Chapter 3

Likelihood Test

The Likelihood Test measures the overall distance between the laser scan and the map representation, assuming the system is located in a given posture. Thus, it is a measure of the error associated to the postures evaluated by the Localisation algorithms. Using the appropriate thresholds, it is also used to validate the candidate postures.

For each sample in the laser scan, the distance to the corresponding point in the map is computed. The process is iterated for all samples and the distribution of the point to point distance is an upper bound of the Localisation distribution error.

Since there is no external reference to benchmark the Localisation results, the error must be measured in terms of the data available to the Localisation algorithms. This solution for performance evaluation embeds several types of error in a common error distribution, which is assigned to the Localisation algorithms. If the map is absolutely accurate and the error induced by the target surfaces is neglected as well as the mechanical errors and inaccuracies associated to the sensor hardware, the error is only due to the laser sensor and the optimal Localisation error distribution is equal to the laser error distribution. In real experiments, though, the other sources of error are relevant and the overall performance for Localisation degrades beyond what would be expected from the algorithm analysis.

The error estimate associated to a posture (x, y, θ) is defined in the form of a distribution of distances, without any possibility of de-coupling the error into the three coordinates.

The Likelihood Test proved to be very sensitive to posture errors: given two postures only two or three centimetres apart, there is a dramatic difference in the error distributions, making clear which one is closest to the optimal posture estimate.

Chapter Organisation

The outline of the method is presented in Section One. The map processing from its original form to a cloud of points, termed "simulated scan", is introduced in Section Two, and the methods for generating the simulated scan from map data are discussed. Section Three presents the formal development of the Likelihood Test. Some experimental results based on Frame Localisation estimates are presented in Section Four, which also serves the purpose of assessing the accuracy of Frame Localisation. Section Five presents the conclusions.

3.1 Algorithm Outline

3.1.1 Foreword

The Frame Localisation algorithm described in Chapter 2 and one of the Approximate Localisation algorithms, described later in Chapter 4 are based on matching a reduced set of elements in the laser scan to the map description. However, it could occur that a limited set of features match the map description and yet the posture is wrong.

In order to assess the accuracy and validate a candidate posture, the whole laser scan must fit the map description with a minimum error. Notwithstanding, it is possible to have a good match between the whole scan and the map at a wrong posture. This usually indicates a symmetric environment description, insufficient map data or poor laser data.

The only ways to disambiguate between different valid postures is to use external data, such as odometry, or to keep track of the system displacements and eliminate the impossible path sequences. Assuming it is a rare occurrence, this exception is not handled. If there are multiple postures validated by the Likelihood Test, they are output to the user in a log file, along with the associated error distribution. In case of an automatic implementation, the posture with the minimum expected value for the error distribution is chosen as the correct one.

3.1.2 Comparing map data to laser data

Given a candidate posture $p_i = (x_i, y_i, \theta_i)$, defined in map coordinates, the elements in the map may be perceived from p_i as a horizontal radial profile centred at p_i. If the physical constraints of the acquisition system are disregarded, this method corresponds to computing the radial field of view from p_i (see Figure 1a). The field of view assumes a 360º free horizon, *i.e.*, the laser scanner is not constrained by the structure of the RESOLV system.

The map features are the *visible* fragments of the map elements (Figure 1b); they define the map-based profile, expressed in radial coordinates, *i.e.*, the same coordinates as the laser scan. If p_i is a close estimate of the system posture, the laser scan should match the map profile in large extents.

In most cases though, some parts of the two profiles will not match. If the laser scan could not measure on a given direction or the target surfaces are beyond the operating range there will be features in the map profile with no counterpart in the laser profile. If the map is not yet built on that direction there will be missing features in the map profile - it is likely that the RESOLV system is moving towards the unknown area to extend an incomplete model. Moreover, if an element in the scene could not be reconstructed from the acquired data, some features in the laser profile will have no correspondence in the

map profile. The constraints on laser based reconstruction are mentioned in Chapter 1, Sections 3 and 5.

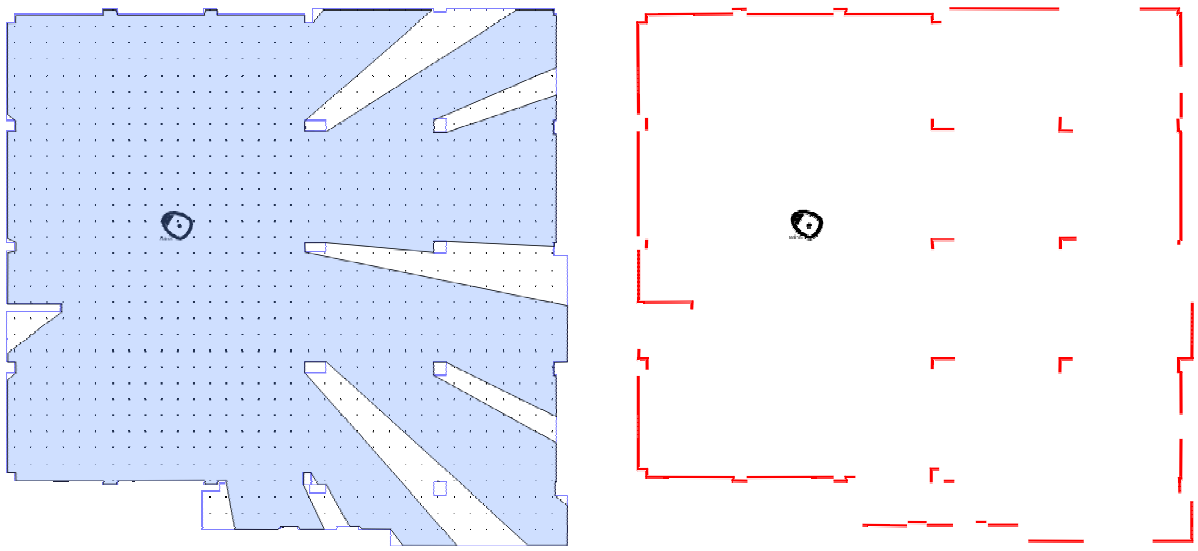

(a) (b)

Figure 1 – The 360º laser field of view (a) and the perceived map-based profile (b)

The algorithm could be built in a symmetrical form, i.e., the laser data could be transposed to the world reference and then a coherent fit would be searched. This option was disregarded for the following reasons:

- Finding a correspondence between the laser and map data is a search and match problem. While the search in the laser scan space has a known fixed dimension (the number of scan points), the search in the map space has a variable growing space, as the number of map features grows during the reconstruction process.
- The map description is not sorted in a useful manner for detecting which are the features close to the expected posture. Thus, selecting the features used for matching in map coordinates would be a time-consuming process. On the other hand, referring the map features to the laser coordinates is a simple process, described in the next section.
- The map information may be expressed in different forms, either lines that are used for Frame Localisation or past range scans, which will be described in Chapter 4. Expressing the laser scan data in the map data formats may require line extraction, reducing the ability to detect errors in the Frame Localisation.
- Since the map data comes from various sources, it must be translated to a common form, in a common coordinate reference. Choosing the laser scanner coordinates avoids one transform, because the laser data is already defined in that reference.

- There are areas in the scan with no correspondence in the map. If one were to look for a map feature to fit those areas, an extensive search of all features would be required, only to yield a void result.

The generation of map profiles, similar to the laser scans, builds on the following hypothesis: any given feature on the map is located at a known distance and bearing of the laser, assuming the laser sensor is located at the candidate posture p_i.

In Figure 2, the robot is represented surrounded by a shaded area denoting its field of view as illustrated in Figure 1a.. The laser sensor and the map feature define a "triangle of visibility", i.e, the laser circular sweep movement will find the map feature between the bearings defined by its end vertexes (Figure 2a).

If no other feature is closer to the laser within the bearing interval, the field of view on those directions will range from the laser to the feature considered (Figure 2a). In case another surface is closer to the laser within the bearing interval, the feature under analysis is occluded, therefore it is not considered when building the range profile (Figure 2b). In mix cases, the range profile is filled to the extent where the considered feature is closer than other features (Figure 2c), hence within the robot's field of view.

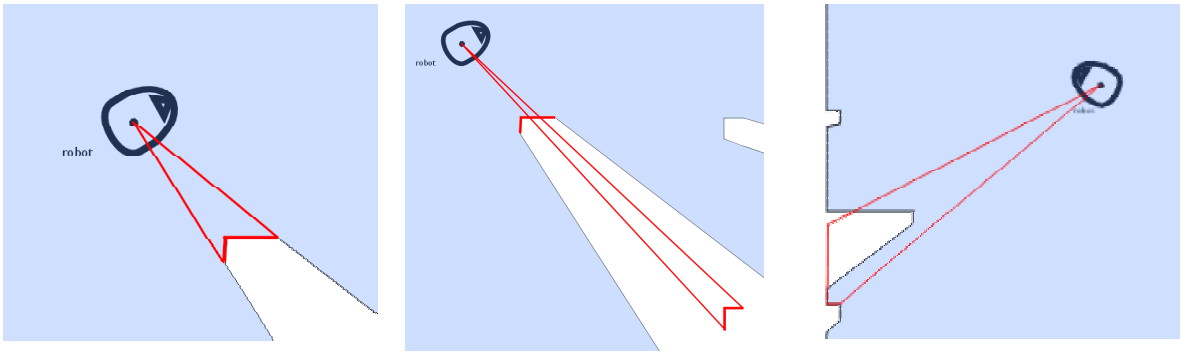

(a) (b) (c)

Figure 2 – Building a map profile around the laser scanner

Once the map data is represented as range profile, similar to the laser scans, the two profiles can be compared.

3.1.3 Point-to-point distance distribution

The map profile represents a theoretical laser scan measured from the candidate posture, p_i. The map profile measures the distance from the origin (where the sensor is located, or it is assumed the sensor is located) to the closest obstacle in each direction. The parameters of the map profile are set to equal the laser scan parameters. Hence, the two profiles cover the same bearing angles with equal angular resolution.

However, the map profile doesn't account for many of the actual characteristics of

the laser device, namely sizeable footprint, range errors and angle of incidence influences. Moreover, it doesn't account for the target surfaces' characteristics such as colour, reflectivity and texture. It would be possible to include these errors in the map profile, based on an experimental model of the laser characteristics. On the other hand, the errors related to the target surfaces are very difficult to model accurately.

Beyond these errors, differences arise when the map profile and/or the map profile have gaps or unexpected features. If the map is incomplete or the laser profile includes a fraction of the robot casing, there will be missing features respectively in the map profile and the laser profile. If there are elements in the map profile which were removed between the generation of the 3D reconstruction model and the laser scan, the extra features will appear with no correspondence in the laser profile.

If the two profiles are compared point-to-point (Figure 3) and the differences are accumulated, the point to point distance defines an error distribution that measures the difference between the two profiles.

For the sake of clarity in Figure 3, the point-to-point distances and the angle shift between adjacent laser samples are much exaggerated.

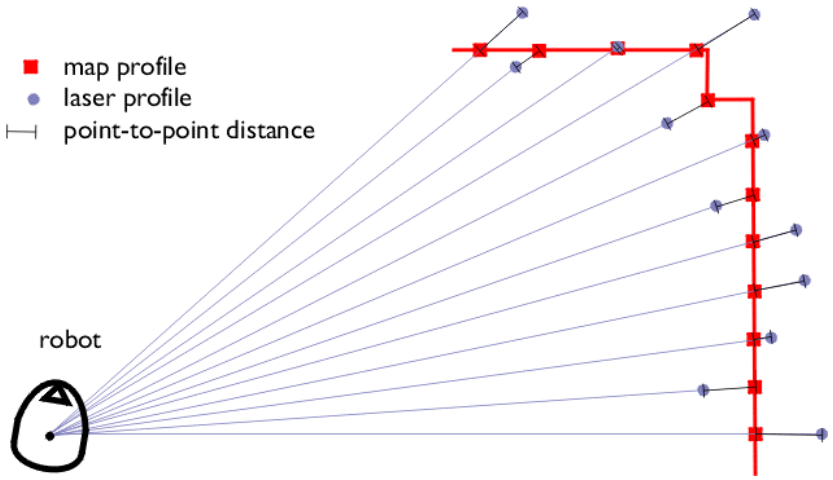

Figure 3 – Point-to-point distance between the laser and map profiles

The point-to-point distances are accumulated in a histogram. Representing the continuous distribution between zero and infinity as a discrete histogram provides a clear insight on the error distribution. The histogram is divided in N slots. The first N-1 slots are equally spaced, whereas the Nth slot gathers all distances above a specified threshold (Figure 4).

Most often, the laser range scanners provide high quality range data, constrained to a narrow band around the true range values. The typical error distribution of a range scanner facing different surfaces resembles the distribution shown in Figure 4, with the exception of the final slot. In spite of the different nature of the errors associated to laser measurements in different environments, it is predictable that the shape and histogram weight will be similar to Figure 4, albeit with some variations due to the environment

features. In the majority of the indoor scenarios, the range error is more relevant than the other errors, outside the exceptional points where the laser data is useless. Thus, it seems acceptable to associate all the mentioned sources of errors as if they were due to the laser only, call it "the laser signature", and consider it fixed. One possible method for validating the results with Likelihood Test is to measure how close to the laser signature is the experimental curve.

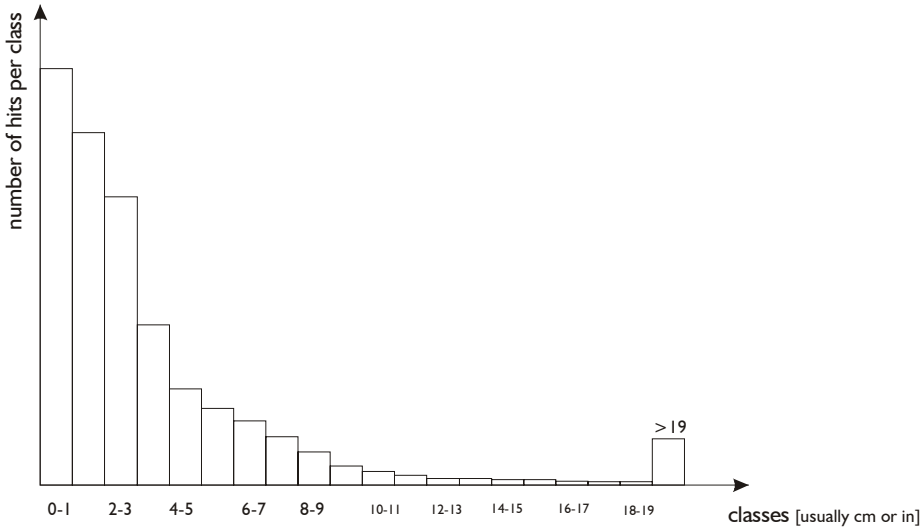

Figure 4 – A sample histogram with unit slots

Given the gaps in the map or the scan, the Likelihood Test analysis must be restricted to the directions where data exists on both profiles. This concern adds a new variable to likelihood test: before testing the point-to-point distance, it is relevant to measure how many pairs of points are being considered. If the map profile encompasses only a narrow field of view, for instance, if it is reduced to a corner, it is possible to obtain an erroneous high quality match, often at multiple postures. In Figure 5a, a sample laser scan is presented; the robot is represented at the scanning location. In Figure 5b, a small map profile is represented with the candidate posture computed by the Localisation algorithm.

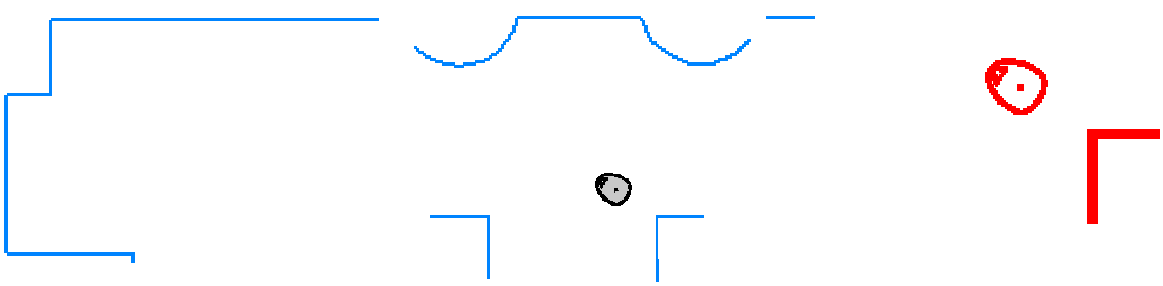

(a) (b)

Figure 5 – (a) Laser profile and (b) Small map profile

Figure 6 shows three different possibilities of matching the small map profile to the laser scan. Clearly, the map data is insufficient to locate the system unambiguously, *i.e.*, if only a small fraction of the world is mapped, the system can locate itself accurately at different spots. Thus, the number of valid pairs must be emphasised to grant liability to the Likelihood Test results. A high quality match based on few points should be regarded as a low confidence estimate.

