

Visually Grounded Language Learning: a Review of Language Games, Datasets, Tasks, and Models

Alessandro Suglia

*School of Mathematical and Computer Sciences
Heriot-Watt University
Edinburgh, Scotland (UK)*

A.SUGLIA@HW.AC.UK

Ioannis Konstas

*School of Mathematical and Computer Sciences
Heriot-Watt University
Edinburgh, Scotland (UK)*

I.KONSTAS@HW.AC.UK

Oliver Lemon

*School of Mathematical and Computer Sciences
Heriot-Watt University
Edinburgh, Scotland (UK)*

O.LEMON@HW.AC.UK

Abstract

In recent years, several machine learning models have been proposed. They are trained with a language modelling objective on large-scale text-only data. With such pretraining, they can achieve impressive results on many Natural Language Understanding and Generation tasks. However, many facets of meaning cannot be learned by “listening to the radio” only. In the literature, many Vision+Language (V+L) tasks have been defined with the aim of creating models that can ground symbols in the visual modality. In this work, we provide a systematic literature review of several tasks and models proposed in the V+L field. We rely on Wittgenstein’s idea of ‘language games’ to categorise such tasks into 3 different families: 1) discriminative games, 2) generative games, and 3) interactive games. Our analysis of the literature provides evidence that future work should be focusing on interactive games where *communication in Natural Language* is important to resolve ambiguities about object referents and action plans and that *physical embodiment* is essential to understand the semantics of situations and events. Overall, these represent key requirements for developing grounded meanings in neural models.

1. Introduction

Symbols of a language acquire meanings when used *to do things in the world* (Clark, 1996). In such cases, language is a cooperative enterprise used by humans to achieve specific goals. During such cooperative activities, humans coordinate meanings (Clark & Brennan, 1991). Meanings are therefore dynamic entities that humans agree upon in conversation. It is important to underline that learning a language cannot happen in isolation. Learning a language is achieved by humans engaged in activities such as *language games* (Wittgenstein, Anscombe, & Rhees, 1953). Such language games involve one or more interlocutors, who use language as a communication protocol used to express preferences, goals, and execute *actions* (Austin, 1975).

There have been many attempts at computationally representing the meanings of words in a language. From symbolic approaches (e.g., (Winograd, 1971)), statistical approaches (e.g., (Landauer & Dumais, 1997)) to distributed representations (e.g., (Elman, 1990)), AI researchers have managed to achieve some significant results. Particularly, after the recent introduction of large-scale neural language models (e.g., BERT (Devlin, Chang, Lee, & Toutanova, 2019a), GPT (Radford, Wu, Amodei, Amodei, Clark, Brundage, & Sutskever, 2019)), the goal of achieving real understanding seems closer than before. However, upon careful inspection and probing, such models demonstrate only a superficial level of understanding of Natural Language. Many researchers in the field of AI and, more specifically Computational Linguistics, have argued that real understanding cannot be achieved by exposing machines to text corpora only (Bender & Koller, 2020). This would mean expecting somebody to learn a language by reading alone (Bisk, Holtzman, Thomason, Andreas, Bengio, Chai, Lapata, Lazaridou, May, Nisnevich, et al., 2020a). They argue that meaning is a result of grounding symbolic representations via multimodal perceptual experiences of concepts in the world. Such a level of perceptual experience can be obtained only when the agent is *embodied* in the environment.

Motivated by the need to expose artificial agents to more sophisticated perceptual information, in this survey we focus on the visual modality as a source of perceptual information. In particular, we are interested in tasks that have been proposed to study the symbol grounding problem (Harnad, 1990) in situated and embodied visual contexts. The main contributions of this survey are the following:

1. We provide a novel categorisation of visually grounded language games based on the skills and capabilities required to solve them;
2. We apply the categorisation to a collection of 50 datasets presented in the last 20 years in the literature on Artificial Intelligence and Natural Language Processing;
3. We report an analysis of 51 recent visually grounded models that have been proposed to tackle the tasks studied in this work;
4. We propose relevant research questions that will guide future research in grounded language learning.

Our survey is divided into several sections that are described as follows: Section 2 presents a discussion of the background topics that are required to understand the importance of language games for V+L research; Section 3 reports our task categorisation that analyses the V+L tasks presented so far in the literature; Section 4 presents a survey of the V+L models that are able to encode the vision and language modalities; finally Section 5 presents a discussion on promising research directions for the future of V+L research; Section 7 concludes the paper.

2. Visually Grounded Language Learning

Teaching agents to understand Natural Language has been the main objective of several research projects since the early days of Artificial Intelligence (Winograd, 1971; Newell,

Simon, et al., 1972). However, understanding the meanings of words in Natural Language can be brought back to the *Symbol Grounding* problem (Harnad, 1990):

“How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols?”

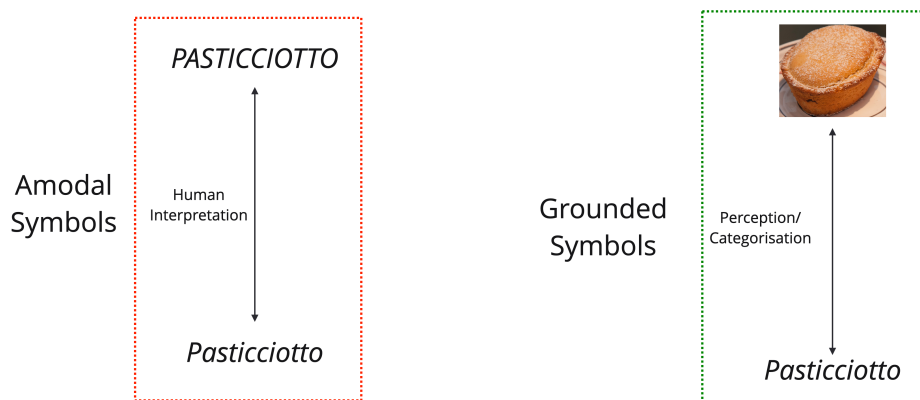


Figure 1: An example of the symbol grounding problem: a non-native speaker coming across the word “pasticciotto” for the first time will struggle to understand it because it is expressed only via *amodal symbols*. When the speaker receives an image associated with it, they are able to perform a *perception/categorisation* step thanks to which the word is grounded in *experience*, finally revealing its *meaning*.

Associating a referent in the real world with a symbolic representation via a specific and contextualised meaning still represents a challenge for artificial agents. For instance in Figure 1, interpreting the meaning of the word “pasticciotto” would not be possible without a *categorisation* step that associates the meaningless symbols of the word to a *concept* which stands for a composition of several perceptual features such as visual, functional, olfactory and gustatory features. These are a result of the agent’s *experience* of the concept in a given physical context.

This problem immediately implies that teaching computers Natural Language just by providing textual information might not be enough to effectively demonstrate that the agent *understands* language just like humans do. Artificial agents should *experience* the object whose symbols refer to, completing the so-called *Semiotic Triad* defined by (Peirce, 1902) as follows:

“I define a *sign* as anything which is so determined by something else, called its *Object*, and so determines an effect upon a person, which effect I call its *interpretant*, that the latter is thereby immediately determined by the former.”

From this statement we recognise the importance of the perceptual experience in the language learning process. The agent has to be exposed to a representation of the entity in the world (Object) that the word used (sign) refers to. The connection between the sign and the object becomes concrete once the object is perceived by the sensory organs of the agent. This process is what we consider as *grounding*.

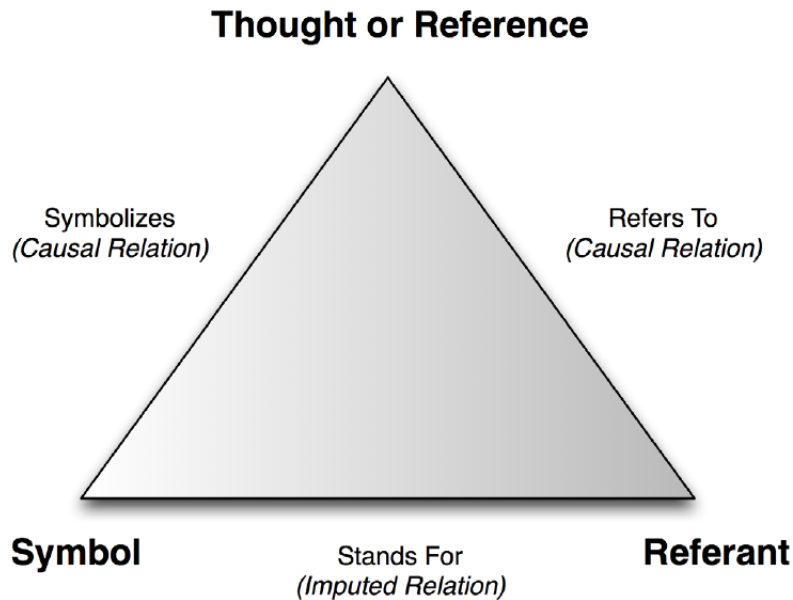


Figure 2: Semiotic Triad defined by (Ogden & Richards, 1923) which describes the relation that holds between symbols and objects in the world as well as the mental process that connects them.

The connection between language and thought was the main topic of interest of the seminal work by (Ogden & Richards, 1923). Particularly, they define *Symbols* as the tools to organise, communicate and direct *Thoughts*. As illustrated in the diagram in Figure 2, a symbol has a causal relationship with a thought that is instantiated when we speak. On the other hand, we use a symbol to refer to a given reference in the external world. Between the abstract *Thought* and the *Referant* (object in the world) a direct (or indirect) causal relation exists as well.

2.1 Compositionality & Meaning

The term ‘compositionality’ is defined as the following property of language: the ability to construct larger linguistic expressions by combining simpler parts. The focus of linguistic studies has been on *semantic compositionality*, i.e., the principle whereby the meaning of a linguistic expression is a function of the meaning of its components and the rules used to combine them (Montague, 1970). This ability to recombine known grounded symbols is what enables the grounding of complex expressions to occur. Following the example in (Harnad, 1990), suppose we are aware that both the symbols “horse” and “stripes” are grounded by appropriate representations learned from experience. Now we can consider a new category “zebra” which is obtained as a combination of the elementary categories as follows: “zebra” = “horse” & “stripes”. The answer by (Harnad, 1990) to the question “What is the representation of a zebra?” is as follows:

“It is just the symbol string ‘horse & stripes’. But because ‘horse’ and ‘stripes’ are grounded, ‘zebra’ inherits the grounding, through its grounded symbolic

representation. In principle, someone who had never seen a zebra (but had seen and learned to identify horses and stripes) could identify a zebra on first acquaintance armed with this symbolic representation alone.”

This ability to combine grounded representations to generate other grounded representations is the result of associating meaning with a given representation. In this sense, compositionality can be considered as a more general learning bias that humans leverage in order to learn new tasks by reusing their knowledge about previously learned tasks and concepts (Fodor & Pylyshyn, 1988; Fodor & Lepore, 2002). Indeed, compositionality represents a key element to cope with the huge amount of labelled data required by classical Deep Learning algorithms (Marcus, 2018) which learn tasks in an *end-to-end* fashion, directly from the raw data (LeCun, Bengio, & Hinton, 2015). Therefore, we are interested in learning representations of concepts that are *grounded in perceptual experience* and also demonstrate some *degree of compositionality*.

