

Increasing the Efficiency of 6-DoF Visual Localization Using Multi-Modal Sensory Data

Ronald Clark¹, Sen Wang¹, Hongkai Wen¹, Niki Trigoni¹, Andrew Markham¹

Abstract—Localization is a key requirement for mobile robot autonomy and human-robot interaction. Vision-based localization is accurate and flexible, however, it incurs a high computational burden which limits its application on many resource-constrained platforms. In this paper, we address the problem of performing real-time localization in large-scale 3D point cloud maps of ever-growing size. While most systems using multi-modal information reduce localization time by employing side-channel information in a coarse manner (eg. WiFi for a rough prior position estimate), we propose to inter-weave the map with rich sensory data. This multi-modal approach achieves two key goals simultaneously. First, it enables us to harness additional sensory data to localise against a map covering a vast area in real-time; and secondly, it also allows us to roughly localise devices which are not equipped with a camera. The key to our approach is a localization policy based on a sequential Monte Carlo estimator. The localiser uses this policy to attempt point-matching only in nodes where it is likely to succeed, significantly increasing the efficiency of the localization process. The proposed multi-modal localization system is evaluated extensively in a large museum building. The results show that our multi-modal approach not only increases the localization accuracy but significantly reduces computational time.

I. INTRODUCTION

6-DoF localisation is a key capability required by autonomous humanoid robots and for human-robot interaction (HRI). Consider, for example, the case where a robot needs to locate and pick-up a cup or mug in a domestic workspace. First, the robot needs to find its 6-DoF location and orientation with respect to the stored map. Once this is achieved, the robot can shift its attention to the last known location of the object, check if it is still there, update its knowledge of the object and then proceed with the task as necessary. Another use-case from the HRI side may involve a user, equipped with a pair of smart-glasses, who can motion his head towards locations or objects of interest and mark these for a robot to attend to. This too requires accurate 6-DoF localisation of the user with respect to a global map. 6-DoF pose estimation, however, is computationally expensive which severely limits its application on resource constrained platforms. A key challenge of operating in real-time using 3D pose estimation is thus to narrow the search through the map to make localisation computationally feasible online. Strategies that exploit side-channel information (eg. WiFi, geo-magnetic field distortions) show a marked improvement over exhaustive matching, but due to noisy sensor information there are many false positives and an overly large candidate subset still needs to be searched.

In this paper, we propose an Efficient Multi-modal Localisation (EM-Loc) system, which leverages multiple modal-



Fig. 1. Localisation against large point clouds comes with a low efficiency and high computational cost making accurate localisation difficult on computationally constrained platforms such as the Nao V5 (with only an Intel Atom @ 1.6 GHz and 48.6 Wh battery). In this paper we reduce the computational load by creating a localisation graph of multiple modalities.

ities to make precise vision-based localisation on resource-constrained platforms more feasible. We exploit the side-channel information and along with an estimated trajectory predict which features will be visible. This intelligent guidance is achieved through a sequential Monte-Carlo process which estimates the posterior distribution of the current location over the nodes which ensures the 3D pose estimation is carried out using only the currently visible features. We show how this dramatically reduces computational time for 6-DoF localisation and achieves high accuracy.

The key contributions of our paper are:

- We present EM-LOC, an accurate localisation system that can be used across a range of platforms
- We propose a localisation policy which exploits side-channel information (WiFi, magnetic) to cut down vision-based localisation time (thus enabling real-time 6-DoF visual localisation)
- We evaluate these contributions on data gathered from a large museum building

The remainder of the paper is organised as follows: Sec. V focuses on the two aspects of the EM-LOC system - graph construction and localisation respectively. Sec. VI presents an evaluation of the proposed system and Sec. VII concludes the paper and highlights ideas for future work.

II. RELATED WORK

In this section we review existing works related to the EM-LOC system and how they relate to our contribution.

Model-based pose estimation: When a 3D model of discriminative feature points is available, such as is obtained

using Structure from Motion (SfM), then the pose of query images can be found using camera-resectioning. This process first establishes feature matches between the image and the 3D model by either 3D-to-2D [26] or 2D-to-3D [21] matching. The camera pose is then found by using RANSAC in combination with a PnP algorithm [14]. Camera localisation techniques where direct matching to the 3D model is performed are generally very computationally expensive due to the exhaustive matching that needs to be carried out between the many local feature descriptors. To reduce the matching cost, heuristic techniques have been proposed based on image retrieval methods [10] to find images with 3D points that are likely to be successfully localized against as well as methods that use mutual visibility information to find matches [15]. Although these methods make strides towards reducing the computational burden of large-scale localisation in point clouds, they do not completely solve the problem due to the large decrease in successful localisations (i.e. they trade localisation efficiency for localisation frequency). Our proposed method is orthogonal to these approaches; we use multiple modalities to increase localisation efficiency with negligible decrease in the number of successful localisations.

