# SeDAR – Semantic Detection and Ranging: Humans can localise without LiDAR, can robots?

Oscar Mendez University of Surrey Guildford GU2 Simon Hadfield University of Surrey Guildford GU2 Nicolas Pugeault University of Exeter Exeter EX4 Richard Bowden University of Surrey Guildford GU2

o.mendez@surrey.ac.uk s.hadfield@surrey.ac.uk n.pugeault@exeter.ac.uk r.bowden@surrey.ac.uk

*Abstract*— How does a person work out their location using a floorplan? It is probably safe to say that we do not explicitly measure depths to every visible surface and try to match them against different pose estimates in the floorplan. And yet, this is exactly how most robotic scan-matching algorithms operate. Similarly, we do not extrude the 2D geometry present in the floorplan into 3D and try to align it to the real-world. And yet, this is how most vision-based approaches localise.

Humans do the exact opposite. Instead of depth, we use high level semantic cues. Instead of extruding the floorplan up into the third dimension, we collapse the 3D world into a 2D representation. Evidence of this is that many of the floorplans we use in everyday life are not accurate, opting instead for high levels of discriminative landmarks.

In this work, we use this insight to present a global localisation approach that relies solely on the semantic labels present in the floorplan and extracted from RGB images. While our approach is able to use range measurements if available, we demonstrate that they are unnecessary as we can achieve results comparable to state-of-the-art without them.

#### I. INTRODUCTION

Indoor localisation is perhaps one of the most crucial aspects for any robotic system. It allows robots to interact with the world and provides a representation and understanding that can be shared with humans and other agents. Traditional Vision-Based Simultaneous Localization and Mapping (VSLAM) systems can provide localisation within a map that is built on-the-fly. However, VSLAM systems are liable to drift in terms of both pose and scale. They can also become globally inconsistent in the case of failed loop closures. Finally, even in the case of no scale drift and correct loop closures, a VSLAM system can only ever guarantee global consistency internally. This means that while pose estimates are globally consistent, they are only valid within the context of the VSLAM system. There are no guarantees, at least in vision-only systems, that we can directly map the reconstruction to the real world (or between agents).

This problem is normally addressed by having a localisation system that can relate the pose of the robot to a preexisting map. Examples of global localisation frameworks include the Global Positioning System (GPS) and traditional Monte-Carlo Localisation (MCL). MCL has the ability to localise within an existing floorplan (which can be safely assumed to be available for most indoor scenarios). This is a highly desirable trait, as it implicitly eliminates drift, is



Fig. 1: A) RGB Image, B) CNN-Based Semantic Labelling and C) Sample SeDAR Scan within floorplan.

globally consistent and provides a way for the created 3D reconstructions to be related to the real world without having to perform expensive post-hoc optimizations. Traditionally, the range-based scans required by MCL have been produced by expensive sensors such as Light Detection And Ranging (LiDAR). These sensors are capable of producing high density measurements at high rates with low noise, making them ideal for range-based MCL. However, in addition to their expense, they are large and require a lot of power.

As a response to this, modern low-budget robotic platforms have used RGB-D cameras as a cheap and lowfootprint alternative. This has made vision-based floorplan localisation an active topic in the literature. However, while many approaches have been proposed, they normally use heuristics to lift the 2D plan into the 3D coordinate system of VSLAM. Examples include Liu *et al.* [15], who use visual cues such as Vanishing Points (VPs) or Chu *et al.* [3] who perform piecemeal 3D reconstructions that can then be fitted back to an extruded floorplan. A common problem with these approaches is that the 3D data extracted from the image is normally orthogonal to the floorplan that it is meant to localise in. This means that assumptions must be made about dimensions not present in the floorplan. These approaches also do not fully exploit the floorplan, ignoring the semantic information.

We propose a fundamentally different approach that is inspired by how humans perform the task. Instead of discarding valuable semantic information, we use a Convolutional Neural Network (CNN)-based encoder-decoder to extract highlevel semantic information. We then collapse all semantic information into 2D in order to reduce the assumptions about the environment. We then use these labels, image geometry and (optionally) depth along with a semantically labelled floorplan to create a state-of-the-art sensing and localisation framework.

Semantic Detection and Ranging (SeDAR) is an innovative human-inspired framework that combines new semantic sensing capabilities with a novel semantic Monte-Carlo Localisation (MCL) approach. As an example, figure 1 shows a sample SeDAR scan localised in the floorplan. We show that SeDAR has the ability to surpass LiDAR-based MCL approaches. SeDAR also has the ability to perform drift-free local, as well as global, localisation. Furthermore, experimental results show that the semantic labels are sufficiently strong visual cues such that depth estimates are no longer needed. Not only does this vision-only approach perform comparably to depth-based methods, it is also capable of coping with map inaccuracies more gracefully than strictly depth-based approaches.

This paper describes the process by which SeDAR is used as a novel human-inspired sensing and localisation framework. In section IV-A, semantically salient elements are extracted from a floorplan. Section IV-B describes how these semantic elements are identified in the robot's camera by using a state-of-the-art CNN-based semantic segmentation algorithm and presented as a novel sensing modality. We then present the three main contributions of this paper. First, section IV-C introduces a novel motion model that includes a "ghost factor" that uses semantic information to influence how particles move through occupied space. Second, section IV-D introduces a novel sensor model that estimates observation likelihoods using semantic information, range and bearing information. Third, section IV-E introduces a second novel motion model that uses semantic and bearing information to allow observation likelihoods to be estimated from an RGB image only. Finally, in section V we present the results obtained by using our approach in multiple sensing modalities.

## II. LITERATURE REVIEW

Monte-Carlo Localisation (MCL) was made possible by the arrival of accurate range-based sensors such as SOund Navigation And Ranging (SONAR) and Light Detection And Ranging (LiDAR). These approaches, which we call Range-Based Monte-Carlo Localisation (RMCL), are robust and reliable and still considered state-of-the-art in many robotic applications. Recent advances in computer vision have made it possible for us to imagine new types of perceptual sensors which are capable of semantic understanding of a scene. Semantic sensing modalities, such as SeDAR, have the ability to revolutionize MCL. RMCL was first introduced by Fox *et al.* [8] and Dellaert *et al.* [6]. RMCL improved the Kalman Filter based state-of-the-art by allowing multi-modal distributions to be represented. It also solved the computational complexity of grid-based Markov approaches. However, these approaches require expensive LiDAR and/or SONAR sensors to operate reliably. Instead, Dellaert *et al.* [5] extended their approach to operate using vision-based sensor models. Vision-based MCL allowed the use of rich visual features and cheap sensors, but had limited performance compared to the more robust LiDAR-based systems.

