Embodied AI-Driven Operation of Smart Cities: A Concise Review

Farzan Shenavarmasouleh¹, Farid Ghareh Mohammadi¹,
M. Hadi Amini², and Hamid R. Arabnia¹

1:Department of Computer Science, Franklin College of arts and sciences,
University of Georgia, Athens, GA, USA

2: School of Computing & Information Sciences, College of Engineering & Computing,
Florida International University, Miami, FL, USA

Emails: fs04199@uga.edu, farid.ghm@uga.edu, amini@cs.fiu.edu, hra@uga.edu

August 22, 2021

Abstract

A smart city can be seen as a framework, comprised of Information and Communication Technologies (ICT). An intelligent network of connected devices that collect data with their sensors and transmit them using wireless and cloud technologies in order to communicate with other assets in the ecosystem plays a pivotal role in this framework. Maximizing the quality of life of citizens, making better use of available resources, cutting costs, and improving sustainability are the ultimate goals that a smart city is after. Hence, data collected from these connected devices will continuously get thoroughly analyzed to gain better insights into the services that are being offered across the city; with this goal in mind that they can be used to make the whole system more efficient. Robots and physical machines are inseparable parts of a smart city. Embodied AI is the field of study that takes a deeper look into these and explores how they can fit into real-world environments. It focuses on learning through interaction with the surrounding environment, as opposed to Internet AI which tries to learn from static datasets. Embodied AI aims to train an agent that can See (Computer Vision), Talk (NLP), Navigate and Interact with its environment (Reinforcement Learning), and Reason (General Intelligence), all at the same time. Autonomous driving cars and personal companions are some of the examples that benefit from Embodied AI nowadays. In this paper, we attempt to do a concise review of this field. We will go through its definitions, its characteristics, and its current achievements along with different algorithms, approaches, and solutions that are being used in different components of it (e.g. Vision, NLP, RL). We will then explore all the available simulators and 3D interactable databases that will make the research in this area feasible. Finally, we will address its challenges and identify its potentials for future research.

keywords: Embodied AI, Embodied Intelligence, Question Answering, Smart Cities, Simulation, Intelligence, Multi-Agent Systems

1 Introduction

A smart city is an urban area that employs Information and Communication Technologies (ICT) [1], an intelligent network of connected devices and sensors that can work interdependently [2, 3] and a distributive manner [4] to continuously monitor the environment, collect data, and share them among the other assets in the ecosystem. A smart city uses all the available data to make real-time decisions about the many individual

components of the city to ease up the livelihood of its citizens, and make the whole system more efficient, more environmentally friendly, and more sustainable [5]. This serves as a catalyst for creating a city with faster transportation, fewer accidents, enhanced manufacturing, more reliable medical services and utilities, less pollution [6] and much more. The good news is any city, even with traditional infrastructures, can be transformed into a Smart City by integrating IoT technologies [7].

An undeniable part of a smart city is its use of smart agents. These agents can vary a lot in sizes, shapes, and functionalities. They can simply be light sensors that along with their controller act as the energy-saving agents or could be more advanced machines, with complicated controllers and interconnected components that are capable of tackling more advanced problems. The latter agents usually come with an embodiment with numerous sensors and controllers built in them that enable them to perform high-level and human-level tasks such as talking, walking, seeing, and complex reasoning along with the ability to interact with the environment. Embodied Artificial Intelligence is the field of study that takes a deeper look into these agents and explores how they can fit into the real-world and how they can eventually act as our future community workers, personal assistants, robocops, and much more.

Imagine arriving home after a long working day and seeing your home robot waiting for you at the entrance door. Although it is not the most romantic thing ever, you then walk up to it and ask it to make a cup of coffee for you and also add two teaspoons of sugar if there is any in the cabinet. For this to become reality, the robot has to have a vast range of skills. It should be able to understand your language and be able to translate questions and instructions to the action. It should be able to see its surroundings and have the ability to recognize objects and scenes. Last but not the least, it must know how to navigate in a big dynamic environment, interact with the objects within it, and be capable of doing long-term planning and reasoning.

In the past few years, there has been significant progress in the fields of computer vision, natural language processing, and reinforcement learning thanks to the advancements in deep learning models. Many things are now possible because of these that seemed impossible a few years ago. However, most of the work has been done in isolation from other lines of work. Meaning that the trained model can only take one type of data (eg. image, text, video) as the input and perform a single task that it is asked for. Consequently, such a model act as a single-sensory machine as opposed to a multi-sensory one. Also, for the most part, they all belong to Internet AI rather than Embodied AI. The goal of Internet AI is learning patterns in text, images, and videos from the datasets collected from the internet.