A prerequisite for learning and developing meanings is to develop grounded representations that can favour systematic generalisation to novel combinations of concepts that we have never seen before. Possessing such flexibility assumes that such learned conceptual representations should, by definition, not be static. They must dynamically change depending on the context and signals that the agent receives when interacting with the environment. Compositionality comes into play when an agent has to dynamically refine its internal representations when tasked to coordinate with another agent to collaboratively achieve a goal (Lake & Murphy, 2021). In such Natural Language interactions, words acquire meanings that are instantiated in the ways in which they are used as part of a systematic communication protocol which is *a language*.

2.2 Language Games

The definition of language as a tool for human communication underlines the social functions of language. Additionally, humans use it to express themselves and to manipulate objects in their environment. Therefore, learning a language means being able to communicate effectively using a predefined communication protocol based on symbols and transformations of them. The meaning associated with each symbol can be acquired by engaging in goal-oriented conversations with other agents. This practice is what (Wittgenstein et al., 1953) defined as “language games”. Therefore, language functions within the active, practical lives of speakers involved, and is deeply connected (i.e., *grounded*) with non-linguistic elements and behaviours that are part of their environment. Indeed, (Wittgenstein et al., 1953) considers language as a system of signs that acquire meaning *in situ*, embedded in speakers’ lives.

When completing real-world tasks, artificial agents must *communicate* in order to understand what the goal and requirements of a certain task are. Therefore, in this work, we underline the importance of the hypothesis that “human communication is thus a fundamentally cooperative enterprise, operating most naturally and smoothly within the context of (1) mutually assumed common conceptual ground, and (2) mutually assumed cooperative communicative motives” (Tomasello, 2010). This definition highlights the importance of using the terminology of language games as an experimental framework for understanding

the complexity, properties, and importance of certain tasks for the overall goal of language understanding and visual grounding.

In contrast to fully symbolic approaches which argue that language can be defined as an abstract symbol manipulation system (Newell et al., 1972; Pylyshyn, 1986; Fodor, 1983), (Wittgenstein et al., 1953) considers language as a system that has a function, and acquires meaning, only when it is actually used to do things in the world. As in the *Builder* example (Wittgenstein et al., 1953), the phrase “pick up” will acquire meaning only when the *Architect* will instruct the *Builder* to pick a specific object up.

In this survey, we underline the importance of the functional and goal-oriented nature of Natural Language, investigating how artificial agents can acquire visually grounded representations *by playing* situated and embodied language games.

3. Visually Grounded Language Learning Tasks

This paper surveys the state of the art in Visually Grounded Language Learning, which is a type of *Interactive Grounded Language Learning*. This can be considered as an instance of *Interactive Task Learning* (Gluck, Laird, & Lupp, 2018), i.e., “any process by which an agent (A) improves its performance (P) on some task (T) through experience (E), when E consists of a series of sensing, effecting, and communication interactions between A , its world, and crucially other agents in the world”. We explore the idea of Wittgenstein’s *language games* (Wittgenstein et al., 1953) as a means to assess the linguistic capabilities of interactive learning agents. Specifically, as in the case of the “Talking Heads” experiments (Steels, 2015), we explore the idea that linguistic capabilities could be assessed by 3 relevant language games: 1) *guessing games* that require the agent to guess an unknown object in a scene; 2) *action games* that involve the execution of actions by an agent as requested by a specific command; and 3) *descriptive games* which require that the agent is able to describe the scene that it observes.

We report an analysis of the state-of-the-art in terms of datasets and tasks that were proposed to study the problem of visual grounding for the English language (summarised in Table 1). We will divide the tasks into 3 macro categories following the different types of language games: 1) *discriminative* tasks; 2) *generative* tasks; and 3) *interactive* tasks. Furthermore, we analyse each paper along the following key dimensions:

- **Embodied:** is the agent able to explore, perceive and act in the environment?
- **Discriminative:** does the agent have to produce a single output by selecting it from a given set of candidates?
- **Generative:** does the agent have to produce a sequence of outputs conditioned on its input?
- **Interaction with environment:** is the agent able to manipulate and change the state of the environment?
- **Interaction with other agents:** does the agent have to communicate with other agents to solve the task?

VISUALLY GROUNDED LANGUAGE LEARNING: A REVIEW

Table 1: Systematic analysis of several state-of-the-art datasets for Grounded Language Learning tasks. The analysis considers specific characteristics of the tasks and environments that are important for learning grounded meanings. We use “x” to mark that a dataset satisfies a given property.

Authors	Dataset	Embodied	Discriminative	Generative	Inter. with Env	Inter. with Agents	Size
(Gordon et al., 2018)	IQA	yes	x	x	x		75K
(Suhr et al., 2019b)	NLVR	no	x				100K
(Suhr et al., 2019a)	CerealBar	yes			x	x	1.2K
(Shridhar et al., 2020)	ALFRED	yes		x	x		25K
(Ku et al., 2020)	RxR	yes			x		126K
(Padmakumar et al., 2021)	TEACH	yes			x	x	3.2K
(Deitke et al., 2020)	RoboTHOR	yes		x	x		
(Wortsman et al., 2019)		yes		x	x		
(Jain et al., 2019)		yes		x	x		
(Kottur et al., 2021)	SIMCC 2.0	no	x			x	11K
(Chen et al., 2019)	Touchdown	yes		x			9.3K
(Haber et al., 2019)	Photobook	no				x	2.5K
(Hermann et al., 2017)		yes		x			
(Yan et al., 2018)	CHALET	yes					
(Mirowski et al., 2019)	StreetLearn	yes					
(De Vries et al., 2018)	Talk The Walk	yes			x	x	10K
(Kim et al., 2019)	CoDraw	no				x	10K
(Brodeur et al., 2017)	HOME	yes					
(Ramakrishnan et al., 2021)	HM3D	yes					
(Savva et al., 2019)	Habitat	yes					
(Kiela et al., 2020)	Hateful Memes	no	x				10K
(Ruis et al., 2020)	gSCAN	no		x	x		300K
(Shekhar et al., 2017)	FOIL	no	x				297K
(Zellers et al., 2019)	VCR	no	x				290K
(Park et al., 2020)	VisualCOMET	no	x				1.46M
(Das et al., 2017a)	VisDial	no	x	x			1.4M
(Bogin et al., 2021a)	COVR	no	x				262K
(Liu et al., 2021)	MaRVL	no	x				5.6K
(Narayan-Chen et al., 2019)	Collab-Minecraft	yes			x	x	509
(Zarri� et al., 2016)	PentoRef	no				x	1.3K
(Pezzelle & Fern�andez, 2019)	MALeViC	no	x				20K
(Pezzelle et al., 2020)	BD2BB	no	x				18K
(Huang et al., 2016)		no		x			20.2K
(Da et al., 2021)	EMU	no	x	x			48K
(Clark et al., 2021)	Iconary	no			x		55K
(Goyal et al., 2017a)	VQA 2.0	no	x				1.1M
(Hudson & Manning, 2019)	GQA	no	x				22M
(Kafke & Kanan, 2017)	TDIUC	no	x				1.6M
(Yang et al., 2021)		no	x				53.2K

Continued on next page

Table 1: Systematic analysis of several state-of-the-art datasets for Grounded Language Learning tasks. The analysis considers specific characteristics of the tasks and environments that are important for learning grounded meanings. We use “x” to mark that a dataset satisfies a given property. (Continued)

Authors	Dataset	Embodied	Discriminative	Generative	Inter. with Env	Inter. with Agents	Size
(Kottur et al., 2019)	CLEVR-dialog	no				x	4.25M
(Zellers et al., 2021)	PIGLeT	yes	x				280K
(Abramson et al., 2020)	Playroom	yes			x	x	
(Ilinykh et al., 2019b)	Tell-me-More	no		x			22K
(Dobnik et al., 2020)	CUPS	no				x	1.3K
(Yu et al., 2017b)	BURCHAK	no				x	2.5K
(Tokunaga et al., 2012)	REX	no			x	x	9.8K
(Chevalier-Boisvert et al., 2018)	BabyAI	yes			x		
(Thomason et al., 2019a)	CVDN	yes			x	x	2K
(Ilinykh et al., 2019a)	MeetUp	yes			x	x	430
(Puig et al., 2020)	Watch-And-Help	yes			x	x	1.3K
(Zheng et al., 2022)	Spot the difference	no			x		95K
(Kiseleva et al., 2022)	IGLU	yes			x	x	509
(Bisk et al., 2016)	BlocksWorld	no		x	x		12K
(Zhong et al., 2021)	SILG	no		x	x		
(Elliott et al., 2016)	Multi30K	no		x			150K
(Wang et al., 2021)	Multisubs	no		x			4.5M
(Kazemzadeh et al., 2014)	ReferitGame	no		x			20K
(Johnson et al., 2016)	DenseCap	no		x			4.1M
(Krishna et al., 2017a)	ActivityNet	no		x			20K
(Thrush et al., 2022)	Winoground	no	x				1.6K
(Bugliarello et al., 2022)	IGLUE	no	x				28K
(Krojer et al., 2022)	ImageCoDe	no	x				22K
(Shridhar et al., 2021)	ALFWorld	yes		x	x		25K
(Gao et al., 2022)	Dialfred	yes			x	x	53K
(Gao et al., 2023)	Alexa Arena	yes			x	x	46K

3.1 Discriminative Tasks

A discriminative task is defined as a language game in which the agent, given an image I and a corresponding textual information t , has to select an option \hat{c} among a set of candidates \mathcal{C} . The agent is successful if \hat{c} matches the target option c^* . Such tasks have been the very first instance of grounded language learning tasks inspired by well-known image classification tasks in the Computer Vision community (Deng, Dong, Socher, Li, Li, & Fei-Fei, 2009).

One of the most famous tasks in this category is **Visual Question Answering (VQA)**. The generic formulation of this task, as shown in Figure 3, assumes that the agent receives a question about an image and has to provide an answer to it. The answer can be in free-form

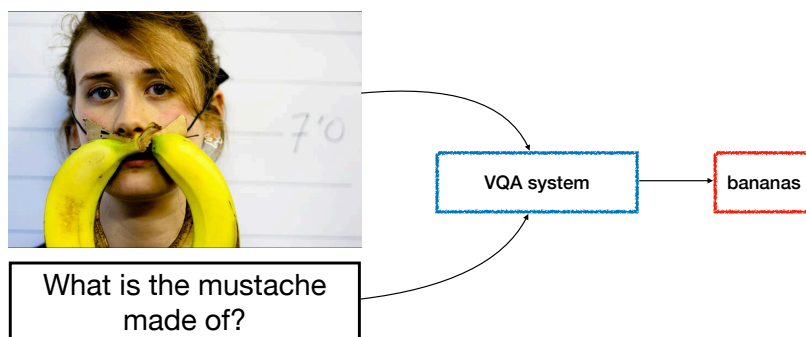


Figure 3: Example from the Visual Question Answering dataset (Antol et al., 2015): the agent should learn to generate the answer “bananas” given the image and the question.

text. However, current implementations of VQA tasks, such as the one presented by (Antol et al., 2015), assume that the agent has a specific answer vocabulary (*vocabulary-based VQA*) that it has to choose from. In this task, the agent should learn to reason about the objects in the image to answer the question. Questions usually require that the agent is aware of specific categories of objects, their colour, or their relative position in the image. However, the first release of the dataset had a bias according to which questions could be answered by relying only on one modality or only by relying on a specific subset of tokens in the question. For this reason, as described by (Goyal, Khot, Summers-Stay, Batra, & Parikh, 2017b), a proper balancing procedure for the dataset is required. This prevents the case where high-capacity models obtain high accuracy only because they learn spurious correlations in the data. Other ways to mitigate this problem have been proposed. For instance, the TDIUC dataset (Kafle & Kanan, 2017) divides the questions into 12 categories to have a better understanding of the model’s capabilities. A related effort in this direction is the GQA dataset (Hudson & Manning, 2019). GQA is a VQA dataset automatically generated by relying on Visual Genome scene graphs (Krishna, Zhu, Groth, Johnson, Hata, Kravitz, Chen, Kalantidis, Li, Shamma, et al., 2017b). Having a structured representation of the scene adds a semantic layer on top of the purely perceptual one. Thanks to this additional level, it is possible to perform a more fine-grained analysis of the model’s ability. For instance, the authors propose measures of *consistency* of the model predictions (i.e., is the answer to semantically similar questions the same?), *validity* (i.e., is the agent responding some colour to a colour question?), and *plausibility* (i.e., does the answer make sense given the question?). However, because the problem is cast as a classification task, models learn to model the head of the answer probability distribution only. Therefore, they perform very poorly when generalising to out-of-distribution examples. To address this problem, specific datasets that test the generalisation abilities of VQA models have been proposed (e.g., (Agrawal, Kembhavi, Batra, & Parikh, 2017)). The community proposed the CLEVR (Johnson, Hariharan, van der Maaten, Fei-Fei, Lawrence Zitnick, & Girshick, 2017) dataset to complete a more in-depth evaluation of the elementary visual reasoning skills of an agent using synthetic images generated using the Blender render engine.