This difficulty arises often at the second iteration of the 3D reconstruction process. The first iteration reconstructs only a small fraction of the world and for the second iteration, the RESOLV system "turned on its heels" to perceive the unknown areas.

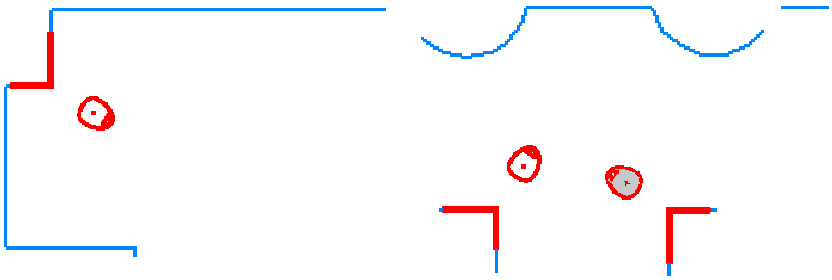

Figure 6 – Matching a small map at different postures

Finally, the main source of difference between the map and the laser profiles is the localisation error associated to the estimation of p_i, which adds to all the previous sources of error.

3.1.4 Error analysis and validation

For any candidate posture computed by a Localisation algorithm, p_i, there must be always some instances near the left of the histogram. These correspond to the scene elements used to locate the system at the proposed posture. The remaining instances define three different types of histogram, illustrated in Figure 7. The data in Figure 7 classifies the trial results in 1 cm slots and adds all distances greater than 19cm in the last slot.

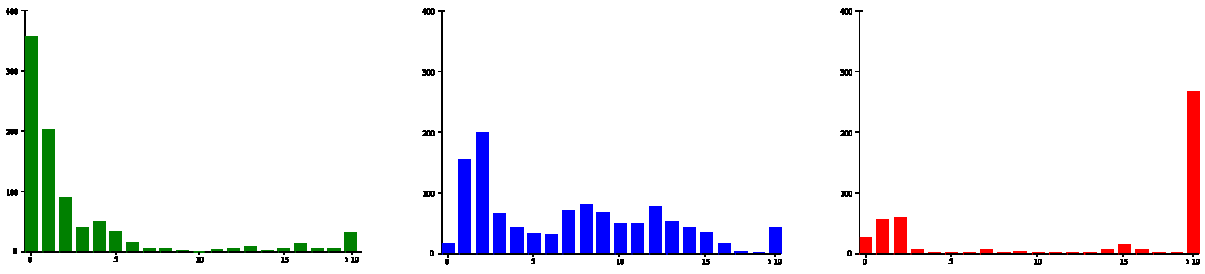

(a) (b) (c)

Figure 7 – Three posture candidates: (a) accurate estimate, (b) may be enhanced, (c) wrong estimate

If the two profiles are close representations of the same features, the errors will be small, filling the left area in the histogram (Figure 7a). If the two profiles are not close, most valid pairs will contribute with a large distance, (Figure 7c); the candidate posture is wrong. The posture estimates computed with Frame Localisation usually lie between these two extreme cases as in Figure 7b. The weight spreads from 0 to 15 cm but the weight of the last class is small. The Approximate Localisation, presented in Chapter 4, is supposed to correct the posture estimate, "pushing" the instances to the left.

This qualitative analysis may be quantified for automatic evaluation and validation. The point-to-point distribution will be summarised by the zero order, first order and second order moments:

1. The zero-th order moment counts the number of valid pairs used in the Likelihood Test. This is a measure of the relevance of the Likelihood Test. In Figure 6, Likelihood Tests were based in a small map and although the first and second order moments acknowledge a good match, the candidate postures are wrong.
2. The first order moment measures the expected value of the point to point distance distribution, assigning to the last class a single value fixed arbitrarily, usually at two times the last slot minimum bound. It is used to select the best posture among all the candidates, provided it passes the relevance test of the zero-th order moment.
3. The second order moment measures the dispersion of the distribution relative to the origin. In case there is more than one posture with first order moments close to the minimum value, the second order moment is used to discriminate between the best candidates. The posture with the lowest dispersion is chosen.

The automatic decision requires the tuning of thresholds and weighing factors. If the moment values fall within the specified thresholds, the posture estimate is validated and the result is made available to the remaining RESOLV modules. The algorithms used for Approximate Localisation embed the Likelihood Test as well. They start from an initial estimate, follow some criteria to find their next best estimates, and validate the results using the Likelihood Test. This is performed in closed loop until the moments computed with the Likelihood Test stabilise at local minima.

3.2 The Simulated Scan

When the Localisation algorithms were first designed, the simulated scan was generated from 3D Reconstruction data, based on the geometric features extracted from the scene. Later, the convenience of postponing the 3D Reconstruction from on site procedures to off-line batch processing introduced the use of pre-existent maps, translated into the current map format.

The origin of the map data is irrelevant for the algorithm development, as long as the map data is expressed in the specified format. However, the map source is crucial to the algorithm performance. The 3D reconstruction based maps describe the scene thoroughly and are update. However, they may be less accurate on corners and complex features that are difficult to reconstruct. On the other hand, pre-existent maps are often very incomplete and outdated.

When using pre-existent maps in the field trials, large differences appear when comparing the laser profile to the map profile, due to the absence of the actual features in the map description. To overcome this problem, two solutions were sought: ignore errors beyond a given (large) threshold if the map data was created beforehand or extend the map data, using the previous laser scans. The first option is immediate. It does not solve the problem but reduces the number of elements in the last class of the histogram, which could conceal a correct solution.

The second option based on using the laser scans to create a new map or extend an off-line map was sketched and tested in a very simple approach. The working principle is as follows: at the end of the iterations, the system evaluated an updated posture estimate, based on a laser profile and possibly other data sources. This scan may be regarded as a description of the environment perceived by the laser sensor from the current posture. The list of points defined by the current laser scan is added to the global list of points associated to previous iterations. The global list is a sort of map, where the clusters of points suggest the presence of obstacles. It should be noticed that the laser-based maps extension is incompatible with Frame Localisation since there are no map features. However, to perform adequately, an elaborate mapping algorithm is required, which proved to be beyond the scope of this dissertation.

3.2.1 Simulated scan based on map data

The algorithm for creating the map profile is based on geometric constraints only, considering the laser beam as a line segment starting on the laser sensor and intersecting the line segment that represents the scene feature. Thus, it does not model some of the real laser or surface characteristics.

The map profile is an array of range data, with the same resolution of the laser scan, defined within the same angular boundaries, resulting in a profile with the same dimensions as the actual laser scan. In the current implementation, the reflectance data for the map profile is not used, although this has been considered as a likely extension. The reflectance data associated to the map profile would not encode the actual reflectivity of the surfaces, but the degree of confidence associated to it, instead.

The map profile is based on the assumption that the laser sensor is located at $p_i = (x_i, y_i, \theta_i)$. The posture p_i is defined in world co-ordinates and is related to the system posture by a known co-ordinate transform.

The first step is to fill all the samples in the simulated scan with an invalid value,

which is a flag denoting the absence of data. This flag should be similar to the one used in the laser scan (see Appendix C, Section 1).

Afterwards, the value of the samples in the simulated scan is evaluated according to the following algorithm:

1. Select from the map line list a line segment (hereinafter called "the map line") that has not been considered before.
2. Compute the bearing coordinates of both ends of the map line, relative to the laser sensor. The result is the first and last array entries included in the triangle of visibility defined by the hypothetical laser sensor and the map feature under analysis.
3. Divide the angular span comprehended between the map line vertexes by the laser sensor angular resolution. The result is the number of array entries included in the triangle of visibility.
4. Define a loop from the bearing associated to one of the map line vertexes to the other vertex.
5. Compute the range from the sensor to the point where the "laser beam" intersects the map line. If the range is shorter than the current entry in the map profile, replace the entry with the new value.
6. Use the laser scan resolution to increment the beam bearing. Return to step 5 until the final bearing is reached.

This method, illustrated in Figure 8, mimics in the laser sweeping movement, discarding the experimental constraints, such as the sizeable footprint and the angle of incidence, among the most relevant. Although these two could be modelled with some additional data about the laser and surface characteristics, their weight in the point-to-point distance distribution is moderate, and the overall effect in the error analysis would be minimal.

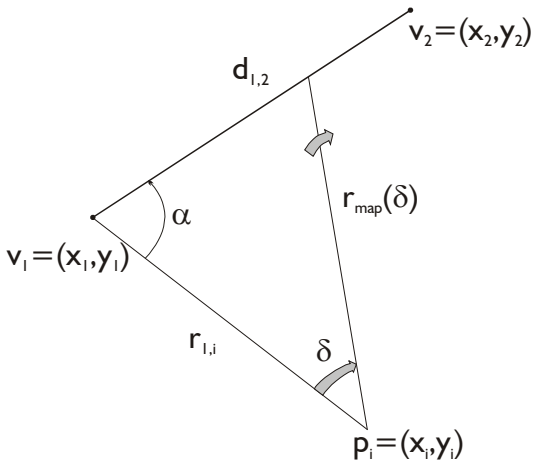

Figure 8 – Definition of the range profile associated to a map line

All the parameters in the simulated scan algorithm are defined in world coordinates. The result of the simulated scan is a laser-like range profile, defined in radial coordinates, relative to p_i.

The bearings b_1 and b_2 of the map line vertexes $v_1 = (x_1, y_1)$ and $v_2 = (x_2, y_2)$ are computed from equations (3.1). They are used to compute the angular span of the "laser beam" in the simulated scan profile. The range from vertex 1 to the laser, $r_{1,i}$ and the length of the map line, $d_{1,2}$ are also required (3.2).

$$\begin{cases} b_1 = \operatorname{atan2}(y_1 - y_i, x_1 - x_i) \\ b_2 = \operatorname{atan2}(y_2 - y_i, x_2 - x_i) \end{cases} \tag{3.1}$$

$$\begin{cases} r_{1,i} = \sqrt{(x_1 - x_i)^2 + (y_1 - y_i)^2} \\ d_{1,2} = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} \end{cases} \tag{3.2}$$

Vertex 1, v_1, acts as the pivot axis of the inner product between the two line segments, $[p_i, v_1]$ and $[v_1, v_2]$. The inner angle, α, is constant throughout the "laser beam" sweep; it is defined from the inner product equation, (3.3).

$$\alpha = \arccos\left[\frac{(x_i - x_1)\cdot(x_2 - x_1) + (y_i - y_1)\cdot(y_2 - y_1)}{r_{1,i}\cdot r_{1,2}}\right] \tag{3.3}$$

For the sweep movement, the "sinus law" can be used, (3.4), defining the unknown range, r_{map} as a function of the sweeping angle, δ. The angle δ starts at 0, when $r_{map} = r_{1,i}$, and is incremented by the angular resolution of the laser scanner. Rearranging the terms in (3.4), the variable $r_{map}(\delta)$ may be defined as simple function of δ (3.5), where $r_{1,i}\cdot\sin(\alpha)$ is a fixed parameter.

$$\frac{r_{map}}{\sin(\alpha)} = \frac{r_{1,i}}{\sin(\pi - \alpha - \delta)} \tag{3.4}$$

$$r_{map}(\delta) = r_{1,i}\sin(\alpha)\cdot\frac{1}{\sin(\alpha + \delta)} \tag{3.5}$$

Once the angle δ reaches vertex v_2, $\delta \geq b_2 - b_1$, the loop ends, and a new map feature is considered. When all the features have been "scanned", the simulated scan is complete.

The use of the "sinus law" is restricted to map features described by line segments.

Should the maps contain other features a more elaborate geometric algorithm would be required. A general purpose algorithm models the "laser beam" as a straight line with known parameters, models the map feature according to its description (usually a biquadratic equation) and finds the point of intersection of the two equations.

3.2.2 Simulated scan based on previous laser range scans

Creating a simulated scan from previous laser scans is a simple, albeit time consuming procedure. This is a last chance method to perform localisation in case no adequate map is available. As before, the simulated scan should have the same resolutions and angular boundaries as the actual laser scan.

In case there is an off-line map, it is transformed into a cloud of points, expressed in world coordinates, similar to the regular map profile. If there is no off-line map, the initial posture estimate is defined by convention or by the user and the Localisation algorithm is skipped, but the laser scan should still be performed.

Quite often, it is easy to create a rough, although incomplete, off-line map. For instance, if there is no map at all, the operator could measure the room with a tape and compute the maximum dimensions of the room envelope.

The algorithm is divided in two parts: simulated map update and scan generation.

Map Update

1. Translate any available off-line map to the line format used in the map definition. In case there are features on the map, which are no longer present in the scene, they should be removed from the map file.
2. Translate the map data into a cloud of points. This may be performed by sampling each line segment at regular intervals, which should approximate the typical interval between adjacent range samples. Using a laser sensor with 2500 points per revolution and scanning at a typical range of 3m, the laser resolution is 0.0075m. The inclusion of the vertexes is necessary. The resulting points should be stored on a (x, y) list, defined in world coordinates.
3. At the end of the iteration, use the laser scan and the evaluated posture to update the map. Compute the xy-position of each individual range sample in the current laser scan based on the estimated posture. These formulæ are simple coordinate transforms, (3.6). By definition, the last posture estimate is $p_{i-1} = (x_{i-1}, y_{i-1}, \theta_{i-1})$, r_k is the k-th range sample, $\theta_{\min}$ is the start pan angle expressed in world coordinates and $\Delta\theta$ is the angle increment.
4. Add the newly computed points to the (x, y) list, building the map incrementally.

$$\begin{cases} x_k = x_{i-1} + r_k \cdot \cos(\theta_{\min} + k \cdot \Delta\theta) \\ y_k = y_{i-1} + r_k \cdot \sin(\theta_{\min} + k \cdot \Delta\theta) \end{cases} \tag{3.6}$$

The procedure in step 3 should be performed iteratively, at the end of each iteration, extending the map as the RESOLV system moves in the scene. The resulting map is a cloud of points; a dense concentration of points indicates a likely feature whereas the free areas show the portions of the map where no features were detected.

This method is not quite a map generator, but a simple accumulator. For long sessions, with more than 10 iterations, the map data sets are very large, increasing the computation requirements. Constraints that are more serious arise, though, from the inconsistencies within the map data: the map features are coarsely defined, isolated and false features appear within the free space, the localisation errors accumulate, and the overall performance degrades.

Simulated Scan

1. For each instance in the map list, (x_m, y_m), compute the distance r_m and the bearing b_m to the candidate posture, p_i, (3.7).
2. If r_m is shorter than the previous range along bearing b_m, replace it with r_m.
3. Return to step 4 for all samples in the simulated laser.

$$\begin{cases} r_m = \sqrt{(x_m - x_i)^2 + (y_m - y_i)^2} \\ b_m = \operatorname{atan2}(y_m - y_i, x_m - x_i) \end{cases} \tag{3.7}$$

An implicit coordinate transform translates the system posture to the laser posture, both in world coordinates. It was mentioned before (Chapter 2, Section 4) that, in order to reduce the computational effort, the algorithms operated on the laser plane reference and only the final posture estimates were translated to the system coordinates using a known coordinate transform. Keeping both values in memory reduces some computation overhead and rounding errors.

If there is no map, the localisation algorithm is skipped in the first iteration.

The simulated scan computed in this form is a poor map: after a few iterations, it has too many instances. The instances are not grouped or classified according to its relevance or the degree of confidence associated to the reflectance data. It is clear this is not a solution, but only the first steps in a long path towards a possible solution, which is beyond the scope of the dissertation.

