

CSC 144V: Autonomous Vehicles

Lecture 9.2

Modern SLAM for AD Exercises Solution

Department of Computer Science,
Hofstra University

Exercise: Reprojection error and scale

Pinhole camera: $f = 800$ px, principal point $(640, 360)$. A map point is at $X = (2.0, -0.5, 10.0)$ m in the camera frame. $u = f \cdot X/Z + c_x$, $v = f \cdot Y/Z + c_y$.

- (a) Projected pixel? (b) The feature is observed at $(803, 318)$: reprojection residual and squared error.
- (c) Scale the scene and the camera position by 2: where does the point project? (d) A second camera is 0.5 m to the right: pixel u_2 , disparity, and depth $Z = f \cdot b/d$.
- (e) The disparity is mismeasured as 41 px: depth and relative error?

Exercise Solution: Reprojection error and scale

- (a) $u = 800 \cdot 2 / 10 + 640 = 800$, $v = 800 \cdot (-0.5) / 10 + 360 = 320$.
- (b) $e = (803 - 800, 318 - 320) = (3, -2) \text{ px}$, $\|e\|^2 = 13$, $\|e\| = 3.6 \text{ px}$. Bundle adjustment minimizes the (robust) sum of such terms.
- (c) $X \rightarrow (4, -1, 20)$: $u = 800 \cdot 4 / 20 + 640 = 800$, $v = 320$: identical. A single camera cannot tell the scale (**scale ambiguity**).
- (d) X in camera 2 = $(1.5, -0.5, 10)$: $u_2 = 800 \cdot 1.5 / 10 + 640 = 760$, disparity **40 px**, $Z = 800 \cdot 0.5 / 40 = 10 \text{ m}$. The known baseline fixes the scale.
- (e) $Z = 400 / 41 = 9.76 \text{ m}$, error $-0.24 \text{ m} = 2.4 \%$ for 1 px. $\Delta Z \approx Z^2 \Delta d / (f \cdot b)$: depth error grows with the **square** of the distance.

Exercise: ORB matching and RANSAC

ORB descriptors (shown with 8 bits instead of 256): query $q = 10110010$; candidates $c1 = 10100010$, $c2 = 01100110$, $c3 = 00001101$.

- (a) Hamming distances and best match; apply the ratio test (accept if $\text{best}/\text{second} < 0.7$).
- (b) In a repeated texture the best two distances are 3 and 4. Accept?
- (c) RANSAC needs $N = \log(1 - p)/\log(1 - w^s)$ iterations ($p = 0.99$). Fill in the table (w = inlier ratio, s = sample size).

Model	s	w	N
P3P (PnP)	3	0.5	
5-point (E)	5	0.5	
8-point	8	0.5	
5-point (E), bad matches	5	0.3	

Exercise Solution: ORB matching and RANSAC

- (a) $d(q, c1) = 1$, $d(q, c2) = 4$, $d(q, c3) = 7$. Best = c1, ratio $1/4 = 0.25 \rightarrow$ accept.
- (b) Ratio $3/4 = 0.75 \geq 0.7 \rightarrow$ reject: the match is ambiguous (repeated structure gives false matches).
- (c) Fewer points per sample need far fewer iterations (P3P vs. 8-point); a low inlier ratio is expensive. This is why guided search and the ratio test (raising w) matter.

Model	s	w	N
P3P (PnP)	3	0.5	35
5-point (E)	5	0.5	146
8-point	8	0.5	1,177
5-point (E), bad matches	5	0.3	1,893

Exercise: LOAM features

LOAM smoothness: $c = \|\sum_{j \in S} (X_i - X_j)\| / (|S| \cdot \|X_i\|)$. Point $X_i = (10, 0)$ m, $|S| = 4$ neighbors.

Case A (wall): neighbors $(10, -0.4)$, $(10, -0.2)$, $(10, 0.2)$, $(10, 0.4)$. Case B (corner): $(10, -0.4)$, $(10, -0.2)$, $(9.4, 0.2)$, $(9.2, 0.4)$.

(a) Compute c for both cases. (b) With threshold 0.02, which is an edge and which planar? (c) Why divide by $\|X_i\|$?

Case	$\sum (X_i - X_j)$	c	Feature
A (wall)			
B (corner)			

Exercise Solution: LOAM features

- (c) Point spacing and noise grow with range: dividing by $\|X_i\|$ makes c **range-independent**, so the same threshold works near and far.
- Edge points are matched to lines, planar points to planes: together they constrain all 6 DoF, while most of the cloud is discarded for speed.

Case	$\Sigma (X_i - X_j)$	c	Feature
A (wall)	(0, 0)	0	planar
B (corner)	(1.4, 0)	$1.4/(4 \cdot 10) = 0.035$	edge

Exercise: Degeneracy of scan matching

Each matched plane with unit normal $\mathbf{n}$ adds $\mathbf{n} \mathbf{n}^T$ to $\mathbf{H}$ (2-D translation only). Point noise $\sigma = 0.05$ m; σ along an eigen-direction $\approx \sigma/\sqrt{\lambda}$.

