On achievable accuracy for pose tracking

Andrea Censi

Ph.D. student,
Control & Dynamical Systems,
California Institute of Technology

Advisor: Richard Murray

On achievable accuracy for pose tracking

"pose tracking" $\triangleq$ recovering the incremental robot displacement given the output of a relative sensor

Central questions:

- How precise can an algorithm be, in principle?
- Do existing algorithms achieve this precision?

Background

- **Long-term goal:** characterizing the performance of dense SLAM (i.e., non-landmark based).
- **Focus:** pose-tracking algorithms in general, "scan matching" (ICP, NDT, ...) in particular.

- Many scan matching algorithms, but few characterizations:
  - Is the algorithm "optimal" in some sense?
  - If not, are there upper/lower bounds on its error?

- Approaches:
  1. Benchmarking (NIST, EURON sig, RAWSEEDS, etc.)
  2. Theoretical analysis:
     - Landmark-based, extended Kalman Filter setting.
       (See works of Roumeliotis, Mourikis, Huang, and many others...)
     - What about algorithm-independent results, infinite-dimensional maps, dense sensors?

Summary

1. Cramér-Rao bound
2. Localization vs. pose tracking vs. mapping vs. SLAM
   - Localization is finite-dimensional and therefore “easy” to analyze
   - For pose tracking, one must deal with the unknown infinite-dimensional world (like SLAM)
3. Main results on pose tracking
   - Cramer-Rao bounds on the accuracy for pose tracking
   - Results independent of the map prior and representation
   - Valid for any relative sensor
4. Experiments with scan matching
   - ICP’s covariance is very close to the bound (almost optimal)
5. Ideas used in the proofs
   - intrinsic estimation, shape spaces, etc.

Cramér-Rao bound

Cramér-Rao bound. Suppose we are estimating the state $\boldsymbol{x}$ based on the observations $\boldsymbol{y}$. Then, for any unbiased estimator $\hat{\boldsymbol{x}}$:

$$cov[\hat{\boldsymbol{x}}] \geq \left(\boldsymbol{\mathcal{I}}[\boldsymbol{x}]\right)^{-1} \tag{1}$$

where $\boldsymbol{\mathcal{I}}\left[\boldsymbol{x}\right]$ is the Fisher Information Matrix (FIM):

$$\boldsymbol{\mathcal{I}}[\boldsymbol{x}] = \mathbb{E}\left\{\frac{\partial \log p(\boldsymbol{y}, \boldsymbol{x})}{\partial \boldsymbol{x}}^{\mathrm{T}} \frac{\partial \log p(\boldsymbol{y}, \boldsymbol{x})}{\partial \boldsymbol{x}}\right\} \tag{2}$$

Notes:

- Elegant, no assumptions on the algorithm.
- It can also be stated for biased estimators, and in a Bayesian context.
- Strict for linear problems, and nonlinear problems at high SNR ratios.

Terminology

(Arbitrary) naming of problems:

Localization: estimate the pose in a <u>known</u> map.
Pose tracking: estimate the incremental displacement in a <u>unknown</u> map.
Mapping: given the poses, estimate the map.
SLAM: estimate both poses and map.

Terminology

(Arbitrary) naming of problems:

Localization: estimate the pose in a known map.

$\uparrow$ finite dimensional $\quad (\mathcal{I}[\boldsymbol{q}] \in \mathbb{R}^{3\times 3})$

$\downarrow$ infinite dimensional $\quad (\mathcal{I}[\boldsymbol{q}, \text{map}]$ infinite dimensional)

Pose tracking: estimate the incremental displacement in a unknown map.
Mapping: given the poses, estimate the map.
SLAM: estimate both poses and map.

Terminology

(Arbitrary) naming of problems:

Localization: estimate the pose in a known map.