Topometric localisation: The act of localisation does not require exact measurements of a robot's pose or position on a map, rather it suffices to estimate the robot's location relative to a set of predefined points. These points take the form of nodes in a graph which represent distinct locations or places in the environment and edges which connect neighbouring locations [24]. Examples of such systems include CAT-SLAM [17] and the real-time topometric localisation system by Kanade et al. [1]. The real-time topometric localisation system works in two phases; in the mapping phase the robot drives around the workspace equipped with a GPS and visual sensor. A graph is then created with nodes at fixed interval with each node storing 3D features garnered from the available sensors and stored in a database. In the localisation phase, a Bayesian method is used to estimate the robot with respect to the stored database using the available sensors. Experience based navigation (EBN) [3] extends this formulation by allowing the robot to collect multiple, disconnected segments of the graph during on-line operation. Each segment (called an experience) can be collected during different times of day, weather conditions or seasons and each segment consists of stereo frames linked by edges with odometry information. These segments are then stitched together, allowing the robot to localize under a range of environmental conditions. In this paper, we adopt the graph-based structure of topometric localisation, by using a map of nodes, edges and 3D points. However, we keep fully metric relations between the nodes which allows for accurate localisation. This structure allows us to partition the monolithic model of 3D points, while still being able to index it for fast localisation.

Place recognition and appearance-based localisation: Appearance-based localisation methods like FAB-MAP [4] and others [19] recognise places purely based on the appearance of the image. In most cases, the image appearance

is described by adopting the bag of visual words (BoVW) model. This process is similar to the large-scale image-retrieval problem. In classical image retrieval, however, the aim is to find as many relevant database images as possible, while appearance-based localisation methods aim at finding co-located images in the database. Thus, approaches such as FAB-MAP are characterised by incredibly high precision with typically low-recall rates. By integrating temporal links and considering places as sequences of images, as is done in SeqSLAM [18], the recall rate of appearance-based localisation methods can be significantly improved. Appearance-based methods can also be used for localizing in model-based maps by rendering synthetic views at a large number of candidate locations and orientations. This approach of rendering synthetic views is used, for example, by [10] as a pre-processing step to speed up the direct-matching process in localisation using an SfM model. Because of their efficiency and robustness, appearance-base recognition methods have been popular in mobile phone localisation systems [28] in order to meet the strict resource constraints. Appearance-based methods are also used for loop-closure detections and re-localisations on robotic platforms as they allow one to globally localize against maps of a very large scale [13]. These approaches, however, have the disadvantage of not providing a 6-DoF pose and being far less accurate compared to model-based localisation systems based on local feature point matching. In our system we consider appearance-based features as one of the side-channels (along with WiFi and geo-magnetic measurements), which are used to find model sections against which to localize. Like SeqSLAM, we utilize temporal/sequential information for global localisation.

Multi-modal localisation: The promise of exploiting multiple modalities for aiding vision-based localisation has been demonstrated in a variety of ways, but not directly for reducing visual processing time. For example, the W-RGBD system [11] uses WiFi in combination with RGBD images to perform localisation using an architectural floorplan. They model the WiFi signal strength across the floorplan using a Gaussian process and through the use of this model intelligently initialize the particles for Monte Carlo localisation (MCL). They show that this improves the convergence rate and final accuracy of the position estimate. Their goal differs from ours; while they are mainly concerned with using WiFi to increase the convergence rate and global accuracy of MCL we are interested in using the side-channel information to guide the expensive 2D-3D feature matching to increase the efficiency of the RANSAC-based pose computation and thus reduce the computational time required on resource-constrained platforms.

III. PROPOSED SYSTEM

Our system consists of a two-tier architecture, as shown in Figure 2. The back-end graph assembler is responsible for creating the localisation graph, while the front-end runs on the target platform and performs on-line localisation. The back-end graph assembler receives as input the data collected from mappers during the mapping session of the system

which are uploaded to the cloud where the server is implemented. The graph assembler constructs the localization graph by merging data traces from different mappers. The localiser carries out state estimation to position the humanoid with respect to the assembled localisation graph.




Fig. 2. EM-LOC consists of a two-tier architecture. The localisation graph is created and updated in the cloud, while the localisation module receives the localisation graph and performs on-line state estimation.