With the rising popularity of RGB-D sensors, more robust vision-based MCL approaches became possible. Paton and Kosecka [17] use a combination of feature matching and Iterative Closest Point (ICP) to perform pose estimation and localisation. Brubaker et al. [2] used visual odometry and pre-existing roadmaps in a joint MCL/closed-form approach in order to localise a moving car. Fallon et al. [7] presented a robust MCL approach that used a low fidelity a priori map to localise in, but required the space to be traversed by a depth sensor beforehand. Winterhalter et al. [24] performed MCL, but based the likelihood of the sensor model on the normals of an extruded floorplan. Chu et al. [3] is the closest to us, they attempted to mimic the human thinking process by creating piecemeal reconstructions of an extruded floorplan, the MCL sensor model was then based on matches against these reconstructions. These MCL-based approaches tend to be robust, but they operate entirely on the geometric information present in the floorplan and therefore require depth images either from sensors and/or reconstructions. By contrast our approach aims to use non-geometric semantic information present in the floorplan in order to perform the localisation.

While the field of MCL evolved in the robotics community, in vision, the non-MCL-based field of floorplan localisation became more popular. Melbouci et al. [16] used extruded floorplans, but performed local bundle adjustments instead of MCL. Shotton et al. [19] used regression forests to predict the correspondences of every pixel in the image to a known 3D scene, they then combined this in a RANdom Sample And Consensus (RANSAC) approach in order to solve the camera pose. Chu et al. [4] use information from the floorplans and Google StreetView in order to reason about the geometry of the building and perform a robust reconstruction. The most similar work to our approach is Wang et al. [23] who use text detection from shop fronts as semantic cues to localise in the floorplan of a shopping centre and Liu et al. [15] who use floorplans as a source of geometric and semantic information, combined with vanishing points, to localise monocular cameras. These visionbased approaches tend to use more of the non-geometric information present in the floorplan. However, a common trend is that assumptions must be made about geometry not present in the floorplan (e.g. ceiling height). The floorplan is then extruded out into the 3<sup>rd</sup> dimension to allow approaches to use the information present in the image. By contrast, our approach aims to extract the information from the image and



Fig. 2: Laser scan matching, the robot is correctly localised when the observations match the geometry of the map [22]

collapse the 3D world down into the 2D floorplan where localisation can be performed. This provides a 3-Degrees of Freedom (DoF) localisation requiring less assumptions about the environment.

Recently, advances in Deep Learning have made robust semantic segmentation models widely available. Approaches like that of Badrinarayanan et al. [1], Kendal et al. [11] and Long et al. [18] have made semantically informed approaches possible. One such approach is Tateno et al. [21] who use the CNN-based depth and semantic label predictions of Laina et al. [13] to aid in their Simultaneous Localization and Mapping (SLAM) pipeline. Lee et al. [14] extend the approach of Badrinarayanan et al. [1] to directly estimate room layout keypoints. While many such approaches exist, they mainly focus on extracting the room layout based on Manhattan world assumptions. Instead, this work proposes to use CNN-based semantic segmentation (that is understandable to humans) in order to extract labels that are inherently present in human-readable floorplans. This allows us to take all that information and collapse it into a 3-DoF problem, making our approach more tractable than competing 6-DoF approaches while avoiding additional assumptions.

#### III. BACKGROUND

While there exist many approaches to perform MCL, RMCL is widely considered to be the state-of-the-art localisation method for pre-existing maps. RMCL is a scanmatching algorithm, it assumes the presence of a sensor that provides range and bearing tuples across a scanline. The problem then becomes finding the pose of the robot that makes the sensor observations match the floorplan. Figure 2a shows a case of the scan being correctly matched for a correctly localised robot. Conversely, figure 2b shows an incorrectly matched scan for an incorrect pose.

State-of-the-art localisation performs this matching in an Sequential Monte-Carlo (SMC) framework, which can be broadly summarised as follows. First, there is a prediction stage where particles are propagated using a motion-model, normally odometry from the robot (with Gaussian noise). Second, an update phase where each particle is weighted according to how accurately the observations align to the map. Finally, a re-sampling step is performed proportional to each particle's weight. The process is then repeated.

More formally, the pose  $\mathbf{x}_t \in \mathbb{X}_t \subset \text{SE}(2)$  can be estimated using a set of pose samples  $\mathbb{S}_t = \{s_t^i; i = 1..N\}$ , odometry measurements  $\mathbb{U}_t = \{u_j; j = 1..t\}$ , sensor measurements  $\mathbb{Z}_t = \{z_j; j = 1..t\}$  and a 2D map  $\mathbb{V}$ . The posterior is calculated as

$$\Pr\left(s_{t}^{i} \middle| \mathbb{Z}_{t}, \mathbb{U}_{t}\right) = \Pr\left(z_{t} \middle| s_{t}^{i\prime}, \mathbb{V}\right) \Pr\left(s_{t}^{i\prime} \middle| u_{t}, s_{t-1}^{i}\right) \Pr\left(s_{t}^{i} \middle| \mathbb{Z}_{t}, \mathbb{U}_{t}\right)$$
(1)

which implies that only the most recent odometry and observations are used [6]. This means that at each iteration the particles from  $\Pr\left(s_{t-1}^{i} | z_{t-1}, u_{t-1}\right)$  are: propagated using a motion model  $\Pr\left(s_{t}^{i\prime} | u_{t}, s_{t-1}^{i}\right)$ , weighted using a sensor model  $\Pr\left(z_{t} | s_{t}^{i\prime}, \mathbb{V}\right)$  and resampled into the posterior  $\Pr\left(s_{t}^{i} | z_{t}, u_{t}\right)$ . Algorithm 1 describes this process in more detail.

1:	function MCL( $\mathbb{S}_{t-1}, u_t, z_t$ )
2:	$\mathbb{S}_t = \mathbb{S}'_t = \emptyset$
3:	for $i = 1 \rightarrow N$ do
4:	$s_t^{i\prime} \leftarrow \text{MOTION\_MODEL}(u_t, s_{t-1}^i)$
5:	$w_t^i \leftarrow \text{SENSOR\_UPDATE}(z_t, s_t^{i\prime}, \mathbb{V})$
6:	$\mathbb{S}'_t \leftarrow \mathbb{S}'_t + \left\langle s^{i\prime}_t, w^i_t \right angle$
7:	end for
8:	for $i = 1 \rightarrow N$ do
9:	$s_t \leftarrow \text{WEIGHTED}\_\text{SAMPLE}(\mathbb{S}'_t)$
10:	$\mathbb{S}_t \leftarrow \mathbb{S}_t + s_t$
11:	end for
12:	$\bar{\mathbb{S}}_t \leftarrow \text{MEAN}(\mathbb{S}_t)$
13:	return $\tilde{\mathbb{S}}_t$
14:	end function

Algorithm 1: Sequential Monte-Carlo Localisation in a known floorplan.

In an MCL context, the motion model is defined by the odometry received from the robot. This propagates the particles according to  $u_t$  with Gaussian noise applied such that

$$\Pr\left(s_{t}^{i\prime} \middle| u_{t}, s_{t-1}^{i}\right) \sim s_{t-1}^{i} + \mathcal{N}\left(u_{t}, \Upsilon_{t}\right)$$
(2)

where  $\Upsilon_t$  is a covariance on the linear and angular components of the odometry. Fundamentally, this allows errors in the odometry to be accounted for during particle propagation.

