

On achievable accuracy for range finder localization

How precise can a localization method be?

Given a sensor, there is an hard limit: the Cramér–Rao bound.

Summary:

1. Definition of the Cramér–Rao bound.
2. Application to range-finder localization.
3. A model for unstructured environments.
4. Experiments: predicting the error of the ICP.

The Cramér–Rao Bound

- A localization algorithm A is an estimator of \mathbf{x} given function A from the sensor data \mathbf{z} : $\hat{\mathbf{x}} = A(\mathbf{z})$
- The estimator is unbiased if $E\{A(\mathbf{z})\} = \mathbf{x}$.


Cramér–Rao inequality: For any unbiased estimator $\hat{\mathbf{x}}$,

$$\text{Cov}(\hat{\mathbf{x}}) \geq (\mathcal{I}(\mathbf{x}))^{-1}$$

The $n \times n$ symmetric matrix $\mathcal{I}(\mathbf{x})$, called Fisher's information matrix, is defined as

$$\mathcal{I}(\mathbf{x}) = E_z \left\{ \frac{\partial \log p(\mathbf{z}, \mathbf{x})}{\partial \mathbf{x}^T} \frac{\partial \log p(\mathbf{z}, \mathbf{x})}{\partial \mathbf{x}^T}^T \right\}$$

Fisher, Mahalanobis, Rao, Cramér



C. R. Rao and Harald Cramér, 1978

- Sir Ronald Ayimer Fisher (1890-1962), England
- Prasanta Chandra Mahalanobis (1893-1972), India
- Harald Cramér (1893-1985), Sweden
- Calyampudi Radhakrishna Rao (1920-), India

The Cramér–Rao lower bound

- The CRB is valid only for unbiased estimators.
 - Consider the localization algorithm

$$A(\mathbf{z}) = (42, 42, 42^\circ)$$


It has a covariance of 0.

- Biased estimators might have a lower Mean Square Error:

$$\text{MSE}_A(x) = \text{var}_A(x) + \text{bias}_A^2(x)$$

- The CRB It is not tight for non-Gaussian problems.
 - Localization is not a problematic problem.
 - Non-linear bounds are very very hard to derive.

CRB for localization




The likelihood function is

$$p(\mathbf{z}|\mathbf{x}) = \prod_i \mathcal{N}(\tilde{\rho}_i - r(t, \theta + \varphi_i), \sigma^2)$$

$r(\mathbf{p}, \psi)$ is the “raytracing function”.

CRB for localization

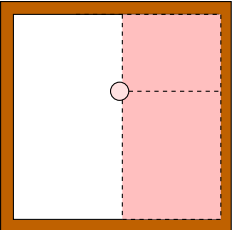
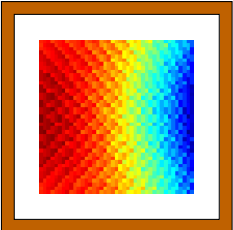
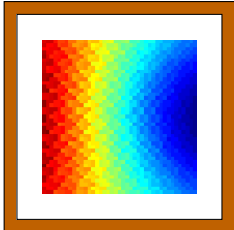
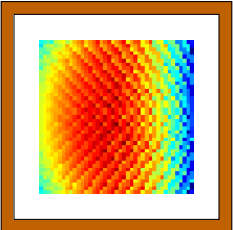
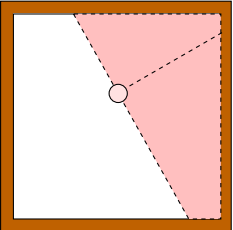
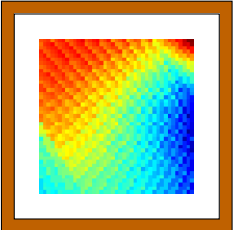
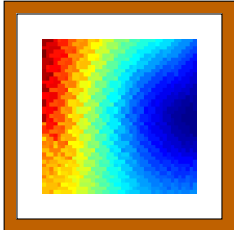
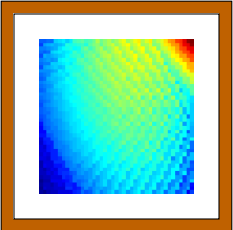
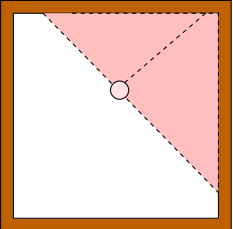
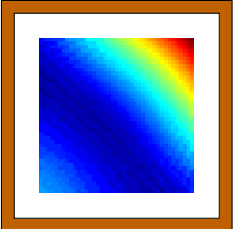
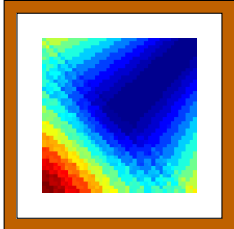
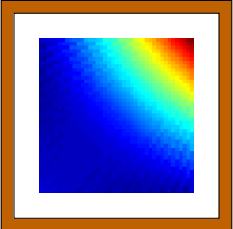


