Landmark Localization using Gaussian Flow

Jay Paek

July 24, 2025

Table of Contents

Introduction: The Localization Problem

The primary goal of this project is to accurately determine the position of a static landmark.

We aim to solve two main challenges:

1. **Estimate the landmark's true position**, denoted by y^*.
2. **Minimize the uncertainty** of our estimate.

But there are two issues.

1. MLE can be intractable in high dimensions.
2. Quantifying uncertainty is hard between two drastically different distributions.

Problem Formulation

System Setup

- **Landmark Ground Truth**: $y^* = \begin{bmatrix} 4.7 & -3.1 \end{bmatrix}^\top$
- **Agent Initial Position**: $x_0 = \begin{bmatrix} 0 & 0 \end{bmatrix}^\top$
- **Prior Estimate of Landmark**: $\mathcal{N}(\mu_0, \Sigma_0)$
  - $\mu_0 = \begin{bmatrix} 1 & 1 \end{bmatrix}^\top$
  - $\Sigma_0 = \begin{bmatrix} 5.5 & -1.5 \\ -1.5 & 5.5 \end{bmatrix}$

Observation Model

The range measurement at step k is given by:

$$z_k = h(x_k) = \|x_k - y^*\| + w_k$$

where the noise w_k is i.i.d. and follows a normal distribution, $w_k \sim \mathcal{N}(0, \sigma^2 = 2)$.

Problem Formulation (cont.)

The simulate the following procedure.

Simulation Steps

1. Draw sample z_k.
2. Update $(\mu_k, \Sigma_k) \to (\mu_{k+1}, \Sigma_{k+1})$
3. Move to x_{k+1}

We are concerned about the optimal first step.

Motivations and Applications

Real World Applications:

- Robotics
- Navigation
- Target Tracking
- Scientific Exploration

There exists other methods:

- Nonlinear Kalman Filter
- Particle Flow Filters

Gaussian Flow via KL-Divergence

We approximate the posterior distribution $p(\cdot|z)$ with a Gaussian $q(\cdot;\mu,\Sigma)$. We then minimize the KL-Divergence $D_{\mathrm{KL}}(q \parallel p)$ using gradient flow.

Gradient Flow Dynamics

The updates for the mean μ and covariance Σ are:

$$\dot{\mu} = -\nabla_\mu D_{\mathrm{KL}} = -\Sigma^{-1}\mathbb{E}_{s\sim q}\left[(s-\mu)\ln p(s,z)\right]$$
$$\dot{\Sigma} = -\nabla_\Sigma D_{\mathrm{KL}} = -\frac{1}{2}\left[\Sigma^{-1} + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(s,z)]\right]$$

Hessian of Log-Likelihood

$$\nabla_s^2 \ln p(z \mid s) = \frac{z - \|x_k - s\|}{\sigma^2} H_h(s) - \frac{1}{\sigma^2} g_h(s) g_h(s)^\top$$

where $g_h(s)$ and $H_h(s)$ are the gradient-transpose and Hessian of the measurement function $h(s) = \|x_k - s\|$.

Key Mathematical Tools

$$D_{\mathrm{KL}} = \mathbb{E}_{s\sim q}\left[\ln q(s) - \ln p(s,z)\right]$$

Theorem (Bonnet's Theorem)

Let $h(s) : \mathbb{R}^d \to \mathbb{R}$ be a locally ACL and continuous function. The following first-order identity holds:

$$\nabla_\mu \mathbb{E}_{s\sim q}[h(s)] = \mathbb{E}_{s\sim q}[\nabla_s h(s)] = \mathbb{E}_{s\sim q}[\Sigma^{-1}(s-\mu)h(s)]$$

Theorem (Price's Theorem)

Let $h(s) : \mathbb{R}^d \to \mathbb{R}$ be continuously differentiable with its derivative $\nabla h(s)$ being locally ACL. The following second-order identity holds:

$$\nabla_\Sigma \mathbb{E}_{s\sim q}[h(s)] = \frac{1}{2}\mathbb{E}_{s\sim q}[\nabla_s^2 h(s)]$$