3.2.3 Differences between the laser scan and the simulated scan

The expression "simulated scan" is an abuse of language, to some extent. The differences between the proposed method and an actual laser scan are due to actual laser features and scene features. The actual laser features were discussed in Chapter 2, Section 5. These include by decreasing importance:

1. Angle of incidence: as the angle of incidence direction approaches the surface direction, the laser samples are subject to increased error distributions.
2. Range measurement errors: this depends on the laser construction and a curve with the range standard deviation as a function of range is usually supplied by the manufacturer.
3. Sizeable laser footprint: this depends on the laser construction and it is usually known. However, the errors induced by mixed footprints and the angle of incidence on the surfaces (see Chapter 2, sub-section 2.5.2) are difficult to model without oversampling.
4. Limited angular resolution: this depends on the device construction. It can not be overridden without changing the device.
5. Mechanical inaccuracies in the sensor set-up: during the life span of the equipment, minor misalignments between parts will occur. This effect can be reduced by careful and periodic calibration and correction of the device.
6. Electric and Control constraints: these relate to the supply voltage, response time of the controllers, etc..

The first and second features could be modelled in the "simulated scan". The sizeable footprint is most relevant in mixed footprints, which are very sensitive to short displacements. The detection of a mixed footprint is immediate but the definition of an accurate model remains a difficult task.

The fourth constraint is embedded in the generation of the "simulated scan". The two latter ones are difficult to model and it is expected that their influence is reduced to a minimum, by choosing the adequate hardware, device drivers and software.

In spite of all the mentioned error factors, the main causes of differences between an actual laser scan and a "simulated scan" taken exactly from the same posture arise from the scene characteristics. These issues were discussed in Chapter 2, Section 5 and will be mentioned again in Appendix C. The main issues are summarised below:

1. There is no reflectance data in the maps. Some work was developed on the extension of the Likelihood Test to reflectance data analysis, using laser scan reflectance only, map reflectance (or confidence) only and both. The results were

inconclusive and the extension will not be effective until the map data includes actual reflectance information. For the 3D reconstruction based maps, this is not a trivial task, though, since the reflectance depends on the laser characteristics. Since the model is reconstructed with multiple views, reflectance filtering and smoothing would be required. In case off-line maps are used, including reflectance would require scanning the environment beforehand and adding the reflectance data manually. The reflectance data is available only in the case of maps based on previous scans. The best solution would be to replace pure reflectance with a confidence measure, computed by reconstruction algorithms on 3D reconstruction maps or modulated by the user around a default value, according to educated guesses from previous trials.

2. The textures, fabrics and materials induce important variations on the measures, which are impossible to model beforehand.
3. The maps are incomplete or inaccurate. An inaccurate map, one with wrong measures or including features that are not in the actual scene, degrades severely the Likelihood Test. Since the Localisation algorithms have no other reference about the world, they will try to fit the current data into the available scene description. On the contrary, an incomplete map is a moderate nuisance because the Localisation algorithms are based on the available map data, regardless of excessive laser data.

Given the constraints described above, the deliberate option of not modelling most of the laser and surface features was taken. An adequate model of the laser features would require a detailed analysis and calibration of each particular sensor device. Modelling the effects of the scene on the simulated scan would require more assumptions and the operator to set environment-related parameters. Both these actions impair the generality of the reconstruction process.

On the other hand, if the effects of the real laser and scene features are ignored in the simulated scan the Localisation algorithms will be accounted for the whole difference between the two scans. Therefore, the computed point-to-point distance distribution will not represent the Localisation error but rather an upper bound. Nevertheless, if the computed point-to-point distance abides to the required accuracy criteria, the Localisation errors will necessarily meet these criteria by excess.

3.3 Formal Development

The elements of the Likelihood Test algorithm are:

- A set of NP postures, $P = \{p_n = (x_n, y_n, \theta_n), n = 1,2,...,NP\}$,
- A laser scan, $L = \{rng_i, rfl_i, i = 1,2,...,N\}$, with N samples, evenly distributed in

the angular interval between bearings $b_{\min}$ and $b_{\max}$, encompassing both range, rng_i, and reflectance data, rfl_i.

- A simulated scan associated to posture p_n, with N range samples, uniformly distributed in the angular interval between bearings $b_{\min}$ and $b_{\max}$, $S = \{s_i, i = 1,2,..., N \mid p_n\}$.

The set-up phase of Likelihood Test involves the definition of an accumulator with NS slots. The operator defines the number of slots (by default 20), as well as the slot size Δ_{slot} (by default 0.01m). All the elements in the accumulator, t_s, $s = 0,1, ..., NS-1$, should be set to zero.

The next phase is to compute the point-to-point distance histogram. For each posture $p_n \in P$, the algorithms follows the steps below, using only the range data:

1. Consider the first pair of samples, rng_i and s_i, with $i = 1$.
2. If rng_i or s_i are invalid ignore the pair and proceed to step 7. Otherwise, add one instance to the "number of valid pairs", NVP.
3. Compute the point-to-point distance, $|rng_i - s_i|$.
4. If $|rng_i - s_i| > T_{Map}$ ignore the pair and proceed to step 7. Otherwise the pair is considered a "matching pair".
5. If $|rng_i - s_i| < T_{Map}$ and $|rng_i - s_i| > (NS-1)\cdot\Delta_{slot}$ add one instance to the last slot of the point-to-point distribution: $t_{NS-1} \leftarrow t_{NS-1} + 1$. Proceed to step 7.
6. Otherwise, compute the respective slot: $k = floor\left(\frac{|rng_i - s_i|}{\Delta_{slot}}\right)$ and add an instance to it: $t_k \leftarrow t_k + 1$.
7. Proceed to the next sample $i \leftarrow i + 1$ and return to step 2 until the last sample.

The threshold T_{Map} corresponds to the maximum distance admitted. This value may be tuned by the operator, and is set by default at 0.5m. This condition is designed to cope with the map errors. Since the candidate postures, p_n, are computed by one of the Localisation algorithms with typical maximum errors below 0.005m, a point-to-point distance greater than T_{Map} (≈ 0.5m) is irrelevant for localisation and would serve only to degrade the likelihood estimation.

A point-to-point distance above T_{Map} is most often due to a false match produced by Frame Localisation, as illustrated in Figure 6, to a scene element featured in the laser scan and missed in the map or to an element in the map that is not detected in the laser scan (e.g., a glass door). The two latter causes affect both Localisation algorithms.

The last phase in Likelihood Test is the computation of the distribution moments

and the result analysis.

The zero-th order moment, (3.8), measures the number of matching pairs considered. The first order moment, (3.9), is the expected value of the point-to-point distribution restrained to the matching pairs. The second order moment, (3.10), is the dispersion of the distribution relative to zero, restrained to the matching pairs.

The last slot is much wider, since it ranges from $NS \cdot \Delta_{slot}$ to T_{Map}. Hence, it should be handled differently from the remaining slots. Its instances are represented by an equivalent distance, $\Delta_{\max}$ defined by convention. This parameter may be tuned by the user and is set by default as $\Delta_{\max} = \frac{3}{2} NS \cdot \Delta_{slot}$.

$$MatchPairs(p_n) = \sum_{s=0}^{NS-1} t_k \tag{3.8}$$

$$ExpectValue(p_n) = \frac{\sum_{k=0}^{NS-2} t_k \cdot \left(\Delta_{slot} \cdot \left(k + \frac{1}{2} \right) \right) + t_{NS-1} \cdot \Delta_{\max}}{MatchPairs(p_n)} \tag{3.9}$$

$$Disp(p_n) = \frac{\sum_{k=0}^{NS-2} t_k \cdot \left(\Delta_{slot} \cdot \left(k + \frac{1}{2} \right) \right)^2 + \left(t_{NS-1} \cdot \Delta_{\max} \right)^2}{MatchPairs(p_n)} \tag{3.10}$$

For development purposes, the point-to-point distribution and the evaluated moments are presented to the user for thorough analysis. However, for automatic implementation of the algorithm, decisions must be taken.

The likelihood validation criterion is simple: select the range of solutions, p_n, which are close to the maximum number of valid pairs, $MatchPairs(p_n), n = 0,1,\ldots$. Among these, select the candidate posture with lowest expected value (3.9). If there are multiple postures with close expected values, the lowest dispersion moment, (3.10), is used to select the best posture estimate.

The difference tolerance within the same moment and the relative weighing of the three moments are moderately important when dealing with well conditioned problems. However, the definition of a suitable cost function is crucial to achieve robustness in the ill-conditioned cases. Moreover, given the variety of the environments visited in the field trials, absolute criteria seem inadequate. Therefore, the decision making was implemented according to the following rules:

1. A posture p_n is valid iff $\frac{MatchPairs(p_n)}{NVP} > MinMatchRatio$. This ratio may be defined

by the user and is 0.3 by default.

2. A posture p_n is valid iff $ExpectValue(p_n) < E_{\max}$, where $E_{\max}$ is the accuracy requirement on position mentioned in Chapter 1, Section 4. This is also available to the user for tuning. The default value is $E_{\max} = 0.05m$.

3. In case the two previous rules hold, the number of match pairs and the expected value are combined into a single cost function, (3.11). A candidate posture is rejected in case it has a cost much higher than the optimal cost. The α parameter has the dimension of a distance and weighs the average cost of non-match pairs.

4. If two postures p_1 and p_2 are considered to be close according to the $MatchPairs(p_n)$ criterion, the weighing criterion emphasises the first order moment at the expense of the zero-th order moment, (3.12). In case the first order moments are clearly different, the lower expected value is chosen. Otherwise, the two zero-th and first order moments are weighed and the lower overall value is chosen.

5. Correspondingly, if two postures p_1 and p_2 are considered to be close according to the $ExpectValue(p_n)$ criterion, the weighing criterion emphasises the second order moment at the expense of the first order moment, (3.13). In case the second order moments are clearly different, the lower dispersion is chosen. Otherwise, the first and second order moments are weighed and the lower overall value is chosen.

$$\cos t(p_n) = ExpectValue(p_n) + \alpha \cdot \frac{ValidPairs(p_n) - MatchPairs(p_n)}{ValidPairs(p_n)} \tag{3.11}$$

$$\frac{MatchPairs(p_1) - MatchPairs(p_2)}{MatchPairs(p_1)} << \frac{ExpectValue(p_1) - ExpectValue(p_2)}{ExpectValue(p_1)} \tag{3.12}$$

$$\frac{ExpectValue(p_1) - ExpectValue(p_2)}{ExpectValue(p_1)} << \frac{Disp(p_1) - Disp(p_2)}{Disp(p_1)} \tag{3.13}$$

The α parameter in equation (3.11) should be tuned by the operator according to the map completeness. If the map represents is very complete, the parameter should put a high penalty if the number of match pairs is inferior to the best posture. On the opposite case, the parameter should be set lower to allow for further analysis.

The exact balance of the two last conditions is also available to the user for tuning. In fact, one of the reasons to provide the extensive data sets to the operator is to help an experienced operator to change the parameters to fit wherever environment the system is

reconstructing. It should be emphasised, though, that the default values performed correctly during the field trials, and these additional parameters are only meant to extend the operation range to difficult environments.

A possible extension would store the posture candidates that comply to the rules 1, 2, 3 and would confirm or discard the candidates in hold as the system progresses along the scene, depending on its corroboration during the next iterations. This requires an auxiliary posture estimate, such as odometry, to relate the data taken in different experiments.

3.4 Experimental results

This section presents the Likelihood Test results of trials performed in different environments. The Localisation process starts always with Frame Localisation eventually followed by Approximate Localisation, which will be discussed in Chapter 4.

These examples were obtained on two office-like rooms: a classroom without windows at the laser height scan and an office with two walls covered with windows (presented in Chapter 2, Section 6). The dimensions in these two rooms are within the laser operating range.

Results from a third trial, performed in a factory like environment, are presented next. This scenario is much wider than the laser operating range. In addition, the reconstruction proved to be quite difficult due to the machinery present there and the large empty spaces where the laser fails to collect any data. Therefore, an off-line line map was used (Figure 1), but it included only the architectural features and not the actual machinery installed.

The three trials were performed with the AEST equipped with the Acuity laser scanner (see Chapter 1, Section 3). This laser has a maximum measurement range of 20m but accurate measurements on medium or low reflectance surfaces are possible only below 10m to 12m. The angular resolution (number of points per revolution) is 2400 points.

3.4.1 The classroom

The classroom is as close to the optimal environment as it was possible to find in the faculty. It's a rectangular room with four walls almost parallel, painted white, without windows at the scanning height or other elements capable of disturbing the laser measurements (Figure 9). The size of the room is within the laser operational range and the target surface is almost uniform all around the field of view.

(a) North wall

(b) East wall

(c) South wall

(d) West wall

Figure 9 – Two classroom walls

The map of the room is presented in Figure 10, superimposed to the Frame Localisation candidate postures and the laser range scan associated to the preferred posture.

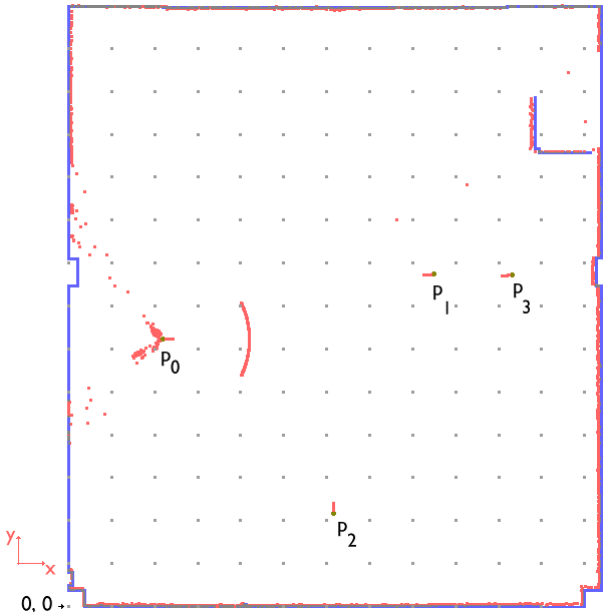

Figure 10 – The classroom map and the Frame Localisation results

The grid spacing is 0.5m. The north wall is on top, with the white board. The x-axis runs horizontally from left to right and the y-axis runs vertically from bottom to top.

The Frame Localisation results are presented in Table 1. The weights suggest the high degree of confidence granted to the preferred posture: it contains more than 60% of the total weight and the next best estimate weighs nine times less.

On the right of Table 1 there is the summary of the Likelihood Test applied to the candidate postures. The number of valid pairs (NVP) is 1749 for all postures. Match pairs (MP) were defined as those with a point-to-point distance lower than 0.4m. From the match pairs to valid pairs ratio one can readily conclude that postures p_0 and p_3 match almost the whole environment (83% and 73%) while p_1 and p_2 match only a fraction of it (51% and 39%). The missing pairs in p_0 are mostly due to the laser samples in the rear of the laser sensor that hit the robot parts and do not reach the target surfaces. These are apparent to the left of Figure 10. The arc at the right represents invalid laser samples (the range is coded as –1.0).

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	1.102	3.111	0.046	0.6281	1458	0.0289	✓
p_1	4.252	3.868	-3.084	0.0681	900	0.0496	✗
p_2	3.082	1.090	1.632	0.0527	676	0.0414	✗
p_3	5.152	3.859	-3.091	0.0508	1276	0.0641	✗
plus 12 clusters with a total weight of 0.2003							

Table 1 – Frame Localisation results with Likelihood Test validation

The point-to-point distance has a low expected value (EV), close to the typical laser range error for the Acuity Laser Scanner (≈ 0.02m). There is the possibility of confusion between two orientations denoted by p_0 and p_3, symmetric with respect to the centre of the classroom (Figure 11). This ambiguity was foreseen during the planning phase and an L-shaped form was added to the north-east corner before the trials (Figure 9b and Figure 10). The L shaped form minimises the ambiguity, reducing the number of match pairs in p_3 (182 less matches) and increasing the point-to-point distance of pairs associated to the L. This degrades the p_3 moments relative to p_0, but the p_3 estimate is still relative accurate due to the remaining pairs still matching the four walls.

The candidate p_3 is rejected according to rule 3, given its higher expected value and lower number of valid pairs. Notwithstanding the candidate posture is also submitted to Approximate Localisation for illustration purposes.