IV. STAGE 1: MAP CREATION

We use three types of information in our map creation phase - RGB data from which sparse SIFT features [16] are extracted and associated with 3D points (these are needed to perform the 6-DoF localization), depth frames (which are used to determine the visibility of the sparse feature points from nodes) and the multi-modal side-channel information. Using this data a map of 3D points associated with SURF features and the relative translation between frames is created using an RGB-D SLAM algorithm [7]¹ The localization graph in our case is a set of nodes, n_1, n_2, \dots, n_N (each frame is considered as a node) which are joined by a set of directed edges $e_{i-1,i}$ that link the nodes. Each edge contains odometry data consisting of the displacement and orientation (computed by the RGB-D SLAM) which describes the metric relation between the nodes. Each edge also contains the side-channel information which includes the distribution of WiFi RSSI values along that edge, geo-magnetic distortion measurements and appearance-based visual features.

The structure of the side-channel features is as follows:

WiFi: The readings are composed of the received signal strength, s , for each access point observed in the environment. As described in [22], the RSS of typical radio-based measurements does not change very rapidly with typical humanoid/human mobility and thus the distribution is collected over a finite duration of time which is used as the feature, $E_i^{wifi} = [s_1, s_2, \dots, s_n]$

Magnetic: Similarly, we capture the magnetic field strength data over a short time segment. The magnetic data consists of the magnitude, m , of the geo-magnetic field calculated from the 3 components in the x, y and z axes, $E_i^{mag} = [m_1, m_2, \dots, m_n]$

Image: We further capture images along each experience and use appearance-based features derived from these images as additional side-channel information. The full images themselves are not stored, rather we detect and compute the distinctive visual features which are present in the image. Each feature is then quantized into BovW

¹If a depth camera is not available, this process could also be carried out using an image-based Structure-from-Motion process followed by a dense multi-view stereo reconstruction.





Fig. 3. Illustration of a 3D model for localization with millions of candidate feature points (green dots). In this paper, we propose EM-Loc (Efficient Multi-modal Localization) that integrates these features in a localization graph and interweaves it with additional (side-channel) sensory data. Our localiser uses this side-channel information to predict which features will give successful visual localization. This cuts down the matching time and enables real-time and fine-grained localization over vast areas.

vector, where the word vocabulary, w_1, w_2, \dots, w_n has been precomputed using k-means clustering. The final representation of the image consists of a histogram of the word occurrence counts $E_i^{img} = [h_1, h_2, \dots, h_n]$ As the quantization incurs additional cost, the inclusion of appearance-based features is entirely optional in our system, and if these features are used the vocabulary is kept very small.

The graph consists of data traces, which may be collected during different times, allowing the map to be readily extended. Many disjoint data traces may be generated by mappers traversing an area. The localisation graph thus consists of short segments, along with their nodes and edges, which need to be linked together at certain co-location points. The graph assembly process harnesses the side-channel information to establish these co-location links across multiple data traces. Given a trace denoted as E_i^n where n is the modality, i is the trace index and t is the timestamp at which the data was recorded, a co-location link established across multiple traces when

$$\exists t_2 \text{ s.t. } \max P(i E_{t_1}^n | i E_{t_2}^n) > \tau. \quad (1)$$

This means is that for each data frame of a certain modality that is assigned to a place, some other modality, already belonging to that place shares a strong similarity with its partner in the current trace. After the co-location links have been established the entire localisation graph is optimized using a robust pose graph optimization through the GT-SAM library [5] with the Vertigo extension [23]. The graph construction process is carried out off-line and

also updated using data collected during additional mapping sessions. The end result of this process is a consistent representation of the world in the form of a linked multi-modal localization graph.

Once the graph has been created and optimized it consists of a large number of nodes each associated with a small number of feature points which represent those that were visible in the frame during the mapping process. However, once the consistent map has been created, many feature points may be visible from a node which is not taken into account by the initial mapping process. We therefore utilise a post-processing step to determine the true point visibility for each node. This is achieved by raycasting from the node to the 3D location of each SIFT feature stored in the map and checking for ray intersections with depth frames registered to their world locations. If no hit occurs, the SIFT feature is visible from that node and is added to the node's feature-set. This process is illustrated in Fig. IV. Each node is associated with a set of 3D features which comprises those features that are visible from the node's location.




Fig. 4. In order to determine the visibility of features from each node location, we make use of the depth frames. The visibility is computed by ray-casting from the node center to the stored 3D location of the feature.

