Localization

RSS Lecture

Monday, March 9, 2009

Prof. Teller

Text: Siegwart and Nourbakhsh Ch. 5
Dudek and Jenkin Ch. 7

Navigation Overview

- Where am I?
  - Localization (Today)
  - Assumes accurate map, but imperfect sensing
- Where have I been?
  - Mapping (Wednesday)
  - Assumes effective localization
- How can I get there from here?
  - Planning (Next M & W)
  - Assumes perfect map, sensing, and actuation

Thought experiment

- Does it make sense to localize in a void (an environment containing absolutely nothing)?

  … not very interesting; We conclude that there has to be some kind of "stuff" in environment

- What if the environment is *isotropic* (space, fog, water, desert, jungle etc.)?

  … again, not very interesting for robot to move or perform tasks within such an environment

  We conclude that environment must contain *features* that can be sensed (distinguished) by bot

Localization Problem Statement

- Given some representation of the environment, to *localize*, robot must, through sensing, determine its pose *with respect to the specified representation*
- Defined with respect to some frame or feature set that is *external* to robot:
  - Global coordinate frame
    - E.g., GPS (Earth) coordinates
  - Local coordinate frame
    - Ceiling or floor tiles
    - Mission starting pose
  - Environment features
    - E.g., nearby walls, corners, markings

Basic Localization

- Open-loop pose estimation:
  - Maintain pose estimate based on expected results of motion commands (no sensing)
- Dead reckoning:
  - Use proprioception (odometry, inertial) to estimate pose w.r.t. *initial* coordinate frame
  - Multiple error sources:
    - Wheel slip, gear backlash
    - Noise (e.g. from encoders)
    - Sensor, processor quantization errors
  - Pose error accumulates with time and motion
  - Typically ~ a few percent of distance traveled

Dead Reckoning Error

- Two hours of slow, rolling motion through MIT main campus corridors at third-floor level
  - Bosse, Leonard, Newman, Teller (IJRR 2004)
- High-precision inertial sensors exist… do they solve problem?

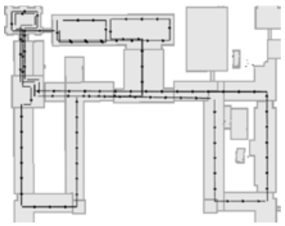

True path topology
(manually overlaid)

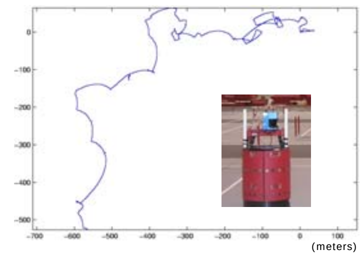

Integrated odometry
(Nomadics B21)

Landmark Attributes

- Is landmark *passive* or *active*?
  - Must sensor emit energy to sense landmark?
- Is landmark *natural* or *artificial*?
  - If placed in env't, how are locations chosen?
- Which sensor(s) can detect it?
  - Vision, sonar, radio, tactile, chemical, …
- What are landmark's geometric properties?
  - Plane, line, segment, point, diffuse source, …
- What is *discriminability* of landmark?
  - (Will discuss this in detail in a minute)

Landmark Types

	Passive	Active
Natural	Wall corner Texture patch River bend Earth's surface	Sun, North star Magnetic dipole Pressure gradient Mineral vent
Artificial	Surveyor's mark Retro-reflector Lighthouse (day) Trail blaze Buoy, channel marker	Chemical marker Radio beacon Lighthouse (night) LORAN GPS

Robot Landmark/Sensor Types

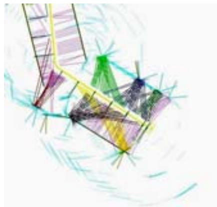

- Range to surface patch, corner
  - Sonar return
- Bearing (absolute, relative, differential)
  - Compass; vision (calibrated camera)
- Range to point
  - RSS, TOF from RF/acoustic beacon
  - Cricket (TDoA acoustic, RF)
- Range and (body-relative) bearing to object
  - Radar return
  - Laser range scanner return
  - Vision (stereo camera rig)
- Distance to sea surface, floor
  - Pressure (depth), bathymetry (depth, altitude)