If we zoom out and look at the way models in Internet AI are being trained, we realize that generally supervised classification is the way to go. For instance, we provide a certain number of dog and cat photos along with the corresponding labels to a perception model and if the number is large enough, the model then can successfully learn the differences that exist between these two animals and discriminate between them. Learning via flashcards falls under the same umbrella for humans.

Extensive amount of time has been devoted in the past years to gather and build huge datasets for the imaging and language communities. A few considerable markers of this can be IMAGENET [8], MS COCO [9], Sun [10], Caltech-256 [11], Places [12] created for vision tasks; Squad [13], Glue [14], Swag [15] built for language objectives; and also Visual Genome [16] and VQA [17] datasets created for joint purposes to name a few.

Apart from playing a pivotal role in the recent advances of the main fields, these datasets also proved to be useful when used with transfer learning methods to help underlying disciplines such as biomedical imaging [18, 19, 20]. However, the aforementioned datasets are prone to restrictions. Firstly, at times it can get extremely costly, both in terms of time and money, to gather all the required data for the collection and label them. Secondly, the collection has to be monitored constantly to assure that they follow certain rules to avoid creating biases that could lead to erroneous results in future works [21] and also make sure that the collected data are all normal and uniform in terms of attributes such as background, size, position of the objects, lighting conditions, etc. while in contrast, we know that in real-world scenarios this cannot be the case and robots have to deal with a mixture of unnormalized noisy irrelevant data plus the relevant well-curated ones. Additionally, the agent would be able to interact with the objects in the wild (e.g. picking it up and looking at the object from another

Figure 1: Embodied AI in Smart Cities

angle) and also use its other senses such as smell and hearing to collect information.

Us humans, we do learn from interactions and it's a must for true intelligence in the real world. In fact, it's not only humans and all the other animals do the same. In, kitten carousel experiment [22], Held and Hein exhibited this beautifully. They studied the visual development of two kittens in a carousel over time. One of which had the ability to touch the ground and control its motions within the restrictions of the device while the other was just a passive observer. At the end of the experiment, they found out that the visual development of the former kitten was normal whereas for the latter one it was not, even though they both saw the same thing. This proves that being able to physically experience the world and interact with it is a key element for learning [23].

The goal of Embodied AI is to bring the ability to interact and being able to use multi senses simultaneously into play to enable the robot to continuously learn in a lightly supervised or even unsupervised way in a rich dynamic environment.

2 Rise of the Embodied AI

In the mid-1980s a major paradigm shift took place towards embodiment and computer science started to become more practical than theoretical algorithms and approaches. Embedded systems started to appear in all kinds of forms to aid humans in everyday life. Controllers for trains, airplanes, elevators, air conditioners, and Softwares for translation and audio manipulation are some of the most important ones to name a few [24].

Embodied Artificial Intelligence is a broad term, and those successes were for sure great ones to start with; yet, it could clearly be seen that it was a huge room for improvement. Theoretically, the ultimate goal of AI is not only to master any given algorithm or task that is given to, but also gain the ability to multitask and get to human-level intelligence, and that as mentioned requires meaningful interaction with the real world. There are many specialized robots for a vast set of tasks out there, especially in large industries, which can do the assigned task to them to perfection, let it be cutting different metals, painting, soldering circuits, and much more, but until one single machine emerges to have the ability to do different tasks or at least a small subset of them by itself and not just by following orders, it cannot be called intelligence.

Humanoids are the main thing that comes to mind when we talk about robots with intelligence. Although it is the ultimate goal, it is not the only form of intelligence on the earth. Other animals, such as insects have their own kind of intelligence and due to being relatively simpler in comparison to humans, they are a very good place, to begin with.

Rodney Brooks has a famous argument that says it took the evolution much longer to create insects from scratch than getting to human-level intelligence from there. Consequently, he suggested that these simpler biorobotics should be first dealt with in the road to make much more complex ones. Genghis, a six-legged walking robot [25] is one of his contributions to this field.

This line of thought was a fundamental change and led researchers to have a change of direction in their work and with that came attention to new domains and topics such as robotics, locomotion, artificial life, bio-inspired systems, and much more. The classical approach did not care about tasks related to interaction with the real world and consequently, locomotion and grasping were the ones to start the journey with.

