



Università di Roma **La Sapienza**
Dipartimento di Informatica e Sistemistica

An accurate closed-form estimate of ICP's covariance

Andrea Censi


- Based on the analysis of the error function.
- Advantages over previous approaches:
 - It does not assume independent point correspondences.
 - Measurements can be correlated.
 - Both **accurate** and **fast** (closed form).

The vanilla ICP algorithm

- Input:
 - a reference surface S_{ref} (created from the first scan z_1)
 - a second sensor scan z_2
 - a starting guess x_0
- Repeat until convergence:
 1. compute a set of correspondences
 2. define an error function $J(z_1, z_2, x)$
 3. adjust roto-translation x to minimize J
- Can be used for scan matching and localization.
- Many flavours available...

The vanilla ICP algorithm

- Reference surface is created with a polyline.
- Three points involved for each correspondence.



Why do I get a wrong solution?

- There is no *right* or *wrong* in statistics.
- Sometimes there just is **not enough information** (under-constrained situations, such as a corridor).

under-constrained situations can be detected using Fisher's matrix

“On achievable accuracy for range-finder localization”

today in the ‘miscellaneous’ *session FrC9 at 14:45*

- Sometimes ICP goes crazy due to a **bad initial guess**.
this is hard to model; we assume that it converges to the right basin.
- And then, there is regular **sensor noise**.
- Why is it hard to estimate the ICP covariance?
 - The correspondence/minimize/repeat loop is hard to analyze.

– The map, created from the first scan, is noisy.


Related work - the naive method

- Associate a covariance matrix to each sensor point.
- Assume each correspondence is an independent observation.

$$\text{cov}(\hat{\boldsymbol{x}})^{-1} = \sum_k (P_k)^{-1}$$

- Limitations:
 - In practice, very optimistic.
 - Correspondences are not independent.
 - Environment structure is important

Pfister et al. (2002); Montesano et al. (2005)



Related work - The “brute force method”

- Monte Carlo approximation to the real covariance. [Bengtsson \(2006\)](#)
- Algorithm:
 1. Approximate a map \tilde{S}_{ref} using the first scan.
 2. Repeat multiple times (> 50):
 - (a) Choose a random displacement \boldsymbol{x}_k .
 - (b) Simulate a sensor scan from \tilde{S}_{ref} .
 - (c) Run ICP and compute the error $\hat{\boldsymbol{x}} - \boldsymbol{x}_k$.
 3. Compute the covariance of the errors.
- Limitations:
 - Computationally expensive.
 - One must simulate using an **imperfect** map.

Related work - the “Hessian” method

- If the problem was linear, and the error function was quadratic:

$$\mathbf{z} = M\mathbf{x} + \sigma^2\epsilon \quad \Rightarrow \quad \text{cov}(\hat{\mathbf{x}}) = \sigma^2 \left(\frac{\partial^2 J}{\partial \mathbf{x}^2} \right)^{-1}$$

- Idea: pretend the problem is linear:

Bengtsson (2006)

$$\text{cov}(\hat{\mathbf{x}}) = \underbrace{2 \frac{J(\mathbf{z}, \hat{\mathbf{x}})}{K - 3}}_{\text{approximation to } \sigma^2} \left(\frac{\partial^2 J}{\partial \mathbf{x}^2} \right)^{-1}$$

- Get a robust Hessian by sampling. Biber and Strasser (2003), etc.
- Limitations:
 - In practice, sometimes very pessimistic.
 - Not sound: the Hessian is just part of the solution.

The covariance of a minimization algorithm

$$\hat{\mathbf{x}} = \hat{\mathbf{x}}(\mathbf{z}) = \arg \min J(\mathbf{z}, \mathbf{x})$$

- Because $\hat{\mathbf{x}}$ is a stationary point of the gradient: $\nabla J(\mathbf{z}, \hat{\mathbf{x}}) = 0$, the implicit function theorem provides the $\mathbf{z} \rightarrow \hat{\mathbf{x}}$ Jacobian.