Discriminability Challenges

- Landmark *Detection*
  - Is landmark distinguishable from *background*?
- Landmark *Measurement, Data Fusion*
  - Sensor gives a noisy, quantized measurement of landmark geometry (bearing and/or range)
  - How accurately can one measurement localize landmark?
  - How can multiple corrupted measurements be combined into one accurate landmark estimate?
- Landmark *Identification*
  - To which element of *representation* (i.e., map) does the detected and measured landmark correspond?
  - To which *previously-observed landmark* (if any) does currently observed landmark correspond?
  - Also known as the "data association" or "feature correspondence" or "matching" problem

Localization Degrees of Freedom

- Model robot/vehicle as a single rigid body
- Aerial, orbital, underwater navigation
  - 6 DOFs: three position + three orientation
- Terrestrial operation (rolling, walking)
  - 3 DOFs: two position + one orientation
  - Used for planar, mildly non-planar terrain
- Underwater surveying (high C. O. B.)
  - 4 DOFs: three position + one orientation

WHOI AUV, Hanu Singh (Aug. 2004)

Localization Examples

- Two dimensions
  - Ideal sensors
  - From measured *ranges* (distances)
  - From measured *bearings* (directions)
- One dimension
  - Real sensor (noisy measurements)
  - From range and odometry
  - Filtering, outlier rejection
- Two dimensions
  - Mobility with RF/acoustic beacons

Triangulation

- Natural geometry for 2D localization
  - Simplest framework combining range, bearing
  - Used by Egyptians, Romans for engineering

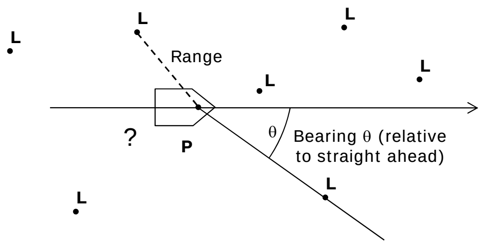

  - Features for which only some geometric attributes can be sensed directly are called "partially observable" features

Triangulation from ranges

- Robot at unknown position **P** measures distances d_1, d_2 to *known* landmarks $\mathbf{L}_1$, $\mathbf{L}_2$
- What are possible solutions for **P**?

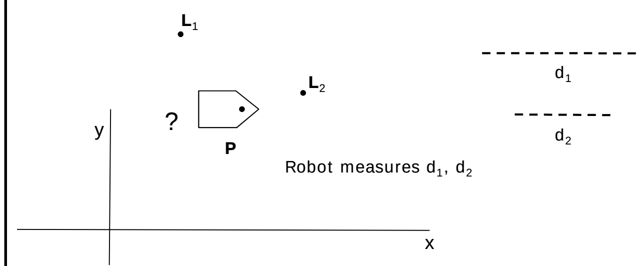

Triangulation from ranges

- Robot must lie on circles of radius d_1, d_2 centered at $\mathbf{L}_1$, $\mathbf{L}_2$ respectively

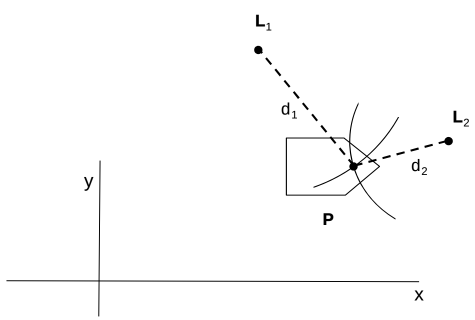

Triangulation from ranges

- Two solutions in general, P and P'
- How to select the correct solution?