Deriving $\nabla_\mu D_{\mathrm{KL}}$

Objective

We want to find the gradient of $D_{\mathrm{KL}}\big(q(.) \parallel p(. \mid z)\big)$ with respect to μ. We start with the objective function, ignoring constants:

$$D_{\mathrm{KL}} = \mathbb{E}_{s\sim q}\big[\ln q(s) - \ln p(s,z)\big]$$

Let $h(s) = \ln q(s) - \ln p(s,z)$. Applying Bonnet's Theorem:

$$\begin{aligned}
\nabla_\mu D_{\mathrm{KL}} &= \nabla_\mu \mathbb{E}_{s\sim q}[h(s)] \\
&= \mathbb{E}_{s\sim q}[\Sigma^{-1}(s-\mu)h(s)] \\
&= \Sigma^{-1}\mathbb{E}_{s\sim q}\big[(s-\mu)(\ln q(s) - \ln p(s,z))\big] \\
&= \Sigma^{-1}\left(\mathbb{E}_{s\sim q}[(s-\mu)\ln q(s)] - \mathbb{E}_{s\sim q}[(s-\mu)\ln p(s,z)]\right)
\end{aligned}$$

Final Result for $\nabla_\mu D_{\mathrm{KL}}$

From the previous slide:

$$\nabla_\mu D_{\mathrm{KL}} = \Sigma^{-1} \left(\mathbb{E}_{s\sim q}[(s-\mu)\ln q(s)] - \mathbb{E}_{s\sim q}[(s-\mu)\ln p(s,z)] \right)$$

Simplification

We use the fact that for a Gaussian $q \sim \mathcal{N}(\mu, \Sigma)$:

- The term $\ln q(s)$ is a quadratic function of $(s-\mu)$.
- The odd moments of a centered Gaussian are zero.

This implies that the expectation $\mathbb{E}_{s\sim q}[(s-\mu)\ln q(s)]$ is zero.

Final Gradient Expression for μ

$$\nabla_\mu D_{\mathrm{KL}} = -\Sigma^{-1}\mathbb{E}_{s\sim q}\left[(s-\mu)\ln p(s,z)\right]$$

Deriving $\nabla_\Sigma D_{\mathrm{KL}}$

Objective

Now we find the gradient of $D_{\mathrm{KL}} = \mathbb{E}_{s\sim q}\left[\ln q(s) - \ln p(s,z)\right]$ with respect to Σ.

Using the linearity of the gradient operator and applying Price's Theorem with $h(s) = \ln q(s) - \ln p(s,z)$:

$$\begin{aligned}\nabla_\Sigma D_{\mathrm{KL}} &= \nabla_\Sigma \mathbb{E}_{s\sim q}[\ln q(s)] - \nabla_\Sigma \mathbb{E}_{s\sim q}[\ln p(s,z)] \\ &= \frac{1}{2}\mathbb{E}_{s\sim q}[\nabla_s^2 \ln q(s)] - \frac{1}{2}\mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(s,z)] \quad \text{(Price's Thm.)}\end{aligned}$$

Final Gradient Expression for Σ

Substituting this identity gives:

$$\nabla_\Sigma D_{\mathrm{KL}} = -\frac{1}{2}\left(\Sigma^{-1} + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(s,z)]\right)$$

Decomposition of the Hessian Term

The final step is to expand the term $\mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(s,z)]$.
The joint log-probability is $\ln p(s,z) = \ln p(s) + \ln p(z|s)$.

$$\nabla_\Sigma D_{\mathrm{KL}} = -\frac{1}{2}\left(\Sigma^{-1} + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(s)] + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(z|s)]\right)$$

Simplifying with the Prior

The prior $p(s)$ is a Gaussian $\mathcal{N}(\mu_0, \Sigma_0)$. Its log-pdf is quadratic in s, so its Hessian is constant:

$$\nabla_s^2 \ln p(s) = \nabla_s^2 \left(C - \frac{1}{2}(s-\mu_0)^\top \Sigma_0^{-1}(s-\mu_0) \right) = -\Sigma_0^{-1}$$

Final Form for Covariance Gradient Flow

$$\dot{\Sigma} = -\nabla_\Sigma D_{\mathrm{KL}} = \frac{1}{2}\left((\Sigma^{-1} - \Sigma_0^{-1}) + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(z|s)]\right)$$

Monte-Carlo Gaussian Flow Algorithm

We approximate the expectations in the gradient flow dynamics using Monte Carlo sampling.