$\uparrow$ finite dimensional $(\mathcal{I}[\boldsymbol{q}] \in \mathbb{R}^{3\times 3})$

$\downarrow$ infinite dimensional $(\mathcal{I}[\boldsymbol{q}, \text{map}]$ infinite dimensional)

Pose tracking: estimate the incremental displacement in a unknown map.
Mapping: given the poses, estimate the map.
SLAM: estimate both poses and map.

- This distinction is useful for the *analysis* of the problems, not the implementation.
- Everything is finite-dimensional if using landmarks.

Setup: one-shot pose tracking

Setup: Take observations $\boldsymbol{y}^a, \boldsymbol{y}^b$ at the poses $\boldsymbol{q}_*, \boldsymbol{q}_* \oplus \boldsymbol{\delta}$:

$$\begin{aligned} \boldsymbol{y}^a &= \boldsymbol{r}(\boldsymbol{q}_*, \boldsymbol{w}) + \text{noise} \\ \boldsymbol{y}^b &= \boldsymbol{r}(\boldsymbol{q}_* \oplus \boldsymbol{\delta}, \boldsymbol{w}) + \text{noise} \end{aligned}$$

Question: How precise can the estimate of $\boldsymbol{\delta}$ be?

Setup: one-shot pose tracking

Setup: Take observations $\boldsymbol{y}^a, \boldsymbol{y}^b$ at the poses $\boldsymbol{q}_*, \boldsymbol{q}_* \oplus \boldsymbol{\delta}$:

$$\begin{aligned} \boldsymbol{y}^a &= \boldsymbol{r}(\boldsymbol{q}_*, \boldsymbol{w}) + \text{noise} \\ \boldsymbol{y}^b &= \boldsymbol{r}(\boldsymbol{q}_* \oplus \boldsymbol{\delta}, \boldsymbol{w}) + \text{noise} \end{aligned}$$

Question: How precise can the estimate of δ be?

The analysis is harder than for localization:

- The world $\boldsymbol{w}$ is unknown and part of the state.
- The world $\boldsymbol{w}$ is infinite dimensional ($\mathcal{I}[\boldsymbol{\delta}, \boldsymbol{w}]$ is an infinite matrix)
- How to get results independent of the <u>representation</u> for $\boldsymbol{w}$?
- How to correctly state the <u>prior</u> for $\boldsymbol{w}$?
    - representation $\simeq$ prior: map composed of segments assigns zero prior to smooth curves, and vice versa.

Approach: look for the best bounds we can give, independently of $\boldsymbol{w}$.

Main results

Cramer-Rao bound: $\mathrm{cov}[\delta] \geq \boldsymbol{\mathcal{I}}[\delta]^{-1}$, but $\boldsymbol{\mathcal{I}}[\delta]$ depends on the world...

Main results

Cramer-Rao bound: $\operatorname{cov}[\delta] \geq \mathcal{I}[\delta]^{-1}$, but $\mathcal{I}[\delta]$ depends on the world...

Theorem: Without additional information on the prior for the world, the sensor model or the size of δ, the best bounds on $\mathcal{I}[\delta]$ are:

$$\mathbf{0} \leq \mathcal{I}[\delta] \leq \left(\mathcal{I}_{\text{loc}}(\boldsymbol{q}_*)^{-1} + \mathcal{I}_{\text{loc}}(\boldsymbol{q}_* \oplus \delta)^{-1}\right)^{-1}$$

where $\mathcal{I}_{\text{loc}}$ is the information matrix for *localization* on a *known* map.
Corollary: this bound is strict for $\delta \to \mathbf{0}$.

Main results

Cramer-Rao bound: $\mathrm{cov}[\delta] \geq \mathcal{I}[\delta]^{-1}$, but $\mathcal{I}[\delta]$ depends on the world...