Since not much computational power was available at the time of this shift, a big challenge for the researchers was the trade-off between simplicity and the potential to operate in complex environments. An extensive amount of work has been done in this area to explore or invent ways to exploit natural body dynamics, materials used in the modules, and their morphologies to make the robots move and become able to grasp and manipulate items without sophisticated processing units [26, 27, 28]. It goes without saying that the ones who could use the physical properties of themselves and the environment to function were more energy-efficient, but they had their own limitations. Not being able to generalize well to complex environments was a major drawback. But, they were fast as the machines with huge processing units needed a reasonable amount of time to think and plan their next action and often move their rigid and non-smooth actuators.

Nowadays, a big part of these issues are solved and we can see extremely fast and smooth natural moving robots capable of doing different types of maneuvers [29], but yet it is foreseen that with the advances of artificial muscles, joints, and tendons this progress can be further improved.

3 Breakdown of Embodied AI

In this section, we try to categorize a broad range of research that has been done under the field of Embodied AI. Due to the huge diversity, each section will necessarily be abstract, selective, and reflect the authors' personal opinion.

3.1 Language Grounding

Machine and human communication has always been a topic of interest. As time goes on, more and more aspects of our lives are controlled by AIs, and hence it is crucial to have ways to talk with them. This is a must for giving new instructions to them or receiving an answer from them, and since we are talking about general day to day machines, we desire this interface to be higher level than programming languages and closer to spoken language. To achieve this, machines must be capable of relating language to actions and the world. Language grounding is the field that tries to tackle this and map natural language instructions to robot behavior.

Hermann et al.'s work show that this can be achieved by rewarding an agent upon successful execution of written instructions in a 3D environment with a combination of unsupervised learning and reinforcement learning [30]. They also argue that their agent can generalize well after training and can interpret new unseen instructions and operate in unfamiliar situations.

3.2 Language plus Vision

Now that we know that machines can understand languages and there exist sophisticated models just for this purpose out there [31], it is time to bring another sense into play. One of the most popular ways to show the

potential of joint training of vision and language is the image and video captioning [32, 33, 34, 35, 36].

More recently, a new line of work has been introduced to take advantage of this connection. Visual Question Answering (VQA) [17] is the task of receiving an image along with a natural language question about that image as an input and attempting to find the accurate natural language answer for it as the output. The beauty of this task is that both the questions and the answers can be open-ended and also the questions can target different aspects of the image such as the objects that are present in them, their relationship or relative positions, colors, and background.

Following this research, Singh et al. [37] cleverly added an OCR module to the VQA model to enable the agent to read the texts available in the image as well and answer questions asked from them or use the additional context indirectly to answer the question better.

One may ask where does the new task stands relative to the previous one. Do agents who can answer questions more intelligent than the ones who deal with captions or not? The answer is yes. In [17] the authors show that VQA agents need a deeper and more detailed understanding of the image and reasoning than models for captioning.

3.3 Embodied Visual Recognition

Passive or fixed agents may fail to recognize objects in scenes if they are partially or heavily occluded. Embodiment comes to the rescue here and gifts the possibility of moving in the environment to actively control the viewing position and angle to remove any ambiguity in object shapes and semantics.

Jayaraman et al. [38] started to learn representations that will exploit the link between how the agent moves and how it will affect its visual surrounding. To do this they used raw unlabeled videos along with an external GPS sensor that provided the agent’s coordinates and trained their model to learn a representation linking these two. So, after this, the agent would have the ability to predict the outcome of its future actions and guess how the scene would look like after moving forward or turning to a side.

This was powerful and in a sense, the agent developed imagination. But, there was an issue here. If we pay attention we realize that the agent is still being fed pre-recorded video as the input and is learning similar to the observer kitten in the kitten carousel experiment explained above. So, following this, the authors went after this problem and proposed to train an agent that takes any given object from an arbitrary angle and then predict or better to say imagine the other views by finding the representation in a self-supervised manner [39].

Up until this point, the agent does not use the sound of its surroundings while humans are all about experiencing the world in a multi-sensory manner. We can see, hear, smell, touch all at the same time, and extract and use the relevant information that could be beneficial to our task at hand. All that said, understanding and learning the sound of objects present in a scene is not easy since all the sounds are overlapped and are being received via a single channel sensor. This is often dealt with as an audio source separation problem and lots of work has been done on it in the literature [40, 41, 42, 43, 44].

Now it was the reinforcement learning turn to make a difference. Policies have to be learned to aid agents move around a scene and this is the task of active recognition [45, 46, 47, 48, 49]). The policy will be learned at the same time it is learning other tasks and representation and it will tell the agent where and how to strategically move to recognize things faster [50, 51].

