

Core Challenges in Embodied Vision-Language Planning

Jonathan Francis

*School of Computer Science, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, USA*

JMF1@CS.CMU.EDU

Nariaki Kitamura

*Komatsu Ltd.,
2-3-6 Akasaka, Minato-ku, Tokyo, Japan*

NARIAKI.KITAMURA@GLOBAL.KOMATSU

Felix Labelle

Xiaopeng Lu

Ingrid Navarro

Jean Oh

*School of Computer Science, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA, USA*

FLABELLE@ALUMNI.CMU.EDU

XIAOPEN2@ALUMNI.CMU.EDU

INGRIDN@CS.CMU.EDU

JEANO@CMU.EDU

Abstract

Recent advances in the areas of multimodal machine learning and artificial intelligence (AI) have led to the development of challenging tasks at the intersection of Computer Vision, Natural Language Processing, and Embodied AI. Whereas many approaches and previous survey pursuits have characterised one or two of these dimensions, there has not been a holistic analysis at the center of all three. Moreover, even when combinations of these topics are considered, more focus is placed on describing, e.g., current architectural methods, as opposed to *also* illustrating high-level challenges and opportunities for the field. In this survey paper, we discuss Embodied Vision-Language Planning (EVLP) tasks, a family of prominent embodied navigation and manipulation problems that jointly use computer vision and natural language. We propose a taxonomy to unify these tasks and provide an in-depth analysis and comparison of the new and current algorithmic approaches, metrics, simulated environments, as well as the datasets used for EVLP tasks. Finally, we present the core challenges that we believe new EVLP works should seek to address, and we advocate for task construction that enables model generalizability and furthers real-world deployment.

1. Introduction

With recent progress in the fields of artificial intelligence (AI) and robotics, intelligent agents are envisaged to cooperate or collaborate with humans in shared environments. Such agents are expected to understand semantic contexts of an environment, e.g., using visual information perceived using sensors, as well as auditory or textual information intended for humans, e.g., presented in natural language. With the goal of developing intelligent agents disposing of these capabilities, embodied AI is a field that studies AI problems situated in a physical environment. Recently, the number of papers and datasets for the tasks that require the agents to use both vision and language understanding has increased markedly (Das et al., 2018a; Gordon et al., 2018; Anderson et al., 2018c; Krantz et al., 2020; Thomason et al., 2019b; Nguyen and Daumé III, 2019; Majumdar et al., 2020; Li et al., 2020b). In this article, we conduct a survey of recent works on these types of problems, which we refer to as *Embodied Vision-Language Planning* (EVLP) tasks. In this article, we aim to provide a

bird’s-eye view of current research on EVLP problems, addressing their main challenges and future directions. **The main contributions of this article are the following:**

1. We propose a taxonomy that unifies a set of related subtasks of EVLP
2. We survey the current state-of-the-art techniques used in the EVLP family to provide a roadmap for the researchers in the field
3. More importantly, we identify and discuss the remaining challenges with an emphasis towards solving real-world problems.

1.1 Scope of this Survey

An “agent” in this article refers to an entity that can make decisions and take actions autonomously (Wooldridge and Jennings, 1995; Castelfranchi, 1998). An embodied agent is situated in a physical or virtual environment which the agent navigates in and interacts with. An embodied agent is generally equipped with sensing capabilities, e.g., via visual or auditory sensing modalities. In this article, we focus on existing works that use visual perception and language understanding as the two major inputs for an embodied agent to make decisions for various tasks in its environment. We note that the topics covered here are mainly for a single agent in a static environment, i.e., topics on multi-agent planning are outside the scope of this article. Moreover, as the majority of tasks discussed in this article are conducted through simulated environments, real-world physical challenges in robotics (e.g., visual affordance learning, proprioceptive control, system identification) are not considered. EVLP problems have previously been studied in the fields of natural language processing, robotics, and computer vision (Oh et al., 2015; Boularias et al., 2015; Duvall et al., 2016). While the focus of this survey is on contemporary works, we will discuss how concepts from classical approaches have inspired recent methodology and how they could be used for future directions, e.g., as in the use of mapping and exploration strategies, search and topological planning, and hierarchical task decomposition (Section 3.1.3).

1.2 Intended Audience and Reading Guide

This paper is tailored to accommodate a broad spectrum of reader backgrounds and perspectives, as this paper is positioned at the intersection of Computer Vision (CV), Natural Language Processing (NLP), and Robotics. For readers that are new to these topics, we provide an in-depth coverage of existing works and methodologies; for readers with significant experience in one or more of these areas, we offer a taxonomy of the broader field of *Embodied Vision-Language Planning* in Section 2.1, provide an analysis of core challenges, and discuss future directions in Section 5.

For readers who are less familiar with implementation and evaluation of embodied agents, it is recommended to read the sections in order, as each section builds upon previous sections. The rest of the article is organized as follows. In Section 2, we formally define the class of EVLP problems discussed in this article and present a taxonomy of the field. After describing the key tasks that compose the EVLP family in Section 2.2, we present modeling approaches including commonly-used learning paradigms, architectural design choices, and techniques proposed to tackle EVLP-specific challenges. Section 4 presents the datasets and evaluation metrics currently used by the research community. Finally, Section 5.1 discusses several open challenges in the field.

1.3 Related Surveys

We review existing survey articles on relevant topics, in order to provide readers with pointers to other papers on more specific topics and to clarify how this article differs from them.

- Ye and Yang (2020) discuss the datasets, methodologies, and common approaches to Vision Indoor Navigation (VIN) tasks. VIN tasks do not necessarily require language, are limited to navigation, and only occur in indoor environments. Our survey touches on planning tasks, of which navigation is a subset, which use language to establish the goal state. Also, EVLP tasks are not limited to indoor environments, but also include outdoor settings.
- Luketina et al. (2019) focus on existing approaches of using language components in the context of reinforcement learning (RL). The authors divide the language-related problem into language-conditional RL and language-assisted RL. In contrast, our work briefly discusses different training paradigms commonly used for tasks in the EVLP domain, in Section 3.
- Mogadala et al. (2019) cover popular tasks and approaches in multimodal research. Our survey, in comparison, specifically focuses on EVLP tasks, and provides more details, such as commonly used datasets, training paradigms, and challenges on this specific domain.
- Bisk et al. (2020) explore the current research in language understanding and its shortcomings. To categorize the differences between human and machine understanding, the authors define the notion of a “world scope.” A worldscope defines the type of real world information that language models capture, such as physical relationships or social concepts. Five worldscopes are used to define a framework that discusses what aspects of language are harder to grasp and issues with the current approaches. We do not set out a general framework to identify the complexity of language. Instead, we survey the low-level research progress of EVLP tasks, discuss the role of language in planning tasks, and further point out possible improvements.
- Tangiuchi et al. (2019) discuss techniques at the intersection of language and robotics research. Their work focuses heavily on the syntactic and semantic structures of language information, and how it can be tied to the low level robotics actions. Instead, our survey focuses primarily on the task level, as opposed to low-level actions, and is suitable for readers with no prior knowledge in the robotics field.
- Uppal et al. (2020) provide discussion of Vision and Language problems, dividing Vision-language tasks into four categories: generation, classification, retrieval, and other tasks. Here, the authors only discuss one EVLP task; however, due to the increased popularity of embodied vision-language planning, we feel that the EVLP family warrants its own specific treatment and careful discussion. Moreover, Uppal et al. (2020) further discuss changes in representation such as the use of transformers, fusion of multiple modalities, architectures, and evaluation metrics. While both surveys discuss changes in multimodal representation, ours focuses on EVLP research, providing a unified problem definition, taxonomy, and analysis of trends for this exciting field.
- Batra et al. (2020) provide a detailed analysis and framework for the rearrangement task, which involves using an embodied agent to change the environment from an initial state to a target state. The authors argue that the manipulation problem can be viewed as two-fold:

agent-centric problem and environment-centric problem. They further analyze the more fine-grained split for each problem. Several evaluation metrics and existing testbeds are presented. While they do introduce a framework which includes one EVLP task, it is not the principle focus of that paper. We pursue a broader discussion of emerging trends and core challenges in the field.

2. Problem Definition

In this article, we discuss a broad set of problems, related to an embodied agent’s ability to make planning decisions in physical environments. We include both stateless problems such as question answering and sequential decision-making problems, such as navigation, under the umbrella of “planning”. Note that stateless problems belong to the special case of single-step decision making. More generally, we refer to sequential decision-making problems as those where an agent must perform actions over a countable time-horizon, given an objective. In the context of navigation, for example, an agent may choose actions, according to a pre-defined action space, that enable an agent to transition from an initial state to a goal state. In the context of manipulation, the agent executes a series of actions, in order to effect a desired change in the environment. More formally, planning problem are defined as follows:

Definition 1 (Planning) *Let S denote a set of states; A , a set of actions. A planning problem is defined as a tuple of its states, actions, initial and goal states: $\Phi = \{S, A, s_{ini}, s_{goal}\}$, where $s_{ini}, s_{goal} \in S$ denote initial and goal states, respectively. A solution $\psi \in \Psi_{\Phi}$ to planning problem Φ is a sequence of actions to take in each state, starting from an initial state to reach a goal state, $\psi = [s_{ini}, a_1, \dots, a_T, s_{goal}]$, where T is a finite time-step and Ψ_{Φ} is a set of possible solutions to Φ .*

EVLP problems require planning in partially-observable environments, i.e., the entire state space may not be known to the agent in advance. Instead, an agent needs to use vision and language inputs to estimate its current and goal states in order to accomplish high-level tasks in a physical environment. These inputs can be given at the start of a task, e.g., a task itself is described in natural language in the form of a question or instruction, or become available as an agent moves through the state space. Visual inputs are an example of such an input as they are primarily online information that can be perceived through sensing. Note this is not always the case, as they can also be provided as part of a task specification, e.g., a view from a goal location.

Definition 2 (Embodied Vision-Language Planning (EVLP)) *Let V and L denote sets of vision and language inputs available to an agent. Given an EVLP problem Φ , state s_t at time step t can be defined in terms of vision and language inputs up to the current time step, such that, $s_t = f(v_1, \dots, v_t, l_1, \dots, l_t)$ where $v_t \in V$ and $l_t \in L$. The objective here is to minimize the difference between an admissible¹ solution $\psi \in \Psi_{\Phi}$ and a predicted one $\bar{\psi}$.*

This definition broadly captures the crux of EVLP problems. A customized definition would be needed for each specific task where additional constraints or assumptions are added to focus on particular subareas of this general problem. For instance, Vision-Language Navigation (VLN) is a natural language direction following problem in an unknown environment, which can be defined

1. Admissibility condition: a solution that is possible, under the constraints of the environment’s transition dynamics as well as the agents’ own system dynamics, which satisfies the task specification.

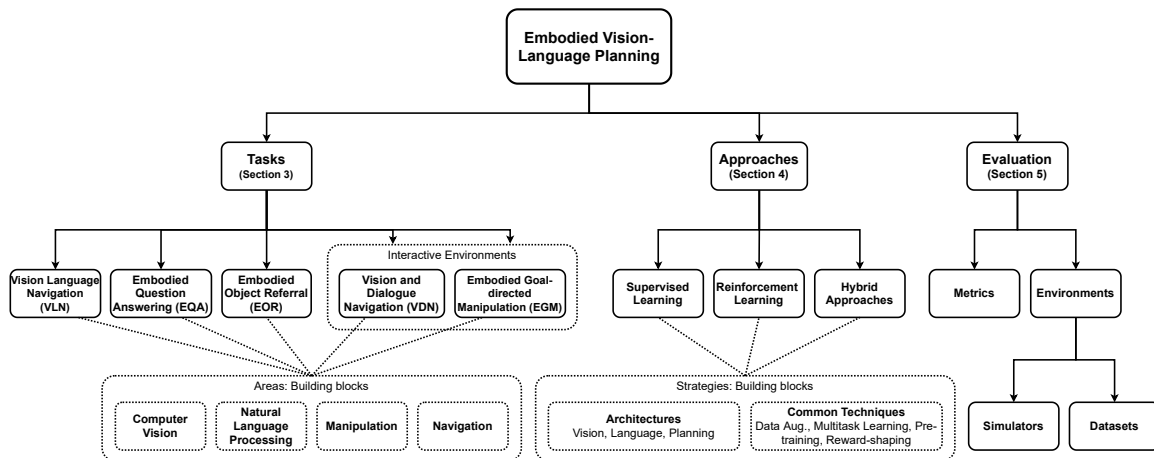


Figure 1: Taxonomy of *Embodied Vision-Language Planning*, aligned with our paper organisation.

as a planning problem where an agent is given an initial state and a solution (or a sequence of actions) represented in natural language, and is equipped with visual perception, e.g., first-person view images.

2.1 Taxonomy

We propose a taxonomy of current EVLP research, illustrated in Figure 1, around which the rest of the paper is organized. The taxonomy subdivides the field into three branches; tasks, approaches, and evaluation methods. The Tasks branch (corresponding to Section 2) proposes a framework to classify existing tasks and to serve as a basis for distinguishing new ones.

The Approaches branch (Section 3) touches on the learning paradigms, common architectures used for the different tasks, as well as common tricks used to improve performance. The technical challenges that underlie all tasks are discussed in Section 3.3.

The right-most branch of the taxonomy, in Figure 1, discusses task Evaluation Methodology (Section 4), which is subdivided into two parts: metrics and environments. The metrics subsection references many of the common metrics and their formulæ, used throughout EVLP tasks, while the environments subsection presents the different simulators and datasets currently used.

2.2 Tasks in Embodied Vision-Language Planning

Many EVLP tasks have been proposed, with each task focusing on different technical challenges and reasoning requirements for agents. Tasks vary on the basis of the action space (types and number of actions possible), the reasoning modes required (e.g., instruction-following, versus exploration and information-gathering), and whether or not the task requires dialogue with another agent. In this survey, we only include tasks where clearly-defined datasets and challenges exist for evaluation benchmarking, i.e., Vision Language Navigation (VLN) (Anderson et al., 2018c; Misra et al., 2018; Hermann et al., 2020; Jain et al., 2019; Ku et al., 2020), Vision and Dialogue History Navigation (VDN) (de Vries et al., 2018; Nguyen and Daumé III, 2019; Thomason et al., 2019b), Embodied Question Answering (EQA) (Das et al., 2018a; Wijmans et al., 2019a), Embodied Object

Referral (EOR) (Qi et al., 2020; Chen et al., 2019b), and Embodied Goal-directed Manipulation (EGM) (Shridhar et al., 2020; Kim et al., 2020; Suhr et al., 2019). In the following sub-sections, we compare and contrast these task families, with the hope that readers gain a solid understanding of the differences in the problem formulations. In this manner, new datasets can be contextualized by existing tasks, and new task definitions can be later positioned alongside those mentioned here. We summarize the attributes for existing tasks in Table 1. We further refer the readers to Section 4.2 for more information about the datasets and simulation environments pertaining each task.

Table 1: Properties of Embodied Vision-Language Planning (EVLP) tasks: ‘Navigation’ describes whether navigation is part of the actions for the task; ‘Object Identification’ indicates whether the agent may only solve the task through identifying specified objects; ‘Environment Interaction’ indicates whether or not the agent can mutate environment state, e.g., through dialogue-based interaction with another agent or through object manipulation; and ‘Primary Reasoning Mode(s)’ describes how agents are intended to interpret the task.

Task	Navigation	Object Identification	Environment Interaction	Primary Reasoning Mode(s)
VLN	✓	✗	✗	Understanding object and scene layout properties; instruction-following
EQA	✓	✓	✗	Exploration and information-gathering
EOR	✓	✓	✗	Understanding spatio-semantic object relations and scene layout properties; instruction-following
VDN	✓	✗	✓	Understanding object and scene layout properties; multiple instruction-following
EGM	✓	✓	✓	Understanding object affordances, environment attributes, and scene layout properties; instruction-following

2.2.1 VISION LANGUAGE NAVIGATION

Vision-Language Navigation (VLN) requires an agent to navigate to a goal location in an environment following an instruction L . A problem in VLN can be formulated as $\Phi_{VLN} = \{S, A, s_1, s_{goal}\}$, where a solution or a path $\psi_{\Phi_{VLN}} = \{s_1, a_1, \dots, a_T, s_{goal}\}$ exists, such that each state $s_t \in S, t \in [1, T]$ is associated with a physical location in the environment leading to the target location. The action space A available to the agent consists of navigation actions between physical states, and a `stop` action which determines the end of a solution. Navigation actions can be discrete, e.g., `turn_left`, `turn_right` or `move_forward` (Anderson et al., 2018c; Krantz et al., 2020), as well as continuous (Irshad et al., 2021). Finally, the goal of the agent is to predict a solution $\tilde{\psi}_{\Phi_{VLN}}$ consisting of a sequence of actions $a_t \in A, t \in [1, T]$ that closest align to the instruction, and thus, to the true solution $\psi_{\Phi_{VLN}}$.

VLN is the most established EVLP task: a number of datasets (See Section 4.2) exist in both, indoor (Anderson et al., 2018c; Jain et al., 2019; Krantz et al., 2020; Irshad et al., 2021), and outdoor environments (Hermann et al., 2020; Misra et al., 2018). Overall, VLN models have seen considerable progress in improving the ability to get closer to the goal and to the ground truth trajectory (Fried et al., 2018b; Tan et al., 2019; Li et al., 2020b; Majumdar et al., 2020; Jain et al., 2019). Nonetheless, Zhu et al. (2021b) show that it is unclear if models are actually aligning the visual modality and that recent work has experienced a slow-down in performance improvements. They suggest that un-

derstanding how VLN agents interpret visual and textual inputs when making navigation decisions is an open challenge.

2.2.2 EMBODIED QUESTION ANSWERING

In Embodied Question-Answering (EQA), an agent initially receives a language-based question L , and must navigate around collecting information about its surroundings to generate an answer. An EQA problem can be formulated as $\Phi_{EQA} = \{S, A, R, s_1, r\}$, where R is the set of possible answers, and $r \in R$ represents the correct answer to the given question. The actions A available to the agent consist of discrete moving actions as those in Subsection 2.2.1. However, as opposed to using a stop action, the ending of an episode is given by an answer action (Das et al., 2018a; Gordon et al., 2018). Lastly, a correct solution to a given question can be expressed as $\Psi_{\Phi_{EQA}} = \{s_1, a_1, \dots, a_T, r\}$.

Unlike the aforementioned tasks, there are some challenges unique to EQA. First, there might not necessarily exist a unique or perfect solution for any given question. Second, EQA also requires an agent to understand the implications of the given questions and to translate them into actions that lead to the correct answer. Furthermore, the EQA task requires an agent’s awareness of its environment and commonsense knowledge on how those environments function or are spatially organized, e.g., food is often located in the fridge, or parked cars are often found in a garage. The level of ambiguity, reasoning required, and knowledge required make EQA one of the more challenging EVLP tasks to date.

2.2.3 EMBODIED OBJECT REFERRAL

In Embodied Object Referral (EOR) tasks, an agent navigates to an object o mentioned in a given instruction L , and has to identify (or select) it upon reaching its location. EOR can be framed as $\Phi_{EOR} = \{S, A, O, s_1, s_{goal}, o\}$. Here, O is the set of possible objects in the environment, which are specified by a class label and a bounding-box or a mask, and $o \in O$ is the object of interest. The set of available actions A includes discrete navigation actions as those described in Subsection 2.2.1. An EOR solution can be expressed as $\Psi_{\Phi_{EOR}} = \{s_1, a_1, \dots, a_T, s_{goal}, o\}$, and represents the path leading to the object of interest o . When the final state s_{goal} is reached, a bounding box or a mask is used to indicate where the object o is located in the viewpoint given at that state. Finally, the goal of the agent is to predict a solution $\tilde{\Psi}_{\Phi_{EOR}}$ that follows the provided instruction and correctly selects the referenced object, and thus, closely matches solution $\Psi_{\Phi_{EOR}}$.