V. STAGE 2: LOCALISATION

In this section we describe our localiser which runs on the humanoid with the objective of providing efficient 6-DoF pose estimates. We adopt a sequential Bayesian approach to perform localisation in the established localisation graph.

$$p(n_i, \mathbf{X}_i, \Theta_i | z_{i:t}) = p(\mathbf{X}_i, \Theta_i | n_i, z_{i:t})p(n_i | z_{i:t})$$

Where n_i are the nodes of the localisation graph as defined before, Θ_i is the orientation (yaw, pitch, roll) of the device, and \mathbf{X}_i is the location of the humanoid *with respect to* the current node n_i . The pose estimates relative to a node are independent of the node, the state estimation problem reduces to

$$p(n_i, \mathbf{X}_i, \Theta_i | z_{i:t}) = p(\mathbf{X}_i, \Theta_i | z_{i:t})p(n_i | z_{i:t})$$

In order to perform localisation, we use a particle filtering approach. This enables us to employ a Rao-Blackwellized approach to performing the localisation on the localisation graph where the term $p(n_i | z_{i:t})$ keeps track of the global localisation of the humanoid in the localisation graph, while the term $p(\mathbf{X}_i, \Theta_i | z_{i:t})$ ensures that the local 6D pose computations remain consistent. In order to ensure that the pose computation remains lightweight, we make the assumption that the local pose estimate produces a uni-modal Gaussian. Our particle-filter is initialized and updated as follows.

A. State and initialization:

The particle structure is split into two components which consists of an estimate of the humanoid's current location in the localisation graph represented by the node n_i and $(x_i, y_i, z_i, \theta_i)$ which is the 6-DoF pose of the humanoid relative to node n_i .

$$\text{Particle: } X_i = [\{n_i\}, \{\mu_{\mathbf{X}_i, \Theta_i}, \Sigma_{\mathbf{X}_i, \Theta_i}\}]$$

We uniformly initialize the particles over the entire localisation graph.

B. Measurement

In order to utilise the incoming sensory measurements for localisation in our localisation graph, we need some means of relating the measurements to the data stored in our graph. This is achieved by utilising likelihood functions for each modality as a function of the system state. In particular, we use the following likelihood functions:

RSSI: For RSSI data, the Kullback-Liebler divergence, D_{KL} , is used as the metric to compare measurements coming from individual access points.

$$p(E_i^{wifi} | E_j^{wifi}) \propto e^{-D_{KL}(E_i^{wifi}, E_j^{wifi})}$$

Magnetic: The magnetic features consist of a series of geomagnetic measurements strung together in a series. As the humanoid may have been moving at different speeds between the stored and query feature, we use dynamic time warping (DTW) as a measure of their similarity

$$p(E_i^{mag} | E_j^{mag}) \propto e^{-DTW(E_i^{mag}, E_j^{mag})}$$

Image: To compute the likelihood of a new image given one in the stored experience we use the FAB-MAP approach [4]. FAB-MAP takes as input the image's BoW vector and produces a normalized likelihood measure for each stored location, $p(E_i^{img} | E_j^{img})$.

We compute the weight by assuming that each of the side-channel likelihoods is independent. The importance weights, w_i^i , for the particles are thus calculated as the product of the individual terms

$$w_i^i = p(E_i^{wifi} | E_j^{wifi})p(E_i^{mag} | E_j^{mag})p(E_i^{img} | E_j^{img})$$

C. Odometry

For the process update of our sequential Bayesian filter, we use odometric measurements. Many humanoid odometry estimation schemes rely on incorporating domain constraints for a specific application, along with a complex dynamics model [25] and state-estimation methods. In this paper, however, we are interested in the case where the IMU is not firmly attached to the body of the subject being tracked with the goal of allowing our approach to be transferred seamlessly from a humanoid to human (eg. human wearing a pair of smart-glasses). In this case such strict and clear-cut assumptions does not exist and thus, rather than relying on a state-space approach and encoder data, most methods extract features from the IMU traces and use these these for odometry estimation in what is know as pedestrian dead reckoning (PDR) [8]. PDR performs 3 main operations; step detection, step-length estimation and heading calculation. By coupling the stride length estimates and step event detections derived from these features, the PDR is able to roughly estimate the relative position of the subject. Step events are detected as maximum of the z-axis accelerometer data and the heading is tracking using the the gyroscope and magnetometer [12]. We use a simple step-frequency which has been shown to work well in practice [20]. This model relates the step length to the step frequency, $l_s = \alpha f + b$, where f is the step frequency and α and b are constants as described in [20].