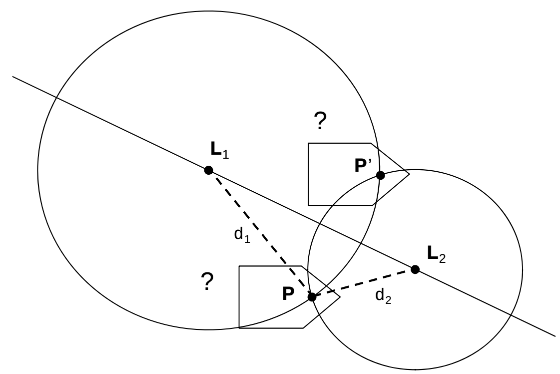

Disambiguating solutions

- *A priori* information (richer map)

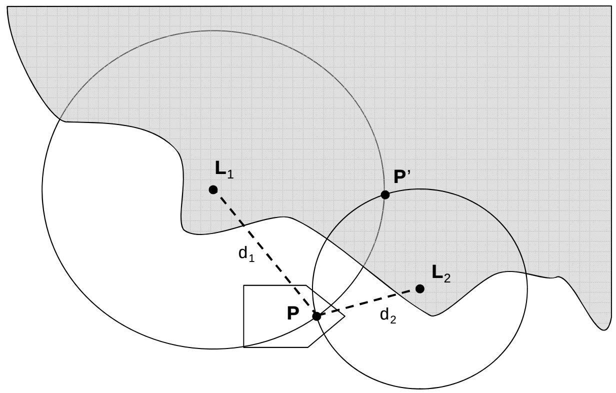

Disambiguating solutions

- Continuity (i.e., spatiotemporal information)

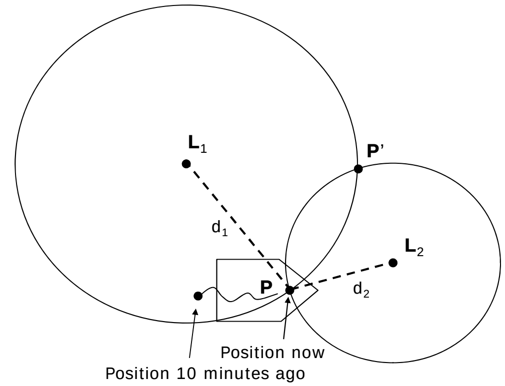

Disambiguating solutions

- Additional landmarks (redundancy)

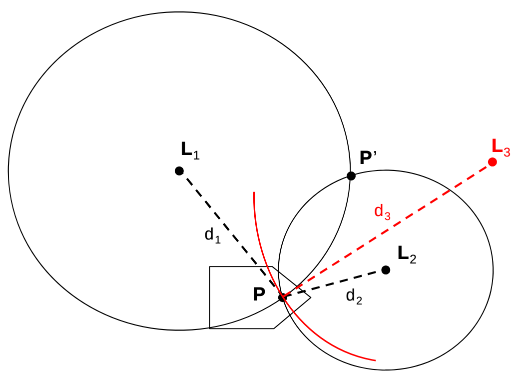

Triangulation from ranges

- Are we done yet, i.e., is pose fully determined?
- No: absolute heading is *not determined*

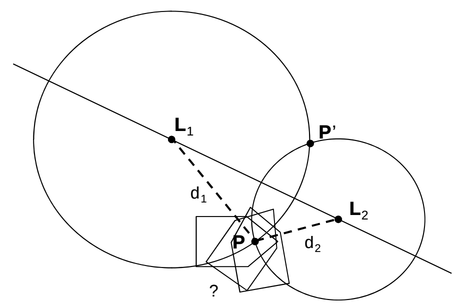

- How to get heading?
  - Motion (difference of positions inferred across *time*)
  - Extent (using two ranges measured over ship *baseline*)

Triangulation from bearings

- *Body-relative* bearings to two landmarks
  - Bearings measured relative to "straight ahead"

Robot observes:
L_1 at bearing θ_1
L_2 at bearing θ_2

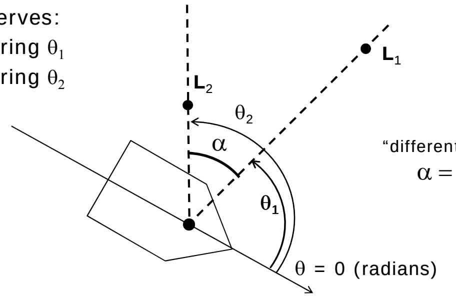

"differential bearing"
$\alpha = \theta_2 - \theta_1$

... are two bearings enough for unique localization?