Similar to VLN, EOR relies on instruction following. In this setting, instructions can be step-by-step (Chen et al., 2019b; Mehta et al., 2020), or they can be under specified (Qi et al., 2020). The latter requires the agent to reason and understand the context of the instruction, as was discussed in Subsection 2.2.2. In addition, EOR introduces the challenge of identifying an object of interest. This object might be visible from different viewpoints, and like previous tasks, multiple instances of such object can exist in the environment. More importantly, while an object might be explicit (Chen et al., 2019b), an agent might also be required to predict what object it should find from a given instruction (Qi et al., 2020).

2.2.4 VISION AND DIALOGUE NAVIGATION

A language instruction can be ambiguous and might require clarification or assistance (Thomason et al., 2019b; Nguyen et al., 2019a). For example, when navigating through a home, an agent might find multiple instances of an object referenced in an instruction. Moreover, an instruction may not specify the goal in the level of granularity that the agent needs to plan and execute actions; the agent may require clarification on intermediate sub-tasks. Unlike VLN and EQA, VDN allows an agent to interact with another agent (e.g., a human collaborator) to resolve these types of uncertainties. This interactive element has been added in two ways: through the use of an ambiguity resolution module (Chi et al., 2020; Nguyen and Daumé III, 2019), or through dialogue (Thomason et al., 2019b; Nguyen et al., 2019a). In the case of the former, if the agent gets confused, it may ask for help from an oracle who is aware of the agent’s state and its goal. In the latter, the agent is given an initial vague prompt and is required to ask an oracle for clarification.

More formally, Vision and Dialogue Navigation (VDN) requires an agent to navigate to a goal location in an environment, but may do so through sequential directives $L = \{l_1, \dots, l_N\}$. These directives may come in the form of sub-instructions upon an agent’s request, or as dialogue-based interaction. Here, each instruction represents a sub-problem $\Phi_{VDN}^{(i)} = \{S, A, s_1^{(i)}, s_{goal}^{(i)}\}_{i=1}^N$, leading the agent closer to its ultimate goal position. The actions A that the agent can execute consist of navigation actions, as in previous tasks, but may additionally include other forms of interactive actions (e.g., to request help). Akin to VLN, a solution to a VDN problem can be denoted as the set $\Psi_{\Phi_{VDN}}^{(i)} = \{s_1^{(i)}, a_1^{(i)}, \dots, a_T^{(i)}, s_{goal}^{(i)}\}_{i=1}^N$, and the goal of the agent is to find $\tilde{\Psi}_{\Phi_{VDN}}^{(i)}, i \in [1, N]$, a feasible set of solutions that best align with the true solutions.

We note that past works in VDN (Thomason et al., 2019b; Nguyen et al., 2019a) use static, multi-turn question-answer pairs to simulate conversation between the ego-agent and another agent. Still, we accommodate, as part of the VDN family, problem definitions that feature active dialogue context as well.

2.2.5 EMBODIED GOAL-DIRECTED MANIPULATION

Unlike previous tasks, Embodied Goal-directed Manipulation (EGM) requires the manipulation of objects in a scene, posing unique challenges for agents (Shridhar et al., 2020; Padmakumar et al., 2021). EGM may combine these manipulation-based environment interactions with requirements from aforementioned tasks, such as navigation and path-planning, state-tracking, instruction-following, instruction decomposition, and object selection (Kim et al., 2020). Due to shared properties between EGM and, e.g., VLN task definitions, EGM also encompasses the mobile manipulation paradigm from previous literature (Tellex et al., 2011; Khatib, 1999).

EGM may require solving multi-step instructions and, as such, it can be framed as a set of multiple sub-problems $\Phi_{EGM}^{(i)} = \{S, A, O, s_1^{(i)}, s_{goal}^{(i)}\}_{i=1}^N$ where the i_{th} sub-problem is associated to the i_{th} sub-instruction. In this setting, the action space A of the agent includes navigation actions as in the previous tasks, in addition to actions involving interactions with objects. These interactions can be sub-divided into multiple types of interactions, e.g. `pick_up`, `turn_on`, `turn_off`, etc. (Shridhar et al., 2020), or used as a single `interact` action (Kim et al., 2020). Then, $\Psi_{\Phi_{EGM}}^{(i)} = \{s_1^{(i)}, a_1^{(i)}, \dots, a_T^{(i)}, s_{goal}^{(i)}\}_{i=1}^N$ represents a solution to an EGM problem. Here, a state s_t not only includes the visual observation of the agent’s location at time step t , but also the state of objects that

the agent interacted with. Lastly, the goal of the agent is to predict $\bar{\psi}_{\Phi_{EGM}}^{(i)}$, $i \in [1, N]$ that follows the provided instruction(s) and correctly selects and interacts with the object(s) referred in it, and thus, best matches the true solution(s).

In EGM, there may be constraints such as which objects can or can't be picked up. To interact with an object, the agent specifies where o is in its view using a bounding box or mask. Furthermore, agents not only need to interpret instructions and recognize objects they are meant to interact with, but also need to understand the consequences of interacting with their environment. Immutability is an important constraint, in real life you cannot un-slice a tomato. This means certain mistakes lead to “un-winnable” states and that there may be an order of operations for certain tasks. This makes EGM one of the most difficult EVLP task.

3. Approaches

This section provides a review of technical approaches used in EVLP. It discusses how current works tackle the different facets of EVLP tasks, namely vision, language, and planning. It then discusses different learning paradigms and tricks used to improve performance.

3.1 Modeling Vision, Language, and Planning

3.1.1 MODELING VISION

Modeling vision is an important part of EVLP tasks, as it is the principle way in which agents build a representation of their environment. Representations of vision used include the use of explicit features (Duvall et al., 2015; Tellex et al., 2011; Duvall et al., 2016) and neural representations. In EVLP models, neural representations are most common, with Convolutional Neural Networks (CNNs) being used as encoders in most works (Mogadala et al., 2019). CNNs are neural network architectures suitable to process structured data such as images, and thus are broadly used in the computer vision tasks, e.g., image classification (Rawat and Wang, 2017).

In existing works, pre-trained neural networks such as ResNet (He et al., 2016) are commonly used for extracting meaningful image features. One limitation of using such pre-trained networks is that these features may lead to over-fitting. Zhang et al. (2020) propose using the logits of the classification layer—i.e., a higher-level feature—in order to reduce the performance gap between the seen and the unseen environments.

Neural feature approaches perform well, but lack the interpretability of using explicit features. Anderson et al. (2018b) proposed a hybrid approach, extracting salient features using a pretrained object detector (He et al., 2017) and encoding them using a pretrained CNN. Majumdar et al. (2020) use this technique in VLN to represent objects in an image rather than the entire image. This lends well to the use of multimodal transformers, which we will discuss in section 3.1.3.

3.1.2 MODELING LANGUAGE

Goal understanding is another critical piece of the EVLP puzzle. Goal information, often in the form of instructions or a statement, are provided in the form of language. Like vision, hand-selected features (Duvall et al., 2015; Tellex et al., 2011) or neural representations can be used. Modeling language is primarily done using Recurrent Neural Networks (RNNs) or Transformers.

RNNs are a commonly used neural network architecture for processing sequential data such as language. For more information on RNNs we refer readers to Lipton et al. (2015). When used as a language encoder RNNs take in tokenized instructions and generates a vectorized representation of language. RNNs have some short comings, for example issues with long-term dependencies. To better handle long-term dependency, gated RNNs such as Long Short-Term Memory (LSTMs) (Hochreiter and Schmidhuber, 1997) and Gated Recurrent Unit (GRUs) (Cho et al., 2014) are widely used for general NLP tasks as well as EVLP (Das et al., 2018a; Krantz et al., 2020; Fried et al., 2018b; Nguyen and Daumé III, 2019; Thomason et al., 2019b; Tan et al., 2019; Zhang et al., 2020) due to their superiority in handling long-term dependencies when compared to vanilla RNNs. Other approaches to tackle this include the use of attention, which allows models to attend over the entire context or memory modules (Zhu et al., 2020c).

Transformers operate on a different principle than RNNs, using stacked attention to combine information from different inputs Vaswani et al. (2017). Transformer-based architectures have shown to be effective across various communities, such as CV (Khan et al., 2021), NLP (Vaswani et al., 2017; Devlin et al., 2019), robotics (Fang et al., 2019; Dasari and Gupta, 2020), and multimodal tasks (Shin et al., 2021). Earlier work has leveraged transformers as language models to improve performance on downstream tasks (Devlin et al., 2019). Several works in EVLP have used transformers. Li et al. (2019) use a BERT encoder to encode language for VLN task which improves upon past models.

3.1.3 MODELING MULTIMODALITY

Representing language and vision can not be done independently from one another. First, there is overlapping information between the two modalities which could help correct for errors in one representation. Moreover, information from each is required to interact with one other for the overall task, for example grounding landmarks from the text to the environment. In contemporary works, combining modalities is done either through the use of attention or transformers.

Attention was originally proposed to improve the performance of Seq2Seq neural machine translation (Bahdanau et al., 2015; Yang et al., 2020). It primarily serves two purposes in EVLP: to fuse modalities; to align modalities (Baltrušaitis et al., 2019). Attention mechanisms come in many forms, but typically can be framed as a process which takes in two inputs, referred to as the key and query, and produces an output using a weighted sum of the input, known as the value (Vaswani et al., 2017). Attention weights the query using the key, which can originate from the same or different sources. In the multimodal context this allows the model to combine vision, language, and the agent’s current state (Wang et al., 2019b; Tan et al., 2019; Fried et al., 2018b; Das et al., 2018a; Qi et al., 2020). This is often in done with the use of multiple attention mechanisms Qi et al. (2020); Fried et al. (2018b); Tan et al. (2019).

More recently, multimodal transformers have also shown to be effective for vision-language tasks (Tan and Bansal, 2019; Lu et al., 2019). For example, Lu et al. (2020a) learn language representations together with image information over 12 different multimodal tasks, outperforming most of the tasks when trained independently. In VLN two works have used transformers to combine different modalities. Hao et al. (2020) pre-train on instruction, image, action triplets across different tasks and then finetune on a target task. Majumdar et al. (2020) use ViBERT (Lu et al., 2019) and pre-train over the conceptual captions dataset (Sharma et al., 2018) to learn multimodal representations

using a masking objective for both text and visual streams. During the task-specific fine-tuning, they fine-tune on the VLN task and that shows significant improvement.

3.1.4 MODELING ACTION-GENERATION AND PLANNING

At a high level, we measure agents' understanding of a given goal, its scene context, and its understanding of how to *satisfy* goals by its ability to generate appropriate sequences of actions to execute in the environment. In fully-observable (known) environments, where the location of the goal position is known (whether in a global or local/relative coordinate system), we expect the agent to generate a sequence of actions for reaching the goal, with allowance for re-planning at arbitrary frequency, based on newly-acquired information. In partially-observable (partially-known) environments, e.g., where the agent only has access to information about the adjacent admissible states, the agent must generate actions over significantly shorter horizons, such as one or two time-steps. Inspired by classical motion-planning techniques in robotics, several planning approaches have been employed in EVLP tasks, such as: mapping and exploration, search and topological planning, and hierarchical task decomposition.

Mapping and exploration strategies. In the robotics literature, mapping is a general concept which refers to a transformation of the agent's observations into a more abstract state representation, wherein, e.g., planning can be performed more efficiently (Filliat and Meyer, 2003). Multiple choices of map representation are possible (metric maps, as in occupancy grids, or topological graphs), depending on the required level of expressivity for agent pose, obstacle locations, and goal position (Cummins and Newman, 2008).

Early mapping strategies relied simply on exploring unknown environments. Here the common objective is to prioritise visiting unmapped states at the extent of the currently-explored regions (frontier nodes), in order to facilitate quick and efficient coverage of the environment Burgard et al. (2005). Agents must estimate the cost of visiting each frontier node and assess whether to proceed in its current exploration heading, versus back-tracking to other frontier nodes that promise more information at lower cost. Ke et al. (2019) propose adding an exploration module to VLN agents, allowing agents to perform local decision-making while utilising global information to back-track when the agent gets confused (i.e., revisits previously-visited states). Ma et al. (2019a,b) propose the *Regretful Agent* (RA), which adds two differentiable modules to the standard VLN architecture: the first, *progress marker* (PM), estimates the agent's progress towards the goal, while the second module, *regret module* (RM), decides whether the agent should back-track by comparing the agent's current observation with its historical information. This module attends over both image and language inputs, and if the weight on the previous image features is larger it will back-track.

Because the overhead from mapping an unseen (unknown) or partially-observable environment can be significant, classical robotics approaches moved quickly to Simultaneous Localization and Mapping (SLAM) techniques (Thrun et al., 1998; Cummins and Newman, 2008), wherein an agent projects its observations to a map that it maintains, while also tracking its location in the map. This serves as a basis for registering objects to specific locations of the map and for generating more efficient plans. Anderson et al. (2019) use a navigation approach similar to SLAM, which consists of three modules: a mapper, a Bayesian filter, and a policy. The mapper generates a semantic map and the Bayesian filter uses RGB images and depth to find the most likely path at each time step. The agent builds up its understanding of the world and assigns probabilities to different locations.

By mapping, an agent can navigate efficiently by avoiding redundant visitations and eventually plan using global information as opposed to local.

Search and topological planning. A popular planning approach is to utilise search algorithms, such as informed search (Hart et al., 1968; Stentz et al., 1995), commonly employed for traversing a graphical state representation (e.g., a tree or topological map). Greedy search algorithms (Black et al., 1998), wherein the most likely action is taken at each state, is of particular interest for partially-observable environments, where the agent has limited information to use for deciding on the next state transition.

In the context of EVLP problems, it is common to implement greedy search by way of neural sequence-to-sequence (Seq2Seq) modeling—serving as one of the earliest action-generation strategies reported for these tasks (Fried et al., 2018b; Das et al., 2018a; Anderson et al., 2018c; Shridhar et al., 2020). One shortcoming of this approach is that it does not account for the likelihood of an entire sequence, meaning that it can generate sub-optimal output sequences. Beam search (Holtzman et al., 2019) is an alternative, which accounts for this by looking at k possible output sequences. Other search-inspired sampling algorithms, such as top- k sampling or nucleus sampling, have shown additional promise, primarily in natural language processing community (Holtzman et al., 2019).

Other approaches couple mapping with greedy search strategies. Deng et al. (2020) propose the *Evolving Graphical Planner*, which creates a graph of all admissible nodes (states), where each node is represented using learnable embeddings of the agent’s scene context. The agent’s global planning module attempts to generate actions based on the current version of the map: the map, trajectory history, and language components are fed into the planner model, which returns a probability distribution over the set of the next node candidates; the agent then samples the next location and builds up a new graph based on the new information at this node, keeping only the top- K locations at each time step.

Hierarchical task decomposition. Early robot planning approaches considered multiple levels of abstraction, defined for an arbitrary planning task: the highest level of abstraction need only contain the goal specification, a lower level of abstraction would contain a sequence of subgoals needed to satisfy the task and reach the goal, and even lower abstraction levels would include the primitive actions required for satisfying each sub-goal (Sacerdoti, 1974; Nau et al., 2003). An agent is said to perform hierarchical task decomposition if it is capable of reasoning at multiple levels of task abstraction, as a consequence of either formal predicate calculus (Nilsson et al., 1984; Kaelbling and Lozano-Pérez, 2010) and/or simply based on the agent’s architectural specification.

Architecturally, it is common to define at least two modules, e.g., a global planner and a local planner, where the former is responsible for generating intermediate sub-tasks towards a given goal, and the latter is responsible for generating satisfactory low-level actions for each sub-task. These ideas manifest in more recent works as, e.g., hierarchical waypoint prediction and navigation (Chen et al., 2020d; Misra et al., 2018; Blukis et al., 2018b) and hierarchical reinforcement learning (Li et al., 2020a; Nachum et al., 2018). Certain datasets (Misra et al., 2018; Shridhar et al., 2020; Suhr et al., 2019; Irshad et al., 2021) pair sub-goal sequences with a high-level goal, allowing for hierarchical approaches to learn to reason at multiple levels of abstraction: agents can predict

abstracted actions, corresponding to sub-goals (e.g., a waypoint to navigate to), and use a lower-level module to control the movement actions to achieve that sub-goal (e.g., motor movements to transition between waypoints). Misra et al. (2018); Blukis et al. (2018b) propose a two-part modular architecture: the first is a Visitation Distribution Prediction (VDP) module, which predicts waypoints using RGB images and instructions on a probabilistic map; and the second stage receives this map from the VDP module, for learning to predict low-level actions, through reinforcement learning.

Generalising beyond the above task decomposition paradigm, where sub-tasks consist of only navigation or only manipulation actions, early work also considered settings where agents were required to dispatch multiple capabilities (e.g., manipulation, navigation, dialogue) within the same procedure or episode. This integrated setting, manifesting in such applications as manufacturing assistance and surgical robotics, represented a new challenge for planning-based approaches: agents quickly needed to contend with heterogeneous action spaces, longer task horizons, and combinatorial complexity in configuring the space of low-level actions. This setting was formalised as the problem of *task and motion planning* (TAMP) (Cambon et al., 2009; Choi and Amir, 2009; Wolfe et al., 2010; Kaelbling and Lozano-Pérez, 2010, 2013; Garrett et al., 2021), wherein algorithmic solutions define hierarchies across the different environment interaction types, with each type having its own unique execution primitive. These execution primitives were configured by way of a task decomposition, between a task-planning module (symbolic planning) and motion and/or manipulation planning modules (geometric planning). Many such solutions were plagued, still, with high search complexity, as early symbolic planners needed to enumerate all possible operations in a state, in order to expand the search tree. On the other hand, motion/manipulation planners were able to deal more effectively with geometry, but struggled when given partial goal specification. Notably, Kaelbling and Lozano-Pérez (2010) utilised *goal regression*, to recursively decompose the planning problem, through aggressive hierarchicality—limiting the length of plans and, thereby, exponentially decreasing the amount of search required.² With the advent of more sophisticated simulators (e.g., AI2-THOR (Kolve et al., 2017)) and benchmarks (e.g., ALFRED (Shridhar et al., 2020), Interactive Gibson (Li et al., 2021a)), recent approaches extend the original TAMP problem in the context of *mobile manipulation*, *interactive navigation*, and *navigation among movable objects*. Li et al. (2020a) pursue interactive navigation by proposing a hierarchical reinforcement learning framework, for learning cross-task hierarchical planning. Sharma et al. (2021) produced a framework for learning hierarchical policies from demonstration, with the intention of identifying reusable robot skills. Modularity as a general design principle in autonomous systems can be most closely attributed to the principles of cross-task hierarchical planning, as studied in this section, which serves as a strong foundation for improving the agents’ sample-efficiency and the tractability of sometimes-conflicting learning objectives (Das et al., 2018b; Chen et al., 2020d).

3.2 Learning Paradigms

3.2.1 SUPERVISED LEARNING

Supervision in EVLP generally refers to a demonstration of a possible solution to a problem. Learning from demonstration approaches are typically focused on matching the behavior of their demonstrator or *expert* (Hester et al., 2018). As such, these approaches can be suitable when there is

2. Garrett et al. (2021) further formalise a class of *task and motion planning* problems and survey solution methodology.

available data from an expert, or when it is easier to collect expert demonstrations than to specify a reward function to train an agent under reinforcement learning. Furthermore, this type of learning paradigm has shown to help in accelerating the learning process in difficult exploration problems (Syed and Schapire, 2007).

