

This is a post-peer-review, pre-copyedit version of an article published in SN Applied Sciences. The final authenticated version is available online at: <http://dx.doi.org/10.1007/s42452-020-2001-3>

Feature-based Visual Simultaneous Localization and Mapping: A Survey

Rana Azzam, Tarek Taha, Shoudong Huang, and Yahya Zweiri

Received: date / Accepted: date

Abstract Visual simultaneous localization and mapping (SLAM) has attracted high attention over the past few years. In this paper, a comprehensive survey of the state-of-the-art feature-based visual SLAM approaches is presented. The reviewed approaches are classified based on the visual features observed in the environment. Visual features can be seen at different levels; low-level features like points and edges, middle-level features like planes and blobs, and high-level features like semantically labeled objects. One of the most critical research gaps regarding visual SLAM approaches concluded from this study is the lack of generality. Some approaches exhibit a very high level of maturity, in terms of accuracy and efficiency. Yet, they are tailored to very specific environments, like feature-rich and static environments. When operating in different environments, such approaches experience severe degradation in performance. In addition, due to software and hardware limitations, guaranteeing a robust visual SLAM approach is extremely challenging. Although semantics have been heavily exploited in visual SLAM, understanding of the

scene by incorporating relationships between features is not yet fully explored. A detailed discussion of such research challenges is provided throughout the paper.

Keywords Robotics · SLAM · Localization · Sensors · Factor Graphs · Semantics

1 Introduction

Following several decades of exhaustive research and intensive investigation, Simultaneous Localization and Mapping (SLAM) continues to dominate a magnificent share of the research conducted in the robotics community. SLAM is the problem of concurrently estimating the position of a robotic vehicle navigating in a previously unexplored environment while progressively constructing a map of it. The estimation is done based on measurements collected by means of sensors mounted on the vehicle including: vision, proximity, light, position, and inertial sensors, to name a few. SLAM systems employ these measurements in a multitude of various methods to localize the robot and map its surroundings. However, the building blocks of any SLAM system include a set of common components such as: map/trajectory initialization; data association; and loop closure. Different estimation techniques can then be used to estimate the robot's trajectory and generate a map of the environment. The implementation details of every SLAM approach relies on the employed sensor(s), and hence on the data collected from the environment. In this paper, we thoroughly review the most recent visual SLAM systems with focus on the feature-based approaches, where conventional vision sensors such as monocular, depth, or stereo cameras are employed to observe the environment. From here on, visual SLAM systems are referred to as monocular SLAM, RGB-D

Rana Azzam
Khalifa University of Science and Technology, Abu Dhabi,
UAE
Tel.: +971563123377
E-mail: rana.azzam@ku.ac.ae

Tarek Taha
Algorithma's Autonomous Aerial Lab, Abu Dhabi, UAE
E-mail: tarek.taha@algorithma.com

Shoudong Huang
University of Technology Sydney, Sydney, Australia
E-mail: Shoudong.Huang@uts.edu.au

Yahya Zweiri
Faculty of Science, Engineering and Computing, Kingston
University London, UK
E-mail: y.zweiri@kingston.ac.uk




Fig. 1: Different visual features extracted from the same visual frame. Left: low-level features (SURF [1]), middle: middle-level features (planes), right: high-level features (semantically labeled objects)

SLAM, or stereo SLAM if they employ a monocular camera, an RGB-D camera, or a stereo camera, respectively.

The non-conventional event-based vision sensor, such as the asynchronous time based image sensor (ATIS) [2] and the dynamic and active pixel vision sensor (DAVIS) [3], can also be used to solve the SLAM problem as proposed in [4], [5], and [6]. Its operation principle is biologically inspired, where instead of capturing frames at a set rate, it asynchronously captures events, which are timestamped changes in brightness of independent pixels. Due to its unique way of acquiring information from the environment, a paradigm shift is necessary to construct algorithms that accommodate such information. Event-based SLAM is beyond the scope of this review paper and interested readers are referred to the comprehensive survey in [7].

Some SLAM systems depend solely on visual measurements, while others augment them with different observations such as range or inertial measurements. Fusion of multiple types of observations might escalate the complexity of the algorithm, require more computational resources, and increase the cost of the platform. However, it makes the system more reliable, robust to outliers, and resilient to failures.

To choose the vision sensor suited for the developed visual SLAM system, the following should be considered. It is not possible to discern the scale of the environment based on observations from a single monocular frame. To compensate for that, monocular SLAM systems adopt different approaches to deduce the depth such as employing a set of one or more other sensors to obtain measurements from which the depth can be deduced, hypothesizing the depth of the observed features using neural networks for example, or by exploiting prior information about the environment, like the size of an observed feature. RGB-D cameras can provide information about depth from a single frame, but they are very sensitive to light, which may limit their appli-

cations or the environments in which they can successfully operate. Stereo cameras overcome the limitations of monocular and RGB-D cameras but they are more expensive and resource extensive. The choice of the vision sensor is also dependent on the robotic platform to be used. For instance, ground vehicles do not have any constraints with regards to the weight of on-board sensors, which makes all the options open. However, if an aerial vehicle is to be used, a monocular camera seems to be the most convenient option since it can be seamlessly accommodated on-board, due to its lightweight, small size, and low power requirements. Nevertheless, the employed algorithms must deal with the scale ambiguity of the obtained visual observations.

Visual measurements can be handled at different levels of detail. Direct SLAM systems, for example: [8],[9], and [10], process the intensities of all or a subset of pixels in the image. Then, based on the brightness consistency constraint [11], correspondences are established between multiple observations. Feature-based SLAM, on the other hand, targets features that exhibit distinctive properties and can be repeatedly detected by the employed detection algorithms. Examples of such systems include [12],[13], and [14]. Features can be classified into different levels; low-level features such as points, corners, and lines, medium-level features such as blobs and planes, and high-level features such as objects as illustrated in Figure 1. A visual SLAM system might employ a single ([15], [16], [17]) or a hybrid ([18], [19], [20]) of different feature levels.

In our review, we classify the state-of-the-art feature-based visual SLAM solutions based on the features used to perform localization and mapping. Within each category, implementation choices of the adopted SLAM pipeline are thoroughly discussed and compared. Strengths and weaknesses of each category are highlighted and open research problems are emphasized at the end.

1.1 Existing Surveys on SLAM:

The proposed approaches to SLAM were surveyed by several researchers in the field and the open research problems to-date were highlighted. In [21], the authors argued that SLAM is entering the robust perception era and thoroughly discussed the main characteristics of state-of-the-art solutions in terms of several performance metrics such as scalability, robustness, and representation. In addition, the paper addressed the recent advancements at the hardware and algorithmic levels and pointed out the research problems that are yet to be solved. A comprehensive review of key-frame based approaches to SLAM was presented in [11] where the general architecture of key-frame based monocular SLAM and the corresponding implementation approaches were presented. The survey conducted in [22] targeted SLAM approaches that omit the assumption that the environment under investigation is static and addressed the underlying techniques adopted to reconstruct a dynamic environment. Along the same lines, the survey presented in [23] studied the SLAM approaches that can operate in dynamic environments and those that employ heterogeneous data that can be obtained through a visual sensor, for instance: color, depth, and semantic information. Visual SLAM approaches that rely on observing primitive features in the scene were surveyed in [24] and classified according to the descriptors used for such features, emphasizing their strengths and weaknesses. An overview of the anatomy of visual odometry and visual SLAM, along with the underlying formulations and implementation choices was provided in [25]. Similarly, in [26], the solutions to visual SLAM were analyzed based on their implementation of the main building blocks of SLAM, and their failure in dynamic environments was analyzed. The SLAM approaches reviewed in [27] were classified into feature-based approaches, direct approaches, and RGB-D based approaches. Comparisons between the state-of-the-art solutions back in 2016 were conducted, followed by a set of open research problems relating to the mentioned categories. Finally, a recent survey on SLAM, with focus on semantics can be found in [28]. In this paper, we contribute a comprehensive survey of the most recent state-of-the-art feature-based visual SLAM systems and we classify the reviewed approaches based on the elements, i.e. features, they extract from visual frames to localize the robot and reconstruct the environment. Such features fall in one of the following categories: low-level, middle-level, or high-level features. So, the reviewed approaches are classified as shown in Figure 2. Our review serves as a thorough reference for researchers interested in investigating the various imple-

mentation options and advances in feature-based visual SLAM. Approaches that fall into the same feature-level category were further grouped based on other goals that they accomplish, like real-time performance, handling scene dynamics, and resilience to data association failures. The techniques that made each of these goals possible were listed and analyzed. This will assist the readers to accurately determine what makes out each of these approaches and what implementation methods they need to adopt and/or improve to develop a system that can achieve a particular set of objectives.




Fig. 2: Classification of feature-based visual SLAM approaches

The rest of this paper is organized as follows. The anatomy of a generic SLAM system is presented in Section 2 where the SLAM building blocks are discussed in detail along with different implementation options. The review and analysis of the feature-based visual SLAM systems and their design choices are provided in Section 3. In Section 4 we highlight the outcomes of our review and identify the open problems that need further investigation.

2 SLAM Building Blocks

Before delving into the implementation details of the current state-of-the-art solutions, the common components of visual SLAM are briefly discussed, including (1) Map/Trajectory Initialization, (2) Data Association, (3) Loop Closure, (4) Relocation, and (5) Estimation Algorithms, as shown in Figure 3. The purpose of each component is first provided, followed by the most prevalent implementation approaches, when applicable.

2.1 Map/Trajectory Initialization

Upon starting a robotic task in a new environment, a map of which is not available a priori, it is necessary to estimate the 3D structure of the surroundings as well as the position of the robot with respect to it. This serves as an initial assessment of the map that will be iteratively updated based on the sensory measurements collected throughout the task. This process is only required to bootstrap the system at startup. There are several ways in which initialization can be carried out




Fig. 3: SLAM Pipeline

when different sensors are employed. For instance, one depth frame or a stereo pair are sufficient to initialize a map, as presented in [29] and [14], since they provide depth and scale information, which monocular frames lack. On the other hand, initialization can be done manually when monocular cameras are in operation, for example [30], where the system is provided with prior information about the observed scene, which include the positions and appearance of four features, resolving the scale ambiguity problem. Examples of other algorithms that are commonly used for map initialization are iterative closest point (ICP) [31], [32], image alignment [33], [34], five-point algorithm [35] together with a model fitting algorithm such as random sample consensus (RANSAC) [36] or MLESAC [37], and inverse depth parameterization relative to the camera, which is used to parameterize observed features [38]. Kinematic models, for example [39], and integration of inertial measurements, as presented in [40], can be used to initialize the trajectory.

2.2 Data Association

While maneuvering in the environment, the robot may sense the same area multiple times. Establishing correspondences between the image frames, collected each time the same scene was observed, is of paramount significance to estimate the map and the robot’s trajectory, and is referred to as data association.