Fisher's information matrix is


$$\mathbf{I}(\mathbf{x}) = \frac{1}{\sigma^2} \sum_i \begin{bmatrix} \frac{\mathbf{v}(\alpha_i)\mathbf{v}(\alpha_i)^T}{\cos^2 \beta_i} & r_i \frac{\tan \beta_i}{\cos \beta_i} \mathbf{v}(\alpha_i) \\ * & r_i^2 \tan^2 \beta_i \end{bmatrix}$$

α is the surface direction, $\beta \triangleq \alpha_i - (\theta + \varphi_i)$

Results in a square environment

Orientation	$\det(\text{cov}(\hat{\mathbf{x}}))$	$\det(\text{cov}(\hat{\mathbf{t}}))$	$\text{var}(\hat{\theta})$
			
			
			

Kernel of $\mathcal{I}(x)$

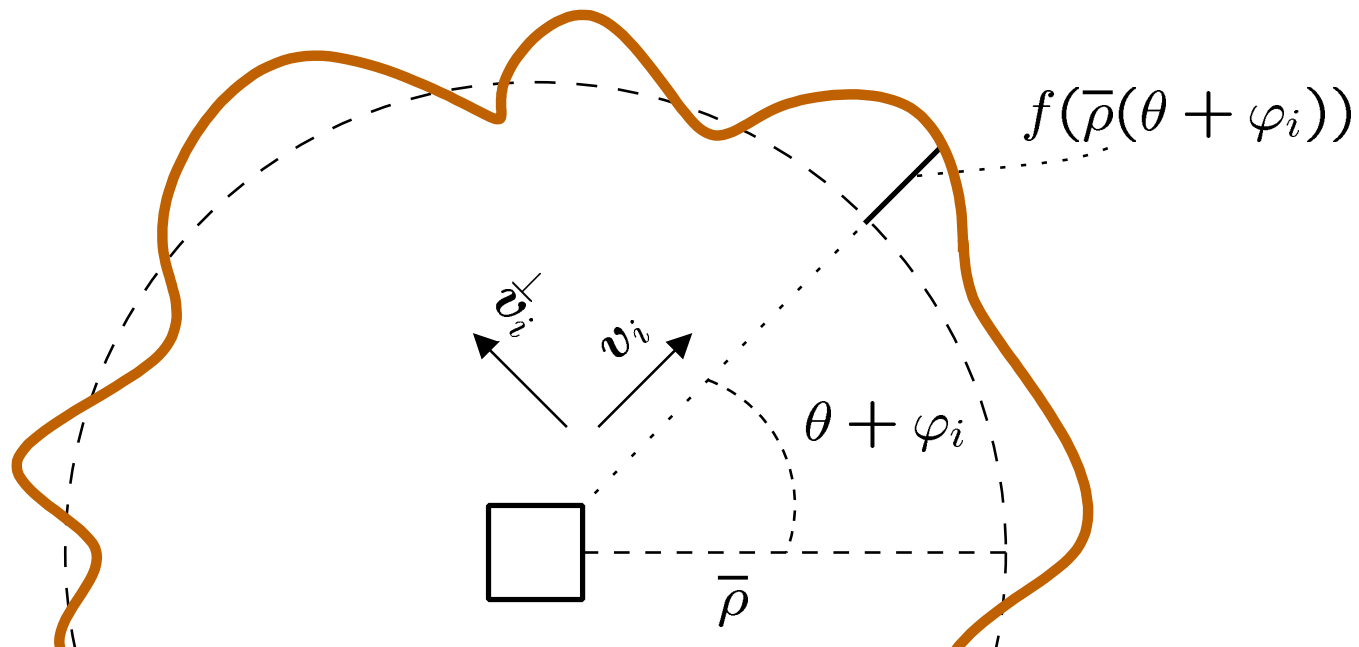


In under-constrained situations, the kernel of $\mathcal{I}(x)$ gives the direction of uncertainty.

A model for unstructured environments

Consider the environment (defined through the raytracing function)

$$r(\mathbf{0}, \phi) = \bar{\rho} + f(\phi\bar{\rho}) \quad |f| \ll \bar{\rho}$$



A model for unstructured environments

For this model, the following approximation holds:


$$\text{var}(\hat{\theta}) \geq \frac{(\sigma/\bar{\rho})^2}{n} \frac{1}{\mathcal{C}}$$
$$\lambda_{\min} \text{cov}(\hat{\mathbf{t}}) \geq 2 \frac{\sigma^2}{n} \frac{1}{1 + \mathcal{C}}$$

where $\mathcal{C} \triangleq E_f \{ f'^2 \}$.

- Localization is easier with features (high values of \mathcal{C}).
- For a circle: $\mathcal{C} = 0$. The uncertainty on θ is infinite, while the uncertainty for \mathbf{t} is still finite.