Theorem: Without additional information on the prior for the world, the sensor model or the size of δ, the best bounds on $\mathcal{I}[\delta]$ are:

$$\mathbf{0} \leq \mathcal{I}[\delta] \leq \left(\mathcal{I}_{\mathrm{loc}}(\boldsymbol{q}_*)^{-1} + \mathcal{I}_{\mathrm{loc}}(\boldsymbol{q}_* \oplus \boldsymbol{\delta})^{-1}\right)^{-1}$$

where $\mathcal{I}_{\mathrm{loc}}$ is the information matrix for *localization* on a *known* map.
Corollary: this bound is strict for $\delta \to \mathbf{0}$.

What makes this result interesting:

- The analysis for pose-tracking (infinite-dimensional) is reduced to localization (finite-dimensional).
- As $\delta \to \mathbf{0}$, the prior of the map does not matter ("the data is the model").
- Valid for any relative sensor (range-finders, sonars, cameras), any algorithm.

Results for scan matching (ICP)

1. The bound is approximately strict for range finders.

Results for scan matching (ICP)

1. The bound is approximately strict for range finders.
2. This bound does not predict statistical bias.

Results for scan matching (ICP)

1. The bound is approximately strict for range finders.
2. This bound does not predict statistical bias.
3. The bound does not predict algorithmic errors (e.g., local minima).

