf{J}(\boldsymbol{\theta}) \triangleq \mathbb{E}\left[\nabla_{\boldsymbol{\theta}}\lambda\left(\nabla_{\boldsymbol{\theta}}\lambda\right)^{'}\right] = \mathcal{I}^{2}\left(\nabla_{\boldsymbol{\theta}}\mathbf{P}(\boldsymbol{\theta})^{'}\right)\mathbf{R}^{-1}\left(\nabla_{\boldsymbol{\theta}}\mathbf{P}(\boldsymbol{\theta})^{'}\right)^{'} \tag{9}$$

- The Cramér Rao lower bound (CRLB) is then

$$\mathcal{P}_{\mathrm{CRLB}} = \mathbf{J}(\boldsymbol{\theta})^{-1} \tag{10}$$

Simulation Results for Target Location Estimation (1/3)

- Parameters Setup: $\sigma = 1$, $s = 1$, $\boldsymbol{\theta} = [\xi \quad \eta]' = [0 \quad 0]'$.[2]
- Assume dead zone width $\mu \in \{0.01, 0.02, 0.03, ..., 0.1\}$ (in units of σ). For a given μ, we vary the pixel size a to see how the CRLB-based standard deviation σ_{CRLB} (in each coordinate) — the target location estimate accuracy — changes.
- The optimal pixel size a_{opt} that minimizes σ_{CRLB} and the corresponding fill factor $\psi(a_{\mathrm{opt}})$ for different μ are summarized in Table 1.

μ	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	0.10
a_{opt}	0.3891	0.4885	0.5575	0.6120	0.6576	0.6972	0.7322	0.7635	0.7912	0.8158
$\psi(a_{\mathrm{opt}})$	0.9505	0.9229	0.9005	0.8811	0.8637	0.8478	0.8331	0.8193	0.8063	0.7935

Table 1. Optimal pixel size for minimum CRLB

- For large pixel size, the accuracy decreases due to coarse granularity.
- For small pixel size, the fill factor is lower and the SNR decreases.

[2] We also randomized the target location within the pixel and dead zone. The resulting CRLBs did not change significantly (less than 1%.)

Simulation Results for Target Location Estimation (2/3)