Different approaches within this paradigm have been used in all EVLP tasks. For example, in VLN, Anderson et al. (2018c) use the *Professor-Forcing* (Lamb et al., 2016) approach, where at each training step the expert action is used to condition later predictions. Here, expert demonstrations generally correspond to the shortest-path from a start to a goal location for any given instruction. Furthermore, this approach is coupled with *Student-Forcing*, a method which samples an agent’s action from a output probability distribution in order to avoid limiting exploration to only the states included in the expert trajectories (Anderson et al., 2018c). Another commonly used approach is DAgger (Ross et al., 2011a), which trains on aggregated trajectories obtained using expert demonstrations (Blukis et al., 2018a; Krantz et al., 2020; Irshad et al., 2021).

In EQA, Wijmans et al. (2019a) train a behavior cloning model using a synthetic dataset of demonstration episodes, e.g., the shortest paths between an agent’s location and the best viewpoint of an object of interest. In VDN, an approach known as Imitation Learning with Indirect Intervention (I3L) (Nguyen et al., 2019b) has been used where, in addition to learning from expert demonstration during training, an agent can also request help from an *assistant* both during training and evaluation in order to navigate to an object specified by the instructions (Nguyen and Daumé III, 2019).

In general, agents who learn from demonstrations do not generalize well, suffering from distribution shift issues due to the greedy nature of imitating expert demonstrations (Reddy et al., 2020). This tendency leads the agent in question to overfit to the seen environments, generally resulting in poor performance in the new, unseen environments.

3.2.2 REINFORCEMENT LEARNING

Although learning from supervision can lead to a faster learning process, agents trained under this paradigm often accumulate significant error due to limiting their exploration to the expert states (Wang et al., 2020b). In Reinforcement Learning (RL) settings, as opposed to having supervisions, an agent learns through interactions with an environment, e.g., by taking actions and receiving feedback from them. Through this setting, agents often learn more general behavior, and are capable of overcoming erroneous actions that may arise in unseen scenarios (Wang et al., 2018).

Policy gradient methods are frequently used in embodied research. This type of algorithms directly models and optimises a policy function, which provides an agent with the guideline for the optimal action that it can take at a given state. Within embodied AI research, two popular types of policy gradient methods used include REINFORCE algorithms (Williams, 1992) and actor-critic algorithms (Mnih et al., 2016; Schulman et al., 2017). The former use episode samples to update an agent’s policy parameters, whereas the latter combine policy learning with value learning (Sutton and Barto, 2018). Here, the policy has the role of the *actor*, choosing the action to take, while the value function plays the role of the *critic*, which criticizes the *actor*’s decisions (Sutton and Barto, 2018).

Actor-critic algorithms generally achieve smoother convergence and generally show superior performance, especially in continuous space problems (Schulman et al., 2017). One such example is

Asynchronous Advantage Actor-Critic (A3C) (Mnih et al., 2016), an algorithm designed for parallel training, where multiple actors interact asynchronously within an environment being controlled by a global network. More recently, Decentralized Distributed Proximal Policy Optimization (DD-PPO) (Wijmans et al., 2019b), an algorithmic extension of PPO (Schulman et al., 2017), has gained popularity within embodied goal-oriented tasks such as PointGoal Navigation (Schulman et al., 2017) and AudioGoal Navigation (Chen et al., 2020b).

Although the aforementioned techniques are used in EVLP research, they have been used jointly with supervised learning methods which we discuss in the following section. Nonetheless, some EVLP approaches employ RL-based approaches exclusively. For instance, leveraging REINFORCE algorithms, Wang et al. (2018) propose *Reinforced Planning Ahead (RPA)* an approach that couples model-free and model-based RL to train a look-ahead model that allows an agent to predict a future state, and thus, plan before taking an action.

For the tasks that require goal understanding and planning rather than instruction following, such as EQA and IQA, Gordon et al. (2018) propose the *Hierarchical Interactive Memory Network (HIMN)* where a model is decomposed into a hierarchy of controllers. For instance, a high-level planner chooses the low-level controllers where each low-level controller operates on a particular task, e.g., navigation, question answering, interacting, etc. Then, Das et al. (2020) propose a Question Answering (QA) model that trains RL agents to use the notion of predictive modeling that stimulates the agents to visit places in the environment that might become relevant in the future. Although RL algorithms do not require labeled data and can endow agents with general skills to solve a task, they suffer from several difficulties including reward specifications, slow convergence, and sample-complexity.

3.2.3 JOINT REINFORCEMENT AND SUPERVISED LEARNING

Combining reinforcement and supervised learning is commonplace in EVLP literature. By doing this, learning agents can leverage both the expert demonstrations and the direct feedback-based interaction with an environment to achieve more generalizable behavior.

In VLN, Tan et al. (2019) use both imitation learning (IL) as a weak supervision to mimic expert trajectories generated using a shortest-path algorithm and on-policy RL to train a more general behavior. Then, Wang et al. (2019b) propose an RL model, *Reinforced Cross-Modal Matching (RCM)*, where intrinsic rewards are used to encourage the learning agent to align the instructions with the trajectories. Then, they combine RCM with a self-supervised imitation learning approach (SIL), where the agent explores the environment by imitating its own past good decisions. Wang et al. (2019c) use both RCM and behavioral cloning to develop a generalized multitask model for natural language grounded navigation tasks including VLN and VDN where learnable parameters are shared between the tasks.

In the context of EQA, both Das et al. (2018b) and Yu et al. (2019) use IL to pre-train a model and RL to finetune it. In particular, Das et al. (2018b) propose *Neural Modular Controller (NMC)*, which uses A3C to fine-tune a set of sub-policies trained through behavioral cloning to mimic expert trajectories from a global policy. In contrast, Yu et al. (2019) uses REINFORCE to fine-tune an EQA agent trained on multi-target questions. This strategy allows agents to learn to recover from errors, which is not possible using expert demonstrations alone.

Blukis et al. (2019) introduce *Supervised Reinforcement Asynchronous Learning (SuReAL)* a framework that learns to map instruction to continuous control of quad-copters. SuReAL asynchronously trains two processes sharing data and parameters. The first one corresponds to a planner, trained via IL to predict a plan, and an action generator, trained using PPO to predict the control actions to solve a given plan. More recently, Krantz et al. (2021) introduce waypoint models for VLN tasks in continuous environments. Motivated by the recent success of PointGoal tasks, Krantz et al. (2021) leverage DD-PPO to extend the VLN-CE (Krantz et al., 2020) task to support language-conditioned waypoint prediction models at varying action-space granularities.

Wang et al. (2020a) propose a module for exploration which enables an agent to decide when and what to explore. They first use IL to guide an initial exploration policy from an expert. Then, they use RL to explore the state-action space outside the demonstration paths and reduce the bias toward copying expert actions. To overcome the issue of reward specification in RL, Wang et al. (2020b) propose a VLN method based on Random Expert Distillation (Wang et al., 2019a), known as Soft Expert Reward Learning (SERL), to learn a reward function by distilling knowledge from expert demonstrations.

Another challenge that has been tackled by using RL and IL jointly is the accurate alignment of the instructions and visual features. Zhu et al. (2020b) propose the Babywalk approach which uses IL to first segment paths and then learn over these finer-grained, shorter instructions by demonstration. Babywalk then uses curriculum-based RL to refine the policy by giving the agent increasingly longer navigation tasks (Bengio et al., 2009). This was shown to improve the model’s ability to extrapolate to longer sequences, although generalization to unseen data does not appear to benefit as much from this approach.

3.3 Common Techniques

The following subsections discuss common improvements beyond architecture, such as: data augmentation, pre-training, the use of additional training objectives, the use of different input representations, and optimization. Note that these can, and often are, used regardless of learning paradigm or architecture.

3.3.1 DATA AUGMENTATION

Data augmentation is commonly used to improve performance and robustness (Shorten and Khoshgoftaar, 2019; Tanner, 2012). In EVLP, back-translation is the most commonly used data augmentation technique. Back-translation was originally introduced by the neural machine translation community (Tan et al., 2019). It augments an existing corpus by generating synthetic training samples for mono-lingual (i.e., single view) data. Due to the difficulty of creating a paired corpus, it is beneficial to train a “backwards” model which creates additional training pairs.

In VLN, paths are sampled throughout the environment. A speaker model is trained to generate synthetic instructions for a given path. These synthetic sentences supplement the training data (Fried et al., 2018b). Improvements on this idea were proposed in Tan et al. (2019), in which the authors introduced Environmental Dropout. Here, entire frames from ground-truth paths are omitted to create “synthetic” environments. New instructions are created for these synthetic environments using

back-translation approach. This approach led to improvements in task success rate and narrowed the gap between the success rates on seen and unseen environments.

Other improvements to back-translation include changing the approach to sampling paths. Yu et al. (2020) proposes using longer paths which are not the shortest path between two points. Shortest paths may favor certain transitions in the training environment, which limit the agent’s need to learn language. Instead, the paper proposes using a random walk which favors similar transition probabilities at each point in the graph. Fu et al. (2020) propose using counter-factual reasoning when sampling augmentation paths. An adversarial path sampler (APS) proposes increasingly hard paths for the model to train on. This module is model-independent and uses the training loss to choose paths the model is struggling with.

3.3.2 ADDITIONAL OBJECTIVES

Additional tasks can also be used to increase the amount of training data and improve performance. In VLN, Ma et al. (2019a) add a completeness objective. Based on the data to date, the agent predicts how close to completion the agent is in meters. Progress monitor (PM) implicitly adds a prior over how instructions are processed. This approach encourages left-right attention of the instructions as the agent progresses through the path (Ma et al., 2019a).

To further improve performance, Zhu et al. (2020a) propose four additional tasks:

- *Trajectory retelling.* Generating a text based on the actions to date from the hidden state. The ground truth is generated using a speaker module (Fried et al., 2018b).
- *Angle prediction.* Generate the angle in degrees of the agent’s heading at the next time step.
- *Instruction-path matching.* Shuffle states within a batch and see if they correspond to given instructions.
- *Progress.* Percentage of steps taken.

Respectively, these auxiliary tasks have the following reasoning objectives: explaining the previous actions, predicting the next orientation, evaluating the trajectory consistency, and estimating the navigation progress (Zhu et al., 2020a).

3.3.3 PRE-TRAINING

Pre-training is done either by training a component of the model, or a full model on another problem entirely. This technique is used across most, if not all, neural EVLP approaches in the form of pre-trained image embeddings (He et al., 2016). Not only are image embeddings used for RGB images (Thomason et al., 2019b; Shridhar et al., 2020; Anderson et al., 2018c; Das et al., 2018a; Gordon et al., 2018), but also for depth information (Wijmans et al., 2019a; Krantz et al., 2020).

Pre-training can also be used for language inputs. Fried et al. (2018b) ablates the use of pre-trained word embeddings, finding that GloVe embeddings slightly improve performance (Pennington et al., 2014). More recent developments in language representation include pre-trained contextual models such as ELMO (Peters et al., 2018), BERT (Devlin et al., 2019) or the GPT models (Radford et al., 2019; Brown et al., 2020). Li et al. (2019) encodes language using either BERT or GPT-2, giving significant improvements compared to a baseline learned LSTM encoding.

Fusion of modalities can also benefit from pre-training. Recent approaches have taken advantage of large multimodal corpora (Sharma et al., 2018; Lin et al., 2014; Krishna et al., 2016; Shen et al., 2021b) to pre-train BERT-like models (Tan and Bansal, 2019; Lu et al., 2019) which capture cross-modal information. Majumdar et al. (2020) leverages pre-training to improve agent performance on VLN. It uses ViLBERT (Lu et al., 2019) as a base model to re-rank paths proposed by other models. Shen et al. (2021b) integrate contrastive language-image pre-training (CLIP) (Radford et al., 2021), a model designed for learning visual representations from natural language supervision, into Vision-and-Language tasks such as Visual Question Answering (VQA) (Antol et al., 2015) and VLN. CLIP leverages a large amount of image-text pairs for training, achieving state-of-the-art performance with zero-shot transfer in a variety of computer vision tasks (Radford et al., 2021). As such, Shen et al. (2021b) improve the performance of VLN agents on two datasets, R2R (Anderson et al., 2018c) and RxR (Ku et al. (2020) using EnvDrop (Tan et al., 2019) as their baseline model and CLIP as a visual encoder.

3.3.4 MULTITASK LEARNING

A closely related topic to pre-training is multitask training. In the age of transformers this approach has grown in popularity for multimodal tasks (Lu et al., 2020a; Nguyen and Okatani, 2019; Akhtar et al., 2019), including EVLP tasks.

Nguyen and Daumé III (2019) trains over three tasks datasets, two VLN datasets and a VDN dataset. The model proposed, *PREVALENT*, is a transformer that fuses modalities in a later layer. It is pre-trained over speaker augmented R2R (Fried et al., 2018b) by converting each step in the path to a *action, image, text triple*. *PREVALENT* is trained to predict the action and masking over the instructions. The models are then specialized and fine-tuned over the R2R, CVDN, and HANNA datasets. This approach benefits CVDN in both seen and unseen environments, while improving performance uniquely over unseen environments for R2R and CVDN. Wang et al. (2020c) train over both R2R and CVDN, but achieve lower performance over both datasets.

Task transfers have only been attempted between instruction following tasks. There have been no transfers between tasks requiring instruction following to those that require interpreting instructions. Doing so may be difficult, but could provide the a fine grained understanding of language, vision, and action required to interact with the world.

3.3.5 LEARNING AND OPTIMIZATION

Another way to improve performance is to modify the loss and optimization functions used. In EVLP tasks, there are two ways to pose the loss function: as a sequence-to-sequence (Seq2Seq) problem; as a re-ranking problem. In a Seq2Seq problem, a series of inputs are used to predict a sequence of ground-truth actions. Current approaches treat this as a classification problem, with the goal of learning a maximum likelihood estimate using cross-entropy loss conditioned on the agent’s observations. In a re-ranking problem, a Seq2Seq model first proposes candidate solutions, and a second model scores them. Majumdar et al. (2020) does this through the use of negative sampling, creating binary labels which reflect the validity of a path. Binary-cross entropy is again used as an objective.

Reasons to choose one over the other include performance, computational cost, and whether or not the environment is known. Re-ranking problems are, at the time of writing, state of the art in VLN

(Majumdar et al., 2020). However they are more computationally expensive and unable to explore the environment without use of an additional model(s).

Two optimization algorithms are commonly used (Mogadala et al., 2019): ADAM (Kingma and Ba, 2015), RMSPROP (Hinton et al., 2012). However the advantages or reasoning behind choosing one or the other are not discussed in EVLP literature.

3.3.6 REWARD SHAPING

Reward-shaping is a widely-used technique in solving Markov decision processes, as with Reinforcement Learning (RL), where domain knowledge is used to define a shaped reward R_{shaped} , for improved agent learning and optimisation (Hu et al., 2020). Typically, this shaped reward is linearly combined with the original reward R , with $R' = R + R_{shaped}$, or used instead. Reasons for incorporating a shaped reward include situations where R is sparse or is deemed not sufficiently informative for the task or for encouraging the desired agent behaviour (Sun et al., 2018).

The EVLP literature generally follows this standard formulation and adopts similar usage considerations for shaped rewards. In EQA, Das et al. (2018b), Gordon et al. (2018) and Yu et al. (2019) use a reward function R of the form $R = R_{terminal} + R_{shaped}$, where $R_{terminal}$ is a sparse reward assigned for correctly answering a question. Gordon et al. (2018) additionally assign a penalty for incorrect answers. R_{shaped} consists of a dense reward determined by the agent’s progress toward its goal: positive if moving closer to the goal and negative otherwise.

VLN and VDN works (Wang et al., 2018; Tan et al., 2019; Wang et al., 2020c,a) have used a similar reward function where $R_{terminal}$ is given if the agent stops within a specified radius of the target, and R_{shaped} is given as progress-based reward. Wang et al. (2019b) use the aforementioned rewards to define R_{extr} , an extrinsic reward to measure the agent’s success ($R_{terminal}$) and progress (R_{shaped}), and further define R_{intr} , an intrinsic reward for measuring the alignment between the instruction and the agent’s trajectory. The latter is obtained as a probability score of generating the original instruction, given sequential encodings of the instruction and the agent’s historical trajectory. The total reward function was defined as $R = R_{extr} + \delta R_{intr}$, where δ is a weighting parameter.

In addition to the sparse extrinsic reward, Jain et al. (2019) and Zhu et al. (2020b) use fidelity of the predicted path with the reference path based on the Coverage weighted by Length Score (CLS) metric (Jain et al., 2019), discussed in Section 4.1.3, to shape the reward function. Similarly, SERL (Wang et al., 2020b) uses two intrinsic rewards, R_{SED} , a soft-expert distillation reward, and R_{SP} , a self-perceiving reward: the former is learned by aligning the VLN agent’s behavior with expert demonstrations (Wang et al., 2019a), and the latter is learned by predicting the agent’s progress toward its goal.

4. Evaluation

Environments and metrics form the backbone of EVLP agent assessment. Proper evaluation systems allow to quantify specific aspects of agent performance (e.g., the ability to navigate toward a goal). Standardization of such metrics and the environments used to train and evaluate agents enables researchers to better coordinate current and future work (Anderson et al., 2018a). This section serves to summarize and compare the evaluation metrics and environments that are most relevant to the EVLP tasks defined in Section 2.2.

4.1 Metrics

The metrics currently used in EVLP can be grouped into five categories, each measuring a different aspect of agent performance: (i) success, (ii) distance, (iii) path-path similarity, (iv) instruction-based metrics, and (v) object selection metrics. We survey and discuss the existing methods for each category. For additional ease of reference, we summarize these metrics in tables 2-5.

4.1.1 SUCCESS

Metrics that assess whether or not an agent successfully completed a task are used, in various forms, across all EVLP tasks. Instruction-following tasks such as VLN and VDN typically define the *success rate (SR)* as how often the agent gets within a threshold distance d_{th} of the goal. Distance d_{th} varies between datasets, with $d_{th} = 3m$ for R2R (Chang et al., 2017), and $d_{th} = 0.47m$ for Lani (Misra et al., 2018). This measure has some weaknesses, particularly for discretized state-action spaces, notably that it is rather dependent on the granularity of said discretization (Thomason et al., 2019a; Krantz et al., 2020). Furthermore, variations in d_{th} also impact SR, which may prompt misleading results where correct executions might be considered wrong (Blukis et al., 2019).

EQA has no notion of a target location; instead, success is measured in terms of accuracy over the output space, i.e., question-answering accuracy (Das et al., 2018a; Gordon et al., 2018). Other common classification metrics that can be used include precision, recall, and *F-score* (Valuations, 2015). Per-class performance for any of these metrics can also be employed to measure if the model is improving accuracy by outputting common responses at the cost of performance over minority classes.

In most cases, a task is considered as “complete” only after the agent is near a final destination, answers a question, or satisfies a manipulation directive. However, certain datasets such as ALFRED (Shridhar et al., 2020), Lani (Misra et al., 2018), and CerealBar (Suhr et al., 2019) report sub-goal completion as well. This addition of sub-goal scoring can be useful, for instance, to exploit modularity and hierarchy of sub-tasks (Shridhar et al., 2020; Jansen, 2020), or to alleviate issues with success metrics, e.g. that early failures early in a task impact later steps in the episodes (Suhr et al., 2019), making it difficult to assess the ability of an agent per timestep. To address this issue, CerealBar evaluates a trajectory T as $|T| - 1$ sub-paths, each shorter than the next: the average success rate across these runs is taken, measuring the model’s ability to recover from errors. A further benefit of sub-goal scoring is that specific agent mistakes can be pinpointed and categorized more easily.

4.1.2 DISTANCE

While success is an intuitive measure, it provides limited information on how efficiently the agent traveled through a space. Distance measures can quantify this type of information, in the context of both navigation and manipulation tasks. Two useful metrics in this respect are navigation (or displacement) error (NE) and overall amount of travel (or motion) required known as the path length (PL). NE measures how far from a target location an agent chooses to stop, giving information about the magnitude of the agent’s inability to get to the target location. In order to further analyze whether agents learn to stop effectively, NE is often coupled with an Oracle Navigation Error (ONE), which measures the distance of the nearest point along the agent’s trajectory to the goal location. Significantly different NE and ONE measures may indicate an issue with stopping.

