

1 Mathematical Foundations

Skew-Symmetric Operator:

$$\mathbf{r}^\times = \begin{bmatrix} 0 & -r_3 & r_2 \\ r_3 & 0 & -r_1 \\ -r_2 & r_1 & 0 \end{bmatrix} \quad \mathbf{x}^T \mathbf{w}^\times \mathbf{x} = 0 \quad \mathbf{C} \mathbf{w}^\times \mathbf{C}^T = [\mathbf{C} \mathbf{w}^\times, \mathbf{C}] \in \text{SO}(3)$$

Singularity/Constraint Dilemma: There is no representation of rotations that has exactly 3 parameters (i.e. no constraints) and also no singularities.

Rotation Matrices: $\mathbf{C}_{21} = \mathcal{F}_2 \cdot \mathcal{F}_1^T \implies \mathbf{r}_2 = \mathbf{C}_{21} \mathbf{r}_1$. $\mathbf{C}^T \mathbf{C} = \mathbf{1}$, $\det \mathbf{C} = 1$.

Principal Rotations: Transforms between this frame and a frame obtained by rotating about an axis by an angle. *Note the notation here is the opposite of what's used in most other courses. The matrices used to rotate a vector about a coordinate axis are the inverse of the matrices here.*

$\mathbf{C}_1(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ r_3 & -\sin \theta & \cos \theta \end{bmatrix}$ $\mathbf{C}_2(\theta) = \begin{bmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}$ $\mathbf{C}_3(\theta) = \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Euler Angles: One of 12 compositions of principal rotations. Has singularities.

Axis-Angle: $\mathbf{C} = \cos(\phi) \mathbf{1} + (1 - \cos(\phi)) \mathbf{a} \mathbf{a}^T - \sin(\phi) \mathbf{a}^\times$ constraint $\mathbf{a}^T \mathbf{a} = 1$.

Quaternion: $\mathbf{q} = q_0 + \mathbf{q} = q_0 + q_1 \mathbf{i} + q_2 \mathbf{j} + q_3 \mathbf{k} \in \mathbb{H}$ are normalized to be unit length on \mathbb{S}^3 , which forms a **double cover** of $\text{SO}(3)$. \mathbf{q} is the same rotation as $-\mathbf{q}$. $\bar{\mathbf{q}}$ (conjugate) is the opposite rotation. $\mathbf{q} = \begin{bmatrix} \eta \\ \boldsymbol{\epsilon} \end{bmatrix} = \begin{bmatrix} \cos(\phi/2) \\ \mathbf{a} \sin(\phi/2) \end{bmatrix}$.

$\mathbf{C} = (\eta^2 - \boldsymbol{\epsilon}^T \boldsymbol{\epsilon}) \mathbf{1} + 2\boldsymbol{\epsilon} \boldsymbol{\epsilon}^T - 2\eta \boldsymbol{\epsilon}^\times$.

Gibbs Vector: $\mathbf{C} = (1 + \mathbf{g}^\times)^{-1} (1 - \mathbf{g}^\times)$, $\mathbf{g} = \mathbf{a} \tan \frac{\phi}{2}$

Infinitesimal Rotations: $\mathbf{C} = 1 - \theta^\times$ where $\theta = \phi \mathbf{a}$.

Velocity in Another Frame: $\dot{\mathbf{r}}_1 = \mathbf{a}^\times + \omega_{21} \times \mathbf{r}_1 \iff \dot{\mathbf{r}}_1 = \mathbf{C}_{12}(\dot{\mathbf{r}}_2 + \omega_{21}^{21} \times \mathbf{r}_2)$

Poisson's Equation: $\dot{\mathbf{C}}_{21} = -\omega_{21}^{21} \times \mathbf{C}_{21}$

Transforms: $\mathbf{r}_1^{pi} = \mathbf{C}_{1v} \mathbf{r}_v^{pi} + \mathbf{r}_1^{vi} \iff \begin{bmatrix} \mathbf{r}_1^{pi} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{1v} & \mathbf{r}_1^{vi} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{r}_v^{pi} \\ 1 \end{bmatrix} = \mathbf{T}_{1v} \begin{bmatrix} \mathbf{r}_v^{pi} \\ 1 \end{bmatrix}$


Note vehicle pose is in \mathbf{T}_{1v} (vehicle-to-inertial), not \mathbf{T}_{vi} !

Inversion: $\mathbf{T}_{v1}^{-1} = \begin{bmatrix} \mathbf{C}_{1v}^T & -\mathbf{C}_{1v}^T \mathbf{r}_1^{vi} \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{vi} & -\mathbf{r}_1^{vi} \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{vi} & \mathbf{r}_1^{vi} \\ 0 & 1 \end{bmatrix} = \mathbf{T}_{vi}$

Combined Velocity: $\boldsymbol{\omega}_v^{vi} = \begin{bmatrix} \nu^{vi} \\ \boldsymbol{\omega}_v^{vi} \end{bmatrix}$, $\dot{\mathbf{T}}_{vi} = \begin{bmatrix} \boldsymbol{\omega}_v^{vi \times} & -\nu^{vi} \\ 0 & 0 \end{bmatrix} \mathbf{T}_{vi}$

2 Vehicle Modelling

2.1 Wheel Kinematics



For a standard wheel with radius r , rotation φ :

$$\begin{bmatrix} \cos(\alpha + \beta) & \sin(\alpha + \beta) & d \sin \beta \\ -\sin(\alpha + \beta) & \cos(\alpha + \beta) & d \cos \beta \end{bmatrix} \dot{\boldsymbol{\xi}} = \dot{\varphi} \mathbf{r}$$

$$\begin{bmatrix} -\sin(\alpha + \beta) & \cos(\alpha + \beta) & d \cos \beta \\ \cos(\alpha + \beta) & \sin(\alpha + \beta) & d \sin \beta \end{bmatrix} \dot{\boldsymbol{\xi}} = 0$$

where $\dot{\boldsymbol{\xi}} = \begin{bmatrix} v & u & \omega \end{bmatrix}^T$ is the vehicle-frame pose rate. Solving for $\dot{\varphi}$ yields **inverse differential kinematics**. Solving for $\boldsymbol{\xi}$ (by pseudoinverse $(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}$) yields **forward differential kinematics**.

For Swedish wheels with rollers at angle γ ($\dot{\varphi}_s$ usually unconstrained):

$$\begin{bmatrix} \cos(\alpha + \beta + \gamma) & \sin(\alpha + \beta + \gamma) & d \sin(\beta + \gamma) \\ -\sin(\alpha + \beta + \gamma) & \cos(\alpha + \beta + \gamma) & d \cos(\beta + \gamma) \end{bmatrix} \dot{\boldsymbol{\xi}} = \dot{\varphi} \cos \gamma + \dot{\varphi}_s \mathbf{r}_s$$

$$\begin{bmatrix} -\sin(\alpha + \beta + \gamma) & \cos(\alpha + \beta + \gamma) & d \cos(\beta + \gamma) \end{bmatrix} \dot{\boldsymbol{\xi}} = -\dot{\varphi} r \sin \gamma$$

Differential Drive Kinematics: Origin halfway between wheels (b from wheels).

$$\begin{bmatrix} \dot{\varphi}_r \\ \dot{\varphi}_l \end{bmatrix} = \frac{1}{r} \begin{bmatrix} 1 & b \\ 1 & -b \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix} \quad \begin{bmatrix} v \\ \omega \end{bmatrix} = \frac{1}{2} \begin{bmatrix} r & r \\ r/b & -r/b \end{bmatrix} \begin{bmatrix} \dot{\varphi}_r \\ \dot{\varphi}_l \end{bmatrix}$$

2.2 Euler-Lagrange

Nonholonomic constraints written as $\mathbf{H}(\mathbf{q})^T \dot{\mathbf{q}} = \mathbf{0}$ where $\mathbf{q} = [x \ y \ \theta]^T$. Null space of $\mathbf{H}(\mathbf{q})^T$ are admissible velocities $\dot{\mathbf{q}} = \mathbf{G}(\mathbf{q}) \boldsymbol{\nu}$ where $\boldsymbol{\nu}$ is generalized velocity and $\mathbf{H}(\mathbf{q})^T \mathbf{G}(\mathbf{q}) = \mathbf{0}$.

Euler-Lagrange Equation: $\frac{d}{dt} \left(\frac{\partial^T L}{\partial \dot{\mathbf{q}}} \right) - \frac{\partial^T L}{\partial \mathbf{q}} = \boldsymbol{\tau} + \mathbf{H}(\mathbf{q}) \boldsymbol{\lambda}$ where $L = T - V$, $\boldsymbol{\tau} = \mathbf{G}(\mathbf{q}) \boldsymbol{\nu}$ (generalized forces), $\boldsymbol{\lambda}$ are Lagrange multipliers. To eliminate the Lagrange multiplier term, premultiply by $\mathbf{G}(\mathbf{q})^T$.

2.3 Vehicle Dynamics Models

Generalized Vehicle Model: $\dot{\mathbf{x}} = \mathbf{A}(\mathbf{x}) \mathbf{x} + \mathbf{B}(\mathbf{x}) \mathbf{u}$ where $\mathbf{x} = [\mathbf{q} \ \mathbf{p}]^T$; with kinematics only, $\mathbf{A} = \mathbf{0}$.