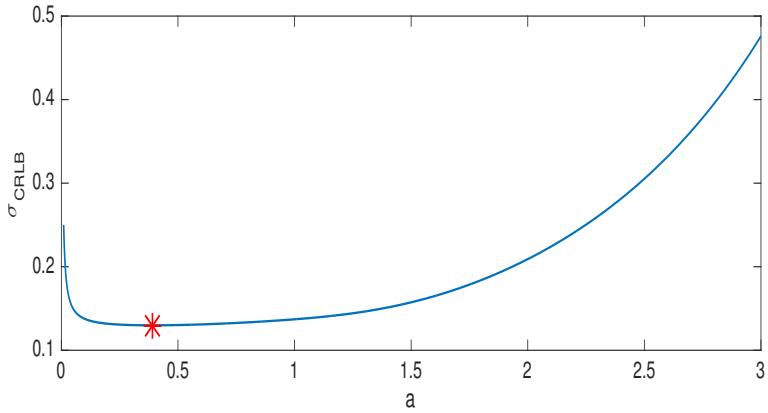

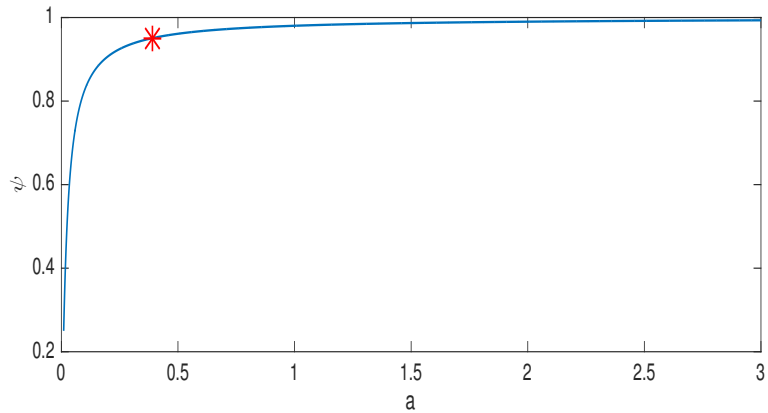

(a) $\mu = 0.01$

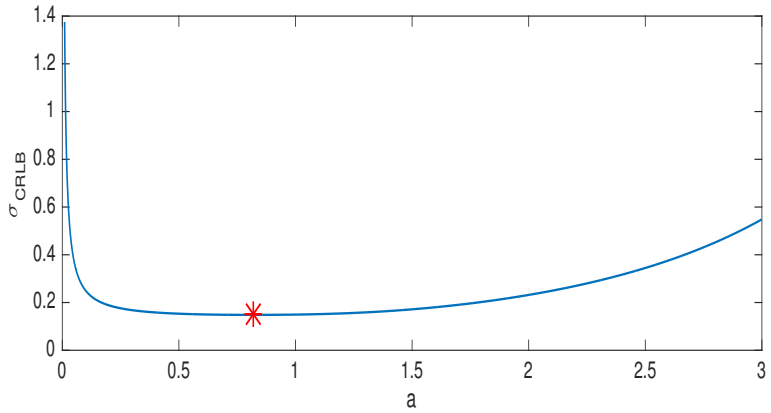

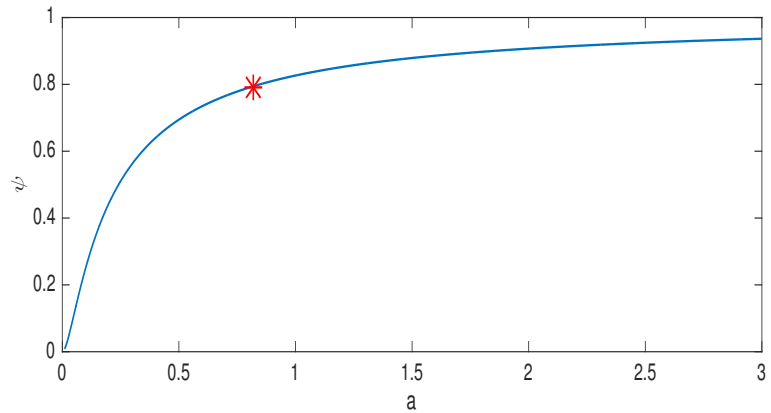

(b) $\mu = 0.10$

Fig. 2. σ_{CRLB} and fill factor ψ vs. a for different dead zone μ (all in units of σ).

The (often used) fill factor $\psi = 0.9$, while not exactly at the minimum, still yields almost the same performance.

Simulation Results for Target Location Estimation (3/3)

Efficiency of MLE vs. signal-to-noise ratio (SNR)

- The raw SNR (prior to matched filtering) is defined as the ratio of the target energy deposited in the $8\sigma \times 8\sigma$ square to the RMS value (SD) of the noise intensity in the same area

$$\gamma = \frac{99.99\%\mathcal{I} \times \psi}{\sqrt{s(8\sigma)^2} \times \psi} = \frac{99.99\%\mathcal{I}}{8\sigma\sqrt{s}} \tag{11}$$

We will convert it to dB as $10 \log_{10} \gamma$ as the above is the ratio of energies.