The sensor model is defined by each range-scanner observation. We can estimate the likelihood of each full range-scan  $(z_t)$  under the assumption that each measurement in the scan is independent of each other. That is,

$$\Pr\left(z_t \middle| s_t^{i\prime}, \mathbb{V}\right) = \prod_{k=1}^{K} \Pr\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right)$$
(3)

is the likelihood of the putative particle  $s_t^{i\prime}$ , where  $z_t = \{\langle \theta_t^k, r_t^k \rangle; k = 1..K\}$  is the set of range and bearing tuples that make up each scan. Calculating the likelihood can be done two ways, the beam model and the likelihood field model.

In the beam model, a raycasting operation is performed. Starting from the current particle's pose, a ray is cast along the bearing angle  $\theta_t^k$  until it intersects map geometry. The likelihood  $\Pr\left(z_t^k | s_t^{i'}, \mathbb{V}\right)$  is then estimated using the difference between the range  $r_t^k$  obtained from the sensor and the range  $r_t^{k*}$  obtained from the raycasting operation.

In the likelihood field model, a distance map is used in order to avoid the expensive raycasting operation. The distance map is a Lookup Table (LUT) of the same size as the map, where each cell contains the distance to the nearest map geometry. This map is estimated similar to the Chamfer distance, where we perform a search in a window around each cell and store the distance to the closest occupied cell in the map. When queried, this distance is converted into a likelihood. Figure 3 shows the estimated distance map for a floorplan, the creation of which will be explored further in section IV-A. This distance map is only estimated once during initialisation. During runtime, the endpoint of each measurement can be estimated directly from the pose, bearing and range. The probability is then simply related to the distance reported by the LUT.

The raycasting method is (strictly speaking) more closely related to the sensing modality, as the closest geometry may not lie along the ray. However, in practice, most robotics systems use the likelihood field model as it is both faster and tends to provide better results. This is because the raycasting operation can report very incorrect measurements due to small pose errors. An example of this is when looking through an open door, an error of a few centimetres can make the rays miss the door. This makes the distribution inherently less smooth.

## IV. METHODOLOGY

The problem with state-of-the-art approaches is that they only use one-dimensional information. More explicitly, using only the range information from the sensors fundamentally limits how discriminative each reading can be. Instead, we present a novel semantic sensing and localisation framework called SeDAR that leverages semantic and, optionally, range information. We will show that we can use our novel SeDAR sensing and localisation framework to outperform traditional RMCL.

#### A. Semantic Floorplans

RMCL requires a floorplan and/or previously created range-scan map that is accurate in scale and globally consistent. A previously created range-scan map requires a robust SLAM algorithm such as GMapping [9] to be run. This is not an ideal situation as it forces the robot to perform an initial exploration, is sensitive to noise and the resulting map is difficult to interpret by humans. Instead of using a metric-accurate reconstruction, a more flexible and feasible alternative is using a human-readable floorplan. However, this would make RMCL less robust due to differences between the floorplan and what the robot can observe (*e.g.* inaccuracies, scale variation and furniture). To overcome this, we augment the localisation with semantic labels extracted from an existing floorplan. We limit the labels to walls, doors and windows. The reason for this limitation is two-fold. First, they are salient pieces of information that humans naturally use to localise. Second, they are simple to automatically extract from a floorplan using image processing. As can be seen in figure 3c, these semantically salient elements have been colour coded in order to represent different labels.

In order to make a labelled floorplan readable by the robot, it must first be converted into an occupancy grid. An occupancy grid is a 2D representation of the world, in which each cell in the grid has an occupancy probability attached to it. Any cell that is above an occupancy threshold is then considered as being occupied. Estimating the occupancy of an existing floorplan is done by taking the normalized greyscale value of each cell.

The map can then be defined as

$$\mathbb{V} = \left\{ v_{\mathbf{m}}; \mathbf{m} \in \mathbb{M} \subset \mathbb{W} \right\}$$
(4)

where  $\mathbb{M}$  is a set of 2D positions and  $\mathbb{W}$  is the set of whole numbers. Assuming  $\mathcal{L} = \{a, d, w\}$  is the set of possible cell labels (wall, door, window), each cell is defined as

$$v_{\mathbf{m}} = \left\langle v_{\mathbf{m}}^{\circ}, v_{\mathbf{m}}^{w}, v_{\mathbf{m}}^{d}, v_{\mathbf{m}}^{a} \right\rangle \tag{5}$$

where  $v_{\mathbf{m}}^{\circ}$  is the occupancy likelihood and  $\ell \in \mathcal{L}$  denotes the label likelihood.

Having incorporated the semantic labels into the standard occupancy grid, it is now necessary to use them in sensing.

### B. SeDAR Sensor

Modern low-cost robotics systems turn RGB-D images into range-based scans  $z_t = \{\langle \theta_t^k, r_t^k \rangle; k = 1..K\}$ . This can be accomplished by looking exclusively at the depth image. It can be shown that the angle along the horizontal axis,  $\theta_t^k$ , can be calculated by

$$\theta_t^k = \operatorname{atan2}\left(\frac{u - c_x}{f_x}\right) \tag{6}$$

where (u, v),  $(c_x, c_y)$ ,  $(f_x, f_y)$  are the pixel coordinates, of the principal point and focal length, respectively, of the camera. While it is possible to estimate a second angle along the vertical axis, this is unnecessary in the case of floorplan localisation. It can also be shown that the range measurement  $r_t^k$  can be calculated as

$$r_{t}^{k} = \sqrt{\left(\frac{d_{t}^{k}(u-c_{x})}{f_{x}}\right)^{2} + \left(\frac{d_{t}^{k}(v-c_{y})}{f_{y}}\right)^{2} + \left(d_{t}^{k}\right)^{2}} (7)$$

where  $d_t^k$  is the current depth measurement at pixel k. At this point, a traditional LiDAR can be emulated. Notice that all the information present in the RGB image is being discarded. On the other hand, a single SeDAR scan consists of a set  $z_t = \{\langle \theta_t^k, r_t^k, \ell_t^k \rangle; k = 1..K\}$  of bearing, range and label tuples.

In order to estimate the labels, CNN-based encoderdecoder network [11] is used. This is trained on the SUN3D



Fig. 3: Original floorplan compared to the likelihood field and the labelled floorplan.

[25] dataset, and can reliably detect doors, walls, floors, ceilings, furniture and windows. This state-of-the-art semantic segmentation runs in real-time, which allows images to be parsed into a SeDAR-scan with negligible latency. The label  $\ell_t^k$  is then simply the label at pixel k.

It is important to note that we extract the labels from the RGB image only. This is by design, as it allows the use of cameras that cannot sense depth. In the following sections we will use this novel sensing modality in a novel MCL formulation with and without the range-based measurements.

#### C. Motion Model

MCL motion models normally assume the motion model to be represented by  $\Pr\left(s_t^{i\prime}|u_t, s_{t-1}^i\right)$ . However, it is well understood in the literature that the actual distribution being approximated is  $\Pr\left(s_t^{i\prime}|u_t, s_{t-1}^i, \mathbb{V}\right)$ . This encodes the idea that certain motions are more or less likely depending on the map (*e.g.* through walls).

