

Toward Holistic Scene Understanding: A Transfer of Human Scene Perception to Mobile Robots

Florenz Graf¹, Jochen Lindermayr¹, Çağatay Odabaşı¹ and Marco F. Huber²

Abstract—The long-term vision for robotics is to have fully autonomous mobile robots that perceive the environment as humans do or even better. This article transfers the core ideas from human scene perception to robot scene perception to contribute toward a holistic scene understanding of robots. The first contribution is to extensively survey and compare state-of-the-art robot scene perception approaches with neuroscience theories and studies of human perception. A step-by-step transfer of the perceptual process reveals similarities and differences between robots and humans. The second contribution represents an analysis of the status quo of holistic robot perception approaches to extract to what extent the perceptual capabilities of humans have been reached. Building on this, the gaps and potentials of robot perception are illustrated to address future research directions.

I. INTRODUCTION

The last years indicated fast technological improvements in the artificial intelligence of robots. The *International Federation of Robotics* records a market growth of 12% in 2020 of professional robots used for various applications, such as transport, inspection, cleaning, medical or hospitality [1]. The report forecasts exponential growth for the upcoming years. Robots will become part of human society. Acting side by side and collaborating within the same environments motivate being consistent with human perception.

In the past, robots mainly fulfilled highly customized tasks in industrial applications independently and separated from humans [2]. People adapted the environments to the application requirements. However, adaptations for the robots' needs are especially undesirable in non-industrial environments. Current products on the market mirror this issue through specialized applications. These products either fulfill a single task autonomously, such as vacuum cleaners and lawnmower robots executing their routine independently from humans, or interact with humans physically with a low autonomy level [1]. Therefore, special attention must be directed to scene perception as enabling technology to break this trade-off between robot proximity to humans and autonomy. Already in 1993, the psychologist Ulric Neisser mentioned that “*without perception there is no knowledge*” [3]. The goal of holistic scene perception is to understand the scene by “*considering the geometric and semantic context of its contents and the intrinsic relationships between them*” [4]. Scene perception is making sense of real-world scenes as a whole by enabling the interaction with the scene [5]. On the one hand, it offers new opportunities for multi-use applications regarding cost per function and physical assistance systems. On the other hand, it increases autonomous capabilities by overcoming challenges, such as occlusion of

objects, instance-specific handling of dynamics, or spatial-temporal reasoning of complex situations. Holistic scene perception will enable safe and complex behaviors for any collaboration.

Increasing attention has been given to specific areas of perception. The approaches achieve excellent results in semantic extraction, such as object detection or image classification [6]. However, in the real world, it is not sufficient that robots solely understand single pieces of the environment. Much more, various scene information needs to be understood concurrently in the context to reason within dynamics, uncertainties, and incompleteness for high-level control [7, Chapter 23]. Little is known about holistic scene perception, aiming to understand the scene in this integrated manner needed for a large spectrum of real-world applications.

Motivated by the performance of human perception, we distanced ourselves in this research from robotic approaches and studied theories and experiments on human perception. Neuroscience research has been extensively studied since Potter and Biederman in the 1970s [8], [9]. It reached international bandwidth since the 2000s [10]. The latest research provides a comprehensive overview of human scene perception [5], [10]. This article transfers the core theories of psychological studies on human scene perception to robots and thereby reveals similarities and differences. The comparison of human perception with artificial perception is not new [11]–[13]. However, we will focus on the scene perception as a whole using a top-down view of the robots' status quo.

II. METHODOLOGY

Figure 1 describes the methodology of this study. First, we analyzed current mobile robot applications where we identified a trend toward robots in everyday life, covering multiple use cases. For future applications, scene perception will play a key role as it provides comprehensive scene knowledge to the cognitive intelligence of robots. Therefore, we analyzed state-of-the-art robot scene perception approaches and theories of human scene perception to elaborate on a transfer. Research in both domains offers a comprehensive overview providing fundamental knowledge as the starting point. Section III summarizes the transfer with a step-by-step analysis of the similarities and differences between human and robot scene perception. Based on this, we addressed the status quo of holistic scene perception approaches of robots within a qualitative evaluation (Section IV) to extract how far robotic scene perception reached the performance of human perception to reveal gaps (Section V) and potentials

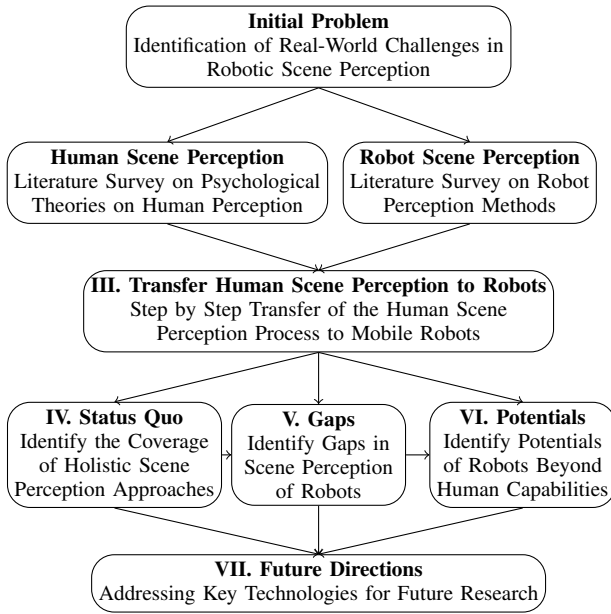


Fig. 1. Methodological approach with section references

(Section VI). Finally, we propose future directions for robotic scene perception (Section VII).

III. TRANSFERRING HUMAN SCENE PERCEPTION TO MOBILE ROBOTS

Humans perceive the environment by recognizing scene information through different senses, such as sight, hearing, smell, touch, or taste. Herewith, researchers discovered that people primarily rely on their visual perception for perceiving the environment due to the richness of information [14]. Although the process of human perception is unknown in detail, most theories define perception as the process of “*recognizing (being aware of), organizing (gathering and storing), and interpreting (binding to knowledge) sensory information*” [15, Chapter 3]. We transfer these three process steps to robot scene perception to reveal similarities and differences to the latest robotic research (Figure 2).

A. Recognition of Information

The recognition step processes sensory information to make it understandable for the subsequent perception tasks. Using the sensory input, humans convert the observations from the scene into understandable information. Whereas vision benefits from a high amount of information, other modalities, such as haptics, are robust to noise, and independent of light exposure to get properties being insufficiently recognized by other senses [16]. For example, the eye serves as a transducer to convert light into the optic nerve by five layers of cells (particularly photoreceptors) within the retina [17]. The nerve cells transmit information by electric signals to the brain for processing [18]. The processing itself divides *pre-attentive* and *post-attentive* recognition. In pre-attentive processing, activities related to low-level vision are usually associated with the extraction of certain physical properties of the visible environment, such as depth, shape,

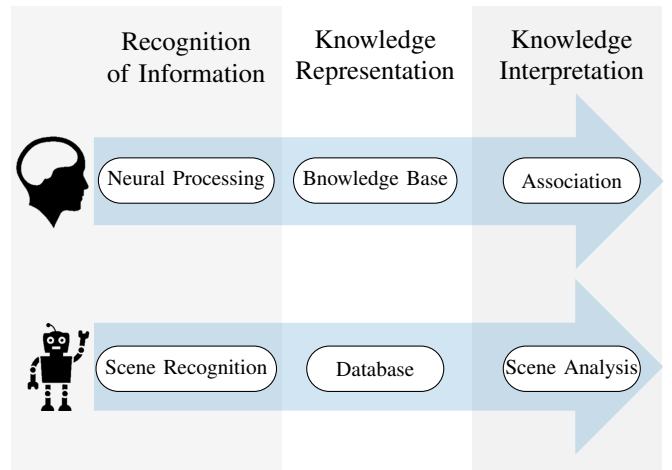


Fig. 2. The process of human perception transferred to mobile robots

color, object boundaries, or surface material properties [10], [19], [20]. Humans are not capable of processing all sensory information at once. Therefore, the post-attentive processing extracts scene information based on attention areas to overcome this issue [15]. It compresses high-level features, such as object classification or re-detection [21], [22]. Herewith, humans receive information from hierarchical levels of abstraction [23], [24] corresponding to a region of interest, triggered by the attention that provides the semantic of sensory input.