Another instance of a discriminative task is **Visual Entailment (VE)** (Xie et al., 2018). It is inspired by Textual Entailment tasks in Natural Language Processing (e.g.,

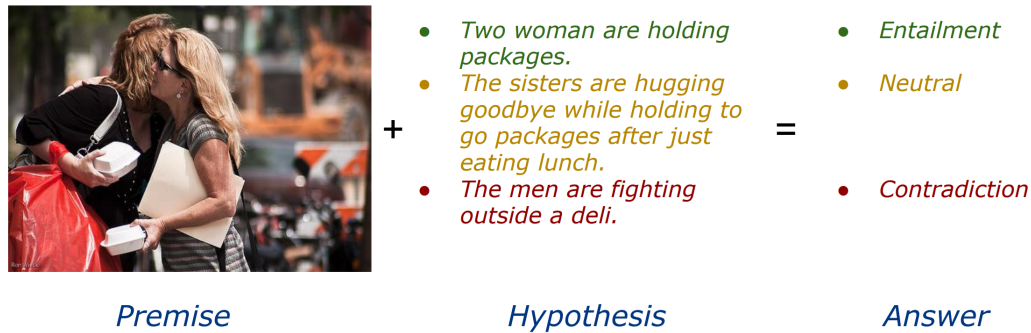


Figure 4: Example from the Visual Entailment task from (Xie et al., 2018). Given the state of the world (exemplified by the image), the agent has to verify, for each hypothesis, the validity of the statement. Figure adapted from (Xie et al., 2018).

(Bowman, Angeli, Potts, & Manning, 2015)) defined as follows: given a text premise P and a text hypothesis H , the goal is to determine if P implies H . This task is usually cast as a 3-way classification problem where the label set includes “entailment”, “neutral”, or “contradiction”, based on the relation conveyed by the (P, H) text pair. As shown in Figure 4, VE replaces premise P with a real-world image. Based on this idea of entailment, several other tasks can be defined. For instance, the *Cornell Natural Language Visual Reasoning (NLVR)* dataset (Suhr, Lewis, Yeh, & Artzi, 2017) is a corpus containing 92,244 sentence-image pairs, whose aim is to teach an agent whether a given statement is true or false for a given image. What makes this dataset interesting is that it includes several semantic phenomena such as cardinality (soft/hard) statements, existential and universal relations, as well as spatial relations. Sentences associated with the images are collected from real users, however, the corresponding images are synthetic and composed of simple coloured blocks. (Suhr et al., 2019b) extended this dataset with real-world images.

Another discriminative task is the **Find One mismatch between Image and Language caption (FOIL)** dataset (Shekhar et al., 2017). It defines 3 different discriminative sub-tasks described as follows (see Figure 5 for an example):

1. *Binary Classification*: Given an image and a caption, the model is asked to mark whether the caption is correct, or wrong. The aim is to understand whether trained models can spot mismatches between their coarse representations of language and visual input;
2. *Foil Word Detection*: Given an image and a caption, the model has to detect the foil word (i.e., a single word that is incorrect for the caption). The aim is to evaluate the understanding of the system at the word level;
3. *Foil Word Correction*: Given an image, a caption and the foil word, the model has to detect the foil and provide its correction. The aim is to check whether the system’s visual representation is fine-grained enough to be able to extract the information necessary to correct the error.

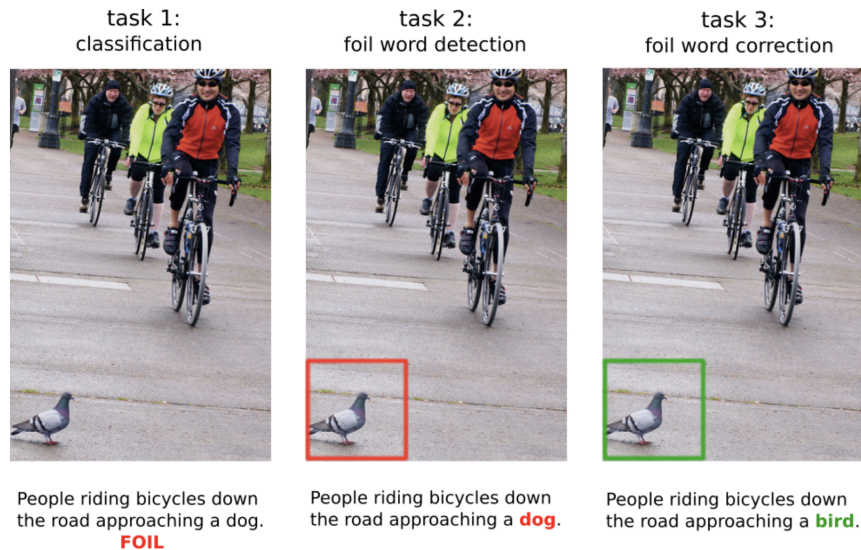


Figure 5: Example extracted from the FOIL dataset. It describes the 3 discriminative tasks: binary classification, foil word detection and foil word correction. Figure adapted from (Shekhar et al., 2017).

Most of these datasets have been annotated relying on image datasets that were derived from an English lexical database (i.e., WordNet (Miller, 1995)). However, this imposes a bias on the data collection because all the elicited concepts will be familiar only to people with a European/American background. Since an over-arching goal in this area is to create general-purpose agents, they must rely on an unbiased source of training data. (Liu et al., 2021) proposed **MaRVL**, a novel dataset for multicultural reasoning over vision and language. First of all, they propose a novel concept hierarchy that can be considered universal and not specific to the English language. After this step, they collect a dataset using a similar approach to (Suhr et al., 2019b). The resulting dataset contains 5.6K captions using 5 typologically, geographically and genealogically diverse languages. As shown in Figure 6, this makes MaRVL a very challenging benchmark for multimodal and multilingual language understanding.

Similar to FOIL, **MALeViC** (Pezzelle & Fernández, 2019) defines a sentence verification task to study how Vision+Language (V+L) models can learn the meaning of gradable adjectives of size from different visual contexts. This benchmark includes several tasks that are specifically designed to study how an agent can learn to identify the reference set of a given statement in a specific visually-grounded context. This is a crucial skill that agents have to master when dealing with complex Embodied AI tasks involving a high level of ambiguity. Another task formulated as a discriminative language game is “**Be Different to Be Better**” (Pezzelle et al., 2020). It is defined as a candidate selection task where the agent, based on a certain intention (i.e., their goal, attitude or feeling), has to choose, among a set of candidate actions, the one that a person would perform. Other instances of these tasks involve matching the correct caption to a given image. This is the case of **Winoground** (Thrush et al., 2022) where an agent has to be able to score

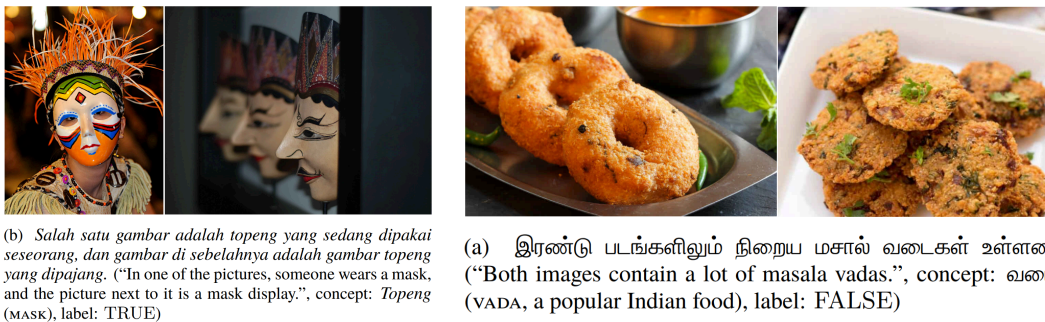


Figure 6: Examples extracted from the MaRVL dataset. The first image shows an example in Indonesian and the second in Tamil. Figures adapted from (Liu et al., 2021).

higher the correct (image, caption) pair instead of the others inspired by the Winograd Schema Challenge (Levesque, Davis, & Morgenstern, 2012). This dataset is extremely challenging because it contains images having an intricate visual structure that have been manually sourced by a team of experts. Humans are able to match the correct caption to its image while state-of-the-art models have a really hard time completing this task with top performance being just above random chance.

The dataset **ImageCoDe** provides an additional benchmark for contextual image retrieval (Krojer et al., 2022) where the agent has to select the image that matches a contextual caption given a set of 9 distractors. As a result of their evaluation, they show that their benchmark is highly multi-modal (i.e., the agent requires both modalities to do well) and that current state-of-the-art V+L models are still far away from human-level performance. Most likely, this is because some of the situations require understanding the unfolding of events over time and cannot be learned from static images only. Similarly to the effort in VQA to inspect the visual grounding ability of the model, (Bogin et al., 2021a) proposes a new benchmark for **compositional language understanding (COVR)**. The task, inspired by NLVR (Suhr et al., 2019b), is formulated as a sentence verification problem having multiple reference images (see Figure 7). The dataset has been created by defining specific splits having different compositional generalisation skills. It uses the GQA (Hudson & Manning, 2019) scene graphs to first find adversarial images to a target image. Then, it instantiates questions from language templates and finally uses manual validation and paraphrasing to generate the actual text. The generated examples require the agent to master the use of quantifiers and other operations, such as counting. This makes the task particularly challenging, especially in zero/few-shot splits where the agent has the chance to learn from a few examples of a given logical operator.

Most of these tasks can be solved by relying on information immediately available in the image (e.g., the properties of the objects or the relationship between them). However, understanding the effects of given actions as well as the goal of specific actors in an image plays an important role in visual commonsense reasoning. To tackle this important problem, (Zellers et al., 2019) propose the **Visual Commonsense Reasoning** challenge. It is defined as a multiple-choice classification task where the agent has to select the correct answer to a given question as well as specify “why” that answer was relevant. **PIGLEt** (Zellers et al., 2021) studies the problem of learning physical commonsense knowledge. The authors

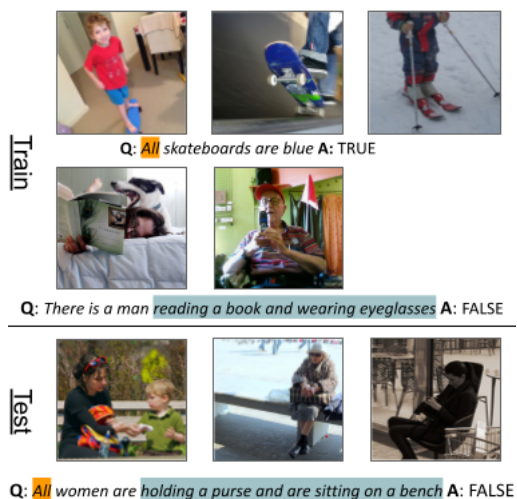


Figure 7: Example extracted from the COVR dataset. At training time, the agent is exposed to specific usages of quantifiers as well as object attributes. At test time, novel, unseen usages are presented so that the model has to *systematically generalise* to novel statements. Figure adapted from (Bogin et al., 2021a).