Under the assumption that the motion of the robot is small, it can be shown that

$$\Pr\left(s_{t}^{i\prime}\big|u_{t}, s_{t-1}^{i}, \mathbb{V}\right) = \kappa \Pr\left(s_{t}^{i\prime}\big|u_{t}, s_{t-1}^{i}\right) \Pr\left(s_{t-1}^{i}\big|\mathbb{V}\right)$$
(8)

(see *e.g.* [22]) where  $\kappa$  is a normalising factor and  $\mathbb{V}$  is the set containing every cell in the map. This allows the two likelihoods to be treated independently.

In an occupancy map, the motion  $\Pr\left(s_{t}^{i\prime}|u_{t}, s_{t-1}^{i}\right)$  is defined in the same way as equation 2. The prior  $\Pr\left(s_{t-1}^{i}|\mathbb{V}\right)$  is simply the occupancy likelihood of the cell that contains  $s_{t}^{i}$ , that is

$$\Pr\left(s_{t-1}^{i}\middle|\mathbb{V}\right) = 1 - \Pr\left(v_{s_{t-1}}^{\circ}\right) \tag{9}$$

which is an elegant solution in the case where the "floorplan" was previously built by the robot.

More explicitly, when a robot builds its own floorplan the observations it uses are noisy. This means that, by definition, the map encodes a measure of occupancy that reflects what the robot can observe. This is because, as the map is built, each cell's occupancy likelihood is updated. Cells that are always observed as occupied (walls) have high likelihood of being occupied. Cells that have always been empty (middle of a room) have low likelihood of being occupied. More importantly, cells that are near walls or windows are uncertain, and this is reflected in their occupancy. This is true for a map built using offline scan-alignment and/or online SLAM techniques (note that map-building is *not* a

However, this approach becomes problematic when using human-made floorplans. Human-made floorplans typically have binary edges and/or edges with image artefacts (in the case where they are scanned into a computer). Either way, this does not reflect what the robot can observe and can cause issues with localisation. Therefore, most approaches tend to assume a binary interpretation of the occupancy. This is done by setting the probability to

$$\Pr\left(v_{s_{t-1}}^{\circ}\right) = \begin{cases} 1 & \text{if } v_{s_{t-1}}^{\circ} \ge \tau_{o} \\ 0 & \text{otherwise} \end{cases}$$
(10)

where  $\tau_o$  is a user defined threshold. While this makes depthbased methods perform reliably, it is a crude estimate of reality. For instance, most humans would not even notice if a door is a few centimetres away from where it should be. Issues like this present real problems when particles propagate though doors, as it is possible that the filter will discard particles as they collide with the edge of the door frame. Instead, we propose to augment this with a *ghost factor* that allows particles more leeway in these scenarios. Therefore the proposed prior is

$$\Pr\left(s_{t-1}^{i} \middle| \mathbb{V}\right) = \left(1 - \Pr\left(v_{s_{t-1}}^{o}\right)\right) e^{-\epsilon_{o} \delta_{a}} \qquad (11)$$

where  $\delta_a$  is the distance to the nearest door. While other labels such as windows can be used, in the case of a ground-based robot doors are sufficient. The distance,  $\delta_a$ , can be efficiently estimated using a lookup table as defined in section IV-D.



Fig. 4: Original floorplan compared to the likelihood field for each label.

More importantly,  $\epsilon_{\rm g}$  is a user defined factor that determines how harshly this penalty is applied. Setting  $\epsilon_{\rm g} = 0$  allows particles to navigate through walls with no penalty, while very high values approximate equation 10. We will explore the effects of  $\epsilon_{\rm g}$  in section V-D. This motion model is more probabilistically accurate than the occupancy model used in most RMCL approaches, and has the added advantage of leveraging the high-level semantic information present in the map.

## D. Sensor Model

The naïve way of incorporating semantic measurements into the sensor model would be to use the beam model. In this modality, the raycasting operation would provide not only the distance travelled by the ray, but also the label of the cell the ray hit. If the label of the cell and the observation match, the likelihood of that particle being correct is increased. However, this approach suffers from the same limitations as the traditional beam model: it has a distinct lack of smoothness. On the other hand, the likelihood field model is significantly smoother, as it provides a gradient between each of the cells. By contrast, the approach presented here uses a joint method that can use likelihood fields to incorporate semantic information in the presence of semantic labels. More importantly, it can also use raycasting within a likelihood field in order to operate without range measurements.

As described in section III, the likelihood field model calculates a distance map. For each cell  $v_m$ , the distance to the nearest occupied cell

$$\delta_o(\mathbf{m}) = \min_{\mathbf{m}'} \|\mathbf{m} - \mathbf{m}'\|, \quad v_{\mathbf{m}'}^o > \tau_o \qquad (12)$$

is calculated and stored. For clarity, we omit the parameter **m** for the remainder of the paper. When a measurement  $z_t^k = \langle \theta_t^k, r_t^k \rangle$  is received, the endpoint is estimated and used as an index to the distance map. Assuming a Gaussian error distribution, the weight of each particle  $s_t^{i\prime}$  can then estimated as

$$\Pr_{\text{\tiny RNG}}\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right) = e^{\frac{-\delta_o^2}{2\,\sigma_o^2}} \tag{13}$$

where  $\delta_o$  is the value obtained from the distance map and  $\sigma_o$  is dictated by the noise characteristics of the sensor. However, this model has three main limitations. First, it makes no use of the semantic information present in the map. Second, the parameter  $\sigma_o$  must be estimated by the user and assumes all measurements within a scan have the same noise parameters. Third, it is incapable of operating in the absence of range measurements.

Instead, this work uses the semantic labels present in the map to create multiple likelihood fields. For each label present in the floorplan, we can calculate a distance map that stores the shortest distance to a cell with the same label. Formally, for each map cell  $v_{\rm m}$  we can estimate the distance to the nearest cell of each label as

$$\delta_{\ell}(\mathbf{m}) = \min_{\mathbf{m}'} \|\mathbf{m} - \mathbf{m}'\|, \quad v_{\mathbf{m}'}^{\ell} > \tau_{o}$$
(14)

where  $\delta_{\ell} = \{\delta_a, \delta_d, \delta_w\}$  are distances to the nearest wall, door and window, respectively. Figure 4 shows the distance maps for each label. This approach overcomes the three limitations of the state-of-the-art, which we will now discuss.