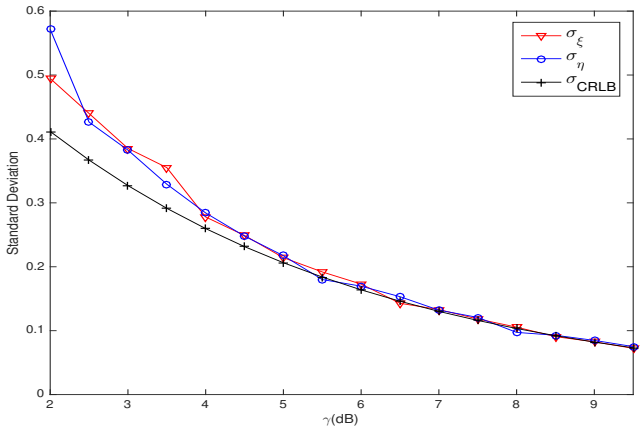

(a) $\mu = 0.01, a = 0.39$

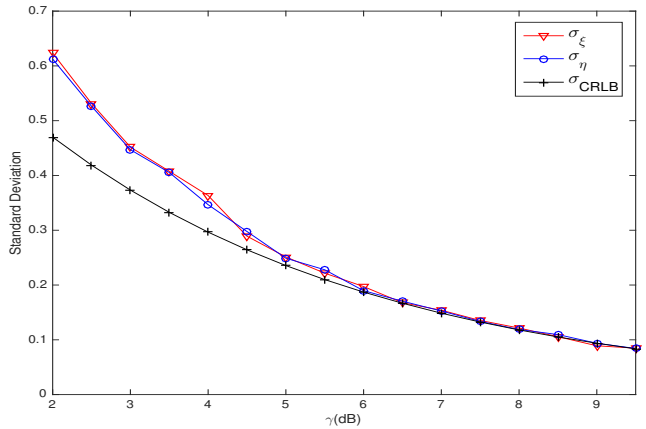

(b) $\mu = 0.1, a = 0.82$

Fig. 3. Sample SDs σ_ξ, σ_η vs. CRLB-based SDs as a function of SNR for some of the optimal (μ, a) pairs over $1,000$ Monte Carlo runs.

The MLE of the target location is efficient down to 6dB.

Target Detection

- To determine whether a target is present, the following hypothesis test problem is used:

$$\begin{aligned} \mathcal{H}_0: &\quad \mathbf{Z} = \mathbf{W} && \text{target absent} \\ \mathcal{H}_1: &\quad \mathbf{Z} = \mathcal{I}\mathbf{P}(\boldsymbol{\theta}) + \mathbf{W} && \text{target present} \end{aligned} \tag{12}$$

- The log-likelihood ratio (LLR) for these two hypotheses is given by

$$L(\mathbf{Z}) = \ln \frac{p(\mathbf{Z}|\boldsymbol{\theta}, \mathcal{H}_1)}{p(\mathbf{Z}|\mathcal{H}_0)} = \ln \frac{\frac{1}{\sqrt{|2\pi\mathbf{R}|}} \exp\left\{-\frac{1}{2}\left(\mathbf{Z} - \mathcal{I}\mathbf{P}(\boldsymbol{\theta})\right)' \mathbf{R}^{-1}\left(\mathbf{Z} - \mathcal{I}\mathbf{P}(\boldsymbol{\theta})\right)\right\}}{\frac{1}{\sqrt{|2\pi\mathbf{R}|}} \exp\left\{-\frac{1}{2}\mathbf{Z}'\mathbf{R}^{-1}\mathbf{Z}\right\}}$$

$$= \mathcal{I}\mathbf{P}(\boldsymbol{\theta})' \mathbf{R}^{-1}\mathbf{Z} - \frac{1}{2}\mathcal{I}^2 \mathbf{P}(\boldsymbol{\theta})' \mathbf{R}^{-1}\mathbf{P}(\boldsymbol{\theta}) \tag{13}$$

- As the true value of $\boldsymbol{\theta}$ is not available, we use the MLE (8) in (13), which leads to the generalized likelihood ratio test (GLRT). Then (13) becomes

$$L_G(\mathbf{Z}) = \mathcal{I}\mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{R}^{-1}\mathbf{Z} - \frac{1}{2}\mathcal{I}^2 \mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{R}^{-1}\mathbf{P}(\hat{\boldsymbol{\theta}}) \tag{14}$$

- As $\mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{R}^{-1}\mathbf{P}(\hat{\boldsymbol{\theta}})$ is not related to $\mathbf{Z}$, we choose the following test statistic

$$T(\mathbf{Z}) = \mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{Z} \tag{15}$$

which is the matched filter (MF) output at the estimated target location.