Unicycle Model: $\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{G}(\mathbf{q}) \\ \mathbf{0} & -\mathbf{M}^{-1} \mathbf{D} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{p} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1} \end{bmatrix} \mathbf{u}$ where

$$\mathbf{p} = \begin{bmatrix} v \\ \omega \end{bmatrix} \quad \mathbf{G}(\mathbf{q}) = \begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ 0 & 1 \end{bmatrix} \quad \mathbf{M} = \begin{bmatrix} m & 0 \\ 0 & I \end{bmatrix} \quad \mathbf{u} = \begin{bmatrix} f \\ g \end{bmatrix} \quad \boldsymbol{\nu} = -\mathbf{D} \mathbf{p} + \mathbf{u}$$

for some damping \mathbf{D} , thrust f , steering torque g .

Bicycle Model: $\begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{G}(\mathbf{q}) \\ \mathbf{0} & -\mathbf{M}(\mathbf{q})^{-1} \mathbf{D}(\mathbf{q}, \mathbf{p}) \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{p} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{M}(\mathbf{q})^{-1} \end{bmatrix} \mathbf{u}$ where

$$\mathbf{q} = \begin{bmatrix} x \\ y \\ \psi \end{bmatrix} \quad \mathbf{p} = \begin{bmatrix} v \\ \omega \\ \varsigma \end{bmatrix} \quad \mathbf{G}(\mathbf{q}) = \begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ \tan \gamma & 0 \\ \frac{d}{a} & 1 \end{bmatrix} \quad \mathbf{M} = \begin{bmatrix} m & 0 & 0 & 0 \\ 0 & m & 0 & 0 \\ 0 & 0 & I + I_f & I_f \\ 0 & 0 & I_f & I_f \end{bmatrix}$$

and $\mathbf{M}(\mathbf{q}) = \mathbf{G}(\mathbf{q})^T \mathbf{M} \mathbf{G}(\mathbf{q})$, $\mathbf{D}(\mathbf{q}, \mathbf{p}) = \mathbf{G}(\mathbf{q})^T \mathbf{M} \mathbf{G}(\mathbf{q}) + \mathbf{D}$ for front wheel steering angle γ , rear wheel velocity v , steering speed ς , distance d between wheels.

3 Path-Tracking Control

Errors: Crosstrack error $\varepsilon_L = y_d - y$ (distance between vehicle and closest point on path) and heading error $\varepsilon_H = \theta_d - \theta$ (based on path tangent at closest point).

Feedback Linearization: Given a nonlinear system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) \mathbf{u}$, define new input $\boldsymbol{\nu}$ and find input mapping $\mathbf{u} = \mathbf{a}(\mathbf{x}) + \mathbf{b}(\mathbf{x}) \boldsymbol{\nu}$ such that $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) (\mathbf{a}(\mathbf{x}) + \mathbf{b}(\mathbf{x}) \boldsymbol{\nu}) = \mathbf{A} \mathbf{x} + \mathbf{B} \boldsymbol{\nu}$, i.e. find \mathbf{A}, \mathbf{B} such that $\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) \mathbf{a}(\mathbf{x}) = \mathbf{A} \mathbf{x}$, $\mathbf{g}(\mathbf{x}) \mathbf{b}(\mathbf{x}) \boldsymbol{\nu} = \mathbf{B} \boldsymbol{\nu}$.

Applied to unicycle model: $\begin{bmatrix} \dot{\varepsilon}_L \\ \dot{\varepsilon}_H \end{bmatrix} = \begin{bmatrix} v \sin \varepsilon_H \\ v \cos \varepsilon_H \end{bmatrix}$ (tangent approximation),

change variables $\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \varepsilon_L \\ v \sin \varepsilon_H \end{bmatrix}$ to get $\begin{bmatrix} \dot{z}_1 \\ \dot{z}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \boldsymbol{\eta}$ where

$\boldsymbol{\eta} = v \omega \cos \varepsilon_H$. Let output $\mathbf{y} = \begin{bmatrix} \alpha & 1 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$ and proportional controller

$\boldsymbol{\eta} = K(y_d - \mathbf{y})$, take $K = 4\alpha$ to get $\boldsymbol{\omega} = -\frac{4\alpha^2}{v} \frac{\varepsilon_L}{\cos \varepsilon_H} - 4\alpha \tan \varepsilon_H$.

This has singularities at $\varepsilon_H = \pm \pi/2$ or $v = 0$, and overshoots turns.

Pure Pursuit: Steer the robot towards a target lookahead point. Lookahead is hard to select for varying speeds; too short results in oscillation; too long results in slow convergence.

Stanley Controller: For bicycle models, using front axle centre as origin, $\delta(t) = \varepsilon_H(t) + \tan^{-1} \left(\frac{k \varepsilon_L(t)}{k_s + v_f(t)} \right)$ for some steering $\delta(t)$ and forward velocity $v_f(t)$, gain k and softening k_s .

Error dynamics: using $\dot{\varepsilon}_H(t) = \frac{-v_f(t) \sin(\delta(t))}{r}$, $\dot{\varepsilon}_L(t) = v_f(t) \sin(\varepsilon_H(t) - \delta(t))$,

crosstrack error has $\dot{\varepsilon}_L(t) = \frac{-k \varepsilon_L(t)}{\sqrt{1 + \left(\frac{k \varepsilon_L(t)}{k_s} \right)^2}}$, exponential decay for small ε_L .

Model Predictive Control: At each step, predict all future errors for some horizon, find optimal sequence of control inputs, apply the first step (or for a short time) and iterate again.

Parametrize control over horizon: $\begin{bmatrix} v(t) \\ \omega(t) \end{bmatrix} = \begin{bmatrix} v_0 \\ \omega_0 \end{bmatrix} + \begin{bmatrix} v_1 \\ \omega_1 \end{bmatrix} t + \begin{bmatrix} v_2 \\ \omega_2 \end{bmatrix} t^2 + \begin{bmatrix} v_3 \\ \omega_3 \end{bmatrix} t^3$ with $v_0, \dots, \omega_0, \dots$ as parameters. Simulate with full model to minimize cost function $J(\mathbf{u}) = \frac{1}{2} \sum_i w_i \varepsilon_{L,i}^2 + \frac{1}{2} \sum_k (w_{v,k} v_k^2 + w_{\omega,k} \omega_k^2)$.

4 Planning

General Approach: A strategic/high-level (global) planner, planning an optimal route based on a priori knowledge, and a tactical/low-level (reactive) planner that handles real-time obstacle avoidance.

Standard Planning Sequence: Compute a simple path, smooth the path to account for kinematics, design a trajectory to account for dynamics (i.e. adding a velocity profile), and design a controller to track the trajectory.

4.1 Reactive Planners

Potential Fields: Define a potential $V(\mathbf{x}_t) = V_{\text{att}}(\mathbf{x}_t) + V_{\text{rep}}(\mathbf{x}_t)$ where $V_{\text{att}}(\mathbf{q}) = K_{\text{att}} \rho(\mathbf{q}, \mathbf{q}^g)$ and $V_{\text{rep}}(\mathbf{q}) = K_{\text{rep}} \sum_{i=1}^n \frac{1}{\rho(\mathbf{q}, \mathbf{O}_i)^2}$ and use gradient descent (move in the direction of $-\nabla V$). Can get stuck in local minima; non-optimal; does not account for vehicle constraints.

Extended Potential Fields: Incorporate vehicle heading by adding a rotation potential, which increases if heading towards obstacles.

Trajectory Rollout: A basic form of MPC; simulates a discrete number of trajectories in a short window, checking each for collisions and scoring based on progress to goal, distance to obstacles, similarity to previous choice, etc. Accounts for vehicle constraints but can still get stuck.

4.2 Global Planners

Ideal graph-based planners should be complete, optimal, and efficient.

Resolution Completeness: For grid-based planners, if a path exists it will always be found given a fine enough grid resolution.

Probabilistic Completeness: For probabilistic planners, if a path exists, the failure probability asymptotically approaches zero as more work is performed.

Anytime Optimality: As more work is performed, better paths are found, i.e. planning can be terminated early for a feasible but less optimal path.

4.2.1 Graph Search Methods

Backward Value Iteration: Working backwards from the final state using dynamic programming to find optimal cost-to-go. At each step for each node iterate as $G^*(x) = \min_u \{ l(x, u) + G^*(f(x, u)) \}$.

Generic Forward Search Algorithm: At each step, take the next state from the queue, mark it as visited, discover all unvisited nodes reachable from this node, and add each reachable unvisited node to the queue. Queue prioritization leads to different algorithms: LIFO (stack) — DFS; FIFO (queue) — BFS; lowest cost-to-come — Dijkstra's algorithm.

A* Algorithm: Prioritize with cost-to-come plus lower bound heuristic (necessary for optimality). After reaching the goal, prune all nodes in queue with a priority worse than the goal and run until queue is empty.

Lifelong Planning A* (LPA*): Result reuses from previous A* runs when only a few edge costs have changed. All cells that have an incoming edge cost change are checked for inconsistency, by recomputing the cost-to-come based on predecessors. Inconsistent cells are set to infinite cost and successors are checked, until all cells are consistent, when the search can be restarted.

Other Improvements: Focused Dynamic A* (D*), D*-Lite are faster than LPA*;

Field D* makes smoother paths.