Results show that policies indeed help the agent to achieve better visual recognition performance and the agents can strategize their future moves and path for better results that are mostly different from shortest paths [52].

3.4 Embodied Question Answering

Embodied Question Answering brings QA into the embodied world. The task starts by an agent being spawned at a random location in a 3D environment and asked a question which its answer can be found somewhere in the

environment. In order for the agent to answer it, it must first strategically navigate to explore the environment, gathers necessary data via its vision, and then answer the question when the agent finds it [53, 54].

Following this, Das et al. [55] also presented a modular approach to further enhance this process by teaching the agent to break the master policy into sub-goals that are also interpretable by humans and execute them to answer the question. This proved to increase the success rate.

3.5 Interactive Question Answering

Interactive Question Answering (IQA) is closely related to the Embodied version of it. The only main issue is that question is designed in a way that the agent must interact with the environment to find the answer. For example, it has to open the refrigerator, or pick up something from the cabinet and then and plan for a series of actions conditioned on the question [56].

3.6 Multi-Agent Systems

Multi-Agent Systems (MAS) is another interesting line of development. The default standpoint of AI has a strong focus on individual agents. MAS research which has its origins in the field of biology tries to change this and studies the emergence of behaviors in groups of agents or swarms instead [57, 58].

Every agent has a set of abilities and is good in them to an extent. The point of interest in MAS is how a sophisticated global behavior can emerge from a population of agents working together. A real-life example of such behavior can be found in insects like ants and bees [59, 60]. One of the interesting goals of this research is to ultimately make agents that could self-repair [61, 62].

The emerging behavior of MAS can be tailored by researchers to let the group of agents tackle various tasks such as rescue missions, traffic control, fun sports events, surveillance, and much more. Additionally, when fused with other fields unexpected outcomes can occur. Take for instance “Talking Heads” experiment by Luc Steels [63, 64] that showed a common vocabulary emerges through the interaction of agents with each other and their environment via a language game.

4 Simulators

Now that we know about the fields and tasks that Embodied AI can shine in, the question is how our agents should be trained. One may say it’s good to directly train in the physical world and expose them to its richness. Although a valid solution, this choice comes with a few drawbacks. First, The training process in the real-world is slow, and the process cannot be sped up or parallelized. Second, it is very hard to control the environment and create custom scenarios. Third, it’s expensive, both in terms of power and time. Fourth, it’s not safe, and improperly trained or not fully trained robots can hurt themselves, humans, animals, and other assets. Fifth, in order for the agent to generalize the training, has to be done in plenty of different environments that is not feasible in this case.

Our next choice is simulators, which can successfully deal with all the aforementioned problems pretty well. In the shift from Internet AI to Embodied AI, simulators take the role that was previously played by traditional datasets. Additionally, one more advantage of using simulators is that the physics in the environment can be tweaked as well. For instance, some traditional approaches in this field [65] are sensitive to noise and for the remedy, the noise in the sensors can be turned off for the purpose of this task.

As a result, agents nowadays are often developed and benchmarked in simulators [66, 67] and once a promising model has been trained and tested, it can then be transferred to the physical world [68, 69].

House3D [70], AI2-THOR [71], Gibson [72], CHALET [73], MINOS [74] and Habitat [75] are some of the popular simulators for the Embodied AI studies. These platforms vary with respect to the 3D environments they

use, the tasks they can handle, and the evaluation protocols they provide. These simulators support different sensors such as vision, depth, touch, and semantic segmentation.

In this paper we mainly focus on MINOS and Habitat since they provide more customization abilities (number of sensors, their positions, and their parameters) and are implemented in a loosely coupled manner to generalize well to new multi-sensory tasks and environments. As their API can be used to define any high-level task and the material, object clutter variation, and much more can be programmatically configured for the environment. They both support navigation with both continuous and discrete state spaces. Also, for the purpose of their benchmarks, all the actuators are noiseless, but they both have the ability to enable noises if desired [76].

In the last section, we saw numerous task definitions and how they each can be tackled by the agents. So, before jumping into MINOS and Habitat simulators and reviewing them, let’s first get more familiarized with the three main goal-directed navigation tasks, namely, PointGoal Navigation, ObjectGoal Navigation, and RoomGoal Navigation.