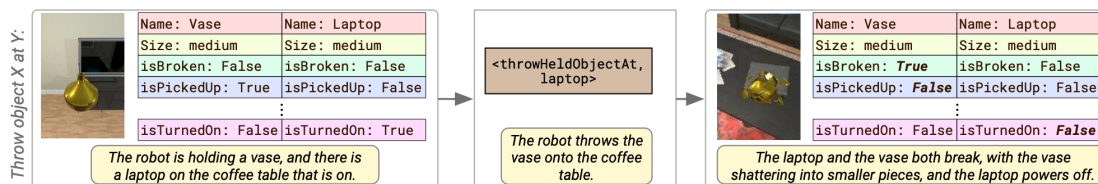


Figure 8: Example extracted from the PIGLeT dataset. Each example contains the representation of the state of the world before and after the action “throw object” happens. The world state is expressed in terms of visual attributes derived from the 3D environment. Language annotations are available for each step of the transition. Figure adapted from (Zellers et al., 2021).

automatically generated trajectories of actions using a planner and then annotated specific transitions where a state change happens. In particular, as shown in Figure 8, for each transition they collect a set of visual attributes for each state as well as a description of the action that triggered the state change, and a description of the final state. One of the main tasks involves predicting the attributes of the final state given the action executed by the agent. **VisualCOMET** (Park et al., 2020) is another dataset for visual commonsense reasoning in static images. Given an image, a person that appears in it is selected and multiple reasoning inferences are annotated that are supposed to describe: 1) what needed to happen before; 2) the intents of the person; and 3) what will happen next.

Another multimodal dataset that requires higher-order reasoning skills is the **Hateful Memes** challenge (Kiela et al., 2020). This dataset was created to develop models that are able to detect internet abuse, an important problem in current social media platforms where sarcasm and other very subtle visual and language cues are used in an offensive manner.

This problem is particularly important because it underlines the fact that words can assume different meanings depending on the multi-modal context in which they are used. Being able to derive their meaning is therefore a highly complex reasoning task.

3.2 Generative Tasks

Generative (or *descriptive*) language games are formally defined by assuming that there is an agent that receives an image I and has to generate a sequence of tokens $\hat{t} = \langle t_1, t_2, \dots, t_n \rangle$. The task is evaluated by considering a specific semantic similarity measure S that determines how close \hat{t} is to the ground-truth target sequence t .



Figure 9: Image captioning example from the MSCOCO dataset (Lin et al., 2014). A possible caption for this image is: “a teddy bear that has been placed on a pile of rocks.”

One of the first tasks that have been proposed in the literature is **Image captioning** in which the main objective is to generate a meaningful description of an input image (see Figure 9 for an example). One of the most used datasets for this task is MSCOCO (Lin et al., 2014) which provides 5 captions for every image as well as segmented objects. A similar dataset is **Flickr30K** (Young, Lai, Hodosh, & Hockenmaier, 2014). In a similar fashion, **video captioning** tasks have been proposed (e.g., (Wang, Wu, Chen, Li, Wang, & Wang, 2019; Lei, Yu, Berg, & Bansal, 2020)) where either the subtitle or a caption has to be generated for a given video clip. A more complex scenario is represented by video summarisation (e.g., (Papalampidi, Keller, Frermann, & Lapata, 2020)). Analogous to text summarisation (Nenkova & McKeown, 2011), video summarisation involves a more structured approach that resembles classic Natural Language Generation pipelines (Reiter & Dale, 2000). It usually involves several steps: 1) selecting the video frames that are worth considering for the summary; 2) aggregating the information in a coherent way to favour sound summaries; and 3) using the aggregated pieces of information to generate the text of an appropriate summary. Another task that involves multiple video frames as part of the input data is visual storytelling (Huang et al., 2016). In this case, the agent is provided with

multiple images that describe the progression of some event over time. This task is more complex than video captioning because it is not only limited to factually describing the image but involves also adding some extra contextual cues (e.g., emotional and time-related aspects). Another generative task that involves describing a visual scene in multiple rounds is **Tell-me-More** (Ilinykh et al., 2019b). In their data collection, they ask their users to imagine they are talking to someone over the phone about an image they can see. This image belongs to a set of candidate images that are available to the other agent only. The receiver has to select which image the speaker is referring to.

Most of the generation tasks described so far assume that the agent has to generate a description for the entire visual input. To evaluate how fine-grained visual representations are it is useful to focus only on certain parts of the image. This task is commonly called *dense captioning* or *referential expression* generation. In this task, an agent has to generate a description that applies to a very specific region of the image (i.e., a bounding box). For instance, (Kazemzadeh et al., 2014) propose the ReferIt game intended as a generic data collection procedure for referential expression generation. This procedure has been then applied to several datasets such as COCO (Lin et al., 2014) and ImageCLEF (Grubinger, Clough, Müller, & Deselaers, 2006). Concurrently, rich semantic annotations of images have been collected via the VisualGenome project (Krishna et al., 2017b). From these annotations, (Johnson et al., 2016) released the **DenseCap** dataset to learn to generate descriptions associated with specific bounding boxes of the image. (Krishna et al., 2017a) collected a dataset for captioning in videos as well. Particularly, this dataset was annotated by dividing the video into key activities (or moments) and then asking the annotator to report a sentence describing each activity.

Another important research direction is to use visual information to boost performance in machine translation. The first dataset to be proposed for this task was **Multi30K** (Elliott et al., 2016). The dataset was created by creating German translations for each image description available in Flickr30K (Young et al., 2014). However, this process was completed without providing images to the annotators introducing an apparent bias in the collected data (Li, Ataman, & Sennrich, 2021). To mitigate this issue, (Wang et al., 2021) proposed **Multisubs**, a large-scale multimodal and multilingual dataset that facilitates research on grounding words to images in the context of their corresponding sentences. This dataset is based on movie subtitles but due to missing copyright, the authors didn't use the images of the video as reference. Instead, they relied on a multilingual knowledge base to find references to *general* images about certain concepts (i.e., *BabelNet* (Navigli & Ponzetto, 2010)).

3.3 Interactive Tasks

In Section 2, we introduced the concept of language games as an experimental framework for studying machine learning tasks aimed at assessing certain language understanding skills of artificial agents. This is important because if we want to aim to create artificial agents that are trained to solve ecologically valid tasks (De Vries, Bahdanau, & Manning, 2020), we need to consider the fact that the real world contains: 1) sophisticated visual stimuli in the form of videos or images; 2) ambiguous, spontaneous, and incremental language that is typically used when communicating in natural language.

The grounded language learning tasks that we described so far did not have a real connection with the concept of language games that we described in previous sections. Although each language game could be considered as *situated* in the visual scene exemplified by the image, there was no notion of *interaction* between the learning agent and the environment or other agents. Each task involved either generating one or multiple elements associated with the input. To get closer to the idea of language games described by (Wittgenstein et al., 1953), we are interested in exploring *interactive* tasks where the agent is either *situated* or *embodied* in the environment and has to interact in particular ways. This can be either by communicating with other agents to solve a given task or by executing actions in the environment to achieve a given goal.

3.3.1 SITUATED INTERACTIVE GAMES

Agents who are able to support a conversation are much more human-like than agents that are able to answer single questions only. For this reason, the VQA task has been extended by the **Visual Dialogue (VisDial)** dataset (Das et al., 2017a). In the VisDial dataset, as depicted in Figure 10, the agent receives in input the dialogue history composed of a sequence of question/answer pairs, the current question and the current reference image. Two different tasks have been proposed for this dataset: 1) *answer selection*, where the agent has to select the correct answer among a set of 100 candidates, and 2) *answer generation*, where the agent should generate a response token by token. However, as highlighted by (Massiceti, Dokania, Siddharth, & Torr, 2018), state-of-the-art results can be achieved by using the Canonical Correlation Analysis algorithm that relies on ad-hoc feature extractors for images (image features extracted from a pretrained *ResNet* (He, Zhang, Ren, & Sun, 2015)) and text (question and answer representations are learned using *FastText* (Bojanowski, Grave, Joulin, & Mikolov, 2017)). Another downside of this task is the relevance of history. (Agarwal, Bui, Lee, Konstas, & Rieser, 2020) show that the number of dialogues actually requiring dialogue history is very low. Therefore, models can do well just by using the current question to generate the answer.

A dataset that proposes to give the role of first-class citizen to dialogue history and common ground is **Photobook** (Haber et al., 2019). It is a dataset of 2,500 human-human goal-oriented English dialogues between two participants that have to identify shared images in their *photobook* by exchanging text messages. In their evaluation, they focus on the task of reference resolution and reference generation in such visually grounded dialogues. However, they do not assess the ability of artificial agents to actually play the game in a collaborative fashion by both generating utterances *as well as* selecting target photos.

3.3.2 LANGUAGE EVOLUTION AND EMERGENT COMMUNICATION

Another line of work is focused on language evolution and emergent communication in referential language games (Lewis, 1969). A specific instance of the game, depicted in Figure 11, can be described as follows: a speaker describes an object in a scene using a set of symbols and the hearer needs to select, among a set of distractors, which one is the object described by the speaker. A recent implementation of the game was studied by several authors (e.g., (Lazaridou, Hermann, Tuyls, & Clark, 2018; Havrylov & Titov, 2017; Bouchacourt & Baroni, 2018)) who use either synthetic images or real-world images

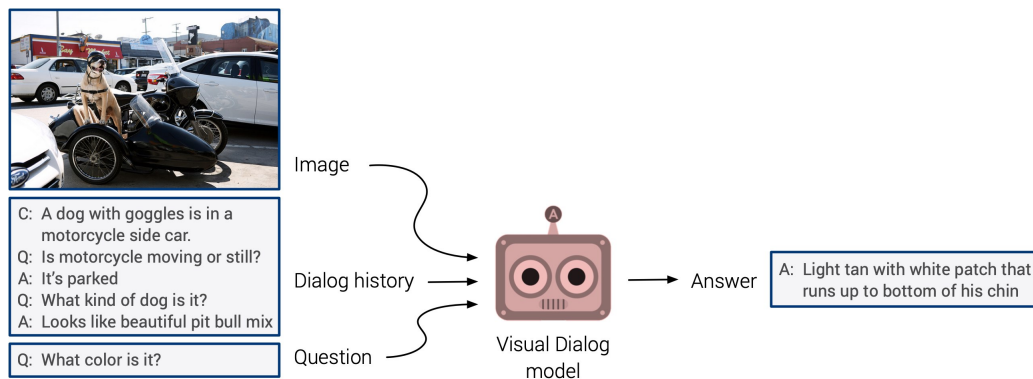


Figure 10: An example from the VisualDialog dataset. The agent has access to a reference image, to the dialogue history composed of the question/answer pairs generated so far, and it has to generate an answer for the current question. Figure adapted from (Das et al., 2017a).

coming from MSCOCO. A downside of these approaches is that they focus on single-turn interactions, and are thus not considering the intricacies of dialogue phenomena which are essential for an agent that has to learn by natural language interaction. It is important to underline the fact that in the language evolution and emergent communication literature, agents start tabula rasa and form communication protocols that maximise task rewards. While this purely utilitarian framework results in agents that successfully learn to solve the task by creating a communication protocol, these emergent communication protocols do not bear core properties of Natural Language (Kottur, Moura, Lee, & Batra, 2017). This is ultimately what agents are required to develop if we want them to be embodied in the environment and join in conversations with humans.