Feature-based approaches target *features*, which are areas in the image that exhibit distinctive properties. Features can be of different scales; low-level features such as geometric primitives, middle-level features such as super-pixels, or high-level features such as semantically labeled objects. The most critical characteristic of a feature is *repeatability*, which makes the feature detectable repeatedly when appearing in multiple frames taken from different viewpoints.

To detect features in an image, several detectors were proposed in the literature for different feature types.

For low-level features, such as points, lines, edges, and corners, Table 1 shows some examples of feature detectors as well as descriptors. After detecting a feature, it is extracted from the image together with its surrounding pixels, then assigned a quantitative measure, referred to as a descriptor, to facilitate matching with other features.

Table 1: Feature Detectors and Descriptors

Feature Detectors	Hessian Corner Detector [41]	Feature Descriptors	BRIEF [42]
	Harris Detector [43]		SURF [1]
	Shi-Tomasi Corners [44]		SIFT [45]
	Laplacian of Gaussian Detector [46]		HoG [47]
	MSER [48]		ORB [49]
	Difference of Gaussian [50]		FREAK [51]
	FAST/AGAST/OAST [52]		BRISK [53]

To detect planes in images, model fitting algorithms, such as RANSAC, are employed. It is also possible to combine modeling and a convolutional neural network (CNN) to identify planes, such as walls, in an image [54]. As for high-level features, several techniques were proposed for detecting objects and semantically labeling them in images including, but not limited to, conditional random fields (CRFs) [55], support vector machines (SVMs) [56], and deep neural networks (for example: single shot multi-box detector [57] and you only look once (YOLO) [58]).

Establishing correspondences between low-level features can be done between features in two images (2D-2D matching), between a point in the 3D map and its projection onto the image frame (3D-2D matching), or between two 3D points in the reconstructed map (3D-3D matching) [59] as depicted in Figure 4(a).

Matching a feature in the current image to a feature in another image (2D-2D) is performed by means of a search within a window in the second image enclosing the location of the feature in the current image. The search is reduced to one dimension if the transformation between both images is known and hence,




Fig. 4: Data association and loop closure examples

epipolar geometry [60] can be established. The similarity between the features’ descriptors can be measured using different quantities depending on their types, such as the sum of squared distance, L1/L2 Norms, or the hamming distance, to name a few. Such measures might hinder the performance of the system due to their high computational requirements and can be replaced by kd-tree search, similar to [61], or bag of binary words approaches such as [62].

3D-2D matching is necessary when the pose of the camera needs to be estimated given the 3D structure of the environment. 3D points surrounding a hypothesized pose are projected onto the current image frame. 2D projections are then matched to 2D features in the image using the previously mentioned techniques.

Upon re-visiting a location, i.e. closing a loop, the corresponding 3D landmarks are matched (3D-3D) yielding a corrected, drift-free path.

Establishing associations between middle-level features, such as planes, is done by comparing plane parameters, such as normals (for example: [19]), the overlap, and the distance between the plane detected in the current frame and those available in the map (such as [20]). If the distance is below a particular threshold, cor-

respondences are established. Otherwise, a new plane is added to the map.

In order to establish correspondences between semantically labeled landmarks, the predicted label is used to associate a detection with a landmark in the map. In case multiple instances of the same object category appear in the environment, a minimum distance threshold between them must be exceeded to consider inserting a new landmark into the map [18]. Otherwise, the detection is associated with its closest landmark. In a recently proposed SLAM solution [63], objects are detected and characterized at the category level rather than just the instance level. This is based on the fact that all objects in one class have common 3D points irrespective of their different categories. Other approaches to data association will be discussed in more detail in the next section.

2.3 Loop Closure

As the robot progresses through its task, errors from several sources accumulate causing the estimation to drift off the real trajectory (An example is illustrated in Figure 4(b)). Such drift may severely affect the reconstruction of the environment and hence lead to failure of the ongoing robotic task. To correct such drift, several techniques were proposed in the literature to detect loop closure, i.e. to detect whether or not the currently observed scene was assessed by the robot earlier, and hence achieve global consistency. Global consistency is the condition where the SLAM estimate matches, approximately, the ground truth and the reconstructed map conforms to the real topological structure of the observed area. However, local consistency refers to the case where the observations are matched locally but, perhaps, not globally [64].

Loop closures usually involve two main steps: visual place recognition and geometric verification. The former can be done using kd-tree search [65], bag of words approaches [66], Bayesian filtering [67], deep learning [68] [69], and visual feature matching [70], [71], while the latter can be achieved through image alignment, and RANSAC [36].

2.4 Re-Localization

Re-localization is the ability of a SLAM system to recover from a fatal localization failure in which the robot is assigned an arbitrary location. This failure can result due to several reasons, such as abrupt motions, motion blur, or absence of features [11]. Moreover, the

robotic vehicle might be re-positioned through an operation that is out of the robot’s control, in which case the robot’s global position is to be determined [72]. These cases are referred to as the Kidnapped Robot problem [73] and can be resolved using several techniques, including but not limited to, matching feature descriptors [74], re-observing semantically labelled objects [75], [76], epipolar geometry [77], or bags of binary words approach [12], [78],

2.5 Estimation Algorithms

Estimation algorithms are needed to resolve the SLAM constraints, and can be classified into batch and incremental algorithms. Batch algorithms, such as global bundle adjustment (GBA) [79] and full graph SLAM [80], process a large set of measurements collected by the robot, over a relatively large period of time, to reconstruct the map of the environment as well as the robot’s trajectory. Incremental algorithms, on the other hand, compute estimates of the map and trajectory upon arrival of new measurements. Some incremental algorithms, such as [81] operate on the entire set of measurements collected throughout the robotic task, while others, such as [82] operate on a subset of those measurements collected over a small time frame, which facilitates operation in an online manner. While batch algorithms succeed in achieving global consistency, they are computationally expensive, and hence, may impede real-time operation. In addition, due to the constrained memory resources, they might not work for large-scale environments or for continuously operating systems, which emphasizes the significance of incremental algorithms that do not suffer from such limitations. Revisiting old data association decisions is not possible when estimation is done through incremental algorithms that do not consider all measurements, which may increase the cumulative error compared to other algorithms. In what follows, batch algorithms, such as GraphSLAM [80] and GBA [79], as well as incremental algorithms, such as extended Kalman filter (EKF) [80], incremental smoothing and mapping [81], [82], and local bundle adjustment (LBA) [83], are briefly presented.

2.5.1 Extended Kalman Filter (EKF) [80]

Given multiple measurements recorded over a period of time, possibly from several sensors, an EKF estimates the state of the system under observation. The state of a system consists of the states of both the environment and the robotic vehicle. The former describes the poses of the landmarks observed in the environment, while the




Fig. 5: The Extended Kalman Filter Algorithm [84]

latter describes the vehicle’s kinematics. The estimation process involves filtering the noise associated with each measurement to reduce the overall uncertainty of the estimated state. Then, EKF estimates the states of the system through several iterations of predictions and updates based on the measurements collected from the environment, as depicted in Figure 5.

2.5.2 Factor Graph SLAM [80]

As the name of this algorithm suggests, a graph is used to reconstruct the map of an environment along with the robot’s trajectory in it. Map features and robot poses are represented as vertices, and are connected using edges that encode two types of nonlinear constraints: motion and measurements as shown in Figure 6(a). The summation of all the constraints makes SLAM a nonlinear least squares problem. To obtain an estimate that is globally consistent, all constraints are first linearized, yielding a sparse information matrix and an information vector. Due to the sparseness of the matrix and for a more efficient computation, the matrix is reduced in size using a variable elimination algorithm. An inference technique is then employed to find the assignment of poses to the nodes of the graph which minimizes the errors imposed by the constraints.

Alternatively, successive robot poses in the environment could be used alone to estimate the location of the robot using a Pose Graph [85]. The graph used in this problem includes the robot poses as nodes and motion constraints as edges between those nodes as depicted in Figure 6(b).

Bundle adjustment (BA) [86] is an instance of factor graph SLAM and can be defined as a refining process that simultaneously optimizes the 3D structure, the camera trajectory, and possibly its calibration parameters using a sequence of images collected from the environment as depicted in Figure 7. A cost function that assesses the error in the system is minimized to yield an improved estimation of the reconstruction. If




Fig. 6: Factor and pose graph examples

all measurements since the beginning of the robotic task were considered in the estimation, the process is referred to as GBA, and is known to be computationally expensive which hinders online operation [87].




Fig. 7: Bundle adjustment illustrative example

A more computationally efficient approach that incrementally adjusts the 3D reconstruction and camera trajectory was proposed in [83] and is referred to as

Local Bundle Adjustment (LBA). Only a window of n recent frames is adjusted upon reception of a new measurement. Using LBA makes it possible to execute SLAM in real time.

ParallaxBA is another variation of BA that was presented in [88] where features are parameterized using parallax angles instead of their Euclidean coordinates or inverse depth, ParallaxBA outperformed traditional BA in terms of accuracy and convergence.

2.5.3 Incremental Smoothing and Mapping

Incremental Smoothing and Mapping is an approach to SLAM that gradually computes estimations of the map and robot trajectory while measurements are being collected from the environment. Several approaches were proposed in the literature, the most prevalent of which are iSAM [89] and iSAM2 [90]. iSAM performs smoothing using QR factorization of the square root information matrix and iSAM2 operates on a novel data structure referred to as the Bayes tree which is obtained from factor graphs.

3 Feature-based Visual SLAM - Design Choices

In this section, an overview of the state-of-the-art feature-based visual SLAM systems is presented. As mentioned earlier, features can be of different granularities; low-level features, middle-level features, or high-level features. A visual SLAM system may be based on the use of either one or a hybrid of two or more feature types as will be discussed in the following sections. The most alarming concern about feature-based approaches is their failure in absence of features. Regardless of their high achievable performance and accuracy in feature-rich environments, if the environment under investigation lacks the features that visual SLAM relies on, be it points, planes, or objects, localization fails and the estimation of the robot's surroundings would not reflect the true structure. In what follows, visual SLAM systems are classified and discussed according to the types of features employed in the system.

3.1 Low-Level Feature-based Approaches

Low-level features are geometric primitives that are observable in abundance in textured scenes. The vast majority of the existing visual SLAM systems, for instance [12], [39], [91], [92], [93] exploit these features, throughout the localization and mapping processes, and have achieved a very high-level of maturity in terms of accuracy and efficiency. However, if the environment in

which the robot is operating is texture-less or lacks the features that the system can track, such methods fail due to the absence of features, hence why the most recent SLAM approaches started to consider features at different levels at the same time.