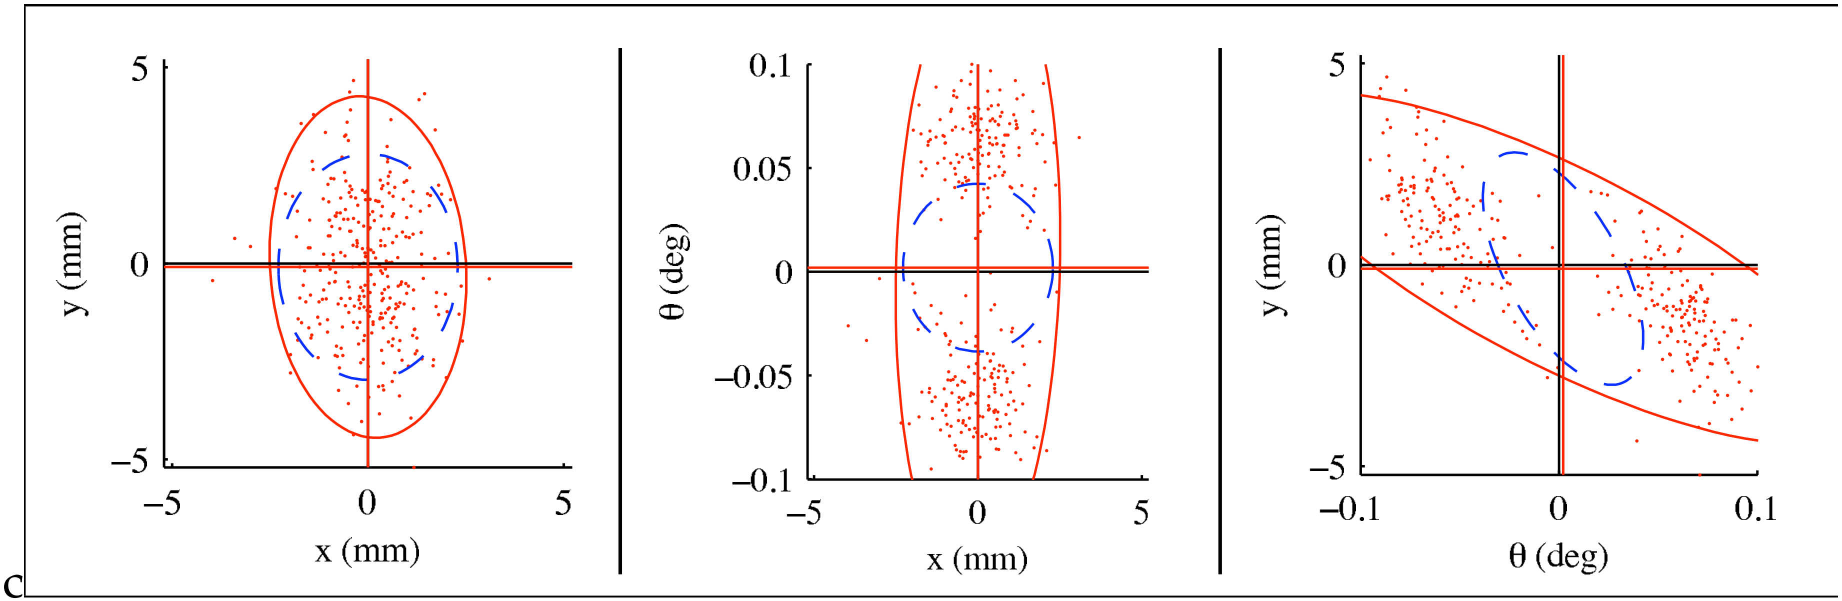

Estimation error versus error bound on the (x, y), (x, θ), (y, θ) planes.

Ideas used in the paper

- Calculus of Fisher Information Matrices
  - (not to be confused with Information Filter, etc.)
- Intrinsic estimation (statistics compatible with differential geometry)
  - Any result must be independent of the parameterization chosen by a particular algorithm.
- Statistical geometry (statistics on shapes)
  - The of worlds can be factorized as

$$\text{World} = \text{Shape} \times \text{Pose}$$

  - ex: one defines formally a "relative sensor" based on this factorization.

Conclusions

- We need stronger analytical results on dense localization/SLAM algorithms:
  - lower bounds on the error (depend on the problem)
  - upper bounds on the error (depend on the solution)
- The analysis for pose tracking has the same issues of SLAM:
  - The world is unknown and infinite-dimensional.
  - How to find results independent of the representation of the world?
  - How to express correctly the prior for the world?
- This paper found a way to sidestep these issues.
  - but these must be dealt with upfront for meaningful analysis of infinite-dimensional SLAM.

Analysis for dense localization

- Localization on known map is a *finite dimensional* problem, $\boldsymbol{x} = \boldsymbol{q} = (x, y, \theta)$.
- What is the achievable accuracy on estimating the pose?

  The FIM for a single pose $\boldsymbol{q} = (x, y, \theta)$ is

$$\boldsymbol{\mathcal{I}}[\boldsymbol{q}] = \sum_i^n \frac{1}{\sigma_i^2 \cos^2 \beta_i} \begin{bmatrix} \boldsymbol{v}(\alpha_i)\boldsymbol{v}(\alpha_i)^{\mathrm{T}} & r_i \sin \beta_i \boldsymbol{v}(\alpha_i) \\ * & r_i^2 \sin^2 \beta_i \end{bmatrix} + \text{prior}$$

  - σ_i is the noise for the *i*-th reading.
  - $v(\alpha_i)$ is the surface normal at the point intercepted by the *i*-th reading.
  - β_i is the incidence angle.

Analysis for dense localization

- Localization on known map is a *finite dimensional* problem, $\boldsymbol{x} = \boldsymbol{q} = (x, y, \theta)$.
- What is the achievable accuracy on estimating the pose?

  The FIM for a single pose $\boldsymbol{q} = (x, y, \theta)$ is

$$\mathcal{I}[\boldsymbol{q}] = \sum_{i}^{n} \frac{1}{\sigma_i^2 \cos^2 \beta_i} \begin{bmatrix} \boldsymbol{v}(\alpha_i)\boldsymbol{v}(\alpha_i)^{\mathrm{T}} & r_i \sin \beta_i \boldsymbol{v}(\alpha_i) \\ * & r_i^2 \sin^2 \beta_i \end{bmatrix} + \text{prior}$$

  - σ_i is the noise for the *i*-th reading.
  - $v(\alpha_i)$ is the surface normal at the point intercepted by the *i*-th reading.
  - β_i is the incidence angle.

- Do algorithms achieve this limit?
  - Not, e.g. ICP, or similar approaches, unless suitably modified.
  - Not the usual implementations of particle filters — they trade accuracy for robustness.

