Chapter 5

Conclusions

The main objective of the research work described in this dissertation was the development of a Localisation methodology based only on laser data that did not require any initial posture estimate. This was achieved with Frame Localisation. In order to access the quality of the posture estimates, the Likelihood Test was developed. The third step in the work described was the development of the Approximate Localisation solutions aiming at refining the posture estimates computed by Frame Localisation.

The three algorithms described met the all the initial requirements and most of the additional requirements that were proposed as the RESOLV project progressed to more challenging operating conditions. The open issues include localisation while the system is moving, localisation associated to map generation based on 2D laser data and localisation in slope planes or uneven grounds.

The future research effort should address the open issues and some other possibilities such as merging the laser data with data from other sensors or extending the current algorithms. Among the interesting extensions are the use of biquadratic curves in Frame Localisation and the adaptive step in Error Descent.

This research work followed a practical approach: the merit of each solution is not associated to its novelty or ingenuity but to its performance during the field trials. Since the Localisation solution was integrated on a real robotic platform expected to work in real world environments, the final appraisal should be granted by the human operators using the RESOLV system and by the partners responsible for the modules that use the estimates computed by Localisation.

Chapter Organisation

In Section One, the main requirements and achievements of this work are listed. In Section Two, the difficulties and open issues are discussed. The directions for future research are outlined in Section Three. Finally, in Section Four, the author elucidates the reader about the reasoning that oriented his research activities as they progressed.

5.1 Requirements and achievements

When the RESOLV project started the specifications for Localisation were based on updated 3D Reconstructed maps and a wide-angle horizontal laser scan (approximately 4.9rad or 280º). The laser scanner featured a typical precision below 0.02m [SequeiraV_1]. The acquisition system should operate in piece-wise horizontal floors. Under these conditions, the Localisation should provide a (x, y, θ) posture estimate with less than 0.1m in position error and 0.017rad (1º) in orientation error. Preferably, the accuracy should be better than 0.02m and 0.0034rad (0.2º). See Chapter 1, Section 3 and Section 4 for details.

During the project development, the experiment conditions changed to increase the efficiency of the acquisition process and the on-line 3D reconstruction was postponed to an off-line phase. Therefore, for Localisation purposes, the maps based on 3D reconstruction data were replaced by maps computed offline, often very incomplete and outdated. In addition, the initial laser scanner was replaced by a new device, which is much faster but less accurate. The construction of the robotic device, the AEST, further reduced the horizontal field of view to 4.4rad (252º).

Notwithstanding all the additional difficulties, the Localisation solutions evolved to match the new operation conditions and the accuracy requirements were matched, if not exceeded. In typical office environments (the classroom and the office in Chapter 3 and Chapter 4), the localisation error associated to the estimated posture, validated with a measuring tape, is below 0.02m. The accurate validation of the orientation estimate requires methods that are more sophisticated.

A second form of validation is obtained from the 3D Reconstruction modules. The environment reconstruction is performed in a succession of stages, all based in 3D laser profiles encompassing a solid angle around the laser. First, the scene visible from each capture point is reconstructed, then it is merged with previously reconstructed models to create a unified model of the environment (see Chapter 1, Section 3 and [SequeiraV_1]). During this process, the map reference is updated and consequently, the system posture relative to the map is updated too. These posture updates are usually below the proposed threshold (0.02m, 0.0034rad (0.2º)).

The Frame Localisation algorithm delivers a posture estimate with a common accuracy better than (0.05m, 0.008rad), based on the map and the 2D laser scan, without any prior knowledge of the system state. The current version is restricted to flat surfaces. It copes with a large majority of unknown or unmapped data, anchoring the data matching on the small fractions in the few areas where the laser scan matches the map. This is apparent in the “factory” experiment where successful Localisation was based on 10% of the field of view only. Another useful feature of Frame Localisation is the generation of a list of weighted posture candidates. In case the first estimate is wrong, the human operator can select the most likely estimate, according to the analysis of the log files and graphical data and input it manually into the RESOLV system.

The line extraction from laser data – described in Appendix C – is an important, albeit autonomous, part of Frame Localisation. The proposed algorithm is preceded by a segmentation stage designed to meet the laser scan characteristics and followed by recursive line extraction based on an iterative form of Least Squares Estimate followed by line merge stage. This sequence aims at maximising the computation efficiency while using the reflectance data to enhance the statistical analysis of the laser data.