Similar to humans, the recognition of robots can be divided into pre- and post-attentive recognition. Equivalent to the human senses, robots use various sensors to recognize information from the environment. The most common sensor modalities are visual sensors, such as cameras, light detection and ranging (LiDAR), ultrasonic, and radar. They are capable of providing color and/or range information of the surrounding. The recognition of visual features is the most popular modality for robots due to information richness. However, few approaches investigate other modalities, such as acoustic perception to recognize objects [25], map the environment [26], [27] or to avoid obstacles [28], and haptic perception to actively touch objects [16]. The linking of action to perceive the information of interest is known as interactive perception, which is a common human exploration strategy. For instance, if humans cannot recognize an object by vision and touch, they will take different interactions to obtain information from other sensory channels [25]. The usage of smelling and tasting for robots is not popular. One rare example presents a navigation approach using smell for odor source localization [29]. A different niche develops a tasting sense using infrared (IR) spectroscopic technologies or chemical sensors to discriminate food [30], [31]. Driven by Computer Vision (CV), two-dimensional image processing is the fundamental pillar for robot recognition tasks. However, three-dimensional methods for scene and object reconstruction gained importance for robots. Various surveys provide an overview of the state-of-the-art robot recognition

tasks [32], [33]. Established robotics approaches focus on specific recognition subareas. The pre-attentive recognition of primitives comprises methods on normal estimation [34], [35], segmentation [36], [37], edge [38], and simple shape (planes, cylinder) [39] detection. The methods usually make use of the entire sensory input. However, robots pre-filter or scale the data to achieve appropriate results on their capabilities. For instance, image-based feature recognition can process high-resolution images by making use of image pyramids [40]. Herewith, the extracted features of a lower image resolution deliver discrete regions of interest by keeping the accuracy of the original scale.

Post-attentive processing comprises higher-level recognition capabilities based on pre-attentive processing. Examples are extracting the semantic by (re)-detection of objects [41], [42], humans [43], [44], places [45], three-dimensional reconstruction of the scene by Simultaneous Localization and Mapping (SLAM) [32], [46], [47], or enriching metric information by semantics, known as semantic segmentation [33], [48], [49].

Equal to humans, robots use pre-processed attention areas within post-attentive processing. For instance, feature-based SLAM methods use sparse feature points to post-attentively register a set of sensor data and to detect loop closures to compensate for the mapping drift [46], [47]. Besides, this example shows how humans and robots use different sensory sources to achieve highly accurate localization within the scene. Humans use walked steps combined with eye information for perceiving ego-motion which improves localization. For illustration, imagine walking a distance with closed eyes. The motion drift sooner or later leads to a loss of localization. Therefore, a series of interconnected spatial-visual features provide the absolute localization within the known scene [50]. Equivalent to humans, the latest robotic SLAM approaches provide similar techniques. Wheel odometry, visual odometry, or an Inertial Measurement Unit (IMU) provides the ego-motion. Laser scanners or camera-based methods simultaneously estimate the pose of landmarks to compensate for the ego-motion drift. Also, some approaches use visual landmarks for place re-detection to close the trajectory loop while mapping [46]. The usage of high-level features enables robots to add semantics to metric information. Knowing the semantics of an area in space helps robots to interpret the scene. It allows, e.g., the exclusion of dynamic objects like people while mapping.

A special challenge of robot perception is the recognition of scene information from different locations and times. For instance, a robot recognizes an object in the scene and simultaneously tracks the object in its field of view (FOV) as long as it is visible. While former research proposed using Kalman Filter [51] or Particle Filter [52], the latest research utilizes deep learning based tracking, such as with a Convolutional Neural Network (CNN) [53], [54] or a Vision Transformer (ViT) [55]. ViT, coming from Natural Language Processing (NLP), splits images into fixed-size patches that gained popularity due to their superior performance on continuous data streams, needed for mobile robots [56]. Since

the object and the robot could move, occlusions, truncation, or invisibility due to sensor noise (as mentioned above) must be handled. When the same object appears again, a Re-Identification (Re-ID) to reallocate the ID is beneficial to better understand the scene. Modern approaches of Re-ID propose similar to tracking the usage of a CNN [57], [58], ViT [59], [60] or end-to-end approaches [61].

B. Knowledge Representation

The second step of the perception process represents environmental information. A knowledge base manages recognized information from different sources and levels of abstraction, times, and places in a centralized and ordered structure. This structure includes understandable information about the scene. The function of the knowledge representation within the human brain has been a controversial topic since the so-called *Gestalt theory*. Modern attempts such as Wagemans et al. [24] or Hommel et al. [62] reveal spatial layouts, organized hierarchically, that represent the human perception memory. A relationship-focused multi-level hierarchical structure of parts represents environmental information.

The transfer of the main functionalities of human knowledge representation to robots requires a complex memory structure focused on flexibility. The knowledge representation must be capable of merging observations and interpretations from different sources and times. For instance, recognized information such as shape, texture, posture, state, probabilities, and trajectory (compare Section III-A) must be managed in real-time within the knowledge base. This issue sets high requirements to the underlying knowledge base as every piece of environmental information needs a known structured representation. Furthermore, the knowledge base includes initial and post-processed knowledge. Robots usually store and represent the scene knowledge in a database [63]–[65], allowing to deploy a human-understandable ontology that conceptualizes multiple entities within a domain and their relationships [66].

The scene representation comprises various perception-related information, such as extracted spatial features for navigation, manipulation, and semantically enhanced maps [67]. The requirements for robotic knowledge representation are high: ideally, it must be real-time capable, generic, scalable, flexible in structure, able to update and extend during the robot's lifetime, as well as easy to connect for access and data sharing. Generally, there are two categories of databases: graph-based and document-based. Both seem suitable for this task as they provide comprehensive features to cover these requirements [68], [69]. Graph-based databases, also known as Relational Database Management Systems (RDBMS), represent knowledge through relations. Herewith, it is necessary to *explicitly* define relations through a common format to link semantics through an ontology [70]. In contrast to relational databases, document-based databases represent data in JSON-like documents without the need for relations or predefined structures. Nevertheless, these databases provide features for querying or indexing the data to model depen-

dencies and relations *implicitly* dynamically. For instance, Kunze et al. propose spatial and temporal indexing for query relations within their document-based database [63]. Besides, linking the knowledge base for decision-making [71] and providing the represented knowledge to robot actions, such as manipulation or moving, enables reactive behaviors.

The comparison of scene knowledge representation indicates that robots have advantages compared to humans. First, the artificial scene knowledge representation has no memory loss due to an almost unlimited storage capacity. Second, robots can easily share and make use of foreign perception, while humans have to transfer knowledge into an appropriate modality, such as verbal communication. Additionally, humans are limited in the range of information exchange without technical assistance. In contrast, robots can share data in its original format over networks. The sharing of perceptual information enables robots to directly exchange data with infrastructure or other robots. Another difference between human and robot knowledge representation is the capability of robots to start with a pre-installed environment perception model. Robots can use the prior information of a Building Information Model (BIM) [72] or a partial or fully pre-mapped environment [64]. Using prior scene knowledge reduces the setup time, especially when using multiple agents.