Due to the complexity introduced by having situated dialogues divided into many turns, there have been several initiatives that tried to study very specific problems in isolation. For example, (Clark et al., 2021) proposes **Iconary**, a collaborative game with two agents: 1) a Drawer that receives a phrase and has to draw it using a specific set of icons; and 2) a Guesser who generates a text to guess the phrase. This dataset is appealing because it involves an asymmetry of information that favours communication between players to guess the correct phrase. **PentoRef** (Zarri   et al., 2016) is another puzzle game where the Instruction Follower has to manipulate specific game pieces to achieve the goal that the Instruction Giver has in mind. In the same spirit, (Kottur et al., 2019) developed **CLEVR-dialog**, a two-player collaborative game where the aim is to reconstruct a hidden scene graph representation known only to one of the players. The focus of this task is on the dialogue history and its importance for reference resolution. Similarly, the **CODRAW** game (Kim et al., 2019) involves two players: a *Teller* and a *Drawer*. The Drawer asks questions to the Teller in order to accurately reconstruct the image that only the Teller can see. The task requires that the agents develop Natural Language understanding and generation skills in order to complete the game. They use the quality of the reconstructed image as a proxy for the quality of the communicative skills of the Drawer.

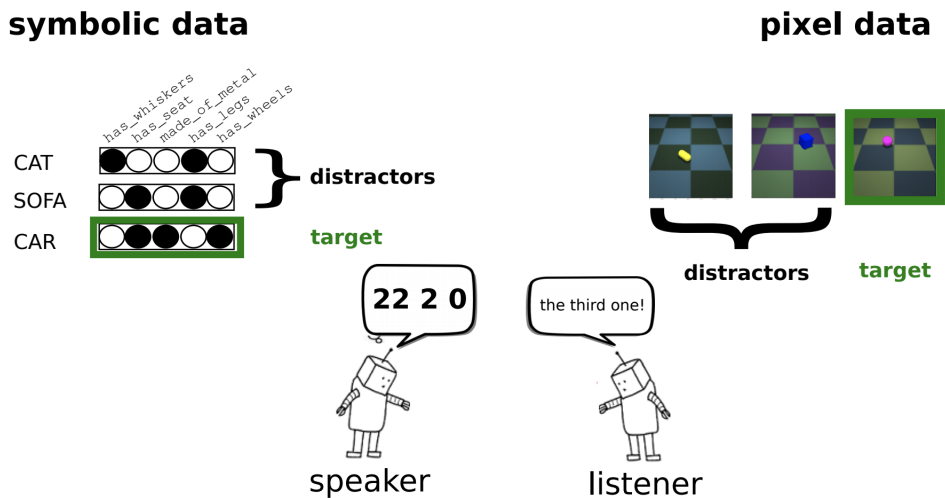
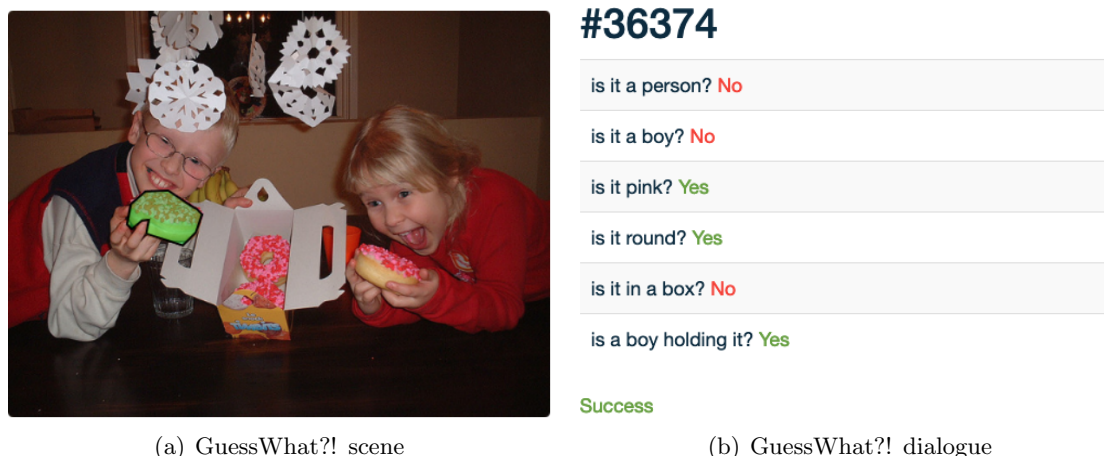


Figure 11: Referential game used by (Lazaridou et al., 2018) as a benchmark for emergent communication between two agents: the Speaker and the Listener. The Speaker, given a representation of the target image—which could be either symbolic or in the form of pixels—generates a message by sampling symbols from a fixed lexicon. Given the message sent by the Speaker, the Listener has to understand which one is the target object among the distractors. Figure adapted from (Lazaridou et al., 2018).

Another situated interaction dataset is **Situated and Interactive Multimodal Conversations (SIMMC)** (Kottur et al., 2021). It represents a collaborative task where a user is looking for an item of clothing and has to ask an AI assistant for help in a simulated clothes shop. Another interactive situated dialogue task is **CUPS** (Dobnik et al., 2020). In this setup, the authors define a situated context represented by a table with multiple coloured cups. They assume there are multiple speakers that have different perspectives of the same scene, each with different objects missing. Their task consists of spotting which objects are missing and different. The difference in perspective favours communication between the players to succeed in the game. In this setup, we assume that the two speakers have knowledge of the language used to communicate. Instead, in **BURCHAK** (Yu et al., 2017b), the authors collected a dataset where one of the agents, the Teacher, is trying to teach new visually grounded words in a completely made-up language to the Student, who is not familiar with it. Again, the concept of asymmetry of information is leveraged to favour communication between agents.

In contrast to the previously described tasks, the **Guess What?!** dataset (De Vries et al., 2017) represents the first benchmark that involves a goal-oriented dialogue that favours: 1) the acquisition of discriminative features associated with the objects in the scene by playing the role of the Questioner – a module which generates questions about a given image; and 2) the development of Natural Language understanding capabilities by playing the role of the Oracle – the module which has to coherently reply to the Questioner. As shown in Figure 12, the Questioner has to *understand* the image in order to *generate* coherent questions about the reference scene and, at the same time, has to *understand* the feedback provided by the Oracle in order to adapt its dialogue strategy. Differently from all other tasks,



(a) GuessWhat?! scene

(b) GuessWhat?! dialogue

Figure 12: An example dialogue extracted from the *GuessWhat?!* dataset (De Vries et al., 2017). Figure 12(a) shows a possible scene and Figure 12(b) shows the game played by the Questioner and the Oracle. The Oracle is aware that the target object is the “Donut” highlighted in green, and has to support the Questioner in guessing it by replying Yes/No to their questions.

GuessWhat?! involves both generation and understanding of Natural Language expressions that are grounded in the reference scene. We consider both capabilities as fundamental in learning grounded language learning. This is because agents able to support a conversation in natural language with other agents should be able to understand and generate coherent responses.

In a scenario in which agents are *divergent* – they have different perspectives, different skills, and different languages – Natural Language communication emerges because it is essential to coordinate upon certain references, and certain goals (Chandu, Bisk, & Black, 2021; Benotti & Blackburn, 2021). The dataset **Spot the Difference** presented by (Zheng et al., 2022) contains 95K simulated dialogues based on two agents trying to solve a referential game. Another example of divergent agents is **MeetUp!** (Ilinykh et al., 2019a). This is a game situated in a simulated environment where two players have to converse and coordinate in order to meet in the same room. Differently from *GuessWhat?!*, in this game the two agents have different visual perspectives. This asymmetry of information forces agents to communicate and clarify their references in order to complete their tasks.

3.3.3 EMBODIED INTERACTIVE GAMES

Differently from *situated* tasks, many approaches can be classified as Embodied AI tasks (Savva et al., 2019) which are focused on the execution of actions either in the real world or within 3D virtual environments (e.g., *Matterport3D* (Chang, Dai, Funkhouser, Halber, Niebner, Savva, Song, Zeng, & Zhang, 2017), *HoME* (Brodeur et al., 2017), *CHALET* (Yan et al., 2018), *HABITAT* (Savva et al., 2019), and *AI2THOR* (Kolve, Mottaghi, Han, Vanderbilt, Weihs, Herrasti, Gordon, Zhu, Gupta, & Farhadi, 2017)). (Misra, Bennett, Blukis, Niklasson, Shatkhin, & Artzi, 2018a) jointly model navigation and plan execution by re-

ceiving visual information from the current scene and generating a plan of a sequence of task-dependent actions such as moving forward, changing the orientation of the camera as well as object interaction. (Hermann et al., 2017) use the *DeepMind Lab* (Beattie, Leibo, Teplyashin, Ward, Wainwright, Küttler, Lefrancq, Green, Valdés, Sadik, et al., 2016) environment to define a language-guided navigation task where the agent has to reach an object indicated in a natural language instruction. (Gordon et al., 2018) present an interactive VQA task in AI2THOR (as shown in Figure 13) where the agent has to answer questions related to specific objects in the environment. The agent has to navigate in the environment to discover the objects and answer questions about them. Therefore the task involves both reasoning and navigation skills to be successfully solved. A navigation task for the CHALET environment, called LANI, has been proposed by (Misra et al., 2018a). In addition, they propose an instruction-following dataset called CHAI that requires the agent to execute a set of actions to complete a given “recipe” in the household domain.

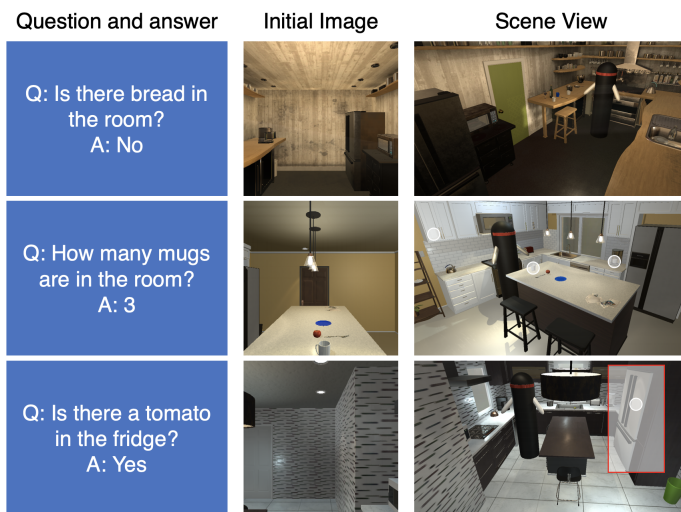


Figure 13: Example from the IQA dataset based on AI2THOR environment presented in (Gordon et al., 2018). The agent starts from an initial position in the environment and receives the question as well. It has to move in the environment in order to find the clues that will help it to answer the question.