One of the first research directions followed by the author was the analysis of the statistical characteristics of the laser and how it propagates throughout the Frame Localisation algorithm. Soon it was realised that the essential contribution to the error characteristics came from the environment, which is very difficult to model with generality. Thus, a major leap was given when the approach based on analysing the error in its standard deviation components, $(\sigma_x, \sigma_y, \sigma_\theta)$ and cross-correlation components was abandoned in favour of the integrated error measure defined in the Likelihood Test.

The Likelihood Test measures the actual difference between the existing map and the laser scan located at the proposed posture estimate. Some cautions must be taken to minimise the effect of missing features in the map or the effect of scan errors, which are not due to the Localisation procedure, but still affect its results.

The Likelihood Test result is characterised by three moments. The measure grants higher relevance to the first two moments (the number of Match Pairs and the Expected Value) with the third moment (Dispersion) being used to disambiguate the cases with similar cost. The moments are intuitive and the histogram distribution provides a good insight on the effect of the updates in the posture estimate.

The author claims that the Likelihood Test, although different from the canonical error estimates, is adequate to the foreseen applications. During the field trials, it proved to be accurate and very sensitive to short displacements. Its main pitfall is over-optimisation regarding pairs with high point-to-point distances, an issue that is minimised to some extent with the Likelihood Distance (see Chapter 3, Section 5 and Chapter 4, Section 5).

The Approximate Localisation algorithms were developed at different times, to meet different operation conditions. While Reference Transform is fast and adequate to regular environments and high quality data, Error Descent requires intense computation while providing reliable estimates. The Error Descent algorithm runs in less than 10 seconds with current PC hardware; it offers a very robust solution, it fits difficult environments such as the “factory” and it always yields solutions that are as good or better than Reference Transform. Therefore, the Error Descent algorithm is chosen as the most adequate solution for Approximate Localisation. Some care should be taken, though, to avoid over-optimisation (see Chapter 4, Section 5).

A lot of effort was put on the human interface of Localisation. Although Localisation is intended for automatic operation, the human operator can manipulate approximately 70 parameters, divided in different sections, to optimise the algorithms' performance. This process is facilitated by numerous figures (GIF images) depicting the various stages of the three modules and a detailed log file, where nearly all the steps taken through the Localisation process are registered and the internal variables are output. Thus, before starting a long campaign in an unknown environment, the operator might perform some Localisation runs to check the quality of the line extraction procedure, matching parameters, accuracy thresholds, *etc.* and/or to adjust the map to the actual scene.

The majority of the figures shown in the dissertation are actual outputs from the Localisation module interface (hence the low definition in some of them). This effort serves a tutorial purpose, too. Other people working over the Localisation module may access and learn the technicalities of the algorithms in detail before actually implementing the code.

The Localisation module was designed to accommodate the widest range of operation conditions. It can be used with 3D reconstructed maps, off-line maps and, if the work proposed in Section 3 is followed, it will be able to build its own maps from 2D laser data and extend off-line maps with self computed maps. It takes advantage of external posture estimates when available (such as the ones provided by odometers) to accelerate the computations and corroborate the results, but requires none to perform correctly.

As a general rule for all the modules in RESOLV project, the data is exchanged by file, which allows for off-line and distributed processing.

The Localisation module currently operates with three different laser devices and can be extended to operate with other equipment. The default accuracy thresholds follow the constraints used in the RESOLV project, but this is not an assumption of performance. The operator defines its accuracy requirements in the configuration file, replacing the default threshold, according to the actual trial conditions and the mission objectives.

5.2 Open issues

Among the many subjects related to Localisation with Laser data, some were considered necessary for the research effort, while others, although relevant, were set aside since they would extend the research beyond the scope of the proposed problem:

1. The localisation with a moving platform was considered an option since the early days of the RESOLV project. Moreover, it is an interesting feature beyond RESOLV. During the field trials, it was verified that the platform movements between acquisition points were short and that large movements were performed

between acquisition sessions and not within the same session. In addition, the first laser scanner used had a very slow pan movement, inhibiting the possibility of creating a wide-angle scan while moving. Therefore, the problem was abandoned in benefit of the extension to the off-line maps and difficult environments.