4.3 Probabilistic Roadmaps

PRM Algorithm: Initialize with start and goal nodes. Repeatedly sample configurations randomly and check if they are feasible, and add as milestones if so. For each milestone try to connect to neighbours; if path is collision-free add it as an edge.

Visibility Set: $V(q) = \{ (q', q'') \in F \}$, i.e. the set of all configurations that can be connected to q by a straight line edge. Size of $V(q)$ is called **expansiveness**. Summing over all $q \in C$ produces expansiveness of the entire configuration space. High expansiveness indicates open spaces, wider, and straighter passages.

Visibility Graph: Each vertex on each obstacle is a milestone and connected to each other possible vertex. Produces theoretically optimal paths in 2D but not higher dimensions.

Workspace-Guided Sampling: Find narrow passages in workspace and map into the configuration space to sample those areas more.

Gaussian Sampling: First sample uniformly, then another configuration with a Gaussian distribution centered at the first; if one sample is in free space and the other is in collision, retain the one in free space, otherwise skip. Samples tend to be near the boundaries of free space.

Bridge Sampling: Sample 2 configurations using the same procedure as Gaussian sampling; if both are in collision, find the midpoint between the samples, and retain that sample if it's in free space. Samples tend to be in narrow passages, which may scale better in higher dimensions.

Bounding Volume Hierarchy: Enclose objects into successively tighter bounding volumes, and checking collisions starting from the largest. If in collision, check again with all sub-volumes enclosed within until the desired level of accuracy is reached. Bounding volumes can be computed in advance in static environments.

Adaptive Collision Checks: Each point checked along the path eliminates a section of the path based on distance to closest obstacle. Check end points first, then bisect any remaining sections.

Lazy Collision Checks: Delay collision checking until path has been found; if in collision, recompute the path.

4.4 RRT and RRT*

RRT Algorithm: Sample a point in space, identify the closest node and the input that drives from the node towards the sampled point. Forward simulate for a short time to determine the new vehicle configuration. If no collision, add the new node to the tree. Probabilistically complete but not optimal.

Probability of expanding each vertex is proportional to its Voronoi region, which heavily favours expansion over refinement. Sample within an expanding box of the origin to encourage more refinement.

RRT* Algorithm: Each node carries a cost-to-come. After connecting the new node, consider nodes within a ball of the new node. If the new node can be connected to another node within the neighbourhood to get a lower cost, connect to this node instead. Then, check all other nodes within the ball to see if going from the new node to the other node results in a lower cost. If so, rewire the node so its parent becomes the new node. Results in asymptotically optimal paths.

Informed RRT*: After finding an initial solution, use a heuristic to form an ellipse where possible path improvements can be, and sample only within this space.

4.5 Planning for Autonomous Driving

Structure: Roadmap level mission planner, action-level behavioural planner (FSM, rule-based, reinforcement learning), path and velocity-level local planner (lattice in structured scenes, sampling-based in unstructured scenes, variational min $_{\delta x} J(x + \delta x)$ for optimality; init with global planner solution).

Conformal Lattice Planner: Fits control actions/grid to the road. Solve with boundary conditions (x, y, θ, κ) with max curvature, min velocity constraints.

Cubic Spirals: Numerically integrate $\kappa(s) = a_3 s^3 + a_2 s^2 + a_1 s + a_0 \implies \theta(s) = \theta_0 + \int_0^s \kappa(s') ds'$, $x(s) = x_0 + \int_0^s \cos(\theta(s')) ds'$, $y(s) = y_0 + \int_0^s \sin(\theta(s')) ds'$. Specify boundary constraints, check curvature at 1/3, 2/3 points. For optim remap params to 4 control points and arc length and use soft constraints.

Objectives: Path length $s_f = \int_{x_i}^{x_f} \sqrt{1 + \left(\frac{dy}{dx} \right)^2} dx$, time $T_f = \int_0^{s_f} \frac{1}{v(s)} ds$, bending energy $\int_0^{s_f} \|\kappa(s)\|^2 ds$ (distributes curvature).

Velocity Generation: Track reference and minimize $\int_0^{s_f} \|v(s) - v_{ref}(s)\| ds$ or Hinge loss (penalize too fast, not too slow). Penalize jerk $\int_0^{s_f} \|\ddot{x}(s)\|^2 ds$ for smoothness. Get desired v_f based on minimum of v_{ref} , lead vehicle speed, and acceleration within friction ellipse/comfort rectangle ($a_{lat} = v^2 \kappa$). Connect to current v_0 with linear ($v_f = \sqrt{2as + v_0^2}$) or trapezoidal ramp.

5 Probability

Total Probability: $\int_{a_1}^{b_1} \dots \int_{a_n}^{b_n} p(x_1, \dots, x_n) dx_1 \dots dx_n = 1$

Bayes' Rule: $p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})}$

Moments:

- Always 1 by total probability.
- Mean: $\boldsymbol{\mu} = \mathbb{E}[\mathbf{x}] = \int \boldsymbol{x} p(\mathbf{x}) d\mathbf{x}$.
- (Co)variance: $\boldsymbol{\Sigma} = \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T]$.
- Skewness
- Kurtosis

Independence: $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{x})p(\mathbf{y})$; implies uncorrelated.

Uncorrelated: $\mathbb{E}[\boldsymbol{x} \boldsymbol{y}^T] = \mathbb{E}[\boldsymbol{x}] \mathbb{E}[\boldsymbol{y}^T] \iff \boldsymbol{\Sigma}_{\boldsymbol{x} \boldsymbol{y}} = \mathbf{0}$

Shannon Information: $H(\mathbf{X}) = -\mathbb{E}[\ln p(\mathbf{x})] = -\int_a^b p(\mathbf{x}) \ln p(\mathbf{x}) d\mathbf{x}$

5.1 Gaussians

PDF: $p(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^N \det \boldsymbol{\Sigma}}} \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right)$

Factoring: $p(\mathbf{x}, \mathbf{y}) = \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu}_x \\ \boldsymbol{\mu}_y \end{bmatrix}; \begin{bmatrix} \boldsymbol{\Sigma}_{xx} & \boldsymbol{\Sigma}_{xy} \\ \boldsymbol{\Sigma}_{yx} & \boldsymbol{\Sigma}_{yy} \end{bmatrix} \right) = p(\mathbf{x}|\mathbf{y})p(\mathbf{y})$

$$p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}_x + \boldsymbol{\Sigma}_{xy} \boldsymbol{\Sigma}_{yy}^{-1} (\mathbf{y} - \boldsymbol{\mu}_y), \boldsymbol{\Sigma}_{xx} - \boldsymbol{\Sigma}_{xy} \boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\Sigma}_{yx})$$

Direct Product: $\boldsymbol{\Sigma}^{-1} = \sum_{k=1}^K \boldsymbol{\Sigma}_k^{-1}, \boldsymbol{\mu} = \boldsymbol{\Sigma} \sum_{k=1}^K \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\mu}_k$

Linearization: $p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) d\mathbf{x} \approx \mathcal{N}(\mathbf{g}(\boldsymbol{\mu}_x), \mathbf{R} + \mathbf{G} \boldsymbol{\Sigma}_{xx} \mathbf{G}^T)$ where

$p(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_x, \boldsymbol{\Sigma}_{xx})$, $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{g}(\mathbf{x}), \mathbf{R})$, and $\mathbf{G} = \left. \frac{\partial \mathbf{g}(\mathbf{x$

6.2 Exteroceptive Sensors

LiDAR: Measures distance using time-of-flight of laser pulses, determined by measuring phase shift of an intensity modulated signal. Beams steered using spinning heads, mirrors, or MEMS. Measurement model: $\rho_{ij} = \sqrt{x_{ij}^2 + y_{ij}^2 + z_{ij}^2} + \Delta\rho$,

$$\theta_{ij} = \tan^{-1} \left(\frac{y_{ij}}{x_{ij}} \right) + \Delta\theta; \alpha_{ij} = \tan^{-1} \left(\frac{z_{ij}}{\sqrt{x_{ij}^2 + y_{ij}^2}} \right) + \Delta\alpha.$$

ToF Camera: Like LiDAR but captures the entire point cloud at once. Much shorter range; suitable for indoor use since sunlight causes problems.


Sonic Rangefinder/SODAR: Measures distance using ToF of ultrasonic pulse. Limited range, resolution, and rate. Popular underwater.

Radar: Measures distance and velocity using radio waves. Good for distance, unreliable for angle.

GNSS/GPS: Measures position using distance to satellites. At least 4 for solution (3 position + 1 clock offset). Satellites use atomic clocks and correct for relativistic effects. Uses fixed WGS84 globe model.

Base carrier frequency at 1575 MHz, modulated to carry a code (chips) at 1.023 Mcps; 1023 chips form a unique orthogonal gold code identifying the satellite (repeated at 1 kHz). Correlation with code determines ToF. 20 code repeats encodes 1 bit (50 bps), carrying orbital ephemeris corrections, time, etc. Raw distances (pseudoranges) are used to solve for position and clock bias using NLLS.

Errors from satellite geometry (dilution of precision), ionospheric/tropospheric effects, multi-path, giving 10m accuracy. DGPS uses known base station to eliminate atmospheric and satellite clock error (10 cm relative to base). RTK GPS also uses carrier phase (2 cm).