Another line of work is interested in studying navigation tasks in simulated environments with photo-realistic scenes. Several datasets involving photo-realistic environments have been proposed: (Mirowski et al., 2019) introduce **StreetLearn** as a simulated environment to foster research in perception, planning, memory and navigation. The **Room-to-Room** dataset (Anderson, Wu, Teney, Bruce, Johnson, Sünderhauf, Reid, Gould, & Van Den Hengel, 2018b) is another relevant benchmark for navigation tasks in which the agent receives an instruction that tells it where to go and the agent has to generate a sequence of actions to reach the destination. (Chen et al., 2019) propose another Google Street View-based environment associated with navigation instructions that should be followed by the agent to reach a given goal position in the world. The **Room-across-Room** dataset (Ku et al., 2020) was designed to study more fine-grained abilities of the agent to

ground the instruction in temporally relevant referents of a photo-realistic 3D environment. This dataset also includes information in different languages to study how people refer to objects in space in different cultures. The dataset **TalkTheWalk** (De Vries et al., 2018) comprises a cooperative task involving a guide and a tourist. The guide is aware of the map and has to give instructions to the tourist about how to reach a specific goal position in New York City. This dataset represents a very complex challenge for AI agents because it requires both navigation and communication skills. **CerealBar** (Suhr et al., 2019a) is another collaborative game situated in a 3D environment where a follower has to collect a set of cards (with specific symbols) following the instructions of another agent. The follower has a limited set of actions that are mostly related to navigating the environment and picking up cards. Another collaborative task that requires natural language communication is **MeetUp** (Ilinykh et al., 2019a). In this task, one agent has to follow instructions and ask for feedback to reach a room that another agent is aware of. The navigator has to provide correct spatial references so that the other agent can give instructions about where to go. A variant of this interactive navigation task has been proposed by (Thomason et al., 2019a), namely **CVDN**. It is a bigger dataset comprising 2050 human-human navigation dialogues of 7k navigation trajectories composed of question-answer exchanges across 83 houses.



Figure 14: An example of a task to be solved in the ALFRED benchmark (Shridhar et al., 2020). The agent has to master several skills to succeed such as navigation, visual memory and object manipulation. Figure adapted from (Shridhar et al., 2020).

Such navigation-oriented tasks have a very important role in developing embodied agents with sophisticated reasoning and memory skills (for a more comprehensive list of these

tasks refer to (Gu, Stefani, Wu, Thomason, & Wang, 2022)). However, learning grounded meaning only by moving in the environment is rather limiting. In this way, the agent is not able to experience the effect of specific actions on specific objects. This is a concept that is related to the theory of affordances (Gibson, 1977) and its importance for learning grounded meanings of objects (Glenberg & Kaschak, 2002). (Chevalier-Boisvert et al., 2018) propose a virtual environment with a teacher that emits synthetic language generated by a context-free grammar. They define a curriculum for the agent based on increasingly complex tasks involving navigation (e.g., “navigate to a specific object”) and manipulation tasks (e.g., “open the door on your left”). The agent has to master preliminary levels to solve successive stages of the game. The game culminates with the *BossLevel* in which all the previous capabilities are required.

Another example of a task where the agent has a richer action set is **Collaborative Minecraft** (Narayan-Chen et al., 2019) where an *Architect* has a goal structure that wants to build and has to communicate specific instructions to the *Builder* that has to move around coloured blocks to create the desired shape. In this case, the resulting dialogues are complex and rich in spatial referential expressions required to specify the position of given game pieces. An up-to-date version of this task is presented in the IGLU challenge (Kiseleva et al., 2022). Another similar task that presents a similar blocks manipulation task is **BlocksWorld** (Bisk et al., 2016). However, the set of actions that the agent is able to execute in these tasks is rather limited. The same is true for another interactive task that was proposed by (Abramson et al., 2020). They define **Playroom**, a 3D environment where a learning agent has to perform several tasks of different complexities such as Q&A and instruction following. Again the action set is limited to navigation actions (e.g., *forward*, *back*, etc.) and one manipulation action (*grab*). This downside makes them an unsuitable benchmark for learning the meaning of specific actions that have an effect on the environment. *AI2Thor* (Kolve et al., 2017) is the only 3D environment with rich simulated physics that models a variety of actions for the agent such as *pick up*, *slice*, *toggle*. Based on this 3D environment, the research community has proposed several challenges. For instance, **ALFRED** (Shridhar et al., 2020) is the first language-guided task completion benchmark that requires both navigation and manipulation actions. As shown in Figure 14, the agent has to complete a goal specified by a Natural Language instruction. In addition to the overall goal instruction, instructions are provided for the agent for every sub-goal. Therefore, it can be intended as an interactive language game where the follower only executes the actions required to complete the task. In ALFRED, every task is very complex, as it is composed of very long action sequences (average trajectory length is 40) compared to other navigation tasks (average trajectory length is around 6). Each trajectory can be divided into a sequence of sub-goals required to achieve the final goal. Manipulating sub-goals might involve irreversible actions (e.g., slicing an apple). It is therefore essential for the agent to learn that specific actions should not be executed on specific objects. Another set of interactive tasks is presented in the **SILG** (Zhong et al., 2021) where they integrate under the same framework grid-like worlds such as NetHack (Küttler, Nardelli, Miller, Raileanu, Selvatici, Grefenstette, & Rocktäschel, 2020) as well as 3D simulated environment like Touchdown (Chen et al., 2019). They all share the same symbolic representation of the scene that incredibly simplifies the visual scene that an agent is supposed to receive in input. For example, the Touchdown environment (Chen et al., 2019) scenes have been

converted to symbols by segmenting its panoramas into semantic grids indicating pixels that belong to specific classes (e.g., sky, bicycle, etc.). This makes the visual grounding problem even harder considering that the model is not able to learn from a variety of signals coming from real-world scenes. **ALFWorld** (Shridhar et al., 2021) is another embodied AI dataset that benefits from language feedback and annotations that are automatically generated using the same game engine used for text adventure games. In this work, they stress the importance of teaching AI agents to “imagine” the sequence of instructions that can be executed to achieve a goal, and decouple this from the actual plan execution which depends on the underlying environment. Thanks to this new environment, they demonstrate that the agent can more efficiently learn language commands in text-world environments, and reliably generalise to unseen environments.



Figure 15: An example of a task to be solved in the TEACH challenge (Padmakumar et al., 2021). The Commander has to support the Follower in completing the task. The Follower is the only agent that can actually execute actions in the environment. Figure adapted from (Padmakumar et al., 2021).

In the ALFRED benchmark, it is assumed that the agent has perfect vision and is always able to find all the objects required to complete the task. Unfortunately, this is not realistic and the current setup does not allow the agent to resolve ambiguities or possible misunderstandings because there is not a real interaction between the agents. In this sense, promoting a *symbiotic relationship* (Rosenthal, Biswas, & Veloso, 2010) between the user and the agent, where both humans and artificial agents can solve tasks and mitigate each other’s limitations is essential for language learning. Natural language offers the user a medium to flexibly express their desired outcomes and provide guidance to avoid task failure; but as we also often provide underspecified, ambiguous, or even incorrect instructions, the agent is challenged to understand known aspects of their environment and pose informative questions when uncertain, in order to minimise the risk of failure. To promote action coordination and object ambiguity resolution, (Gao et al., 2022) created **Dialfred**, a set of 53K question/answer pairs that can be used to resolve ambiguities in the ALFRED environment. To the idea of collaboration between agents even further, (Padmakumar et al., 2021) propose the **TEACH** challenge, a two-player game situated in the AI2Thor environment. Inspired by ALFRED, there are multiple tasks of different complexity to be solved. As shown in Figure 15, there are two players with different roles: Follower and Commander. The Follower executes the instructions provided by the Commander while the

Commander knows the task to be solved and provides assistance to the Follower without actually executing any actions. The ultimate challenge of this dataset is to train an agent that can support a human in completing daily tasks. Finally, to advance the state-of-the-art in Embodied Conversational AI, (Gao et al., 2023) proposed the **Alexa Arena**, a novel multi-room simulated environment that was used as the main benchmark for the Amazon Simbot challenge, a university competition to advance the state-of-the-art in Embodied Conversational AI (Shi, Ball, Thattai, Zhang, Hu, Gao, Shakiah, Gao, Padmakumar, Yang, et al., 2023).

4. Grounded Language Learning Models

In this section, we will provide a critical survey and analysis of ML models that have been proposed to tackle some of the grounded language learning tasks reviewed in the previous sections. The search of the studies was focused on the topic of grounded language learning and was limited to the time period 2015-2022. We used Scopus¹ to complete the search, considering only published works in relevant Machine Learning or Natural Language Processing conferences. From the original search completed in 2019 we selected a total of 32 papers. Such papers were then extended with additional references published in major conferences in the field such as ACL, EMNLP, etc. Overall, we analysed a total of 50 studies according to specific criteria. These criteria have been selected because we consider them essential for artificial agents able to support the visually grounded interactive language games that we described in this work. The selected criteria are reported as follows:

- *Compositionality*: is the representation of a complex entity the result of the combination of simpler ones learned by the model?
- *Dialogue*: is the model able to support a conversation in natural language with another agent?
- *Architecture*: is the proposed approach designed as an end-to-end neural architecture or as a modular architecture?
- *Visual representation*: which kind of visual representation is used?
- *Reasoning*: does the model support any kind of reasoning mechanism able to leverage the learned representations?
- *Concept representation*: which kind of representations are learned and used by the system?
- *Few-shot learning*: Is the model able to solve a task by having access to a few learning examples?
- *Real-world vision*: is the model exposed to real-world images?
- *Natural Language*: is the model exposed to Natural Language (i.e., English content generated by real users)?

1. <https://www.scopus.com/>

A summary of the comparison is reported in Table 2. A more detailed explanation of the work is reported in the following section.

From the completed literature review it emerged that in studying grounded language learning, often one of the two modalities (i.e., vision and language) is lacking or is not considered at all. Indeed, several of the analysed papers either use real-world vision or synthetic images while focusing on more complex language structures. Bridging this gap is really important in order to understand how well current neural models are able to handle complex visual scenes, as well as interesting Natural Language phenomena available in human language annotations. Running experiments in the real world would be ideal to simulate such complex visual scenes—although this is not always practical. One way to cope with this limitation is by designing rich 3D environments that expose models to photo-realistic visual scenes. Instead, when analysing the language dimension, another simplification is that the tasks typically involve discriminative or generative tasks involving a single interaction. Only very few approaches consider interactive tasks where the dialogue history is actually relevant for the task at hand (Agarwal et al., 2020). In some cases, the developed models are actually able to learn a dialogue strategy to complete the task whereas others assume that the dialogue is somehow scripted in advance via a hand-crafted procedure.

4.1 Learned Multi-modal Representations

We are interested in understanding the type of learned representations these systems are learning, and the way they are using them for specific grounded language learning tasks.