2. The generation of maps based in 2D laser profiles. The need for this feature became apparent during the project course. Some experiments were made using the laser scan acquired at iteration $k-1$ as a source of map data to be used in iteration k. Extending this principle to the n previous iterations was also experimented, with poor results (see Chapter 3, Section 2). A proper mapping algorithm is required to substitute the clouds of points.
3. One of the most severe constraints in the use of the Localisation solution is the requirement of horizontal ground. Although many environments are strictly flat, it was realised that Localisation can not operate in numerous heritage sites due to slopes in the floor (such as theatres) or large differences in height in large halls (balconies, stairs, etc.). In the current implementation, the height of the acquisition relative to the $z=0$ plane is defined by the user but the system does not perform any measurement to compute this value.

In addition, two questions within the project scope were left unanswered:

4. The Frame Localisation algorithm uses straight lines only, although the 3D surface description defined in the 3D Reconstruction maps includes biquadratic curves too. This option is explained in Chapter 2, Section 2. To extend the current algorithm to the biquadratic features it is necessary to develop the curve extraction parameter, which might consider the straight line as a special case, and to add the curve parameters to the Frame definition to allow proper frame matching.
5. There is not a formal proof that a global minimum error solution is reached in each of the proposed algorithm. The four algorithms yield local minima, although it can be readily validated by visual inspection of the diagrams. On the other hand, the Frame Localisation and Error Descent algorithms provide a list of local minima, sorted by decreasing likelihood, which can be very useful in practical applications.

5.3 Directions for future research

From the author’s point of view, the open issues offer some lines of research, although there are probably more exciting ones.

For instance, extending the Localisation to three dimensions and six variables would be more interesting. Two different approaches are foreseen. One could repeat the process rotated to the XZ plane, instead of the XY plane, or perform true localisation with 3D features. The former should be based on two 2D-laser scans while the latter would require a new pattern of laser data acquisition.

If the condition of using laser as the primary source is abandoned a broad field of possibilities is open. Combining laser with Global Positioning System could enable the RESOLV system to reconstruct the exterior of buildings. Combining laser with stereo video and/or ultrasound would create an additional free space representation to fill in the areas that the laser can not acquire (dark surfaces, detailed textures, windows and mirrors). All these would require major changes in the Localisation system but the author believes that the work herein described could be useful in these new contexts.

The RESOLV project experience taught the author that the lines of research embedded in real world projects are driven by the actual assignment requirements and not by *a priori* expectations.

5.4 History of the Localisation solution

The Localisation module was developed in a server-client spirit, as it is mentioned in Chapter 1, Section 4. The Localisation should adapt to the available conditions and perform as good as possible without adding new constraints. In the end, the horizontal floor constraint was added, but it was found acceptable by the other RESOLV partners.

Beyond this request, the Localisation solution adapted to the requests presented by the other partners: they are the clients; Localisation serves them. When asked if the reconstructed maps could be bypassed, a new version with the off-line map was developed. When the operation was extended to rooms with minimal off-line maps, the possibility of extending the maps with past laser scan was considered, albeit with poor results. When the laser data quality was reduced by the constrained field of view and reduced accuracy in the device, the algorithm accommodated these changes with revised parameters. Finally, when the operation was extended to areas where the extracted features were only a small fraction of the environment, the Error Descent was implemented to operate without models based on the data, but rather operate with the raw data.

The three main algorithms proposed (Frame Localisation, Likelihood Test and Error Descent) were designed with efficiency in mind, not ingenuity or theoretical novelty.

When the problem of laser based localisation was first considered, the author was surprised when he realised that all the solutions found in the literature (the most relevant are mentioned in Chapter 1, Section 4) assumed an auxiliary posture estimate to initiate the algorithms. In spite of the complex appearance, the author felt that the localisation without prior posture estimate ought to be a simple problem. Before diving deeply into sophisticated methods, the author joggled with the laser data, the map structure and some basic geometry principles and thus, Frame Localisation was created, yielding a simple, intuitive and yet powerful solution to the proposed problem.

The Likelihood Test was developed with the same concern for efficiency and simplicity in mind. The concept of the histogram is straightforward, and the associated three moments are also easy to understand. The combination of both, illustrated by the diagrams, provides a deep insight on the algorithm performance, the cause of errors and the sensitivity of the solution to estimate updates computed by the algorithms.

Finally, the Error Descent was considered, then abandoned in favour of a more efficient and elegant solution, and then recovered when the practical constraints of the project required robustness and higher accuracy under adverse environments.

The algorithms were submitted to difficult environments and their performance was accessed by the members of the RESOLV project. Abiding to the server-client spirit, the author’s effort did not cease until they were fully satisfied.