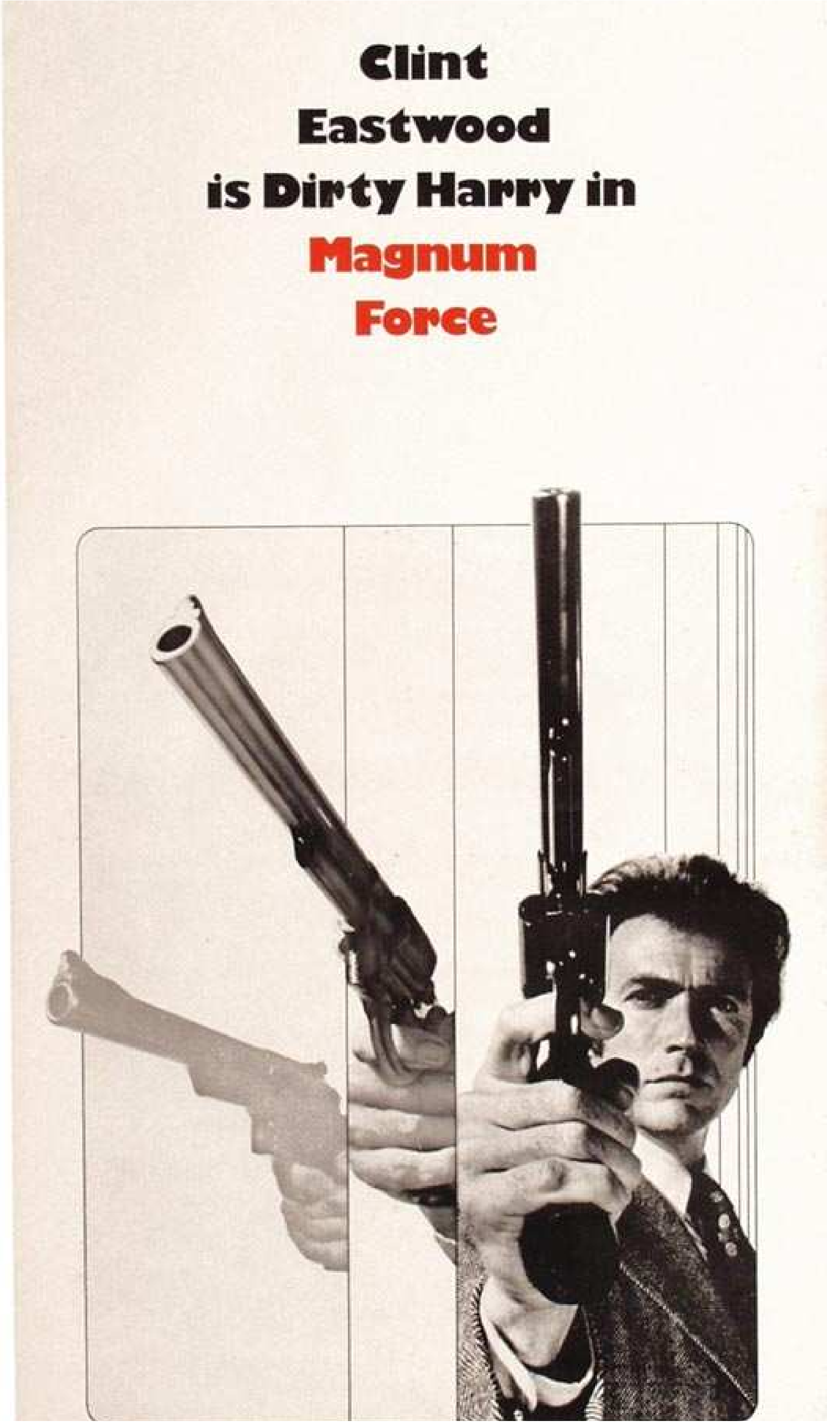

"A good man
knows his limitations"

— Harry Callahan