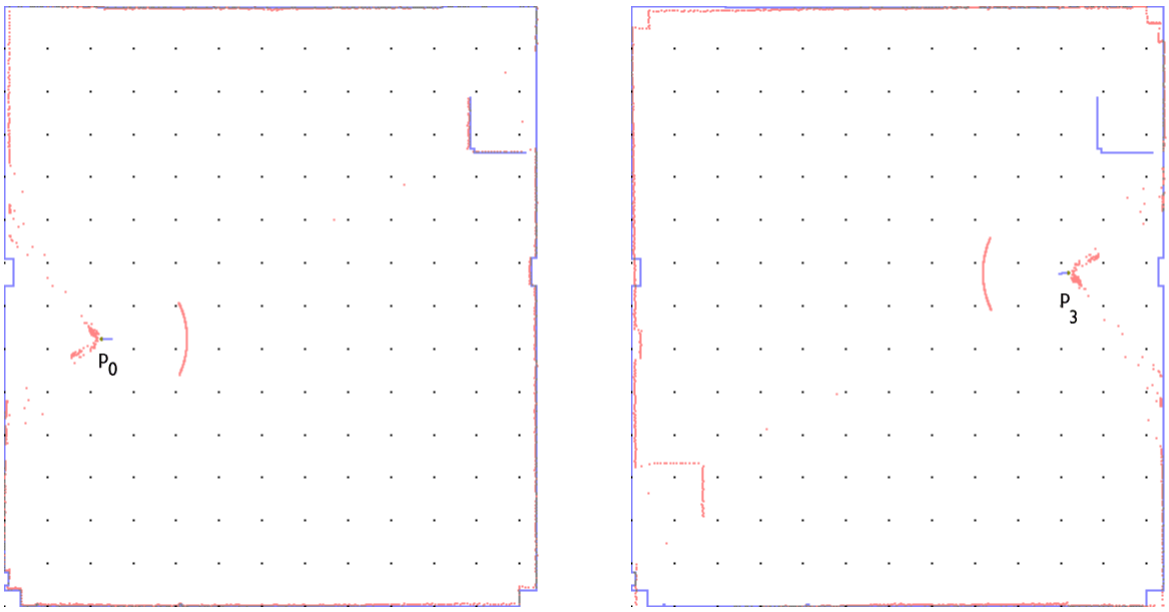

Figure 11 – The postures p_0 (left) and p_3 (right) are symmetric

The candidate postures p_0 and p_3 were refined with Approximate Localisation. The algorithm used was Error Descent with an initial 0.04m/0.004rad step and 7 iterations (see Chapter 4, Section 3 for details). Running the algorithm with such fine parameters - the final step examines postures separated by 0.625×10^{-3}m - procures high quality results at the expense of longer computation times.

The results of Approximate Localisation are summarised in Table 2. The correction column contains the coordinate difference between the Frame Localisation and the Approximate Localisation in millimetres and 10^{-3}rad (mili-rads). The last column presents the square root of the second order moment, *i.e.*, the deviation from the expected value. Smaller values mean histograms that are more concentrated.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	1.1261	3.1119	0.04632	24.4, 0.63, -0.07	1456	0.02197	0.04369
p_3	5.0924	3.8836	-3.08823	60, 25, 2.75	1284	0.03388	0.04572

Table 2 – Approximate Localisation results

The point-to-point distributions are presented in Figure 12. In the background (dark columns) the point-to-point distance distribution associated to the Frame Localisation estimate is presented. Its Approximate Localisation counterpart is presented in the foreground (lighter columns). The horizontal axis depicts the histogram classes and the vertical axis the number of instances in each class (see also Section 1).

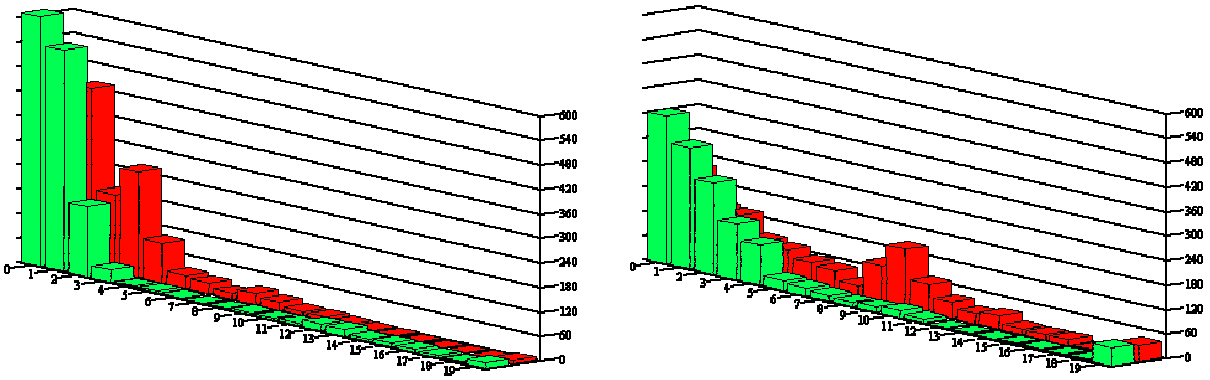

Figure 12 – The point-to-point distance distributions for p_0 (left) and p_3 (right)

The refined posture p_0 has reached an optimal value. Its point-to-point distance distribution is similar to the laser range error histogram of the laser device used. Beyond its own merits, the Localisation algorithms benefit of the accurate map description, wide flat surfaces, white walls and the limited size of the room.

In case of posture p_3, the results are also enhanced but its intrinsic error remains, particularly on the corners and the L shaped feature, responsible for the reduction on the number of valid pairs. This is apparent in the last class that gathers the points that are more than 0.19m apart. The weight of this class is not reduced, in fact it is increased by four points (41 to 45).

From the localisation point of view, this experiment was conducted in an optimal room. It illustrates the Localisation methods and their possibilities when the actual constraints match the requirements stated in the design phase. The next experiments show how Localisation copes with the difficulties encountered in regular environments, which are usually more difficult.

3.4.2 The office

This room was introduced in Chapter 2, Section 6, where the Frame Localisation results for two different trials were discussed. The trial methodology and the presentation of results are the same as in the classroom.

Two images of the room were presented there. In Figure 13 there are four additional images. The north wall is on the right of the maps (see Chapter 2, Section 6) and the door indicated in the map is on the west wall, next to the north-west corner.

(a) North wall

(b) East wall

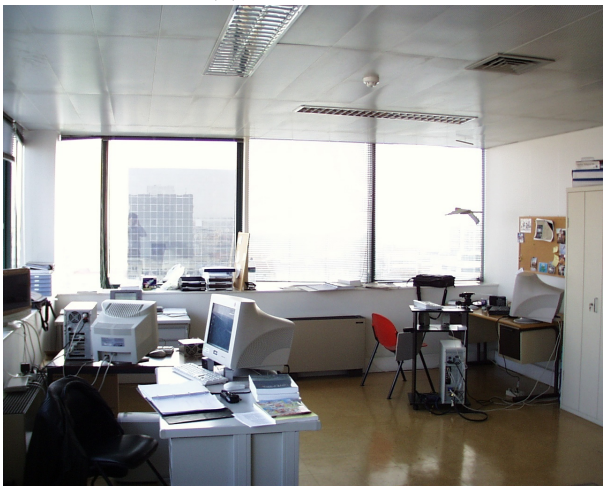

(c) South wall

(d) West wall

Figure 13 – The four office walls

As it was discussed in Chapter 2, Section 6, the Frame Localisation algorithm succeeded in computing a good posture estimate in spite of the two walls covered with windows, taking advantage of its ability to "anchor" on the good features, ignoring the directions without useful data.

The results of trials (a) and (b), performed at different locations in the office, are presented in Table 3 and Table 4, respectively, including the validation information computed with the Likelihood Test. The maximum match distance, T_{Map}, was set to 0.4m.

In Trial (a) the number of valid pairs (NVP) is 1732. In Trial (b) the number of valid pairs is 1736.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	4.137	2.481	-0.019	0.3936	955	0.0614	✓
p_1	4.156	2.909	-0.010	0.1405	402	0.0484	✗
p_2	3.718	1.154	1.551	0.1206	313	0.0927	✗
p_3	2.921	1.154	1.551	0.0978	431	0.0961	✗
p_4	2.453	1.154	1.550	0.0931	524	0.0695	✗
p_5	4.192	7.821	3.136	0.0779	308	0.0378	✗
p_6	10.827	2.010	1.570	0.0497	306	0.0328	✗
plus 2 clusters with a total weight of 0.0268							

Table 3 – Trial (a): Frame Localisation results with Likelihood Test validation

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	3.593	3.000	1.541	0.2265	1191	0.1064	✓
p_1	3.575	3.425	1.537	0.1182	540	0.2030	✗
p_2	3.331	3.002	1.543	0.0807	1060	0.0785	✓
p_3	4.771	7.257	-1.507	0.0705	446	0.0496	✗
p_4	3.352	10.153	-0.032	0.0439	292	0.0884	✗
plus 22 clusters with a total weight of 0.4602							

Table 4 – Trial (b): Frame Localisation results with Likelihood Test validation

Only one candidate posture in Trial (a), p_0, and two candidate postures in Trial (b), p_0 and p_2, passed the first validation rule, regarding match pairs (MP): $MatchPairs(p_n)/NVP > MinMatchRatio$. Therefore, the expected value (EV) criteria were considered only in these three cases.

In Trial (a) there is only one candidate. Since it complies to the accuracy rule, it is validated. In Trial (b) there are two candidates; they both comply to the accuracy rule and to the relative cost rule, (3.11), so they are both valid.

The fourth rule, (3.12), should lead to the selection of the best candidate among the valid ones. Applying the rule to the figures in Table 4 yields:

$$\frac{MatchPairs(p_1) - MatchPairs(p_2)}{MatchPairs(p_1)} = \frac{1191 - 1060}{1191} = 0.11$$

$$\frac{ExpectValue(p_1) - ExpectValue(p_2)}{ExpectValue(p_1)} = \frac{0.1064 - 0.0785}{0.1064} = 0.262$$

Although the second candidate is based on less match pairs than the first (11% less), the expected value of the point-to-point distance is significantly lower than the first candidate (26% less).

Deciding which is the best posture estimate is a matter of thresholds and tolerance. Under the current parameters, the 26% reduction on the expected value is more important than the 11% loss in the number of match pairs. Therefore, the best posture estimate is p_2. These results require further analysis. The two candidate postures are almost parallel, $\theta_2 - \theta_0 = 0.0022rad$; they have approximately the same Y coordinate, $y_2 - y_0 = -0.0026m$, and they are separated by $x_2 - x_0 = 0.262m$. As explained in Chapter 2, Section 6, posture p_2 is induced by the match of the laser samples associated to the blinds on the window surface, approximately 25cm behind it (compare left and right in Figure 14). Thus, the number of pairs with a high point-to-point distance is reduced while some others are removed from the distribution when the maximum match threshold, T_{Map} is exceeded. This explains the lower expected value in p_2.

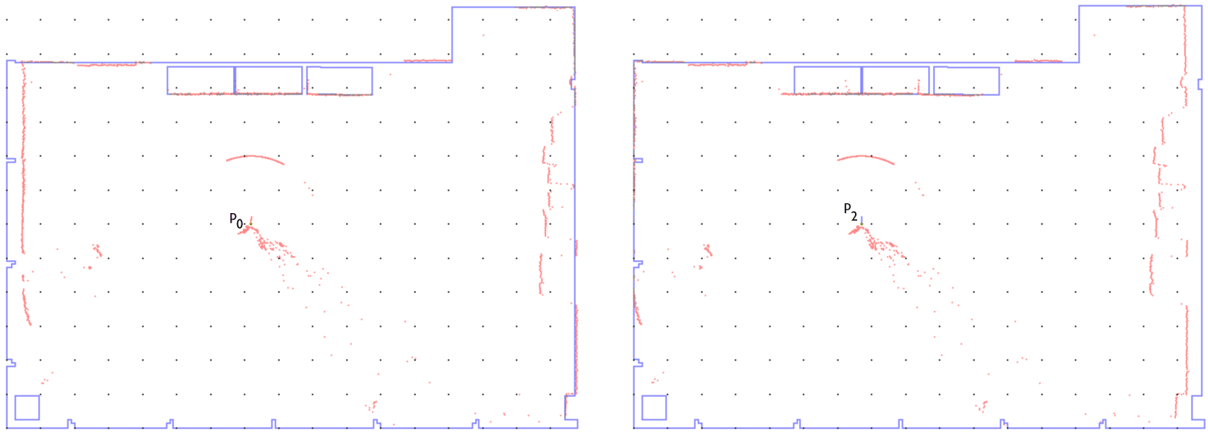

Figure 14 – Trial (b): The valid Frame Localisation candidates: p_0 (left) and p_2 (right)

Submitting the valid posture candidates to the Error Descent algorithm for Approximate Localisation yields the refined results. The Error Descent parameters are the same for the classroom: initial 0.04m/0.004rad step and seven iterations (see Chapter 4, Section 3). The results for Trial (a) are summarised in Table 5.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	4.1363	2.45510	-0.01543	-0.66, -25, 3.18	966	0.0546	0.1113

Table 5 – Trial (a): Approximate Localisation results

In spite of the modest posture correction, the point-to-point distribution is clearly pushed to the left, meaning a lower expected value (Figure 15, the Frame Localisation distribution is in the background while the Approximate Localisation distribution is in the foreground). However, the EV figure is not reduced further, due to the samples accumulated in the last slot, which can not be enhanced with a little correction. The samples in the last slot account for 67% of the distribution weight. Reducing T_{Map} would eliminate most instances in the last slot, strengthening the meaning of the expected value as a measurement of the algorithm error.

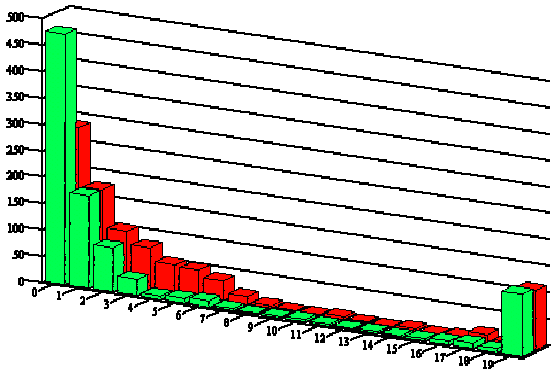

Figure 15 – The point-to-point distance distributions for Trial (a)

The instances in the last slot are not enhanced by the Approximate Localisation algorithm because they are caused by the differences between the laser scan and the map, which are apparent in Figure 16. The samples inside the ovals hit objects that were not present in the map.

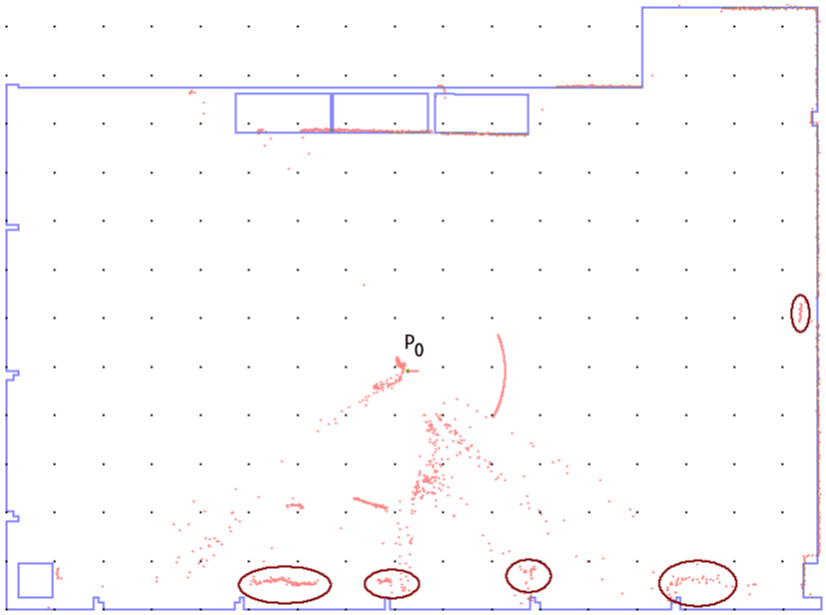

Figure 16 – Trial (a): the laser scan centred at the refined posture

If the errors in the map were minimised by setting $T_{Map} = 0.25$, the expected value would be close to 0.024m, which is a fair accuracy given the laser range error distribution. It is clear that the map inaccuracy or incompleteness can seriously hamper the algorithm performance. Notwithstanding, the default parameters should be set by excess at the expense of some apparent performance degradation to cope with the broadest range of operation environments. A specialised operator can then tune the parameters on site according to the features found and the log files.