1) Semantic Information: First, SeDAR uses the semantic information present in the map. When we receive an observation  $z_t^k = \langle \theta_t^k, r_t^k, \ell_t^k \rangle$ , we use the bearing  $\theta_t^k$  and range  $r_t^k$  information to estimate the endpoint of the scan. We then use the label  $\ell_t^k$  to decide which semantic likelihood field to use. Using the endpoint from the previous step, the label-likelihood can be estimated similarly to equation 13,

$$\Pr_{\scriptscriptstyle LBL}\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right) = e^{\frac{-\delta_\ell^2}{2\,\sigma_\ell^2}} \tag{15}$$

where  $\delta_{\ell}$  is the distance to the nearest cell of the relevant label and  $\sigma_{\ell}$  is the standard deviation (which we will define using the label prior). The probability of an observation given the map and pose can then be estimated as

$$\Pr\left(z_{t}^{k} \middle| s_{t}^{i\prime}, \mathbb{V}\right) = \epsilon_{o} \Pr_{\text{\tiny RNG}}\left(z_{t}^{k} \middle| s_{t}^{i\prime}, \mathbb{V}\right) + \epsilon_{\ell} \Pr_{\text{\tiny LBL}}\left(z_{t}^{k} \middle| s_{t}^{i\prime}, \mathbb{V}\right)$$
(16)

where  $\epsilon_o$  and  $\epsilon_\ell$  are user defined weights. When  $\epsilon_\ell = 0$  the likelihood is the same as standard RMCL. On the other hand, when  $\epsilon_o = 0$  the approach is using only the semantic

information present in the floorplan. These weights are properly explored and defined in section V-C. Unlike range scanners,  $\sigma_{\ell}$  cannot be related to the physical properties of the sensor. Instead, this standard deviation is estimated directly from the prior of each label on the map. Defining  $\sigma_{\ell}$  this way has the benefit of not requiring tuning. However, there is a much more important effect that must be discussed.

2) Semantically Adaptive Standard Deviation: When a human reads a floorplan, unique landmarks are the most discriminative features. The more unique a landmark, the easier it is to localise using it (because there aren't many areas in the map that contain it). It then follows that the more rare a landmark, the more discriminative it is for the purpose of localisation. Indeed, it is easier for a person to localise in a floorplan by the configuration of doors and windows than it is by the configuration of walls. This translates into the a simple insight: *lower priors are more discriminative*. Therefore,  $\sigma_{\ell}$  is tied to the prior of each label not only because it is one less parameter to tune, but because it implicitly makes observing rare landmarks more beneficial than common landmarks.

Relating  $\sigma_{\ell}$  to the label prior  $\Pr(\ell)$  controls how smoothly the distribution decays w.r.t. distance from the cell. The smaller  $\Pr(\ell)$  is, the smoother the decay. In essence, the localisation algorithm should be more *lenient* on sparser labels.

## E. Range-less Semantic Scan-Matching

The final, and most important, strength of this approach is the ability to perform all of the previously described methodology in the complete absence of range measurements. So far, we have formalised this approach on the assumption that we received either  $\langle \theta_t^k, r_t^k \rangle$  tuples (existing approaches) or  $\langle \theta_t^k, r_t^k, \ell_t^k \rangle$  tuples (SeDAR-based approach). However, this approach is capable of operating directly on  $\langle \theta_t^k, \ell_t^k \rangle$  tuples. In other words, depth measurments are *explicitly* not added to this approach.

Incorporating range-less measurements is simple. The beam and likelihood field models are combined in a novel approach that avoids the degeneracies that would happen in traditional RMCL approaches. In the standard approach, the raycasting operation terminates when an occupied cell is reached and the likelihood is estimated as

$$\Pr\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right) = e^{\frac{-\left(r_t^k - r_t^{k*}\right)^2}{2 \sigma_o^2}}$$
(17)

where  $r_t^k$  is the range obtained from the sensor and  $r_t^{k*}$  is the distance travelled by the ray. Unfortunately, in the absence of a range-based measurement  $r_t^k$  this is impossible. Using the standard distance map is also impossible, since we can't estimate the endpoint of the ray. Using raycasting in the distance map fails similarly. The raycasting terminates on an occupied cell, implying  $\delta_o = 0$  for every ray cast.

On the other hand, the semantic likelihood fields can still be used as  $\delta_{\ell}$  will still have a meaningful and discriminative value. We call this operation semantic raycasting. For every  $z_t^k = \langle \theta_t^k, \ell_t^k \rangle$ , the raycasting is performed as described in

section III. However, instead of comparing  $r_t^k$  and  $r_t^{k*}$  or using  $\delta_o$ , the label  $\ell_t^k$  is used to decide what likelihood field to use. The cost can then be estimated as

$$\Pr\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right) = \Pr_{\text{\tiny LBL}}\left(z_t^k \middle| s_t^{i\prime}, \mathbb{V}\right) \tag{18}$$

where  $\Pr_{\text{LBL}}(z_t^k | s_t^{i'}, \mathbb{V})$  is defined in equation 15. This method is essentially a combination of the beam-model and the likelihood field model.

In the absence of range-measurements to estimate an endpoint from, this hybrid approach uses semantic raycasting to find the nearest occupied cell. The distances are then used to provide smoothness to equation 18. More explicitly, it would be possible to assign binary values to this likelihood (*i.e.* label matches or not). However, this would be a naïve solution that provides no smooth gradient to the correct solution. While this approach uses distances from the likelihood field, this should not imply that the likelihood of an observation is dependant on distances to labels. Instead, equation 18 implies that the observation likelihood is directly proportional to the angular distribution of labels (*i.e.* how closely the bearing/label tuples match the observation).

To summarise, this section presented several important concepts. We introduced the idea of a semantic floorplan that contains information that is salient to humans. We also introduced a new sensing modality, SeDAR, that adds semantic labels to the traditional LiDAR information. We then incorporated these two ideas into a novel MCL-based approach. This approach is capable of using the semantic information present in the map to define a novel motion model. It is also capable of using the labels from a CNNbased segmentation to localise within the map. Our approach can do all of the above both in the presence, and absence, of range measurements. In the following section, we show that our approach is capable of outperforming standard RMCL approaches when using depth, and that it provides comparable performance in its absence.

## V. RESULTS

This section will demonstrate that SeDAR-based MCL is capable of reliably out-performing the state-of-the-art when using range measurements. It will also show that our approach it is capable of comparable performance even in the absence of range. First, the experimental setup is described. This consists of creating a dataset of a trajectory within a floorplan, as well as establishing error metrics. Then a comparison of several approaches is performed. The comparison is done in terms of room-level and global localisation, both quantitative and qualitative. Finally, we show the effects of our parameters.

# A. Experimental Setup

In order to evaluate this approach, we require a dataset that has several important characteristics. The dataset should consist of a robot navigating within a human-readable floorplan. Human-readability is required to ensure semantic information is present. The trajectory should be captured with an RGB-D camera in order to extract all the possible tuple



Fig. 5: Sample Trajectory used for evaluation.

combinations (range, bearing and label). Finally, we expect the trajectory of the robot to happen on the same plane as the floorplan. Unfortunately, most of the MCL datasets in the literature do not contain a floorplan, opting instead for laser-scans. RGB-D SLAM datasets are more appropriate, but they either do not move on the floorplan plane or simply do not contain ground-truth trajectory estimation.

Therefore, we are forced to use our own dataset - which we will make publicly available. We use the floorplan in figure 3a because it is large enough to provide multiple trajectories with no overlap. The dataset was collected using the popular TurtleBot platform, as it has a front-facing Kinect that can be used for emulating both LiDAR and SeDAR.