Algorithm Monte-Carlo-based Gaussian Flow

```
Input: Prior N(μ_0, Σ_0), observation z, number of steps N, number of samples n, step size η.
for k = 0, ..., N − 1 do
    Draw samples s_1, ..., s_n ~ N(μ_k, Σ_k)
    ∇_μ D_KL ≈ −Σ_k^{-1} (1/n) Σ_{i=1}^n [(s_i − μ_k) ln p(s_i, z)]
    ∇_Σ D_KL ≈ −(1/2) [(Σ_k^{-1} − Σ_0^{-1}) + (1/n) Σ_{i=1}^n ∇_s^2 ln p(z | s_i)]
    μ_{k+1} ← μ_k − η ∇_μ D_KL
    Σ_{k+1} ← Σ_k − η ∇_Σ D_KL
end for
Return: μ_N, Σ_N
```

Where:

- **Input:** Prior $\mathcal{N}(\mu_0, \Sigma_0)$, observation z, number of steps N, number of samples n, step size η.
- **for** $k = 0, \ldots, N-1$ **do**
  - Draw samples $s_1, \ldots, s_n \sim \mathcal{N}(\mu_k, \Sigma_k)$
  - $\nabla_\mu D_{\mathrm{KL}} \approx -\Sigma_k^{-1} \frac{1}{n} \sum_{i=1}^{n} \left[(s_i - \mu_k) \ln p(s_i, z) \right]$
  - $\nabla_\Sigma D_{\mathrm{KL}} \approx -\frac{1}{2} \left[(\Sigma_k^{-1} - \Sigma_0^{-1}) + \frac{1}{n} \sum_{i=1}^{n} \nabla_s^2 \ln p(z \mid s_i) \right]$
  - $\mu_{k+1} \leftarrow \mu_k - \eta \nabla_\mu D_{\mathrm{KL}}$
  - $\Sigma_{k+1} \leftarrow \Sigma_k - \eta \nabla_\Sigma D_{\mathrm{KL}}$
- **end for**
- **Return:** μ_N, Σ_N

Result: Gaussian Flow Estimation

The Gaussian flow update successfully fuses the prior information with the new measurement to produce a more accurate posterior estimate.

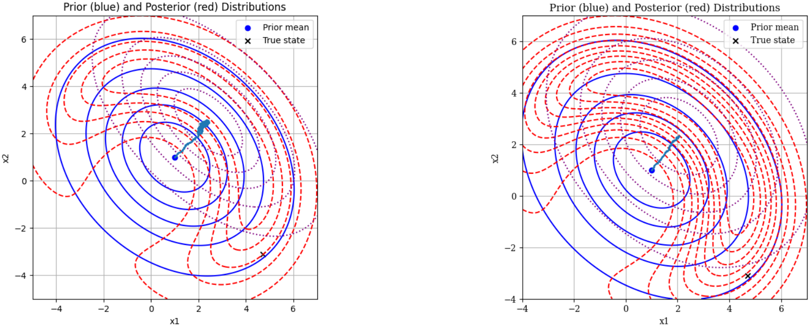

Figure: Two trajectories where Comparison of the prior (blue), the MLE estimate (red), and the Gaussian flow posterior estimate (purple). Plot on the right has increased noise.

Result: The Challenge of Ambiguity

The Gaussian flow can converge to incorrect estimates, highlighting the challenge of local minima in non-linear problems.

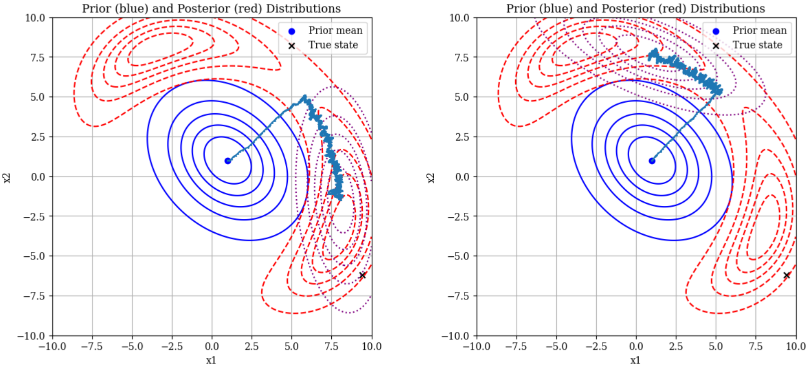

Figure: Two trajectories where the ground truth landmark is at $2y^*$. The flow converges to two different locations, demonstrating ambiguity.

The Optimal Control Question

Question: Using our current model, is there a movement x_1, I can make AFTER my first flow update $(\mu_0, \Sigma_0) \rightarrow (\mu_1, \Sigma_1)$, that is better than other movements (on average).

