Tracking System

Paul Dubrulle - MRG

Where the Tracking System Sits

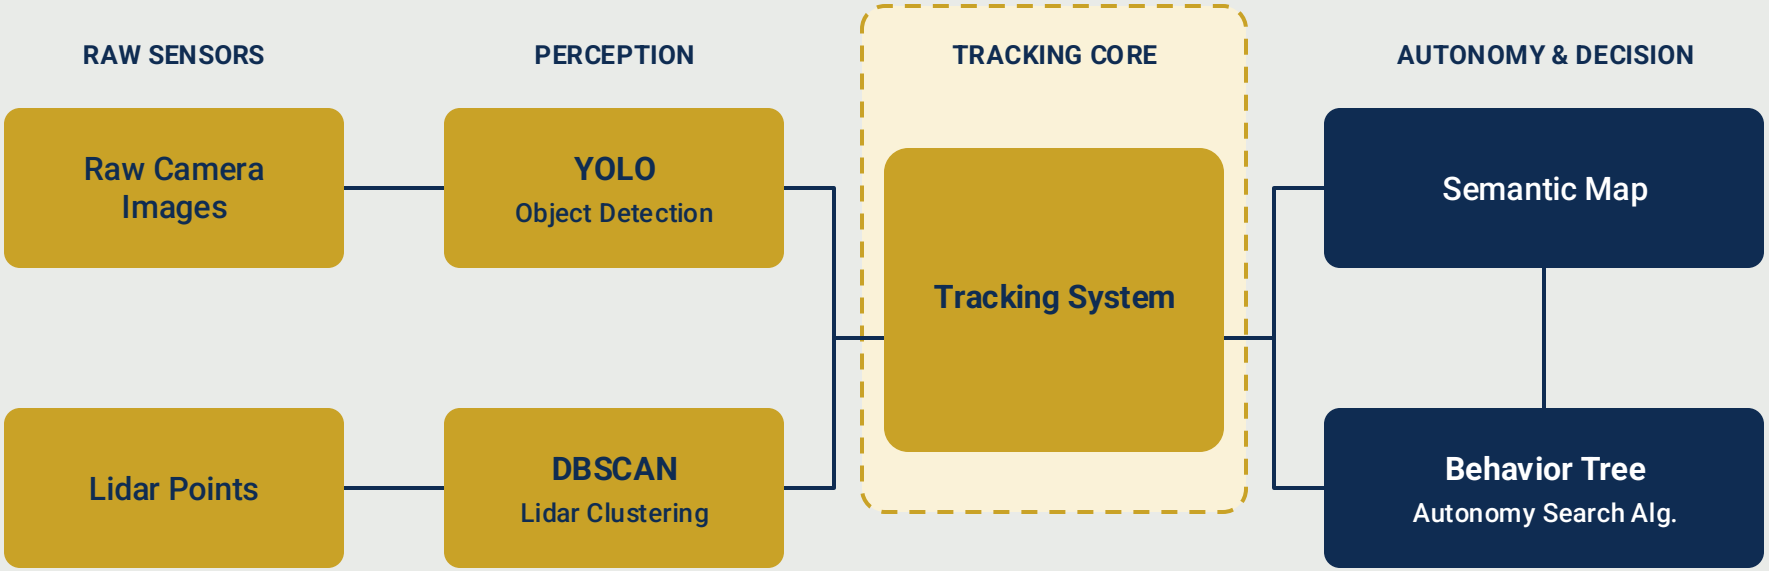

The tracker's job: turn a stream of one-off detections into a small set of persistent objects, each with a position, a color, and a confidence that it is real.

What is a Track?

A Track is a probabilistic estimate of an object such as a buoy or floating debris. Every track carries four attributes:

01. Position

Where the object is: the center point of a LiDAR cluster.

Initial State

The first detection's location.

Uncertainty

Inherited from how spread-out that first cluster was.

02. Velocity

Assumed zero, because buoys are anchored, so apparent drift is treated as noise rather than motion.

Initial State

Not moving (0 m/s).

Prior Uncertainty

Held nearly certain, so the tracker never invents motion.

03. Existence

How sure we are that the object is real. Sightings raise it, unexplained misses lower it.

Initial State

Low (15%)

A single detection could be noise, so trust has to be earned.

04. Class Belief

What the object is. Camera detections vote for its color over time.

Classes Tracked

Red and Green Buoys.

Initial State

Undecided: 50% red / 50% green.

Which Measurement Belongs to Which Track?

Our answer is Probabilistic Data Association (PDA): never pick one measurement for a track, score every nearby candidate instead and update the track by all of them in proportion to how likely each one is. The next two slides grow PDA into the full method the boat runs, JIPDA.

1. **The question every cycle:** something has to decide which measurement updates which track, whether that measurement is a LiDAR cluster centroid or a camera bearing.

2. **Why picking the closest is dangerous:** the nearest measurement is usually the right one and occasionally it is a false detection. One wrong pick drags the track onto that false detection.

3. **How PDA answers it:** PDA measures how well each nearby measurement fits the track's prediction, then turns those fits into probabilities that add up to 100%: candidate A 70%, candidate B 20%, candidate C 5%, object missed this cycle 5%.