$$\underbrace{\text{cov}(\hat{\mathbf{x}})}_{\substack{\text{solution} \\ \text{covariance}}} = \underbrace{\left(\frac{\partial^2 J}{\partial \mathbf{x}^2} \right)^{-1} \frac{\partial^2 J}{\partial \mathbf{x} \partial \mathbf{z}}}_{\frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{z}} \text{ Jacobian}} \underbrace{\text{cov}(\mathbf{z})}_{\substack{\text{input} \\ \text{covariance}}} \frac{\partial^2 J}{\partial \mathbf{x} \partial \mathbf{z}}^T \left(\frac{\partial^2 J}{\partial \mathbf{x}^2} \right)^{-1}$$

- All is evaluated at the minimum $\hat{\mathbf{x}}$.
- Contains the Hessian **and** the mixed derivative $\partial^2 J / \partial \mathbf{x} \partial \mathbf{z}$: how the shape changes with respect to the measurements.
- Reduces to a familiar formula if J is quadratic (try it).

Application to ICP

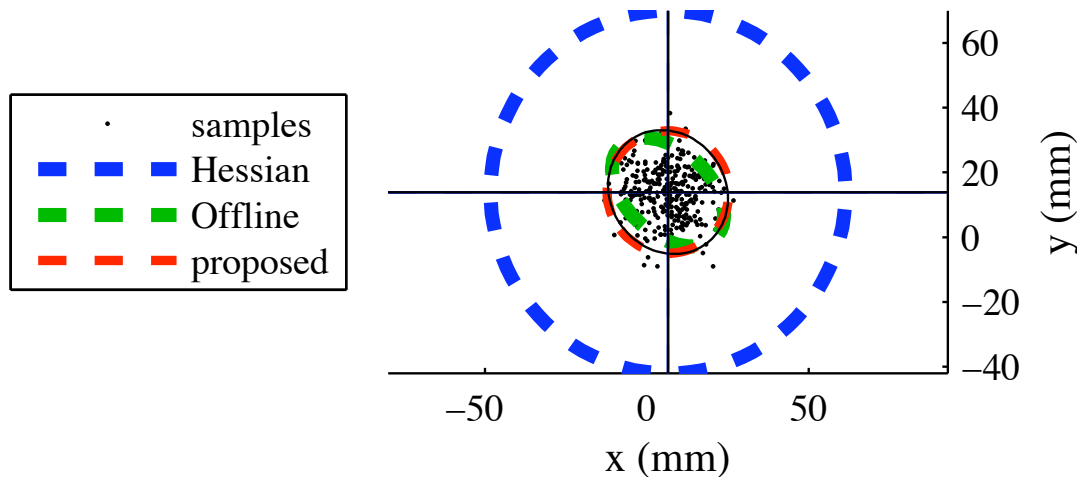
- How to implement this:
 1. run ICP to get \hat{x}
 2. evaluate the derivatives at \hat{x}
 - closed form is possible (lengthy but simple formulas)
- Note that:
 - If the same measurements contributes to two different correspondences, that is taken into account in $\frac{\partial^2 J}{\partial x \partial z}$.

$$\text{Error function} = \sum_{\text{correspondences}} \text{term involving 3 measurements}$$

- The measurement matrix $\text{cov}(z)$ can be a full matrix.

Experiments - scan matching

- the Hessian method is very pessimist.
- the proposed method is very accurate
 - better than the Offline method (!)



Scan matching errors ($mm, mm, ^\circ$)			
	$\sigma(x)$	$\sigma(y)$	$\sigma(\theta)$
true	7.6	7.8	0.058
Hessian	20.0	20.3	0.171
Offline	7.0	6.8	0.086
proposed	7.7	7.7	0.060


- An observability analysis is needed in under-constrained situations; proposed \simeq Hessian, slightly optimistic.

Experiments - localization

- The method can also be used for localization
 - the map is assumed to be perfect.


$$\text{cov}(\mathbf{z}) = \begin{pmatrix} \text{cov}(\mathbf{z}_1) & 0 \\ 0 & \text{cov}(\mathbf{z}_2) \end{pmatrix} \quad \text{becomes} \quad \begin{pmatrix} 0 & 0 \\ 0 & \text{cov}(\mathbf{z}_2) \end{pmatrix}$$

- ... one has the same results as the Cramér–Rao bound.



Correlation among successive poses

- In scan matching, each sensor scan is used twice.




Hence, the estimated displacements x_k are not independent.