Our Strategy: Choose the next location x_1 that maximizes the initial expected reduction in uncertainty $\|\Sigma(T) - \Sigma(0)\|_2$, to quantify "reduction in uncertainty"

But $\|\Sigma(T) - \Sigma(0)\|_2$ is an an easily accessible value. We perform numerical experiments to find a trend for $\min_{x \in \mathcal{X}} \mathbb{E}_{z_1}[\|\dot{\Sigma}(0; x_1)\|_2]$

Result: Chasing minimum $\|\dot{\Sigma}\|$

We fix $\mu_0 = \begin{bmatrix} 1 & 1 \end{bmatrix}^\top, \Sigma_0 = 0.1I$. The following are normalized heatmaps of $\mathbb{E}_{z_1}[\|\dot{\Sigma}(0; x_1)\|_2]$ for positions in $x_1 \in [-2, 2] \times [-2, 2]$.

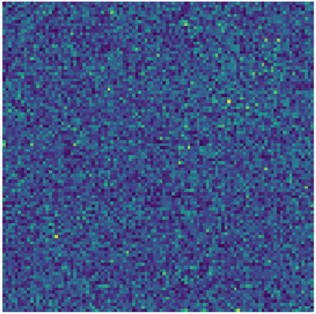

Figure: $y = \mu_0$

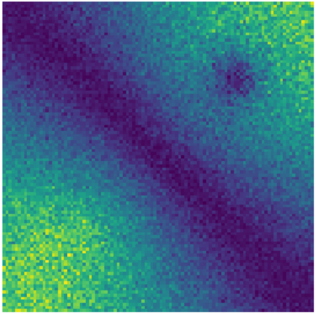

Figure: $y = -\mu_0$

Result: Chasing minimum $\|\dot{\Sigma}\|$

We fix $\mu_0 = \begin{bmatrix} 1 & 1 \end{bmatrix}^\top$, $\Sigma_0 = 0.1I$. We are evaluating for positions in $[-2, 2] \times [-2, 2]$.

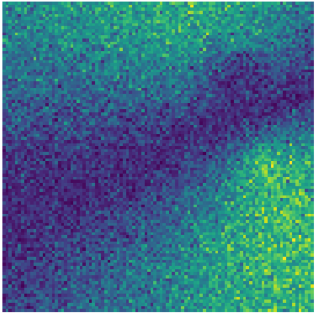

Figure: $y = e_1$

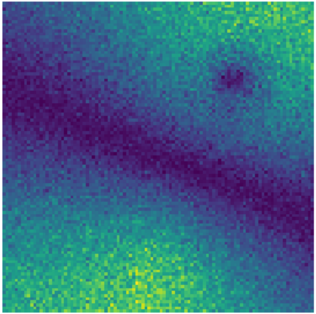

Figure: $y = -e_2$

Step 1: Initial Rate of Change

The Gaussian flow dynamics for the covariance matrix Σ are given by:

$$\dot{\Sigma}(z_1) = -\nabla_\Sigma D_{\mathrm{KL}} = \frac{1}{2}\left((\Sigma^{-1} - \Sigma_0^{-1}) + \mathbb{E}_{s\sim q}[\nabla_s^2 \ln p(z_1|s)]\right)$$

At time $t = 0$

We are interested in the *initial* change. At this point, our variational distribution q is exactly the prior distribution $q_0 = \mathcal{N}(\mu_0, \Sigma_0)$.

Substituting $\Sigma = \Sigma_0$ into the equation, the first term vanishes:

$$\dot{\Sigma}(z_1)|_{t=0} = \frac{1}{2}\left((\Sigma_0^{-1} - \Sigma_0^{-1}) + \mathbb{E}_{s\sim q_0}[\nabla_s^2 \ln p(z_1|s)]\right)$$

This simplifies to:

$$\dot{\Sigma}(z_1)|_{t=0} = \frac{1}{2}\mathbb{E}_{s\sim q_0}[\nabla_s^2 \ln p(z_1|s)]$$

Step 1 (cont.): Expectation over Measurements

The Hessian of the log-likelihood $\ln p(z_1|s) = C - \frac{(z_1 - \|x_1 - s\|)^2}{2\sigma^2}$ is:

$$\nabla_s^2 \ln p(z_1|s) = \frac{1}{\sigma^2}\left[\frac{z_1 - \|x_1 - s\|}{\|x_1 - s\|} I - \frac{z_1}{\|x_1 - s\|^3}(s - x_1)(s - x_1)^\top\right]$$