6.3 Cameras

6.3.1 Pinhole Camera Model

Basic Projective Map: Normalized by focal length,

$$\tilde{x} = \mathcal{P}_z(\mathbf{p}) = \mathbf{p}/p_z = \begin{bmatrix} x/z & y/z & 1 \end{bmatrix}^T$$

Intrinsic Matrix: For focal lengths f_x, f_y , center c_x, c_y (pixels), skew s :

$$\begin{bmatrix} x_s \\ y_s \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & s & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x/z \\ y/z \\ 1 \end{bmatrix} = \mathbf{K} \tilde{\mathbf{p}}$$

This projects 3D to pixel coords. Note one can also normalize after projection. One can obtain a ray of points that project to the same pixel by multiplying the augmented pixel vector by \mathbf{K}^{-1} .

FoV, Sensor Width, Focal Length: $\tan \frac{\theta}{2} = \frac{W}{2f}$

Extrinsic Matrix: $\mathbf{E}_{CW} \in SE(3)$ contains the world frame to camera frame transform. $\mathbf{E}_{CW} = \mathbf{E}_{WC}^{-1}$ where \mathbf{E}_{WC} is the camera pose in world frame.

Projective Matrix: $\tilde{\mathbf{P}} = \begin{bmatrix} \mathbf{K} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{C} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} = \tilde{\mathbf{K}}\mathbf{E}$. Project a world point to image plane as $\tilde{x}_s = \tilde{\mathbf{P}}\mathbf{p}_w$, then normalize by the third element.

6.3.2 Distortion & Optical Effects

Plumb Bob Distortion Model: From normalized image coordinates to distorted coordinates, where $r = \sqrt{x_n^2 + y_n^2}$ and x_d, y_d normalized by f_x, f_y :


$$\begin{bmatrix} x_d \\ y_d \end{bmatrix} = (1 + \kappa_1 r^2 + \kappa_2 r^4 + \kappa_3 r^6) \begin{bmatrix} x_n \\ y_n \end{bmatrix} + \begin{bmatrix} 2\tau_1 x_n y_n + \tau_2 (r^2 + 2x_n^2) \\ 2\tau_2 x_n y_n + \tau_1 (r^2 + 2y_n^2) \end{bmatrix}$$

Note distortion is applied after extrinsics & projective map, but before doing intrinsics.

Radial Distortion: Barrel (corner points pushed in), pincushion (corner points pulled out), or mustache (combination).

Tangential Distortion: Pixels shifted in one direction (further pixels shifted more), due to the lens and imaging plane not being parallel.

Unwarping: Undoing distortion by precomputing distorted location of each pixel, then using bilinear interpolation.




6.3.3 Image Operations & Transformations

Histogram Equalization: $T(r_k) = \sum_{j=0}^k p_r(r_j) = \sum_{j=0}^k \frac{T_j}{n}$ where $T(r_k)$ is the CDF and $p_r(r_j)$ is the PDF (normalized histogram) as a function of pixel intensity. Each pixel value is mapped to the proportion of pixels that it's brighter than, resulting in a flat-ish histogram.

Linear Filtering: $g(i, j) = f * h = \sum_{k,l} f(i-k, j-l)h(k, l)$ where $h(k, l)$ is the kernel. Note how the kernel is flipped. Separable kernels can be expressed as $\mathbf{K} = \mathbf{v}\mathbf{h}^T$ and is equivalent to convolution by \mathbf{v} then \mathbf{h} .

Geometric Transform: Applies to domain, $g(\mathbf{x}) = f(\mathbf{h}(\mathbf{x}))$. Better to compute the inverse transformation and map each pixel in the output to a location in the input, and bilinearly interpolate.

6.3.4 Stereo Vision



Epipolar Geometry: Form the epipolar plane from a point and the cameras' optical centres; the intersection of the epipolar plane with each imaging plane forms an epipolar line, which is where correspondences lie.

Stereo Rectification: Transforming and unwarping the images so the imaging planes become parallel, at the same depth, and rotationally aligned (fronto-parallel or standard rectified geometry). After rectification all epipolar lines are horizontal.

Disparity: Difference in x values of a point between cameras. $Z = \frac{f_b}{d}$ where b is baseline and d is disparity. Larger disparity indicates a point is closer.

$$\text{Stereo Camera Model: } \begin{bmatrix} x_0 \\ y_0 \\ d \end{bmatrix} = \mathbf{f}(\mathbf{p}_c) = \frac{f}{Z} \begin{bmatrix} X \\ Y \\ b \end{bmatrix} + \begin{bmatrix} c_x \\ c_y \\ 0 \end{bmatrix}$$

Baseline: Distance between cameras. Larger baseline improves accuracy, but makes matching more difficult as overlapping FOV decreases and foreshortening is increased (oblique objects appear as different sizes). Wider baseline also makes errors more Gaussian-like (shorter tail); for short baseline, using Gaussian errors biases the points to be closer.

Combined Stereo Model: Forward model produces pixel locations in both cameras. Frame origin is halfway between two cameras.

$$\text{Forward: } \mathbf{f}(\mathbf{p}) = \mathbf{M}^{-1}\mathbf{p} = \begin{bmatrix} f_u & 0 & c_u & f_u b/2 \\ 0 & f_v & c_v & 0 \\ f_u & 0 & c_u & -f_u b/2 \\ 0 & f_v & c_v & 0 \end{bmatrix} \frac{1}{Z} \mathbf{p} = \begin{bmatrix} [u] \\ [v] \\ [u] \\ [v] \end{bmatrix}$$

$$\text{Inverse model: } \mathbf{g}(\mathbf{y}) = \frac{b}{u_l - u_r} \begin{bmatrix} \frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) \\ \frac{1}{2}(u_l + u_r) - c_u \\ f_u \\ 0 \end{bmatrix}$$

$$\text{Jacobians: forward } \frac{\partial \mathbf{f}}{\partial \mathbf{p}} = \mathbf{M}^{-1} \frac{\partial \mathbf{f}}{\partial \mathbf{p}} = \begin{bmatrix} 1 & 0 & -p_1/p_3 & 0 \\ 0 & 1 & -p_2/p_3 & 0 \\ 0 & 0 & -p_1/p_3 & 1 \\ 0 & 0 & -p_2/p_3 & 1 \end{bmatrix} \text{inverse } \frac{\partial \mathbf{g}}{\partial \mathbf{y}} =$$

$$\frac{b}{(u_l - u_r)^2} \begin{bmatrix} -\frac{f_u}{f_v} (-u_r + c_u) & -\frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) & \frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) & 0 \\ -\frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) & -\frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) & -\frac{f_u}{f_v} (\frac{1}{2}(u_l + u_r) - c_u) & \frac{f_u}{f_v} \\ 0 & 0 & -p_1/p_3 & 1 \\ 0 & 0 & -p_2/p_3 & 1 \end{bmatrix}$$

Pixel-Disparity Stereo Model: Forward model produces a pixel location and disparity. Frame origin is at the left camera.

$$\text{Forward: } \mathbf{y} = \mathbf{f}(\mathbf{p}) = \mathbf{M}^{-1}\mathbf{p} = \begin{bmatrix} f_u & 0 & c_u & 0 \\ 0 & f_v & c_v & 0 \\ 0 & 0 & f_u b & p_3 \end{bmatrix} \frac{1}{p_3} \mathbf{p} = \begin{bmatrix} [u] \\ [v] \\ [d] \end{bmatrix}$$

$$\text{Inverse model: } \mathbf{g}(\mathbf{y}) = \frac{b}{d} \begin{bmatrix} u_l - c_u \\ \frac{f_u}{f_v} (v_l - c_v) \\ f_u \end{bmatrix}$$

$$\text{Forward Jacobian: } \frac{\partial \mathbf{f}}{\partial \mathbf{p}} = \frac{1}{p_3} \mathbf{M} \begin{bmatrix} 1 & 0 & -p_1/p_3 & 0 \\ 0 & 1 & -p_2/p_3 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -p_1/p_3 & 1 \end{bmatrix}$$

$$\text{Inverse Jacobian: } \frac{\partial \mathbf{g}}{\partial \mathbf{y}} = \frac{b}{d^2} \begin{bmatrix} d & 0 & c_u - u_l \\ \frac{f_u}{f_v} d & \frac{f_u}{f_v} d & -u_l \\ 0 & \frac{f_u}{f_v} d & \frac{f_u}{f_v} (c_v - v_l) \\ 0 & -f_u & -f_u \end{bmatrix}$$

6.3.5 Image Features

Features should be salient (distinctive), local, repeatable (can be found in other images) and compact.

Weighted Summed Squared Difference: Images I_0, I_1 , displacement \mathbf{u}

$$E_{WSSD}(\mathbf{u}) = \sum_w w(\mathbf{x}_i)(I_1(\mathbf{x}_i + \mathbf{u}) - I_0(\mathbf{x}_i))^2$$

For $I_0 = I_1$ this is **autocorrelation** for a small $\Delta\mathbf{u}$; high autocorrelation means the area is distinctive. $E_{AC}(\Delta\mathbf{u}) \approx \Delta\mathbf{u}^T \mathbf{A} \Delta\mathbf{u}$. As convolution, $\mathbf{A} = w * \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$.