- If you just “sum” covariances, you would be pessimist, as scan matching errors tend to cancel out.
- Problem solved by [Mourikis and Roumeliotis \(2006\)](#): you just need the Jacobians $\frac{\partial \mathbf{x}_k}{\partial \mathbf{z}_k}$ and $\frac{\partial \mathbf{x}_k}{\partial \mathbf{z}_{k+1}}$ which we computed:

$$\frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{z}} = \begin{bmatrix} \frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{z}_1} & \frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{z}_2} \end{bmatrix} = \left(\frac{\partial^2 J}{\partial \mathbf{x}^2} \right)^{-1} \frac{\partial^2 J}{\partial \mathbf{x} \partial \mathbf{z}}$$

Example: 1-dimensional scan matching

- In 1-dimensional scan matching, errors tend to cancel out.



$$\begin{aligned}\text{final pose} &= \mathbf{x}_1 \oplus \mathbf{x}_2 \oplus \cdots \oplus \mathbf{x}_n && \text{(sum of deltas)} \\ &= (z_1 - z_0) + (z_2 - z_1) + \cdots + (z_n - z_{n-1}) = z_n - z_1\end{aligned}$$

- The real final covariance is very small:

$$\text{var}(\text{final pose}) = 2 \text{var}(z) \ll 2n \text{var}(z)$$

- The more correlated the measurements are, the more you are being pessimist if you ignore the correlation.

Conclusions

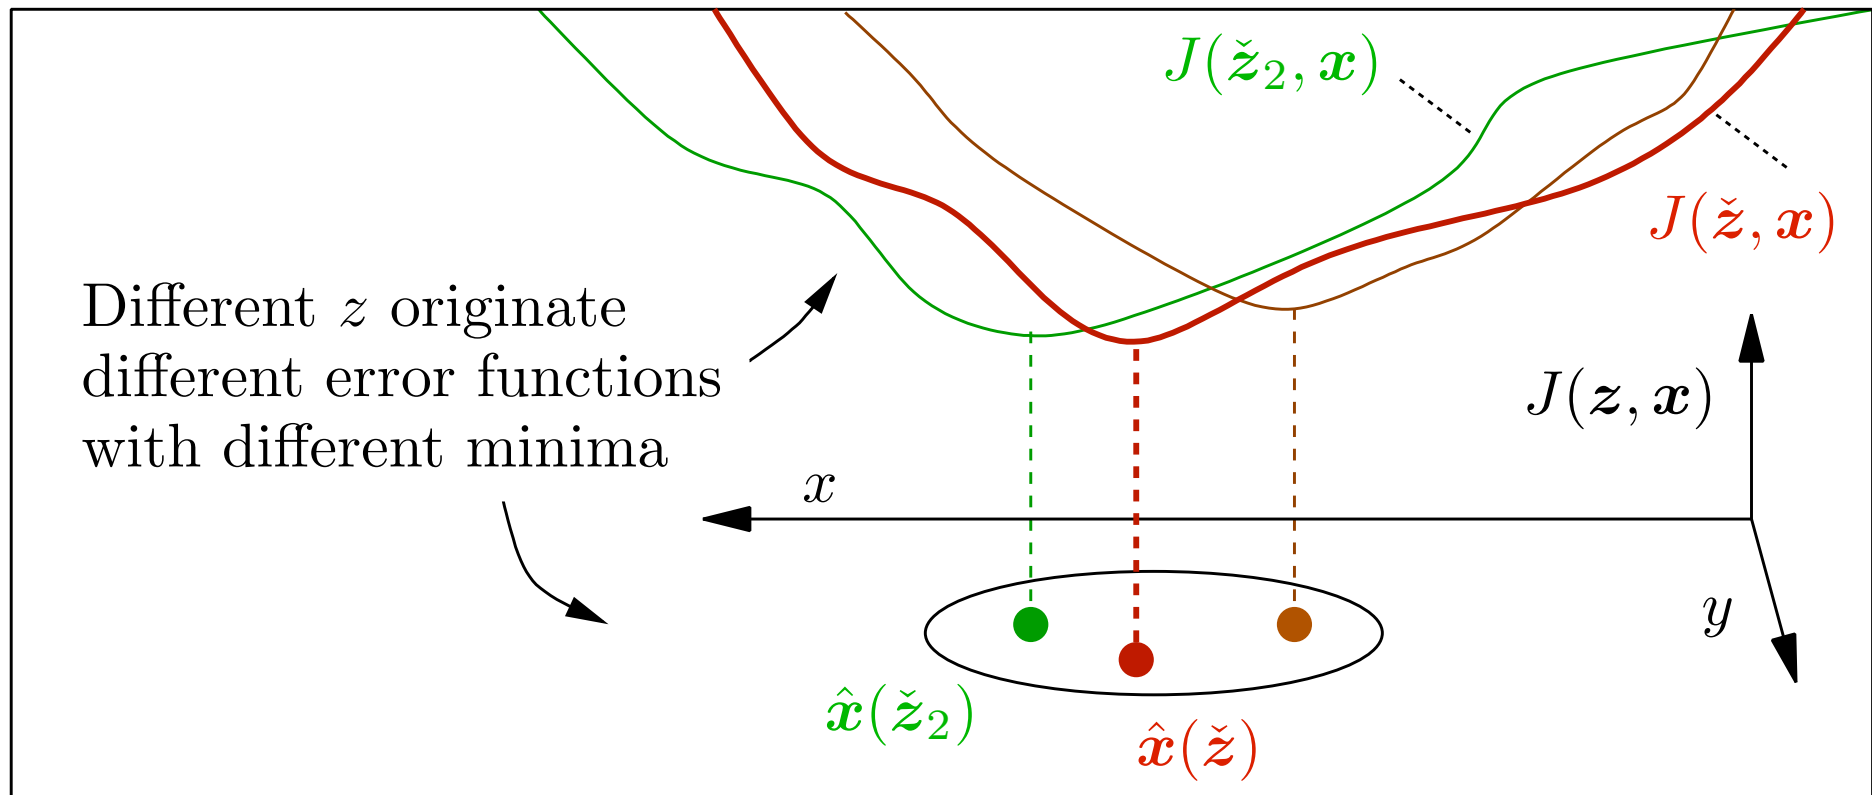
- **A good trick to know:** The covariance of an minimization algorithm only depends on the error function.
- Advantages over previous methods for estimating ICP's covariance:
 - mathematically sound
 - accurate (also more than simulations-based methods)
 - fast: closed form
 - also solves the problem of correlated estimates
- For more on the observability analysis, Fisher's information matrix, Cramér–Rao bound, please see