3.1.1 Multiple Feature Types to Aid Robustness:

Feature-based visual SLAM systems that depend on a single type of features are susceptible to failure when such features do not exist in the environment under investigation. To circumvent this issue, the work presented in [94] proposes using points and lines together to perform monocular SLAM in a poorly textured environment. Lines are parameterized by their endpoints to facilitate integration with point-based approaches. In absence of point features, this work proposes a novel technique to initialize the system using lines only. The same set of landmarks were also adopted to perform stereo SLAM in [95]. Stereo visual odometry is used to track points and lines, and Gauss Newton optimization is then employed to estimate the motion of the camera by minimizing the projection errors of the corresponding features. In [96], observations of point features are combined with laser scans and used in a factor graph to estimate the pose of the robot. A new map representation combining both an occupancy grid map and point features was proposed. By matching observed features to landmarks in the map, loop closure and localization can be achieved efficiently. Hence, the flexibility regarding what type of feature to adopt while estimating the robot’s trajectory in the environment greatly benefits the robustness of visual SLAM.

3.1.2 Facilitating Real-Time Performance

The maps generated by low-level features are sparse yet require large computational and memory resources. This is attributed to the fact that the process of detecting, extracting, and matching features is one of the most computationally expensive blocks in the SLAM pipeline.

In order to achieve real-time performance, some systems [14], [97], [98], [99] heavily exploit parallelism to perform tracking and mapping as originally proposed in PTAM [13]. Two threads are concurrently run to localize the robot and map its surroundings [13], [14]. Unlike tracking, delays are tolerable in the mapping thread where most of the heavy computations take place. To further reduce computations, [98] limited the number of features to be extracted, and used a local map through which feature matching is performed. In order to maximize parallelism, a separate thread was employed to




Fig. 8: Techniques to facilitate real-time performance

perform loop closing and a synchronization process was proposed where access to map points is granted to a thread only if the points are not currently being processed by another thread.

In [97], three parallel modules are employed; *scene flow* for feature detection, extraction, and matching, *visual odometry* for camera motion estimation, and *global SLAM* for loop closing and global consistency.

Localization and mapping can also be done in a distributed manner by multiple robotic vehicles while exploiting parallelism as proposed in [99] where tracking and image acquisition, which are lightweight processes, are run on-board all MAVs in parallel while mapping is done off-board by a powerful computer due to its computational demands. A recent monocular SLAM system was proposed in [100] where EKF and BA were exploited together to achieve real-time robust performance. ORB features and inertial measurements were used in a visual inertial odometry (VIO) framework based on EKF which is capable of estimating the camera motion with minimal delays. To further assist real-time performance, not all ORB features are extracted from visual frames in the VIO framework which operates on all incoming frames. In addition, to circumvent the estimation errors resulting from EKF, a globally consistent map estimated using BA is frequently updated based on selected keyframes and fed to EKF to correct any estimation errors. The selected keyframes go through another round of feature extraction and matching since the features extracted for VIO are not sufficient for building a robust map. Loop closure is run in a parallel thread to correct accumulated error by performing place recognition and ORB feature matching. Once a loop is detected, pose graph optimization as well




Fig. 9: Techniques to resolve scale ambiguity

as GBA are carried out. Due to the fusion of visual and inertial measurements, the approach is robust to abrupt motions and is capable of resolving scale ambiguity. It also combines the advantages of EKF and BA to achieve real-time performance and robustness respectively.

Figure 8 summarizes the techniques that can be adopted to speed up the localization and mapping processes and get the estimation done in real-time.

3.1.3 Resolving Scale Ambiguity

When using a monocular camera, a SLAM system needs to handle the inherent scale ambiguity challenge which results from the difficulty to discern depth from a single frame. An EKF based approach was proposed in [91] where scale ambiguity and intermittent absence of features are compensated for by fusion of monocular vision, ultrasonic and atmospheric pressure measurements. Fusion of multiple sensors was also seen in [101] where vision, inertial, and range measurements were employed to achieve the objectives of SLAM. Scale ambiguity in [77] is circumvented by two-view initialization. A pair of images is selected according to their relative rotation, Euclidean distance, and the time difference between them. Then, epipolar geometry is used to estimate the scale based on the matched features between these frames. In another monocular SLAM approach [102], the depth of ORB features was computed based on their distance to the vanishing points identified in the scene. Furthermore, inverse depth parameterization was used in [103] to recover the scale of the scene.

While not required for RGB-D and stereo SLAM, adopting a technique to resolve the scale of the map is essential for monocular SLAM. Figure 9 illustrates the techniques that can be used to resolve scale ambiguity.

3.1.4 Resilience to Feature Detection/Association Failure

Failure to observe or match low-level features in an environment is equivalent to operating in texture-less environments in which feature-based visual SLAM systems fail. In both cases, the system suffers from absence of measurement constraints, causing severe performance degradation. A vision system fails to detect or match features between frames in case of abrupt sensor motions or in presence of dynamics in the scene.

One of the limitations of the original EKF-SLAM, which is described in [80], is its inability to handle abrupt motions. To overcome this, the approach proposed in [105] employs visual input in both phases of the filter; prediction and update. Optical flow and epipolar geometry are used to estimate the state transition of the camera. Using images in the prediction stage made the system robust against abrupt motions and infrequent data acquisition. This has also eliminated the need for dynamic models and resulted in a faster and more efficient performance. Although this EKF variant improves the robustness and efficiency of SLAM in particular cases, it still fails if there are no features in the scene. Another variation of EKF-SLAM is proposed in [109] IMU measurements are used in the prediction phase and RGB-D images are used in the update phase. To achieve global consistency, pose graph optimization is performed. Fusion of IMU measurements made it possible for the system to successfully operate in texture-less and dynamic environments.

ORB-SLAM2 [12] is a state-of-the-art visual SLAM system that performs tracking, mapping, and loop closing based solely on ORB features in real time while running on standard CPUs. Due to its dependence on visual features, ORB-SLAM2 fails in absence of ORB features in the scene. To this end, a tightly-coupled fu-

Table 2: Implementation choices adopted by Low-level Feature-based Approaches. Abbreviations: factor graph (FG), pose graph (PG) bundle adjustment (BA), extended Kalman filter (EKF), maximum-a-posteriori(MAP), vocabulary tree (VT), particle filter (PF), feature matching (FM), bags of binary words (BBW), Gauss Newton (GN), Levenberg Marquardt (LM), first frame (FF), prior map (PM), stereo initialization (SI), Epipolar geometry (EG), Inverse depth parameterization (IDP), Odometry (Odom), small blurry image relocation (SBI), least square optimization (LSO), image moment invariants (IMI)

Ref	SLAM Component				
	Initialization	Data Association	Loop Closure	Estimation	Relocation
[101]	IMU measurements	Sweep matching & approximate nearest neighbor	LSO	Algorithm in [104]	-
[91]	Range measurements	Visual measurement models	-	EKF	-
[39]	Kinematic model	Geometric matching then ICP	-	EKF	-
[92]	-	FM	BBW	FG	BBW
[93]	SI	FM	VT	RANSAC & non-linear refinement	-
[77]	SI	Optical flow	IMI	GBA	EG & BA
[94]	Trifocal tensor eq.	Relational graph strategy	Essential graph	FG and GBA	EPnPL
[105]	IDP	FM	-	EKF	-
[13]	SI & RANSAC	FM	Mapping module	GBA	-
[97]	-	VT and EG	VT	GBA	VT
[106]	PM	Edge alignment using GN	-	GN Optimization + EKF	-
[14]	SI	FM	BA	LM algorithm	-
[98]	SI	FM	PG	Co-visibility graph	-
[12]	FF	2D-3D points matching	BBW	FG and GBA	BBW
[107]	FF	2D-3D points matching	BBW	FG and GBA	Odom
[99]	SI+ RANSAC	FM	GBA	GBA	SBI [108]
[109]	FF	FM	FM	FG	-
[110]	Odom	FM	BBW	FG	-
[111]	PM	Multi-hypotheses via PF	-	PF	-
[112]	PM	Distance function	-	EKF	-
[113]	FF	FM	FM	Efficient PnP and BA	-
[100]	VIO EKF	ORB FM	PG & GBA	EKF & BA	-

(-) indicates that implementation details about the corresponding element/block are not provided

sion of odometry and ORB-SLAM2 was proposed in [107] where the motion model is replaced by odometry, which supports the estimation when no features can be detected in the scene.

Similarly, the approach proposed in [92], exploits tightly-coupled fusion of inertial and visual measurements to perform visual inertial odometry. Global consistency is then achieved by means of loop closure detection and global pose graph optimization. Another variation of ORB-SLAM2 can be found in [114], where ORB features were replaced by learnt point features, referred to as GCNv2. It was demonstrated that the proposed approach has comparable performance to ORB-SLAM2 in most scenarios, but performs slightly better in case of fast rotations.

Failure to associate features across subsequent frames can also result from dynamics in the scene. The work proposed in [113] demonstrates the ability to successfully perform RGB-D SLAM in a dynamic environment while observing low level features only. Using the fundamental matrix, feature points belonging to moving

parts of the scene are extracted. Then, efficient PnP was used to estimate the pose of the camera in the environment. The re-projection errors are then further optimized by means of BA. The proposed approach was successfully used in real experiments but only under the assumptions that there is small parallax and more than 24 point matches between consecutive frames. Hence, the approach fails to work in presence of abrupt motions and in absence of low-level features in the environment.

In order to enhance the performance of visual SLAM in dynamic environments, the approach proposed in [115] employs a sparse motion removal scheme. A Bayesian filter is used to compute the similarities and differences between consecutive frames to determine the dynamic features. After eliminating such features, the scene is fed to a classical visual SLAM approach to perform pose estimation. This approach works only in presence of features in the scene and fails otherwise.

Another low-level feature based visual SLAM approach that is robust to false data association occurring in dynamic scenes is found in [116]. The approach




Fig. 10: Illustration of reconstructed map based on planar features

is based on a novel filter where poses are encoded as dual quaternions. Association of ORB feature observations and map landmarks is done through an optical flow-based approach which makes it robust to dynamics in the scene.

In summary, lack of features, abrupt camera motions, and dynamics in the observed scene are the main reasons behind failure to perform data association. Some techniques that are adopted in the literature to solve these issues include employing multiple sensors that observe different information in the scene and eliminating observations that involve dynamics.

The implementation details of the reviewed low-level feature-based approaches are provided in 2.

3.2 Middle-Level Feature-based Approaches

Middle-level features are planes or blobs that are observed in the environment. Using such features as landmarks improves the SLAM performance in texture-less environments where it is challenging to observe low-level features; in corridors for example. To observe such features, model fitting approaches are employed. Hence, there is a trade-off between the estimation accuracy and the time needed to compute accurate models from the environment. Using those features alone is not common since fusing them with low- and high-level features results in better accuracy as discussed in Section 3.4. In [117], a SLAM approach based solely on RGB-D data is proposed. A 3D map of the environment is constructed using planes representing walls and floors while removing all other objects from the scene. RANSAC is employed to estimate planar surfaces which are then refined by estimating their normals and extracting the corresponding convex. Then, an l_0 norm minimization algorithm is used to maintain planes that are highly

likely to represent walls or floors while minimizing the inclusion of smaller ones. Using this approach, it was possible to reconstruct a map of the walls and floor as illustrated in Figure 10. However, no other features are present in the map which makes it unusable for the majority of SLAM application. This motivates the need for considering high-level features, as presented in the next section.