Harris Corner Detector: det $\mathbf{A} - \alpha \text{tr} \mathbf{A}^2 = \lambda_0 \lambda_1 - \alpha(\lambda_0 + \lambda_1)^2$. Largest when both eigenvalues are large, indicating distinctiveness in both directions, i.e. a corner. **SIFT:** Scale and rotation invariant, capable but slow. Construct scale space representation by repeated blurring and downsampling. Find keypoints with Difference of Gaussians (similar to LoG). Reject keypoints using Hessians similar to Harris. Assign each keypoint an orientation using gradients of points around the keypoint in a histogram. Assemble final descriptor using 4x4 grid of subpatches, each subpatch having a gradient orientation histogram.

SURF: Like SIFT but faster. Uses box filters instead of Gaussians (can use integral images). Smaller descriptor size (half as SIFT, 64-dimensional vector).

FAST: Very fast but not scale or rotation invariant. For each pixel look at the 16 pixels in a ring around it, if there are 12 contiguous pixels that are brighter or darker than the centre it is a feature. Use non-maximum suppression (suppress non-maximum features in an area around each maximum).

BRISK: Sample pairs of points in a circular pattern around the centre (Gaussian applied around each sampled point); compare the pairs and generate one bit in the feature for each pair.

BRIEF: Draw points randomly (one of several distributions) around the centre, compare them similarly to BRISK to form descriptor. Sampling pattern is decided once and used consistently across images. 512-bit features. Fast but not scale/rotation invariant and sensitive to noise.

ORB: Combines FAST and BRIEF to get a fast and reliable feature. Downsample similarly to SIFT and apply FAST at all scales to find keypoints. Compute orientation using intensity centroid. Use rotated BRIEF (rotate sampling pattern or image subpatch) to form descriptor.

Learning-Based Detectors: Often use classic features to generate training data for supervised methods, or image + pose sequences for unsupervised methods. LIFT is one such method using 3 CNNs to detect, estimate orientation, and generate descriptor, trained using contrastive learning.

7 Odometry

7.1 Dead Reckoning/Wheel Odometry

Key Characteristics: Mean is unbiased, variance is unbounded.

Derivation: Get $\hat{q}(\varphi)$ from kinematics (transform (v, u, ω) to inertial frame first), discretize to get $q(t+h)$ and split $q(t) = q(t) + \delta q(t)$.

Linearize about $q(t)$ and drop product of Gaussians to get $q(t+h) + \delta q(t+h)$. Take $q(t+h) = \mathbb{E}[q(t+h)] + \delta q(t+h)$.

Subtract $q(t+h)$ to get linear update equation $\delta q(t+h) = \mathbf{A}(t)\delta q(t) + \mathbf{B}(t)w(t)$. Take $\Sigma(t+h) = \mathbb{E}[\delta q(t)\delta q(t)^T] = \mathbf{A}(t)\Sigma(t)\mathbf{A}(t)^T + \mathbf{B}(t)\mathbf{Q}\mathbf{B}(t)^T$ where \mathbf{Q} is covariance of $w(t)$.

Differential Drive: Let $\mathbf{R} = \begin{bmatrix} r/2 & r/2 \\ r/(2b) & -r/(2b) \end{bmatrix}$

$$\hat{q}(t+h) = \hat{q}(t) + \begin{bmatrix} \cos \theta(t) & 0 \\ \sin \theta(t) & 0 \\ 0 & 1 \end{bmatrix} \mathbf{R} \begin{bmatrix} \Delta\varphi_r(t) \\ \Delta\varphi_l(t) \end{bmatrix}$$

$$\delta q(t+h) = \delta q(t) + \begin{bmatrix} -\sin \theta(t) & 0 \\ \cos \theta(t) & 0 \\ 0 & 0 \end{bmatrix} \mathbf{R} \begin{bmatrix} \Delta\varphi_r(t) \\ \Delta\varphi_l(t) \end{bmatrix} \delta \theta(t) + \begin{bmatrix} \cos \theta(t) & 0 \\ \sin \theta(t) & 0 \\ 0 & 1 \end{bmatrix} \mathbf{R} \begin{bmatrix} w_r(t) \\ w_l(t) \end{bmatrix}$$

Params: $r = 2 \frac{x_{\text{true}}}{\sum_k (\Delta\varphi_r(kh) + \Delta\varphi_l(kh))}; b = \frac{r}{2} \frac{\sum_k (\Delta\varphi_r(kh) - \Delta\varphi_l(kh))}{2\pi N}$

7.2 Feature-Based Stereo Visual Odometry

Given images captured in frame a and frame b , find T_{ba} . Use where high accuracy short-term localization is needed, e.g. indoors, Mars, or high wheel slip environments.



Point Cloud Alignment: Given $\mathbf{p}_a^j, \mathbf{p}_b^j$, minimize cost

$$E(C_{ba}, \mathbf{r}_a) = \frac{1}{2} \sum_{j=1}^J (\mathbf{p}_b^j - C_{ba}(\mathbf{p}_a^j - \mathbf{r}_a))^T \Sigma^j (\mathbf{p}_b^j - C_{ba}(\mathbf{p}_a^j - \mathbf{r}_a))$$

with matrix weights $\Sigma^j = (G_b^j R_b^j G_b^{jT} + C_{ba} G_a^j R_a^j G_a^{jT} C_{ba}^T)^{-1}$ where $G_b^j = \frac{\partial \mathbf{g}}{\partial \mathbf{y}} \mathbf{f}(\mathbf{p}_b^j)$, $G_a^j = \frac{\partial \mathbf{g}}{\partial \mathbf{y}} \mathbf{f}(\mathbf{p}_a^j)$, and R_b^j, R_a^j are pixel measurement covariances.

Closed-form solution for scalar weights:

- Align centroids: $\mathbf{p}_a = \sum_{j=1}^J w^j \mathbf{p}_a^j, \mathbf{p}_b = \sum_{j=1}^J w^j \mathbf{p}_b^j$
- Outer product: $\mathbf{W}_{ba} = \frac{1}{\sum_{j=1}^J w^j} \sum_{j=1}^J w^j (\mathbf{p}_b^j - \mathbf{p}_b)(\mathbf{p}_a^j - \mathbf{p}_a)^T$
- Compute SVD $\mathbf{V} \mathbf{S} \mathbf{U}^T = \mathbf{W}_{ba}$, then compute rotation $C_{ba} = \mathbf{V} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \det(\mathbf{U}) \det(\mathbf{V}) \end{bmatrix} \mathbf{U}^T$, and translation $\mathbf{r}_a = -C_{ba}^T \mathbf{p}_b + \mathbf{p}_a$.

Solution is unique if rank $\mathbf{W}_{ba} \geq 2$. If rank $\mathbf{W}_{ba} = 2$ points are coplanar (unique if we enforce no reflection). If rank $\mathbf{W}_{ba} = 1$ points are colinear.

7.3 LiDAR Scan Matching

Given reference $M = \{\mathbf{p}_1^M, \dots, \mathbf{p}_M^M\}$, data $D = \{\mathbf{p}_1^D, \dots, \mathbf{p}_D^D\}$, where $N_M = N_D$ (i.e. one-to-one correspondence), find $T_{MD} = \{C_{MD}, \mathbf{t}_{MD}^D\}$.

ICP Algorithm: Until convergence, do:

- Transform $\mathbf{p}_M^j = C_k(\mathbf{p}_M^j - \mathbf{t}_k)$ to form \tilde{D}
- Find correspondences $\{I, J\} = \text{NN}(M, \tilde{D})$
- Error: $\mathcal{J} = \frac{1}{2} \sum_{j=1}^J w^j (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))^T (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))$
- Gradient descent: $\mathbf{t}_{k+1} = \mathbf{t}_k + \alpha \nabla \mathcal{J}$

Subsampling: *Uniform* (voxel-based): take one point per voxel, preserves structure; *random*: preserves density, any lose structure; *feature-based*: identifies structural features to preserve environment traits; *normal-space*: sample a subset with surface normals uniformly distributed, good for smooth areas with sparse features.

Surface Normal Estimation: For a neighbourhood of local points, try to fit a plane $\mathbf{a}^T \mathbf{p} = c$ by minimizing $\frac{1}{k} \sum_{i=1}^k (\mathbf{a}^T \mathbf{p}_i - c)^2$; equivalent to PCA of $M = \frac{1}{k} \sum_{i=1}^k (\mathbf{p}_i - \bar{\mathbf{p}})(\mathbf{p}_i - \bar{\mathbf{p}})^T$, where the eigenvector of the smallest eigenvalue (least variance) is the normal.

Data Association: *Closest compatible point*: check normals/colours/etc. within range; *normal shooting*: first point along the surface normal direction; *outlier rejection* (by checking distances of neighbours).

Point-to-Plane Matching: Only penalize errors in the normal direction;

$$\mathcal{J} = \frac{1}{2} \sum_{j=1}^J w^j (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))^T \mathbf{n}^j \mathbf{n}^{jT} (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))$$

Plane-to-Plane Matching: Uses normals in both sets;

$$\mathcal{J} = \frac{1}{2} \sum_{j=1}^J (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))^T (\Sigma_M^j + C_k \Sigma_D^j C_k^T)^{-1} (\mathbf{p}_M^j - C_k(\mathbf{p}_D^j - \mathbf{t}_k))$$

$\Sigma_k^M = C(\mathbf{n}_k^M) \Sigma_c C(\mathbf{n}_k^M)^T, \Sigma_D^j = C(\mathbf{n}_D^j) \Sigma_c C(\mathbf{n}_D^j)^T, \Sigma_c = \text{diag}(\epsilon, 1, 1)$.