In PointGoal Navigation, an agent is appeared at a random starting position and orientation in a 3D environment and is asked to navigate to target coordinates which are given relative to the agent’s position. The agent can access its position via an indoor GPS. There exists no ground-truth map and the agent must only use its sensors to do the task. The scenarios start the same for ObjectGoal Navigation, and RoomGoal Navigation as well, however, instead of coordinates, the agent is asked to find an object or go to a specific room.

4.1 MINOS

Minos simulator provides access to 45,000 three-dimensional models of furnished houses with more than 750K rooms of different types available in the SUNCG [77] dataset and 90 multi-floor residences with approximately 2,000 annotated room regions that are in the Matterport3D [78] dataset by default. Environments in Matterport3D are more realistic looking than the ones in SUNCG. MINOS simulator can approximately reach hundreds of frames per second on a normal workstation.

In order to benchmark the system, the authors studied four navigation algorithms; three of which were based on asynchronous advantage actor-critic (A3C) approach [79] and the remaining one was Direct Future Prediction (DFP) [80].

The most basic one among the algorithms was Feedforward A3C. In this algorithm, a feedforward CNN model is employed as the function approximator to learn the policy along with the total value function that is the expected sum of rewards from the current timestamp until the end of the episode. The second one was LSTM A3C that used an LSTM model with the Feedforward A3C act as a simple memory. Next was UNREAL, an LSTM A3C model boosted with auxiliary tasks such as value function replay and reward prediction. Last but not the least, the DFP algorithm was employed that can be considered as Monte Carlo RL [81] with a decomposed reward.

The authors benchmarked these algorithms on PointGoal and RoomGoal tasks and found out that firstly, the naive feedforward algorithm fails to learn any useful representation; secondly, in small environments, DFP performs better while in big and more complex environments UNREAL beat the others.

4.2 Habitat

Habitat was designed and built in a way to provide the maximum customizability in terms of the datasets that can be used and how the agents and the environment can be configured. That being said, Habitat works with all the major 3D environment datasets without a problem. Moreover, it’s extremely fast in comparison to other simulators. AI2-THOR and CHALET can get to an fps of roughly ten, MINOS and Gibson can get to around a hundred, and House3D yields 300 fps in the best case, while Habitat is capable of getting up to 10,000 frames per second. Habitat also provides a more realistic collision model in which if a collision happens, the agent can be moved partially or not at all in the intended direction.

To benchmark Habitat, the owners employed a few naive algorithm baselines, Proximal Policy Optimization (PPO) [82] as the representer of learning algorithms versus ORB-SLAM2 [83, 84] as the chosen candidate for non-learning agents and tested them on the PointGoal Navigation task on Gibson and Matterport3D. They used Success weighted by Path Length (SPL) [85] as the metric for their performance. The PPO agent was tested with different levels of sensors (e.g. No visual sensor, only depth, only RGB, and RGBD) to perform an ablation study and find the proportion in which each sensor helps the progress. SLAM agents were given RGBD sensors in all the episodes.

The authors found out that first, PPO agents with only RGB perform as bad as agents with no visual sensors. Second, all agents perform better and generalize more on Gibson rather than Matterport3D since the size of environments in the latter is bigger. Third, agents with only depth sensors generalize across datasets the best and can achieve the highest SPL. But most importantly, they realized that unlike what has been mentioned in the previous work, if the PPO agent learns long enough, it will eventually outperform the traditional SLAM pipeline. This finding was only possible because the Habitat simulator was fast enough to train PPO agents for 75million time steps as opposed to only 5million time steps in the previous investigations.

5 Future of Embodied AI

5.1 Higher Intelligence

Consciousness has always been considered as the ultimate characteristic for true intelligence. Qualia [86, 87] is the philosophical view of consciousness and it is related to the subjective sensory qualities like "the redness of red" that humans have in their mind. If at some point machines can understand this concept and objectively measure such things, then the ultimate goal can be marked as accomplished.

Robots still struggle at performing a wide spectrum of tasks effortlessly and smoothly, and this mainly due to actuator technology as currently mostly electrical motors are used. Advances in artificial muscles and skin sensors that could cover the entire embodiment of the agent would be essential to fully mitigate the human experience in the real world and eventually unlock the desired cognition [88].

5.2 Evolution

One more key component for cognition is the ability to grow and evolve over time [89, 90, 91]. It's easy to evolve the agent's controller via an evolutionary algorithm but it's not enough. If we aim to have completely different agents, we might as well give them the ability to evolve in terms of embodiment and the sensors as well. This again requires the above mentioned artificial cell organism to encode different physical attributes in them and flip them slightly over time. Of course, we are far from this to become reality, but it is always good to know the furthestmost step that has to be done one day.