Distribution of the test statistic under $\mathcal{H}_0$ (1/3)

- Under $\mathcal{H}_0$, $\mathbf{Z} = \mathbf{W}$. The test statistic is

$$T(\mathbf{Z}) = \mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{Z} = \max_{\boldsymbol{\theta}} \left\{ \mathbf{P}(\boldsymbol{\theta})' \mathbf{W} \right\} \tag{16}$$

- Consider that the $\boldsymbol{\theta}$ that maximizes $\mathbf{P}(\boldsymbol{\theta})' \mathbf{W}$ is chosen from the list with M elements $\boldsymbol{\Theta} \triangleq \left\{ \boldsymbol{\theta}^j \right\}_{j=1}^{M}$. Therefore, the CDF of $T(\mathbf{Z})$ at value t is

$$\Pr\left(T(\mathbf{Z}) < t | \mathcal{H}_0\right) = \Pr\left(\max_{\boldsymbol{\theta}^j \in \boldsymbol{\Theta}} \left\{ \mathbf{P}(\boldsymbol{\theta}^j)' \mathbf{W} \right\} < t | \mathcal{H}_0 \right) = \prod_{j=1}^{M} \Pr\left(\mathbf{P}(\boldsymbol{\theta}^j)' \mathbf{W} < t | \mathcal{H}_0 \right) \tag{17}$$

  where it is assumed that all the M MF outputs $\mathbf{P}(\boldsymbol{\theta}^j)' \mathbf{W}$ are mutually independent.

- For any $\boldsymbol{\theta}$ in the FPA (except those on the edges), if the FPA is large enough to cover all the target energy,

$$\mathbf{P}(\boldsymbol{\theta})' \mathbf{P}(\boldsymbol{\theta}) = \sum_i \sum_j (p_{i,j}(\boldsymbol{\theta}))^2 \approx C(\mu, a) \tag{18}$$

  where $C(\mu, a)$ is a constant depending on the μ and a.

Distribution of the test statistic under $\mathcal{H}_0$ (2/3)

- We assume the M MF outputs are independently identically distributed (i.i.d.) as

$$\mathbf{P}(\boldsymbol{\theta}^j)' \mathbf{W} \sim \mathcal{N}(0, C(\mu, a)\sigma_w^2) \tag{19}$$

  The number M is determined later.

- Therefore, the CDF of the test statistic is

$$\Pr\left(T(\mathbf{Z}) < t | \mathcal{H}_0\right) \approx \left(\Phi\left(\frac{t}{\sigma_w \sqrt{C(\mu, a)}}\right)\right)^M \tag{20}$$

- Given a false alarm probability P_{fa}, the threshold τ is determined from

$$\left(\Phi\left(\frac{\tau}{\sigma_w \sqrt{C(\mu, a)}}\right)\right)^M \approx 1 - P_{\mathrm{fa}} \tag{21}$$

  which yields

$$\tau \approx \sigma_w \sqrt{C(\mu, a)}\, \Phi^{-1}\left((1 - P_{\mathrm{fa}})^{1/M}\right) \tag{22}$$

Distribution of the test statistic under $\mathcal{H}_0$ (3/3)

- The only unknown variable in (22) is M. An empirical way of choosing M that we found satisfactory is

$$M = \left(\left\lceil \frac{N(a+\mu)}{\sigma} \right\rceil \right)^2 \tag{23}$$

which means that we assume there are $\left\lceil \frac{N(a+\mu)}{\sigma} \right\rceil$ equivalent independent matched filter outputs in each dimension (N is the number of pixels in each dimension of the FPA).