Normal Distribution Transform (NDT): Divide into cells and model each as $\rho_{c_i}(\mathbf{p}) = \exp(-\frac{1}{2}(\mathbf{p} - \boldsymbol{\mu}_i)^T \Sigma_i^{-1}(\mathbf{p} - \boldsymbol{\mu}_i))$; score each point in D based on the Gaussian in its cell: $\Lambda(T) = -\sum_{\mathbf{p}_D^j \in D} \rho_{c_i}(T(\mathbf{p}_D^j))$.

Practical Considerations: Motion compensation; anisotropic (e.g. distance-based) filtering to remove low-information points; coarse-to-fine alignment (starting with smaller subsample); quality estimation via final association covariance, overlap between sections, etc.

8 Localization and Mapping

Map Types by Application: *Planning:* tracks occupancy, safe/unsafe places; *application:* tracks spatially distributed properties, e.g. topography; *navigation:* tracks localization features, may be sparse.

Map Types by Information: *Metric:* retains scale, metric consistency; *topometric:* locally metric, globally topographic; *topological:* no distances, just interconnectedness.

8.1 Occupancy Grid Mapping

Assumptions: Cell occupancy independent from other cells (no space connectivity) or robot movement, measurements independent given robot states and map.

Probabilistic Update: $p(c_k|y^t, x^t) = \frac{p(y^t|c_k, x^t)p(c_k|y^{t-1}, x^{t-1})}{p(y^t)}$

Log-Odds: $l(x) = \ln\left(\frac{p(x)}{1-p(x)}\right) \iff p(x) = \frac{\exp(l(x))}{1+\exp(l(x))}$

Log-Odds Update: $l(c_k|y^t, x^t) = l(c_k|y^{t-1}, x^{t-1}) + \ln\left(\frac{p(y^t|c_k, x^t)}{1-p(y^t|c_k, x^t)}\right)$ where the last term is α if c_k near measurement, $-\beta$ if c_k closer than measurement, 0 otherwise.

8.2 Terrain Assessment

Determine which parts of terrain are safe to traverse and which parts are not by fitting planes to local sections. Use RANSAC to reject outliers. Use info such as plane height and slope to determine traversability.

Regular Grid: Patch j : let $A_j = \begin{bmatrix} x_{1j} & y_{1j} \\ \vdots & \vdots \\ 1 & x_{nj} \end{bmatrix}, x_j = \begin{bmatrix} a_j \\ b_j \\ c_j \end{bmatrix}, b_j = \begin{bmatrix} z_{1j} \\ \vdots \\ z_{nj} \end{bmatrix}$

minimize error $e_j = A_j x_j - b_j$, with solution $x_j = (A_j^T A_j)^{-1} A_j^T b_j$.

Recursive Formulation: Let $p_{ij} = \begin{bmatrix} 1 & x_{ij} & y_{ij} & z_{ij} \end{bmatrix}^T$ and define

$$P_j = \sum_i p_{ij} p_{ij}^T = \begin{bmatrix} A_j^T A_j & A_j^T b_j \\ (A_j^T b_j)^T & b_j^T b_j \end{bmatrix} = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \\ p_{41} & p_{42} & p_{43} & p_{44} \end{bmatrix} = \begin{bmatrix} Q^{-1} & \\ & r \end{bmatrix} \quad P_{44}$$

- Number of points: $n_j = p_{11}$
- Plane fit: $[a_j \ b_j \ c_j]^T = Qr$
- Residuals: $e_j = p_{44} - r^T Qr$
- Average height: $h_j = \frac{p_{14}}{n_j}$
- Max slope: $\alpha_j = \cos^{-1} \frac{1}{\sqrt{a_j^2 + c_j^2 + 1}}$

8.3 Filtering-Based Localization

Markov Assumption: The conditional probabilities of future states are independent of any past states, given the current state; i.e. the state contains enough information to predict the future.

Bayes Filter: $p(x_k|\tilde{x}_0, v_{1:k}, y_{0:k}) = \int p(x_k|x_{k-1}, v_k)p(x_{k-1}|\tilde{x}_0, v_{1:k-1}, y_{0:k-1})d\tilde{x}_0$

Particle Filtering: Representing the robot state PDF with particles, where each particle is a hypothesis:

1. Prediction: $\begin{bmatrix} \tilde{x}_{k-1,m} \\ w_{k,m} \end{bmatrix} \sim p(x_{k-1}|\tilde{x}_0, v_{1:k-1}, y_{0:k-1})p(w_k)$ from prior and motion noise; generate predictions $\tilde{x}_{k,m} = f(\tilde{x}_{k-1,m}, v_k, w_{k,m})$.
2. Correction: Assign weights $w_{k,m} = \eta p(y_k|\tilde{x}_{k,m})$ by using expected measurements $y_{k,m} = g(\tilde{x}_{k,m}, 0)$ and assuming $p(y_k|\tilde{x}_{k,m}) = p(y_k|\tilde{y}_{k,m})$; resample based on weights to get new particles $x_{k,m}$.

Practical considerations: Maintaining sample variance by adding more noise, adding samples drawn uniformly from sample space, or resampling less often; adapting particle count online using weight sum heuristic.

Madow Systematic Sampling: Create M bins where bin m has size $w_{k,m}$; starting from random, take steps of $1/M$ and sample from each visited bin. Helps preserve particle diversity.

Kalman Filter: Linear model $x_k = A_{k-1}x_{k-1} + v_k + w_k, y_k = C_k x_k + n_k$:

1. Prediction: $\begin{cases} \hat{P}_k = A_{k-1}\hat{P}_{k-1}A_{k-1}^T + Q_k \\ \hat{x}_k = A_{k-1}\hat{x}_{k-1} + v_k \end{cases}$
2. Kalman gain: $K_k = \hat{P}_k C_k^T (C_k \hat{P}_k C_k^T + R_k)^{-1}$
3. Correction: $\begin{cases} \hat{P}_k = (1 - K_k C_k) \hat{P}_k \\ \hat{x}_k = \hat{x}_k + K_k (y_k - C_k \hat{x}_k) \end{cases}$

With correct modeling KF is *consistent* and *unbiased* (correct mean and covariance).

Extended Kalman Filter: Linearize nonlinear models:

$$x_k = f(x_{k-1}, v_k, w_k) \approx \hat{x}_k + F_{k-1}(x_{k-1} - \hat{x}_{k-1}) + w'_k$$

$$y_k = g(x_k, n_k) \approx \hat{y}_k + G_k(x_k - \hat{x}_k) + n'_k$$

where $F_{k-1} = \frac{\partial f}{\partial x_{k-1}}, G_k = \frac{\partial g}{\partial x_k}, w'_k = \frac{\partial f}{\partial w_k} w_k, n'_k = \frac{\partial g}{\partial n_k} n_k, f$ Jacobians evaluated at $(\hat{x}_{k-1}, v_k, 0)$ and g Jacobians at $(\hat{x}_k, 0)$.

1. Prediction: $\begin{cases} \hat{P}_k = F_{k-1}\hat{P}_{k-1}F_{k-1}^T + Q'_k \\ \hat{x}_k = f(\hat{x}_{k-1}, v_k, 0) \end{cases}$
2. Kalman gain: $K_k = \hat{P}_k G_k^T (G_k \hat{P}_k G_k^T + R'_k)^{-1}$
3. Correction: $\begin{cases} \hat{P}_k = (1 - K_k G_k) \hat{P}_k \\ \hat{x}_k = \hat{x}_k + K_k (y_k - g(\hat{x}_k, 0)) \end{cases}$

where $Q'_k = \frac{\partial f}{\partial w_k} Q_k \left(\frac{\partial f}{\partial w_k}\right)^T, R'_k = \frac{\partial g}{\partial n_k} R_k \left(\frac{\partial g}{\partial n_k}\right)^T$. EKF can be biased and inconsistent for severe nonlinearity since it linearizes about the estimate.

Sigma Point Transform: For L -dimensional state and parameter κ :

1. Compute $2L + 1$ sigma points, where $x_0 = \mu_x, x_i = \mu_x + \sqrt{L + \kappa} \text{col}_i L$, and $x_{i+L} = \mu_x - \sqrt{L + \kappa} \text{col}_i L$, where $LL^T = \Sigma_x$.
2. Pass each through the nonlinearity: $y_i = g(x_i)$.
3. Output Gaussian: $\mu_y = \sum_{i=0}^{2L} \alpha_i y_i, \Sigma_y = \sum_{i=0}^{2L} \alpha_i (y_i - \mu_y)(y_i - \mu_y)^T$, with weights $\alpha_i = \frac{\kappa}{L + \kappa}, \alpha_0 = \frac{1}{2} \frac{1}{L + \kappa}$.