6 Conclusion

Embodied AI is the field of study that takes us one step closer to the true intelligence. It is a shift from Internet AI towards embodiment intelligence that tries to exploit the multi-sensory abilities of agents such as vision, hearing, touch, and together with language understanding and reinforcement learning attempts to interact with real-world in a more sensible way. In this paper, we tried to do a concise review of this field, and its current advancements, subfields, and tools hoping that this would help and accelerate future researches in this area.

References

- [1] Jong Hyuk Park, Muhammad Younas, Hamid R Arabnia, and Naveen Chilamkurti. Emerging ict applications and services—big data, iot, and cloud computing, 2021.
- [2] M Hadi Amini, Ahmed Imteaj, and Panos M Pardalos. Interdependent networks: A data science perspective. *Patterns*, page 100003, 2020.
- [3] Farid Ghareh Mohammadi and M Hadi Amini. Promises of meta-learning for device-free human sensing: learn to sense. In *Proceedings of the 1st ACM International Workshop on Device-Free Human Sensing*, pages 44–47, 2019.
- [4] M Hadi Amini, Javad Mohammadi, and Soumya Kar. Promises of fully distributed optimization for iot-based smart city infrastructures. In *Optimization, Learning, and Control for Interdependent Complex Networks*, pages 15–35. Springer, 2020.
- [5] M Hadi Amini, Hamidreza Arasteh, and Pierluigi Siano. Sustainable smart cities through the lens of complex interdependent infrastructures: panorama and state-of-the-art. In *Sustainable interdependent networks II*, pages 45–68. Springer, 2019.
- [6] Ditsuhi Iskandaryan, Francisco Ramos, and Sergio Trilles. Air quality prediction in smart cities using machine learning technologies based on sensor data: A review. *Applied Sciences*, 10(7):2401, 2020.
- [7] Michael Batty, Kay W Axhausen, Fosca Giannotti, Alexei Pozdnoukhov, Armando Bazzani, Monica Wachowicz, Georgios Ouzounis, and Yuval Portugali. Smart cities of the future. *The European Physical Journal Special Topics*, 214(1):481–518, 2012.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [9] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [10] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010.
- [11] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007.
- [12] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. *Advances in neural information processing systems*, 27:487–495, 2014.
- [13] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [14] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [15] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. Swag: A large-scale adversarial dataset for grounded commonsense inference. *arXiv preprint arXiv:1808.05326*, 2018.

- [16] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- [17] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [18] Farzan Shenavarmasouleh and Hamid R Arabnia. Drdr: Automatic masking of exudates and microaneurysms caused by diabetic retinopathy using mask r-cnn and transfer learning. *arXiv preprint arXiv:2007.02026*, 2020.
- [19] Farzan Shenavarmasouleh, Farid Ghareh Mohammadi, M Hadi Amini, and Hamid R Arabnia. Drdr ii: Detecting the severity level of diabetic retinopathy using mask rcnn and transfer learning. *arXiv preprint arXiv:2011.14733*, 2020.
- [20] Farzan Shenavarmasouleh, Farid Ghareh Mohammadi, M. Hadi Amini, Thiab Taha, Khaled Rasheed, and Hamid R. Arabnia. Drdrv3: Complete lesion detection in fundus images using mask r-cnn, transfer learning, and lstm. *arXiv preprint arXiv:2108.08095*, 2021.
- [21] F. Shenavarmasouleh and H. Arabnia. Causes of misleading statistics and research results irreproducibility: A concise review. In *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 465–470, 2019.
- [22] Richard Held and Alan Hein. Movement-produced stimulation in the development of visually guided behavior. *Journal of comparative and physiological psychology*, 56(5):872, 1963.
- [23] Hans Moravec. Locomotion, vision and intelligence. 1984.
- [24] Matej Hoffmann and Rolf Pfeifer. The implications of embodiment for behavior and cognition: animal and robotic case studies. *arXiv preprint arXiv:1202.0440*, 2012.
- [25] Rodney A Brooks. New approaches to robotics. *Science*, 253(5025):1227–1232, 1991.
- [26] Steven H Collins, Martijn Wisse, and Andy Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research*, 20(7):607–615, 2001.
- [27] Fumiya Iida and Rolf Pfeifer. Cheap rapid locomotion of a quadruped robot: Self-stabilization of bounding gait. In *Intelligent autonomous systems*, volume 8, pages 642–649. IOS Press Amsterdam, The Netherlands, 2004.
- [28] Tomoyuki Yamamoto and Yasuo Kuniyoshi. Harnessing the robot’s body dynamics: a global dynamics approach. In *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, volume 1, pages 518–525. IEEE, 2001.
- [29] Gerardo Bledt, Matthew J Powell, Benjamin Katz, Jared Di Carlo, Patrick M Wensing, and Sangbae Kim. Mit cheetah 3: Design and control of a robust, dynamic quadruped robot. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2245–2252. IEEE, 2018.
- [30] Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojciech Marian Czarnecki, Max Jaderberg, Denis Teplyashin, et al. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*, 2017.