“On achievable accuracy for range-finder localization”

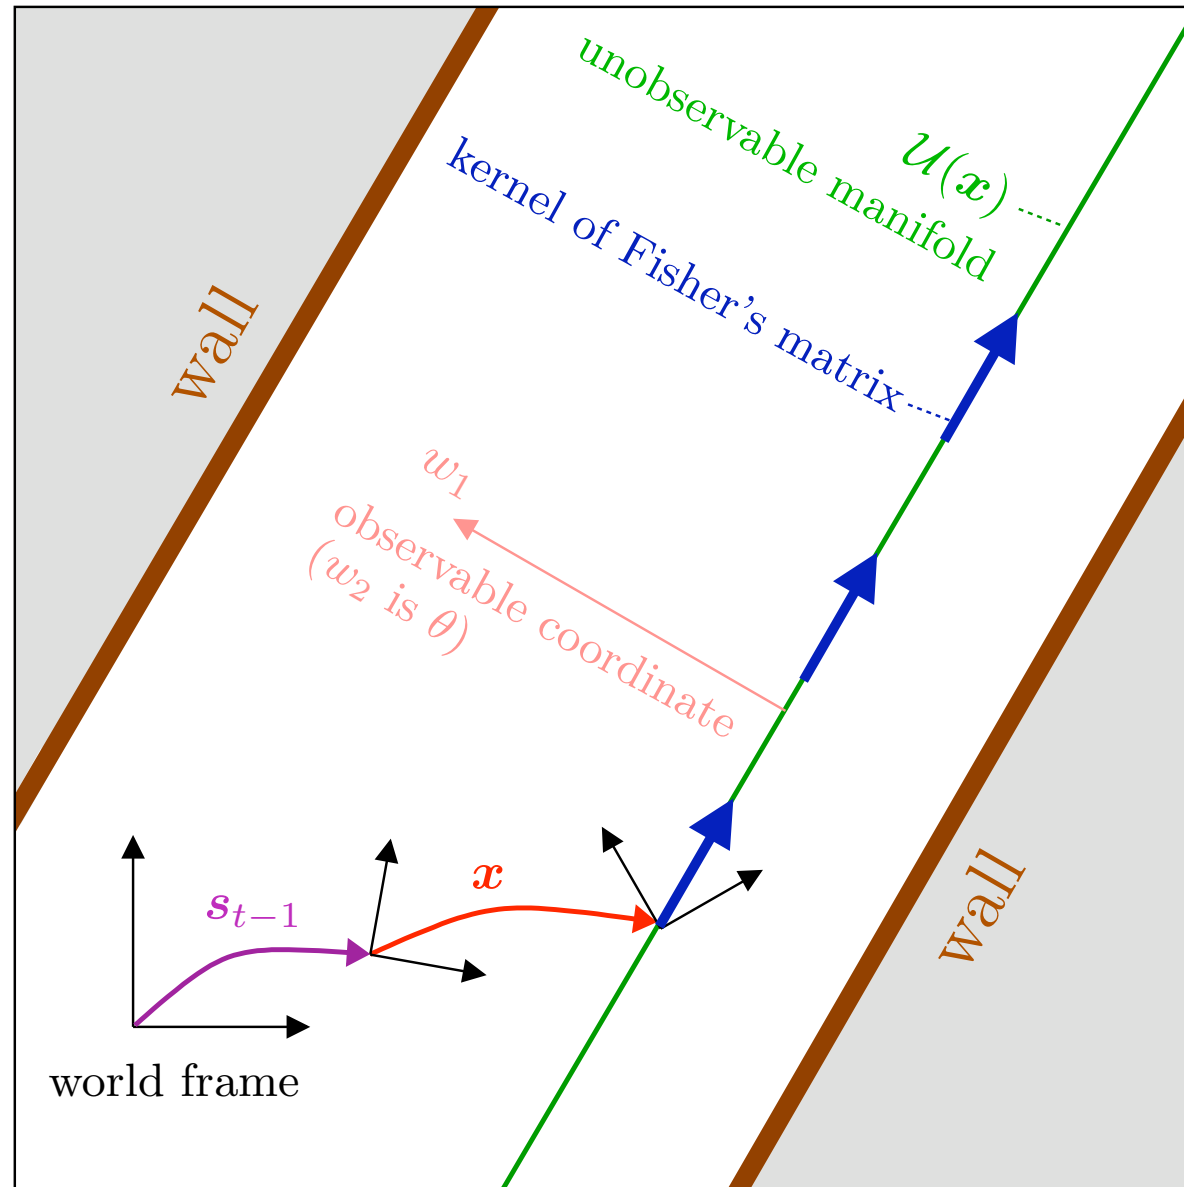
today in the ‘miscellaneous’ *session FrC9 at 14:45*