C. Knowledge Interpretation

Based on available scene knowledge, this perception step interprets existing knowledge to make sense using cognitive capabilities. It has been proven that humans interpret the scene, however, it is unclear how this is executed within the brain. The research by Isik et al. found that the human brain starts recognizing view-invariant observations, such as human actions, fast, in around 200 ms [73]. This suggests that the brain uses the form, as well as the motions, to represent states. Furthermore, previous work [74] proposes that the human brain benefits from causal relations, such as temporal continuity and spatial relations, between objects, humans, and their actions. The prerequisite is that a known structure represents the knowledge (see Section III-B). A high-level scene analysis encodes spatial and temporal relations between instances [75]. For most examples, the interpretation of scene knowledge is a trivial task for humans due to lifelong learning. The interpretation of perceived information has been trained with pre-knowledge, dependent on the culture [76, Chapter 14], context, and situation itself [77]. Therefore, each person has a unique perception system based on and enriched by their environment. Thus, human interpretation is neither predictable nor always the same. The famous duck-rabbit illusion [78] indicates that even the season could differ the interpretation of the scene. Therefore, perception is influenced by environmental factors [77].

For robots, the interpretation of the scene is still an unsolved problem [79]. The challenge is to generate and make use of high-level semantic knowledge to reason about the present scene. There is no commonsense model that could be applied to every environment without adaptation

of the interpretations. The association of scene knowledge across multiple dimensions, such as time and space, with or without relational dependencies between single pieces of information, leads to this new knowledge that improves scene understanding. For instance, a spatio-temporal scene analysis could reveal daily habits, such as when, how, and how often we go to the kitchen to fetch a coffee. This example shows that the kind of required high-level information is environment-specific as well as use-case-specific. The goal is to understand what is involved and what to do when with which objects. The use of high-level semantics within the scene is fundamental for complex robot behavior tasks. Herewith, we identify two types of interpretations. On the one hand, there is research aimed at reconstructing and interpreting the *structured part* of the scene, such as room segmentation, junction detection [32], or occlusion reasoning for simple shapes [80]. On the other hand, there is research focusing on the *unstructured part* that goes deeper into handling dynamics. The approach presented in [81] is one of the rare examples in which already collected scene information is used to gather new information. Observations of objects are anchored in the scene model, which provides basic tracking functionality. In combination with knowledge about the whole scene, including other objects and their spatial and semantic relations, this is used for reasoning about the state of occluded objects, which improves tracking and hence the whole scene state. The tracking of instances over multiple observations enables further interpretations, such as detecting its action [82]. The actions of the people together with environmental semantics are valuable input for a robot since they usually share an environment. So, they need to understand the actions and fulfill their tasks proactively. For instance, if the person is cooking in a room, it is very likely that this room is a kitchen in which the robot can act accordingly [83]. Previous work [84] focuses on learning human actions by observing humans in their daily life. They merge joint motions and locations concerning the landmark points on the map. Kostavelis et al. [85] propose using object recognition and skeleton-based action recognition to make their social robot understand human actions. They deploy the robot in real home environments to test their system. The previous work [86] employs Long Short Term Memory (LSTM) [87] network on the robot to greet the user. The background may be misleading for the algorithms which are using appearance features. That is why generating action proposals may improve performance, and this is essential for mobile robots because the background in the robot view is dynamic [88]. In contrast to the spatio-temporal interpretations of a single instance, there are a few approaches that interpret relations within multiple instances. For example, Philipp et al. [89] propose to use Bayesian Networks to estimate on which object the user's attention focuses. The attention-sensitive functionalities as high-level scene interpretations are essential for improving human-robot interaction. More complex interpretations go into the affordance estimation [90]. Herewith, robots know what can be done with objects based on past observations,

probabilities, and personal data (emotions, preferences, relations). For instance, knowing that a pod can be used for cooking offers novel capabilities for robots. Affordances link perception into the cognitive capabilities that are fundamental for interpreting scenes.

D. Implications

Recent research on robot perception displays similarities to fundamental theories on human perception. Figure 3 visualizes the transferred scene perception process. Humans and robots recognize information from the scene by making sensor data understandable. The human brain and the robot storage, respectively, represent the knowledge in multiple layers in a known structure. Compared to humans, robots are capable of using external perceived data within its original format whereas humans have to exchange perceived data through verbal, visual, or written communication. The interpretation of knowledge over multiple dimensions, such as time, space, and relations, enables a high-level scene understanding that improves the cognitive intelligence of robots. Herewith humans highly benefit from life-long learning, whereby robots can benefit from shared and initial data.

In recent years, robotic recognition and interpretation use deep learning, such as CNNs and Generative Adversarial Networks (GANs) [91], based on ANNs, to solve a very specific task [92] as described above. Inspired by biological neural networks [93], ANNs consist of multiple layers of connected artificial neurons [94]. The weight assigned to these neurons relies on an initial ANN training using a high amount of labeled data. This in turn enables the ANN to compute labels of unknown data via inference. In particular, deep learning based approaches, such as vision transformers, gained popularity in robotics due to their superior performance on continuous data streams. Vision transformers, coming from Natural Language Processing (NLP), split images into fixed-size patches [56]. Machine learning, such as deep learning, achieves outstanding results for many recognition tasks [95], [96]. However, the division of pre-attentive and post-attentive feature processing can neither be strictly adhered to nor easily extended. The inseparability of ANNs challenges their flexibility and re-usability due to the “black box” characteristic. Additionally, they require enhanced tensor computation power that needs special attention when pushing mobile robots to the real world [92].

IV. STATUS QUO OF HOLISTIC SCENE PERCEPTION

The transfer reveals similarities and differences in the perception process between humans and robots. But how much does the status quo of robotic scene perception cover holistic capabilities? To answer this question, we first deal with the challenges of fundamental technologies to extract possible boundaries for human-like perception. Afterward, we analyze the scope of most holistic robotic approaches for everyday environments to answer how much they already cover the holistic scope.

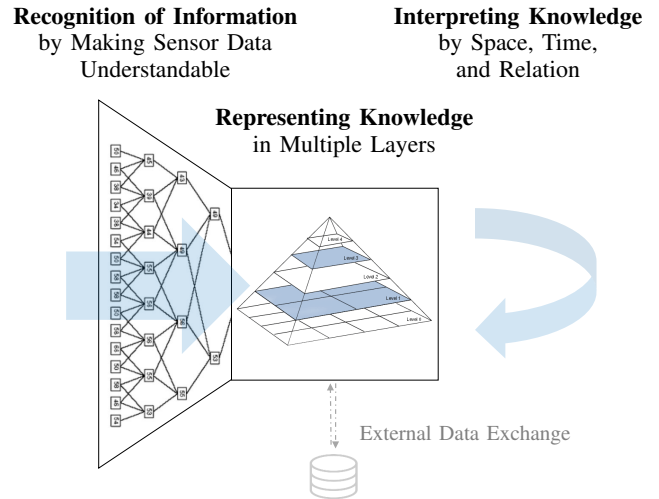


Fig. 3. Summary of the three steps of the transferred scene perception

A. From Narrow Toward Holistic Scene Perception

Research on robot scene perception focuses on sensor-close processing by affordable sensors. The availability of cheap RGB-cameras and two-dimensional laser scanners set the entry barrier in robotic developments, e.g., for student or hobby projects, relatively low. Investigations mainly focus on a deep learning based detection of common objects in images or on mapping the environment using SLAM. Two-dimensional methods in object detection benefit from matured research in CV. It reached high accuracy due to challenges, such as the Pascal Visual Object Classes (VOC) Challenge [97] and the Large Scale Visual Recognition Challenge [98] by public datasets of labeled data just as the Common Objects in Context (COCO) dataset [99]. For the application on robots, these two-dimensional image-based approaches must be transferred to time-continuous three-dimensional processing merged with SLAM techniques providing the ego-motion of the robot to reconstruct the environment.