- The accuracy for θ depends on the “normalized” noise $\sigma/\bar{\rho}$ (invariance to scale)
- $\text{cov}(\hat{\mathbf{t}})$ does not depend on $\bar{\rho}$.

Experiments

The CRB predicts very well the uncertainty of the ICP in localization, while it is a weak bound for scan matching.



Asymmetric situation: square environment (5m side),
FOV = 180° , $\mathbf{x} = (-2m, 2m, 30^\circ)$

Future work

- Localization is a finite-dimensional problem, while SLAM is infinite-dimensional.
- First explore *mapping*: the key is to choose a good representation for the map.
 - Polygonal environment: very easy to obtain the CRB for the map.
 - Occupancy grids: inference is tricky.
 - Splines.
 - Gaussian processes.

Predicting localization accuracy in a filter

Update equation of a Bayesian filter:

$$p(\mathbf{x}_t|t) \propto \underbrace{p(\mathbf{z}_t|\mathbf{x}_t)}_{\text{likelihood}} \underbrace{\int p(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{u}_t)p(\mathbf{x}_{t-1}|t-1)d\mathbf{x}_{t-1}}_{p(\mathbf{x}_t|t-1)}$$

If everything is Gaussian:

$$\Sigma_{\mathbf{x}_t|t} = \left(\Sigma_{\mathbf{z}_t|\mathbf{x}_t}^{-1} + \left(\Sigma_{\mathbf{x}_{t-1}|t-1} + \Sigma_{\mathbf{u}} \right)^{-1} \right)^{-1}$$

If your localization algorithm is efficient:

$$\Sigma_{\mathbf{x}_t|t} \simeq \left(\mathcal{I}(\mathbf{x}_t) + \left(\Sigma_{t-1} + \Sigma_{\mathbf{u}} \right)^{-1} \right)^{-1}$$