3.3 High-Level Feature-based Approaches

Perceiving high-level features is paramount when robots are expected to perform tasks that require scene understanding such as searching for a victim after a catastrophe, building meaningful maps, and grasping or operating on particular objects in the environment. This is very challenging to achieve with maps reconstructed using low-level features since they lack expressive representation which makes it harder for humans to understand [118], [119]. High-level features add critical information about the structure of the scene and convey the semantic meaning of every part of the reconstructed map. They are environment-specific and may vary in size, shape, and dynamicity. In a city-scale application, possible landmarks include trees, buildings, streets, or sidewalks. On the other hand, furniture, office supplies, and home appliances may serve as landmarks for indoor applications. In this section, different approaches to data association in high-level feature-based SLAM approaches will be thoroughly discussed. Then, techniques to achieve real-time performance and handle dynamics in the scenes will be presented.

3.3.1 Associating High-Level Feature Observations With Landmarks

Although high-level features are detected and semantically annotated, data association in the event that multiple instances of the same object category exist in the environment poses a fundamental challenge in high-level feature-based visual SLAM systems [15].

In [63], objects are detected and characterized at the category level rather than just the instance level. This is based on the fact that all objects in one class have common 3D points, irrespective of their different categories. Such points, referred to as keypoints, are used to distinguish between the different categories of the same class. Input monocular frames are passed to an object detector, YOLO9000 [120], and 3D keypoints in the resulting bounding boxes are localized by means of another convolutional neural network. Shapes and poses are optimized using the Ceres solver. Instead of performing object and keypoint detection on every




Fig. 11: High-level features data association techniques

frame, objects are tracked in successive frames leading to higher efficiency and speed.

Another novel data association approach was presented in [17] to localize a robot in a prior map. First, a query graph is computed for each image where a vertex represents an object's class and centroid, and undirected edges between vertices indicate the fulfillment of a proximity requirement. A merged graph for all the images is then created by connecting vertices from consecutive images using the Euclidean distances between them. Vertices that are too close to each other are merged to avoid duplicates. The second step is the generation of a random walk descriptor for each vertex. That is, an $n \times m$ matrix containing the labels of m visited vertices in n random walks. Third, the query graph is to be matched to the global database graph based on a similarity score. The similarity score of two vertices indicates the number of identical rows in their descriptors. The highest k matches are then used to localize the query graph in the database graph.

In [121], semantically labeled objects as well as their inter-relationships are employed in the process of es-

tablishing correspondences between input monocular frames. An RGB frame is first passed to a Faster R-CNN to detect objects. Then, the transformation between consecutive images is computed by first generating multiple cuboids that lie along the line, formed by the camera center and the center of the bounding box, and projecting them onto the detected bounding box. Generation of cuboids is done at discrete distances and angles. After that, coordinate descent is performed to minimize the difference between the corners of the detected bounding box and the projection of each cuboid into the image plane. Redundant cuboids are then removed. Each of the remaining cuboids is then used as a seed to generate a scene, which is a set of cuboids each corresponding to a detected bounding box, based on contextual constraints. To find correspondences between the generated set of scenes, a sampling-based approach is used. Every pair of scenes is searched for correspondences based on semantic labels. Three correspondences from every pair are picked and frames of reference for each scene are constructed. The transformation between the scenes is computed accordingly and scored

based on how well the remaining correspondences fit using the computed transformation. The sample with the highest rank is then used to estimate the transformation between camera poses.

In [122], an object hypothesis is generated if the same object segment is observed in multiple frames and is represented using 3D feature descriptors which facilitate loop closure. Inlier correspondences between the current object and the objects in the map are computed, then, the object is associated with the hypothesis with which it achieved the highest number of correspondences. If the number of correspondences falls below a threshold, a new object representation is added. Only one or a few static instances of an object category are assumed to be in the environment. A prior estimation of the robot pose based on odometry and ICP is computed using OmniMapper [123]. Based on that, the current frame's segments are projected into a common frame of reference with all previously segmented objects. The centroid of each segment is matched to the closest segment centroid in the map. To verify the match, the bounding box of the current segment and that of the segment to which it was matched are compared. If there is not enough overlap between the bounding boxes, a new object is initialized. The final object model is created by aggregating all the corresponding segments after transforming them according to the relative camera poses. Spatial constraints between the object model and the robot poses are then added to the SLAM system.

In [15], SLAM and data association are addressed as tightly coupled problems and a novel approach is proposed to simultaneously estimate a robot's position and associate its observations with landmarks. A back-end approach was used to jointly solve the object detection and SLAM problems. After being detected, an object is represented by the centroid of its point cloud obtained from RGB-D data. Neither data association nor the total number of landmarks in the environment are known a priori. A probabilistic model based on the Dirichlet Process was hence introduced to establish proper data association. Overall, a mixed-integer nonlinear problem is set up to estimate the robot poses, landmark locations, and data association given the robot's relative poses and observations.

The most common approach to data association in presence of multiple instances of the same object category is the distance threshold as presented in [16]. Each robot in the proposed distributed SLAM framework performs SLAM through OmniMapper [123] based on visual and odometry measurements. Each input RGB image is passed to a YOLO object detector. Detected objects are segmented and the PFH/RGB features in their point cloud and in the corresponding model are

extracted and matched. If at least 12 correspondences were detected, generalized iterative closest point (GICP) [124] is performed to compute a refined pose of the object. Data association is then performed by searching for instances of the same detected object category within a distance threshold. Figure 11 summarizes the main approaches found in the literature to perform data association of high-level features.

3.3.2 Facilitating Real-Time Performance

Performing real-time localization and mapping is very critical for some robotic tasks, especially those performed in harsh environment for search and rescue purposes. However, the processing time for some blocks in the pipeline, such as object detection and segmentation, goes beyond that. In this section, focus will be devoted to the techniques used to facilitate real-time performance in high-level feature-based SLAM approaches.

The work proposed in [63] proposes not performing object detection on all the incoming frames. Rather, after detecting an object in a keyframe, it is tracked in successive frames, which significantly reduces the time needed to process the data.

For the same purpose, the system proposed in [122] pre-processes the scene by dividing it into planar and non-planar (object) segments. After removing planar segments, object segments are refined and associated to already existing landmarks in the map.

Representing objects using quadrics is an alternative technique to reduce computations while employing semantically labeled landmarks in a visual SLAM system. The work proposed in [125] uses an object detector as a sensor where the detected bounding boxes are used to identify the parameters of the quadric representing the corresponding object. A quadric provides information about the size of the object, its position, and its orientation, encoded in ten independent parameters. A geometric error formulation was proposed to account for the spatial uncertainty of object detections, resulting from occlusions for example. Using quadrics instead of detailed object models enhances the speed of the system at the expense of reconstructing information-rich maps which are useful in a wide range of applications. An illustration of the discussed techniques that aid the efficiency of high-level feature-based visual SLAM systems is provided in Figure 12.

3.3.3 Handling Dynamics In The Scene

The majority of SLAM systems are developed under the unrealistic assumption that the environment is static. Only a few systems were proposed in the literature




Fig. 12: Techniques to achieve real-time performance by high-level feature-based visual SLAM approaches

where scene dynamics are accounted for. Most of these systems detect the non-stationary parts of the observed scene, eliminate them, then perform SLAM based on the remaining static environment. An example of such an approach can be found in [126] where moving objects were tracked and stationary ones were used to generate a static map of the environment under investigation. Observations were done using a laser scanner and data association was carried out using a multi-level RANSAC approach.

Differently, the work presented in [127] uses cuboids as objects' representation, where an object SLAM system is proposed. The system relies on observations from a monocular camera and exploits dynamic objects in the scene to improve localization by adding motion model constraints to the multi-view BA formulation that is used to solve the optimization problem. Objects and feature points belonging to them are tracked in successive frames and motion models are estimated and employed to improve the accuracy of trajectory and map estimation.

Exploiting motion models of dynamic objects rather than ignoring them imposes additional constraints on the systems and hence improves the accuracy of the estimation.

A summary of all the reviewed high-level feature-based approaches in the previous sections is provided in Table 3.

3.4 Hybrid Feature-based Approaches

In the previous sections, SLAM systems that employ features of a single type were discussed and analyzed. Features at each level enhance the outcome of SLAM in a distinct way. For instance, localization methods based on observation of low-level features have achieved a high level of maturity in terms of accuracy and efficiency. The maps they produce, however, are of high

sparsity without any semantic indications. Taking advantage of middle-level features, such as planes, in the scene makes it possible to attain higher reconstruction density as well as more robustness in texture-less environments. To create meaningful maps that humans can easily perceive, recent SLAM approaches make effective use of the emerging object detection techniques and employ semantically labeled observations throughout the localization and mapping processes. To make the most out of what can be visually observed in a scene and to enhance their overall outcome, SLAM systems have lately started to employ features at two or more levels as discussed in this section. In this section, feature-based visual SLAM approaches that adopt features from multiple levels will be reviewed. The reviewed systems are classified based on the features used to perform SLAM into three categories; low- and middle-level feature based approaches, low- and high-level feature based approaches, and low-, middle-, and high-level feature-based approaches. Table 4 presents a summary of the implementation choices adopted by the reviewed approaches.

3.4.1 Low- and Middle- Level Feature-based Approaches

The systems presented in [20], [129], [130], and [132] employ low and middle-level features to achieve the objectives of SLAM.

In some environments, such as corridors, plane SLAM becomes unconstrained. Fusing planes and points can greatly enhance the robustness of SLAM in such environments as proposed in [20] where planes, detected in monocular frames using a pop-up 3D model, are used to estimate the camera trajectory and the 3D map of the environment. Across different frames, planes are associated based on a weighted sum of three quantities: the difference between their normals, the distance between them, and the overlap between their projections. For each incoming monocular frame, ORB descriptors are computed and a bag of words approach is used to detect loops. Upon detection of a loop, corresponding plane pairs are determined and the factor graph is modified accordingly.

Geometric primitives and planes were employed differently in [129]. A least-squares optimization using a graph formulation, where planar constraints are involved, is used to solve the SLAM problem. The detected points are constrained to belong to a particular plane, parameterized by its normal and depth with respect to the camera, in the environment. Angles between planes in the environment are also considered as constraints. All constraints are coupled into a cost function and the resulting non-linear least squares problem is solved.