Only comparatively new research focuses on three-dimensional multi-frame scene recognition solving the high computational performance with mobile graphic accelerators [49], [100], [101]. However, the representation as well as high-level interpretation of scene knowledge, are not focused. Indeed, as described by Neisser [3], solely the knowledge that was recognized can be represented or interpreted. The focus on sensor-close recognition in combination with visual and mobile robotic challenges, such as changing visual appearance, and limited computation resources, explains this observation. These aspects also explain why research on the representation and interpretation of scene knowledge is comparatively rare. As a result, robotic scene perception in real-world applications is focused on a concrete use-case utilizing highly optimized bottom-up perception pipelines

with narrow functionality. For instance, industrial environments have adapted to the perception capabilities of robots. Herewith, robots are enabled to fulfill narrow perception tasks with high accuracy. Especially deep learning-based recognition systems are often used as monolithic black box systems, being hard to combine efficiently without probing deep into a technical level. Research papers mostly bury these challenges by describing specific techniques [21]. However, robots in everyday environments need to fulfill multiple recognition tasks simultaneously as a requirement for complex and extensive behaviors. Open-source frameworks such as the Robot Operating System (ROS) [102] provide, thanks to its large community, many tutorials, and open-source basic functionalities that are suitable for creating a powerful overall performance based on single software pieces. Standardized communications between various software components enable robots to use, fuse, and analyze multiple data streams. However, on the one hand, paralyzing multiple narrow perception pipelines is challenging regarding performance. On the other hand, splitting pipelines into multiple steps, to reuse, e.g., pre-attentive information, is not easy when using monolithic black box models. Nevertheless, few researchers worked on integrated top-down approaches that aim to perceive the scene as a whole by combining multiple methods.

B. Holistic Scope of Latest Integrated Approaches

There are few approaches in research that integrate multiple scene perception tasks into a single framework. In contrast to the narrow perception techniques, they achieve a more holistic understanding of the scene. These approaches simultaneously recognize environmental information over a long time. Storing this data in a known structure offers new potential for a more complex spatial and temporal interpretation of the scene. We looked deeper into these approaches to extract how much they cover a holistic scene understanding. In the following, a comparison of the perceptual capabilities should answer this question. As criteria, we select research approaches covering more than a tabletop scene, reconstruct in real-time, and represent scene data by multiple types of instances, such as semantics. However, each approach set a different focus starting at sensor-close processing, such as on paralyzing, fusing, and handling of dynamics going to ontology-based reasoning. Table I shows the perceptual properties of these approaches, divided into the three steps of the perception transfer presented in Section III. If we could not find details to a criterion, we marked it either as not available (N/A) or not specified (NS). The approaches of Table I use a three-dimensional camera as sensory input providing color and depth information of its FOV. Additionally, some approaches use an IMU to support visual odometry for a more precise estimate of the ego-motion, e.g., needed for drones and wheeled-based robots. Based on the sensory input stream, the presented approaches combine several recognition techniques to reconstruct a virtual scene model. However, they cover the recognition differently. *KnowRob* [65], a knowledge representation and reasoning

framework, solely offers an interface for individual visual recognition systems. Wyatt et al. [67] limit the reconstruction to a metric map without recognizing semantics by vision sensors. Its scene recognition is trained from visual properties by a human tutor using supervised learning. Similarly, the *SOMA* framework restricts the reconstruction to a metric map but enhances objects by a CNN for color image-based object detection. *SOMA* aims at understanding changes in everyday environments by perceiving geometries and semantics. Alternatively, the approach of Suchan et al. [103] enhances the metric map by detecting walls, which are used for a clustering algorithm to generate a floor plan. The other approaches investigate further into a fully three-dimensional metric-semantic reconstruction of the scene. The *SLAM++* project of Salas et al. [104] is an early approach from 2013 that concentrates on semantic mapping. It consists of an object-based SLAM that uses object recognition trained on a database of scanned object models. It is capable of detecting changes in the environment, such as moving objects. *Fusion++* [100] set its focus on semantic mapping. It is similar to *SLAM++* but runs a Mask Region-Based Convolutional Neural Networks (R-CNN) object segmentation to initialize a Truncated Signed Distance Field (TSDF) reconstruction for each object. Rosinol et al. recently published *Kimera* [105], a multi-layer spatial scene perception framework aiming to close the gap between human and robot scene perception. *Kimera* uses a metric-semantic SLAM to perform a full mesh reconstruction by TSDF volumes with its semantic on top of localization. It recognizes building structures as well as objects from Computer-Aided Design (CAD) model match. In addition, human detection and pose estimation extend dynamic scene information [106].

The recognition techniques feed its scene information into the knowledge base where they are represented and connected in different ways. *Fusion++* and *SLAM++*, which concentrate on the semantic mapping, are not providing details of their knowledge representation. Contrastly, the approach of Wyatt et al. focuses on the representation of knowledge gaps and uncertainties [67]. A layered structure of proxies, unions, and beliefs represents the spatial scene inside a relational database. They validate by experiments in a lab that a human tutor is capable of helping a robot to fill a knowledge gap through verbal conversation. The robot asks for missing visual features, such as the color and shape, to prove an object is believed to close a knowledge gap. *Kimera*, which extensively covers the recognition, is also concentrating on spatial knowledge representation by multiple hierarchical layers, separated by the semantic. The spatial layers comprise the metric-semantic mesh, objects, structures, rooms, and buildings. Dynamic scene graphs simultaneously update scene information by linking the spatial scene information within the layers. The knowledge representation of *SOMA* comprises three layers inside a document-based database: observation layer, semantic layer, and interpretation layer. The scene information is organized similarly to *Kimera* with a spatial layout of hierarchical structures connected by graphs. Suchan et al. [103] propose using ontologies and

TABLE I
OVERVIEW OF INTEGRATIVE ROBOTIC SCENE PERCEPTION APPROACHES

	Recognition of Information				Knowledge Representation			Knowledge Interpretation	
	Sensory Input	Reconstruction	Static Instances	Dynamic Instances	Database	Knowledge Structure	Scene Representation	Spatial-Temporal	Reasoning
Wyatt et al. [67]	3D-cam	metric	objects from user input	N/A	relational	multi-layer spatial representation by proxies, unions, and beliefs	point map	place classification	belief verification by human
<i>SLAM++</i> Salas et al. [104]	3D-cam	metric-semantic	N/A	objects by scan model match	NS	single-layer spatial object graph	metric-semantic mesh	N/A	N/A
Suchan et al. [103]	3D-cam	metric-semantic	walls by planes	human pose detection	NS	spatial-temporal representation of entities	metric map, semantic by primitives	human activities, spatial relations and pattern	human-centered commonsense
<i>Fusion++</i> McCormac et al. [100]	3D-cam	metric-semantic	objects by R-CNN	N/A	NS	single-layer spatial object graph	metric-semantic mesh	N/A	N/A
<i>SOMA</i> Kunze et al. [63]	3D-cam, 2D-laser scanner	metric	N/A	objects, people by CNN	document-based	observation, semantic and interpretation layer	point map, objects by pose and BB	human activities	N/A
<i>KnowRob</i> Beetz et al. [65], [90]	NS	NS	NS	NS	relational	ontology graph, multi-level of metric-semantic, logic and episodic memories	mesh, poses	episodic memory for reasoning	hypotheses verification, inner world, motion control
<i>Kimera</i> Rosinol et al. [105], [106]	3D-cam, IMU	metric-semantic	building structures, objects by CAD model match	people by pose detection	NS	hierarchical graph connects spatial layers	metric-semantic mesh	place and room classification	N/A

formal characterization to represent knowledge by space and motion. The knowledge representation of *KnowRob* is the most powerful due to the usage of description logic for representing and connecting knowledge within multiple levels by the Web Ontology Language (OWL). The framework links an inner world, virtual, logical, and episodic memories [65].