Unscented Kalman Filter:

1. State Prediction:
 - (a) Extract sigma points z_i from $\mu_z = \begin{bmatrix} \hat{x}_{k-1} \\ 0 \end{bmatrix}, \Sigma_{zz} = \begin{bmatrix} \hat{P}_{k-1} & 0 \\ 0 & R_k \end{bmatrix}$.
 - (b) Nonlinear motion model: $z_i = \begin{bmatrix} \hat{x}_{k-1,i} \\ w_{k,i} \end{bmatrix}, \tilde{x}_{k,i} = f(\hat{x}_{k-1,i}, v_k, w_{k,i})$.
 - (c) Recombine sigma model: $\hat{x}_k = \sum_{i=0}^{2L} \alpha_i \tilde{x}_{k,i}, \hat{P}_k = \sum_{i=0}^{2L} \alpha_i (\tilde{x}_{k,i} - \hat{x}_k)(\tilde{x}_{k,i} - \hat{x}_k)^T$.
2. Measurement Prediction:
 - (a) Extract sigma points z_i from $\mu_z = \begin{bmatrix} \hat{x}_k \\ 0 \end{bmatrix}, \Sigma_{zz} = \begin{bmatrix} \hat{P}_k & 0 \\ 0 & R_k \end{bmatrix}$.
 - (b) Nonlinear measurement model: $z_i = \begin{bmatrix} \tilde{x}_{k,i} \\ y_{k,i} \end{bmatrix}, \tilde{y}_{k,i} = g(\tilde{x}_{k,i}, n_{k,i})$.
 - (c) Recombine sigma points: $\mu_{y,k} = \sum_{i=0}^{2L} \alpha_i \tilde{y}_{k,i}, \Sigma_{yy,k} = \sum_{i=0}^{2L} \alpha_i (\tilde{y}_{k,i} - \mu_{y,k})(\tilde{y}_{k,i} - \mu_{y,k})^T, \Sigma_{xy,k} = \sum_{i=0}^{2L} \alpha_i (\tilde{x}_{k,i} - \hat{x}_k)(\tilde{y}_{k,i} - \mu_{y,k})^T$.
3. Update: $K_k = \Sigma_{xy,k} \Sigma_{yy,k}^{-1} \begin{cases} \hat{P}_k = \hat{P}_k - K_k \Sigma_{yy,k}^T \\ \hat{x}_k = \hat{x}_k + K_k (y_k - \mu_{y,k}) \end{cases}$

UKF deals with nonlinearities better since the approximation is third-order and does not need analytical Jacobians.

9 Simultaneous Localization and Mapping

Given frame-to-frame odometric measurements and measurements of landmarks in robot frames, localize the robot and compute a consistent map of landmarks in the inertial frame.

9.1 Filtering-Based SLAM

Filtering-based approaches marginalize update on past states, leading to difficulties with large loop closures.

EKF-SLAM: Form augmented state x_k by stacking robot states and landmark states. Landmarks have identity motion model with no process noise. Use EKF to estimate robot and landmark positions simultaneously. Scans poorly with number of landmarks. Can use sparse extended information filter formulation, which sparsifies P^{-1} and uses sparsity to quickly invert.

Information Filter: Alternative KF formulation using P^{-1} and $P^{-1}z$:

1. Prediction: $\begin{cases} \hat{x}_k = A_{k-1}\hat{x}_{k-1} + v_k \\ \hat{P}_k = Q_k + A_{k-1}\hat{P}_{k-1}A_{k-1}^T \end{cases}$
2. Correction: $\begin{cases} \hat{P}_k^{-1} = \hat{P}_k^{-1} + C_k^T R_k^{-1} C_k \\ \hat{P}_k^{-1} \hat{x}_k = \hat{P}_k^{-1} \hat{x}_k + C_k^T R_k^{-1} y_k \end{cases}$

Rao-Blackwellized Particle Filter/FastSLAM: Factor the belief as $p(x_{1:k}, m|y_{1:k}, u_{1:k}) = p(m|x_{1:k}, y_{1:k})p(x_{1:k}|y_{1:k}, u_{1:k})$, represent the second piece (robot states) with particle filter, marginalize out first piece (e.g. Gaussians, occupancy grid). Each particle has its own map; weights are calculated based on measurements and each particle's map. On loop closure the particles with the most consistent map (best explaining measurements) will remain while others drop out.

9.2 Batch SLAM

All poses are kept and globally optimized along with landmark positions, similar to bundle adjustment. Consider $x_k = f(x_{k-1}, u_k, w_k)$ and $y_{kl} = g(x_k, m_l, n_k)$, define errors $e_k = x_k - f(x_{k-1}, u_k, 0) \sim \mathcal{N}(0, Q)$

and $e_{kl} = y_{kl} - g(x_k, m_l, 0) \sim \mathcal{N}(0, R)$. Minimize error:

$$J = \frac{1}{2} \sum_k e_k^T Q^{-1} e_k + \frac{1}{2} \sum_{k,l} e_{kl}^T R^{-1} e_{kl}$$

$$= \frac{1}{2} \sum_k (\tilde{e}_k + \delta x_k - F_{k-1} \delta x_{k-1})^T Q^{-1} (\tilde{e}_k + \delta x_k - F_{k-1} \delta x_{k-1})$$

$$+ \frac{1}{2} \sum_{k,l} (\tilde{e}_{kl} - G_{x,kl} \delta x_k - G_{m,kl} \delta m_l)^T R^{-1} (\tilde{e}_{kl} - G_{x,kl} \delta x_k - G_{m,kl} \delta m_l)$$

$$= \frac{1}{2} (\tilde{e} + E \delta z)^T T^{-1} (\tilde{e} + E \delta z)$$

where $F = \frac{\partial f}{\partial x}, G_x = \frac{\partial g}{\partial x}, G_m = \frac{\partial g}{\partial m}, T$ has Q, R repeated along the diagonal, δz collects perturbations to robot states and landmark positions, and \tilde{e} collects all odometry and measurement errors. Solve for optimal update $E^T T^{-1} E \delta z^* = -E^T T^{-1} \tilde{e} \iff A \delta z^* = b$ and iterate until convergence. $E^T T^{-1} E$ has a sparse arrowhead structure for fast inversion.

Levenberg-Marquardt/Damped Least Squares: Solve $(E^T T^{-1} E + \mu \text{diag}(E^T T^{-1} E) \delta z^* = -E^T T^{-1} \tilde{e}$ where μ is tunable and $\text{diag}(A)$ is a matrix with the diagonal elements of A . Helps when A is noninvertible and prevents large jumps for highly nonlinear functions.

Line Search: Search in the direction of δz^* ; do $\tilde{z} \leftarrow \tilde{z} + \alpha \delta z$ where we pick the largest $\alpha \in [0, 1]$ that decreases J .

Sliding Window Filter: Marginalize out old poses and invisible landmarks beyond the sliding window, but keeps poses within the window; similar to a hybrid of batch and filtering-based SLAM, but still has issues with large loop closures.

Adaptive Relative Bundle Adjustment: Represent landmark poses relative to observation pose, and robot poses relative to the last pose to avoid a privileged global frame. Optimize only landmarks and poses within the active window. Active window size adapts to incorporate more poses and landmarks on loop closure until error is small enough.

9.3 Visual SLAM

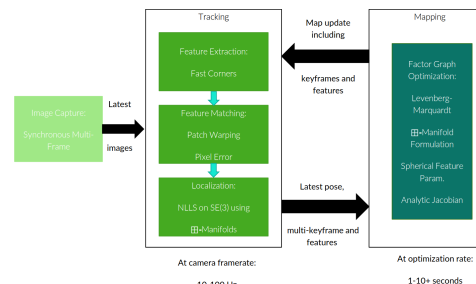
9.3.1 Feature-Based SLAM

Usually consists of a frontend for feature extraction and association (including loop closure), and a backend for map optimization.

MonoSLAM: First realtime monocular SLAM, using Shi-Tomasi features and constant velocity models. Builds probabilistic feature map through per-frame bundle adjustment.

Parallel Tracking and Mapping (PTAM): Separate threads for tracking (against existing point cloud) and mapping (adding to point cloud) into parallel threads. Coarse-to-fine feature matching and local window bundle adjustment for tracking.

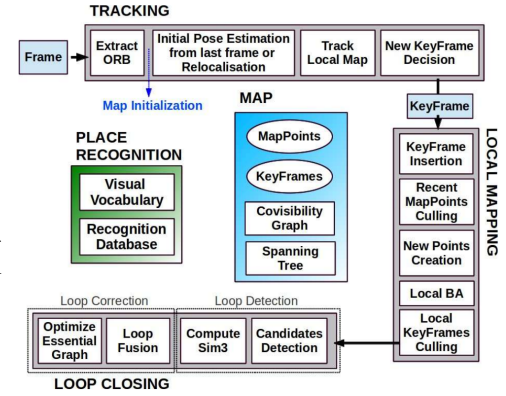
Multi-Camera PTAM (MCPTAM): Extension of PTAM to multiple cameras with non-overlapping FOVs. Spherical coordinate feature parametrization to limit depth uncertainty to a single axis. Use inverse depth to improve numerical optimization. Keyframe selection by estimating the amount of covariance reduction in features.



ORB-SLAM: Improving PTAM with loop closure ORB features. Scale-drift aware loop closure optimizes over and allows variation of the scale parameter when matching features. Covisibility graph uses viewpoints where features were observed to filter out implausible features when attempting matching. Bag of Binary Words extracts BoW descriptor from each frame for relocalization and loop closures. Parallel thread architecture with tracking, local bundle adjustment, and global loop closure/pose graph optimization, and full bundle adjustment after loop closure. Grid-based ORB feature extraction to spread out features. Aggressive adding of keyframes, call when 90% of points are visible in 3 other frames.