Normally, the ground-truth trajectory for floorplan localisation is either manually estimated (as in [24]) or estimated using Motion Capture (MoCap) systems (as in [20]). However, both of these approaches are limited in scope. Manual ground-truth estimation is time-consuming and impractical. MoCap is expensive, difficult to calibrate, and normally cannot remain in the public areas required for floorplan localisation. In order to overcome these limitations, well established RGB-D SLAM systems are used instead. The excellent approach by Labbe et al. [12] provides very accurate pose estimation in complex environments. While it does not localise within a floorplan, it does provide an accurate reconstruction and trajectory for the robot, which can then be registered into the floorplan. Figure 5a shows a sample trajectory and map estimated by [12], while figure 5b shows them overlaid on the floorplan.

To quantitatively evaluate the presented approach against ground truth, the Absolute Trajectory Error (A) metric presented by Sturm *et al.* [20] is used. A is estimated by first registering the two trajectories using the closed form solution of Horn [10], who finds a rigid transformation  ${}^{\mathbb{G}}\mathbf{T}_{\mathbb{X}}$  that registers the trajectory  $\mathbb{X}_t$  to the ground truth  $\mathbb{G}_t$ . At every time step t, the A can then be estimated as

$$e_{\mathbf{g}} = \mathbf{g}_t^{\mathbf{1}\,\mathbb{G}} \mathbf{T}_{\mathbf{x}} \mathbf{x}_t \tag{19}$$

where  $\mathbf{g}_t \in \mathbb{G}_t$  and  $\mathbf{x}_t \in \mathbb{X}_t$  are the current time-aligned poses of the ground truth and estimated trajectory, respectively. The Root Mean Square Error (RMSE), mean and median values of this error metric are reported, as these are indicative of performance over room-level initialisation. In



Fig. 6: Semantic Floorplan Localisation, room-level initialisation.

order to visualise the global localisation process, the error of each successive pose is shown (error as it varies with time). These metrics are sufficient to objectively demonstrate the systems ability to globally localise in a floorplan, while also being able to measure room-level initialisation performance.

We compare the work presented here against the extremely popular MCL approach present in Robot Operating System (ROS), called Adaptive Monte Carlo Localisation (AMCL) [6]. AMCL is the standard MCL approach in the robotics community. Any improvements over this approach are therefore extremely valuable. In all experiments, any overlapping parameters (such as  $\sigma_o$ ) are kept the same. The only parameters varied are  $\epsilon_{\ell}$ ,  $\epsilon_o$  and  $\epsilon_o$ .

### B. Room-Level Initialisation

For this evaluation, a room-level initialisation with standard deviations of 2.0m in (x, y) and 2.0rad in  $\theta$  is given to both AMCL and the proposed approach. The systems then ran with a maximum of 1000 particles (minimum 250) placed around the covariance ellipse. We record the error as each new image in the dataset is added.

1) Quantitative Results: Figure 6 compares four distinct scenarios against AMCL. Of these four scenarios, two use the range measurements from the Microsoft Kinect (blue lines) and two only use the RGB image (red lines).

The first range-enabled scenario uses the range measurements to estimate the endpoint of the measurement (and therefore the index in the distance map) and then sets ( $\epsilon_o = 0.0, \epsilon_{\ell} = 1.0$ ). This means that while the range information is used to inform the lookup in the distance map, the costs are always directly related to the labels. The second range-enabled scenario performs a weighted combination ( $\epsilon_o = 0.25, \epsilon_{\ell} = 0.75$ ) of both the semantic and traditional approaches.

In terms of the ray-based version of our approach, we use equation 18. This means there are no parameters to set.

Average Trajectory Error (m)						
Approach	RMSE	Mean	Median	Std. Dev.	Min	Max
AMCL	0.24	0.21	0.20	0.11	0.04	0.95
Range (Label Only)	0.19	0.16	0.14	0.10	0.02	0.55
Range (Combined)	0.22	0.19	0.17	0.11	0.04	0.62
<b>Rays</b> ( $\epsilon_{\rm G} = 3.0$ )	0.40	0.34	0.27	0.22	0.07	1.51
<b>Rays</b> $(\epsilon_{\rm g} = 7.0)$	0.58	0.45	0.38	0.37	0.02	2.23

(a) Room-Level Initialisation (b) AMCL



TABLE I: Room-Level Initialisation



(c) SeDAR (Ray-Based) Path (b) SeDAR (Range-Based) Path

Fig. 7: Qualitative view of Localisation in different modalities.

Instead, a mild ghost factor ( $\epsilon_{\rm g} = 3.0$ ) and a harsh one ( $\epsilon_{\rm g} =$ (7.0) are shown.

Since room-level initialisation is an easier problem than global initialisation, the advantages of the range-enabled version of our approach are harder to see compared to stateof-the-art. However, it is important to notice how closely the ray-based version of the approach performs to the rest of the scenarios, despite using no depth data. Apart from a couple of peaks, we essentially perform at the same level as AMCL. This becomes even more noticeable in table I, where it is clear that range-based semantic MCL (using only the labels) outperforms state of the art, while the ray-based  $\epsilon_{\rm g} = 3.0$  version lags closely behind. The reason  $\epsilon_{\rm g} = 3.0$ performs better than  $\epsilon_{\rm g} = 7.0$  is because small errors in the pose can cause the robot to "clip" a wall as it goes through the door. Since  $\epsilon_{\rm g} = 3.0$  is more lenient on these scenarios, it is able to outperform the harsher ghost factors. We will explore this relationship further in section V-D.

2) *Qualitative Results:* In terms of qualitative evaluation, we show both the convergence behaviour and the estimated path.

The convergence behaviour can be seen in figure 7. Here, figure 7a shows how the filter is initialised to roughly correspond to the room the robot is in. As the robot starts moving,

Fig. 8: Estimated path from room-level initialisations.

we can see how AMCL (7b), the range-based version of SeDAR (7c) and the ray-based version (7d) converge. Notice that while the ray-based approach has a predictably larger variance on the particles, the filter has successfully localised. This can be seen from the fact that the reconstructed Kinect pointcloud is properly aligned with the floorplan. It is important to note that although the Kinect pointcloud is present for visualisation in the ray-based method, the depth is not used.

The estimated paths can be seen in figure 8, where the red path is the estimated path and green is the ground truth. Figure 8a shows the state-of-the-art, which struggles to converge at the beginning of the sequence (marked by a blue circle). It can be seen that the range-based approach in figure 8b (combined label and range), converges more quickly and maintains a similar performance to AMCL. It only slightly deviates from the path at the end of the ambiguous corridor on the left, which also happens to AMCL. It can also be seen that the ray-based approach performs very well. While it takes longer to converge, as can be seen by the estimated trajectory in figure 8c, it corrects itself and only deviates from the path in areas of large uncertainty (like long corridors).

These experiments show that SeDAR-based MCL is capable of operating in a room-level initialised scenario. It is



Fig. 9: Semantic Floorplan Localisation, global initialisation.

now important to discuss how discriminative SeDAR is when there is no initial pose estimate provided to the system.