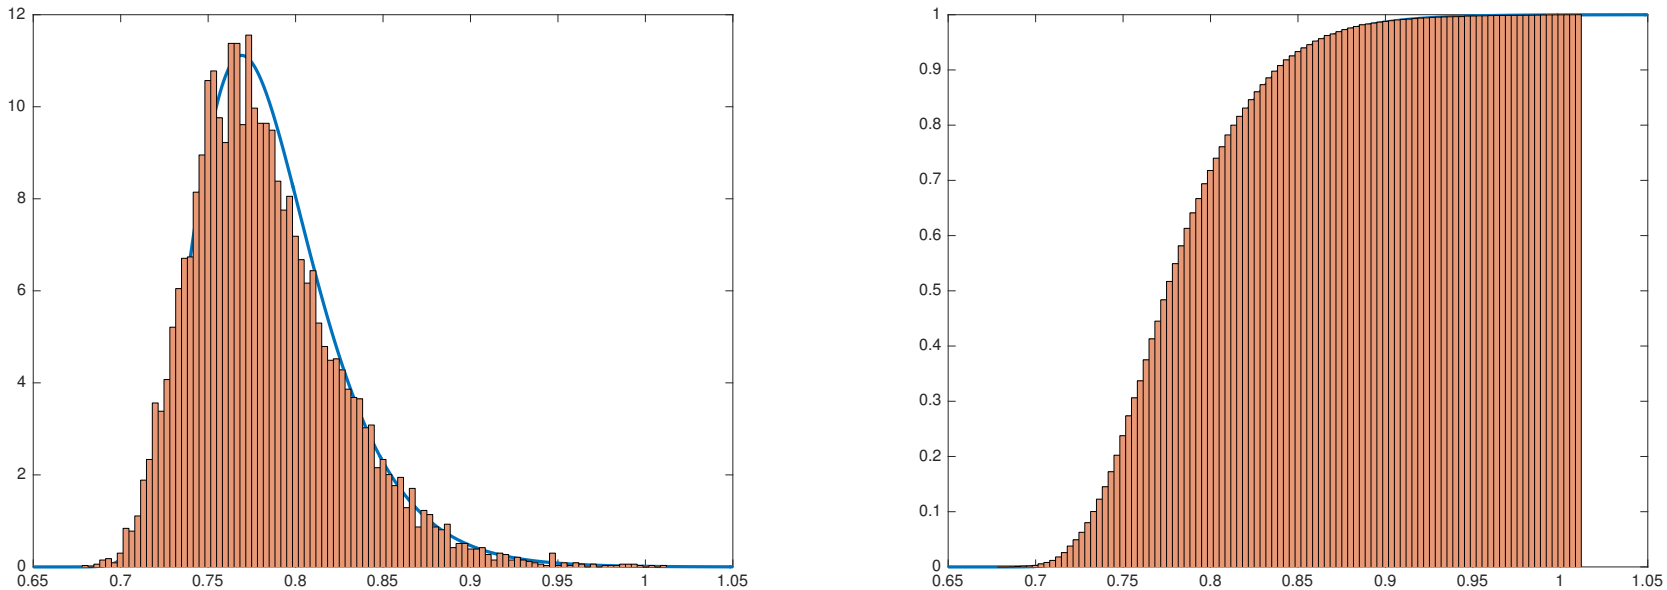

(a) PDF (b) CDF

Fig. 4. Histograms from 10^4 Monte Carlo runs and corresponding PDF and CDF of $T(\mathbf{Z})$ based on (20) under $\mathcal{H}_0$ when $N = 1000$ (1Mpx FPA), $\mu = 0.1$, $a = 0.82$ and $M = 840,000$.