Table 3: Implementation choices adopted by high-level feature-based SLAM approaches

Ref	SLAM Component			
	Initialization	Data Association	Loop Closure	Estimation
[15]	Open loop predictions	A probabilistic model based on the Dirichlet process	Data association approach	Factor graph
[16]	-	Object class/label within a distance threshold	objects detected by multiple robots	Factor Graph
[17]	Prior Map	Matching of Random-Walk Descriptors	-	Graph-based technique then Maximim A Posteriori
[126]	-	Multi-level RANSAC and Mahalanobis distance	ML-RANSAC	EKF
[121]	Relative poses are estimated then refined using coordinate descent	Contextual Relevance based on the face centric geometric descriptors	Sampling-based approach to search for correspondences	Coordinate descent
[128]	-	Object class/label and greedy object tracker	Hungarian Algorithm (Cost is estimated object shape and pose)	Factor Graph
[127]	Object detections used to initialize depth of points	Matching point features belonging to detected objects, 2-D KLT sparse optical flow algorithm for tracking dynamic points, and visual object tracking for dynamic objects	No explicit loop closure module used	BA
(-) indicates that implementation details about the corresponding element/block are not provided				

A third variation was proposed in [130] where an RGB-D SLAM approach based on planes and points was proposed. Each incoming image is divided into intervals, then labeled, based on the planes present in it. The orientation of a frame is estimated based on the orientation of the most dominant plane in it while translation between frames is computed based on matched SIFT features and RANSAC. Global alignment and loop closure are carried out based on a fusion of the low- and middle-level features, which aids the robustness of the proposed approach.

A recent RGB-D SLAM was proposed in [131] where points and planes are exploited to estimate the pose of a camera and a map of its surroundings. ORB features are extracted from RGB frames and handled by the RGB-D version of ORB-SLAM2. On the other hand, depth frames are used to extract planes, along with their contour points from the scene. Contour points are employed to construct spatial and geometric constraints between planes in the reconstructed map. A novel data association technique for planes was used, where the angle between two planes was used to judge whether they are perpendicular or parallel, while accounting for measurement noise. Two planes are matched if the distance between the observed plane's point and the plane in the map is below a particular threshold. Imaginary planes that are perpendicular to planes appearing in the scene are also exploited and treated as the other observed features in the pose estimation process. A factor graph is constructed and solved by means of the Levenberg-

Marquardt optimizer. The proposed plane data association method is more robust than approaches considering plane normals and/or plane distances because it takes into account the measurement noise, which is inevitable.

Super-pixels are middle-level features seen as planar regions exhibiting similar intensities in input frames. Employing super-pixels comes with the advantage of being able to reconstruct poorly-textured scenes. However, there isn't a robust descriptor of such features, which makes it hard to match them in different images. In [132], a feature-based monocular SLAM approach was proposed, integrating super-pixels with PTAM, where PTAM keyframes are divided into super-pixels of irregular sizes. The map state, that is to be estimated, consists of the pose of all keyframes, the Euclidean coordinates of point features, and the parameters of the planar super-pixels. Two keyframes, the pose of which is already computed using PTAM, are used to initialize a super-pixel. All super-pixels in the keyframes are extracted and matched using a Monte Carlo approach. BA is used to optimize the camera and 3D points' states, which are then used to estimate the parameters of the super-pixels. On every new keyframe, all super-pixels are re-projected to search for matches. When the re-projection error drops below a threshold, the match is added to the optimization problem as a constraint.

Another work that exploits the fusion of point features and planar regions, represented as squared fiducial markers in this case, in an environment can be found in

Table 4: Implementation choices adopted by hybrid feature-based SLAM approaches. Abbreviations: factor graph (FG), bundle adjustment (BA), feature matching (FM), bags of binary words (BBW), bag of words (BoW) first frame (FF), Inverse depth parameterization (IDP), least square optimization (LSO), nearest neighbor search (NNS)

Ref	SLAM Component					Approach		Features Level		
	Initialization	Data Association	Loop Closure	Estimation	Relocation	Filter	Key-frame	Low	Mid	High
[20]	Ground truth pose	Plane normals' difference, distance between planes and projection overlap	BoW	FG	-	x	x	✓	✓	x
[129]	-	Points belonging to planes are known a priori	-	FG	-	x	x	✓	✓	x
[130]	FF	Joint compatibility branch and bound test	FM	FG	-	x	✓	✓	✓	x
[131]	FF	ORB FM and angles between planes	BoW	FG	FM	x	✓	✓	✓	x
[132]	Monte Carlo Approach, SI & RANSAC	Super-pixel contour re-projection	Done by the mapping module	GBA	-	x	✓	✓	✓	x
[75]	Triangulation	FM between input image and object model using k-d tree search	object matching	GBA	object matching	x	✓	✓	x	✓
[133]	-	SIFT FM + object detection	Implicitly using object detection	Semantic-BA	-	x	x	✓	x	✓
[76]	-	Matching detected text in door signs	Matching detected text in door signs	FG	Matching detected text in door signs	x	x	✓	x	✓
[29]	FF	K-d tree search based on Euclidean distance between 3D points of detection and landmarks	BBW	FG & GBA	BBW	x	✓	✓	x	✓
[103]	IDP	FM then RANSAC	Visual recognition	EKF	Real Time SLAM Relocalization [134]	✓	x	✓	x	✓
[78]	Cheap optimization of the camera position	FAST FM	-	GBA	BBW & RANSAC	x	✓	✓	x	✓
[74]	FM	NNS in a Kd tree and Place Recognition	NNS in a Kd tree and Place Recognition	non-linear LSO	Matching DAISY Descriptors	✓	x	✓	x	✓
[18]	-	Expectation maximization for soft data association	Mahalanobi's distance to semantic measurements	FG	-	x	✓	✓	x	✓
[19]	-	ORB FM, semantic labels & NNS, distance & normal difference between planes	BoW	FG	-	x	✓	✓	✓	✓
[127]	Depth of points initialized based on objects and planes' detections	ORB FM, and matching points of objects and planes	FM	BA	-	x	✓	✓	✓	✓
[135]	Homography or essential matrix and Marker-based initialization [136]	FM, markers matching	Marker detection using a BoW approach	BA	Marker detection OR RANSAC PnP approach	x	✓	✓	✓	x

(-) indicates that implementation details about the corresponding element/block are not provided
(✓) indicates that the corresponding option is adopted
(x) indicates that the corresponding option is not adopted

[135]. Besides the robustness achieved due to employing point features, utilizing fiducial markers in this system comes with several advantages such as eliminating scale ambiguity, robustness in repetitive environments where distinguishing point features can be challenging, and feature invariance over time.

3.4.2 Low- and High-Level Feature-based Approaches

A multitude of different SLAM approaches were proposed based on the use of a combination of low- and high-level features in [75], [133], [76], [29], [103], [78], [18], [74], and [137]. Such approaches demonstrate high-level expressiveness while maintaining robustness.

The system proposed in [75] does tracking, object recognition, and mapping while mainly operating on monocular RGB frames. Frames exhibiting distinctive geometrical and/or semantic information are selected as keyframes. A semantically labelled object is added to the map after being detected in multiple frames that contain at least 5 point correspondences, have a minimum parallax angle of 3° , and must exhibit acceptable geometric conditioning. To distinguish between instances of the same object model in the scene, the pose of the detected instance in the world frame is hypothesized given the map scale, and the overlap with previously detected instances is computed. If no overlap is

detected, a new object instance is added to the map. If the scale of the map is not yet known, objects detected sequentially are assumed to belong to the same object instance in the map. Correspondences are established between measurements and object models using a k-d tree search. For more robustness, ORB features in input images are computed and 2D-3D correspondences are established.

Instead of employing low-level features independently, geometric features can be used to detect objects in the scene as proposed in [133] where object detection and SLAM were done jointly for 2D and 3D sensors using a novel BA formulation, referred to as Semantic BA. Upon reception of a new image, features are extracted and matched to those in the objects model database. A validation graph is then created for each set of correspondences to an object. The frames in which the features are matched along with the model from the database are then transformed into a common pose and the cost of the corresponding semantic feature is the re-projection error of the detected features weighted by the confidence of matches. In the 3D case, when an object is detected multiple times, the cost function of the semantic edge includes the re-projection of one detected feature into the other. Frames in which features were matched to a common point in the model are said to have a virtual match represented by an edge in the graph. For consistency purposes, geometric constraints obtained from SLAM are added to the graph. The resulting validation graph is optimized to obtain the minimum re-projection error for all constraints.

In some environments, such as educational entities and hospitals, each room is assigned a unique identifier which can serve as a landmark in a SLAM system as presented in [76]. After eliminating the points that corresponds to walls, a door-sign detector, based on an SVM classifier, is employed. Characters contained in a door sign are recognized using Optical Character Recognition (OCR). Lines extracted from laser data along with measurements from the door-sign detector are then passed to a mapper to map the environment.

Observations of generic objects were used to extend RGB-D ORB-SLAM2 in [29]. Objects are detected, segmented, and associated to landmarks in the map by means of a k-d tree. The pose of the objects is determined using ORB-SLAM. Detected objects are stored with three pieces of information: the RGB point cloud of the object, their pose from ORB-SLAM, and the accumulated detection confidence. The class label is determined based on the entire history of detection of an object. A sparse map of the environment can be built explicitly by projecting the point cloud based on the latest trajectory estimate. Finally, object points are in-

serted into the SLAM state vector as Euclidean coordinates and hence are tracked and further refined upon reception of new data in the following frames.

EKF-Monocular-SLAM, Structure from Motion (SfM), and Visual Recognition were combined in the system proposed in [103]. Objects are detected by associating SURF points in an images to object models in a database. Such associations are then geometrically verified using RANSAC. Afterwards, the PnP algorithm or DLT algorithm are employed to compute the transformation or Homography matrices for non-planar and planar models, respectively, which are then used to refine the pose of the object. Matched points are fed into the monocular SLAM module which is based on EKF-Monocular-SLAM where the state vector to be estimated consists of the camera motion parameters and the point features along with the geometry of the detected objects.

On a different note, some scenes in the environment under observation may exhibit dynamicity which if not accounted for, hinders the overall performance of SLAM systems. Hence, most SLAM systems assume a scene where objects remain static throughout the localization and mapping processes. The SLAM system presented in [78] eliminates this assumption by removing dynamic objects from the observed scene before operation. More specifically, every RGB-D frame is processed to mask out regions in which a person was detected using an RGB-D based method [138]. The remaining data image a static environment which can be processed using a standard visual SLAM algorithm. A similar approach can be found in [137] where dynamic objects are segmented out of the scene by means of a computationally efficient step-wise approach to detect the object and extract its contour. The static environment is then mapped based on point features using a novel look-up table approach that targets using a large amount of distinct, evenly-distributed point features from the environment, which enhances the accuracy of mapping and localization.

Along the same lines, an online method for extracting non-static objects from the observed scene, and hence improving the performance of RGB-D SLAM in non-static environments was proposed in [139]. The approach consists of three main stages, starting with image differencing to detect any moving objects in the scene. A particle filter is then employed to track motion patches in consecutive RGB-D frames, which makes it more general than approaches that track particular object models. Finally, maximum-a-posteriori is used to identify the scene's foreground, after segmenting the moving objects by means of vector quantization. To op-

erate reliably, the approach requires the observed scene to consist mainly of static objects and to contain planes.