Triangulation from two bearings

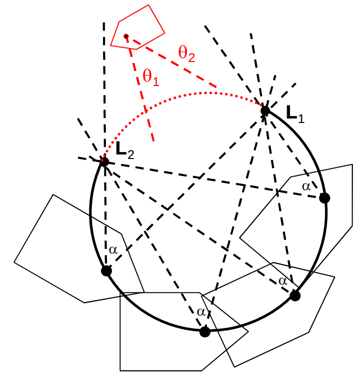

- Robot somewhere on circular arc shown
  - Can it be *anywhere* on circle? (No; ordering constraint)

Triangulation from bearings

- Measure bearing to third landmark
  - Yields robot position *and* orientation
  - Also called robot *pose* (in this case, 3 DoFs)

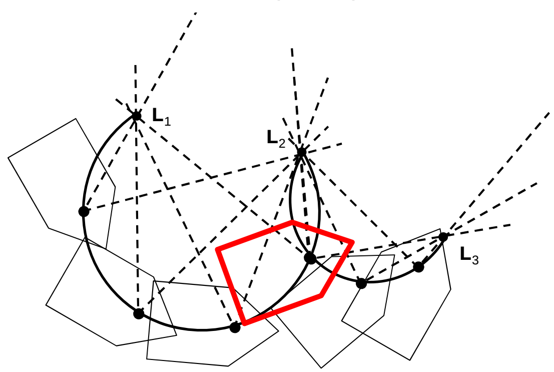

Measurement Uncertainty

- Ranges, bearings are typically *imprecise*
- Range case (estimated ranges $\sim d_1$, $\sim d_2$)

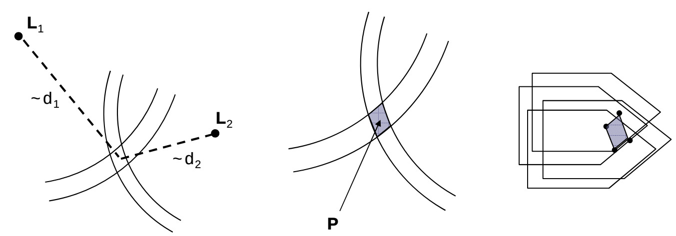

Locus of likely **positions**

Measurement Uncertainty

- Two-bearing case (estimated bearings $\sim\theta_1$, $\sim\theta_2$)
- What is *locus* of recovered vehicle poses?
- Solve in closed form? Is there an alternative?

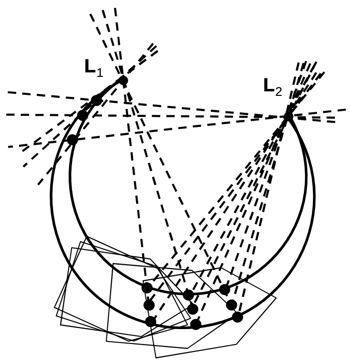

Measurement Uncertainty

- Bearing case (measurements $\sim\theta_1$, $\sim\theta_2$, $\sim\theta_3$)

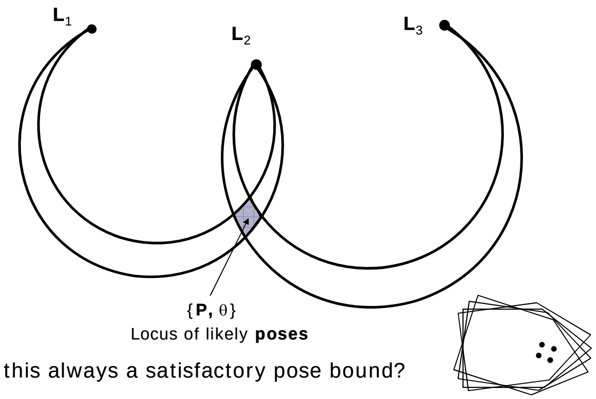

- ... is this always a satisfactory pose bound?