## C. Global Initialisation

We now focus on SeDAR-based MCL's ability to perform global localisation. In these experiments, the system is given no indication of where in the map the robot is. Instead, we use a maximum 50,000 particles (minimum 15,000) placed over the floorplan.

1) Quantitative Results: Figure 9 shows the same four scenarios as in the previous section. For the range-based scenarios (blue lines) it can be seen that using only the label information ( $\epsilon_o = 0.0, \epsilon_{\ell} = 1.00$ ) consistently outperforms the state of the art, both in terms of how quickly the values converge to a final result and the actual error on convergence. This shows that SeDAR used in an MCL context is more discriminative than standard occupancy maps in RMCL. The second range-based measurement ( $\epsilon_o = 0.25, \epsilon_{\ell} = 0.75$ ) significantly outperforms all other approaches. This is probably because, in principle, the occupancy maps can be considered another "label" in the semantic floorplan. This makes sense because setting  $\epsilon_o = 0.25$  is equivalent to weighting all labels equally, as it is a third of  $\epsilon_{\ell} = 0.75$  which is the weight of 3 labels.

In terms of the ray-based version of our approach (red lines), we compare two scenarios. A mild ghost factor ( $\epsilon_{\rm g} = 3.0$ ) and a harsh one ( $\epsilon_{\rm g} = 7.0$ ). These versions of the approach both provide comparable performance to the state-of-the-art. It is important to emphasise that this approach uses absolutely no range and/or depth measurements. As such, comparing against depth-based systems is inherently unfair. Still, SeDAR ray-based approaches compare favourably to AMCL. In terms of convergence, the mild ghost factor  $\epsilon_{\rm g} = 3.0$  gets to within several meters even quicker than AMCL, at which point the convergence rate slows down and is overtaken by AMCL. The steady state performance is also comparable. While the performance temporarily degrades, it

manages to recover and keep a steady error rate throughout the whole run. On the other hand, the harsher ghost factor  $\epsilon_{\rm g} = 7.0$  takes longer to converge, but remains steady and eventually outperforms the milder ghost factor. Table II shows the RMSE, error along with other statistics.

2) *Qualitative Results:* Similar to the previous section, we can provide qualitative analysis by looking at the convergence behaviour and the estimated paths.

In order to visualise the convergence behaviour, figure 10a shows a series of time steps during the filters' initialisation. On the first image, the particles have been spread over the ground floor of a  $(49 \times 49)$ m office area. In this dataset, the robot is looking directly at a door during the beginning of the sequence. Therefore, in figure 10b the filter converges with particles looking at doors that are a similar distance away. The robot then proceeds to move through the doors. Going through the door makes the filter converge significantly faster as it implicitly uses the ghost factor in the motion model. It also gives the robot a more unique distribution of doors (on a corner), which makes the filter converge quickly. This is shown in figure 10c, where the filter converges.

The estimated paths can be seen in figure 11, where the blue circle denotes the point of convergence. It can be seen that AMCL takes longer to converge (further away from the corner room) than the range-based approach. More importantly, it can be seen that the range-based approach suffers no noticeable degradation in the estimated trajectory over the room-level initialisation. On the other hand, the ray-based method's performance degrades more noticeably. This is because the filter converges in a long corridor with ambiguous label distributions (doors left and right are similarly spaced). However, once the robot turns around the system recovers and performs comparably to the range-based approach.

As mentioned previously, entering or exiting rooms helps the filter converge because it can use the ghost factor in the motion model. The following experiments, evaluate how the ghost factor affects the performance of the approach.

## D. Ghost Factor

The effect of the ghost factor can be measured in a similar way to the overall filter performance. We show that the ghost factor provides more discriminative information when it is *not* defined in a binary fashion. This is shown in the label-only scenario for both the range-based and ray-based approaches, in both the global and room-level initialisation.

1) Global Initialisation: Figure 12 shows the effect of varying the ghost factor during global initialisation. It can be seen that not penalising particles going through walls,  $(\epsilon_{\rm g} = 0)$ , is not a good choice. This makes sense, as there is very little to be gained from allowing particles to traverse occupied cells without any consequence. It follows that we should set the ghost factor as high as possible. However, setting the ghost factor to a large value ( $\epsilon_{\rm g} = 7.0$ ), which corresponds to reducing the probability by 95% at 0.43m, does not provide the best results.

Average Trajectory Error (m)						
Approach	RMSE	Mean	Median	Std. Dev.	Min	Max
AMCL	7.31	2.26	0.20	6.95	0.028	35.45
Range (Label Only)	6.71	2.59	1.31	6.20	1.15	38.60
Range (Combined)	4.78	1.69	0.69	4.47	0.43	31.19
<b>Rays</b> ( $\epsilon_{\rm G} = 3.0$ )	7.74	4.36	2.46	6.40	1.07	27.55
<b>Rays</b> $(\epsilon_{\rm G} = 7.0)$	8.09	4.49	2.22	6.73	1.61	28.47



TABLE II: Global Initialisation

Fig. 10: Qualitative view of Localisation in different modalities.



Fig. 11: Estimated path from global initialisations.

While it might seem intuitive to assume that a higher  $(\epsilon_{\rm a})$  will always be better, this is not the case. High values of the ghost factor correspond to a binary interpretation of occupancy which makes MCL systems unstable in the presence of discrepancies between the map and the environment. This happens because otherwise correct particles can clip door edges and be completely eliminated from the system. A harsh ghost factor also exacerbates problems with limited number of particles. In fact,  $\epsilon_{\rm a} = 3.0$ , corresponding to a 95% reduction at 1.0m, consistently showed the best results in all of the global initialisation experiments, as can be seen in table III.

2) Room-Level Initialisation: In terms of room-level initialisation, having an aggressive ghost factor is more in line with our initial intuition. Table IV shows that for both of the range-based scenarios,  $\epsilon_{\rm g} = 7.0$  provides the best results. This is because room-level initialisation in the presence of range-based measurements is a much easier problem to solve. As such, the problem of particles "clipping" edges of doors is a smaller issue.

On the other hand, the ray-based scenario still prefers a milder ghost factor of  $\epsilon_{\rm g} = 3.0$ . In this scenario, inaccuracies in both the map and the sensing modalities allow for otherwise correct particles to be heavily penalised by an aggressive ghost factor. Both of these results are reflected in figures 13a and 13b.

These results allow us to come to a single conclusion. The ghost factor must be tuned to the expected amount of noise in the map and sensing modality. Aggressive ghost factors can be used in cases where the pre-existing map is accurate and densely sampled, such as the case where the map was collected by the same sensor being used to localise



Fig. 12: Different ghost factors ( $\epsilon_{\rm g}$ ), global initialisation.

Average Trajectory Error (RMSE)							
<b>Ghost Factor</b> $(\epsilon_{\rm G})$	Range (Labels)	Range (Weighted)	Rays				
0.0	10.88	10.13	11.71				
3.0	6.71	4.78	7.74				
5.0	6.97	6.30	9.54				
7.0	7.19	6.10	8.09				

TABLE III: Global A for Different Ghost Factors

(*i.e.* SLAM). On the other hand, in the case where there are expected differences between what the robot is able to observe (*e.g.* furniture, scale errors, etc.), it is more beneficial to provide a milder ghost factor in order to be more lenient on small pose errors.