As the scene to be re-constructed by visual SLAM grows larger, matching features to points becomes more challenging because some places exhibit similar appearances. To circumvent this, the work presented in [74] employs a coarse place recognition module where frames containing common points are grouped together under location classes using an overlapping view clustering algorithm. Matching features is then done based on the Hamming distance between BRIEF descriptors of Harris corners.

Data association and SLAM are tightly coupled problems that were not considered jointly except in a few research work where they were solved as two optimization sub-problems. Data association for each observation-landmark pair is estimated then used to estimate the sensor and landmark poses. Using this approach, the accuracy of sensor and landmark pose estimation is critically degraded by incorrect data association. In addition, measurements that are discarded due to their ambiguity cannot be reconsidered when more refined measurements of the same landmarks are obtained.

These limitations motivated the changes in the SLAM algorithm proposed in [18] where data association, and estimations of sensor and landmarks poses were considered in a single optimization problem. Instead of associating each observation to a single landmark, Expectation Maximization was employed to account for the entire density of the data association while estimating the sensor and landmark poses, which was referred to as soft data association. Estimation is based on inertial measurements, ORB features, and semantic information obtained from an object detector. The depth of an observed landmark is the median of the ORB features detected within the bounding box of that landmark. In case multiple instances of the same object exist in the environment, Mahalanobi's distance is used to decide data association. An extension of this work was presented in [140], where semantic structure was inferred differently. Instead of relying on ORB features, a stacked hourglass convolutional network was used to detect semantic features of the object found within each bounding box. Structure constraints are used to relate each semantic feature to the corresponding landmark and Kabsch Algorithm is then used to estimate the orientation of the object. A very similar approach can be found in [141] with the distinction that it employs non-Gaussian sensor models as opposed to majority of the proposed approaches, where Gaussian model are always assumed.

The system proposed in [142] combines high-level semantically labeled features and low-level CNN features

to localize a mobile robot by means of a coarse to fine approach. Observations are matched to visual frames in the map by first comparing the objects appearing in the image. A finer search is then carried out based on CNN features of the image. The estimated poses of the camera as well as the features are finally refined using BA.

3.4.3 Low-, Middle-, and High-level Feature-based Approaches

In [143] and [127], SLAM systems are developed based on features from all three levels; points, planes, and objects.

The system proposed in [143] employs an RGB-D sensor to observe features in the environment. Real-time, efficient performance of this system is achievable because objects are represented by means of quadrics which do not require highly detailed representation. The SLAM problem is formulated as a factor graph where various types of factors are used, including observations of points, objects, and planes as well as point-plane, plane-plane, and object-plane relationships. A variation of ORB-SLAM2 is used to detect points in the environment which are then matched among frames in a coarse-to-fine pyramid. Faster R-CNN is used to detect objects in incoming frames and the corresponding ellipsoids representing the objects are then computed. Across frames, semantic labels are used to associate observations to objects if a single instance of the object appears in the environment. Otherwise, data association is achieved by means of nearest-neighbor matching. The point cloud representing the scene is segmented to extract planes using an organized point cloud segmentation technique. Planes are associated using thresholds on the distance between them and the difference between their normals. Factors are added between planes and points that belong to them, objects and the corresponding planes they lie on, and between multiple planes assuming a Manhattan World. A bag of words approach is adopted to detect loop closures.

Using points, planes, and objects, observed through a monocular camera, the work presented in [127] achieves improved localization, especially in absence of loop closure, compared to state-of-the-art SLAM systems. This is attributed to the long-range observability of objects and planes which facilitates more associations between old and new measurements. Objects are represented as cuboids, plane edges are detected then back-projected to obtain their parameters, and points are added to further constrain the camera poses. BA formulation is used with four types of constraints, camera-plane, camera-object, object-plane, and point-plane. The maps gen-

erated are dense and exhibit a high level of expressive representation.

4 Conclusion

Simultaneous Localization and Mapping is the most predominant research problem in the robotics community where tremendous amounts of effort are put into generating novel approaches that maximize its robustness and reliability. Upon acquisition of the first set of measurements from the environment to be reconstructed, the trajectory of the robot and the map are initialized. Subsequent measurements go through a pipeline of different processes that are implemented differently in each SLAM system but do achieve the same purpose. Such processes include data association, loop closure, re-localization, and trajectory and map estimation.

In this paper, we surveyed most of the state-of-the-art visual SLAM solutions that employ features to localize the robot and map its surroundings. We classified feature-based visual SLAM approaches into categories based on the types of features they rely on; low-level, middle-level, high-level, or hybrid features. The strengths and weaknesses of each category were thoroughly investigated and the challenges that each solution overcomes were highlighted, when applicable. Comparisons between approaches in the same category were provided in tables, comparing the methods that were adopted to implement each component of the SLAM pipeline.

Based on our intensive review, we believe that the following challenges remain unsolved.

1. *Generality*: Current SLAM solutions lack the ability to adapt to the environment in which the robot is operating. Because they depend on a certain type of features. Failure to detect such features in the environment leads to catastrophic degradation in the accuracy of the SLAM outcome. This could be due to the intermittent presence of features in the environment or the inability of the employed vision system to detect them. The former happens if the SLAM system depends on a very limited set of features, for instance the set of objects that a neural network can detect, while not utilizing other elements in the image like planes, geometric primitives, or new objects that the network was not trained to detect. The latter might occur in challenging environments or due to abrupt motions. To cope with such challenges, the vision system employed by SLAM should be flexible to accommodate various types of features based on the environment in which the robot is operating,
2. *Robustness*: In presence of noise from several sources in the SLAM pipeline, it is sometimes hard for the estimation algorithm to generate optimum estimates of the map and trajectory. Very limited research work has been done to guarantee the optimality of a SLAM estimate or at least verify whether or not the estimate is optimal [144], [145], [146], [147], [148], [149]. To that end, post-processing SLAM estimates by means of a neural network, for example, might result in significant improvements to the estimated trajectory and reconstructed map, and hence a more robust SLAM system.
3. *Scene Understanding and Expressive Representation*: Ever since the deep learning breakthrough in 2012, object detectors have been heavily exploited in SLAM. However, the current object detectors do not exploit any temporal or spatial relationships between the detections [150]. If such constraints are accounted for, an increase in the efficiency and reliability of the detections is expected.

The advances in software and hardware technology that we currently witness should be directed towards developing an environment-aware, error-free, general visual SLAM algorithm that is capable of circumventing all of these challenges.

Acknowledgements This publication is based upon work supported by the Khalifa University of Science and Technology under Award No. RC1-2018-KUCARS

Conflict of Interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

1. Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer Vision and Image Understanding*, 110(3):346 – 359, 2008. Similarity Matching in Computer Vision and Multimedia.
2. C. Posch, D. Matolin, and R. Wohlgenannt. A qvga 143 db dynamic range frame-free pwm image sensor with lossless pixel-level video compression and time-domain cds. *IEEE Journal of Solid-State Circuits*, 46(1):259–275, Jan 2011.
3. C. Brandli, R. Berner, M. Yang, S. Liu, and T. Delbruck. A 240 x 180 130 db 3 μ s latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, Oct 2014.

4. D. Weikersdorfer, D. B. Adrian, D. Cremers, and J. Conradt. Event-based 3d slam with a depth-augmented dynamic vision sensor. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 359–364, May 2014.
5. Hanme Kim, Ankur Handa, Ryad Benosman, Sio-Hoi Ieng, and Andrew Davison. Simultaneous mosaicing and tracking with an event camera. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014.
6. David Weikersdorfer, Raoul Hoffmann, and Jörg Conradt. Simultaneous localization and mapping for event-based vision systems. In Mei Chen, Bastian Leibe, and Bernd Neumann, editors, *Computer Vision Systems*, pages 133–142, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
7. Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. 2019.
8. M. J. Milford, G. F. Wyeth, and D. Prasser. Ratslam: a hippocampal model for simultaneous localization and mapping. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 1, pages 403–408 Vol.1, April 2004.
9. M. J. Milford, F. Schill, P. Corke, R. Mahony, and G. Wyeth. Aerial slam with a single camera using visual expectation. In *2011 IEEE International Conference on Robotics and Automation*, pages 2506–2512, May 2011.
10. Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsd-slam: Large-scale direct monocular slam. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 834–849, Cham, 2014. Springer International Publishing.
11. Georges Younes, Daniel Asmar, Elie Shammas, and John Zelek. Keyframe-based monocular slam: design, survey, and future directions. *Robotics and Autonomous Systems*, 98:67 – 88, 2017.
12. R. Mur-Artal and J. D. Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, Oct 2017.
13. G. Klein and D. Murray. Parallel tracking and mapping for small ar workspaces. In *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 225–234, Nov 2007.
14. T. Pire, T. Fischer, J. Civera, P. De Cristóforis, and J. J. Berles. Stereo parallel tracking and mapping for robot localization. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1373–1378, Sept 2015.
15. B. Mu, S. Y. Liu, L. Paull, J. Leonard, and J. P. How. Slam with objects using a nonparametric pose graph. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4602–4609, Oct 2016.
16. Siddharth Choudhary, Luca Carlone, Carlos Nieto, John Rogers, Zhen Liu, Henrik I. Christensen, and Frank Dellaert. Multi robot object-based slam. In Dana Kulić, Yoshihiko Nakamura, Oussama Khatib, and Gennaro Ventura, editors, *2016 International Symposium on Experimental Robotics*, pages 729–741, Cham, 2017. Springer International Publishing.
17. A. Gawel, C. D. Don, R. Siegwart, J. Nieto, and C. Cadena. X-view: Graph-based semantic multi-view localization. *IEEE Robotics and Automation Letters*, 3(3):1687–1694, July 2018.
18. S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas. Probabilistic data association for semantic slam. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1722–1729, May 2017.
19. Mehdi Hosseinzadeh, Yasir Latif, Trung Pham, Niko Sünderhauf, and Ian D. Reid. Towards semantic SLAM: points, planes and objects. *CoRR*, abs/1804.09111, 2018.
20. S. Yang, Y. Song, M. Kaess, and S. Scherer. Pop-up slam: Semantic monocular plane slam for low-texture environments. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1222–1229, Oct 2016.
21. C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, Dec 2016.
22. Muhamad Risqi U. Saputra, Andrew Markham, and Niki Trigoni. Visual slam and structure from motion in dynamic environments: A survey. *ACM Comput. Surv.*, 51(2):37:1–37:36, February 2018.
23. Nathan Piasco, Désiré Sidibé, Cédric Demonceaux, and Valérie Gouet-Brunet. A survey on visual-based localization: On the benefit of heterogeneous data. *Pattern Recognition*, 74:90 – 109, 2018.
24. Emilio Garcia-Fidalgo and Alberto Ortiz. Vision-based topological mapping and localization methods: A survey. *Robotics and Autonomous Systems*, 64:1 – 20, 2015.
25. Khalid Yousif, Alireza Bab-Hadiashar, and Reza Hoseinnezhad. An overview to visual odometry and visual slam: Applications to mobile robotics. *Intelligent Industrial Systems*, 1(4):289–311, Dec 2015.
26. Jorge Fuentes-Pacheco, José Ruiz-Ascencio, and Juan Manuel Rendón-Mancha. Visual simultaneous localization and mapping: a survey. *Artificial Intelligence Review*, 43(1):55–81, Jan 2015.
27. Takafumi Taketomi, Hideaki Uchiyama, and Sei Ikeda. Visual slam algorithms: a survey from 2010 to 2016. *IPSJ Transactions on Computer Vision and Applications*, 9(1):16, Jun 2017.
28. Muhammad Sualeh and Gon-Woo Kim. Simultaneous localization and mapping in the epoch of semantics: A survey. *International Journal of Control, Automation and Systems*, 17(3):729–742, Mar 2019.
29. N. Sünderhauf, T. T. Pham, Y. Latif, M. Milford, and I. Reid. Meaningful maps with object-oriented semantic mapping. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5079–5085, Sept 2017.
30. A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse. Monoslam: Real-time single camera slam. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):1052–1067, June 2007.
31. Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, pages 2724–2729 vol.3, Apr 1991.
32. P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, Feb 1992.
33. Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI’81,