Distribution of the test statistic under $\mathcal{H}_1$ (1/2)

- Under $\mathcal{H}_1$, the test statistic becomes

$$T(\mathbf{Z}) = \mathbf{P}(\hat{\boldsymbol{\theta}})' \mathbf{Z} = \mathbf{P}(\hat{\boldsymbol{\theta}})' \left(\mathcal{I}\mathbf{P}(\boldsymbol{\theta}) + \mathbf{W}\right) \tag{24}$$

- Taking the first order Taylor expansion of $\mathbf{P}(\hat{\boldsymbol{\theta}})$ around $\boldsymbol{\theta}$ yields

$$\mathbf{P}(\hat{\boldsymbol{\theta}}) \approx \mathbf{P}(\boldsymbol{\theta}) + \left(\nabla_{\boldsymbol{\theta}} \mathbf{P}'\right)' \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right) \tag{25}$$

- Define the estimation error as $\tilde{\boldsymbol{\theta}} \triangleq \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}$. If $\hat{\boldsymbol{\theta}}$ is efficient,

$$\tilde{\boldsymbol{\theta}} \sim \mathcal{N}\left(\mathbf{0}, \mathcal{P}_{\mathrm{CRLB}}\right) \tag{26}$$

  which is assumed to be independent of the measurement noise $\mathbf{W}$ ($\hat{\boldsymbol{\theta}}$ depends on the small peak MF region while $\mathbf{W}$ contains the noise in all the pixels).
- The test statistic (24) is approximated as

$$\begin{aligned} T(\mathbf{Z}) &\approx \left(\mathbf{P}(\boldsymbol{\theta}) + \left(\nabla_{\boldsymbol{\theta}} \mathbf{P}'\right)' \tilde{\boldsymbol{\theta}}\right)' \left(\mathcal{I}\mathbf{P}(\boldsymbol{\theta}) + \mathbf{W}\right) \\ &= \mathcal{I}\mathbf{P}(\boldsymbol{\theta})' \mathbf{P}(\boldsymbol{\theta}) + \mathbf{P}(\boldsymbol{\theta})\mathbf{W} + \mathcal{I}\tilde{\boldsymbol{\theta}}' \left(\nabla_{\boldsymbol{\theta}} \mathbf{P}'\right) \mathbf{P}(\boldsymbol{\theta}) + \tilde{\boldsymbol{\theta}}' \left(\nabla_{\boldsymbol{\theta}} \mathbf{P}'\right) \mathbf{W} \end{aligned} \tag{27}$$

Distribution of the test statistic under $\mathcal{H}_1$ (2/2)

- As $\left(\nabla_{\boldsymbol{\theta}}\mathbf{P}^{'}\right)\mathbf{P}(\boldsymbol{\theta}) \approx \mathbf{0}$ (this is an even function $\times$ an odd function), the test statistic is then

$$T(\mathbf{Z}) \approx \mathcal{I}\mathbf{P}(\boldsymbol{\theta})^{'}\mathbf{P}(\boldsymbol{\theta}) + \mathbf{P}(\boldsymbol{\theta})\mathbf{W} + \tilde{\boldsymbol{\theta}}^{'}\left(\nabla_{\boldsymbol{\theta}}\mathbf{P}^{'}\right)\mathbf{W} \tag{28}$$

- The first two moments of the above (under $\mathcal{H}_1$) are

$$m_T \triangleq \mathbb{E}\left[T(\mathbf{Z})\,|\mathcal{H}_1\right] \approx \mathcal{I}\mathbf{P}(\boldsymbol{\theta})^{'}\mathbf{P}(\boldsymbol{\theta}) \approx \mathcal{I}C(\mu, a) \tag{29}$$

$$\sigma_T^2 \triangleq \mathbb{E}\left[\left(T(\mathbf{Z}) - \mathbb{E}\left[T(\mathbf{Z})\,|\mathcal{H}_1\right]\right)^2 |\mathcal{H}_1\right]$$

$$\approx \mathbb{E}\left[\mathbf{P}(\boldsymbol{\theta})^{'}\mathbf{W}\mathbf{W}^{'}\mathbf{P}(\boldsymbol{\theta})\right] + \mathbb{E}\left[\tilde{\boldsymbol{\theta}}^{'}\nabla_{\boldsymbol{\theta}}\mathbf{P}^{'}\mathbf{W}\mathbf{W}^{'}\left(\nabla_{\boldsymbol{\theta}}\mathbf{P}^{'}\right)^{'}\tilde{\boldsymbol{\theta}}\right] \approx \left(sa^2C(\mu, a) + \frac{2s^2a^4}{\mathcal{I}^2}\right)\mathbf{I}_2 \tag{30}$$