Based on the represented knowledge, few approaches look deeper into the interpretation. For instance, the second pillar of *KnowRob* is knowledge-based reasoning to learn general knowledge. An object class-related affordance extraction reasons what to do with objects [90]. Suchan et al. [103] show by temporal human detection and object recognition how to use reasoning to enhance the scene understanding. The use of logical spatial relations between instances (e.g., an object is located left of another object) or the detection of human activities represent valuable information about the scene.

The presented approaches indicate that they already perform more than one perception process step on an advanced level. Due to the different focus of every approach, none sufficiently covers holistic scene perception. The degree of coverage is difficult to quantify due to missing metrics in research. However, the overview (see Table I) offers a starting point to identify gaps and potentials compared to human perception.

V. GAP TO HUMAN SCENE PERCEPTION

How well does robot scene perception mimic human-level performance nowadays? Former research compared the perceptual performance of robots with children’s age, such as Szeliski who claimed that computer vision reached the level of a two years old child [107]. This comparison might not be beneficial since robot perception is not quantifiable in human terms as there is rather a specific than a broad perceptual skill development of robots. Following this hypothesis, this paper details major gaps in robot perception in everyday scenes.

A. Non-Usage of Sensory Modalities

The first identified gap in robotic scene perception is the missing usage of multiple sensor modalities as the input source. Humans use different senses providing the opportunity to fuse and rely on the optimal sensor since each sensor modality is affected, such as vision by illumination, color, occlusion, and posture of objects [25] (see Section III-A). The hardware design of robots allows them an optimal sensory input due to flexible sensor choice, amount, and alignment. On the one hand, this underlines the statement of Prenebida [33] that robotic recognition tasks, such as vision-based object classification, could deliver higher performance than humans. However, on the other hand, all robotic approaches of Section IV-B limit the scene perception to visual sensors, except for ego-motion sensors, as a single modality. Although vision sensors provide the highest information content of the scene, using a single modality is insufficient for the situations mentioned above (localization of acoustics, haptic feedback of obstacles, darkroom problem). Therefore, current integrated scene perception approaches are not able to deploy different senses in the way that humans do.

B. Different Perceptual Learning

A fundamental difference between robots and humans is the acquisition of perceptual understanding. Perceptual learning has been first defined by Gibson as “any relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array” [108]. Thus, humans learn an individual perception without initial knowledge that is constantly adjusted and influenced by society and the environment. The learning enables humans to achieve an optimal perception within the surrounding since especially high-level scene interpretations depend on the subjective impression that is difficult to generalize (see duck-rabbit illusion [78]). In contrast to humans, popular perceptual techniques of robots use rigid learning strategies, which do not offer adaptation. It would require re-training or parameter adjustment of the initially deployed

models. Moreover, these approaches provide neither the flexibility for extension nor adaptation of the perception capabilities over time. All presented integrated robotic scene perception approaches (see Table I) build up pre-learned skills except for the method of Wyatt et al. [67] that allows modifications of the perception after deployment (see Section IV-B). Especially in the future, when robots will be deployed for a long time, the initial perception will become obsolete unless the perception is adjusted continuously to the environment. Therefore, a gap concerns the learning of perceptual capabilities during run-time. Although research on robot learning, such as domain adaptation [109] and continuous learning (lifelong learning, perceptual learning, and never-ending learning) [110], reaches back almost 30 years [111], it is still rarely used in practical applications. A common strategy is to learn from demonstration [112]. For instance, a human could teach the robot to adjust a generic perception to changes in the environment. The human can approve or declines estimates of the perception system via a user interface, such as classifying new objects or teaching novel actions. A possible solution for supervised learning by a human tutor has been presented by Wyatt et al. [67]. However, enabling robots to adapt to the environment requires new techniques. On the one hand, as proposed by Wyatt et al. [67], this technique could close knowledge gaps. On the other hand, learning environment-adapted perception skills once with a tutor will not avoid adjusting the perception within the deployed environment.

C. Lack of Commonsense Perception

Our third identified gap is the lack of commonsense perception caused by specific recognition capabilities. Fed by a high amount of training data, the artificial system becomes an expert for narrow recognition tasks. The previous arguments reveal several applications that hit the performance for specialized tasks. Although the recognition performance of a specific perception task can outperform humans, the fundamental issue is the lack of flexibility due to required re-training or missing capability to extend recognition. Therefore, even the latest approaches achieve just a low overall recognition performance (see Section IV-B).

Humans have individual perception skills depending on various influences, such as profession, culture, and age, since humans learn from practice and experience [108]. However, in a society, there are perception skills simplifying a commonsense understanding. It enables humans to execute trivial tasks, such as knowing how to open a door or how to use an elevator. Nowadays, trivial commonsense perception skills, such as detecting a door and its handle to move through or detecting the elevator buttons, are not default functionalities of mobile robots. Thus, robot recognition is not making sense of the whole scene. Similarly, the interpretation of scene data indicates specific capabilities. There are a few advanced approaches for interpreting a complex scene, such as *Knowrob* [65] (see Section III-C). However, they provide narrow interpretations independent of the available scene information, whereas robots can not instantly interpret an

unseen scene. Therefore, the knowledge interpretation may not retrieve sufficient scene knowledge.

VI. POTENTIALS OF ROBOTS

This section highlights potentials in robot scene perception going beyond the perceptual capabilities of humans.

A. Flexible Sensor Design

The first presented gap in robot scene perception (using few or solely a single sensor modality) can be overcome by the usage of flexible sensor design. Human perception is limited to recognized information of its senses, their range, and accuracy, e.g., fog or darkness negatively influences scene recognition. Contrastively, robots benefit from scene-adjusted sensor modalities and their flexible configuration. They can overcome limitations of human senses, such as through ultrasonic or radar recognition techniques [113]. Robots can be equipped with sensor modalities, such as radar, that enables them to freely adapt to their environment. The flexibility of the robot design allows the adaptation to the application. Thus, robots can use sensors with the desired amount, properties, and alignments. For instance, multiple visual sensors could enable robots to recognize the scene in 360 degrees. Herewith blind spots can be avoided, which is especially important for safe usage.

B. Initial Perception Capabilities

In contrast to the highlighted gap of missing perceptual learning, robots can be deployed with pre-learned perception capabilities and with initial scene knowledge. This trivial fact allows reducing the setup time of the robot to a minimum. For instance, mirroring the perception skills of an existing robot in an environment enables a new robot to perceive the scene equivalently. In contrast to robots, humans would have to learn from scratch.

C. Cooperative Perception

Speaking the “same language” in terms of understanding and exchanging information enables the sharing of perceptual data independent of the robot, assistive system, or infrastructure. Therefore, artificial systems have fewer restrictions than humans as the number of collaborators and the communication range and bandwidth are not hard-restricted. The sharing of scene information in real-time with multiple artificial agents enables the fusion of scene observations from various perspectives. The data exchange enhances the scene knowledge of every agent, which offers a comprehensive scene perception from its scene overview. A popular research area presents cooperative approaches that share navigation data, e.g., to distribute the mapping [64] or for reactive path planning of multiple robots [114]. Moreover, possible decentralized computation, e.g., cloud computing, can save resources of mobile robots [115]. Therefore, cooperative perception is a key technique for reducing the setup time of large environments, safe navigation, and better economy of larger robot fleets through its inter-communication [114], [116].