- pages 674–679, San Francisco, CA, USA, 1981. Morgan Kaufmann Publishers Inc.
34. H.-Y. Shum and R. Szeliski. *Construction of Panoramic Image Mosaics with Global and Local Alignment*, pages 227–268. Springer New York, New York, NY, 2001.
 35. Henrik Stewénus, Christopher Engels, and David Nistér. Recent developments on direct relative orientation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 60(4):284 – 294, 2006.
 36. Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981.
 37. P.H.S. Torr and A. Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1):138 – 156, 2000.
 38. J. Civera, A. J. Davison, and J. M. M. Montiel. Inverse depth parametrization for monocular slam. *IEEE Transactions on Robotics*, 24(5):932–945, Oct 2008.
 39. Hyunhak Cho, Eun Kyeong Kim, and Sungshin Kim. Indoor slam application using geometric and icp matching methods based on line features. *Robotics and Autonomous Systems*, 100:206 – 224, 2018.
 40. A. Concha, G. Loianno, V. Kumar, and J. Civera. Visual-inertial direct slam. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1331–1338, May 2016.
 41. P. R. Beaudet. Rotationally invariant image operators. In *Proceedings of the 4th International Joint Conference on Pattern Recognition*, pages 579–583, Kyoto, Japan, November 1978.
 42. M. Calonder, V. Lepetit, M. Ozuysal, T. Trzcinski, C. Strecha, and P. Fua. Brief: Computing a local binary descriptor very fast. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7):1281–1298, July 2012.
 43. Chris Harris and Mike Stephens. A combined corner and edge detector. In *In Proc. of Fourth Alvey Vision Conference*, pages 147–151, 1988.
 44. Jianbo Shi and C. Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, Jun 1994.
 45. D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157 vol.2, 1999.
 46. Tony Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):79–116, Nov 1998.
 47. N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 886–893 vol. 1, June 2005.
 48. J. Matas, O Chum, M Urban, and T Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761 – 767, 2004. British Machine Vision Computing 2002.
 49. E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *2011 International Conference on Computer Vision*, pages 2564–2571, Nov 2011.
 50. David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, Nov 2004.
 51. A. Alahi, R. Ortiz, and P. Vandergheynst. Freak: Fast retina keypoint. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 510–517, June 2012.
 52. Elmar Mair, Gregory D. Hager, Darius Burschka, Michael Suppa, and Gerhard Hirzinger. Adaptive and generic corner detection based on the accelerated segment test. In Kostas Daniilidis, Petros Maragos, and Nikos Paragios, editors, *Computer Vision – ECCV 2010*, pages 183–196, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
 53. Stefan Leutenegger, Margarita Chli, and Roland Y. Siegwart. Brisk: Binary robust invariant scalable keypoints. In *Proceedings of the 2011 International Conference on Computer Vision, ICCV ’11*, pages 2548–2555, Washington, DC, USA, 2011. IEEE Computer Society.
 54. Shichao Yang, D. Maturana, and S. Scherer. Real-time 3d scene layout from a single image using convolutional neural networks. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2183–2189, May 2016.
 55. Xuming He, R. S. Zemel, and M. A. Carreira-Perpinan. Multiscale conditional random fields for image labeling. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 2, pages II–695–II–702 Vol.2, June 2004.
 56. Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, Sep 1995.
 57. Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015.
 58. Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018.
 59. Georges Younes, Daniel C. Asmar, and Elie A. Shamas. A survey on non-filter-based monocular visual SLAM systems. *CoRR*, abs/1607.00470, 2016.
 60. Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, USA, 2 edition, 2003.
 61. Marius Muja and David G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *In VISAPP International Conference on Computer Vision Theory and Applications*, pages 331–340, 2009.
 62. D. Galvez-López and J. D. Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, Oct 2012.
 63. Parv Parkhiya, Rishabh Khawad, J. Krishna Murthy, Brojeshwar Bhowmick, and K. Madhava Krishna. Constructing category-specific models for monocular object-slam. *CoRR*, abs/1802.09292, 2018.
 64. M. Mazuran, G. Diego Tipaldi, L. Spinello, W. Burgard, and C. Stachniss. A statistical measure for map consistency in slam. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3650–3655, May 2014.
 65. Yang Liu and H. Zhang. Indexing visual features: Real-time loop closure detection using a tree structure. In *2012 IEEE International Conference on Robotics and Automation*, pages 3613–3618, May 2012.
 66. Nishant Kejriwal, Swagat Kumar, and Tomohiro Shibata. High performance loop closure detection using bag of word pairs. *Robotics and Autonomous Systems*, 77:55 – 65, 2016.

67. A. Angeli, S. Doncieux, J. Meyer, and D. Filliat. Real-time visual loop-closure detection. In *2008 IEEE International Conference on Robotics and Automation*, pages 1842–1847, May 2008.
68. X. Gao and T. Zhang. In *2015 34th Chinese Control Conference (CCC)*, pages 5851–5856, July 2015.
69. Y. Xia, J. Li, L. Qi, H. Yu, and J. Dong. An evaluation of deep learning in loop closure detection for visual slam. In *2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pages 85–91, June 2017.
70. Kin Leong Ho and Paul Newman. Loop closure detection in slam by combining visual and spatial appearance. *Robotics and Autonomous Systems*, 54(9):740 – 749, 2006. Selected papers from the 2nd European Conference on Mobile Robots (ECMR '05).
71. S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19, Feb 2016.
72. Iksan Bukhori and Zool Hilmi Ismail. Detection of kidnapped robot problem in monte carlo localization based on the natural displacement of the robot. *International Journal of Advanced Robotic Systems*, 14(4):1729881417717469, 2017.
73. S. P. Engelson and D. V. McDermott. Error correction in mobile robot map learning. In *Proceedings 1992 IEEE International Conference on Robotics and Automation*, pages 2555–2560 vol.3, May 1992.
74. H. Lim, S. N. Sinha, M. F. Cohen, and M. Uyttendaele. Real-time image-based 6-dof localization in large-scale environments. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1043–1050, June 2012.
75. Dorian Gálvez-López, Marta Salas, Juan D. Tardós, and J.M.M. Montiel. Real-time monocular object slam. *Robotics and Autonomous Systems*, 75:435 – 449, 2016.
76. J. G. Rogers, A. J. B. Trevor, C. Nieto-Granda, and H. I. Christensen. Simultaneous localization and mapping with learned object recognition and semantic data association. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1264–1270, Sept 2011.
77. S. Maity, A. Saha, and B. Bhowmick. Edge slam: Edge points based monocular visual slam. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 2408–2417, Oct 2017.
78. L. Riazuelo, L. Montano, and J. M. M. Montiel. Semantic visual slam in populated environments. In *2017 European Conference on Mobile Robots (ECMR)*, pages 1–7, Sept 2017.
79. Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment — a modern synthesis. In Bill Triggs, Andrew Zisserman, and Richard Szeliski, editors, *Vision Algorithms: Theory and Practice*, pages 298–372, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
80. Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.
81. M. Kaess, A. Ranganathan, and F. Dellaert. isam: Incremental smoothing and mapping. *IEEE Transactions on Robotics*, 24(6):1365–1378, Dec 2008.
82. Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John J Leonard, and Frank Dellaert. isam2: Incremental smoothing and mapping using the bayes tree. *The International Journal of Robotics Research*, 31(2):216–235, 2012.
83. E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd. Real time localization and 3d reconstruction. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 1, pages 363–370, June 2006.
84. Roberto Sabatini, Subramanian Ramasamy, Alessandro Gardi, and Leopoldo Rodriguez Salazar. Low-cost sensors data fusion for small size unmanned aerial vehicles navigation and guidance. *International Journal of Unmanned Systems Engineering*, 1:16–47, 08 2013.
85. N. Sünderhauf and P. Protzel. Towards a robust backend for pose graph slam. In *2012 IEEE International Conference on Robotics and Automation*, pages 1254–1261, May 2012.
86. Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment — a modern synthesis. In Bill Triggs, Andrew Zisserman, and Richard Szeliski, editors, *Vision Algorithms: Theory and Practice*, pages 298–372, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
87. A. Eudes and M. Lhuillier. Error propagations for local bundle adjustment. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2411–2418, June 2009.
88. Liang Zhao, Shoudong Huang, Yanbiao Sun, Lei Yan, and Gamini Dissanayake. Parallaxba: bundle adjustment using parallax angle feature parametrization. *The International Journal of Robotics Research*, 34(4-5):493–516, 2015.
89. M. Kaess, A. Ranganathan, and F. Dellaert. isam: Incremental smoothing and mapping. *IEEE Transactions on Robotics*, 24(6):1365–1378, Dec 2008.
90. Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John J Leonard, and Frank Dellaert. isam2: Incremental smoothing and mapping using the bayes tree. *The International Journal of Robotics Research*, 31(2):216–235, 2012.
91. Sarquis Urzua, Rodrigo Munguía, and Antoni Grau. Vision-based slam system for mavs in gps-denied environments. *International Journal of Micro Air Vehicles*, 9(4):283–296, 2017.
92. A. Kasyanov, F. Engelmann, J. Stückler, and B. Leibe. Keyframe-based visual-inertial online slam with relocalization. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6662–6669, Sept 2017.
93. H. Lim, J. Lim, and H. J. Kim. Real-time 6-dof monocular visual slam in a large-scale environment. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1532–1539, 2014.
94. A. Pumarola, A. Vakhtov, A. Agudo, A. Sanfeliu, and F. Moreno-Noguer. Pl-slam: Real-time monocular visual slam with points and lines. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4503–4508, May 2017.
95. Ruben Gomez-Ojeda, Francisco Angel Moreno, Davide Scaramuzza, and Javier González Jiménez. PL-SLAM: a stereo SLAM system through the combination of points and line segments. *CoRR*, abs/1705.09479, 2017.
96. Guolai Jiang, Lei Yin, Shaokun Jin, Chaoran Tian, Xinbo Ma, and Yongsheng Ou. A simultaneous localization and mapping (slam) framework for 2.5d map building based on low-cost lidar and vision fusion. *Applied Sciences*, 9(10), 2019.