Backup / Change in the error function

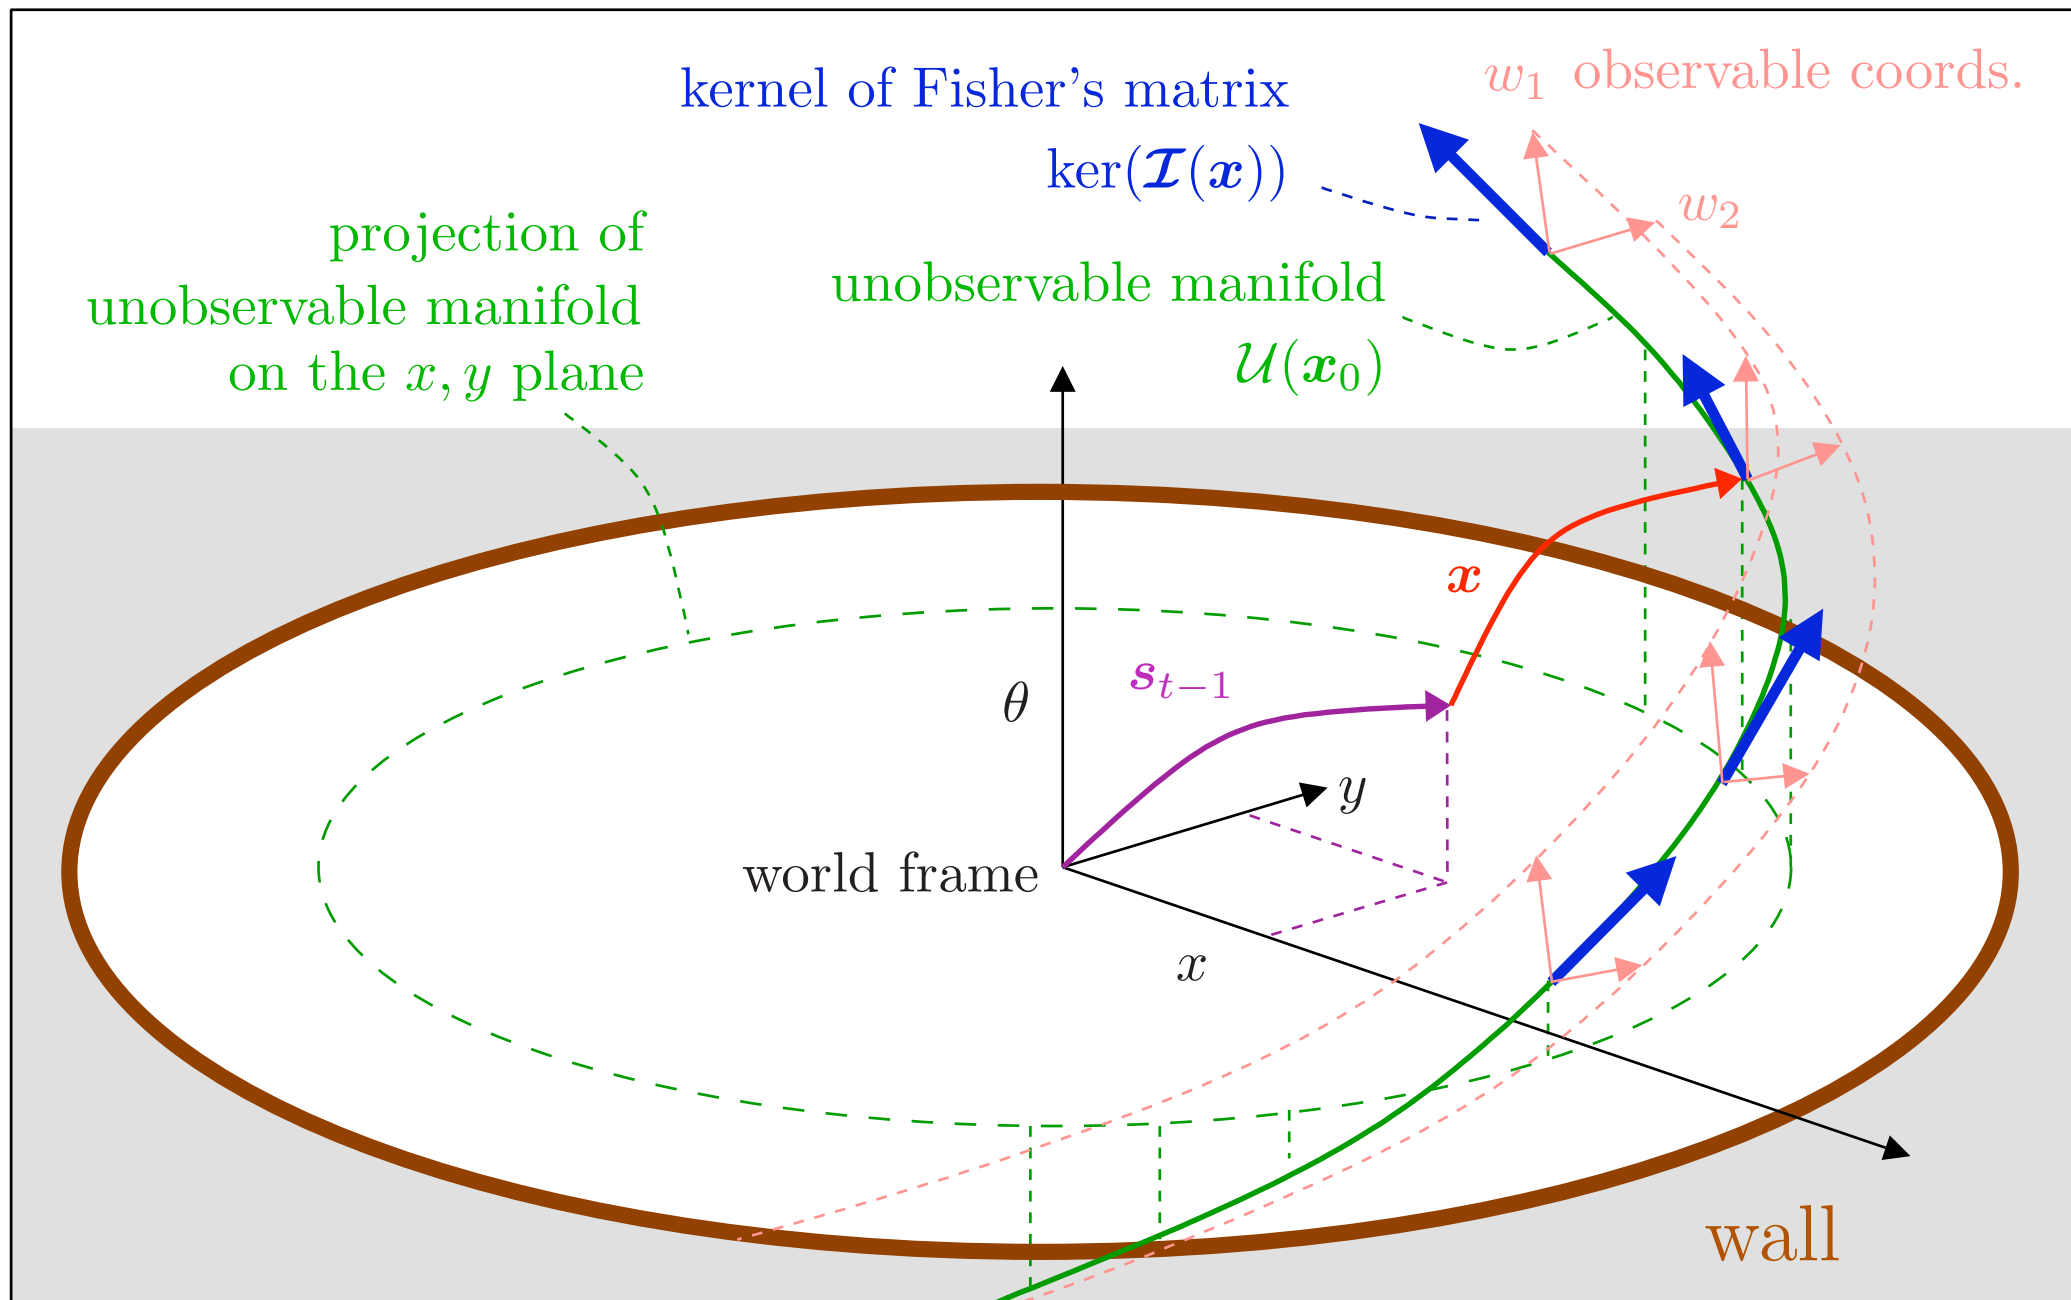
- The term $\frac{\partial^2 J}{\partial \mathbf{x} \partial \mathbf{z}} = \frac{\partial}{\partial \mathbf{z}} \nabla J$ accounts for the change in the shape of the error function.



Backup / Observable manifold – corridor



Backup / Observable manifold – circle



References

- Ola Bengtsson. *Robust Self-Localization of Mobile Robots in Dynamic Environments Using Scan Matching Algorithms*. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology, Göteborg, Sweden, 2006. ISBN 91-7291-744-X.
- Peter Biber and Wolfgang Strasser. The normal distributions transform: A new approach to laser scan matching. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- L. Montesano, J. Minguez, and L. Montano. Probabilistic scan matching for motion estimation in unstructured environments. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Canada, 2005.

Anastasios Mourikis and Stergios Roumeliotis. On the treatment of relative-pose measurements for mobile robot localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, FL, 2006.

S.T. Pfister, K.L. Kriechbaum, S.I. Roumeliotis, and J.W. Burdick. Weighted range sensor matching algorithms for mobile robot displacement estimation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2002.