ORB-SLAM2: Extension of ORB-SLAM for RGB-D input, semi-dense and dense mapping.

ORB-SLAM3: Map atlases for large scale mapping.



9.3.2 Direct SLAM

Match images and produce a pose between frames based on image intensities only, without extracting features, and using (almost) every pixel to construct the map, i.e. solving pixel correspondence and camera pose estimation at once, instead of decoupling them like in sparse VO.

Comparison to Feature-Based SLAM:

1. More robust to common failure modes due to inclusion of global information, e.g. blur, self-similar textures, defocus
2. Uses almost all information in the image (as opposed to only pixels around features)
3. Dense maps enable scene understanding and path planning
4. Implicit data association by matching pixel intensities
5. Very high dimensionality (10^5 to 10^6 pixels per image), computationally expensive
6. Contains a lot more local minima (due to implicit data association), sensitive to initialization (works only for small camera movements); IMU data needed for initialization
7. No modelling of geometric noise (e.g. rolling shutter)

Image Warping: Given the location of a pixel in frame $k - 1$ and pose transformation $T_{k,k-1}, \begin{bmatrix} u_k \\ v_k \\ d_k \end{bmatrix} = f\left(T_{k,k-1} g \begin{bmatrix} u_{k-1} \\ v_{k-1} \\ d_{k-1} \end{bmatrix}\right)$.

Photometric Error: $e_i = I_{k-1}(u_{k-1}^i) - I_k^i(u_k^i)$ where I_k^i is interpolated from I_k , indexed with continuous coordinates u_k^i , and $y_k^i = \begin{bmatrix} u_k^i \\ d_k^i \end{bmatrix} = f(T_{k,k-1} g(y_{k-1}^i))$ from


image warping. Solve for optimal pose transform $\hat{T}_{k,k-1} = \arg \min_{T_{k,k-1}} \sum_{i=1}^N \left(\frac{1}{\sigma} e_i\right)^2$.

Coarse-to-Fine Optimization: Starting with a blurred and subsampled image, then repeatedly optimizing while increasing resolution, using the previous optimization result as the initial guess for the next iteration. First iteration can use IMU/odometry. This smoothes out local minima and increases region of convergence

Selective Pixel Matching: Performing matching/optimization only over "informative" pixels, based on gradient magnitude. Divide image into a grid, and for each cell take only the pixel with the largest gradient magnitude for optimization.

Photometric Calibration: Dense mapping is affected by gamma, exposure, vignetting, etc., so photometric calibration can be used to build a model correcting these effects.

LSD-SLAM: Seminal work on direct SLAM.



Direct Sparse Odometry (DSO): Integrates photometric calibration and optimizes over a sliding window of 7 keyframes, fixed number of 2000 active points. Region-adaptive gradient threshold for pixel selection. Try to match pixel along epipolar line based on motion estimate.

9.4 LiDAR SLAM

Lidar Odometry and Mapping (LOAM): Uses a keypoint extractor to look for planes and edges, while avoiding oblique planes, and edges caused by occlusion; matches only feature points. High-frequency thread performs odometry and does motion compensation, low frequency thread builds map with fine alignment.


Visual LOAM (VLOAM): Fuses IMU, camera, and LiDAR; high frequency pose propagation using IMU, medium frequency feature-based visual inertial odometry, then low-frequency corrections from LOAM.

LIO-SAM: Tightly coupled LiDAR-inertial odometry and SLAM, fusing LiDAR along with IMU, GPS, magnetometer, with a sliding window factor graph for motion correction.

LiDAR-Camera-Inertial SLAM Approaches: *LiDAR-enhanced visual SLAM:* using LiDAR to initialize depths for bundle adjustment in visual SLAM; *vision-enhanced LiDAR SLAM:* using visual SLAM output poses to initialize scan registration and pose graph optimization; *tightly coupled:* simultaneous optimization of LiDAR and camera measurements (BA or BA+PG).

9.5 Dense RGB-D SLAM

Signed Distance Fields: Represent surfaces implicitly by storing the signed distance to the nearest surface in each cell in a voxel grid. Max distance often truncated (capped). Recover surfaces from zero crossings using marching cubes. Noise naturally cancels over time when fusing observations, and surfaces can be represented at sub-voxel accuracy.



Dense Tracking and Mapping (DTAM): TSDF fusion from monocular camera. Builds highly detailed maps but requires many images and suffers with lighting changes and textureless areas. Initialization with feature-based tracking.

KinectFusion: RGB-D dense mapping and odometry using Kinect. Uses TSDF fusion and ICP. Can only map a fixed volume and relies on slow camera movements.

Kintinuous: Extension of KinectFusion to large-scale maps and adds loop closures. KinectFusion builds small submaps attached to poses which are loaded dynamically. When loop closure is detected via place recognition, the submap poses are changed, but their contents are fixed.

ElasticFusion: RGB-D SLAM using surfel maps. Each surfel has a position, normal, radius, colour and confidence. On loop closure a polygonal mesh is created from the surfel map, elastically deformed, and surfel positions are recalculated using interpolation.

Neural Radiance Field (NeRF) Methods: Implicit representation of the scene in a neural network. Sacrifices geometric accuracy but allows rendering realistic views from unseen angles. **NICE-SLAM** does SLAM using NeRFs in real time. **iSDF** uses SDFs stored implicitly in a neural network.

10 Useful Identities

$$\sin\left(x + \frac{\pi}{2}\right) = \cos(x) \quad \sin\left(x - \frac{\pi}{2}\right) = -\cos(x) \quad \sin\left(\frac{3\pi}{2} - x\right) = -\cos(x)$$

$$\cos\left(x + \frac{\pi}{2}\right) = -\sin(x) \quad \cos\left(x - \frac{\pi}{2}\right) = \sin(x) \quad \cos\left(\frac{3\pi}{2} - x\right) = \sin(x)$$

$$\sin(\hat{\theta} + \delta\theta) \approx \sin \hat{\theta} + \cos(\hat{\theta})\delta\theta \quad \cos(\hat{\theta} + \delta\theta) \approx \cos \hat{\theta} - \sin(\hat{\theta})\delta\theta$$

$$\sin(x + y) = \sin x \cos y + \cos x \sin y \quad \sin(x - y) = \sin x \cos y - \cos x \sin y$$

$$\cos(x + y) = \cos x \cos y - \sin x \sin y \quad \cos(x - y) = \cos x \cos y + \sin x \sin y$$

$$\sin(\tan^{-1} x) = \frac{x}{\sqrt{1+x^2}} \quad \cos(\tan^{-1} x) = \frac{1}{\sqrt{1+x^2}}$$

$$\sin(\cos^{-1} x) = \cos(\sin^{-1} x) = \sqrt{1-x^2}$$

$$\operatorname{atan2}(-y, x) = -\operatorname{atan2}(y, x) \quad \operatorname{atan2}(y, -x) = \pi - \operatorname{atan2}(y, x) \\ \forall a > 0, \operatorname{atan2}(ay, ax) = \operatorname{atan2}(y, x)$$

$$\operatorname{atan2}\left(x, \sqrt{1-x^2}\right) = \sin^{-1} x \quad \operatorname{atan2}\left(\pm\sqrt{1-x^2}, x\right) = \pm \cos^{-1} x$$

$$\operatorname{atan2}\left(x, -\sqrt{1-x^2}\right) = \begin{cases} \pi - \sin^{-1} x & x > 0 \\ -\pi - \sin^{-1}(x) & x < 0 \end{cases}$$

$$v_1 \cdot v_2 = \|v_1\| \|v_2\| \cos \theta \quad v_1 \times v_2 = \|v_1\| \|v_2\| \sin(\theta) \hat{n}$$

$$\mathbf{A}^{-1} = \frac{1}{\det \mathbf{A}} (\operatorname{cof} \mathbf{A})^T \quad \begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

$$\frac{d}{dx} \sin^{-1} x = \frac{1}{\sqrt{1-x^2}} \quad \frac{d}{dx} \cos^{-1} x = -\frac{1}{\sqrt{1-x^2}} \quad \frac{d}{dx} \tan^{-1} x = \frac{1}{1+x^2}$$

Common Derivatives: (Denom. layout/gradients are *column vectors*)

$$\frac{\partial}{\partial \mathbf{x}} \mathbf{A} \mathbf{x} = \mathbf{A} \quad \frac{\partial}{\partial \mathbf{x}} \mathbf{x}^T \mathbf{A} = \mathbf{A}^T \quad \frac{\partial}{\partial \mathbf{x}} \mathbf{x}^T \mathbf{b} = \frac{\partial}{\partial \mathbf{x}} \mathbf{b}^T \mathbf{x} = \mathbf{b}$$

$$\frac{\partial}{\partial \mathbf{x}} \mathbf{x}^T \mathbf{A} \mathbf{x} = (\mathbf{A} + \mathbf{A}^T) \mathbf{x} \quad \frac{\partial f}{\partial \mathbf{x}} = \left(\frac{\partial \mathbf{u}}{\partial \mathbf{x}} \right)^T \frac{\partial f}{\partial \mathbf{u}} \quad \frac{\partial \|\mathbf{x}\|}{\partial \mathbf{x}} = \frac{1}{\|\mathbf{x}\|} \mathbf{x}^T$$