A: corridor, 100 points on walls with $\mathbf{n} = (0, 1)$. B: A plus 20 points on an end wall $\mathbf{n} = (1, 0)$. C: A plus 30 points on a 45° wall $\mathbf{n} = (1, 1)/\sqrt{2}$.

(a) Fill in $\mathbf{H}$, its eigenvalues, and the weakest-direction σ . (b) What happens if an attacker removes or replaces the end-wall points of B?

Case	$\mathbf{H}$	Eigenvalues	Weak-direction σ
A			
B			
C			

Exercise Solution: Degeneracy of scan matching

- (b) B collapses to A: $\lambda_{\min} = 0$, the x direction is unconstrained, and the estimate is free to slide along the corridor. The filter must rely on IMU and odometry; a spoofer who controls the few informative points controls the result.
- Monitoring the **eigenvalues of H** (or the match residuals) detects both accidental degeneracy and attacks that shrink the information.

Case	H	Eigenvalues	Weak-direction σ
A	$[[0, 0], [0, 100]]$	0 and 100	∞ (x unobservable)
B	$[[20, 0], [0, 100]]$	20 and 100	$0.05/\sqrt{20} = 0.011 \text{ m}$
C	$[[15, 15], [15, 115]]$	12.8 and 117.2	$0.05/\sqrt{12.8} = 0.014 \text{ m}$

Exercise: Lateral budget and latency

Lane width 3.5 m, car width 1.8 m. Cruise at $v = 25$ m/s (90 km/h).

- (a) One-sided margin and the largest localization σ if 3σ must stay within the margin.
- (b) If $\sigma_{\text{loc}} = 0.1$ m, how much margin is left for control and perception?
- (c) A pose arrives 80 ms late: how far does the car travel? How small must the delay be to stay below 0.1 m? What residual error remains if the pose is extrapolated with IMU and the car accelerates at 2 m/s²?
- (d) 500 m tunnel, no GNSS, uncompensated gyro bias $0.02^\circ/\text{s}$: lateral error $\approx \frac{1}{2} \cdot v \cdot b \cdot t^2$. Compare with the margin; what bias is allowed?

Exercise Solution: Lateral budget and latency

- (a) $(3.5 - 1.8)/2 = 0.85 \text{ m} \rightarrow \sigma \leq 0.85/3 = 0.28 \text{ m}$. (b) $3\sigma = 0.30 \text{ m} \rightarrow 0.55 \text{ m}$ left.
- (c) $25 \cdot 0.08 = 2.0 \text{ m}$. For 0.1 m: $\leq 4 \text{ ms}$. With extrapolation only the unmodeled acceleration matters: $\frac{1}{2} \cdot 2 \cdot 0.08^2 = 6.4 \text{ mm}$. Latency must be **compensated**, not just minimized.
- (d) $t = 500/25 = 20 \text{ s}$; $b = 0.02^\circ/\text{s} = 3.5 \times 10^{-4} \text{ rad/s} \rightarrow \frac{1}{2} \cdot 25 \cdot 3.5 \times 10^{-4} \cdot 400 = 1.75 \text{ m} > 0.85 \text{ m}$. Allowed bias: $0.85 / (\frac{1}{2} \cdot 25 \cdot 400) = 1.7 \times 10^{-4} \text{ rad/s} = 0.0097^\circ/\text{s}$. Tunnels need good IMU calibration and odometry or map constraints.

Exercise: Slow GNSS drift vs. the gate

1-D fusion filter in steady state: Kalman gain $K = 0.2$, innovation variance $S = 1.25 \text{ m}^2$. Gate: $\text{NIS} = y^2/S \leq 3.84$ (95 %, 1 DoF).

- (a) Gate half-width in meters. A sudden 5 m GNSS jump: NIS and decision?
- (b) The spoofer instead adds a ramp of $\delta = 0.12 \text{ m}$ per step. In steady state the innovation carries a bias δ/K . Compute it, the mean NIS, and the chance of exceeding the gate (non-central χ^2 , $\lambda = \text{bias}^2/S$; answer 8.4 %).
- (c) Position error after 100 steps (error = $b - \delta(1 - K)/K$)?
- (d) Averaging the innovation over n steps has std $\sqrt{S/n}$: detectable at 3σ for $n = 10$ or $n = 50$?

Exercise Solution: Slow GNSS drift vs. the gate

- (a) half-width $\sqrt{3.84 \cdot 1.25} = 2.19 \text{ m}$; jump: $\text{NIS} = 25/1.25 = 20 > 3.84 \rightarrow \text{rejected}$.
- (b) bias = $0.12/0.2 = 0.6 \text{ m}$, mean $\text{NIS} = 1 + 0.36/1.25 = 1.29$, exceedance **8.4 %** vs. 5 % without attack: barely visible.
- (c) $b = 12 \text{ m}$, lag = $0.12 \cdot 0.8/0.2 = 0.48 \text{ m} \rightarrow$ error **11.5 m**: the filter follows the spoofer.
- (d) $n = 10$: $0.6/0.354 = 1.7\sigma$ (not detectable); $n = 50$: $0.6/0.158 = 3.8\sigma$ (**detectable**). A **windowed** check catches what the per-step gate misses, at the price of delay.