4.1.1 VISUAL REPRESENTATIONS

A few approaches rely on low-level sensory features modelling RGB colour channels or shape information (Antunes et al., 2017; Patki & Howard, 2018). Based on the assumption that an image can be expressed as a bag of visual words (Csurka, Dance, Fan, Willamowski, & Bray, 2004; Grauman & Darrell, 2005; Sivic, Russell, Efros, Zisserman, & Freeman, 2005), several works (e.g., (Thomason et al., 2019b; Yu, Eshghi, & Lemon, 2017a)) use a set of visual classifiers that are manually mapped to specific classes that represent visual concepts. Due to the expressivity of deep neural networks, a common practice is to use latent features learned by a Convolutional Neural Network (CNN) trained on a large-scale image dataset (e.g., ImageNet (Deng et al., 2009)). Specifically, the features belonging to the last layer of a CNN, have been widely adopted due to their ability to be re-used for other tasks other than image classification. With such features, it is possible to represent an entire scene with a single distributed representation. However, as demonstrated in the VisDial challenge, using an object detector to predict object bounding boxes (e.g., *FastRCNN* (Anderson, He, Buehler, Teney, Johnson, Gould, & Zhang, 2018a)) really improves performance. Having such object-centric representations allows the agent to learn the task by reasoning over objects and their relationships which are potentially invariant across scenes. In this setup, the model receives in input a high-dimensional feature vector for every detected object (Region Of Interest (ROI) vector). Additionally, to provide the model with positional information about the objects, bounding box coordinates are encoded as additional features. To incorporate class information associated with each object, either class labels are encoded (Tan &

Paper	Compositionality	Dialogue	Architecture type	Vision module	Concept Representation	Extreme Generalisation	Real world vision	Natural Language
Agrawal, Gurumurthy, Sharma, Lewis, & Spector, 2019)		yes	end-to-end		Distributed	zero	yes	yes
Ahmed, Sajwan, Naveed, Ahsan, Sami, Vaseem, & Ghafoor, 2017)		hand-crafted	end-to-end	pretrained VGG-net features	Distributed	few	yes	synthetic
Balsharan, Hill, Leake, Hughes, Hasseni, Kohli, & Gerasidze, 2018a)			end-to-end	multiple SQA	Distributed	zero	yes	synthetic
(Balsharan, Hill, Leake, Hughes, Hasseni, Kohli, & Gerasidze, 2018b)			end-to-end	CNN	Distributed	zero	yes	synthetic
(Pati & Horned, 2018)		hand-crafted	end-to-end	Replative vision algorithms	Set of features	zero	yes	synthetic
(Chenille, 2019)			end-to-end	pretrained VGG-net features	Distributed	zero	yes	synthetic
(Yu, Zhang, & Xu, 2018a)			end-to-end	CNN	Distributed	zero	yes	synthetic
(Shikhar, Venkatesh, Baumgartner, Brunl, Plank, Bernardi, & Fernández, 2018)		yes	end-to-end	pernethal YQC/ResNet152	Distributed	zero	yes	synthetic
(Pillai & Matsuzek, 2018)		yes	end-to-end	CNN	Distributed	zero	yes	synthetic
(Ding, Wu, Wu, Hu, Luv, & Tan, 2018)		yes	end-to-end	RGB colour	visual classifiers	zero	yes	synthetic
(Yu, Wu, Cao, Torralba, Kohli, & Fernandez, 2018)		yes	end-to-end	pretrained VGG	distributed	zero	yes	yes (only one dataset)
(Wu, Baheti, Gholami, Lee, Farhadi, & Barrera, 2018)			end-to-end	Mask-CNN	Distributed, attribute-based	zero	yes	synthetic
(Wu, Baheti, Gholami, Lee, Farhadi, & Barrera, 2018)			end-to-end	CNN	Distributed, attribute-based	zero	yes	synthetic
(Chee-Bros, Gupta, Sanjey, Akari, De'Nevo, Ahmed, & Levine, 2018)			end-to-end	CNN + FLML	hand-crafted	zero	yes	synthetic
(Lai, Yang, Barrat, & Parikh, 2018)			end-to-end	CNN, Mask-CNN	distributed	zero	yes	synthetic
(Zhang, Yu, & Xu, 2018)		yes	end-to-end	CNN	Distributed	zero	yes	synthetic
(Mardziel & Abbeel, 2018)		multi-agent communication	end-to-end		Distributed	zero	yes	synthetic
(Shah, Fiser, Fiser, Koo, & Hukkeri-Tur, 2018)		yes	end-to-end	CNN (dataloader and segmentation map)	Distributed	zero	yes	synthetic
(Miera, Gristack, Ferrás, Heller, Gupta, & van der Maaten, 2018a)			end-to-end	CNN resnet	Distributed	zero	yes	yes (only one dataset)
(Miera, Gristack, Ferrás, Heller, Gupta, & van der Maaten, 2018b)			end-to-end	CNN	Distributed	zero	yes	yes
(Khan, Chatterjee, Hara, 2018)			end-to-end	CNN	Distributed	zero	yes	yes
(Berrou-Blanc, Christensen, & Janczarek-Lopez, 2018)		yes	end-to-end	protog-representation hand-crafted	Horn classes	yes	yes	yes
(Kafory, Lee, Kamma, & Barrá, 2018)			end-to-end	CNN	Distributed	zero	yes	yes
(Yu, Liu, Zhang, & Xu, 2018a)			end-to-end	CNN	Distributed	zero	yes	yes
(Thomson, Pathakumar, Shajoy, Walker, Jiang, Venkatesh, Hart, Stone, & Mooney, 2019b)		yes	end-to-end	visual classifiers	visual classifiers	zero	yes	emerged language
(Choi, Lazaridou, & de Freitas, 2018)			end-to-end	CNN	Distributed	zero	yes	emerged language
(Xie et al., 2018)			end-to-end	CNN	Distributed	zero	yes	synthetic
(Gardar et al., 2018)			end-to-end	CNN	Distributed	zero	yes	synthetic
(Zhang et al., 2018)			end-to-end	CNN	Distributed	zero	yes	synthetic
(Mansour, Shrivastava, Lee, Anderson, Parikh, & Barrá, 2018)		yes	end-to-end	CNN or Mask-CNN	Distributed	both	yes	synthetic
(Zhang Li, Hu, Yang, Zhang, Wang, Choi, & Cao, 2021)			end-to-end	FastRCNN	distributed	both	yes	synthetic
(Padoletti, Schmid, & Sun, 2021)			end-to-end	CNN	distributed	both	yes	yes
(Min, Chagnel, Revellam, Bisk, & Sakshidhar, 2021)			end-to-end	CNN	distributed	both	yes	yes
(Magregorán, Sagun, & Kwon, 2021)			end-to-end	hybrid	distributed	both	yes	yes
(Tan & Bansal, 2019a)			end-to-end	FastRCNN	distributed	both	yes	yes
(Zhang, Hu, Qiu, Shaw, & Sun, 2021)		yes	end-to-end	FastRCNN	distributed	both	yes	yes
(Tan, Wang, & Bansal, 2019)		yes	end-to-end	FastRCNN	distributed	both	yes	yes
(Bian, Gu, Murphy, & Fain, 2019)		yes	end-to-end	CNN	distributed	both	yes	yes
(Zellers et al., 2021)		no	end-to-end	attribute-representation	distributed	both	yes	yes
(Zellers, Liu, Hessel, Yu, Park, Cho, Fathall, & Choi, 2021)		yes	end-to-end	CNN	distributed	both	yes	yes
(Simple, Deng, & Ahn, 2021)		no	end-to-end	CNN	distributed	both	yes	yes
(Radford, Kim, Hallacy, Ramesh, Goh, Agarwal, Sidor, Askel, Mishkin, Clark, et al., 2021)		no	end-to-end	CNN	distributed	both	yes	yes
(Ranzini, Pavlov, Goh, Gray, Vass, Radford, Chen, & Sutskever, 2021)		no	end-to-end	CNN	distributed	both	yes	yes
(Lazaridou et al., 2018)		no	end-to-end	CNN	distributed	both	yes	yes
(Lazaridou, Charatan, Ferrettes, & Barrat, 2016)		no	end-to-end	CNN	distributed	both	yes	yes
(Lazaridou, Charatan, Ferrettes, & Barrat, 2016)		no	end-to-end	CNN	distributed	both	yes	yes
(Lazaridou, Charatan, Ferrettes, & Barrat, 2016)		no	end-to-end	CNN	distributed	both	yes	yes
(Lazaridou, Charatan, Ferrettes, & Barrat, 2016)		no	end-to-end	CNN	distributed	both	yes	yes
(Zhou, Niu, & Bai, 2021)		no	end-to-end	CNN	distributed	both	yes	yes

Table 2: Comparison of several multimodal models developed for several grounded language learning tasks.

Bansal, 2019a) or a probability distribution over the object detector label set (Zhou et al., 2020). This provides the model with more fine-grained information associated with the scene as well as a better representation of the objects in it. This approach was used by several recent models (e.g., (Xie et al., 2018; Lu et al., 2018; Yi et al., 2018; Tan & Bansal, 2019a; Zhou et al., 2020)). However, extracting features from an object detector is a very time-consuming operation and limits the ability of the model to encode specific objects that are not part of the object detector training set. Therefore, some approaches revisited grid features generated by a CNN to improve the efficiency of multi-modal models (e.g., (Jiang, Misra, Rohrbach, Learned-Miller, & Chen, 2020; Kim, Son, & Kim, 2021)).

4.1.2 LANGUAGE REPRESENTATIONS

Language representations, on the other hand, are typically trained separately, either using the reference task dataset or by pretraining using self-supervised techniques. Such representations are either in the form of word embeddings (e.g., (Mikolov, Yih, & Zweig, 2013; Pennington, Socher, & Manning, 2014)) or contextual embeddings (e.g., (Devlin, Chang, Lee, & Toutanova, 2018)). The latter have had tremendous success in both Natural Language Understanding, and Natural Language Generation tasks (e.g., GPT-3 (Brown, Mann, Ryder, Subbiah, Kaplan, Dhariwal, Neelakantan, Shyam, Sastry, Askell, et al., 2020)).

4.1.3 VISION+LANGUAGE REPRESENTATIONS

Once we have a representation for each modality, how do we derive a *multimodal representation*? A multimodal representation is a representation of data that is obtained by combining multiple sources of input such as visual, perceptual or symbolic information (Baltrušaitis, Ahuja, & Morency, 2018). Representing multiple modalities poses many difficulties regarding the way we combine heterogeneous sources and how to deal with noise in them (i.e. when they are just noisy or completely missing). Therefore, designing mechanisms to meaningfully represent such data becomes a crucial problem in multimodal representation learning. In recent work by (Baltrušaitis et al., 2018), several multimodal fusion mechanisms have been reviewed:

- **early**: the model learns multi-modal representations as part of the learning objective by jointly minimising a loss function that optimises a posterior involving both modalities;
- **middle**: the model uses independent training data and objectives for specific modalities. A parameterised similarity function can be used to learn the relationships between features in multi-modal spaces;
- **late**: it is the converse of the previous approach because it first computes the similarity scores between the representations and then it aggregates them together.

As highlighted by (Kiela, 2017), a combination between *early* and *middle* fusion can be considered plausible from a cognitive perspective because “it allows for learning unimodal representations independently but which also allows for combining said representations into an overall multimodal one that takes all modalities into account, possibly in varying degrees”.