D. Pose Estimation using 3D features

6-DoF pose estimation using 3D features is performed using the EPnP algorithm [14]. Using a set of 2D features detected in an image at locations, u_i, v_i matched to 3D features x_i^N, y_i^N, z_i^N stored relative to the coordinate system of node n . The pose, \mathbf{X}, θ , of the the camera which recorded the image relative to the world coordinate system, is computed using a PnP solver. We use the standard EPnP solver wrapped in a RANSAC loop as implemented in OpenCV [2].

E. EKF Process and Measurement Update

Unlike the side-channel information which gives inexpensive, but noisy and outlier-prone estimates of the current location, the 6-DoF pose computation gives outlier-free measurements. An internal EKF is used to keep track of the pose mean $\mu_{\mathbf{X}_i, \Theta_i}$ and the covariance $\Sigma_{\mathbf{X}_i, \Theta_i}$, relative to a particular node. The process update of the EKF is carried out by using odometric measurements.

In particular, process update is carried out as follows

$$\begin{aligned} x_n^t &= x_n^t + l_s \times \cos(\theta_\phi) \\ y_n^t &= y_n^t + l_s \times \sin(\theta_\phi) \end{aligned}$$

Where the step lengths l_s and step detections are obtained from the PDR. For the EKF measurement, a 6-DoF pose computation is done using the 3D features, this yields an observation $z = \hat{x}_i^N, \hat{y}_i^N, \hat{z}_i^N, \Theta_i$ relative to the 3D node, n_i , to which those features belong. Using this measurement, the Kalman gain is calculated using the standard gain equations

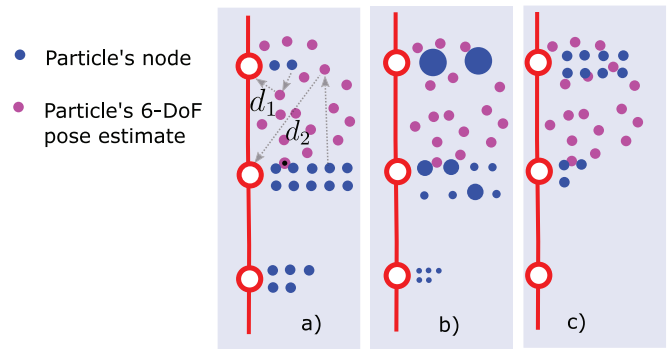


Fig. 5. a) Before particle re-weighting many particles contain 6-DoF pose estimates which are far from their current graph node. b) After weight modification, particles with with nodes far from their 6-Dof estimates are given a lower weight. c) After re-sampling more particles have closely-related node and 6-DoF pose estimates.

[9] and the mean and covariance updated using this gain. We concurrently perform the measurement update for the weights of each particle sharing the same node n_i , which significantly reduces the computational effort.

F. Update and resampling

The weights consider only the likelihoods of the particles at the discretised node-level. However, our particles contain more information i.e. in the form of a continuous 6-DoF pose estimate (maintained by the EKF). As the dead reckoning and 6-DoF pose estimate may lead to the particle deviating from the current nearest node and coming closer to another one, we further adjust the weight of the particles such that those closer to the predicted pose have a greater weight.

$$w_k^{i'} = w_k^i e^{-d}.$$

The quantity $d = \sqrt{(x_k^{i'})^2 + (y_k^{i'})^2 + (z_k^{i'})^2}$ represents the distance of the particle's 6-DoF pose estimate to the location its node i' . The transformation to a particular node's coordinate system is done using the information stored along the edges in the graph. This weight adjustment is illustrated in Fig. 5.

G. Localisation policy

A major bottleneck in using full 3D localization for the purpose of mobile device localisation is the severe limitation which the computational constraints of the mobile device places on the number of 3D features which can be localized against while still allowing for real-time operation. To address this, we use a localisation policy which performs the expensive 2D-3D matching only in nodes where it is essential, guided by the current distribution of particles on the graph.

For reliable localisation, we would like to ensure that 2 criteria are met; 1) the particle distribution should faithfully represent the posterior distribution of the humanoid's location on the graph and 2) we only want to try to localize on a limited number of nodes and these should have a high probability of being the humanoid's actual location.

In order to satisfy 1) we use a metric based on the importance weights obtained from the previous time instant, w_{t-1}^i . These weights are used to calculate the “effective number of particles” $N_{eff} = \frac{1}{\sum \omega_{t-1}^i}$, which is a metric that gauges how well the particles represent the posterior [6].

The localisation policy then uses this metric to determine when localisation should be attempted. If the effective number of particles is high enough, the localiser attempts to localize in the k most likely nodes. Typically, the 3D matching only has to be done on 1 or 2 nodes and each of these frames *naturally* contains only a fraction of the most relevant 3D features needed for localisation. The full localisation algorithm is detailed in Alg. 1 where our contributions compared to standard MCL are highlighted in red.