VII. FUTURE DIRECTIONS

We propose two comparatively non-popular future directions which contribute to deploying robots within our society. First, developing frameworks to combine multiple perception modules. The scope of robotic scene perception is low since fundamental commonsense skills are missing. Commonsense skills can be generated with scene perception frameworks that combine several specific methods to achieve a holistic understanding. Second, although existing research on robot perception achieves good performance on specific tasks, flexibility for adjusting the perception within the robot usage is missing. It must be possible to extend initially deployed perception skills and to adjust the perception at any time to the scene. Therefore, we address especially our second and third identified gaps to future research that underlines the importance of perceptual learning being fundamental in everyday life. It is a central technique toward and beyond human-like perceptual performance.

VIII. CONCLUSION

This article transfers human scene perception to mobile robots for comparison since the performance of human perception is superior in many tasks. Current research in robotics presents specific perception skills evaluated in a predefined and constricted manner, whereby they already achieve promising results for everyday applications. However, for lifelong unsupervised and autonomous usage, especially multi-functional mobile robots need to perceive the scene holistically to reliably handle unknown and changing situations, as well as uncertainties and dynamics. Human perception has been extensively studied since the 1970s, offering mature neuroscience studies and theories that define human perception as the process of recognizing, representing, and interpreting scene information. This article uses this three-fold division for a transfer from human to robot scene perception to identify similarities and differences in the process. The transfer revealed that the robotic approaches partly mirror human-like perception. However, much research investigates specific methods, such as object classification, which outperform human perception. This prerequisite towards and beyond human-like perceptual performance is promising. A new research area integrates multiple state-of-the-art methods in frameworks that aim toward a holistic scene understanding. However, these frameworks lack trivial commonsense perception skills and therefore can not substantiate a holistic scene understanding. Moreover, only non-popular research contributes to the perceptual learning of robots, needed to learn, adjust and customize their perception. Therefore, these two major gaps need to be addressed by future research.

ACKNOWLEDGMENT

This research has received funding from the German Ministry of Economic Affairs, Labour and Tourism Baden-Württemberg by the AI Innovation Center *Learning Systems and Cognitive Robotics* under grant agreement No 017-180036 (841022).

REFERENCES

- [1] C. Müller, B. Graf, and K. Pfeiffer, "World Robotics 2021 - Service Robots," IFR Statistical Department, Tech. Rep., 2021.
- [2] A. Gasparetto and L. Scalera, "A Brief History of Industrial Robotics in the 20th Century," in *Advances in Historical Studies*, vol. 8, 2019.
- [3] U. Neisser, "Without Perception, there is no Knowledge: Implications for Artificial Intelligence," *Natural and Artificial Minds*, 1993.
- [4] M. Naseer, S. Khan, and F. Porikli, "Indoor Scene Understanding in 2.5/3D for Autonomous Agents: A Survey," *IEEE Access*, vol. 7, 2019.
- [5] G. L. Malcolm, I. I. A. Groen, and C. I. Baker, "Making Sense of Real-World Scenes," *Trends in cognitive sciences*, vol. 20, no. 11, 2016.
- [6] S. Garg, N. Sünderhauf, F. Dayoub, D. Morrison, A. Cosgun, G. Carneiro, Q. Wu, T.-J. Chin, I. Reid, S. Gould, P. Corke, and M. Milford, "Semantics for Robotic Mapping, Perception and Interaction: A Survey," *Foundations and Trends® in Robotics*, vol. 8, 2020.
- [7] H. Levesque and G. Lakemeyer, "Cognitive Robotics," in *Handbook of Knowledge Representation*, ser. Foundations of Artificial Intelligence. Elsevier, 2008, vol. 3.
- [8] M. Potter, "Meaning in Visual Search," *Science*, vol. 187, no. 4180, 1975.
- [9] I. Biederman, "Perceiving Real-World Scenes," *Science*, vol. 177, no. 4043, 1972.
- [10] R. A. Epstein and C. I. Baker, "Scene Perception in the Human Brain," *Annual Review of Vision Science*, vol. 5, no. 1, 2019.
- [11] J. Beck, B. Hope, and A. Rosenfeld, *Human and Machine Vision: 1st Edition*. Springer Verlag, 1983.
- [12] A. Oliva and A. Torralba, "Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope," *International Journal Computer Vision*, vol. 42, no. 3, 2001.
- [13] C. M. Funke, J. Borowski, K. Stosio, W. Brendel, T. S. A. Wallis, and M. Bethge, "Five Points to Check when Comparing Visual Perception in Humans and Machines," *Computer Vision and Pattern Recognition*, 2020.
- [14] L. San Roque, K. Kendrick, E. Norcliffe, P. Brown, R. Defina, M. Dingemans, T. Dirksmeyer, N. Enfield, S. Floyd, J. Hammond, G. Rossi, S. Tufvesson, S. van Putten, and A. Majid, "Vision Verbs Dominate in Conversation Across Cultures, but the Ranking of Non-Visual Verbs Varies," *Cognitive Linguistics*, vol. 26, 2015.
- [15] M. Ward, G. Grinstein, and D. Keim, *Interactive Data Visualization: Foundations, Techniques, and Applications*. USA: A. K. Peters, Ltd., 2015.
- [16] L. Seminara, P. Gastaldo, S. Watt, K. Valyear, F. Zuher, and F. Mastrogiovanni, "Active Haptic Perception in Robots: A Review," *Frontiers in Neurobotics*, vol. 13, 2019.
- [17] S. Deutsch and A. Deutsch, *The Eye as a Transducer*. IEEE Press, 1993.
- [18] D. Purves, *Neuroscience*, 6th ed. Oxford University Press, 2018.
- [19] P. Gärdenfors, *Conceptual Spaces - the Geometry of Thought*. MIT Press, 2000.
- [20] P. Gärdenfors, *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press, 2014.
- [21] A. Dix, M. Pohl, and G. Ellis, *Perception and Cognitive Aspects*. Eurographics Association, 2010, ch. 7.
- [22] J. Wolfe, N. Klempen, and K. Dahlen, "Postattentive Vision," *Journal of experimental psychology. Human perception and performance*, vol. 26, 2000.
- [23] S. E. Palmer, "Hierarchical Structure in Perceptual Representation," *Cognitive Psychology*, vol. 9, no. 4, 1977.
- [24] J. Wagemans and R. Kimchi, *The Perception of Hierarchical Structure*. Oxford University Press, 2014.
- [25] S. Jin, H. Liu, B. Wang, and F. Sun, "Open-Environment Robotic Acoustic Perception for Object Recognition," *Frontiers in Neurobotics*, vol. 13, 2019.
- [26] C. Evers and P. A. Naylor, "Acoustic SLAM," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 9, 2018.
- [27] M. Kreković, I. Dokmanić, and M. Vetterli, "EchoSLAM: Simultaneous localization and mapping with acoustic echoes," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016.