97. B. Clipp, Jongwoo Lim, J. M. Frahm, and M. Pollefeys. Parallel, real-time visual slam. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3961–3968, Oct 2010.
98. M. A. Nitsche, G. I. Castro, T. Pire, T. Fischer, and P. De Cristóforis. Constrained-covisibility marginalization for efficient on-board stereo slam. In *2017 European Conference on Mobile Robots (ECMR)*, pages 1–6, Sept 2017.
99. R. Williams, B. Konev, and F. Coenen. Scalable distributed collaborative tracking and mapping with micro aerial vehicles. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3092–3097, Sept 2015.
100. M. Quan, S. Piao, M. Tan, and S. Huang. Accurate monocular visual-inertial slam using a map-assisted ekf approach. *IEEE Access*, 7:34289–34300, 2019.
101. T. Lowe, S. Kim, and M. Cox. Complementary perception for handheld slam. *IEEE Robotics and Automation Letters*, 3(2):1104–1111, April 2018.
102. A. S. Zhang, B. S. Liu, C. J. Zhang, D. Z. Wang, and E. X. Wang. Fast initialization for feature-based monocular slam. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 2119–2123, Sept 2017.
103. J. Civera, D. Gálvez-López, L. Riazuelo, J. D. Tardós, and J. M. M. Montiel. Towards semantic slam using a monocular camera. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1277–1284, Sept 2011.
104. M. Bosse, R. Zlot, and P. Flick. Zebedee: Design of a spring-mounted 3-d range sensor with application to mobile mapping. *IEEE Transactions on Robotics*, 28(5):1104–1119, Oct 2012.
105. Jigang Liu, Dongquan Liu, Jun Cheng, and Yuanyan Tang. Conditional simultaneous localization and mapping: A robust visual slam system. *Neurocomputing*, 145:269 – 284, 2014.
106. K. Qiu, T. Liu, and S. Shen. Model-based global localization for aerial robots using edge alignment. *IEEE Robotics and Automation Letters*, 2(3):1256–1263, July 2017.
107. B. A. C. Caldato, R. A. Filho, and J. E. C. Castanho. Orb-odom: Stereo and odometer sensor fusion for simultaneous localization and mapping. In *2017 Latin American Robotics Symposium (LARS) and 2017 Brazilian Symposium on Robotics (SBR)*, pages 1–5, Nov 2017.
108. Georg Klein and David Murray. Improving the agility of keyframe-based slam. In David Forsyth, Philip Torr, and Andrew Zisserman, editors, *Computer Vision – ECCV 2008*, pages 802–815, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
109. U. Qayyum, Q. Ahsan, and Z. Mahmood. Imu aided rgb-d slam. In *2017 14th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, pages 337–341, Jan 2017.
110. A. Annaiyan, M. A. Olivares-Mendez, and H. Voos. Real-time graph-based slam in unknown environments using a small uav. In *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1118–1123, June 2017.
111. X. Wang, S. Vozar, and E. Olson. Flag: Feature-based localization between air and ground. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3178–3184, May 2017.
112. J. Unicom, L. Dantanarayana, J. Arukgoda, R. Ranasinghe, G. Dissanayake, and T. Furukawa. Distance function based 6dof localization for unmanned aerial vehicles in gps denied environments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5292–5297, Sept 2017.
113. Runzhi Wang, Wenhui Wan, Yongkang Wang, and Kaichang Di. A new rgb-d slam method with moving object detection for dynamic indoor scenes. *Remote Sensing*, 11:1143, 2019.
114. Jiexiong Tang, Ludvig Ericson, John Folkesson, and Patric Jensfelt. Gcnv2: Efficient correspondence prediction for real-time SLAM. *CoRR*, abs/1902.11046, 2019.
115. J. Cheng, C. Wang, and M. Q. Meng. Robust visual localization in dynamic environments based on sparse motion removal. *IEEE Transactions on Automation Science and Engineering*, pages 1–12, 2019.
116. Simon Bultmann, Kailai Li, and Uwe Hanebeck. Stereo visual slam based on unscented dual quaternion filtering. In *Proceedings of the 22nd International Conference on Information Fusion (Fusion 2019)*, 07 2019.
117. L. A. V. Souto and T. P. Nascimento. Object subtraction planar rgb-d slam. In *2016 XIII Latin American Robotics Symposium and IV Brazilian Robotics Symposium (LARS/SBR)*, pages 19–24, Oct 2016.
118. A. P. Gee, D. Chekhlov, A. Calway, and W. Mayol-Cuevas. Discovering higher level structure in visual slam. *IEEE Transactions on Robotics*, 24(5):980–990, Oct 2008.
119. A. Flint, C. Mei, I. Reid, and D. Murray. Growing semantically meaningful models for visual slam. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 467–474, June 2010.
120. Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. *CoRR*, abs/1612.08242, 2016.
121. J. Li, D. Meger, and G. Dudek. Context-coherent scenes of objects for camera pose estimation. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 655–660, Sept 2017.
122. S. Choudhary, A. J. B. Trevor, H. I. Christensen, and F. Dellaert. Slam with object discovery, modeling and mapping. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1018–1025, Sept 2014.
123. A. J. B. Trevor, J. G. Rogers, and H. I. Christensen. Omnimapper: A modular multimodal mapping framework. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1983–1990, May 2014.
124. Aleksandr Segal, Dirk Hähnel, and Sebastian Thrun. Generalized-icp. In Jeff Trinkle, Yoky Matsuoka, and José A. Castellanos, editors, *Robotics: Science and Systems*. The MIT Press, 2009.
125. Lachlan Nicholson, Michael Milford, and Niko Sünderhauf. Quadricslam: Constrained dual quadrics from object detections as landmarks in semantic SLAM. *CoRR*, abs/1804.04011, 2018.
126. Masoud S. Bahraini, Mohammad Bozorg, and Ahmad B. Rad. Slam in dynamic environments via ml-ransac. *Mechatronics*, 49:105 – 118, 2018.
127. S. Yang and S. Scherer. Monocular object and plane slam in structured environments. *IEEE Robotics and Automation Letters*, 4(4):3145–3152, Oct 2019.
128. Parv Parkhiya, Rishabh Khawad, J. Krishna Murthy, Brojeshwar Bhowmick, and K. Madhava Krishna. Constructing category-specific models for monocular object-slam. *CoRR*, abs/1802.09292, 2018.

129. M. Henein, M. Abello, V. Ila, and R. Mahony. Exploring the effect of meta-structural information on the global consistency of slam. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1616–1623, Sept 2017.
130. P. H. Le and J. Košečka. Dense piecewise planar rgb-d slam for indoor environments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4944–4949, Sept 2017.
131. Xiaoyu Zhang, Wei Wang, Xianyu Qi, Ziwei Liao, and Ran Wei. Point-plane slam using supposed planes for indoor environments. In *Sensors*, 2019.
132. A. Concha and J. Civera. Using superpixels in monocular slam. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 365–372, May 2014.
133. N. Fioraio and L. Di Stefano. Joint detection, tracking and mapping by semantic bundle adjustment. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1538–1545, June 2013.
134. B. Williams, G. Klein, and I. Reid. Real-time slam relocalisation. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8, Oct 2007.
135. Rafael Muñoz-Salinas and Rafael Medina Carnicer. Ucoslam: Simultaneous localization and mapping by fusion of keypoints and squared planar markers. *CoRR*, abs/1902.03729, 2019.
136. Rafael Muñoz-Salinas, Manuel J. Marín-Jimenez, and R. Medina-Carnicer. Spm-slam: Simultaneous localization and mapping with squared planar markers. *Pattern Recognition*, 86:156 – 171, 2019.
137. Zemin Wang, Qian Zhang, Jiansheng Li, Shuming Zhang, and Jingbin Liu. A computationally efficient semantic slam solution for dynamic scenes. *Remote Sensing*, 11(11):1363, Jun 2019.
138. O. H. Jafari, D. Mitzel, and B. Leibe. Real-time rgb-d based people detection and tracking for mobile robots and head-worn cameras. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5636–5643, May 2014.
139. Yuxiang Sun, Ming Liu, and Max Q.-H. Meng. Improving rgb-d slam in dynamic environments: A motion removal approach. *Robotics and Autonomous Systems*, 89:110 – 122, 2017.
140. Nikolay Atanasov, Sean L. Bowman, Kostas Daniilidis, and George J. Pappas. A unifying view of geometry, semantics, and data association in slam. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 5204–5208. International Joint Conferences on Artificial Intelligence Organization, 7 2018.
141. K. Doherty, D. Fourie, and J. Leonard. Multimodal semantic slam with probabilistic data association. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2419–2425, May 2019.
142. W. Zhang, G. Liu, and G. Tian. A coarse to fine indoor visual localization method using environmental semantic information. *IEEE Access*, 7:21963–21970, 2019.
143. Mehdi Hosseinzadeh, Yasir Latif, Trung Pham, Niko Sünderhauf, and Ian D. Reid. Towards semantic slam: Points, planes and objects. *CoRR*, abs/1804.09111, 2018.
144. S. Huang, Yingwu Lai, U. Frese, and G. Dissanayake. How far is slam from a linear least squares problem? In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3011–3016, Oct 2010.
145. S. Huang, H. Wang, U. Frese, and G. Dissanayake. On the number of local minima to the point feature based slam problem. In *2012 IEEE International Conference on Robotics and Automation*, pages 2074–2079, May 2012.
146. Shoudong Huang and Gamini Dissanayake. A critique of current developments in simultaneous localization and mapping. *International Journal of Advanced Robotic Systems*, 13(5):1729881416669482, 2016.
147. L. Carlone and F. Dellaert. Duality-based verification techniques for 2d slam. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4589–4596, May 2015.
148. L. Carlone. A convergence analysis for pose graph optimization via gauss-newton methods. In *2013 IEEE International Conference on Robotics and Automation*, pages 965–972, May 2013.
149. L. Carlone, D. M. Rosen, G. Calafiore, J. J. Leonard, and F. Dellaert. Lagrangian duality in 3d slam: Verification techniques and optimal solutions. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 125–132, Sep. 2015.
150. Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, and Peter Corke. The limits and potentials of deep learning for robotics. *The International Journal of Robotics Research*, 37(4-5):405–420, 2018.