Algorithm 1 Localizer

Require: localisation graph G , Particles X_i

```

1: function LOCALIZE( $G, X_i$ )
2:   Calculate  $N_{eff}$ 
3:   for  $m = 1 : k$  do
4:     if  $N_{eff} > \tau_{eff}$  then
5:       Select  $m^{th}$  most common node
6:       Pose computation using  $m$ 's 3D features
7:     end if
8:   end for
9:   return  $X_t, \theta_t$       ▷ Proceed with EKF update
10: end function

```

The quantity k is a very important parameter as it determines the number of nodes in which localisation is attempted. It therefore directly affects the number of feature matches that are carried out during the 6-Dof pose estimate. By adjusting this quantity, the number of feature matching can be controlled as required and thus allows the localisation time to be scaled according to the requirements of the platform.

VI. EXPERIMENTS

In order to adequately evaluate EM-LOC, we gauge its performance along three performance axes. The first is the **localisation time**. As the 6-DoF localisation is the most computationally expensive component which we seek to minimize, we define as the localisation time as the time it takes from receiving the current visual frame from the camera to the time that the 6-DoF pose for the frame is determined by the RANSAC-based pose computation. The second factor is the **localisation efficiency**. The localisation efficiency is defined as the number of successfully localised divided by the total number of images for which localisation was attempted. This is important as the RANSAC-based pose computation may not localise a 6-DoF pose if it does not find enough inliers. In our case we use 8 inliers as the threshold for successful localisation. The final aspect we evaluate is the **localisation accuracy** which is simply the




Fig. 6. Visualization of the magnetic and WiFi similarities for the data used in the museum experiment. a) shows the ground-truth metric distance between the nodes b) shows the magnetic signal similarity and c) the WiFi similarity

difference between the system’s current estimate of the angle and orientation and the ground truth pose.

For fair comparison, we compare to two existing methods in terms of both accuracy and processing time. The first system, which we use as a baseline for comparison is the EBN framework [3]. The EBN method stores frames along with 3D feature points against which incoming images can be localised. We consider two variants of the framework. The first carries out an exhaustive matching against 3D points belonging to all available nodes stored in the map - in order to vary the number of nodes that are used in the localisation, we randomly select k nodes from the graph (we label this method A1). In the second method for comparison, we make use of the side-channel data to rank the nodes which are selected (A2). The ranking is computed by using the likelihoods of Sec. V-B as features and using the independence assumption - ranking the nodes as the product of the features in the same way as the weights of the particles are calculated in EM-LOC. Our third method for comparison is based on Travi-Navi [27], which is a state-of-the-art teach-and-repeat pedestrian bi-pedal navigation system from the mobile computing literature. Travi-Navi simply uses an SVM to produce an estimate of the current location from the WiFi, magnetic and appearance-based visual data. For our purposes, we use the SVM to rank the nodes on which the 3D pose estimation is performed (A3).

Our evaluation data was gathered from a large museum area. The area was chosen for numerous reasons; firstly, its size ($\approx 5000m^2$) allows for thorough evaluation of our approach; secondly it is representative of a practical implementation setting of our system - a popular tourist attraction and public site with much commercial interest for a viable localisation service. Furthermore, the museum interior is a very complex space - filled with vast open rooms, narrow corridors and dim-lighting which makes it a challenge for vision-based localisation methods. In order to evaluate such a large area and considering the locomotion speed of current humanoid robots, the data for our experiments were collected by two humans equipped with a pair of Google Glasses. This is reasonable as our odometry does not rely on modelling intricate dynamics. For offline procedure of map creation, RGB-D frames were collected using a Project Tango device and run through an RGB-D Slam algorithm [7]. The same procedure was used to collect ground-truth locations for the online localisation tests. We stress, however,

that all the reported online localisation results are obtained using the data from Google Glass. We test at a total of 167 test locations across the museum. Our experiments were run on an Intel Core i5-2467M @ 1.60GHz, which is similar in performance to that of the Nao v5.




Fig. 7. Comparison between the processing times for EM-LOC system, and comparison methods A1, A2 and A3 (see text)

The error distribution for EM-LOC, A2 and A3 at 3 values of k is reported in the left subplots of Fig. 10 (A1 is not plotted in the graph as the random node decimation performs significantly worse than the other methods at these low k values) and the CDF in Fig. 9. The EM-LOC system clearly outperform methods A2 and A3 in terms of localisation accuracy. This can be explained through two observations. Firstly, EM-LOC, has the ability to better propose matching candidates for the localisation. This means that nodes where localisation may be possible, but of low accuracy (such as distant nodes) are automatically excluded. Secondly, by leveraging multi-modal data, EM-LOC is able to resolve visual ambiguities such as corridors of similar appearance which are common very common in indoor environments and in many cases cannot be resolved using visual data alone. These two factors contribute to the 6-DoF pose estimation in EM-LOC producing fewer outliers and more accurate pose estimates.