In Trial (b) the map issues mislead the algorithm, leading it to wrong postures as it will be shown. The Approximate Localisation results are summarised in Table 6 and the point-to-point distributions are presented in Figure 17.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	3.4333	3.0	1.53790	-160, 0, -3.12	1104	0.0737	0.1687
p_2	3.3508	3.0	1.53924	20, -2.5, -4	1071	0.0738	0.1370

Table 6 – Trial (b): Approximate Localisation results

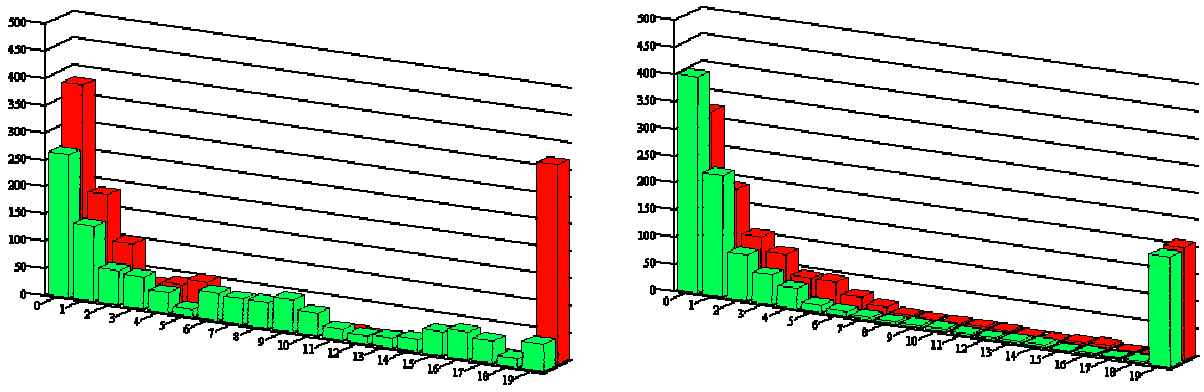

Figure 17 – Trial (b): the point-to-point distance distributions for p_0 (left) and p_2 (right)

The histogram analysis highlights the difficulties encountered. While posture p_2 was refined with a small correction, mainly on the X-axis (0.02m) and on the orientation (-0.004rad), posture p_0 moved 0.16m along the X-axis (to the left). The reasoning behind this dramatic difference is apparent from Figure 17.

On the left, (p_0), the majority of the samples in the last slot were redistributed among the other slots, especially between 8 and 19cm while the remaining points were discarded (the number of match pairs is reduced from 1191 to 1104). This effectively reduces the expected value, but it does not correspond to any match. The resulting posture is represented on the left of Figure 18, with the laser scan centred on it.

On the right of Figure 17, the distribution associated to p_2, shows a typical

enhancement by Approximate Localisation, while keeping the map errors in the last slot. Again, the elements in the last slot, which are mainly due to the poor map, account for the majority of the expected value, (77%).

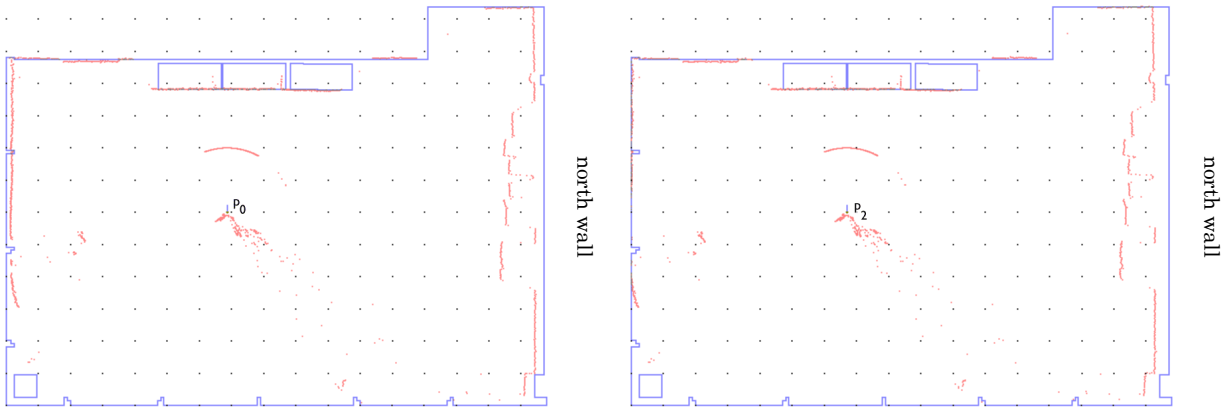

Figure 18 – Trial (b): the refined postures: p_0 (left) and p_2 (right)

What should be done to cope with these errors? After all, the Frame Localisation computed the correct posture as the best candidate and the Approximate Localisation was driven into two wrong solutions. How can this be explained?

Looking to Figure 18 one sees that on the north wall (to the right) a long section of the scan does not fit the map. This section corresponds to the printers and computer hardware installed there, which are visible in Figure 13a. The line extraction procedure transformed these samples into short line segments, which are of little relevance for Frame Localisation. On the south wall (to the left), the blinds were lowered during the experiment, whereas the map did not include any blinds (Figure 13c). The blinds are made of aluminium blades, a highly reflective material. Again, the line extraction procedure transformed the laser samples representing the blinds into medium or short lines with moderate confidence. Therefore, the Frame Localisation algorithm found the p_0 solution with high confidence and also the p_2, albeit with a lower confidence. The posture p_2 results of a frame match based on the blinds samples instead of the north wall samples.

When the two solutions are analysed with the Likelihood Test, the map-to-scan comparison is thorough and all laser samples have equal relevance – in the current version used in the field trials the Likelihood Test does not account for reflectance variations. Thus, the method measures only the number of match pairs and the expected value to assess the quality of a posture candidate and these numeric pitfalls are possible.

Should this be accounted as a Localisation error? In case of p_0 it is clearly an error and one might be tempted to correct it by tuning the parameters, especially T_{Map}. In case of p_2 it is not strictly a Localisation error. The working principle for Localisation was "fitting the laser scan to the map as good as possible" and the map is the ultimate truth about the environment. Nevertheless, it is not the expected result.

One possible solution would be to lower the weight of the samples in the last slot. This is almost equivalent to reduce T_{Map}. Another solution would be to raise the blinds, but this would solve this very particular case, only. Finally, the preferred solution ought to be *to include the blinds in the map because they are actually there.*

This is one of the main reasons for creating different mapping modes and mixed map modes and storing the maps with very simple syntax. Adding a new map feature is almost immediate: it suffices to measure its boundary at the scanning height, compute the vertex coordinates in the map reference and add the feature vertexes to the map file.

The Trial (b) was repeated as Trial (c) with a new line added to the map, corresponding to the blind present in the laser scan. The experiment parameters and the presentation of results are the same as before. The results are summarised in Table 7. Comparing Table 7 with Table 4 shows minor adjustments in the clusters. It also shows increased weight granted to the preferred solution p_0: 0.287 against 0.227. On the contrary, the Likelihood Test validation shows significant improvements in the expected value for the preferred solution (p_0) and a strong degradation of the expected value associated to posture p_2. This is a direct result of the newly inserted feature: in p_0, the distance from the blinds samples to the new feature is negligible while the windows that induced p_2 are no longer a part of the simulated scan because the newly inserted feature is barring the "laser field of view" towards them.

According to the thresholds used in the trial, p_2 would be discarded because its cost (3.11) is much higher than the cost of p_0. However, for the sake of illustration both candidates are submitted to Approximate Localisation.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	3.601	3.014	1.548	0.2870	1185	0.0591	✓
p_1	3.584	3.438	1.542	0.1218	540	0.0921	✗
p_2	3.331	3.002	1.543	0.0536	1061	0.1416	✗
p_3	3.365	0.607	3.111	0.0469	338	0.1027	✗
p_4	4.771	7.257	-1.607	0.0468	446	0.0496	✗
p_5	5.021	7.255	-1.607	0.0406	445	0.0518	✗
plus 26 clusters with a total weight of 0.4034							

Table 7 – Trial (c): Frame Localisation results with Likelihood Test validation

In Figure 19 the map is shown with clusters p_0, p_1 and p_2. The laser scan is also represented, centred on p_0.

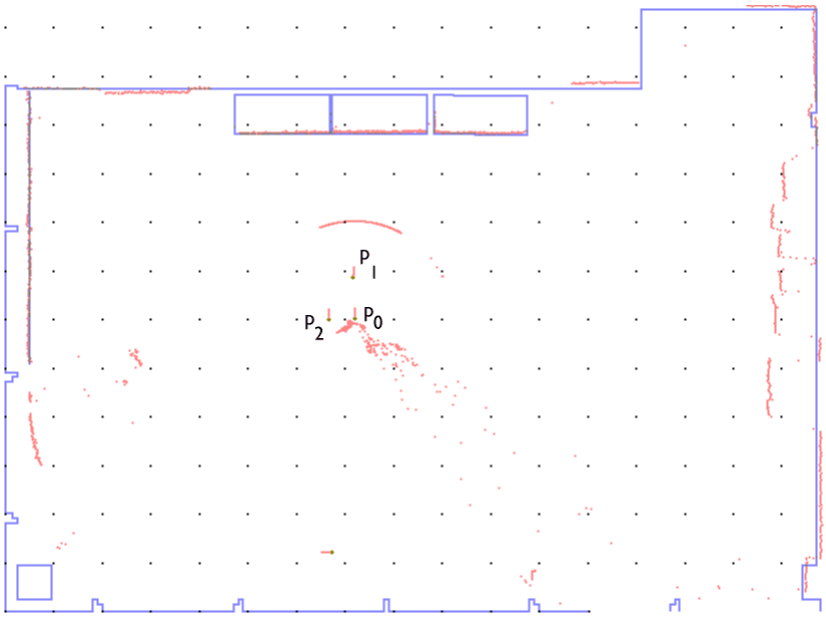

Figure 19 – Trial (c): the extended map with the Frame Localisation clusters and the laser scan centred on the preferred posture

The Approximate Localisation was run with Error Descent with an initial step of 0.04m and 0.004rad and 7 iterations. The results are summarised in Table 8 and the distributions are presented in Figure 20.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	3.5707	3.0014	1.5373	-0.03, -0.0125, -0.01025	1163	0.0477	0.0982
p_2	3.5708	2.9950	1.5361	0.24, -0.0075, -0.00713	1167	0.0476	0.0988

Table 8 – Trial (c): Approximate Localisation results

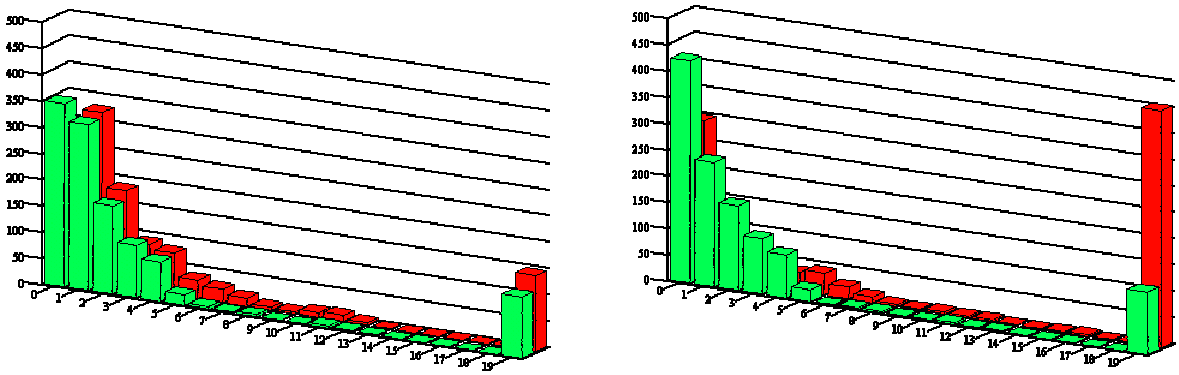

Figure 20 – Trial (c): the point-to-point distance distributions for p_0 (left) and p_2 (right)

The most noteworthy fact is that the algorithm fed with two different initial postures converged to a single posture. Initially, $p_0^{FL} - p_2^{FL} = (0.2699m, 0.0114m, 0.00427rad)$

and after Approximate Localisation, $p_0^{AL} - p_2^{AL} = (-0.00014m, 0.00645m, 0.00115rad)$. The histogram analysis illustrates the move of most instances in the rightmost slot of p_2 to the first slots. Although the first two slots have different weights in the two histograms, the statistical distribution is similar as shown by the first three moments in Table 8.

The insertion of a new map feature describing the blind suppressed most disturbances in the Localisation process. Usually, adding a single feature has a less dramatic effect; in this case, this happens also because there is a wide fraction of the scan useless because of the windows. The simple office map ignored the hardware elements on the north side and the blinds on the south side and the definition of the x-posture was very "loose", since it was based on very little samples. From the moment a good feature has been found, the algorithm anchors on it and provides a sound posture estimate (Figure 21).

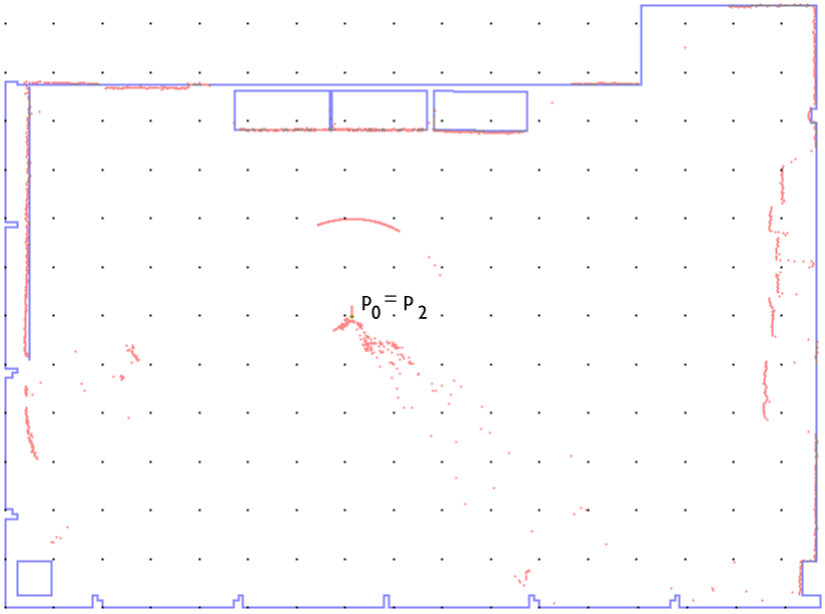

Figure 21 – Trial (c): the final posture estimate, computed by Approximate Localisation

The remaining elements on the rightmost slot are most likely due to the hardware on the north wall and other minor map inaccuracies. They still contribute with 64% of the overall expected value. Reducing T_{Map} to minimise this cost and enhance the overall performance would probably prevent the convergence of the Approximate Localisation from p_2^{FL} to $p_2^{AL} \cong p_0^{AL}$. This illustrates the advantages of using large parameters to accommodate different working conditions. In regular trials, the initial estimate fed into Approximate Localisation could be distant from the best posture, making it necessary to accept large point-to-point distances that will be reduced as the algorithm converges. Otherwise, the number of match pairs will be cut, reducing the relevance of the estimated postures.

3.4.3 The “factory”

The factory like environment is in fact a laboratory in the faculty. It is equipped for tests with concrete structures, seismic simulations; there is a wind tunnel and installations for evaluation of corrosion and humidity effects on materials and many other issues in construction engineering. The images in Figure 22 were captured with the video camera embarked in the AEST during an acquisition campaign.

For simplicity sake, the laboratory will be named the “factory” throughout the remaining chapter.