- [31] Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline, 2019.
- [32] Yingwei Pan, Ting Yao, Houqiang Li, and Tao Mei. Video captioning with transferred semantic attributes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6504–6512, 2017.
- [33] Soheyla Amirian, Khaled Rasheed, Thiab R. Taha, and Hamid R. Arabnia. Automatic image and video caption generation with deep learning: A concise review and algorithmic overlap. *IEEE Access*, 8:218386–218400, 2020.
- [34] Soheyla Amirian, Khaled Rasheed, Thiab R. Taha, and Hamid R. Arabnia. Automatic generation of descriptive titles for video clips using deep learning. In *Springer Nature - Research Book Series: Transactions on Computational Science & Computational Intelligence*, page Springer ID: 89066307, 2020.
- [35] Lianli Gao, Zhao Guo, Hanwang Zhang, Xing Xu, and Heng Tao Shen. Video captioning with attention-based lstm and semantic consistency. *IEEE Transactions on Multimedia*, 19(9):2045–2055, 2017.
- [36] Yang Yang, Jie Zhou, Jiangbo Ai, Yi Bin, Alan Hanjalic, Heng Tao Shen, and Yanli Ji. Video captioning by adversarial lstm. *IEEE Transactions on Image Processing*, 27(11):5600–5611, 2018.
- [37] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8317–8326, 2019.
- [38] Dinesh Jayaraman and Kristen Grauman. Learning image representations tied to egomotion from unlabeled video. *International Journal of Computer Vision*, 125(1-3):136–161, 2017.
- [39] Dinesh Jayaraman, Ruohan Gao, and Kristen Grauman. Shapecodes: self-supervised feature learning by lifting views to viewgrids. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 120–136, 2018.
- [40] Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–53, 2018.
- [41] Sanjeel Parekh, Slim Essid, Alexey Ozerov, Ngoc QK Duong, Patrick Pérez, and Gaël Richard. Guiding audio source separation by video object information. In *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pages 61–65. IEEE, 2017.
- [42] Jie Pu, Yannis Panagakis, Stavros Petridis, and Maja Pantic. Audio-visual object localization and separation using low-rank and sparsity. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2901–2905. IEEE, 2017.
- [43] Sanjeel Parekh, Slim Essid, Alexey Ozerov, Ngoc QK Duong, Patrick Pérez, and Gaël Richard. Motion informed audio source separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6–10. IEEE, 2017.
- [44] Ehsan Asali, Farzan Shenavarmasouleh, F. Mohammadi, P. Suresh, and H. Arabnia. Deepmsrf: A novel deep multimodal speaker recognition framework with feature selection. *ArXiv*, abs/2007.06809, 2020.
- [45] John Aloimonos, Isaac Weiss, and Amit Bandyopadhyay. Active vision. *International journal of computer vision*, 1(4):333–356, 1988.
- [46] Dana H Ballard. Animate vision. *Artificial intelligence*, 48(1):57–86, 1991.