- [28] A. Balachandran, M. Brown, S. Erlien, and J. Gerdes, "Predictive Haptic Feedback for Obstacle Avoidance Based on Model Predictive Control," *IEEE Transactions on Automation Science and Engineering*, vol. 13, 2015.
- [29] C. Lytridis, G. Virk, and E. Kadar, "Co-operative smell-based navigation for mobile robots," in *Climbing and Walking Robots*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 1107–1117.
- [30] H. Shimazu, K. Kobayashi, A. Hashimoto, and T. Kameoka, "Tasting Robot with an Optical Tongue: Real Time Examining and Advice Giving on Food and Drink," in *Human Interface and the Management of Information. Methods, Techniques and Tools in Information Design*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007.
- [31] B. Ciui, A. Martin, R. K. Mishra, T. Nakagawa, T. J. Dawkins, M. Lyu, C. Cristea, R. Sandulescu, and J. Wang, "Chemical Sensing at the Robot Fingertips: Toward Automated Taste Discrimination in Food Samples," *ACS Sensors*, vol. 3, no. 11, 2018.
- [32] G. Pintore, C. Mura, F. Ganovelli, L. Fuentes Perez, R. Pajarola, and E. Gobbetti, "State-of-the-art in Automatic 3D Reconstruction of Structured Indoor Environments," *Computer Graphics Forum*, vol. 39, 2020.
- [33] C. Premebida, R. Ambrus, and Z. Marton, "Intelligent Robotic Perception Systems," in *Applications of Mobile Robots*, 2018.
- [34] K. Klasing, D. Althoff, D. Wollherr, and M. Buss, "Comparison of Surface Normal Estimation Methods for Range Sensing Applications," *IEEE International Conference on Robotics and Automation*, 2009.
- [35] R. Bormann, J. Hampp, M. Hägele, and M. Vincze, "Fast and Accurate Normal Estimation by Efficient 3D Edge Detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.
- [36] A. Nguyen and B. Le, "3D Point Cloud Segmentation: A Survey," in *IEEE Conference on Robotics, Automation and Mechatronics (RAM)*, 2013.
- [37] S. Minaee, Y. Y. Boykov, F. Porikli, A. J. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image Segmentation Using Deep Learning: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [38] J. Canny, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, 1986.
- [39] A. Kaiser, J. A. Y. Zepeda, and T. Boubekeur, "A Survey of Simple Geometric Primitives Detection Methods for Captured 3D Data," *Computer Graphics Forum*, vol. 38, no. 1, 2019.
- [40] P. J. Burt and E. H. Adelson, "The Laplacian Pyramid as a Compact Image Code," in *Readings in Computer Vision*. San Francisco (CA): Morgan Kaufmann, 1987.
- [41] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'12. Red Hook, NY, USA: Curran Associates Inc., 2012.
- [42] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [43] D. T. Nguyen, W. Li, and P. O. Ogunbona, "Human detection from images and videos: A survey," *Pattern Recognition*, vol. 51, 2016.
- [44] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh, "OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [45] 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*, vol. 32, no. 1, 2016.
- [46] M. Labbé and F. Michaud, "RTAB-Map as an Open-Source Lidar and Visual Simultaneous Localization and Mapping Library for Large-Scale and Long-Term Online Operation," *Journal of Field Robotics*, vol. 36, 2018.
- [47] 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*, vol. 33, no. 5, 2017.
- [48] I. Kostavelis and A. Gasteratos, "Semantic Mapping for Mobile Robotics Tasks: A Survey," *Robotics and Autonomous Systems*, vol. 66, 2015.
- [49] F. Poux, "The Smart Point Cloud: Structuring 3D intelligent point data," Ph.D. dissertation, Université de Liège, 2019.
- [50] R. Wang and J. Brockmole, "Human Navigation in Nested Environments," *Journal of experimental psychology. Learning, memory, and cognition*, vol. 29 3, 2003.
- [51] G. Welch and G. Bishop, "An Introduction to the Kalman Filter," 1995.
- [52] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund, "Particle Filters for Positioning, Navigation, and Tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, 2002.
- [53] Y. Zhang, C. Wang, X. Wang, W. Zeng, and W. Liu, "A Simple Baseline for Multi-Object Tracking," *CoRR*, 2020.
- [54] N. Wojke, A. Bewley, and D. Paulus, "Simple Online and Realtime Tracking with a Deep Association Metric," *CoRR*, 2017.
- [55] P. Chu, J. Wang, Q. You, H. Ling, and Z. Liu, "TransMOT: Spatial-Temporal Graph Transformer for Multiple Object Tracking," *CoRR*, 2021.
- [56] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *CoRR*, 2020.
- [57] X. Li and Z. Zhou, "Object Re-Identification Based on Deep Learning," in *Visual Object Tracking with Deep Neural Networks*. Rijeka: IntechOpen, 2019, ch. 5.
- [58] V. Bansal, S. James, and A. Del Bue, "Re-Obj: Jointly Learning the Foreground and Background for Object Instance Re-Identification," in *Image Analysis and Processing – ICIAP 2019*. Berlin, Heidelberg: Springer-Verlag, 2019.
- [59] S. He, H. Luo, P. Wang, F. Wang, H. Li, and W. Jiang, "TransReID: Transformer-based Object Re-Identification," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [60] Y. Zhao, S. Zhu, D. Wang, and Z. Liang, "Short Range Correlation Transformer for Occluded Person Re-Identification," *CoRR*, 2022.
- [61] H. Liu, J. Feng, M. Qi, J. Jiang, and S. Yan, "End-to-End Comparative Attention Networks for Person Re-Identification," *IEEE Transactions on Image Processing*, vol. PP, 2016.
- [62] B. Hommel, J. Gehrke, and L. Knuf, "Hierarchical Coding in the Perception and Memory of Spatial Layouts," *Psychological Research*, vol. 64, 2000.
- [63] L. Kunze, H. Karaoguz, J. Young, F. Jovan, J. Folkesson, P. Jensfelt, and N. Hawes, "SOMA : A Framework for Understanding Change in Everyday Environments Using Semantic Object Maps," in *Conference on Artificial Intelligence (AAAI-18)*, 2018.
- [64] M. Labbé and F. Michaud, "Long-Term Online Multi-Session Graph-Based SPLAM with Memory Management," *Autonomous Robots*, vol. 42, no. 6, 2018.
- [65] M. Beetz, D. Beßler, A. Haidu, M. Pomarlan, A. Bozcuoglu, and G. Bartels, "Know Rob 2.0 — A 2nd Generation Knowledge Processing Framework for Cognition-Enabled Robotic Agents," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [66] L. Prieto González, V. Stantchev, and R. Colomo-Palacios, "Applications of Ontologies in Knowledge Representation of Human Perception," *International Journal of Metadata, Semantics and Ontologies*, vol. 9, 2014.
- [67] J. Wyatt, A. Aydemir, M. Brenner, M. Hanheide, N. Hawes, P. Jensfelt, M. Kristan, G.-J. Kruijff, P. Lison, A. Pronobis, K. Sjöö, A. Vrecko, H. Zender, M. Zillich, and D. Skočaj, "Self-Understanding and Self-Extension: A Systems and Representational Approach," *IEEE Transactions on Autonomous Mental Development*, vol. 2, 2011.
- [68] R. Ravichandran, E. Prassler, N. Huebel, and S. Blumenthal, "A Workbench for Quantitative Comparison of Databases in Multi-Robot Applications," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [69] A. Dietrich, S. Mohammad, S. Zug, and J. Kaiser, "Ros meets cassandra: Data management in smart environments with nosql," in *11th International Baltic Conference DB and IS*, 2014.
- [70] M. Tenorth, U. Klank, D. Pangercic, and M. Beetz, "Web-enabled Robots – Robots that Use the Web as an Information Resource," *Robotics & Automation Magazine*, 2011.
- [71] P. Philipp, M. Maleshkova, A. Rettinger, and D. Katic, "A Semantic Framework for Sequential Decision Making," in *Engineering the Web in the Big Data Era*. Cham: Springer International Publishing, 2015.
- [72] C. Follini, M. Terzer, C. Marcher, A. Giusti, and D. T. Matt, "Combining the Robot Operating System with Building Information Modeling for Robotic Applications in Construction Logistics," in

Advances in Service and Industrial Robotics. Springer International Publishing, 2020.