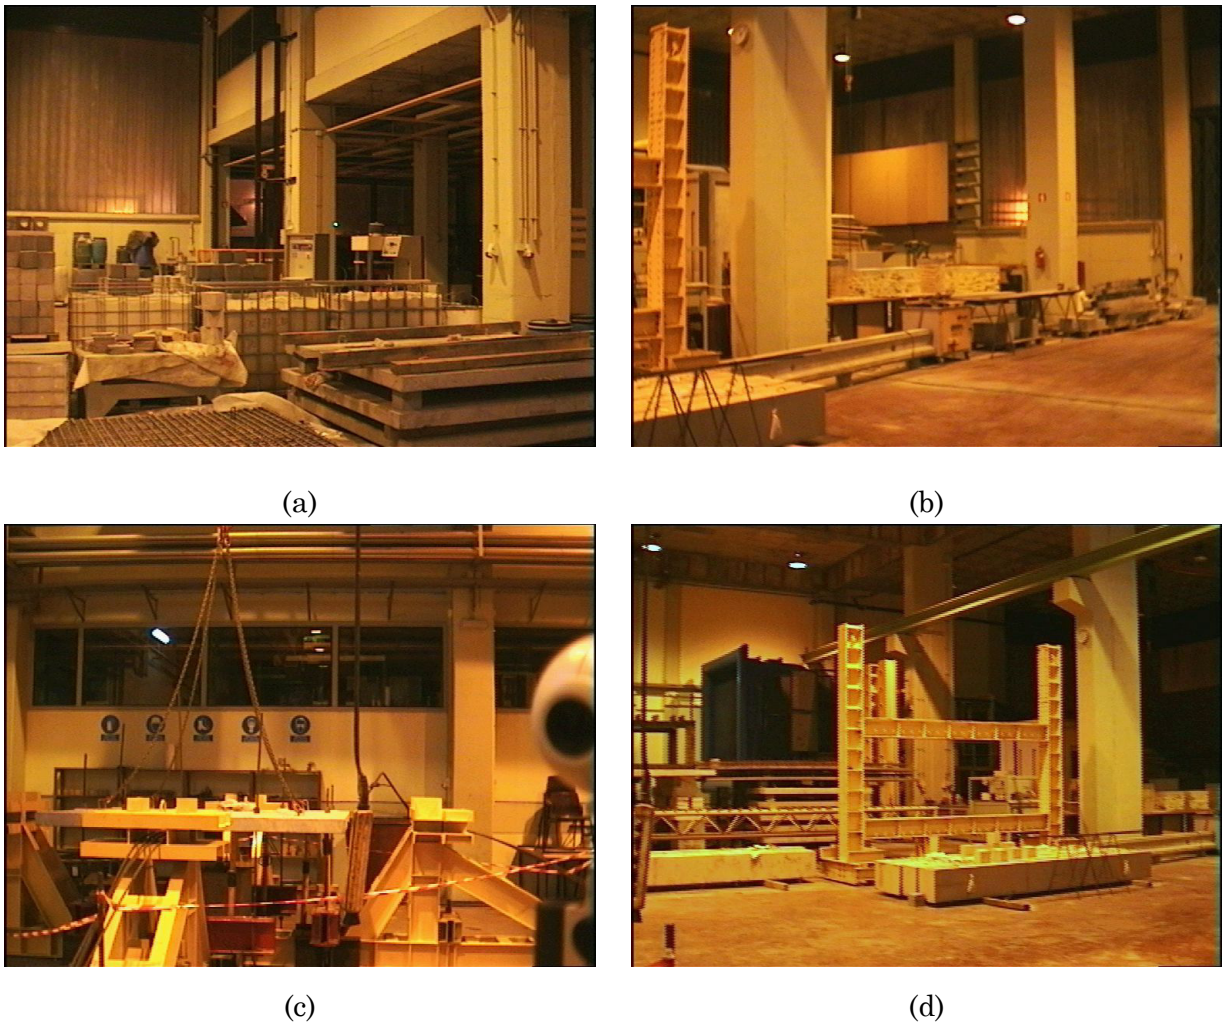

(a) (b)

(c) (d)

Figure 22 - Partial images from the “factory”

The map of the “factory” was introduced in Figure 1, to illustrate the laser field of view. The room is 34.4m long and 32.1m wide; the distance between the columns is approximately 8.2m. The “factory” map was based on the architectural plant where only the walls and the columns are described.

All around there are different types of hardware and materials; both on the ground

and suspended from the ceiling by cranes or other devices. The navigation of the AEST within this scenario is difficult, as it has to circumvent any obstacles while trying to explore the entire environment and keep close to the walls to keep track of localisation and enhance the 3D reconstruction process.

The reconstruction of the "factory" was performed in two sessions. On each session, the AEST followed a path given by the human operator, step by step, *i.e.*, when an iteration is complete the goal for the next iteration is set by the human operator and the AEST travels toward the destination. Each iteration consists of Localisation, laser data acquisition, video acquisition and movement. The 3D reconstruction process was postponed to a later phase. Therefore, an off-line map was required, and the map in Figure 23 was defined, based on the architectural plans.

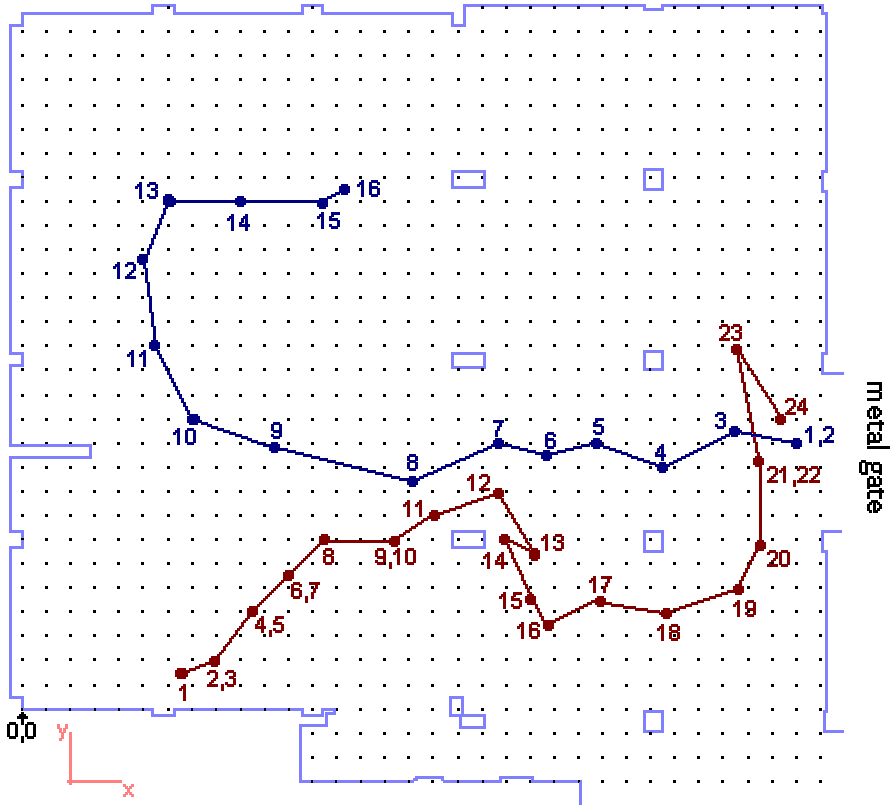

Figure 23 – The "factory" map with the proposed trajectory and goals

First Session

The first session consists of 24 iterations, starting from the left-lower corner of the map (Figure 23). A few iterations were repeated due to operation errors or hardware failures. On iteration 1, the AEST posture was measured both with a measuring tape and with Localisation and the difference was within the accuracy requirements. It should be

recalled that, for reconstruction purposes, the maximum posture error should be less than 0.1m and 0.02 rad (1°).

The results presented below were obtained at iteration 2, 15 and 23. The Localisation algorithm starts with Frame Localisation, followed by Likelihood Test validation and Approximate Localisation, just like the two previous examples in the classroom and in the office. The presence of an external localisation estimate is ignored in order to assess the Localisation performance.

The data from the iteration 2 is presented in Table 9. The number of valid pairs varies from 1547 to 1684, depending on the number of undefined simulated scan samples (on the bottom and right of Figure 23 there are some undefined boundaries).

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	7.834	2.243	0.730	0.0794	608	0.0502	✓
p_1	1.741	20.443	-0.832	0.0739	493	0.0560	✓
p_2	24.959	26.838	-2.412	0.0613	502	0.0590	✓
p_3	7.837	1.972	0.725	0.0569	552	0.2242	✗
p_4	25.302	20.578	-0.827	0.0471	43	0.0850	✗
p_5	7.498	8.522	2.315	0.0471	146	0.0709	✗
p_6	7.864	20.637	0.724	0.0471	153	0.0485	✗
p_7	7.477	18.907	2.313	0.0415	115	0.1146	✗
plus 129 clusters with a total weight of 0.5458							

Table 9 – Iteration 2: Frame Localisation results with Likelihood Test validation

The first relevant notice is the reduced number of match pairs and the large number of posture estimates with similar weight. This is a result of the reduced information conveyed by the laser scan and the environment symmetry.

In Figure 24a, the range data is presented. It is apparent that only the walls at both sides of the AEST provide useful data since the central field of view returns no information about the room topology. In Figure 24b (the grid spacing is 2.0m), the first 8 candidate clusters are presented with the range data superimposed to the first posture estimate, p_0, showing clearly that this is the correct estimate.

Because there is a reduced number of match pairs, the first validation rule must be relaxed and the *MinMatchRatio* must be set to 0.2.

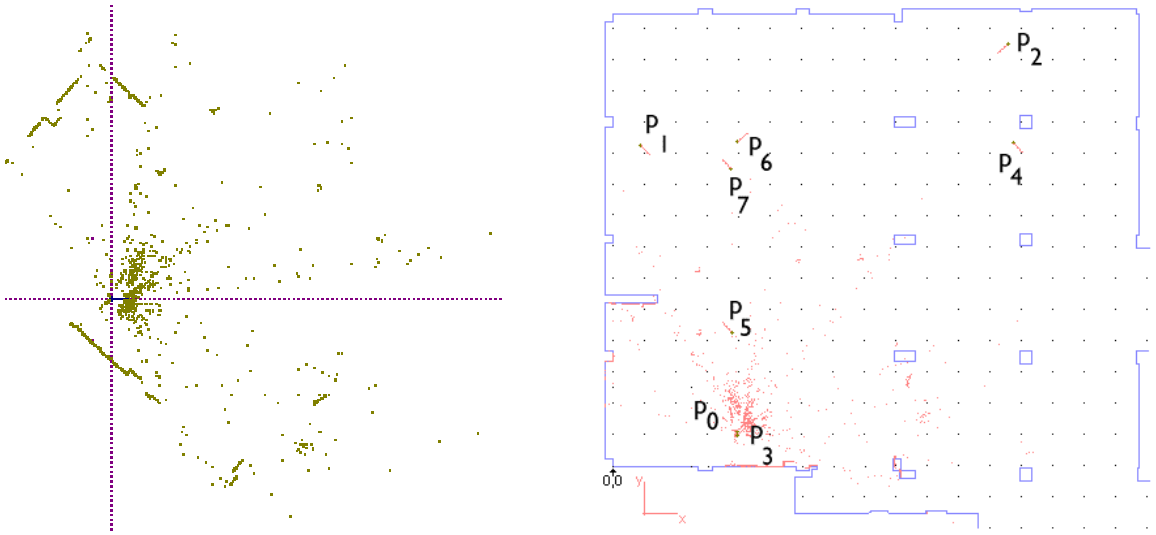

(a) (b)

Figure 24 – Iteration 2: Frame Localisation: (a) range data, (b) localisation clusters

The three valid postures were submitted to Approximate Localisation, using the Error Descent algorithm with an initial step of 0.05m and 0.005rad and 7 iterations. The results are summarised in Table 10. The point-to-point distance histograms for p_0 and p_1 are presented in Figure 25 and the associated postures are presented in Figure 26.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	7.8414	2.2341	0.7225	7.81, -8.59, -7.74	614	0.0414	0.0679
p_1	1.7411	20.4433	-0.8324	-2.5, -140.63, 5	493	0.0345	0.0647
p_2	24.8967	26.8801	-2.4049	-62.5, 42.18, 7.19	469	0.0413	0.0858

Table 10 – Iteration 2: Approximate Localisation results

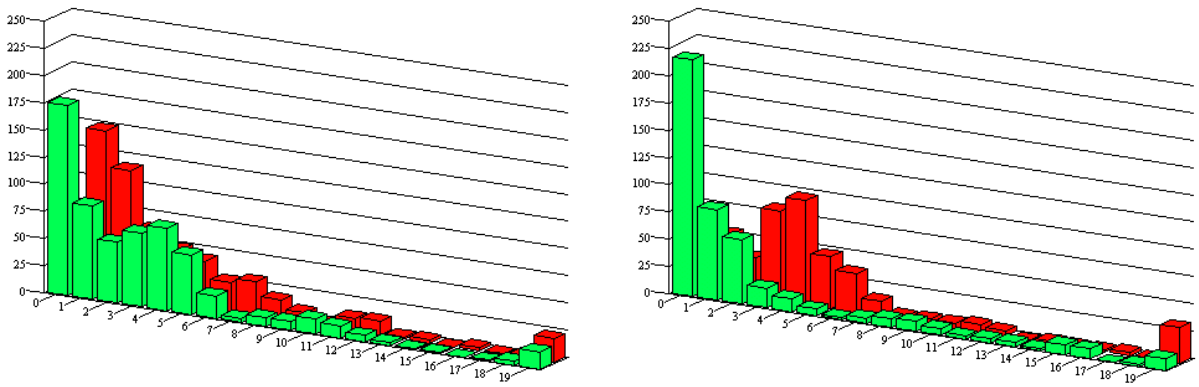

Figure 25 – Iteration 2: distance distribution for p_0 (left) and p_1 (right)

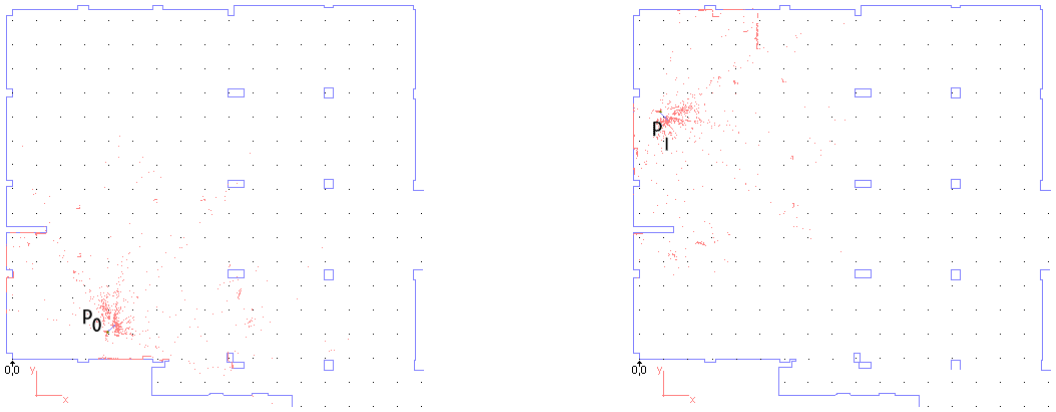

Figure 26 – Iteration 2: refined estimates for p_0 (left) and p_1 (right)

Again, the best posture estimate, p_0, is selected by a combination of number of match pairs and expected value, according to (3.12) and (3.13).

From iteration 2 to iteration 15, the AEST has travelled through the open space towards the central corridor (Figure 22b) and entered a cluttered area where blocks of concrete are manufactured and tested (Figure 22a).

The Frame Localisation results are summarised in Table 11. Figure 27 shows the range data and the two posture estimates.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	21.139	4.863	-0.782	0.0887	393	0.1656	✓
p_1	21.145	5.322	-0.803	0.0437	413	0.1264	✓
plus 348 clusters with a total weight of 0.8676							

Table 11 – Iteration 15: Frame Localisation results with Likelihood Test validation

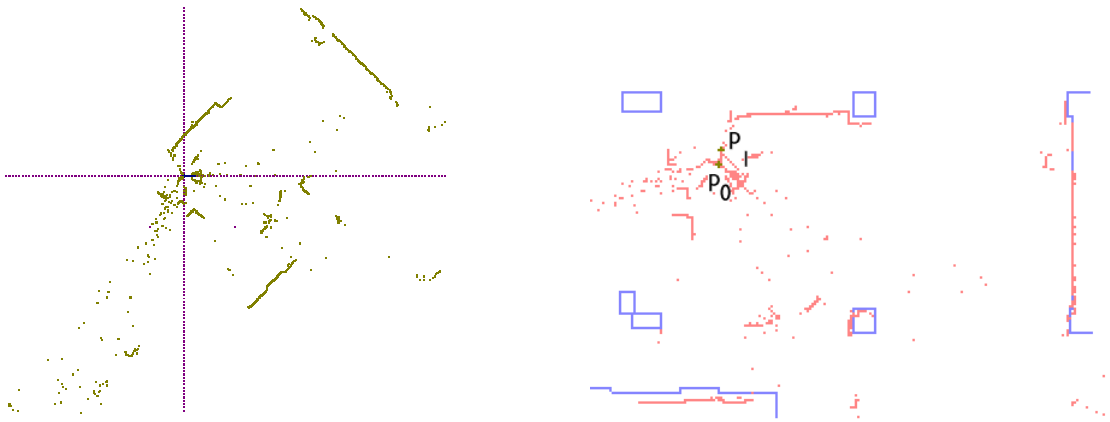

Figure 27 – Iteration 15: Range data (left) and Frame Localisation clusters (detail of map on the right)

The number of valid pairs is 1428. This experiment is similar to the previous: there is a small number of match pairs (less than 30%) and a large number of laser samples that are not related to the map features. The laser scan comprehends two walls at perpendicular angles (right and bottom of Figure 27 right) that are used to define the frames matches. Additionally, there are objects around the AEST and a wall between the two pillars on top. This wall does not fit the map and is responsible for a large fraction of the clusters detected (348) with a total weight of 87%. The two validated clusters show little weight and large expected values as a result of the poor data definition.