- Therefore, $T(\mathbf{Z})|\mathcal{H}_1 \sim \mathcal{N}\left(m_T, \sigma_T^2\right)$. The matched filter output SNR is then defined as, similarly to (11)

$$\gamma_{\mathrm{MF}} = \frac{m_T}{\sigma_T} \tag{31}$$

- The detection probability given the threshold τ (22) is

$$P_{\mathrm{d}} = \mathrm{Pr}\left(T(\mathbf{Z}) > \tau|\mathcal{H}_1\right) = 1 - \Phi\left(\frac{\tau - m_T}{\sigma_T}\right) \tag{32}$$

Simulation results for the detection performance (1/2)

- First, for a given dead zone width μ, raw SNR γ and false alarm probability P_{fa}, the relationship between the detection probability P_{d} and the pixel size a is shown below.
  - There is a small mismatch between the analytic P_{d} (32) and simulation-based P_{d} for small SNR.
  - P_{d} increases and then decreases as a increases.

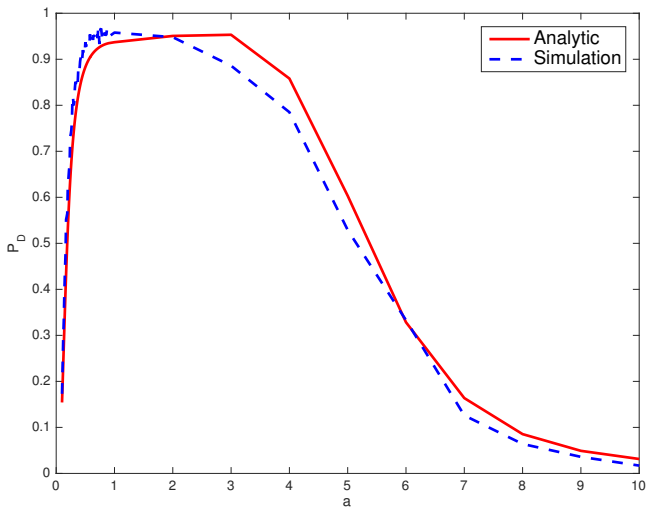

(a) $\gamma = 5$dB

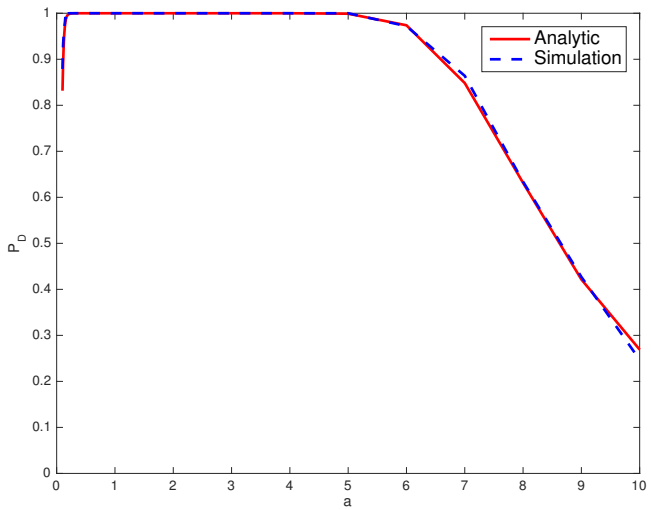

(b) $\gamma =$7dB

Fig. 5. Detection Probability P_{d} vs. pixel size a for different (raw) SNRs when $P_{\mathrm{fa}} = 10^{-4}$ and $\mu = 0.1$.

Simulation results for the detection performance (2/2)

- Second, the receiver operating characteristic (ROC) curves (P_{d} vs. P_{fa}) are plotted for given μ, a at different (raw) SNRs.