4. **That 70% is called β:** β is the probability that one specific measurement came from one specific track. The track moves 70% of the way toward A and only 20% toward B, and its uncertainty grows because the answer was not certain.

Two Tracks, One Measurement (the J in JIPDA)

PDA on its own has a blind spot. Two buoys sit close together, a single measurement lands between them, and each track scores it without knowing the other track exists. Both claim it, the same evidence is spent twice, and the two tracks slide into each other.

1. **Why running PDA twice fails:** each track only asks whether the measurement suits itself. Neither one can notice that its neighbor is claiming the very same return, so both claim it in full and the two tracks drift together until they merge.

2. **What the J adds, joint scoring:** JIPDA looks at the whole scene at once and lists every way the measurements could be assigned, this one is A's, or it is B's, or it is noise. No single assignment gives one measurement to two tracks, and each is scored by how well it fits both tracks together.

3. **How that becomes β:** add up the assignments where the measurement belongs to A and you get A's β, then do the same for B. The two now share one budget of 100%: A 70%, B 25%, noise 5%. Neither track can absorb it fully, so nearby buoys stay distinct instead of collapsing into one.

Is the Track Even Real? (the I in JIPDA)

Even joint scoring still assumes every track is a real object. On water that is false, because a single false detection can start a track that nothing real supports.

1. **Existence probability r:** each track carries r, estimated inside the same matching math rather than by a separate rule added afterward.

2. **r shapes the assignments:** a doubtful track makes a weak claim on a cluster, so an established track wins a contested measurement over a newborn ghost.

3. **The assignments update r:** a match raises r, and being visible but unmatched lowers it. Ghost tracks fade on their own, while real buoys climb toward confirmation.

JIPDA = PDA + joint scene assignments + integrated existence

Camera Detections: the Same JIPDA in Angle Space

A camera detection is what gives a track its color, and it runs through the identical machinery. The only change is the space the comparison happens in: position space for LiDAR, angle space for the camera.

1. **There is no ray intersection test:** the detection is reduced to one number, the map-frame bearing from the camera to the box center. Each track answers with its own predicted bearing, and the surprise is simply the gap between the two angles.

2. **The tolerance is a wedge, not a circle:** the gate combines the camera's bearing noise, roughly 2 degrees, with the track's position uncertainty seen from the camera. The same 0.3 m subtends a wide angle at 3 m and a sliver at 30 m, so near tracks get generous gates and far ones tight gates.

3. **Ambiguity is broken by visibility:** if one bearing gates into two tracks, JIPDA scores the same assignments, weighted by existence, detection probability and angular likelihood. The nearer track occludes the farther one, whose detection probability falls to nearly zero, so the near track wins.

4. **What β delivers here:** color evidence is deposited into every gated track in proportion to β, position is nudged sideways only, because a bearing carries no depth, and existence rises on a match. A bearing that gates into nothing is dropped, since bearings never birth tracks.

Putting It Together: One Update Cycle

01. Predict

The motion model carries position and velocity forward to the measurement's time. Existence and color are not predicted. Then compare:

surprise = newest measurement − predicted position

Velocity is held at zero, so the position stays put and only its uncertainty grows while unseen.

02. Position Update

new position = predicted + K × surprise

K = track uncertainty ÷ (track uncertainty + measurement uncertainty): a confident track barely moves for a noisy measurement; a vague new track leaps toward a precise one.

03. Existence Update

Sightings raise existence r. Misses lower it, but only if the object was visible:

on a miss: r ← r × (1 − P_see) ÷ (1 − r × P_see)

P_see = chance the sensor could see it (in view, in range, not blocked by a closer object). Not visible → P_see ≈ 0 → r unchanged.

04. Class Belief Update

Each matched detection votes for its color, weighted by the match probability β from JIPDA:

red tally += β × s^1.5

β = match probability, s = detection confidence (0 to 1). Belief = red tally ÷ total tally, and old votes slowly fade so early mistakes wash out.

A Track's Life Over Time

01. Birth Criterion

An unclaimed LiDAR cluster becomes a new track: position adopted from the cluster, existence starting low, color undecided.

Unclaimed means:
Every story agrees it came from no existing track, so it must be something new.

02. Two Ways to Die

Tracks are removed when the evidence stops supporting them.

Existence collapses: too many “should have seen it” misses.

Silence: six seconds pass with no matched measurement.

03. Promotion to the Permanent Map

With repeated sightings existence climbs, and at around 90% the track is CONFIRMED. A small buffer below that keeps it from flickering in and out.

Once a confirmed track is also confidently one color (about 75%), it is written into the semantic map, the boat's permanent memory, which outlives the track itself.

Recap: One Cycle, End to End

Every cycle repeats the same four moves, and a buoy on the map is simply a track that survived them often enough.

1. **Detect:** LiDAR clusters give positions, the camera detector gives colors and bearings. Every measurement arrives with no idea which object it belongs to.

2. **Match:** JIPDA scores every way the scene could be assigned, then hands each track a match probability β for each nearby measurement, including the possibility that it is only noise.

3. **Update:** position, existence r, and color belief each move in proportion to β, so strong matches move a track a lot and weak ones barely nudge it.

4. **Commit:** tracks that go unexplained lose existence and die. Tracks that reach about 90% existence and 75% color confidence are written into the semantic map the boat steers by.