The two candidates are almost parallel: the main difference is along the y-axis, where they differ by 0.46m. Submitting p_0 and p_1 to Approximate Localisation with the parameters used on iteration 2 yields the results summarised in Table 12. The two candidates are refined until they stabilise at a common posture: the difference between the refined postures is (18.6mm, 2mm, 4.3 x 10^{-3}rad).

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	21.1949	5.1857	-0.7927	39, -135.9, 10	449	0.0556	0.0989
p_1	21.1763	5.1877	-0.7884	37.5, 325, -6.25	452	0.0571	0.1041

Table 12 – Iteration 15: Approximate Localisation results

Although the two estimates are very close, they do not coincide due to two factors: since the step used is a rational number, an infinitesimal step is required to match precisely the other candidate. However, the main factor is the local minima in the Error Descent algorithm. Since the pairs with large point-to-point distance represent a significant fraction of the expected value (approx. 40%), moving a small step may discard large distant samples, thus creating a local EV minimum.

The distance distributions and the refined postures are presented in Figure 28. The large displacement of p_1 reduced the samples in the last slot from 173 to 28, emphasising the convenience of using a large match threshold, T_{Map}.

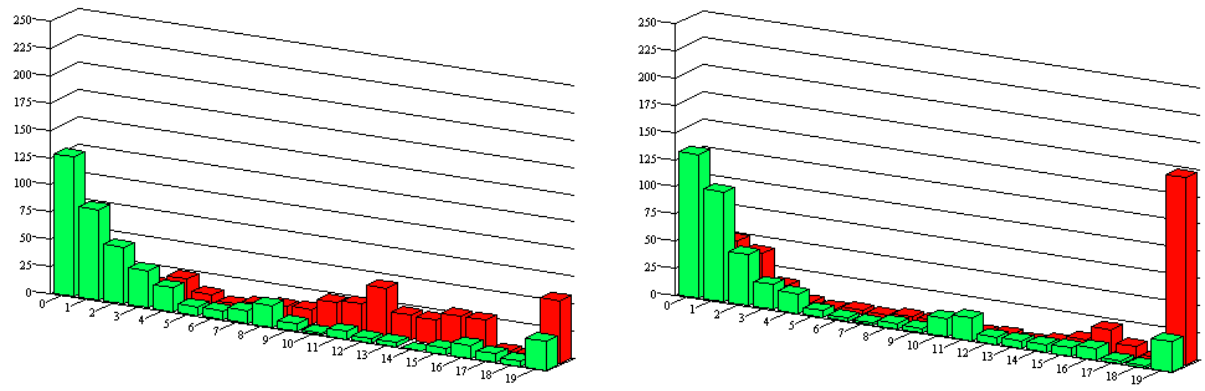

Figure 28 – Iteration 15: distance distribution for p_0 (left) and p_1 (right)

The best posture estimate is presented in Figure 29. The posture estimate correction is visible when compared to Figure 27.

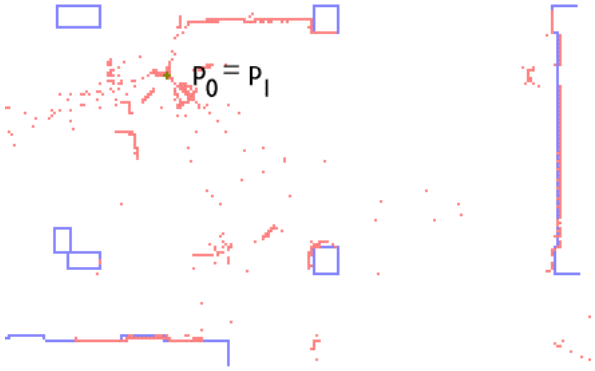

Figure 29 – Iteration 15: best refined estimate (p_0), map detail

When iteration 23 is reached, the AEST has travelled to the gate at the far right of the map and turned around to face the "factory" interior. The mosaic in Figure 30 illustrates the view from the gate, when the AEST is heading approximately 2.5rad (140º). This image is composed by the individual snapshots captured with the video camera used for texture mapping.

Figure 30 – The "factory" viewed from iteration 23

The range data captured at iteration 23 is transformed in laser scan lines (Figure 31a). The scan lines comprehend two walls in perpendicular directions and some minor features associated to the pillars and other obstacles in the "factory." However, one of the main walls are fractionated in small segments, due to the large distance from the laser sensor (14m to 18m) and obstacles in the way, thus hampering Frame Localisation.

The Frame Localisation results are summarised in Table 13 and the candidate postures are shown in Figure 31b. The number of valid pairs varies from 1665 to 1514. The laser scan also represented, centred on the preferred posture, p_0.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	30.324	14.598	2.320	0.0791	627	0.0614	✓
p_1	2.474	21.508	-0.817	0.0482	492	0.0705	✓
p_2	25.808	26.134	-2.388	0.0465	553	0.0494	✓
p_3	30.338	23.638	2.324	0.0439	464	0.0818	✗
p_4	4.962	25.820	-2.388	0.0408	447	0.0753	✗
p_5	2.452	5.477	-0.812	0.0408	452	0.0826	✗
plus 314 clusters with a total weight of 0.7005							

Table 13 – Iteration 23: Frame Localisation results with Likelihood Test validation

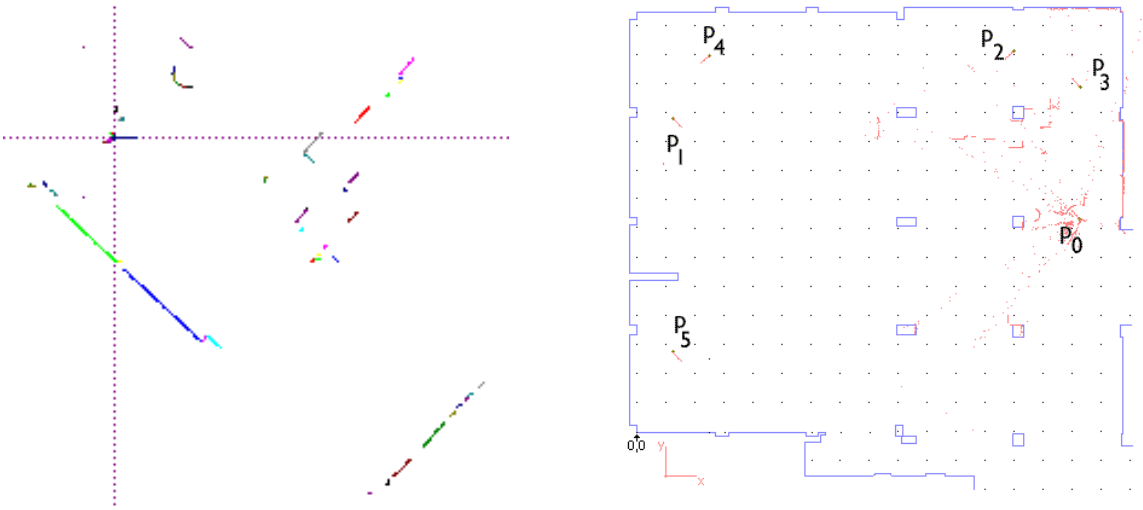

(a) (b)

Figure 31 –Frame Localisation at iteration 23: (a) laser scan lines, (b) localisation clusters

The algorithm's performance is similar to the previous iterations: many solutions and, although the preferred solution is the correct one, there are other candidates almost as good. The line split results of irregular and poor reflectance due to using the laser close to its operation limits. This adds a new strain on Frame Localisation, augmenting the number of frame matches required to identify a pattern. It should be noticed too the large number of small objets standing in front of the AEST (Figure 31a) that are not present in the map (Figure 31b). Although these are ignored by Localisation, they reduce the field of view to detect the features behind them that could be referenced in the map.

Should the Likelihood Test use the laser reflectance information, the relevance of the distant samples would be reduced regardless of the pair-to-pair distance. This is of little importance in case there are reliable samples closer to the AEST but here,

Localisation is based primarily on the distant walls and the far walls are the reliable features in spite of the poor reflectance.

The valid postures are submitted to Approximate Localisation for refinement, using the same parameters as before. The results are presented in Table 14 and the point-to-point distributions are shown in Figure 32.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	30.3460	14.6602	2.3284	21.87, 62.5, 7.97	641	0.0432	0.0803
p_1	25.7896	26.1498	-2.3818	-36.03, 17.72, 5.94	554	0.0436	0.0838
p_2	2.46136	21.5139	-0.8122	22.76, 6.25, 4.68	487	0.0659	0.1273

Table 14 – Iteration 23: Approximate Localisation results

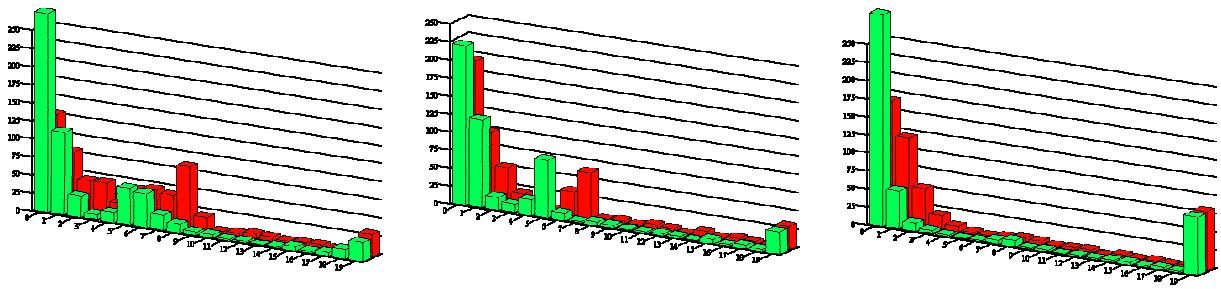

Figure 32 – Iteration 23: distance distribution for p_0 (left), p_1 (middle) and p_2 (right)

There was a substantial refinement, since the match pairs (MP) increased, the distributions were "pushed to the left" and, consequently, the expected value (EV) was reduced. A qualitative analysis of the distance distributions shows how the algorithm adapted to the pairs that do not match exactly. In the refined p_0 and p_1 estimates some samples remain in the 6 to 8cm slots while in p_2 these were either ignored (reduced number of valid pairs) ore stored on the last slot. Looking at the moments, the difference between the correct (p_0) and the false candidates (p_1 and p_2) is not clear. This experiment emphasises the limitations due to the laser measurement range: few reliable samples and a symmetric environment lead to relaxed validations and consequently many false candidates appear.

Second session

The second session consists of 16 iterations, starting from the large metal gate in the middle of the right wall. It starts close to the location where the first session ended. The posture was initialised with a measuring tape. At the initial posture, the Frame

Localisation fails due to lack of solid walls within the laser field of view. Two iterations are presented: iteration 8 at the middle of the central corridor and iteration 10 when the AEST enters the top-left section of the "factory".

The Frame Localisation results for iteration 8 are summarised in Table 15. The number of valid pairs (NVP) varies from 1684 for p_0 to 1572 for p_1.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	15.966	9.340	-3.087	0.0522	472	0.1213	✓
p_1	8.884	12.246	1.617	0.0451	236	0.1498	✗
p_2	16.746	19.716	0.046	0.0451	174	0.1824	✗
plus 255 clusters with a total weight of 0.8575							

Table 15 – Iteration 8: Frame Localisation results with Likelihood Test validation

The range scan is shown in Figure 33a and the candidate clusters are shown in Figure 33b. It is apparent from the images and the data in Table 15 that the results are below any reasonable confidence threshold. The weight threshold for Frame Localisation is set at 4% of the total weight and the preferred solution "weighs" only 5.2%, followed by two other candidates with 4.5% and 255 candidates with 86% of the weight. Moreover, the number of match pairs for the preferred solution is only 472 (28% of NVP) and the expected value of the distance distribution is 0.12m. The two remaining candidates are eliminated according to rule 1 (insufficient number of valid pairs).

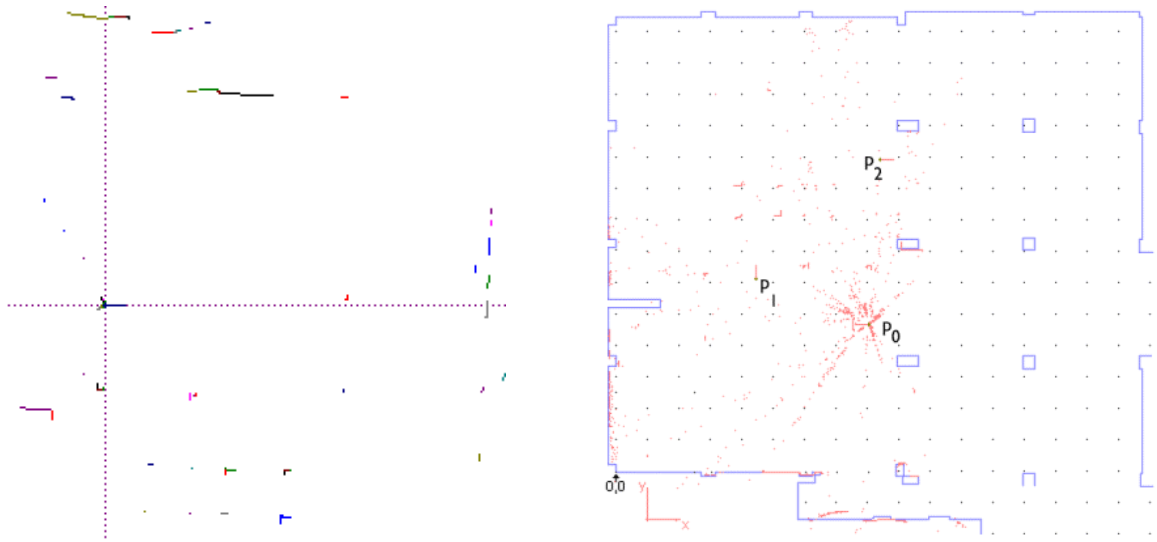

(a) (b)

Figure 33– Iteration 8: (a) scan lines, (b) Frame Localisation solutions

Submitting p_0 to Approximate Localisation with the parameters used in the first session produces a significant refinement (Table 16). The point-to-point distance distribution is shown in Figure 34a (notice the different scale from previous iterations) and the refined posture estimate is presented on the map with the laser scan (Figure 34b).