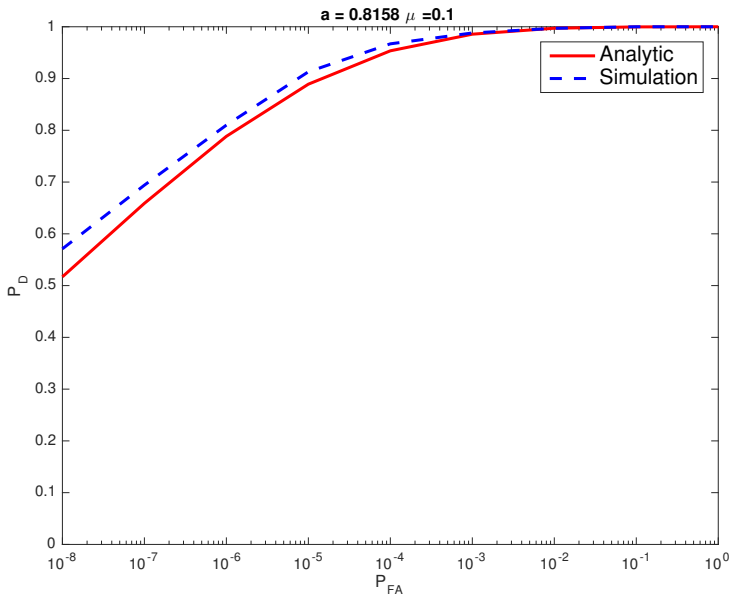

(a) $\gamma = 5$dB

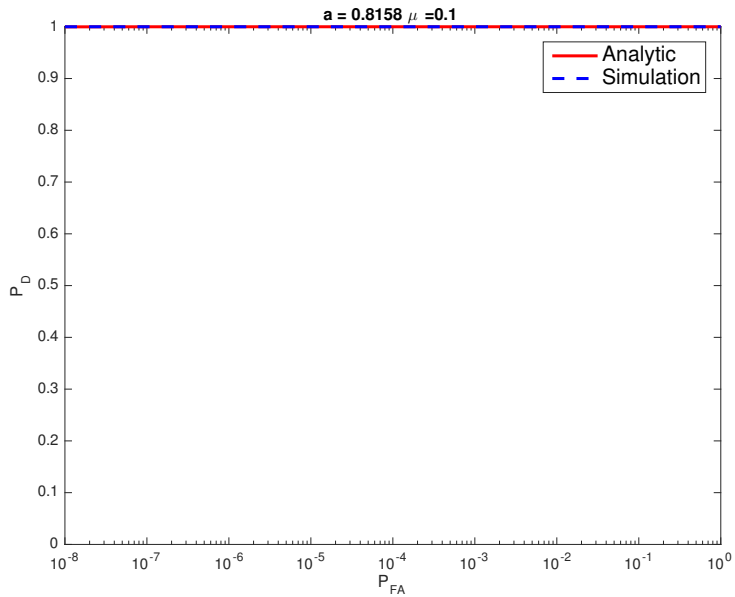

(b) $\gamma =$7dB

Fig. 6. ROC curves for $\mu = 0.1, a = 0.82$.

Conclusions

- This work derived the measurement extractor for a point target with unknown location in an optical sensor's FPA.
- Firstly, we derived the CRLB for covariance of the estimated location based on the assumptions that energy density of the target deposited in the FPA conforms to a Gaussian PSF and that the noise variance is proportional to the pixel area. Simulation results show that there is an optimal pixel size that minimizes the CRLB for a given dead zone width. The MLE is shown to be efficient at higher SNRs.
- Secondly, we investigated target detection: the GLRT statistic uses the maximum MF output in the FPA. The distributions of the test statistic under both (target absent/present) hypotheses are explored. It is shown that the detection probability is remarkably flat around the optimal pixel size for minimum estimation errors for a certain false alarm probability and SNR.