Marginalizing out the measurement z_1

We take the expectation with respect to the measurement model $z_1 \sim \mathcal{N}(\|x_1 - y\|, \sigma^2)$, for which $\mathbb{E}[z_1] = \|x_1 - y\|$. Let's define:

- $d_y = \|x_1 - y\|_2$ (distance from agent to true landmark)
- $d(s) = \|x_1 - s\|_2$ (distance from agent to a hypothesised landmark s)

The expected initial rate of change is then:

$$\mathbb{E}_{z_1}[\dot{\Sigma}] = \frac{1}{2\sigma^2}\mathbb{E}_{s\sim q_0}\left[\frac{d_y - d(s)}{d(s)} I - \frac{d_y}{d(s)^3}(s - x_1)(s - x_1)^\top\right]$$

Step 2: Approximating the Expectation

The expression for $\mathbb{E}_{z_1}[\dot{\Sigma}]$ is still complex due to the expectation over $s \sim q_0$. We can simplify this with a key assumption.

Assumption 1

The prior distribution is concentrated far away from the agent's position x_1.

Justification: If the prior belief μ_0 is far from x_1, then for most samples s from q_0, the vector $s - x_1$ points in roughly the same direction as $\mu_0 - x_1$. In particular

$$\mathbb{E}_{s \sim q_0}[d(s)] \approx d(\mu_0)$$

Step 2 (cont.): The Approximated Gradient

Applying the approximation $\mathbb{E}_{s\sim q_0}[f(s)] \approx f(\mu_0)$ to our expression for $\mathbb{E}_{z_1}[\dot{\Sigma}]$:

$$\begin{aligned}\mathbb{E}_{z_1}[\dot{\Sigma}] &\approx \frac{1}{2\sigma^2}\left[\frac{d_y - d(\mu_0)}{d(\mu_0)} I - \frac{d_y}{d(\mu_0)^3}(\mu_0 - x_1)(\mu_0 - x_1)^\top\right] \\ &= \frac{1}{2\sigma^2}\left[\left(\frac{d_y}{d(\mu_0)} - 1\right) I - \frac{d_y}{d(\mu_0)} \frac{(\mu_0 - x_1)(\mu_0 - x_1)^\top}{\|\mu_0 - x_1\|^2}\right] \\ &= \frac{1}{2\sigma^2}\left[\frac{d_y}{d(\mu_0)}\left(I - \frac{(\mu_0 - x_1)(\mu_0 - x_1)^\top}{d(\mu_0)^2}\right) - I\right]\end{aligned}$$

where $d(\mu_0) = \|x_1 - \mu_0\|_2$.

Finding the Optimal Condition

Let's analyze the approximated rate of change:

$$\mathbb{E}_{z_1}[\dot{\Sigma}] \approx C \left[\frac{\|x_1 - y\|_2}{\|x_1 - \mu_0\|_2} P^{\perp}_{\mu_0} - I \right]$$

where $P^{\perp}_{\mu_0} = I - \frac{(\mu_0 - x_1)(\mu_0 - x_1)^{\top}}{\|x_1 - \mu_0\|_2^2}$ is a projection matrix onto the space perpendicular to the direction $(\mu_0 - x_1)$.

Assumption 2

To avoid trivial solutions (e.g., moving directly to μ_0), we assume that the true landmark y and the prior mean μ_0 are both sufficiently far from the agent.

The Optimal Condition for Agent Movement

Because we are working in $\mathbb{R}^2$, algebraically computing the eigenvalues of $\|\mathbb{E}_{z_1}[\dot{\Sigma}(z_1)]\|$ is possible.

$$\|\mathbb{E}_{z_1}[\dot{\Sigma}(z_1)]\| = \max\left(1, \left|\frac{d_y}{d(\mu)} - 1\right|\right)$$

So we have approximate optimality if

$$\left|\frac{\|x - y\|}{\|x - \mu_0\|} - 1\right| < 1$$

In particular

$$\frac{\|x - y\|}{\|x - \mu_0\|} = 0 \implies \|x_1 - y\|_2 = \|x_1 - \mu_0\|_2$$

Future Steps

- More insight in $\mathbb{R}^d$.
- Extend to different kinds of observation models or a general model.
- Deduce a optimal control algorithm for more time stamps without oracle assumption (using the previous result).