- [73] L. Isik, A. Tacchetti, and T. Poggio, "A Fast, Invariant Representation for Human Action in the Visual System," *Journal of neurophysiology*, vol. 119, no. 2, 2018.
- [74] Y. Peng, "Causal Action: A Framework to Connect Action Perception and Understanding," Ph.D. dissertation, UCLA, 2019.
- [75] J. Henderson and A. Hollingworth, "High-Level Scene Perception," *Annual review of psychology*, vol. 50, 1999.
- [76] M. H. Segall, D. T. Campbell, and M. J. Herskovit, "The Influence of Culture on Visual Perception," in *Social Perception*, 1966.
- [77] S. Torresin, G. Pernigotto, F. Cappelletti, and A. Gasparella, "Combined Effects of Environmental Factors on Human Perception and Objective Performance: A Review of Experimental Laboratory Works," *Indoor Air*, vol. 28, no. 4, 2018.
- [78] P. Brugger and S. Brugger, "The Easter Bunny in October: Is it Disguised as a Duck?" *Perceptual and Motor Skills*, vol. 76, no. 2, 1993.
- [79] M. Ozturk, M. Ersen, M. Kapotoglu, C. Koc, S. Sariel, and H. Yalçın, "Scene Interpretation for Self-Aware Cognitive Robots," *AAAI Spring Symposium - Technical Report*, 2014.
- [80] Z. Jiang, B. Liu, S. Schuster, Z. Wang, and M. Chandraker, "Peek-a-Boo: Occlusion Reasoning in Indoor Scenes With Plane Representations," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [81] A. Persson, P. Z. Dos Martires, L. De Raedt, and A. Loutfi, "Semantic Relational Object Tracking," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 12, no. 1, 2019.
- [82] X. Zhang, Z. Wu, and Y. G. Jiang, "SAM: Modeling Scene, Object and Action with Semantics Attention Modules for Video Recognition," *IEEE Transactions on Multimedia*, 2021.
- [83] L. Piyathilaka and S. Kodagoda, "Human Activity Recognition for Domestic Robots," in *Field and Service Robotics*. Springer, 2015.
- [84] P. Duckworth, M. Alomari, Y. Gatsoulis, D. C. Hogg, and A. G. Cohn, "Unsupervised Activity Recognition Using Latent Semantic Analysis on a Mobile Robot," in *IOS Press Proceedings*, 2016.
- [85] I. Kostavelis, M. Vasileiadis, E. Skartados, A. Kargakos, D. Giakoumis, C.-S. Bouganis, and D. Tzovaras, "Understanding of Human Behavior with a Robotic Agent Through Daily Activity Analysis," *International Journal of Social Robotics*, vol. 11, no. 3, 2019.
- [86] K. Li, J. Wu, X. Zhao, and M. Tan, "Real-Time Human-Robot Interaction for a Service Robot Based on 3D Human Activity Recognition and Human-Mimicking Decision Mechanism," in *IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems*. IEEE, 2018.
- [87] F. A. Gers, J. Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *International Conference on Artificial Neural Networks ICANN*, 1999.
- [88] F. Rezazadegan, S. Shirazi, B. Uprofit, and M. Milford, "Action Recognition: From Static Datasets to Moving Robots," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017.
- [89] P. Philipp, M. Bommersheim, S. Robert, and J. Beyerer, "Probabilistic Estimation of Human Interaction Needs in Context of a Robotic Assistance in Geriatrics," *Current directions in biomedical engineering*, vol. 5, no. 1, 2019.
- [90] D. Beßler, R. Porzel, P. Mihai, M. Beetz, R. Malaka, and J. Bateman, "A Formal Model of Affordances for Flexible Robotic Task Execution," in *Proc. of the 24th European Conference on Artificial Intelligence (ECAI)*, 2020.
- [91] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014.
- [92] L. Alzubaidi, J. Zhang, A. J. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma, J. Santamaría, M. A. Fadhel, M. Al-Amidie, and L. Farhan, "Review of Deep Learning: Concepts, CNN Architectures, Challenges, Applications, Future Directions," *Journal of big data*, vol. 8, no. 1, 2021.
- [93] W. S. McCulloch and W. Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity," *The bulletin of mathematical biophysics*, vol. 5, no. 4, 1943.
- [94] D. Purves, G. Augustine, D. Fitzpatrick, W. Hall, A. Lamantia, and L. White, *Neuroscience 5th Edition*. Sunderland: Sinauer, 2011.
- [95] L. Jiao and J. Zhao, "A Survey on the New Generation of Deep Learning in Image Processing," *IEEE Access*, vol. 7, 2019.
- [96] K. O'Shea and R. Nash, "An Introduction to Convolutional Neural Networks," *ArXiv*, 2015.
- [97] M. Everingham, L. Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal Visual Object Classes (VOC) Challenge," *International Journal of Computer Vision*, vol. 88, 2009.
- [98] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, 2015.
- [99] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *Computer Vision – ECCV 2014*. Cham: Springer International Publishing, 2014.
- [100] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, "Fusion++: Volumetric Object-Level SLAM," in *2018 International Conference on 3D Vision (3DV)*, 2018.
- [101] Y. Lin, J. Tremblay, S. Tyree, P. Vela, and S. Birchfield, "Multi-view Fusion for Multi-level Robotic Scene Understanding," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [102] A. Koubaa, *Robot Operating System (ROS): The Complete Reference (Volume 1)*, 1st ed. Springer Publishing Company, Incorporated, 2016.
- [103] J. Suchan and M. Bhatt, "Commonsense Scene Semantics for Cognitive Robotics: Towards Grounding Embodied Visuo-Locomotive Interactions," in *IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017.
- [104] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly, and A. J. Davison, "SLAM++: Simultaneous Localisation and Mapping at the Level of Objects," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [105] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an Open-Source Library for Real-Time Metric-Semantic Localization and Mapping," *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [106] A. Rosinol, A. Gupta, M. Abate, J. Shi, and L. Carlone, "3D Dynamic Scene Graphs: Actionable Spatial Perception with Places, Objects, and Humans," in *Robotics: Science and Systems (RSS)*, 2020.
- [107] R. Szeliski, *Computer Vision: Algorithms and Applications*, 1st ed. Berlin, Heidelberg: Springer-Verlag, 2010.
- [108] E. J. Gibson, "Perceptual Learning," *Annual Review of Psychology*, vol. 14, no. 1, 1963.
- [109] A. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, "A Brief Review of Domain Adaptation," in *Advances in Data Science and Information Engineering*. Cham: Springer International Publishing, 2021.
- [110] H. Qin and D. Zhang, "A Perpetual Learning Algorithm That Incrementally Improves Performance With Deliberation," *IEEE Access*, vol. 8, 2020.
- [111] S. Thrun and T. M. Mitchell, "Lifelong robot learning," *Robotics and Autonomous Systems*, vol. 15, no. 1, 1995, the Biology and Technology of Intelligent Autonomous Agents.
- [112] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A Survey of Robot Learning from Demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, 2009.
- [113] C.-C. Carbon, "Understanding Human Perception by Human-Made Illusions," *Frontiers in Human Neuroscience*, vol. 8, 2014.
- [114] S. Dörr, *Cloud-based Cooperative Long-term SLAM for Mobile Robots in Industrial Applications*, ser. Stuttgarter Beiträge zur Produktionsforschung. Fraunhofer Verlag, 2020.
- [115] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg, "A Survey of Research on Cloud Robotics and Automation," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 2, 2015.
- [116] A. Miller, K. Rim, P. Chopra, P. Kelkar, and M. Likhachev, "Cooperative Perception and Localization for Cooperative Driving," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.