These measures assume that the task has a specified goal or target location. In the context of VLN and VDN this is a sensible assumption. However for EQA, or even EGM, that is not necessarily the case. For example, in EQA, certain questions such as those about the state of an object, the answer could be given from different points in the room. Rather than assess the distance from a goal, one could look at the length of the path taken commonly referred to as PL. In a real-world scenario, minimizing the length of travel is an important consideration. To only account for successful paths, SR weighted by normalized inverse Path Length (SPL) can be calculated.

In essence, distance can be used not only to derive success metrics for VLN, VDN, but also to quantify the error and efficiency of the agent’s path. In the context of navigation, this measures how much does the agent meander from its target. This adds a dimension when evaluating and comparing model performance.

4.1.3 PATH-PATH SIMILARITY

Distance metrics fail to capture whether or not the agent adhered to the ground-truth path. For instruction-following tasks such as VLN and VDN, where the user specifies a given path with instructions, we may want to explicitly follow those directions and not find the shortest path. Jain et al. (2019) discuss desiderata for path similarity metrics:

1. *Metrics should gauge the fidelity between the path taken and the reference path, rather than only the goal*
2. *Mistakes should not be hard penalties, metrics should favor a path that deviates and follows the ground truth path over one that short cuts it entirely*
3. *Maxima should be unique and only occur in the case of a perfect match*
4. *The underlying scale of the dataset should not impact the value of the output*
5. *Metrics should allow fast automated performance evaluation*

We believe these desiderata can be applied more widely to any instruction-following task. One of the early metrics to measure path following is edit distance (Chen et al., 2019b). Edit Distance (ED) measures the number of changes required to match two graphs. ED does not satisfy criterion #2, as it only penalizes based on absolute deviation. Jain et al. (2019) introduce Coverage weighted by Length Score (CLS), a metric which satisfies all 5 of these criteria. CLS is given by the product of the path coverage, a measure of average per point overlap between the target and reference path, and the length score, which penalizes paths both shorter and longer than the reference. Jain et al. (2019) verify its utility by using it as an objective for a RL-based agent.

Iiharco et al. (2019) build on this work and propose two measures: normalized Dynamic Time Warping (nDTW) and Success weighted by normalized Dynamic Time Warping (SDTW). The former is a similarity function that identifies an optimal warping of the elements from a reference path and a query path, while the latter constrains nDTW to successful episodes. These metrics have a number of desired properties: they respect the desiderata outlined above, they can be used effectively as a reward signal for RL-based agents, and can be compared with human judgment. Such properties make them consistently preferred compared to proposed alternatives such as CLS, SPL, SR.

4.1.4 INSTRUCTION-BASED

Different approaches can be used to measure the alignment between paths in the form of images and text: Semantic Propositional Image Caption Evaluation (SPICE) (Anderson et al., 2016) and Consensus-based Image Description Evaluation (CIDEr) (Vedantam et al., 2015). While those were originally used in the context of VQA, they have been recently adapted to VLN as well (Zhao et al., 2021).

An alignment metric can also be used as part of training. Certain approaches (Fried et al., 2018b; Tan et al., 2019) generate synthetic instructions to augment the data available, of varying quality. Huang et al. (2019) train an alignment model in which instructions and paths are encoded using BiLSTMs and fed into a binary classifier. Positive outputs are assigned to matched path instruction pairs and negative outputs are mismatched instruction path pairs. This binary classifier is then used as a scoring function to rank paths. Low scoring paths are removed from training leading to a statistically significant improvement in performance while reducing the total dataset size.

4.1.5 OBJECT REFERRAL

EGM and EOR require object selection, either through masking or selecting from ground-truth bounding boxes. Chen et al. (2019b) and Shridhar et al. (2020) use Intersection over Union (IoU) as an evaluation metric, a commonly-used measure in CV, specifically in object detection tasks (Padilla et al., 2020). However, it is not the only metric nor the most informative one.

Table 2: List of *success* metrics used in EVLP tasks with a brief description, type of metric, types and references to tasks they are typically used in, and the formula if applicable. The nomenclature used is as follows: P is the predicted path, R is the reference path corresponding to P , p and r are nodes in P and R , respectively. d is the appropriate distance, S is the set of instruction-path pairs. d_{th} is threshold distance, BB is the bounding box, \hat{y} is a prediction value, and y the corresponding ground-truth value.

Metric	Description	Type	Typical Tasks	Formulas
SR (Anderson et al., 2018c)	Percentage of episodes an agent stops within a threshold distance of goal	Success	VLN, VDN, EOR, EGM	$\frac{1}{ S } \sum_{P,R \in S} \mathbb{1}[NE(P,R) < d_{th}]$
OSR (Anderson et al., 2018c)	SR if the agent had stopped at it's nearest point to the goal	Success	VLN, VDN, EOR	$\frac{1}{ S } \sum_{P,R \in S} \mathbb{1}[ONE(P,R) < d_{th}]$
Sub-goal SR (Shridhar et al., 2020)	SR for a specific sub-goal	Success	VLN, EGM	SR
Accuracy (Das et al., 2018a)	Measures whether a prediction value matches the ground-truth value	Success	EQA	$\mathbb{1}(\hat{y} = y)$
EM	The agent perfectly replicates the target path	Success	All	$\mathbb{1}(P = R)$

Recently, Kim et al. (2020) proposed three metrics to evaluate how well objects are selected: Collected Target Correctness (CTC), Placement Target Correctness (PTC), and reciprocal Placement Object Distance (rPOD). CTC measures if the correct object was selected; PTC measures if the object was correctly placed. Note that CTC and PTC can both be framed as thresholded IoU problems, as that is how some agents (Shridhar et al., 2020) decide whether or not to pick up an object. rPOD is a normalized measure that gives a near zero score if the object is far from it's ideal placement and a score of 1 if the placement is perfect.

Table 3: List of *distance* metrics used in EVLP tasks with a brief description, type of metric, types and references to tasks they are typically used in, and the formula if applicable, where P is the predicted path, R is the reference path corresponding to P , p and r are nodes in P and R , respectively.

Metric	Description	Type	Typical Tasks	Formulas
NE (Misra et al., 2018)	Geodesic distance between an agent’s final position and the goal	Distance	VLN, VDN	$d(p_{ P }, r_{ R })$
ONE (Misra et al., 2018)	Distance to target if the agent has stopped at the nearest point	Distance	VLN, VDN	$\min_{p \in P} d(p, r_{ R })$
PL (Anderson et al., 2018c)	How far the agent traveled	Distance	VLN, VDN, EOR	$\sum_{i=1}^{ P -1} d(p_i, p_{i+1})$
SPL (Anderson et al., 2018a)	Average of the path lengths where the agent succeeded	Distance	VLN, VDN, EOR, EGM	$\mathbb{1}[NE(P, R) < d_{th}] \frac{d(p_1, r_{ R })}{\max(PL(P), d(p_1, r_{ R }))}$

Table 4: List of *path-path similarity* metrics used in EVLP tasks with a brief description, type of metric, types and references to tasks they are typically used in, and the formula if applicable; here, P is the predicted path, R is the reference path corresponding to P , p and r are nodes in P and R , respectively. d is the appropriate distance, d_{th} is threshold distance, and $lev(\cdot, \cdot)$ is the Levenshtein edit distance.

Metric	Description	Type	Typical Tasks	Formulas
CLS (Jain et al., 2019)	Measure of similarity between paths using path coverage and length-score	Path-Path Similarity	VLN	$PC(P, R) \cdot LS(P, R);$ $PC(P, R) = \frac{1}{ R } \sum_{r \in R} \exp\left(-\frac{d(r, P)}{d_{th}}\right);$ $LS(P, R) = \frac{PC(P, R) \cdot PL(R)}{PC(P, R) \cdot PL(R) + PC(P, R) \cdot PL(R) - PL(P) }$
nDTW (Ilharco et al., 2019)	Dynamic time warping normalized between target and reference path	Path-Path Similarity	VLN	$\exp\left(-\frac{DTW(R, P)}{ R \cdot d_{th}}\right)$
SDTW (Ilharco et al., 2019)	nDTW weighted by success rate	Path-Path Similarity	VLN	$\mathbb{1}[NE(P, R) < d_{th}] \cdot nDTW(P, R)$
SED (Chen et al., 2019b)	Success weighted edit distance	Path-Path Similarity	VLN	$\mathbb{1}[NE(P, R) < d_{th}] \cdot \left(1 - \frac{lev(P, R)}{\max(P , R)}\right)$

Table 5: List of *path-instruction fidelity* metrics used in EVLP tasks with a brief description, type of metric, types and references to tasks they are typically used in, and the formula if applicable. The nomenclature used is as follows: S is the set of instruction-path pairs, BB is the bounding box, \hat{y} is a prediction value, and y the corresponding ground-truth value. Entries with a single asterisk (*) are learned approaches. For entries with double-asterisks (**), we refer the reader to the referenced work.

Metric	Description	Type	Typical Tasks	Formulas
BLEU (Papineni et al., 2002)	Compare generation caption to some reference	Path-Instr. Fidelity	Instr. generation	$BP \cdot \exp\left(\frac{1}{ N } \sum_{n=1}^N \frac{\text{Count}(n, s)}{\text{Count}(n, \hat{s})}\right), s \in S;$ $BP = \exp\left(1 - \frac{ s }{ \hat{s} }\right)$ if $ s \leq s $ else 1
SPICE (Anderson et al., 2016)	Measure of image-caption alignment	Path-Instr. Fidelity	Instr. generation	*
CIDER (Vedantam et al., 2015)	Parsing based method to align captions to images	Path-Instr. Fidelity	Instr. generation	**
ROUGE-N (Lin, 2004)	Compare caption generation to some reference	Path-Instr. Fidelity	Instr. generation	$\frac{\sum_{n-\text{gram} \in S} \text{Count}_{\text{match}}(n-\text{gram})}{\sum_{n-\text{gram} \in S} \text{Count}(n-\text{gram})}, s \in S$
Path-Instr discriminator (Huang et al., 2019)	Discriminator trained to find matching vs synthetically sampled datasets	Path-Instr. Fidelity	VLN	**
Path-Instr compatibility (Zhao et al., 2021)	Compares path-instruction alignment	Path-Instr. Fidelity	VLN	*
IoU (Everingham et al., 2010)	Measures overlap between 2 bounding boxes	Object Referral	EQA, EGM, EOR	$\frac{ BB_1 \cap BB_2 }{ BB_1 \cup BB_2 }$

While only applicable to EGM and EOR, measuring the quality of object selection interaction is important. It is another aspect of the task and provides a good way of understanding what aspects of language are harder for the model. Recent work in object recognition (Padilla et al., 2020) reveals that there are many metrics and variants. Even more traditional metrics such as precision, recall, and F-score are not currently reported in those tasks. Doing so would provide more information about the bounding box accuracy. In addition we could report per-class accuracy, which would be useful to identify which objects are poorly detected by the model.

4.2 Simulation Environments and Datasets

Solving EVLP tasks generally involves using a simulation environment and a dataset. Simulation platforms and datasets facilitate the reproduction and evaluation of embodied systems. Simulators aim to replicate aspects of the real-world and model agents capable of solving complex tasks while abstracting away the difficulties of designing and supervising a real-world agent. In contrast, datasets play a crucial role in articulating how each task is framed. They provide examples of *how* agents should behave in response to certain multimodal stimuli.

In this section, we provide a per-task summary and comparison of literature pertaining to environments and datasets used in, or designed for EVLP research. Furthermore, we provide a brief discussion of other recent and relevant multimodal platforms and datasets outside the EVLP tasks discussed in this survey article. Finally, we summarize important aspects of these simulators and datasets in Table 6 and Table 7, respectively.

4.2.1 SIMULATORS

Early simulation platforms for embodied research typically leverage video game environments to create and train neural controllers. Human performance was quickly achieved on some of these platforms (Bellemare et al., 2015; Kempka et al., 2016; Synnaeve et al., 2016; Beattie et al., 2016; Johnson et al., 2016; Nichol et al., 2018) as simplified environments generally lack the diversity and complexity of real-world settings (Savva et al., 2017; Kolve et al., 2017; Deitke et al., 2020).

Recent works have addressed this lack of realism through the use of photo-realism (Savva et al., 2019; Xia et al., 2018; Anderson et al., 2018c; Herman et al., 2021) and the use of interactive contexts where agents are able to modify the states of objects in the environment (Kolve et al., 2017; Xia et al., 2019; Yan et al., 2018; Brodeur et al., 2018). Toward this end, there is also interest in developing frameworks focused on simulation-to-real transfer and evaluation (Deitke et al., 2020), allowing the study of discrepancies between real settings and simulated ones. Finally, other platforms have also focused on encouraging reproducibility of work, flexibility of design, and benchmarking (Savva et al., 2019; Weihs et al., 2020).

VLN Simulators. Three environments are relevant to VLN tasks: Matterport3DSim (Chang et al., 2017; Anderson et al., 2018c), Habitat (Savva et al., 2019), and StreetLearn (Mirowski et al., 2018). In both Matterport3DSim and Habitat, VLN is set in indoor environments, while StreetLearn is set outdoors. Matterport3DSim allows an agent to navigate in a discrete environment framed as an undirected navigation graph which associates positions with the corresponding viewpoints, subject to admissible (i.e., collision-free) intra-node transitions. Agents navigating in this setting *teleport* between viewpoints, through high-level atomic actions.

In contrast, Krantz et al. (2020); Irshad et al. (2021) leverage Habitat to provide a more realistic setting for VLN, where agents can navigate in a discrete (Krantz et al., 2020), as well as, continuous (Irshad et al., 2021) unconstrained state-space, as opposed to a discretized graph. In the unconstrained setting, agents execute low-level actions over longer horizon sequences (Irshad et al., 2021), and more realistically, risk colliding with obstacles along the way. Tangentially, VLN has also been studied in the context of simulation-to-real transfer and evaluation. Anderson et al. (2020) ports a VLN agent trained in Matterport3DSim using Robot Operating System (ROS) as a publish-subscribe framework for robot control. They also evaluated the gap in performance between an agent moving through high-level actions compared to a robot taking low-level actions in the real world. They found that, although there has been significant progress in simulation environments for VLN, simulation-to-real transfer is still not reliable, particularly when the agent does not have access to a map.

EQA Simulators. Both, indoor, synthetic (Das et al., 2018a; Gordon et al., 2018) and photo-realistic (Wijmans et al., 2019a) environments are used for EQA. Das et al. (2018a) use House3D (Wu et al., 2018), a synthetic large-scale environment to train EQA agents, whereas Gordon et al. (2018) used AI2-THOR (Kolve et al., 2017). In both settings, the agent navigates in a discrete unconstrained environment and perceives the world through ego-centric RGB images. While AI2-THOR is not a large-scale environment as House3D, it enables the study of interactive tasks something crucial in IQA.

Wijmans et al. (2019a) argue that although the above-mentioned environments are semantically realistic in terms of layout and types of surrounding objects, they lack visual realism in terms of graphics and intra-class variation. Thus, they propose MP3D-EQA, an EQA task set in a photorealistic simulator, and built upon the MINOS simulator (Savva et al., 2017) and the Matterport3D dataset (Chang et al., 2017). In addition leveraging RGB images, they add 3D point-cloud observations, which they empirically showed to be more effective for navigation, in contrast to previous works. More recently, MP3D-EQA has also been developed in Habitat. Savva et al. (2019) showed that Habitat’s performance in terms of frames-per-second (FPS) is significantly superior to that of MINOS, and other platforms, allowing to shift bottlenecks in simulation training.

EOR Simulators. Both indoor and outdoor settings exist for the EOR task. REVERIE (Qi et al., 2020) is an indoor-based simulator built upon Matterport3DSim. In particular, it adds object-level bounding boxes to Matterport3DSim allowing for object retrieval. Since a target object may be observed from multiple viewpoints, and because it is expected that the agent reaches the goal within a short distance, they only preserve object bounding boxes within a 3-meter radius of any given viewpoint.

TouchDown (Chen et al., 2019b) was made publicly available by transferring it to the StreetLearn environment (Mirowski et al., 2018; Mehta et al., 2020). StreetLearn contains panoramas from New York City and Pittsburgh. Like Matterport3DSim, navigation in StreetLearn is based on an undirected graph, and agents move through discrete actions.

VDN Simulators. The VDN task family has also been established in Matterport3DSim (Nguyen and Daumé III, 2019; Thomason et al., 2019b; Chi et al., 2020). Typically the VLN task is extended with crowd-sourced dialog or sub-instructions that provide the agents with assistance to complete the task. We refer the reader for further information to Section 4.2.2.

EGM Simulators. As discussed in Section 2.2.5, interaction with the environment is a crucial component of the EGM task family. Currently, photo-realistic simulation environments are limited in providing interactive or dynamic interaction prospects. Consequently, EGM tasks have mainly leveraged synthetic environments (Shridhar et al., 2020; Misra et al., 2018; Suhr et al., 2019; Kim et al., 2020).

Shridhar et al. (2020) used the AI2-THOR 2.0 simulation environment for the ALFRED task. In contrast to IQA, which also leverages this simulator, ALFRED does not require agents to answer contextualized questions. Instead, agents must complete tasks that require changing the state of objects (e.g., cleaning, picking, stacking, and heating objects).

In contrast, ArraMon (Kim et al., 2020), and CerealBar (Suhr et al., 2019) are both framed in outdoor synthetic environments. ArraMon uses a custom two-phase environment consisting of synthetic street-layouts and buildings, used for the navigation phase, and a grid-based room layout used during the assembly phase. CerealBar was designed to study of collaborative interactions between agents.

Other Simulators. We briefly discuss other platforms that have recently shown significant promise, in the definition of multimodal embodied tasks that outside conventional EVLP families. Some of the features proposed by these simulators include designing interactive tasks with improved physics (Xia et al., 2019; Szot et al., 2021; Xiang et al., 2020; Gan et al., 2020) and object states (Li et al., 2021a), additional modalities (Chen et al., 2020b; Shen et al., 2021a), low-level articulation (Xiang et al., 2020; Szot et al., 2021), and integraton with virtual reality interfaces (Li et al., 2021a; Gan et al., 2020).

iGibson 1.0 (Xia et al., 2019; Shen et al., 2021a) is a recent simulation environment that allows for the development of interactive navigation and manipulation tasks in large-scale realistic scenes, in contrast to its precursor Gibson (Xia et al., 2018), which only provides support for robot navigation tasks (where agent interaction is limited to collisions with scene meshes). Among the most valuable aspects of this platform are the physics-based simulation capabilities, which enable high-quality interactive tasks. More recently, iGibson 2.0 (Li et al., 2021a) was provided as an extension to iGibson 1.0; it includes a set of extended states and logical descriptions, which simulate such underlying physical processes as temperature, wetness, and cleanness of objects—thus motivating a wider range of household tasks and perceptual modes. Moreover, to facilitate modeling household tasks, iGibson 2.0 introduces a virtual reality (VR) interface to allow humans to collect demonstrations, e.g., for developing solutions that leverage imitation-based learning.

Szot et al. (2021) present Habitat 2.0, a simulation platform for training virtual robots in interactive environments and physics-enabled synthetic scenarios. It was built upon Habitat (Savva et al., 2019) and, as such, it prioritises high-performance and speed. Szot et al. (2021) introduce the Home Assistant Benchmark (HAB) a suite of tasks for assistive robots, such as, preparing groceries and cleaning the house, designed for testing the manipulation capabilities of agents. Also extending Habitat, SoundSpaces (Chen et al., 2020b) integrates environment acoustics to enable audio-based navigation tasks in indoor environments. To do so, they leverage algorithms for modeling room acoustics as well as sound reflections based on room geometry.