Landmark, sensor geometry

- Consider off-axis and near-axis bearing measurements to two known landmarks (simplification: assume absolute heading is known)

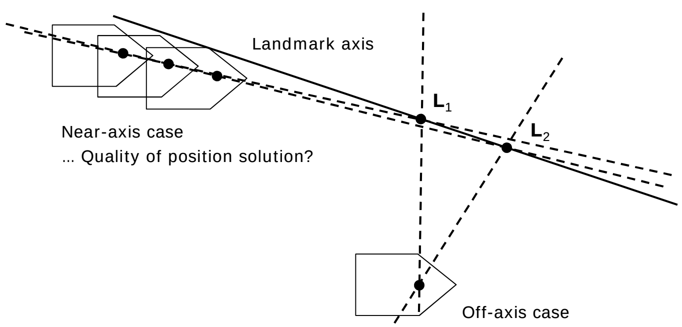

Dilution of Precision

- General phenomenon that sensor, landmark, and motion geometry can *degrade* solution quality, even for a *fixed set* of observed landmarks
- Geometric DOP = GDOP
  - Also Vertical DOP, Horizontal DOP etc.
- How to take GDOP into account?
  - If sufficiently many landmarks are available, *select* those with minimal GDOP
  - Decouple pose, solve separately, recombine

Sensing uncertainty

- Time series of round-trip-time to one acoustic beacon for an underwater autonomous vehicle

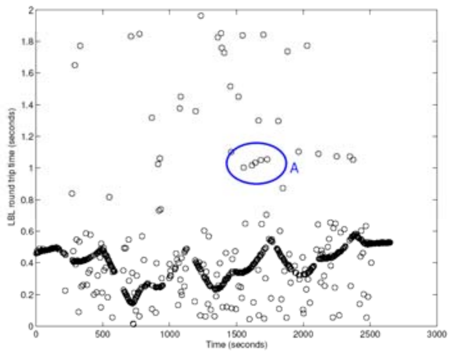

(Olson, Leonard, Teller, Robust Range-Only Beacon Localization, Proc. IEEE AUV, June 2004)

Gaussian noise model

- Measurement returns a corrupted value

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

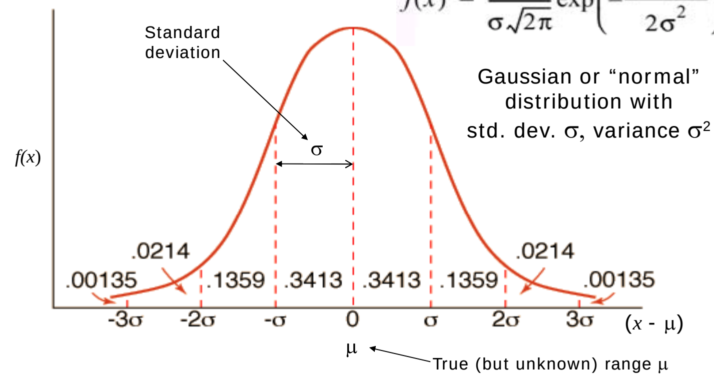

Outliers

- Many measurements are *outliers*; their frequency is not well-modeled by a Gaussian distribution

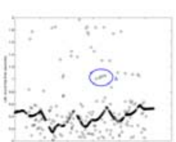

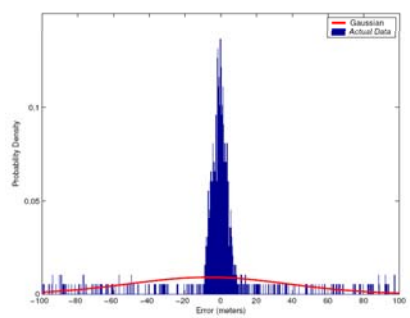

... what to do?

Filtering

- Consider one-dimensional localization:
  - Robot measures range r(i) at i^{th} time step
  - Ranges *corrupted* by Gaussian noise, outliers

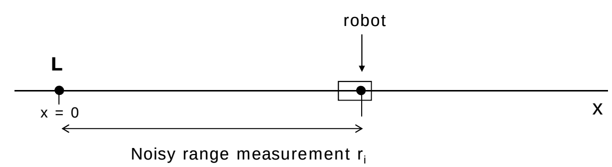

- *Filter* measurements; combine over time
  - Deal with each measurement as it arrives
  - *Recursive* or *on-line* filtering (contrast *batch*)

Filtering with no outliers

- Suppose robot does not move
  - What should our filtering strategy be?
  - Call x(t) our *estimate* of x after t time steps

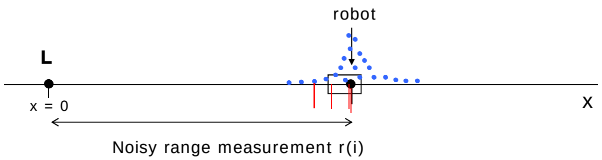

- Take mean (arithmetic average)
  - x(i) = (r(1) + r(2) + ... + r(i)) / i (batch)
  - x(i) = [x(i-1) * (i-1) / i] + [r(i) / i] (recursive / on-line)
  - ... if no outliers, no change over time, filter is optimal
- Computational complexity of each update?