- [47] Dana H Ballard and Christopher M Brown. Principles of animate vision. *CVGIP: Image Understanding*, 56(1):3–21, 1992.
- [48] Ruzena Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.
- [49] Sumantra Dutta Roy, Santanu Chaudhury, and Subhashis Banerjee. Active recognition through next view planning: a survey. *Pattern Recognition*, 37(3):429–446, 2004.
- [50] Hsiao-Yu Fish Tung, Ricson Cheng, and Katerina Fragkiadaki. Learning spatial common sense with geometry-aware recurrent networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2595–2603, 2019.
- [51] Dinesh Jayaraman and Kristen Grauman. End-to-end policy learning for active visual categorization. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1601–1614, 2018.
- [52] Jianwei Yang, Zhile Ren, Mingze Xu, Xinlei Chen, David Crandall, Devi Parikh, and Dhruv Batra. Embodied visual recognition, 2019.
- [53] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2054–2063, 2018.
- [54] Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, and Dhruv Batra. Embodied question answering in photorealistic environments with point cloud perception. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6659–6668, 2019.
- [55] Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Neural modular control for embodied question answering. *arXiv preprint arXiv:1810.11181*, 2018.
- [56] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4089–4098, 2018.
- [57] Michael G Hinchey, Roy Sterritt, and Chris Rouff. Swarms and swarm intelligence. *Computer*, 40(4):111–113, 2007.
- [58] Junjue Wang, Ziqiang Feng, Zhuo Chen, Shilpa George, Mihir Bala, Padmanabhan Pillai, Shao-Wen Yang, and Mahadev Satyanarayanan. Bandwidth-efficient live video analytics for drones via edge computing. In *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, pages 159–173. IEEE, 2018.
- [59] Scott Camazine, Peter K Visscher, Jennifer Finley, and Richard S Vetter. House-hunting by honey bee swarms: collective decisions and individual behaviors. *Insectes Sociaux*, 46(4):348–360, 1999.
- [60] Christopher G Langton. Artificial life: An overview cambridge. *Mass. MIT*, 1995.
- [61] Fumio Hara and Rolf Pfeifer. *Morpho-functional machines: The new species: Designing embodied intelligence*. Springer Science & Business Media, 2003.
- [62] Satoshi Murata, Akiya Kamimura, Haruhisa Kurokawa, Eiichi Yoshida, Kohji Tomita, and Shigeru Kokaji. Self-reconfigurable robots: Platforms for emerging functionality. In *Embodied Artificial Intelligence*, pages 312–330. Springer, 2004.
- [63] Luc Steels. Language games for autonomous robots. *IEEE Intelligent systems*, 16(5):16–22, 2001.

- [64] Luc Steels. Evolving grounded communication for robots. *Trends in cognitive sciences*, 7(7):308–312, 2003.
- [65] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.
- [66] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017.
- [67] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3357–3364. IEEE, 2017.
- [68] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.
- [69] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.
- [70] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209*, 2018.
- [71] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.
- [72] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9068–9079, 2018.
- [73] Claudia Yan, Dipendra Misra, Andrew Bennet, Aaron Walsman, Yonatan Bisk, and Yoav Artzi. Chalet: Cornell house agent learning environment. *arXiv preprint arXiv:1801.07357*, 2018.
- [74] Manolis Savva, Angel X Chang, Alexey Dosovitskiy, Thomas Funkhouser, and Vladlen Koltun. Minos: Multimodal indoor simulator for navigation in complex environments. *arXiv preprint arXiv:1712.03931*, 2017.
- [75] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9339–9347, 2019.
- [76] Samyak Datta, Oleksandr Maksymets, Judy Hoffman, Stefan Lee, Dhruv Batra, and Devi Parikh. Integrating egocentric localization for more realistic point-goal navigation agents. *arXiv preprint arXiv:2009.03231*, 2020.
- [77] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017.
- [78] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017.

- [79] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- [80] Alexey Dosovitskiy and Vladlen Koltun. Learning to act by predicting the future. *arXiv preprint arXiv:1611.01779*, 2016.
- [81] Satinder P Singh and Richard S Sutton. Reinforcement learning with replacing eligibility traces. *Machine learning*, 22(1-3):123–158, 1996.
- [82] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [83] Raul Mur-Artal and Juan D. Tardos. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, Oct 2017.
- [84] Dmytro Mishkin, Alexey Dosovitskiy, and Vladlen Koltun. Benchmarking classic and learned navigation in complex 3d environments. *arXiv preprint arXiv:1901.10915*, 2019.
- [85] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [86] Frank Jackson. Epiphenomenal qualia. *The Philosophical Quarterly (1950-)*, 32(127):127–136, 1982.
- [87] Michael Tye. Qualia. 1997.
- [88] Koh Hosoda. Robot finger design for developmental tactile interaction. In *Embodied Artificial Intelligence*, pages 219–230. Springer, 2004.
- [89] Dario Floreano, Francesco Mondada, Andres Perez-Urbe, and Daniel Roggen. Evolution of embodied intelligence. In *Embodied artificial intelligence*, pages 293–311. Springer, 2004.
- [90] Dario Floreano, Phil Husbands, and Stefano Nolfi. Evolutionary robotics. Technical report, Springer Verlag, 2008.
- [91] Rolf Pfeifer and Fumiya Iida. Embodied artificial intelligence: Trends and challenges. In *Embodied artificial intelligence*, pages 1–26. Springer, 2004.