VirtualHome (Puig et al., 2018) is an environment designed for simulating household activities by generating programs containing sequences of symbolic instructions representing atomic actions. Modalities supported by this platform include RGB, depth, semantic segmentation and natural language.

ThreeDWorld (TDW) (Gan et al., 2020) enables the development of interactive navigation and manipulation tasks in both outdoor and indoor environments. It also allows users to procedurally create custom environments (and the objects to populate them), as well as interact with objects through VR. TDW also supports complex physical dynamics, with audio support, which is relevant for exploring topics such as physical reasoning (Zhang et al., 2017) and audio-visual learning (Chen et al., 2020b; Gan et al., 2019). Finally, Sapien (Xiang et al., 2020) is a synthetic-based platform that provides rich physics support for developing low-level interactive manipulation tasks.

Table 6: A comparison of simulation environments, related to EVLP research. **Supported Tasks:** we include PointNav (Wijmans et al., 2019b), RoomNav(Wu et al., 2018), AVN (Audio-Visual Navigation; (Chen et al., 2020b)), HAB (Szot et al., 2021), ArmPointNav (Ehsani et al., 2021), general navigation, and manipulation, in addition to conventional EVLP tasks. **Agent-World Interaction:** C: continuous actions / forces, D: discrete actions. **Environments:** S: synthetic, P: photorealistic. **Scenes:** I: indoors, O: outdoors. **Modalities:** A: audio, C: color, D: depth, F: flow, L: LiDAR, NL: natural language, SN: surface normals, SS: semantic segmentations. **Scene Scale:** B: buildings, R: rooms, H: houses, RO: realistic outdoors, SO: synthetic outdoors O: object-level.

Environment	Supported Task	Agent-World Interaction	Environments	Scenes	Scene Scale	Modalities	Multi-Agent	Specialty	Reference
Matterport3DSim	VLN, VDHN, EOR	D	P	I	B	C,D,NL		Topological planning	Chang et al. (2017)
MINOS	EQA, PointNav, RoomNav	C,D	PS	I	B	C,D,NL,SN,SS		High-speed, navigation	Savva et al. (2017)
A12-THOR	VDHN, EQA, EGM	C,D	S	I	R	A,C,D,SS	✓	Actionable objects, task planning	Kolve et al. (2017)
ManipulaTHOR	ArmPointNav	D	S	I	R	C,D		Mobile manipulation	Ehsani et al. (2021)
House3D	EQA, RoomNav	D	S	I	B,R	C,D,SS,NL		Environment diversity, navigation	Wu et al. (2018)
CerealBar	EGM	D	S	O	SO	C,NL	✓	Collaborative interactions	Suhr et al. (2019)
ArraMon	EGM	D	S	O	SO	C,NL		Joint navigation and manipulation	Kim et al. (2020)
StreetLearn	VLN, EOR	D	P	O	RO	C,NL		Urban Navigation	Mirowski et al. (2018)
Gibson	Nav.	C	PS	I	H	C,D,SN,SS		Navigation	Xia et al. (2018)
iGibson 1.0	Nav., Manip.	C	PS	I	H	C,D,F,L,SN,SS		Physics Interaction	Shen et al. (2021a)
iGibson 2.0	Nav., Manip.	C	PS	I	H	C,D,F,L,SN,SS		Extended object states, VR	Li et al. (2021a)
Habitat	VLN, EQA, PointNav	C,D	PS	I	B	C,D,SS		High-speed, navigation, customizable	Savva et al. (2019)
Habitat 2.0	HAB	C	S	I	H	C,D		High-speed, task planning, articulated dynamics	Szot et al. (2021)
SoundSpaces	AVN	C,D	PS	I	B	A,C,D,SS		High-speed, audio-visual navigation	Chen et al. (2020b)
VirtualHome	Nav., Manip.	D	S	I	H,R	C,D,F,SS		Task Planning	Puig et al. (2018)
ThreeDWorld	Manip.	C	S	I,O	R,H,SO	A,C,D,SS	✓	Audio, Physics Interaction, VR	Gan et al. (2020)
Sapien	Manip.	C	S	I	O	C,D,SN,SS		Low-level articulation	Xiang et al. (2020)

4.2.2 DATASETS

EVLP datasets vary across three main dimensions: visual observations, natural language prompts, and navigation demonstrations. Visual observations, in general, consist of RGB images often paired with depth data or semantic masks. These observations can represent both indoor and outdoor environments from both, photo-realistic or synthetic-based settings. In contrast, language varies in the type of prompt. Language prompts may come in the form of questions (Das et al., 2018a; Wijmans et al., 2019a), step-by-step instructions (Anderson et al., 2018c), or ambiguous instructions that require some type of clarification through dialog or description (Nguyen and Daumé III, 2019; Thomason et al., 2019b). Language can also vary in terms of complexity of language sequences and scope of vocabulary. Finally, navigation traces differ in aspects like the granularity (or discretization) of the action-space and the implicit alignment that a provided action sequence (or trajectory) has with the other two dimensions.

VLN Datasets. Both indoor (Anderson et al., 2018c; Jain et al., 2019; Zhu et al., 2020b; Hong et al., 2020) and outdoor (Chen et al., 2019b; Mehta et al., 2020; Hermann et al., 2020) datasets have been

designed for the VLN task. Anderson et al. (2018c) introduced Room2Room (R2R), the first and most widely known indoor VLN dataset. A number of R2R variants have been developed since its introduction. Some of these variants propose longer and more complex navigation sequences. For instance, R4R (Jain et al., 2019), R6R, and R8R (Zhu et al., 2020b) are formed by concatenating sequences of two, three and four paths from the R2R dataset, respectively. Others aim to provide better alignment between instructions and visual observations. For instance, Fine-Grained Room-to-Room (FGR2R) (Hong et al., 2020), proposes finer alignment between instructions and visual inputs. To do so, the authors provide additional supervision for the agent by dividing R2R instructions into sub-instructions and matching them with visual observations along the corresponding path. Room-Across-Room (RxR) (Ku et al., 2020) uses the PanGEA annotation tool (Ku et al., 2021) to also tackle path-instruction alignment by densely recording and matching 3D pose traces with audio-based instructions.

Multilingual datasets also exist in VLN. Two such examples of this type of dataset are Cross-Lingual Room-to-Room (XL-R2R) (Yan et al., 2020) with instructions in Chinese in addition to English, and RxR (Ku et al., 2020) which uses instructions in Hindi, Telugu, and English.

The preceding datasets are based on Matterport3DSim where, as mentioned in Section 4.2.1, navigation is constrained to an undirected graph. Nonetheless, making the underlying navigation space more realistic is another direction into which R2R has evolved. Krantz et al. (2020) proposed VLN in Continuous Environments (VLN-CE), which transfers the graph-based trajectories of R2R to an unconstrained discrete state-space in AI Habitat. Furthermore, Irshad et al. (2021) proposed RoboVLN, a dataset which extends the high-level discrete actions of VLN-CE by adding low-level ones that represent the agent’s linear and angular velocities between high-level actions. Beyond continuity, there are other shortcomings in existing datasets that hamper realism. These include biases in path distribution (Ku et al. (2020)). Some datasets, such as R2R, have limited variability in path length. In addition they only contain paths that go directly from the origin to a target location. Both these biases impact generalization to new environments and the agent’s ability to follow instructions, favoring goal finding rather than instruction following.

In addition to R2R and its variants, other datasets explore outdoor navigation settings. For instance, LANI (Misra et al., 2018) is a dataset designed for outdoor navigation in a synthetic environment. LANI evaluates an agent’s ability to follow instructions referring to multiple landmarks. Hermann et al. (2020) propose the StreetNav Suite which leverages the StreetLearn environment (Mirowski et al., 2018) built using RGB panoramas from Google Street View. Similar to R2R, in this dataset navigation is based on an undirected graph but unlike the previous datasets it uses synthetic instructions rather than human-generated ones.

EQA Datasets. Datasets in EQA mainly differ based on the type of environment used (Section 4.2.1) and the type of questions asked. For instance, IQUADv1 (Gordon et al., 2018) is an indoor-based dataset that contains three types of questions: existence questions (*Is there an [OBJECT] in [PLACE]?*), counting questions (*How many [OBJECT_1] are on the [OBJECT_2]?*), and spatial relationship questions (*Are there [OBJECT_1] in the [OBJECT_2]?*). These questions along with their corresponding answers are generated automatically using templates.

EQAv1 (Das et al., 2018a) also implements template-generated questions and answers. The type of questions involve scene recognition (*What [ROOM] is the [OBJECT] located in?*), color recognition

(*What color is the [OBJECT] ?*), and spatial reasoning (*What is [on/above/below/next-to] the [OBJECT] in the [ROOM] ?*). The object in a question is single-target, and the possible answer options are room names, color names, and object names. MP3D-EQA (Wijmans et al., 2019a) is the extension of EQAv1 to photo-realistic environments. However, it only defines questions about scene and color recognition. Finally, MT-EQA (Yu et al., 2019) proposes a generalization of EQA where the types of questions involve multiple target objects and rooms. They propose 6 types of questions which compare attribute properties (color, size and distance) between multiple targets (objects and rooms). For instance, *Is [ROOM_1] [bigger/smaller] than [ROOM_2] ?*, or, *Does [OBJECT_1] in [ROOM_1] share the same color as [OBJECT_2] in [ROOM_2] ?*.

VDN Datasets. Current indoor-based VDN datasets were built on the R2R dataset. JustAsk (Chi et al., 2020) uses R2R and extends the VLN task by adding an interactivity module where a confused agent can ask an oracle to resolve any ambiguities. In turn, the oracle provides an action to take with some probability of being incorrect. In HANNA (Nguyen and Daumé III, 2019), an agent receives traditional VLN instructions and can also request assistance when it is lost or confused. Here, in contrast to JustAsk, an oracle named ANNA provides the agent with a sub-task to help correct its path. As such, the dataset builds a route system simulating human assistance (sub-tasks) on top of the R2R dataset.

Thomason et al. (2019b) study dialog interactions between agents and human assistants. To do so, they designed the Cooperative Vision-Dialog Navigation (CVDN) dataset which consists of ambiguous and underspecified instructions along with human-based dialog guidance between a navigator and an oracle. Compared to the step-by-step instructions in R2R, their instructions do not initially specify a goal location and require clarification. Thus, resolution is done through the use of dialog between a navigator, who is given the instruction, and an oracle, who knows how to resolve ambiguities.

Similarly, Talk the Walk (de Vries et al., 2018) consists of a *tourist* and a *guide* that interact with each other using dialog to get the tourist to a target location. The dataset consists of RGB panoramic images of New York City which are used to build a 2D grid-like environment in which the tourist moves. The dialogs for this dataset were crowd-sourced and correspond to trajectories followed from randomly sampled start and target locations.

EOR Datasets. There are two datasets in EOR: REVERIE (Qi et al., 2020) and Touchdown (Chen et al., 2019b; Mehta et al., 2020). REVERIE is an indoor-based dataset built upon Matterport3DSim. It differs from the R2R dataset in two main aspects: (1) instructions are intended to be more realistic, and (2) the dataset contains object-level annotations. In R2R, instructions are often long and provide step-by-step guidance toward a goal location. Qi et al. (2020) argue that it is unrealistic for humans to give such instructions to a robot. Thus, they propose a dataset containing simplified instructions such as *Bring me a mug*. Additionally, the REVERIE dataset includes object annotations in the form of 2D bounding boxes since the task requires the agent to localize the target object to which the instruction refers (e.g., the mug).

Touchdown is set in an urban environment. The dataset leverages RGB panoramic images of New York City in Google Street View. It has recently been added to StreetLearn (Mehta et al., 2020) and made publicly available to the research community. Like StreetNav, navigation is based on an undirected graph. The goal of the agent is to find the location of a hidden teddy bear named

Touchdown. Touchdown only appears if the correct location was selected. Compared to REVERIE, the instructions in Touchdown represent a step-by-step guidance toward the goal, similar to VLN. Moreover, these instructions clearly specify where to find the teddy bear, whereas in REVERIE, the agent needs to infer both the object of interest and its probable location in the environment.

EGM Datasets. All of the contemporary EGM datasets use interactive synthetic environments (Subsection 4.2.1). EGM datasets are crowd-sourced and validated by human annotators which interact with these environments to verify the alignment between instructions, visual inputs, and interactions.

The CHAI dataset (Misra et al., 2018) is based on household instructions (e.g., *Go to the kitchen and put the cup on the sink*). A single instruction can be divided into sub-tasks requiring multiple navigation and manipulation steps. ALFRED (Shridhar et al., 2020) is also based on household instructions with multiple goals. The main differences between ALFRED and CHAI are that ALFRED includes expert demonstrations and allows the agent to perform several manipulation actions (e.g., *pick up, turn on, open, etc.*), whereas CHAI only uses an *interact* action.

ArraMon (Kim et al., 2020) is set on an outdoor environment and consists of two types of instructions: navigation and assembly. Navigation instructions provide step-by-step guidance toward a goal object which the agent requires to find and collect. Assembly instructions specify where does the agent need to place each the collected item. These directives come in turns, i.e., a navigation instruction is followed by an assembly one. Each episode is defined by two turns. Lastly, CerealBar (Suhr et al., 2019) also set outdoors, considers a turn-based game between two agents, a *leader*, and a *follower*. The leader gives step-by-step instructions to collect game cards of different colors, shapes and count, and the follower has to execute them.

Table 7: Summary of EVLP dataset statistics. *MP3D* refers to the Matterport3D simulator (Chang et al., 2017). *CS* refers to crowd-sourced instructions and *G* refers to generated instructions. We use *EN* for English, *CH* for Chinese, *HI* for Hindi and *TE* for Telugu. Then, we use *N* for navigation and *M* for manipulation. The vocabulary sizes provided are taken from the original works and account for all unique words generated using that work’s tokenizer. Finally, missing fields, indicated by hyphens (-), were not computed by the originating works.

Dataset	Task	Simulator	Navigation Information			Language Information				Reference	
			Avg. # Steps	# Trajectories	Navigation	# Instr. / Dialog / Questions	Avg. Words Per Instruction	Instruction Type	Vocabulary Size		Language
R2R	VLN	MP3D	5	7,189	graph	21,567	29	CS	3,100	EN	Anderson et al. (2018c)
R4R	VLN	MP3D	20.5	7,189	graph	279,810	58.4	CS	3,100	EN	Jain et al. (2019)
R6R	VLN	MP3D	-	-	graph	12,5409	91.2	CS	3,100	EN	Zhu et al. (2020b)
R8R	VLN	MP3D	-	-	graph	138,004	121.6	CS	3,100	EN	Zhu et al. (2020b)
FG2R2	VLN	MP3D	-	7,189	graph	21,567	29	CS	3,100	EN	Li et al. (2020b)
XL-R2R	VLN	MP3D	-	5,798	graph	17,394	-	CS	EN: 1,583, CH: 1,134	EN, CH	Yan et al. (2020)
RXR	VLN	MP3D	7	16,522	graph	126,000	EN: 101, HI: 76, TE: 56	CS	-	EN, HI, TE	Ku et al. (2020)
VLN-CE	VLN	Habitat	55.8	4,475	discrete	13,425	-	CS	-	EN	Krantz et al. (2020)
RoboVLN	VLN	Habitat	326	3,177	continuous	9,533	-	CS	-	EN	Irshad et al. (2021)
LANI	VLN	Unity3D	24.6	6,000	grid	28,204	12.1	CS	2,292	EN	Misra et al. (2018)
StreetNav	VLN	StreetLearn	125	613,000	graph	~2,5M	7	G	-	EN	Hermann et al. (2020)
IQAv1	EQA	A12-THOR	-	-	discrete	75,000	-	G	70	EN	Gordon et al. (2018)
EQAv1	EQA	House3D	-	-	discrete	5,000	-	G	-	EN	Das et al. (2018a)
MT-EQA	EQA	House3D	-	-	discrete	19,287	-	G	-	EN	Yu et al. (2019)
MP3D-EQA	EQA	Habitat	-	-	discrete	~1,100	-	G	-	EN	Wijmans et al. (2019a)
CVDN	VDHN	MP3D	17.4	7,000	graph	2,050	81.6	CS	~1,100	EN	Thomason et al. (2019b)
HANNA	VDHN	MP3D	-	-	graph	21,594	-	CS	2,532	EN	Nguyen and Daumé III (2019)
Talk the walk	VDHN	Talk the Walk	62	-	grid	10,000	-	CS	10,000+	EN	de Vries et al. (2018)
REVERIE	EOR	MP3D	-	-	graph	21,702	18	CS	1,600	EN	Qi et al. (2020)
(Re)Touchdown	EOR	StreetLearn	-	-	graph	9,326	108	CS	5,625	EN	Chen et al. (2019b); Mehta et al. (2020)
CHAI	EGM	Chalet	54.5	-	discrete	13,729	10.1	CS	1,018	EN	Misra et al. (2018)
ALFRED	EGM	A12-THOR	50	8,055	graph	25,726	50	CS	~100	EN	Shridhar et al. (2020)
ArraMon	EGM	ArraMon	75.78 (N), 13.68 (M)	7,692	discrete/grid	30,768	48 (N), 21 (M)	CS	-	EN, HI	Kim et al. (2020)
CerealBar	EGM	CerealBar	8.5	1,202	grid	23,979	14	CS	3,641	EN	Suhr et al. (2019)

5. Open Challenges in Embodied Vision-Language Planning

5.1 New Directions in EVLP Research

Research in Embodied Vision-Language Planning has experienced a quick rise in popularity, due to recent advances from the robotics, computer vision, and natural language processing communities on vision-language grounding and situated learning. As the EVLP research field evolves, we highlight three promising directions, in the pursuit of more ubiquitous human-robot interaction and better agent generalisation. Firstly, we describe *social interaction* (Section 5.1) as a necessary progression from static dialogue contexts to more dynamic ones, where new collaborative and assistive capabilities of agents can more effectively emerge. Next, we discuss the need for developing agents in *dynamic environments* (Section 5.1), which encourage agents to incorporate reasoning strategies that are robust to environment uncertainty and non-stationarity. Finally, we discuss a vision for *cross-task robot learning* (Section 5.1) for cross-task transfer, wherein agents may acquire experience from related modality-centric tasks, before deployment to shared multimodal settings with significant task overlap.

5.1.1 SOCIAL INTERACTION

Social agent interaction is born out of a desire for an agent to carry out tasks and make decisions collaboratively with other agents in a complex, multimodal world. Particularly compelling application formats are those that require the intelligent system to perform EVLP tasks, wherein agents not only perceive the world to extract knowledge from it but also actuate it, to create change that is beneficial to others. Of particular importance is a shared understanding of the physical space and task structure, throughout the interactions with other agents in the environment—a capability that is absent from disembodied dialogue interaction systems. Such elevated understanding gives the agent access to previously-inaccessible interaction scenarios, which can manifest in assistive technologies for those with physical disabilities or allow for the deployment of search-and-rescue systems in which a layperson can ask for help and explain their circumstances.

Despite the importance of collaboration for real-world deployment, current Embodied Vision-Language Planning tasks (Section 2.2) tend not to include significant social interaction. Even in the case of the Vision and Dialogue Navigation (VDN) and Embodied Goal-directed Manipulation (EGM) tasks, the dialogue interaction components are *static*, i.e., pre-generated datasets of dialogue histories are provided to the agent, during training and testing. Without tasks and simulators that enable more *dynamic* interaction with agents, it is challenging to develop approaches that manifest behaviour typical of online collaborative settings (Fried et al., 2021; Suhr et al., 2019), such as: (i) modeling other agent(s) conceptual schema of the world, (ii) incorporating direct feedback about task progress, (iii) and generating statements and questions (e.g., for informing other agent(s) about task progress, requesting assistance, resolving ambiguity in received instructions, asking for alternative commands, etc.) (Padmakumar et al., 2021).