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	15.8724	9.4086	-3.0805	-93.75, 68.75, 6.25	480	0.0766	0.1325

Table 16 – Iteration 23: Approximate Localisation results

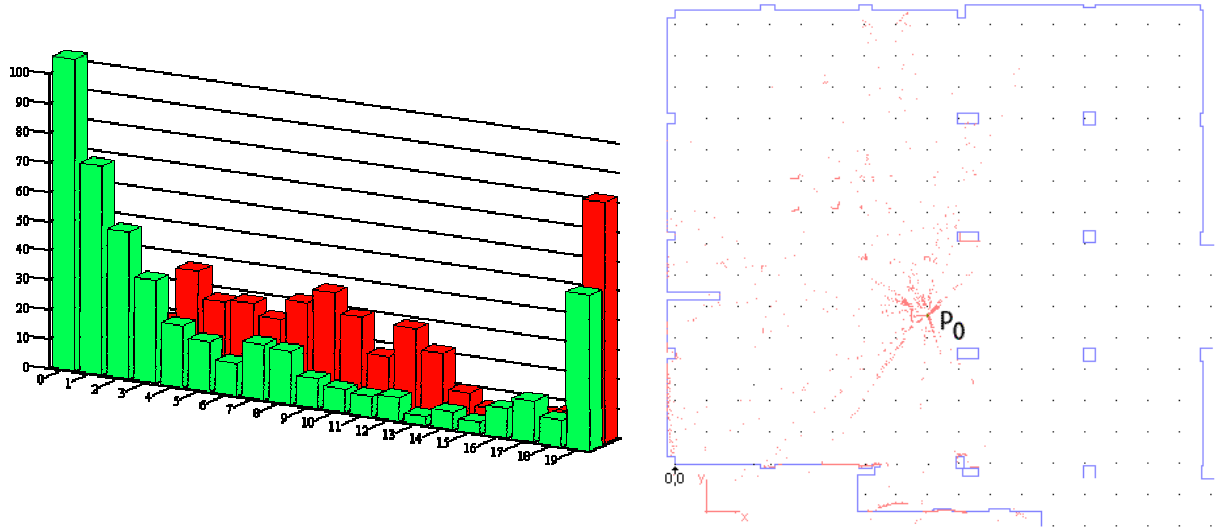

(a) (b)

Figure 34 – Iteration 8: (a) distance distribution, (b) approximate localisation

The expected value is reduced by 0.045m although many samples lay in the middle slots in the distribution (Figure 34a). By carefully looking at Figure 34b one detects some samples in the pillars and the bottom wall too far to the right in order to align the front wall samples to the map. Since the wall stands at 16 to 20m away from the AEST, the laser range error increases, but again the reflectance data was not used due to the limitations in the map description.

After travelling along the central corridor, the AEST reaches iteration 10. From this point onwards, it enters a new working area, where complex structures of brick and mortar are built and tested. In addition, there are large extensions of walls hidden by metal cabinets.

The results of Frame Localisation with Likelihood Test validation are summarised in Table 17. The number of valid pairs is 1641. The scan lines are shown in Figure 35a and the candidate postures are drawn over the map in Figure 35b. The two candidates,

p_0 and p_1, are almost parallel, separated by 0.52m in the y-direction.

Frame Localisation results					Likelihood Test validation		
posture	x [m]	y [m]	θ [rad]	weight	MP	EV [m]	valid
p_0	7.271	12.417	2.296	0.0929	467	0.1039	✓
p_1	7.276	11.893	2.3	0.0402	249	0.1029	✗
plus 272 clusters with a total weight of 0.8669							

Table 17 – Iteration 10: Frame Localisation results with Likelihood Test validation

The main difference is the number of match pairs. Although the number of match pairs is low on both cases, p_1 is below the threshold $MinMatchRation = 0.2$.

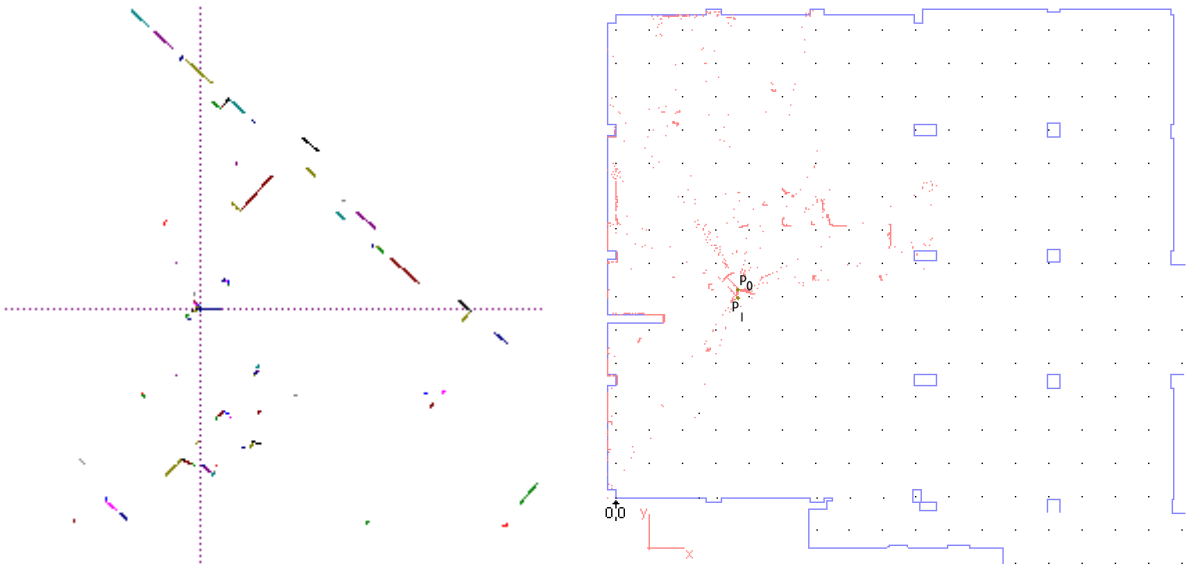

(a) (b)

Figure 35– Iteration 10: (a) scan lines, (b) Frame Localisation solutions

Submitting p_0 to Approximate Localisation using the same parameters as before, provides the refined estimate indicated in Table 18. The point to point distance distribution is illustrated in Figure 36a and the refined estimate drawn over the map is shown in Figure 36b.

	x [m]	y [m]	θ [rad]	correction [mm, mm, 10^{-3}rad]	MP	EV [m]	$\sqrt{\text{Dispersion}}$ [m]
p_0	7.2871	12.5447	2.2974	15.63, 128.13, 1.25	424	0.0767	0.1078

Table 18 – Iteration 10: Approximate Localisation results

Although there is some refinement (Figure 36a) the main cause for the enhancement of the expected-value is the rejection of 33 samples on the last slot. These 33 samples account for account for 21% of the expected value.

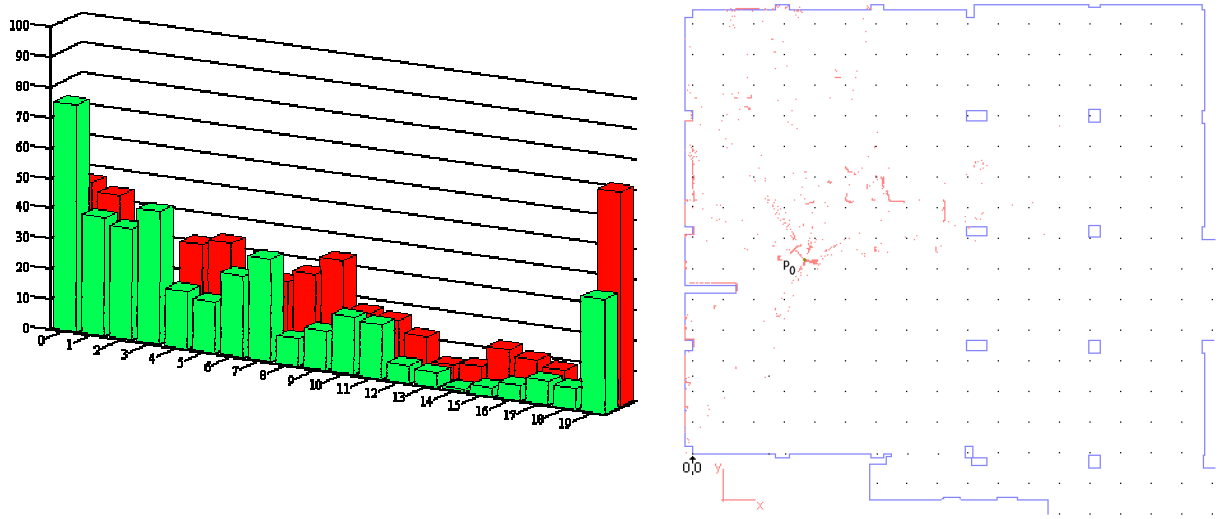

(a) (b)

Figure 36 – Iteration 10: (a) distance distribution, (b) approximate localisation

The distribution shows the matching difficulties: few matches and low quality matches due to differences between map and the scene.

Many more experiments could be shown. In this environment, the Localisation algorithms working autonomously (without using external posture estimates) succeeded in 8 out of 19 iterations in the first session and 7 out of 15 iterations on the second session. The examples presented were selected to highlight different issues and difficulties encountered in the “factory”.

This detailed example of mitigated success experiments intends to highlight the strong points of Likelihood Test and the Localisation under odd or difficult operation scenarios:

1. When the working conditions degrade the Likelihood Test and the algorithms degrade smoothly as well. This is a result of anchoring the algorithms to the few reliable features.
2. The high number of parameters adds flexibility to the Likelihood Test. The defaults suit most indoor environments, alleviating the beginner operator of understanding the core of the algorithm. As he gains experience he will be able to proceed to more difficult scenarios, by relaxing the constraints and learning to read the wealth of information provided in the log files. The log files include all the posture candidates, the diagrams of the postures on the map with their

associated range scan, the point to point distance distribution, the zero-th to second order moments and the cost functions depicting all Likelihood Test steps in detail.

3. The algorithms still perform satisfactory when the operating requirements are not met. In such cases the operator may choose to remove the validation and elimination conditions (relaxing the thresholds) to analyse all candidates and, in case the Likelihood Test chooses a wrong solution, he/she can replace it with the appropriate one taken from the log file.

For the sake of illustration, five sample images of the reconstructed models are shown in Figure 37, where white indicates the non-reconstructed areas.

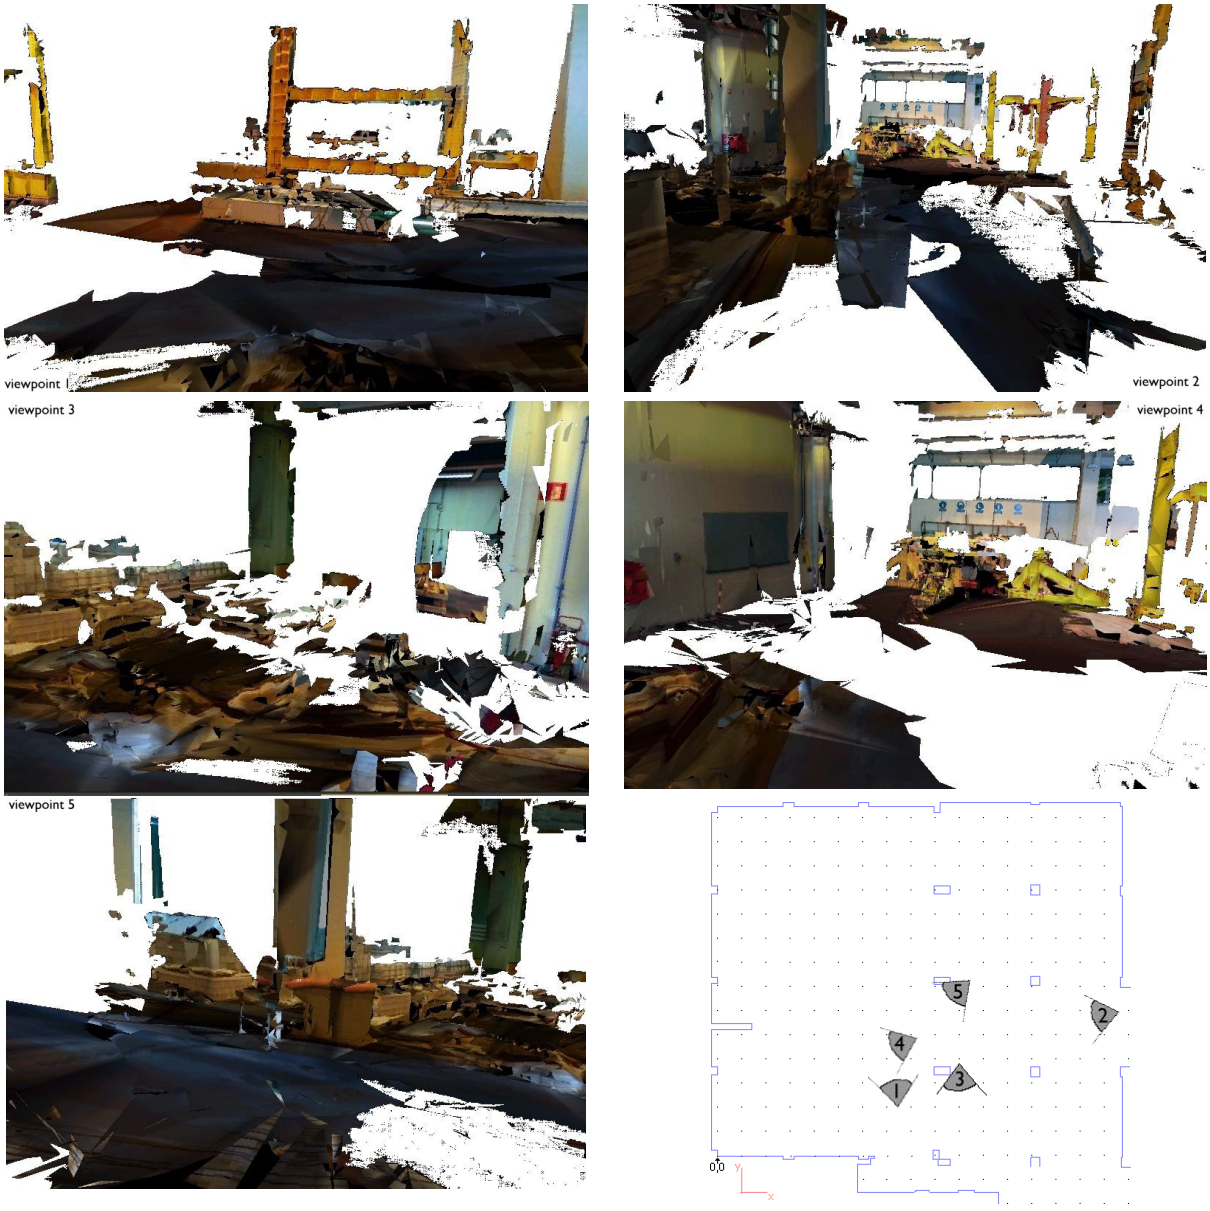

Figure 37 – Five images from the reconstructed model and their relative position

A better laser or a different approach is required to regard this environment as a normal office room. The laser requires a larger measurement range and higher angular resolution to be able to scan from larger distances. A different operation approach would be to scan close to the surfaces; however, this would require removing all the equipment and materials from the scene, or using a large extensible robotic arm.

3.5 Conclusions and extension

Through the Likelihood Test, all candidate postures have been validated with the whole laser scan, which is the most complete set of data available. The result of the test is a best posture estimate associated with an *a posteriori* Localisation error estimate.

However, it may occur that a thorough scan is not possible because the map is incomplete, the scan is incomplete or both are incomplete. Moreover, it has no relevance to measure the distance between one map feature and its assumed counterpart in the laser scan when they are too apart. In most cases, this distance corresponds to a wrong posture but it may also correspond to inaccurate maps or laser scans.

Therefore, it is preferable to measure the number of valid pairs and the number of match pairs to assess *the relevance of a posture estimate* and measure the first and second order moments of the distance distribution to measure *accuracy of the posture estimate.*

The Likelihood Test suits the two proposed Localisation Algorithms, supplying an effective method to assess the Localisation error and merge the position and orientation error into a single coherent measurement, the point to point distance distribution. This is most useful for automatic implementations, where decisions are based on quantitative relations.

The Likelihood Test proved to be very sensitive to minor posture changes, allowing for precision posture updates, although it may be too sensitive to the pairs with higher distances. This disturbance may be minimised by reducing the maximum distance, T_{Map}, or the equivalent value of the last slot, $\Delta_{\max}$.

The Likelihood Test degrades smoothly on poor operation conditions, extending the operation capabilities to new scenarios. The Likelihood Test provides a wealth of information on the experiment conditions, allowing the expert user to understand what's going wrong and what should be done to enhance the results.

The discrete nature of the histograms was created for the benefit of the human operator. In fact, the discrete data is used for visual feedback only, since the Likelihood Test moments restore the continuous nature of data. Therefore, an extension of the Likelihood Test was created, where the moments are based on the continuous data. This variant, termed Likelihood Distance, yields marginally better results because the

equivalent of a histogram slot is its middle point while usually most of the samples in one slot are closer to the lower boundary. However, the greatest interest of Likelihood Distance is the possibility of modulating the weight granted to the higher point-to-point distances avoiding the abrupt threshold that either includes or excludes a high cost sample by a matter of millimetres.

With Likelihood Distance, all the valid pairs are considered, albeit the pairs with higher point –to-point distance are included with very low weight. Thus, the variations in the Likelihood Distance parameters are always smooth, an important benefit for Approximate Localisation.