Dealing with outliers

- Suppose a fraction of r(i) are wildly wrong
  - Classify r(i) as *inliers* or *outliers*
  - How to do this?

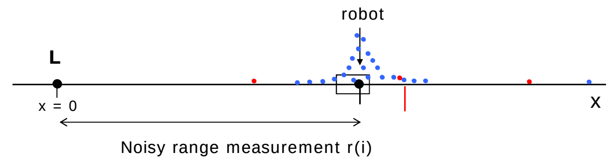

Modeling measurement noise

- Estimate sample *variance* as well as mean

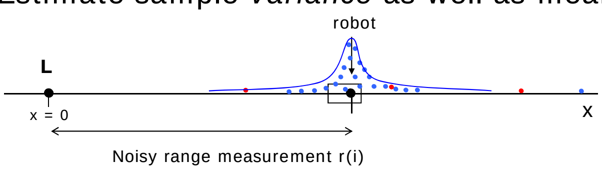

- Reject unlikely samples (e.g., $p < 1\%$)
  - Filter only inliers, by averaging as before
- ... But where does variance come from?
  - Determine it *a priori* (e.g. from bench tests)
  - Or, estimate it *on-line*, in addition to mean
    - Chicken-and-egg problem (could be unlucky)
    - If "outliers" become frequent, what can you do?

Estimating variance

- Define $\sigma^2(i)$ as variance after i steps
- Batch computation:
  - As before, x(i) is the mean after i steps
  - Then variance $\sigma^2(i)$ is $[\sum(r(i)-x(i))^2] / i$
- Recursive (on-line) computation:
  - Estimate x(i) recursively as before
  - Define $\sigma^2(1) = 0$; then for $i > 1$:

$$\sigma^2(i) = (i-1)/i * \sigma^2(i-1) + 1/(i-1) * (r(i) - x(i))^2$$

Robustness, Validation

- Additional measurements can be used
- Increase robustness to noise:
  - Average measurements as shown earlier
  - Require more than minimum # of landmarks
  - Drawbacks? Takes more time, or restricts space in which method works. These are fundamental tradeoffs in localization
- Enable validation w.r.t. gross error:
  - Decompose into subsets, solve independently; compare solutions
  - Predict additional landmarks from observed

Localization challenges

- Partial observability
- Measurement noise
  - Amplified by GDOP
- Outlier measurements
- ... Are those all we have to worry about?

Light at bearing b_1,
Light at bearing b_2,
Light at bearing b_3

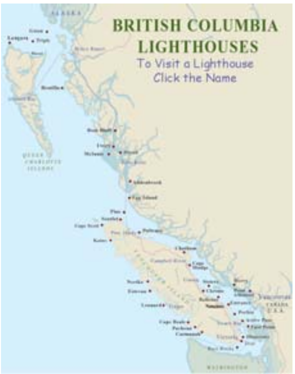

Data association problem

- General problem: determining how an observation corresponds to a map feature, or to a previously observed feature (also called *correspondence problem*)
- How to tackle?
  - Initialization and continuity
  - Identify distinguishing features among landmarks
  - Combinatorial testing / cross-validation
    - RANSAC, Random Sampling and Consensus, 1981

Localization: Summary

- Localization: from a map and its sensors, robot must determine its pose with respect to map
- Challenging problem in general, due to:
  - Partial observability
  - Data association
  - Noise & GDOP
  - Outliers
- Strategies for robust localization
  - Geometric decoupling
  - Landmark selection
  - Initialization, continuity, combinatorial search
  - Filtering
  - On-line variance estimation, outlier rejection