Fig. 8. Number of features visible for each node in the museum environment test.

In terms of efficiency, a comparison between EM-LOC and the three competing approaches are reported in the right three subplots of Fig. 10. As is evident from the results of A2, the simple means of using the side-channel information as a prior estimate for selecting nodes in which to localise

is not very effective. At low values of k localisation is attempted in many nodes, leading to low efficiency and slow performance, while at high k , nodes in which localisation would have been accurate are missed. The reason for this can easily be seen from 6 where it is evident that the WiFi feature similarity does not perfectly mirror the true metric similarity. From the figure it is clear that EM-LOC successfully localises significantly more incoming images than the competing approaches at the same k values. Another noticeable trend is that even at $k = 1$, EM-LOC is still able to localise most images, meaning that the localisation can be performed extremely efficiently (every incoming frame only has to be matched against the 3D features contained by a single node). The localisation efficiency of EM-LOC is extremely high, and thus the efficiency is significantly higher for EM-LOC even compared to the A3 where a trained SVM is used to select candidate nodes.




Fig. 9. Cdf plot of the orientation and translation error of EM-LOC from various values of k .

In Fig. 7 we report a comparison and breakdown of the on-line localisation processing times for EM-LOC and the three competing methods. Due to the efficiency EM-LOC significantly cuts down the processing time spent on 2D- 3D feature matching compared to A1 which uses all nodes in the matching process.

TABLE I
EFFICIENCY OF THE 6-DOF LOCALISATION FOR THE 3 METHODS

Method	$k = 10$	$k = 5$	$k = 1$
EM-LOC	0.4783	0.4534	0.4534
A1	0.1043	0.0519	0.0187
A3	0.1317	0.1257	0.1078

In summary, by localising only against a select subset of nodes we are able to perform 6-DoF pose updates more frequently than exhaustive matching and also avoid occasional visual false positives, suppressing outlier pose estimates and resulting in a higher localization accuracy. Although the side-channel information is an ideal means of selecting candidate nodes for localisation, the noisy nature of this data means that existing methods such as using independent features (as in A2) or even an SVM that models correlations between features (as in A3) is not able to provide sufficiently accurate candidate nodes. By making use of a stream of side-channel information along with rough odometry estimates the EM-LOC system can significantly boost the accuracy of the node proposals.

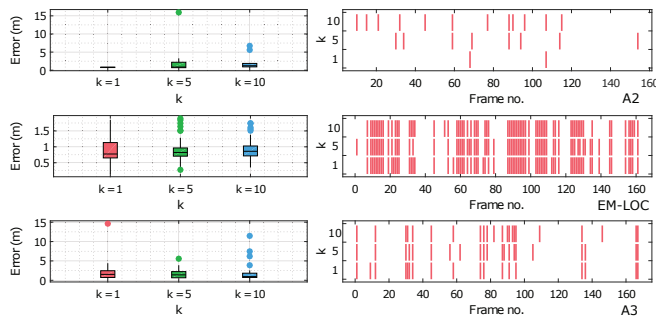


Fig. 10. Evaluation of the localisation accuracy and the localisation efficiency. The right subplots show the successful localisations for each of the input test images.

VII. CONCLUSION

In this paper, we proposed EM-LOC, an indoor localisation system which leverages multiple modalities for efficient global localisation. The key challenge here is the extremely high computational cost incurred when trying to match features to the millions of 3D features stored on a localisation graph. We address this by intelligently integrating multiple modalities in our localisation system, allowing us to significantly reduce the matching cost and thereby achieve real-time performance. Unlike other methods, our system is able to utilise a temporal stream of side-channel information in order to select candidate nodes in which to localise. The results show that our multi-modal approach not only increases the localisation accuracy but significantly reduces computational time compared to other approaches. For future work, we plan to focus on simultaneous, lightweight human and humanoid localisation.