### E. Timing

The approach presented here makes the conscious decision to collapse the 3D world into a 2D representation. This has very noticeable effects to the computational complexity, and therefore speed, of the approach.

The speed of our approach was evaluated on a machine equipped with an Intel Xeon X5550 (2.67GHz) and an NVidia Titan X (Maxwell). We used OpenMP for threading expensive for-loops (such as raycasting). During room-level initialisation, or once the system has converged, our approach can run with 250 particles in 10ms, leaving us more than enough time to process the images from the Kinect into an SeDAR scan. Transforming the RGB images into semantic labels is the most extensive operation, taking on average 120ms. This means that a converged filter can run at 8-10fps. When performing global localisation, we can integrate a new sensor update, using 50,000 particles, in 2.25 seconds. This delay does not impact the ability of the system to converge, as most MCL-based approaches require motion between each sensor integration, meaning that the effective rate is much lower than the sensor output.

# VI. CONCLUSION

In conclusion, this work has presented a novel approach that is capable of localising a robotic platform within a known floorplan using human-inspired techniques. We first extracted the semantic information that is naturally present and salient in a floorplan. Our first novelty was using the semantic information present in a standard RGB image to extract labels and present them as a new sensing modality called SeDAR. The semantic information present in the



Fig. 13: Different ghost factors ( $\epsilon_{\rm g}$ ), room-level initialisation.

Average Trajectory Error (RMSE)							
<b>Ghost Factor</b> $(\epsilon_{G})$	Range (Labels) Range (Weighter		Rays				
0.0	0.25	0.27	1.20				
3.0	0.24	0.25	0.40				
5.0	0.22	0.24	0.70				
7.0	0.19	0.22	0.58				

TABLE IV: Room-Level A for Different Ghost Factors

floorplan and the SeDAR scan were then used in a novel semantic MCL approach. This approach presented three main novelties. First, the semantic information present in the floorplan was used to define a novel motion model for MCL Second, the SeDAR scan was used to localise in a floorplan using a combination of range and label information. Finally, SeDAR was used in the absence of range data to localise in the floorplan using only an RGB image.

In future work, we would like to explore several areas of interest. Using odometry estimated directly from the camera would be a priority. We would also like to use CNN-based depth estimation to both populate the range in SeDAR and inform the scale of the visual odometry. The Bayesian probabilities from the semantic segmentation can also be used to augment the sensor model. Another avenue of research would be to use the detected semantic labels that are *not* present in the floorplan to help localisation and/or augment the floorplan. Finally, we would like to integrate SeDAR to SLAM algorithms that operate on the scan-matching principle.

#### REFERENCES

- Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. pages 1–14, 2015.
- [2] Marcus A. Brubaker, Andreas Geiger, and Raquel Urtasun. Lost! leveraging the crowd for probabilistic visual self-localization. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 3057–3064, 2013.
- [3] Hang Chu, Dong Ki Kim, and Tsuhan Chen. You are here: Mimicking the Human Thinking Process in Reading Floor-Plans. In *Proceedings* of the IEEE International Conference on Computer Vision, pages 2210–2218, 2015.
- [4] Hang Chu, Shenlong Wang, Raquel Urtasun, and Sanja Fidler. Housecraft: Building houses from rental Ads and street views. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9910 LNCS, pages 500–516, 2016.
- [5] F Dellaert, W Burgard, D Fox, and S Thrun. Using the CONDENSA-TION algorithm for robust, vision-based mobile robot localization. *Computer Vision and Pattern Recognition*, 1999. IEEE Computer Society Conference on., 2:594 Vol. 2, 1999.

- [6] Frank Dellaert, Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Monte Carlo localization for mobile robots. In *International Conference on Robotics and Automation (ICRA)*, number May, pages 1322– 1328, Detroit, 1999. IEEE.
- [7] Maurice F. Fallon, Hordur Johannsson, and John J. Leonard. Efficient scene simulation for robust monte carlo localization using an RGB-D camera. *Proceedings - IEEE International Conference on Robotics* and Automation, pages 1663–1670, 2012.
- [8] Dieter Fox, Wolfram Burgard, Frank Dellaert, and Sebastian Thrun. Monte Carlo Localization: Efficient Position Estimation for Mobile Robots. *Aaai-99*, (Handschin 1970):343–349, 1999.
- [9] Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improved techniques for grid mapping with Rao-Blackwellized particle filters. *IEEE Transactions on Robotics*, 23(1):34–46, 2007.
- [10] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629, 1987.
- [11] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian SegNet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding. 2015.
- [12] Mathieu Labbe and François Michaud. Online global loop closure detection for large-scale multi-session graph-based slam. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 2661–2666. IEEE, 2014.
- [13] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016*, pages 239–248, 2016.
- [14] Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. RoomNet: End-to-End Room Layout Estimation. 2017.
- [15] Chenxi Liu, Alexander G. Schwing, Kaustav Kundu, Raquel Urtasun, and Sanja Fidler. Rent3D: Floor-plan priors for monocular layout estimation. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 07-12-June:3413–3421, 2015.
- [16] Kathia Melbouci, Sylvie Naudet Collette, Vincent Gay-Bellile, Omar Ait-Aider, and Michel Dhome. Model based RGBD SLAM. Proceedings - International Conference on Image Processing, ICIP, 2016-Augus:2618–2622, 2016.
- [17] Michael Paton and Jana Kosecka. Adaptive RGB-D localization. Proceedings of the 2012 9th Conference on Computer and Robot Vision, CRV 2012, pages 24–31, 2012.
- [18] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):640–651, 2017.
- [19] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images. In *Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2013.
- [20] Jurgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of RGB-D SLAM systems. *IEEE International Conference on Intelligent Robots and Systems*, pages 573–580, 2012.
- [21] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. CNN-SLAM: Real-time dense monocular SLAM with learned depth prediction. 2017.
- [22] Sebastian Thrun. Probabilistic robotics. *Communications of the ACM*, 45(3), 2002.
- [23] Shenlong Wang, Sanja Fidler, and Raquel Urtasun. Lost Shopping ! Monocular Localization in Large Indoor Spaces. Proceedings of the IEEE International Conference on Computer Vision, pages 2695–2703, 2015.
- [24] Wera Winterhalter, Freya Fleckenstein, Bastian Steder, Luciano Spinello, and Wolfram Burgard. Accurate indoor localization for RGB-D smartphones and tablets given 2D floor plans. In *IEEE International Conference on Intelligent Robots and Systems*, volume 2015-Decem, pages 3138–3143, 2015.
- [25] Jianxiong Xiao, Andrew Owens, and Antonio Torralba. SUN3D: A database of big spaces reconstructed using SfM and object labels. In Proceedings of the IEEE International Conference on Computer Vision, pages 1625–1632, 2013.