

Viewpoint - Localisation

Every so often, a new term has to be coined for an old concept. I have the impression that the French are to blame for the introduction of the term 'localisation'. When I designed autopilots over forty years ago, an assortment of gyroscopes and accelerometers plus compass signals gave information the aircraft's attitude, while altimeter and a variety of radio navigation aids such as VOR, ILS Localiser and radar altimeter fixed its position in three dimensions.

Now, at mobile-robot scale, we have sonar and scanning rangefinders to add to odometry, plus GPS in a variety of precisions and as many other beacons as we care to provide. Machine vision provides a rich tapestry for further sensing, with triangulation and visual streaming among its many techniques. The core task remains the same, to combine the signals to give a best estimate of where we are and how to go where we want.

The game becomes more interesting when we are in unknown surroundings. SLAM, simultaneous localisation and mapping, concentrates on our relationship to what we can see and remember around us. As humans, we are quite comfortable to remain in total ignorance of our global coordinates, as long as we can detect and avoid the nearest obstacle while heading in the right direction to find our parked car.

Combining the signals can be performed in countless ways. GPS will only give signals at fixed intervals, though these can be as short as a fraction of a second. An inertial package is therefore an ideal candidate to interpolate between readings if our movement involves high-speed dynamics. Signals can be blended together with some form of linear processing such as a Kalman filter, or something more adventurous can be tried.

Each sensor has its individual way of introducing uncertainty and noise. Sensors on the wheels can theoretically be integrated to give position and heading, but the slightest slip will see the position estimate drifting away in a wrong direction. A GPS reading can leap aside at the whim of a multipath error or the setting of a satellite. Code-phase pseudorange measurements suffer from conventional noise and carrier phase ranges suffer from ambiguity. Sensor signals can vote on 'alternate realities', calculating a set of distinct positions with various probabilities of truth.

When paired off, sensors can lean on each other to compensate for their shortcomings. Odometry can be propped up by a compass signal, though some occasional 'absolute fix' is needed to remove the integrated position error. On its own, a video camera can do little more than recognise objects and estimate the angle between them. But put the camera into motion, with displacement measured by odometry, and most of the benefits of stereo vision become available without the drudgery of camera calibration. Subsets of sensors can be combined to check for consistency and probable accuracy, or to vote out a faulty reading.

It is refreshing when a novel research paper is found, possibly combining signals in a new nonlinear way or maybe finding a measurement technique that is novel in itself.

So often, however, papers are mere 'variations on a theme', perhaps following a 'fuzzy' or 'neural' paradigm that is a disguised echo of methods half a century old.

Pedestrian papers can cling onto a 'traditional favourite' device, so that the ORP12 photoresistor and 741 operational amplifier still survive their much-deserved demise. Robots have long been ringed with ultrasonic rangefinders, originally justified because they were a cheap spin-off from a camera focussing system. A clunky optical time-of-flight rangefinder has enjoyed a more recent popularity. Designed for intruder detection, its data output seems to have been intended merely for diagnostic purposes. Something much better could be designed to present signals with the bandwidth of a vision interface.

In all its guises and by any other name, localisation will never cease to be a 'hot topic'.