REFERENCES

- [1] Hernán Badino, Daniel Huber, and Takeo Kanade. Real-time topometric localization. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1635–1642. IEEE, 2012.
- [2] G. Bradski. *opencv. Dr. Dobb's Journal of Software Tools*, 2000.
- [3] W. Churchill and P. Newman. Experience-based navigation for long-term localisation. *The International Journal of Robotics Research*, 32(14):1645–1661, 9 2013.
- [4] Mark Cummins and Paul Newman. Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 27(6):647–665, 2008.
- [5] Frank Dellaert. Factor graphs and gtsam: A hands-on introduction. 2012.
- [6] Arnaud Doucet, Nando De Freitas, Kevin Murphy, and Stuart Russell. Rao-blackwellised particle filtering for dynamic bayesian networks. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, pages 176–183. Morgan Kaufmann Publishers Inc., 2000.
- [7] Felix Endres, Jürgen Hess, Jürgen Sturm, Daniel Cremers, and Wolfram Burgard. 3-d mapping with an rgb-d camera. *IEEE Transactions on Robotics*, 30(1):177–187, 2014.
- [8] T Gadeke, Johannes Schmid, Wilhelm Stork, and KD Muller-Glaser. Pedestrian dead reckoning for person localization in a wireless sensor network. In *Proc. 2011 International conference on indoor positioning and indoor navigation*, pages 1–4, 2011.
- [9] Simon S Haykin, Simon S Haykin, and Simon S Haykin. *Kalman filtering and neural networks*. Wiley Online Library, 2001.
- [10] Arnold Irschara, Christopher Zach, Jan-Michael Frahm, and Horst Bischof. From structure-from-motion point clouds to fast location recognition. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2599–2606. IEEE, 2009.
- [11] Satoshi Ito, Felix Endres, Markus Kuderer, Gian Diego Tipaldi, Cyrill Stachniss, and Wolfram Burgard. W-rgb-d: floor-plan-based indoor global localization using a depth camera and wifi. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 417–422. IEEE, 2014.
- [12] Alexander Kleiner and Dali Sun. Decentralized slam for pedestrians without direct communication. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1461–1466. IEEE, 2007.
- [13] M. Labbe and F. Michaud. Appearance-Based Loop Closure Detection for Online Large-Scale and Long-Term Operation. *IEEE Transactions on Robotics*, 29(3):734–745, 2013.
- [14] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate $o(n)$ solution to the pnp problem. *International journal of computer vision*, 81(2):155–166, 2009.
- [15] Yunpeng Li, Noah Snavely, and Daniel P Huttenlocher. Location recognition using prioritized feature matching. In *European Conference on Computer Vision*, pages 791–804. Springer, 2010.
- [16] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [17] Will Maddern, Michael Milford, and Gordon Wyeth. Cat-slam: probabilistic localisation and mapping using a continuous appearance-based trajectory. *The International Journal of Robotics Research*, 31(4):429–451, 2012.
- [18] Michael J. Milford and Gordon F. Wyeth. Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 1643–1649, 2012.
- [19] David Nistér and Henrik Stewénius. Scalable recognition with a vocabulary tree. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2161–2168, 2006.
- [20] Valérie Renaudin, Melania Susi, and Gérard Lachapelle. Step length estimation using handheld inertial sensors. *Sensors*, 12(7):8507–8525, 2012.
- [21] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast image-based localization using direct 2d-to-3d matching. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 667–674. IEEE, 2011.
- [22] Yuanchao Shu, Cheng Bo, Guobin Shen, Chunshui Zhao, Liqun Li, and Feng Zhao. Magicol: indoor localization using pervasive magnetic field and opportunistic wifi sensing. *Selected Areas in Communications, IEEE Journal on*, 33(7):1443–1457, 2015.
- [23] Niko Sünderhauf and Peter Protzel. Switchable constraints for robust pose graph slam. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1879–1884. IEEE, 2012.
- [24] Iwan Ulrich and Illah Nourbakhsh. Appearance-based place recognition for topological localization. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 2, pages 1023–1029. Ieee, 2000.
- [25] Katsu Yamane. Systematic derivation of simplified dynamics for humanoid robots. In *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*, pages 28–35. IEEE, 2012.
- [26] Wei Zhang and Jana Kosecka. Image based localization in urban environments. In *3D Data Processing, Visualization, and Transmission, Third International Symposium on*, pages 33–40. IEEE, 2006.
- [27] Yuanqing Zheng, Guobin Shen, Liqun Li, Chunshui Zhao, Mo Li, and Feng Zhao. Travi-navi : Self-deployable indoor navigation system. In *Proceedings Mobicom 2014*, number 2. ACM, 2014.
- [28] Yuanqing Zheng, Guobin Shen, Liqun Li, Chunshui Zhao, Mo Li, and Feng Zhao. Travi-navi: Self-deployable indoor navigation system. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking, MobiCom '14*, pages 471–482, New York, NY, USA, 2014. ACM.