In order for EVLP agents to perform complex tasks that require dynamic interaction with other agents, the underlying simulation environments and task structures must support the desired interaction settings. Primarily, environments must enable the representation of others’ anticipated actions, mental state, and previous behaviour, which has been shown to be critical in related areas, such as social navigation (Tsai and Oh, 2020; Vemula et al., 2018; Mavrogiannis et al., 2021), natural language processing (Fried et al., 2018a, 2021; Zhu et al., 2021a), and human-machine in-

teraction (Newman et al., 2018; Jeon et al., 2020; Charalampous et al., 2017). We highlight these related fields as inspiration for this new direction in EVLP research.

5.1.2 DYNAMIC ENVIRONMENTS

Most contemporary tasks and agents assume that the underlying environments are stationary and unchanging. Instead, we wish to highlight the notion of *dynamic* environments, which we define as those that change due to events both inside and outside of the ego-agent’s locus of control. The RobustNav task (Chattopadhyay et al., 2021) extended the RoboTHOR visual navigation environment (Deitke et al., 2020), to introduce structured perturbations on perception and transition dynamics, with the intention of encouraging agents to learn robust navigation policies in lieu of these corruptions. These perturbations are static with respect to the agent’s task-execution, however, and do not change overtime. Other tasks, such as ALFRED (Shridhar et al., 2020) and TEACH (Padmakumar et al., 2021), introduced entity-based interaction and unrecoverable states. However, these interventions can only be caused by the agent and cannot occur independently.

In real-world environments, the state of the environment can change without any intervention from the agent. Whether these changes should result from another agent (cognitive or otherwise) acting on the environment, from physical dynamics of the environment (e.g., a book sliding off a shelf), or from temporal dependencies on available interactions (e.g., interactions with certain objects are only available at certain times of day), the planning and reasoning complexities introduced by dynamic environments are challenging, yet largely unexplored, properties. We advocate for the introduction of these dynamic elements to EVLP tasks, in order to advance one step closer to enabling agents that accommodate the complexities of real-world deployment.

5.1.3 CROSS-TASK ROBOT LEARNING

EVLP tasks are currently evaluated separately from one another, which fails to capture the underlying skills shared between tasks and the overall progress made as a field. Being able to follow instructions (e.g., as in the VLN task family) and also answer questions about an environment (e.g., as in the EQA task family) are not and should not be treated as mutually-exclusive tasks. Rather than propose and evaluate tasks independently, tasks could be combined according to more unifying requirements on agent capability. Similar notions of cross-task learning gave rise to seminal modeling strategies (e.g., hierarchical task decomposition) and problem definitions (e.g., mobile manipulation, under EGM), which remain quite relevant to this day, despite their classical foundations (see Section 3.1.3). Combinations of existing task families may, likewise, yield new insights and advances.

Recent works have begun to test over several environments, but often within the same task family (Hao et al., 2020). Extending these types of benchmarks could allow for the use of transfer learning to bootstrap models. Moreover, it would allow for shared metrics that measure underlying skills of a particular agent. The idea that pre-trained models have specific abilities required for downstream tasks is not new and has been explored by the NLP (Talmor et al., 2020) and robotics communities (He et al., 2018). Extending this to EVLP tasks could prove beneficial to evaluating progress in the field as a whole rather than over very specific benchmarks. Zero-shot learning between tasks, with the idea that a model which can ground language should be as capable at manipulation as it is at navigation could be used to this end. Otherwise, specific benchmarks could also be designed to test for specific abilities, such as the ability to ground left and right, ability to recognize if an object

is invalid, ability to recognize objects. In this vein, Zhu et al. (2021b) test the impact of masking direction and object information in text on performance to evaluate how current VLN models use these types of information. They find that the type of information that led to the biggest decrease, either object or directional information, was dataset dependent. This finding means that shifts in information type in instruction could impact model performance.

Another test was conducted to evaluate the ability of models to ground text to vision by masking objects in the images found in the instructions. Doing so resulted in a limited decrease in performance, casting doubt on the ability of current models to ground. More works in this line could help understand and evaluate potential weaknesses in current models and how to improve their ability to generalize. We provide further discussion on the need for improved cross-task agent generalizability in Section 5.4.

5.2 Use of Domain Knowledge

Using domain knowledge to guide the learning progress of models has seen a recent resurgence in various domains, such as: autonomous driving (Park et al., 2020; Chen et al., 2019a; Li et al., 2020c), autonomous racing (Herman et al., 2021; Chen et al., 2021; Francis et al., 2022), robot navigation (Das et al., 2018b; Yang et al., 2019a), natural language processing (Andreas et al., 2016; Ravichander et al., 2017; Ma et al., 2019c, 2021; Oltramari* et al., 2020; Li et al., 2021b; Lu et al., 2020b), sensing and control (Jiang et al., 2018; Francis et al., 2019; Chen et al., 2020a; Munir et al., 2017), and many others. Indeed, domain knowledge comes in many forms, such as in graphical models, physics-based constraints, admissibility conditions, auxiliary objectives, knowledge distillation, pre-training steps, symbolic commonsense knowledge, and many others. While domain knowledge holds the promise of improving agents’ sample-efficiency, interpretability, safety, and generalizability, the challenge exists in how to effectively express and utilize the domain knowledge in an arbitrary problem. In this section, we highlight *pre-training* and *commonsense knowledge*, in particular, as these are two manifestations of domain knowledge that show incredible promise for endowing agents with the aforementioned attributes.

5.2.1 PRE-TRAINING

A common strategy for injecting agents with external domain knowledge is to borrow information from other tasks, models, and domains through generalized pre-training steps. Stemming from principles of non-convex optimization, the idea behind pre-training is that task-specific models and agents can start with a more optimal parameterization for a given task (for subsequent training or “fine-tuning”), after having already undergone partial training on other similar tasks. However, the selection of these pre-training tasks is not to be taken lightly: pre-training models on tasks that contain counterfactual or causally-confusing samples, with respect to those in the intended downstream task, can cause models to get stuck in sub-optimal local minima during the fine-tuning process. In response to these challenges, pre-training tasks have been carefully designed and coupled with popular high-capacity models, in such domains as image classification (He et al., 2016; Mogadala et al., 2019) and natural language processing (Devlin et al., 2019; Yang et al., 2019b; Ma et al., 2021), in attempts maximize the generalizability of transferred or fine-tuned approaches. While there is some progress in the context of specific multimodal problems (Majumdar et al., 2020; Hao et al., 2020; Lu et al., 2019), where the cross-modal reasoning is performed primarily at a single time-step, challenges remain for developing generalizable pre-training strategies that encompass the scope of

the broader EVLP task family, wherein cross-modal reasoning must occur over arbitrary temporal extents and alongside complex transition dynamics.

5.2.2 COMMONSENSE KNOWLEDGE

In the real world, much of the background information that governs the interactions amongst objects, cognitive agents, and the environment is *implicit*—that is to say that this knowledge is learned (by humans) through experience and is not commonly expressed explicitly, as we undertake our actions in the world. Tasks that require this tacit or “commonsense” knowledge³ are notoriously difficult for machines; even for tasks, in which models have recently enjoyed substantial performance improvements by other means, researchers find that commonsense knowledge still represents a fundamental cognitive gap in model- versus human-level performance, in scene understanding capability (Ma et al., 2019c, 2021; Oltramari* et al., 2020). Challenges arise, however, if the *type* of commonsense knowledge to be used (e.g., declarative, relational, procedural, sentiment, metaphorical, etc.) is not chosen to align well with the representation of that knowledge or with the implicit semantic characteristics of the downstream task.

Commonsense knowledge acquisition and injection in models remains an active research area in NLP (Talmor et al., 2019; Ma et al., 2019c, 2021; Li et al., 2021b), with some works proposing to ground observations with structured commonsense knowledge bases, directly, thereby improving downstream performance on relevant tasks. However, the use of commonsense knowledge, in the context of EVLP tasks remains largely unexplored. As the ultimate goal of EVLP tasks is to develop intelligent agents that are capable of solving real-world problems in realistic environments, it is reasonable to consider providing models with structured external knowledge from the world (Yang et al., 2019a). In general, EVLP tasks can be viewed as a series of tasks that test agents’ commonsense reasoning, where the agent is required to learn general skills that can be transferred to unseen contexts.

5.3 Agent Training Objectives

Selecting the appropriate training objective(s) for agents undertaking a given task has been a long-standing problem in machine learning and artificial intelligence; this selection depends on the nature of the available training signal(s) (reward/cost, full supervision, limited or non-existent supervision) and on the degree to which external knowledge (e.g., auxiliary objectives, constraints) is deemed necessary for effectively biasing agent behaviour. For EVLP tasks, specifically, the selection of training objectives is made more challenging by the complex nature of the environments, often necessitating frameworks that consist of more than one biasing strategy. For example, many state-of-the-art models in VLN use self-supervised learning or multi-task learning in a pre-training stage then supervised learning in a fine-tuning stage, in order to borrow information from large external datasets (Majumdar et al., 2020). Compelling approaches in the EQA task family often start with supervised (or imitation) learning to first learn from available expert traces, then use reinforcement learning in a fine-tuning stage, with the hope of finding a more generalisable model parameterisation for unseen environments (Das et al., 2018a; Wijmans et al., 2019a). Of the pop-

3. Common sense is broad and inherits several concepts from related domains, such as cognitive science and psychology. For a brief treatment of the various types of commonsense knowledge, in the context of neural prediction on AI tasks such as question answering in NLP, we refer interested readers to the works by Ma et al. (2019c, 2021).

ular learning paradigms in use, for a given task, it is challenging to define which permutation(s) yield the best downstream and cross-domain results. However, given the underlying motivation of optimising for generalisability and interpretability, explicit treatment should be given to finding the learning paradigm(s) that most effectively integrate information for various related sources and generalises agents’ inductive biases to new environments; the training paradigms should include explicit mechanisms (constraints, regularisation strategies) for encouraging these properties we hope to imbue.

5.4 Model Evaluation Paradigms

Datasets and simulation environments are the primary driving forces behind EVLP research, since one’s ability to measure model efficacy relies on the availability of strong testing scenarios and the appropriate evaluation criteria. Current EVLP tasks are implemented as a set of goals and metrics, atop pre-existing simulators or datasets. In this section, we urge the community to consider and prioritise the deployment of EVLP agents to real-world settings. As examples of potential task applications, we may imagine VLN models to be used for object-delivery and retrieval, EQA models for search and reconnaissance scenarios, VDN models for shared autonomy and human-machine collaboration, EOR models for scene understanding and surveillance, and EGM models for such applications as household robotics and manufacturing. Here, task metrics can have conflicting objectives, where they need to resemble measures of what “good” agent behaviour might look like in the real-world, despite being subject to the practical limitations of the dataset or simulator itself (e.g., physical fidelity, semantic saliency, phenomenological coverage). Specifically, we assert that various EVLP tasks and metrics may be improved on the basis of three dimensions: (i) simulator realness, (ii) dataset realness, and (iii) tests for model generalisability.

5.4.1 SIMULATOR REALNESS

Simulation-based training is intended to serve as a proxy for training agents in the real-world, particularly in situations where operating or fully-supervising real-world agents would be prohibitively expensive or impossible. Furthermore, simulation-based training and execution are especially attractive when modeling a sequential learning problem (i.e., in which the environment’s state and an agent’s observations are functions of the agent’s previous action), since offline datasets do not allow for such recursive interaction with an environment. There are limitations, however, in how effectively scientists and practitioners can encourage the desired model behaviour to emerge for real-world use-cases—chief among these limitations are those of the simulation environment itself as well as discrepancies in how task metrics are defined. Compared to the real-world, simulators have reduced physical fidelity (e.g., discretized navigation graphs, synthetic visual contexts) and limited semantic/phenomenological coverage (e.g., time-invariant weather and obstacles). As a result, a number of real-world physical challenges in robotics and computer vision (e.g., visual affordance learning, proprioceptive control, system identification) are abstracted away in simulated EVLP tasks. While this may allow for a concerted effort for pursuing other aspects of advanced reasoning in agents, agents trained in these environments risk overfitting to the incomplete transition dynamics provided by the simulated environment, which are only approximations of the *true* dynamics that govern corresponding natural generative processes (Alami et al., 1998). Furthermore, the task execution primitives (e.g., `pick`, `place`, `move`), which have become commonplace in simulated EVLP tasks as abstracted actions, may not always translate to the same abstractions

in the real-world, e.g., due to noisy execution (Chattopadhyay et al., 2021). Because this dissonance reduces models’ immediate viability for real-world deployment, we assert the importance of increased attention from the computer vision and robotics communities on the topics of simulation-to-real transfer, unseen generalisation, out-of-distribution prediction, and domain adaptation. In situations where task metrics are inappropriately defined (e.g., only end-goal success rate metrics are available, ignoring collisions and sharp movements), trained models are optimised according to incomplete objectives, which serve only to distinguish between entries on AI task leaderboards rather than fundamentally assess model readiness for real-world deployment. We encourage the definition of metrics that assess intermediate agent behaviours and task efficiency, as opposed to simply indicating task completion.

5.4.2 DATASET REALNESS

Dataset-based agent training thrives in many situations, including those that dictate modeling the space of all observations and agent actions as probability distributions (as in normalising flow-based approaches, variational Bayesian methods, etc.), when performing knowledge distillation with privileged information (as in learning from demonstrations), when warm-starting models before policy refinement (e.g., pre-training, data augmentation), or when introducing auxiliary objectives to the training procedure (e.g., learning skills, supervising stopping criteria).⁴ Despite these momentary advantages over simulation, physical fidelity and semantic coverage requirements are even more pronounced in dataset-based training, since, now, limited (or no) observation of the environment’s transition dynamics is available to agents as additional supervision. Indeed, rare-event coverage is particularly challenging for models in dataset-based training settings, due to the imbalanced support that leads to unrealistic priors and ill-prepared posteriors. Another limiting factor in the use of datasets for real-world model deployment is the adoption of unrealistic assumptions about the design or collection of data, e.g., the inclusion of programmatically-generated template instructions that would likely not be observed in the real-world. Next, the modalities provided by various datasets reflect only a subset of the sensory information that a cognitive agent may utilise for solving EVLP tasks in the real world. For example, few datasets have the notion of ambient environmental sounds (e.g., running sink faucets, vibrations from home appliances), despite this information being crucial to how humans interact with the world (Chen et al., 2020c). Moreover, dataset metrics vary in the agent behaviours they seek to encourage, i.e., by way of the model training objectives they necessitate (see Section 5.3). We want agent behaviour to manifest, free from the spurious variation in the visual observations, natural language, and expert traces; we want agent behaviour to remain invariant to irrelevant features (e.g., environment-specific scene backgrounds, linguistic synonyms and variance in sentence structure, etc.) and, accordingly, we encourage the emergence of EVLP datasets that enable agents to be simulator-agnostic. The same metrics may be defined across *multiple* environments, enabling “purer” skill affordances. Finally, special care should be given to the formats in which datasets are released and reported: characterising (i) the action space,

4. One may reasonably argue in favour of dispensing with the datasets and, instead, using a simulator (should one be available). However, we offer a reminder that using labeled datasets constitutes performing knowledge distillation with privileged information (Lopez-Paz et al., 2016), as in Imitation Learning (IL), which has provably faster convergence time and reduced model capacity requirements, compared to learning *without* privileged information. The challenges facing IL (Ross et al., 2011b; de Haan et al., 2019) result primarily from the lack of comprehensiveness and robustness in the expert (teacher) traces, not in the notion of using an expert in itself. Furthermore, IL approaches have become the *de facto* alternative to formulating EVLP tasks as POMDPs with sparse rewards.

(ii) instruction length (average and range), (iii) vocabulary size, (iv) data collection procedure, and (v) data availability (e.g., open-source or academic use permitted).

5.4.3 TESTS FOR GENERALISABILITY

EVLP datasets, such as those listed in Table 7 of this paper, have been conventionally separated across domain and task lines (Hermann et al., 2020; Chen et al., 2019b; Anderson et al., 2018c; Das et al., 2018a). In the context of real-world deployment, however, models are expected to encompass such properties as transferability to unseen environments and robustness to model and environmental uncertainty. Agents should be able to perform the same EVLP task, regardless of the nuances in the visual background; agents should be capable of driving in both urban and highway scenarios; agents trained on a relevant set of tasks should enjoy an advantage on similar unseen tasks. We assert that evaluation paradigms should assess agents using explicit tests for: generalisability across domains; generalisability across tasks. The need for these additional evaluation assessments is not new. As the ubiquity of deployed approaches increased while the challenges in robust multimodal alignment continued to abound, other communities, such as visual question answering (VQA), have faced similar changes in requirements. Agrawal et al. (2018) produced the VQA-CP dataset, an extension of the original VQA task (Antol et al., 2015) which contained multiple validation splits from different environments. While, in EVLP tasks, there is already the notion of seen and unseen environment splits (Anderson et al., 2018c; Das et al., 2018a; Gordon et al., 2018; Shridhar et al., 2020), the multiple dimensions of possible variation in the unseen split (distribution of objects, distribution of backgrounds, distribution of indoor layouts) remain largely unstudied; there are no guarantees that these splits are representative of the types of desired real-world scenarios. Having a test bed with different test splits would be beneficial, as it would allow researchers to verify if their models are learning spurious (or non-transferable) correlations in the observations or if the models are actually learning to plan. The authors also advocate for the assessment of agents’ abilities to generalise to longer path lengths and to other tasks. The EVLP community already has several datasets that would facilitate this analysis, which broadly fall into two general categories: path concatenation and path decomposition. Path concatenation, such as R4R, R6R, R8R (Jain et al., 2019; Zhu et al., 2020b) works by joining paths that start and terminate near one another to generate longer paths. Agents can be trained on these longer paths or evaluated over them (Jain et al., 2019). Rather than building longer paths from the same dataset, path decomposition breaks down the path into fine-grained instructions (Li et al., 2020b; Zhu et al., 2020b). The agent is trained over those and then evaluated over the larger dataset with longer instructions. Note that path concatenation and decomposition are not mutually exclusive (Zhu et al., 2020b). Extending these techniques to other datasets or developing new techniques would be one way of evaluating how length and complexity impact performance.

6. Conclusion

As a field in development, EVLP has never been fully documented. In this paper we proposed a taxonomy which covers: tasks in EVLP, learning paradigms and training techniques used, evaluations, open challenges. We provided a framework to discuss existing and future tasks based on the skills required to solve them. We discussed usage of training paradigms such as reinforcement learning and supervised learning, common training tricks such as data augmentation, and commonly

used loss functions. Evaluation of models and a framework to discuss what aspects of performance each of these metrics covers. Finally the challenges currently being tackled in the field and what challenges still remain to be tackled. Specifically we discuss issues that could prevent real world deployment, such as a lack of generalization, robustness, simulator realness, and lack of interactions. With the speed of improvements in the field we feel that these challenges could be tackled and allow for real world deployment of these EVLP agents.

Acknowledgements

The authors thank Alessandro Oltramari, Yonatan Bisk, Eric Nyberg, and Louis-Philippe Morency for insightful discussions; we thank Mayank Mali for support throughout the editing process, and we thank the JAIR reviewers for their valuable feedback. This work was supported, in part, by a doctoral research fellowship from Bosch Research, by the U.S. Air Force Office of Scientific Research, under award number FA2386-17-1-4660, and by Jonathan’s dual affiliation with the Human-Machine Collaboration group at Bosch Research Pittsburgh. The views expressed in this article do not necessarily represent those of the aforementioned entities. All authors contributed equally to the manuscript; author list in alphabetical order. Professor Jean Oh is the corresponding author, reachable at jeanoh@cmu.edu.

References

- Agrawal, A., Batra, D., Parikh, D., and Kembhavi, A. (2018). Don’t just assume; look and answer: Overcoming priors for visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4971–4980. IEEE Computer Society.
- Akhtar, M. S., Chauhan, D., Ghosal, D., Poria, S., Ekbal, A., and Bhattacharyya, P. (2019). Multi-task learning for multi-modal emotion recognition and sentiment analysis. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 370–379, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alami, R., Chatila, R., Fleury, S., Ghallab, M., and Ingrand, F. (1998). An architecture for autonomy. *The International Journal of Robotics Research*, 17(4):315–337.
- Anderson, P., Chang, A. X., Chaplot, D. S., Dosovitskiy, A., Gupta, S., Koltun, V., Kosecka, J., Malik, J., Mottaghi, R., Savva, M., and Zamir, A. R. (2018a). On evaluation of embodied navigation agents. *CoRR*, abs/1807.06757.
- Anderson, P., Fernando, B., Johnson, M., and Gould, S. (2016). Spice: Semantic propositional image caption evaluation.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., and Zhang, L. (2018b). Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. IEEE Computer Society.

- Anderson, P., Shrivastava, A., Parikh, D., Batra, D., and Lee, S. (2019). Chasing ghosts: Instruction following as bayesian state tracking. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E. B., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 369–379.
- Anderson, P., Shrivastava, A., Truong, J., Majumdar, A., Parikh, D., Batra, D., and Lee, S. (2020). Sim-to-real transfer for vision-and-language navigation. *CoRR*, abs/2011.03807.
- Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I. D., Gould, S., and van den Hengel, A. (2018c). Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3674–3683. IEEE Computer Society.
- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. (2016). Neural module networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 39–48. IEEE Computer Society.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., and Parikh, D. (2015). VQA: visual question answering. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2425–2433. IEEE Computer Society.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Baltrušaitis, T., Ahuja, C., and Morency, L. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443.
- Batra, D., Chang, A. X., Chernova, S., Davison, A. J., Deng, J., Koltun, V., Levine, S., Malik, J., Mordatch, I., Mottaghi, R., et al. (2020). Rearrangement: A challenge for embodied ai. *arXiv preprint arXiv:2011.01975*.
- Beattie, C., Leibo, J. Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., Schrittwieser, J., Anderson, K., York, S., Cant, M., Cain, A., Bolton, A., Gaffney, S., King, H., Hassabis, D., Legg, S., and Petersen, S. (2016). Deepmind lab. *CoRR*, abs/1612.03801.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. (2015). The arcade learning environment: An evaluation platform for general agents (extended abstract). In Yang, Q. and Wooldridge, M. J., editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 4148–4152. AAAI Press.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). Curriculum learning. In Danyluk, A. P., Bottou, L., and Littman, M. L., editors, *Proceedings of the 26th Annual International*

- Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, volume 382 of *ACM International Conference Proceeding Series*, pages 41–48. ACM.
- Bisk, Y., Holtzman, A., Thomason, J., Andreas, J., Bengio, Y., Chai, J., Lapata, M., Lazaridou, A., May, J., Nisnevich, A., Pinto, N., and Turian, J. (2020). Experience grounds language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.
- Black, P. E. et al. (1998). Dictionary of algorithms and data structures.
- Blukis, V., Brukhim, N., Bennett, A., Knepper, R. A., and Artzi, Y. (2018a). Following high-level navigation instructions on a simulated quadcopter with imitation learning. In Kress-Gazit, H., Srinivasa, S. S., Howard, T., and Atanasov, N., editors, *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*.
- Blukis, V., Misra, D., Knepper, R. A., and Artzi, Y. (2018b). Mapping navigation instructions to continuous control actions with position-visitation prediction. In *Conference on Robot Learning*, pages 505–518. PMLR.
- Blukis, V., Terme, Y., Niklasson, E., Knepper, R. A., and Artzi, Y. (2019). Learning to map natural language instructions to physical quadcopter control using simulated flight. In *Conference on Robot Learning (CoRL)*.
- Boularias, A., Duvallet, F., Oh, J., and Stentz, A. (2015). Grounding spatial relations for outdoor robot navigation. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1976–1982. IEEE.
- Brodeur, S., Perez, E., Anand, A., Golemo, F., Celotti, L., Strub, F., Rouat, J., Larochelle, H., and Courville, A. C. (2018). Home: a household multimodal environment. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Burgard, W., Moors, M., Stachniss, C., and Schneider, F. E. (2005). Coordinated multi-robot exploration. *IEEE Transactions on robotics*, 21(3):376–386.
- Cambon, S., Alami, R., and Gravot, F. (2009). A hybrid approach to intricate motion, manipulation and task planning. *The International Journal of Robotics Research*, 28(1):104–126.
- Castelfranchi, C. (1998). Modelling social action for ai agents. *Artificial intelligence*, 103(1-2):157–182.

- Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., and Zhang, Y. (2017). Matterport3d: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)*.
- Charalampous, K., Kostavelis, I., and Gasteratos, A. (2017). Recent trends in social aware robot navigation: A survey. *Robotics and Autonomous Systems*, 93:85–104.
- Chattopadhyay, P., Hoffman, J., Mottaghi, R., and Kembhavi, A. (2021). Robustnav: Towards benchmarking robustness in embodied navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15691–15700.
- Chen, B., Francis, J., Oh, J., Nyberg, E., and Herbert, S. L. (2021). Safe autonomous racing via approximate reachability on ego-vision.
- Chen, B., Francis, J., Pritoni, M., Kar, S., and Bergés, M. (2020a). Cohort: Coordination of heterogeneous thermostatically controlled loads for demand flexibility. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*.
- Chen, C., Jain, U., Schissler, C., Gari, S. V. A., Al-Halah, Z., Ithapu, V. K., Robinson, P., and Grauman, K. (2020b). Soundspaces: Audio-visual navigation in 3d environments. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J., editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VI*, volume 12351 of *Lecture Notes in Computer Science*, pages 17–36. Springer.
- Chen, C., Jain, U., Schissler, C., Gari, S. V. A., Al-Halah, Z., Ithapu, V. K., Robinson, P., and Grauman, K. (2020c). Soundspaces: Audio-visual navigation in 3d environments. In *ECCV*.
- Chen, C., Majumder, S., Al-Halah, Z., Gao, R., Ramakrishnan, S. K., and Grauman, K. (2020d). Learning to set waypoints for audio-visual navigation. *arXiv preprint arXiv:2008.09622*.
- Chen, D., Zhou, B., Koltun, V., and Krähenbühl, P. (2019a). Learning by cheating.
- Chen, H., Suhr, A., Misra, D., Snavely, N., and Artzi, Y. (2019b). TOUCHDOWN: natural language navigation and spatial reasoning in visual street environments. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12538–12547. Computer Vision Foundation / IEEE.
- Chi, T., Shen, M., Eric, M., Kim, S., and Hakkani-Tür, D. (2020). Just ask: An interactive learning framework for vision and language navigation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2459–2466. AAAI Press.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.

- Choi, J. and Amir, E. (2009). Combining planning and motion planning. In *2009 IEEE International Conference on Robotics and Automation*, pages 238–244. IEEE.
- Cummins, M. and Newman, P. (2008). Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 27(6):647–665.
- Das, A., Carnevale, F., Merzic, H., Rimell, L., Schneider, R., Abramson, J., Hung, A., Ahuja, A., Clark, S., Wayne, G., and Hill, F. (2020). Probing emergent semantics in predictive agents via question answering. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 2376–2391. PMLR.
- Das, A., Datta, S., Gkioxari, G., Lee, S., Parikh, D., and Batra, D. (2018a). Embodied question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 1–10. IEEE Computer Society.
- Das, A., Gkioxari, G., Lee, S., Parikh, D., and Batra, D. (2018b). Neural modular control for embodied question answering. *arXiv preprint arXiv:1810.11181*.
- Dasari, S. and Gupta, A. (2020). Transformers for one-shot visual imitation.
- de Haan, P., Jayaraman, D., and Levine, S. (2019). Causal confusion in imitation learning. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11693–11704.
- de Vries, H., Shuster, K., Batra, D., Parikh, D., Weston, J., and Kiela, D. (2018). Talk the walk: Navigating new york city through grounded dialogue.
- Deitke, M., Han, W., Herrasti, A., Kembhavi, A., Kolve, E., Mottaghi, R., Salvador, J., Schwenk, D., VanderBilt, E., Wallingford, M., Weihs, L., Yatskar, M., and Farhadi, A. (2020). Rethor: An open simulation-to-real embodied AI platform. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 3161–3171. IEEE.
- Deng, Z., Narasimhan, K., and Russakovsky, O. (2020). Evolving graphical planner: Contextual global planning for vision-and-language navigation. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Duvallet, F. (2015). Natural language direction following for robots in unstructured unknown environments. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA ROBOTICS INST.
- Duvallet, F., Walter, M. R., Howard, T., Hemachandra, S., Oh, J., Teller, S., Roy, N., and Stentz, A. (2016). Inferring maps and behaviors from natural language instructions. In *Experimental Robotics*, pages 373–388. Springer.
- Ehsani, K., Han, W., Herrasti, A., VanderBilt, E., Weihs, L., Kolve, E., Kembhavi, A., and Mottaghi, R. (2021). Manipulathor: A framework for visual object manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 4497–4506. Computer Vision Foundation / IEEE.
- Everingham, M., Gool, L. V., Williams, C. K. I., Winn, J. M., and Zisserman, A. (2010). The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338.
- Fang, K., Toshev, A., Li, F., and Savarese, S. (2019). Scene memory transformer for embodied agents in long-horizon tasks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 538–547. Computer Vision Foundation / IEEE.
- Filliat, D. and Meyer, J.-A. (2003). Map-based navigation in mobile robots:: I. a review of localization strategies. *Cognitive systems research*, 4(4):243–282.
- Francis, J., Chen, B., Ganju, S., Kathpal, S., Poonganam, J., Shivani, A., Vyas, V., Genc, S., Zhukov, I., Kumskey, M., Oh, J., Nyberg, E., and Herbert, S. L. (2022). Learn-to-race challenge 2022: Benchmarking safe learning and cross-domain generalisation in autonomous racing.
- Francis, J., Quintana, M., von Frankenberg, N., Munir, S., and Berges, M. (2019). Occutherm: Occupant thermal comfort inference using body shape information. In *Proceedings of the 6th International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '19, New York, NY, USA*. ACM.
- Fried, D., Andreas, J., and Klein, D. (2018a). Unified pragmatic models for generating and following instructions. In *Proceedings of NAACL-HLT*, pages 1951–1963.
- Fried, D., Chiu, J. T., and Klein, D. (2021). Reference-centric models for grounded collaborative dialogue. *arXiv preprint arXiv:2109.05042*.
- Fried, D., Hu, R., Cirik, V., Rohrbach, A., Andreas, J., Morency, L., Berg-Kirkpatrick, T., Saenko, K., Klein, D., and Darrell, T. (2018b). Speaker-follower models for vision-and-language navigation. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3318–3329.
- Fu, T.-J., Wang, X. E., Peterson, M. F., Grafton, S. T., Eckstein, M. P., and Wang, W. Y. (2020). Counterfactual vision-and-language navigation via adversarial path sampler. In *European Conference on Computer Vision*, pages 71–86. Springer.

- Gan, C., Schwartz, J., Alter, S., Schrimpf, M., Traer, J., Freitas, J. D., Kubilius, J., Bhandwaldar, A., Haber, N., Sano, M., Kim, K., Wang, E., Mrowca, D., Lingelbach, M., Curtis, A., Feigelis, K. T., Bear, D. M., Gutfreund, D., Cox, D. D., DiCarlo, J. J., McDermott, J. H., Tenenbaum, J. B., and Yamins, D. L. K. (2020). Threedworld: A platform for interactive multi-modal physical simulation. *CoRR*, abs/2007.04954.
- Gan, C., Zhao, H., Chen, P., Cox, D. D., and Torralba, A. (2019). Self-supervised moving vehicle tracking with stereo sound. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7052–7061. IEEE.
- Garrett, C. R., Chitnis, R., Holladay, R., Kim, B., Silver, T., Kaelbling, L. P., and Lozano-Pérez, T. (2021). Integrated task and motion planning. *Annual review of control, robotics, and autonomous systems*, 4:265–293.
- Gordon, D., Kembhavi, A., Rastegari, M., Redmon, J., Fox, D., and Farhadi, A. (2018). IQA: visual question answering in interactive environments. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4089–4098. IEEE Computer Society.
- Hao, W., Li, C., Li, X., Carin, L., and Gao, J. (2020). Towards learning a generic agent for vision-and-language navigation via pre-training. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 13134–13143. IEEE.
- Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R. B. (2017). Mask R-CNN. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2980–2988. IEEE Computer Society.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- He, Z., Julian, R., Heiden, E., Zhang, H., Schaal, S., Lim, J. J., Sukhatme, G., and Hausman, K. (2018). Zero-shot skill composition and simulation-to-real transfer by learning task representations. *arXiv preprint arXiv:1810.02422*.
- Herman, J., Francis, J., Ganju, S., Chen, B., Koul, A., Gupta, A., Skabelkin, A., Zhukov, I., Kumskey, M., and Nyberg, E. (2021). Learn-to-race: A multimodal control environment for autonomous racing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9793–9802.
- Hermann, K. M., Malinowski, M., Mirowski, P., Banki-Horvath, A., Anderson, K., and Hadsell, R. (2020). Learning to follow directions in street view. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in*

- Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 11773–11781. AAAI Press.
- Hester, T., Vecerík, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Horgan, D., Quan, J., Sendonaris, A., Osband, I., Dulac-Arnold, G., Agapiou, J. P., Leibo, J. Z., and Gruslys, A. (2018). Deep q-learning from demonstrations. In McIlraith, S. A. and Weinberger, K. Q., editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3223–3230. AAAI Press.
- Hinton, G., Srivastava, N., and Swersky, K. (2012). Neural networks for machine learning lecture 6a overview of mini-batch gradient descent.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. (2019). The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Hong, Y., Rodriguez, C., Wu, Q., and Gould, S. (2020). Sub-instruction aware vision-and-language navigation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3360–3376, Online. Association for Computational Linguistics.
- Hu, Y., Wang, W., Jia, H., Wang, Y., Chen, Y., Hao, J., Wu, F., and Fan, C. (2020). Learning to utilize shaping rewards: A new approach of reward shaping. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Huang, H., Jain, V., Mehta, H., Baldrige, J., and Ie, E. (2019). Multi-modal discriminative model for vision-and-language navigation. In *Proceedings of the Combined Workshop on Spatial Language Understanding (SpLU) and Grounded Communication for Robotics (RoboNLP)*, pages 40–49, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ilharco, G., Jain, V., Ku, A., Ie, E., and Baldrige, J. (2019). General evaluation for instruction conditioned navigation using dynamic time warping. *CoRR*, abs/1907.05446.
- Irshad, M. Z., Ma, C.-Y., and Kira, Z. (2021). Hierarchical cross-modal agent for robotics vision-and-language navigation. *arXiv preprint arXiv:2104.10674*.
- Jain, V., Magalhaes, G., Ku, A., Vaswani, A., Ie, E., and Baldrige, J. (2019). Stay on the path: Instruction fidelity in vision-and-language navigation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1862–1872, Florence, Italy. Association for Computational Linguistics.
- Jansen, P. (2020). Visually-grounded planning without vision: Language models infer detailed plans from high-level instructions. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4412–4417, Online. Association for Computational Linguistics.

- Jeon, H. J., Losey, D. P., and Sadigh, D. (2020). Shared autonomy with learned latent actions. *arXiv preprint arXiv:2005.03210*.
- Jiang, Z., Francis, J., Sahu, A. K., Munir, S., Shelton, C., Rowe, A., and Bergés, M. (2018). Data-driven thermal model inference with armax, in smart environments, based on normalized mutual information. In *2018 Annual American Control Conference (ACC)*, pages 4634–4639.
- Johnson, M., Hofmann, K., Hutton, T., and Bignell, D. (2016). The malmo platform for artificial intelligence experimentation. In Kambhampati, S., editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 4246–4247. IJCAI/AAAI Press.
- Kaelbling, L. P. and Lozano-Pérez, T. (2010). Hierarchical planning in the now. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- Kaelbling, L. P. and Lozano-Pérez, T. (2013). Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9-10):1194–1227.
- Ke, L., Li, X., Bisk, Y., Holtzman, A., Gan, Z., Liu, J., Gao, J., Choi, Y., and Srinivasa, S. S. (2019). Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6741–6749. Computer Vision Foundation / IEEE.
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Jaskowski, W. (2016). Vizdoom: A doom-based AI research platform for visual reinforcement learning. In *IEEE Conference on Computational Intelligence and Games, CIG 2016, Santorini, Greece, September 20-23, 2016*, pages 1–8. IEEE.
- Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., and Shah, M. (2021). Transformers in vision: A survey. *CoRR*, abs/2101.01169.
- Khatib, O. (1999). Mobile manipulation: The robotic assistant. *Robotics and Autonomous Systems*, 26(2-3):175–183.
- Kim, H., Zala, A., Burri, G., Tan, H., and Bansal, M. (2020). ArraMon: A joint navigation-assembly instruction interpretation task in dynamic environments. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3910–3927, Online. Association for Computational Linguistics.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kolve, E., Mottaghi, R., Gordon, D., Zhu, Y., Gupta, A., and Farhadi, A. (2017). AI2-THOR: an interactive 3d environment for visual AI. *CoRR*, abs/1712.05474.
- Krantz, J., Gokaslan, A., Batra, D., Lee, S., and Maksymets, O. (2021). Waypoint models for instruction-guided navigation in continuous environments. *CoRR*, abs/2110.02207.

- Krantz, J., Wijmans, E., Majumdar, A., Batra, D., and Lee, S. (2020). Beyond the nav-graph: Vision-and-language navigation in continuous environments. *CoRR*, abs/2004.02857.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L., Shamma, D. A., Bernstein, M. S., and Li, F. (2016). Visual genome: Connecting language and vision using crowdsourced dense image annotations. *CoRR*, abs/1602.07332.
- Ku, A., Anderson, P., Patel, R., Ie, E., and Baldridge, J. (2020). Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4392–4412, Online. Association for Computational Linguistics.
- Ku, A., Anderson, P., Pont-Tuset, J., and Baldridge, J. (2021). Pangea: The panoramic graph environment annotation toolkit. *CoRR*, abs/2103.12703.
- Lamb, A., Goyal, A., Zhang, Y., Zhang, S., Courville, A. C., and Bengio, Y. (2016). Professor forcing: A new algorithm for training recurrent networks. *CoRR*, abs/1610.09038.
- Li, C., Xia, F., Martín-Martín, R., Lingelbach, M., Srivastava, S., Shen, B., Vainio, K., Gokmen, C., Dharan, G., Jain, T., Kurenkov, A., Liu, C. K., Gweon, H., Wu, J., Fei-Fei, L., and Savarese, S. (2021a). igibson 2.0: Object-centric simulation for robot learning of everyday household tasks. *CoRR*, abs/2108.03272.
- Li, C., Xia, F., Martin-Martin, R., and Savarese, S. (2020a). Hrl4in: Hierarchical reinforcement learning for interactive navigation with mobile manipulators. In *Conference on Robot Learning*, pages 603–616. PMLR.
- Li, D., Opazo, C. R., Yu, X., and Li, H. (2020b). Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 1448–1458. IEEE.
- Li, J., Yang, F., Tomizuka, M., and Choi, C. (2020c). Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Li, X., Li, C., Xia, Q., Bisk, Y., Celikyilmaz, A., Gao, J., Smith, N. A., and Choi, Y. (2019). Robust navigation with language pretraining and stochastic sampling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1494–1499, Hong Kong, China. Association for Computational Linguistics.
- Li, Y., Goel, P., Rajendra, V. K., Singh, H. S., Francis, J., Ma, K., Nyberg, E., and Oltramari, A. (2021b). Lexically-constrained text generation through commonsense knowledge extraction and injection. In *Common Sense Knowledge Graphs at the 35th AAAI Conference on Artificial Intelligence (CSKGs@AAAI-21)*.
- Lin, C.-Y. (2004). ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In Fleet, D., Pajdla, T., Schiele, B., and Tuytelaars, T., editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Lipton, Z. C., Berkowitz, J., and Elkan, C. (2015). A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*.
- Lopez-Paz, D., Bottou, L., Schölkopf, B., and Vapnik, V. (2016). Unifying distillation and privileged information. In Bengio, Y. and LeCun, Y., editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Lu, J., Batra, D., Parikh, D., and Lee, S. (2019). Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23.
- Lu, J., Goswami, V., Rohrbach, M., Parikh, D., and Lee, S. (2020a). 12-in-1: Multi-task vision and language representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10434–10443. IEEE.
- Lu, X., West, P., Zellers, R., Bras, R. L., Bhagavatula, C., and Choi, Y. (2020b). Neurologic decoding: (un)supervised neural text generation with predicate logic constraints.
- Luketina, J., Nardelli, N., Farquhar, G., Foerster, J. N., Andreas, J., Grefenstette, E., Whiteson, S., and Rocktäschel, T. (2019). A survey of reinforcement learning informed by natural language. In Kraus, S., editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 6309–6317. ijcai.org.
- Ma, C., Lu, J., Wu, Z., AlRegib, G., Kira, Z., Socher, R., and Xiong, C. (2019a). Self-monitoring navigation agent via auxiliary progress estimation. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Ma, C., Wu, Z., AlRegib, G., Xiong, C., and Kira, Z. (2019b). The regretful agent: Heuristic-aided navigation through progress estimation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6732–6740. Computer Vision Foundation / IEEE.
- Ma, K., Francis, J., Lu, Q., Nyberg, E., and Oltramari, A. (2019c). Towards generalizable neuro-symbolic systems for commonsense question answering. In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 22–32, Hong Kong, China. Association for Computational Linguistics.
- Ma, K., Ilievski, F., Francis, J., Bisk, Y., Nyberg, E., and Oltramari, A. (2021). Knowledge-driven data construction for zero-shot evaluation in commonsense question answering. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-21)*.

- Majumdar, A., Shrivastava, A., Lee, S., Anderson, P., Parikh, D., and Batra, D. (2020). Improving vision-and-language navigation with image-text pairs from the web. *CoRR*, abs/2004.14973.
- Mavrogiannis, C., Baldini, F., Wang, A., Zhao, D., Trautman, P., Steinfeld, A., and Oh, J. (2021). Core challenges of social robot navigation: A survey. *arXiv preprint arXiv:2103.05668*.
- Mehta, H., Artzi, Y., Baldridge, J., Ie, E., and Mirowski, P. (2020). Retouchdown: Adding touch-down to streetlearn as a shareable resource for language grounding tasks in street view.
- Mirowski, P., Grimes, M. K., Malinowski, M., Hermann, K. M., Anderson, K., Teplyashin, D., Simonyan, K., Kavukcuoglu, K., Zisserman, A., and Hadsell, R. (2018). Learning to navigate in cities without a map. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2424–2435.
- Misra, D., Bennett, A., Blukis, V., Niklasson, E., Shatkhin, M., and Artzi, Y. (2018). Mapping instructions to actions in 3D environments with visual goal prediction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2667–2678, Brussels, Belgium. Association for Computational Linguistics.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In Balcan, M. and Weinberger, K. Q., editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1928–1937. JMLR.org.
- Mogadala, A., Kalimuthu, M., and Klakow, D. (2019). Trends in integration of vision and language research: A survey of tasks, datasets, and methods. *CoRR*, abs/1907.09358.
- Munir, S., Arora, R. S., Hesling, C., Li, J., Francis, J., Shelton, C., Martin, C., Rowe, A., and Berges, M. (2017). Real-time fine grained occupancy estimation using depth sensors on arm embedded platforms. In *2017 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 295–306. IEEE.
- Nachum, O., Gu, S., Lee, H., and Levine, S. (2018). Data-efficient hierarchical reinforcement learning. *arXiv preprint arXiv:1805.08296*.
- Nau, D. S., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., and Yaman, F. (2003). Shop2: An htn planning system. *Journal of artificial intelligence research*, 20:379–404.
- Newman, B. A., Aronson, R. M., Srinivasa, S. S., Kitani, K., and Admoni, H. (2018). Harmonic: A multimodal dataset of assistive human-robot collaboration. *arXiv preprint arXiv:1807.11154*.
- Nguyen, D. and Okatani, T. (2019). Multi-task learning of hierarchical vision-language representation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 10492–10501. Computer Vision Foundation / IEEE.

- Nguyen, K. and Daumé III, H. (2019). Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 684–695, Hong Kong, China. Association for Computational Linguistics.
- Nguyen, K., Dey, D., Brockett, C., and Dolan, B. (2019a). Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12527–12537.
- Nguyen, K., Dey, D., Brockett, C., and Dolan, B. (2019b). Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12527–12537. Computer Vision Foundation / IEEE.
- Nichol, A., Pfau, V., Hesse, C., Klimov, O., and Schulman, J. (2018). Gotta learn fast: A new benchmark for generalization in rl. *arXiv preprint arXiv:1804.03720*.
- Nilsson, N. J. et al. (1984). Shakey the robot.
- Oh, J. H., Suppé, A., Duvallet, F., Boularias, A., Navarro-Serment, L. E., Hebert, M., Stentz, A., Vinokurov, J., Romero, O. J., Lebiere, C., and Dean, R. (2015). Toward mobile robots reasoning like humans. In Bonet, B. and Koenig, S., editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, pages 1371–1379. AAAI Press.
- Oltamari*, A., Francis*, J., Henson, C., Ma, K., and Wickramarachchi, R. (2020). Neuro-symbolic architectures for context understanding.
- Padilla, R., Netto, S. L., and da Silva, E. A. B. (2020). A survey on performance metrics for object-detection algorithms. In *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*, pages 237–242.
- Padmakumar, A., Thomason, J., Shrivastava, A., Lange, P., Narayan-Chen, A., Gella, S., Piramuthu, R., Tur, G., and Hakkani-Tur, D. (2021). Teach: Task-driven embodied agents that chat. *arXiv preprint arXiv:2110.00534*.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Park, S. H., Lee, G., Bhat, M., Seo, J., Kang, M., Francis, J., Jadhav, A. R., Liang, P. P., and Morency, L.-P. (2020). Diverse and admissible trajectory forecasting through multimodal context understanding. In *European Conference on Computer Vision (ECCV)*.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.

- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Puig, X., Ra, K., Boben, M., Li, J., Wang, T., Fidler, S., and Torralba, A. (2018). Virtualhome: Simulating household activities via programs. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 8494–8502. IEEE Computer Society.
- Qi, Y., Wu, Q., Anderson, P., Wang, X., Wang, W. Y., Shen, C., and van den Hengel, A. (2020). REVERIE: remote embodied visual referring expression in real indoor environments. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9979–9988. IEEE.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language models are unsupervised multitask learners.
- Ravichander, A., Manzini, T., Grabmair, M., Neubig, G., Francis, J., and Nyberg, E. (2017). How would you say it? eliciting lexically diverse dialogue for supervised semantic parsing. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 374–383, Saarbrücken, Germany. Association for Computational Linguistics.
- Rawat, W. and Wang, Z. (2017). Deep convolutional neural networks for image classification: A comprehensive review. *Neural Computation*, 29(9):2352–2449. PMID: 28599112.
- Reddy, S., Dragan, A. D., and Levine, S. (2020). SQIL: imitation learning via reinforcement learning with sparse rewards. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ross, S., Gordon, G. J., and Bagnell, D. (2011a). A reduction of imitation learning and structured prediction to no-regret online learning. In Gordon, G. J., Dunson, D. B., and Dudík, M., editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011*, volume 15 of *JMLR Proceedings*, pages 627–635. JMLR.org.
- Ross, S., Gordon, G. J., and Bagnell, J. A. (2011b). A reduction of imitation learning and structured prediction to no-regret online learning.
- Sacerdoti, E. D. (1974). Planning in a hierarchy of abstraction spaces. *Artificial intelligence*, 5(2):115–135.

- Savva, M., Chang, A. X., Dosovitskiy, A., Funkhouser, T. A., and Koltun, V. (2017). MINOS: multimodal indoor simulator for navigation in complex environments. *CoRR*, abs/1712.03931.
- Savva, M., Malik, J., Parikh, D., Batra, D., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., and Koltun, V. (2019). Habitat: A platform for embodied AI research. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 9338–9346. IEEE.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *CoRR*, abs/1707.06347.
- Sharma, P., Ding, N., Goodman, S., and Soricut, R. (2018). Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.
- Sharma, P., Torralba, A., and Andreas, J. (2021). Skill induction and planning with latent language. *arXiv preprint arXiv:2110.01517*.
- Shen, B., Xia, F., Li, C., Martín-Martín, R., Fan, L., Wang, G., Pérez-D’Arpino, C., Buch, S., Srivastava, S., Tchapmi, L., Tchapmi, M., Vainio, K., Wong, J., Fei-Fei, L., and Savarese, S. (2021a). igibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic, September 27 - Oct. 1, 2021*, pages 7520–7527. IEEE.
- Shen, S., Li, L. H., Tan, H., Bansal, M., Rohrbach, A., Chang, K., Yao, Z., and Keutzer, K. (2021b). How much can CLIP benefit vision-and-language tasks? *CoRR*, abs/2107.06383.
- Shin, A., Ishii, M., and Narihira, T. (2021). Perspectives and prospects on transformer architecture for cross-modal tasks with language and vision.
- Shorten, C. and Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60.
- Shridhar, M., Thomason, J., Gordon, D., Bisk, Y., Han, W., Mottaghi, R., Zettlemoyer, L., and Fox, D. (2020). ALFRED: A benchmark for interpreting grounded instructions for everyday tasks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10737–10746. IEEE.
- Stentz, A. et al. (1995). The focussed d^* algorithm for real-time replanning. In *IJCAI*, volume 95, pages 1652–1659.
- Suhr, A., Yan, C., Schluger, J., Yu, S., Khader, H., Mouallem, M., Zhang, I., and Artzi, Y. (2019). Executing instructions in situated collaborative interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2119–2130, Hong Kong, China. Association for Computational Linguistics.

- Sun, F., Chang, Y., Wu, Y., and Lin, S. (2018). Designing non-greedy reinforcement learning agents with diminishing reward shaping. In Furman, J., Marchant, G. E., Price, H., and Rossi, F., editors, *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES 2018, New Orleans, LA, USA, February 02-03, 2018*, pages 297–302. ACM.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Syed, U. and Schapire, R. E. (2007). A game-theoretic approach to apprenticeship learning. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T., editors, *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, pages 1449–1456. Curran Associates, Inc.
- Synnaeve, G., Nardelli, N., Auvolat, A., Chintala, S., Lacroix, T., Lin, Z., Richoux, F., and Usunier, N. (2016). Torchcraft: a library for machine learning research on real-time strategy games. *CoRR*, abs/1611.00625.
- Szot, A., Clegg, A., Undersander, E., Wijmans, E., Zhao, Y., Turner, J., Maestre, N., Mukadam, M., Chaplot, D. S., Maksymets, O., et al. (2021). Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in Neural Information Processing Systems*, 34.
- Talmor, A., Elazar, Y., Goldberg, Y., and Berant, J. (2020). olympics - on what language model pre-training captures. *Trans. Assoc. Comput. Linguistics*, 8:743–758.
- Talmor, A., Herzig, J., Lourie, N., and Berant, J. (2019). CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tan, H. and Bansal, M. (2019). LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Tan, H., Yu, L., and Bansal, M. (2019). Learning to navigate unseen environments: Back translation with environmental dropout. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2610–2621, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tangiuchi, T., Mochihashi, D., Nagai, T., Uchida, S., Inoue, N., Kobayashi, I., Nakamura, T., Hagiwara, Y., Iwahashi, N., and Inamura, T. (2019). Survey on frontiers of language and robotics. *Advanced Robotics*, 33(15-16):700–730.
- Tanner, M. A. (2012). *Tools for statistical inference: observed data and data augmentation methods*, volume 67. Springer Science & Business Media.

- Tellex, S., Kollar, T., Dickerson, S., Walter, M. R., Banerjee, A. G., Teller, S. J., and Roy, N. (2011). Understanding natural language commands for robotic navigation and mobile manipulation. In Burgard, W. and Roth, D., editors, *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*. AAAI Press.
- Thomason, J., Gordon, D., and Bisk, Y. (2019a). Shifting the baseline: Single modality performance on visual navigation & QA. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1977–1983, Minneapolis, Minnesota. Association for Computational Linguistics.
- Thomason, J., Murray, M., Cakmak, M., and Zettlemoyer, L. (2019b). Vision-and-dialog navigation. *CoRR*, abs/1907.04957.
- Thrun, S., Burgard, W., and Fox, D. (1998). A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots*, 5(3):253–271.
- Tsai, C.-E. and Oh, J. (2020). A generative approach for socially compliant navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2160–2166. IEEE.
- Uppal, S., Bhagat, S., Hazarika, D., Majumdar, N., Poria, S., Zimmermann, R., and Zadeh, A. (2020). Emerging trends of multimodal research in vision and language. *arXiv preprint arXiv:2010.09522*.
- Valuations, E. (2015). A review on evaluation metrics for data classification evaluations.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Vedantam, R., Zitnick, C. L., and Parikh, D. (2015). Cider: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.
- Vemula, A., Muelling, K., and Oh, J. (2018). Social attention: Modeling attention in human crowds. In *2018 IEEE international Conference on Robotics and Automation (ICRA)*, pages 4601–4607. IEEE.
- Wang, H., Wang, W., Shu, T., Liang, W., and Shen, J. (2020a). Active visual information gathering for vision-language navigation. *CoRR*, abs/2007.08037.
- Wang, H., Wu, Q., and Shen, C. (2020b). Soft expert reward learning for vision-and-language navigation. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J., editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part IX*, volume 12354 of *Lecture Notes in Computer Science*, pages 126–141. Springer.

- Wang, R., Ciliberto, C., Amadori, P. V., and Demiris, Y. (2019a). Random expert distillation: Imitation learning via expert policy support estimation. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6536–6544. PMLR.
- Wang, X., Huang, Q., Çelikyilmaz, A., Gao, J., Shen, D., Wang, Y., Wang, W. Y., and Zhang, L. (2019b). Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6629–6638. Computer Vision Foundation / IEEE.
- Wang, X., Jain, V., Ie, E., Wang, W. Y., Kozareva, Z., and Ravi, S. (2019c). Natural language grounded multitask navigation. In *Visually Grounded Interaction and Language (ViGIL), NeurIPS 2019 Workshop, Vancouver, Canada, December 13, 2019*.
- Wang, X., Jain, V., Ie, E., Wang, W. Y., Kozareva, Z., and Ravi, S. (2020c). Environment-agnostic multitask learning for natural language grounded navigation. *arXiv preprint arXiv:2003.00443*.
- Wang, X., Xiong, W., Wang, H., and Wang, W. Y. (2018). Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y., editors, *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XVI*, volume 11220 of *Lecture Notes in Computer Science*, pages 38–55. Springer.
- Weih, L., Salvador, J., Kotar, K., Jain, U., Zeng, K., Mottaghi, R., and Kembhavi, A. (2020). Allenact: A framework for embodied AI research. *CoRR*, abs/2008.12760.
- Wijmans, E., Datta, S., Maksymets, O., Das, A., Gkioxari, G., Lee, S., Essa, I., Parikh, D., and Batra, D. (2019a). Embodied question answering in photorealistic environments with point cloud perception. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6659–6668. Computer Vision Foundation / IEEE.
- Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M., and Batra, D. (2019b). Decentralized distributed PPO: solving pointgoal navigation. *CoRR*, abs/1911.00357.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8:229–256.
- Wolfe, J., Marthi, B., and Russell, S. (2010). Combined task and motion planning for mobile manipulation. In *Twentieth International Conference on Automated Planning and Scheduling*.
- Wooldridge, M. and Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2):115–152.
- Wu, Y., Wu, Y., Gkioxari, G., and Tian, Y. (2018). Building generalizable agents with a realistic and rich 3d environment. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net.

- Xia, F., Shen, W. B., Li, C., Kasimbeg, P., Tchapmi, M., Toshev, A., Martín-Martín, R., and Savarese, S. (2019). Interactive gibbon: A benchmark for interactive navigation in cluttered environments. *CoRR*, abs/1910.14442.
- Xia, F., Zamir, A. R., He, Z., Sax, A., Malik, J., and Savarese, S. (2018). Gibson env: Real-world perception for embodied agents. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 9068–9079. IEEE Computer Society.
- Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A. X., Guibas, L. J., and Su, H. (2020). SAPIEN: A simulated part-based interactive environment. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11094–11104. IEEE.
- Yan, A., Wang, X. E., Feng, J., Li, L., and Wang, W. Y. (2020). Beyond monolingual vision-language navigation.
- Yan, C., Misra, D. K., Bennett, A., Walsman, A., Bisk, Y., and Artzi, Y. (2018). CHALET: cornell house agent learning environment. *CoRR*, abs/1801.07357.
- Yang, S., Wang, Y., and Chu, X. (2020). A survey of deep learning techniques for neural machine translation. *arXiv preprint arXiv:2002.07526*.
- Yang, W., Wang, X., Farhadi, A., Gupta, A., and Mottaghi, R. (2019a). Visual semantic navigation using scene priors. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J. G., Salakhutdinov, R., and Le, Q. V. (2019b). Xlnet: Generalized autoregressive pretraining for language understanding. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.
- Ye, X. and Yang, Y. (2020). From seeing to moving: A survey on learning for visual indoor navigation (vin).
- Yu, F., Deng, Z., Narasimhan, K., and Russakovsky, O. (2020). Take the scenic route: Improving generalization in vision-and-language navigation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14-19, 2020*, pages 4000–4004. IEEE.
- Yu, L., Chen, X., Gkioxari, G., Bansal, M., Berg, T. L., and Batra, D. (2019). Multi-target embodied question answering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6309–6318. Computer Vision Foundation / IEEE.
- Zhang, Y., Tan, H., and Bansal, M. (2020). Diagnosing the environment bias in vision-and-language navigation. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 890–897. ijcai.org.

- Zhang, Z., Li, Q., Huang, Z., Wu, J., Tenenbaum, J., and Freeman, B. (2017). Shape and material from sound. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1278–1288.
- Zhao, M., Anderson, P., Jain, V., Wang, S., Ku, A., Baldrige, J., and Ie, E. (2021). On the evaluation of vision-and-language navigation instructions. *CoRR*, abs/2101.10504.
- Zhu, F., Zhu, Y., Chang, X., and Liang, X. (2020a). Vision-language navigation with self-supervised auxiliary reasoning tasks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10009–10019. IEEE.
- Zhu, H., Neubig, G., and Bisk, Y. (2021a). Few-shot language coordination by modeling theory of mind. In *International Conference on Machine Learning*, pages 12901–12911. PMLR.
- Zhu, W., Hu, H., Chen, J., Deng, Z., Jain, V., Ie, E., and Sha, F. (2020b). BabyWalk: Going farther in vision-and-language navigation by taking baby steps. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2539–2556, Online. Association for Computational Linguistics.
- Zhu, W., Qi, Y., Narayana, P., Sone, K., Basu, S., Wang, X. E., Wu, Q., Eckstein, M. P., and Wang, W. Y. (2021b). Diagnosing vision-and-language navigation: What really matters. *CoRR*, abs/2103.16561.
- Zhu, Y., Zhu, F., Zhan, Z., Lin, B., Jiao, J., Chang, X., and Liang, X. (2020c). Vision-dialog navigation by exploring cross-modal memory. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10727–10736. IEEE.