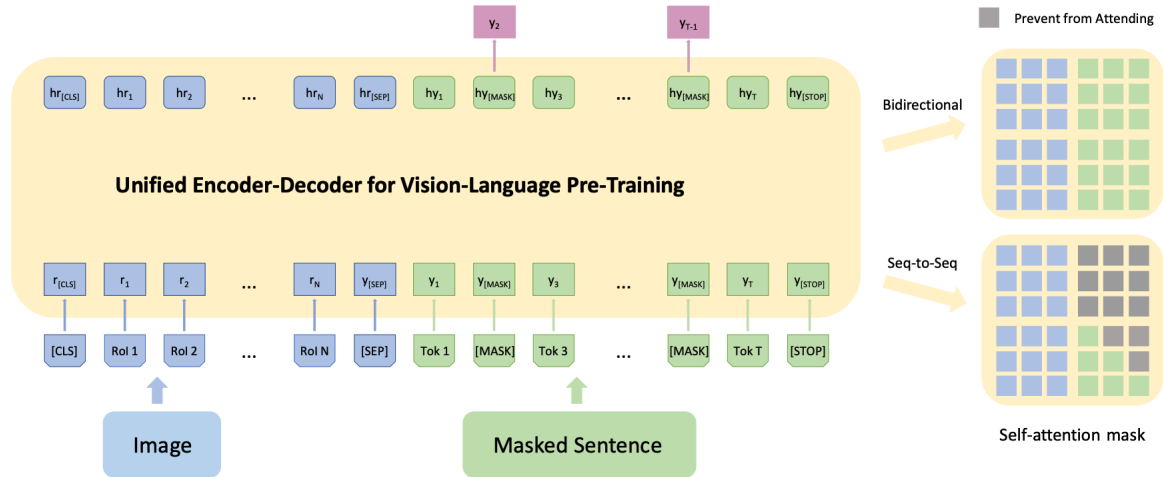


Figure 16: An overview of the VLP architecture presented by (Zhou et al., 2020). It uses a single Transformer stack to encode both visual tokens and textual tokens. Then, it is pretrained on Conceptual Captions using both masked language modelling and sequence-to-sequence loss. Figure adapted from (Zhou et al., 2020).

For instance, early fusion models have been relatively successful in multimodal tasks thanks to the development of large-scale unimodal encoders. Concatenating, averaging or summing unimodal representations are some of the methods used to generate multimodal representations (e.g., (Shekhar et al., 2019; Das et al., 2018)). More powerful multimodal architectures have been proposed that use specific co-attention mechanisms (e.g., (Deng et al., 2018; Yu et al., 2018a; Lu et al., 2018)). More recently, Transformers have been used to learn multimodal representations by combining modality-specific inputs via self- and cross-attention. Such models are indeed closer to the early/middle fusion models identified by (Kiela, 2017).

Transformer-based models have been particularly successful in most of the grounded language learning tasks that we reviewed. Following (Bugliarello, Cotterell, Okazaki, & Elliott, 2021), we can define two different types of multimodal architectures: 1) *single-stream* Transformers; and 2) *dual-stream* Transformers. Single-stream Transformers assume that there is a single Transformer stack that receives in input both visual and textual inputs. Dual-stream Transformers instead assume that each modality is encoded by a modality-specific encoder whose outputs are fused by using a dedicated cross-modal encoder. This distinction was similarly described by (Baltrušaitis et al., 2018) as *joint encoder representations* and *coordinated encoder representations*, respectively. However, (Bugliarello et al., 2021) demonstrates dual-stream attention functions act as restricted versions of the attention function in any single stream Transformer architecture. So dual-stream architectures have an inductive bias that favours interactions between modalities.

Single-stream Multimodal Transformers One of the first single-stream multimodal models, together with many others (e.g., (Chen, Li, Yu, Kholy, Ahmed, Gan, Cheng, & Liu, 2020; Li, Yatskar, Yin, Hsieh, & Chang, 2019; Li, Yin, Li, Hu, Zhang, Zhang, Wang,

Hu, Dong, Wei, Choi, & Gao, 2020)), is VLP (Zhou et al., 2020). Its architecture is depicted in Figure 16. First of all, the architecture mimics the one used by BERT (Devlin et al., 2019a). Following BERT, it represents the language input as Wordpieces (Schuster & Nakajima, 2012). It introduces *visual tokens* which are derived from object-centric representations generated by FastRCNN. Such visual tokens are then projected to a dimensionality comparable to textual tokens. In this way, they can be passed in input to the stack of Transformer layers to generate multimodal representations of the input. Following the success of self-supervised learning for text-based Transformers, VLP is initialised with a pretrained BERT and then is pretrained using several multimodal losses on the Conceptual Captions dataset (Sharma, Ding, Goodman, & Soricut, 2018). In order to support both discriminative and generative tasks, VLP is trained with both masked language modelling loss as well as with a sequence-to-sequence loss (they are randomly alternated during training). In this way, it is one of the few models that can be used for both VQA and image captioning. Other models have also proposed additional self-supervised losses for the visual modality as well. For instance, *UNITER* (Chen et al., 2020) uses masked region modelling which consists in reconstructing the original visual tokens or predicting their class. By learning to predict missing objects, the model learns to model the visual context in a better way which can enhance language understanding. Another model that belongs to this family is FLAVA (Singh, Hu, Goswami, Couairon, Galuba, Rohrbach, & Kiela, 2022), which takes advantage of the advancements in Vision Transformers (Dosovitskiy, Beyer, Kolesnikov, Weissenborn, Zhai, Unterthiner, Dehghani, Minderer, Heigold, Gelly, et al., 2020) to learn representations of the image directly from pixels without any object detector features. Both language and visual features are then fed into modality-specific Transformers that derive representations for each modality. Then, a single cross-modal Transformer derives multimodal representations. The main novelty of FLAVA is represented by its ability to perform both multimodal and unimodal tasks. In particular, it is pretrained with both multimodal and unimodal datasets so that it can simulate scenarios where only some modalities are available. Additionally, it can leverage the additional signal that is derived from both unimodal and multimodal task-specific losses.

Dual-stream Multimodal Transformers In contrast with single-stream multimodal Transformers, many dual-stream transformers have been proposed in the literature. As highlighted before, these models have two dedicated modality-specific encoders that are used to encode language and vision separately. An additional cross-modal Transformer block is then used to “fuse” the hidden representations generated for the two modalities. One instance of these models is LXMERT (Tan & Bansal, 2019b) which has a “Object-relationship Encoder”, a dedicated Transformer block that encodes object representations derived from a pretrained object detector, and a “Language encoder” implemented as a Transformer block that encodes language tokens. Additionally, a “Cross-modality Encoder”, implemented as another Transformer block, is used to fuse the two modalities and favours the development of cross-modal representations that are used for the downstream tasks. This process is performed in a symmetric way. First, language representations are used as queries in the self-attention operations against the vision token representations. Then, vision token representations are used as queries in the multi-head self-attention operation. Concurrently, ViLBERT (Lu, Batra, Parikh, & Lee, 2019) defined an alternative way of

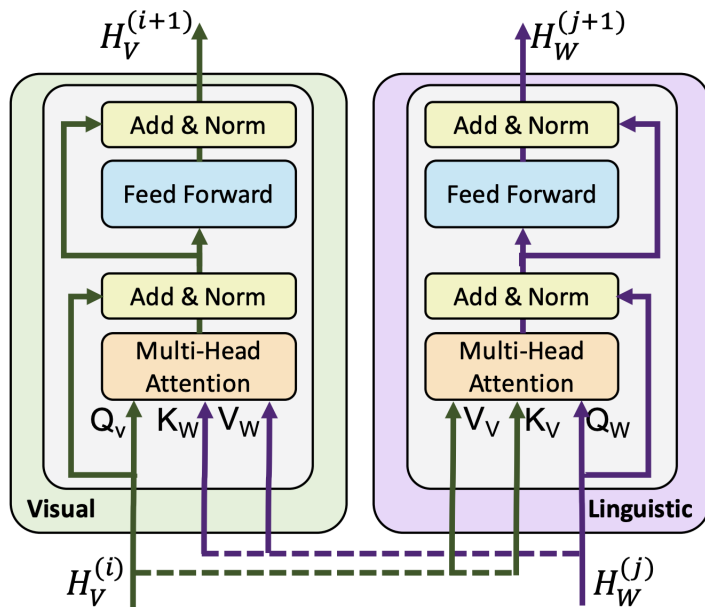


Figure 17: The co-attention layer implemented in ViLBERT to fuse vision and language representations. The co-attention layer relies on multi-head self-attention, and is applied in parallel to both vision and language representations to derive multimodal representations. Figure adapted from (Lu et al., 2019).

using cross-attention to “fuse” the vision and language modalities. This is implemented as a two-step approach where, after using modality-specific encoders to generate latent representations for language and vision respectively, the model uses a ‘co-attention’ transformer layer (showed in Figure 17) for a target modality that performs multi-head self-attention using as queries the representations from the target modality, and as keys and values the representations associated with the other modality.

4.2 Compositionality

Natural Languages are characterised by immense *productivity* meaning that they can generate an infinite number of expressions. According to linguists, this is possible thanks to the property of compositionality. Of specific interest is *semantic compositionality*, the principle whereby the meaning of a linguistic expression is a function of the meaning of its components and the rules used to combine them (Frege, 1892; Montague, 1970). In this paper, we intend a broader definition of compositionality that is related to the ability to combine simpler features to obtain novel objects/entities. This broader definition can be extended to compositionality in vision where an object can be seen in terms of its attributes (i.e., cat = has fur, has eyes, has 4 legs, etc.). When studying the generalisation power of artificial agents, it is fundamental to consider compositionality because it can be conjectured to be a landmark not only of language but of human thought in general (Fodor & Pylyshyn, 1988; Lake, Ullman, Tenenbaum, & Gershman, 2017). Following (Harnad, 1990), this broader definition of compositionality is the key ingredient of learning grounded representations. As

demonstrated in the GROLLA evaluation framework (Suglia, Konstas, Vanzo, Bastianelli, Elliott, Frank, & Lemon, 2020), learning grounded representations should result in high performance for in-domain examples as well as generalisation to examples never seen during training.

Several papers take into account only some aspects of compositionality which are not fully integrated into the learned representation nor in the model reasoning mechanism. The work presented by (Choi et al., 2018) studies how, in a referential game, two agents evolve their linguistic skills to generate a compositional language. In particular, they develop a two-player image description game: the speaker and hearer receive two random images. The speaker needs to generate a sequence of characters to describe the image that they see. The hearer looks at its image and needs to decide if they are looking at the same image or not. They base their model on the *Obverter model* (Batali, 1998). The obverter technique motivates an agent to search over messages and generate the ones that maximise their own understanding. They study the properties of the language that evolved from the communication game. Additionally, they complete a zero-shot evaluation by holding out five objects from the dataset during the training and observing how agents describe them during the test phase. Discriminative features in this case are just colour and shape (with a strong bias towards the colour) so the task cannot tell to what extent these networks are able to compose the learned representations. (Mordatch & Abbeel, 2018) study the properties of the evolved communication protocol in a population of agents that can interact with each other via either verbal or non-verbal communication (i.e., pointing or guiding). The environment contains a fixed number of agents and landmarks. Every agent can either execute an action (i.e., go to location, look at location and do nothing) or utter a symbol. Every agent has a goal that is not revealed to the other agents requiring communication skills in order to be successful.

Differently from these models that try to analyse to what extent neural networks are able to learn grounded compositional language, the work presented by (Thomason et al., 2019b; Kottur et al., 2019; Yi et al., 2018) includes logical form as a structural bias in their models. (Thomason et al., 2019b) uses a semantic parser—separately trained from the agent—to derive semantic representations of the language in the form of predicative structures.

Powerful generalisation capabilities can be obtained when compositional representations are directly integrated in the model. For instance, (Yi et al., 2018) presents a model that learns to answer questions about images by “executing” a sequence of predefined functional programs. First, it learns to generate a rich scene graph for the reference image and, as a second step, generates a functional program for the given question that, once executed against the scene graph, returns the answer. (Kottur et al., 2019) uses CorefNMN (Kottur, Moura, Parikh, Batra, & Rohrbach, 2018) for the CLEVR-dialog task with an approach similar to (Yi et al., 2018). CorefNMN is a neural model equipped with a *program decoder* that generates functional programs given the dialogue history and the current question. Then, it derives the answer by executing the extracted functional programs. A specific program “Refer” can be used by the model to have access to a pool of previously mentioned entities in order to complete visual coreference resolution over the turns of the dialogue. Integrating these functional programs in the model guarantees a high level of generalisation

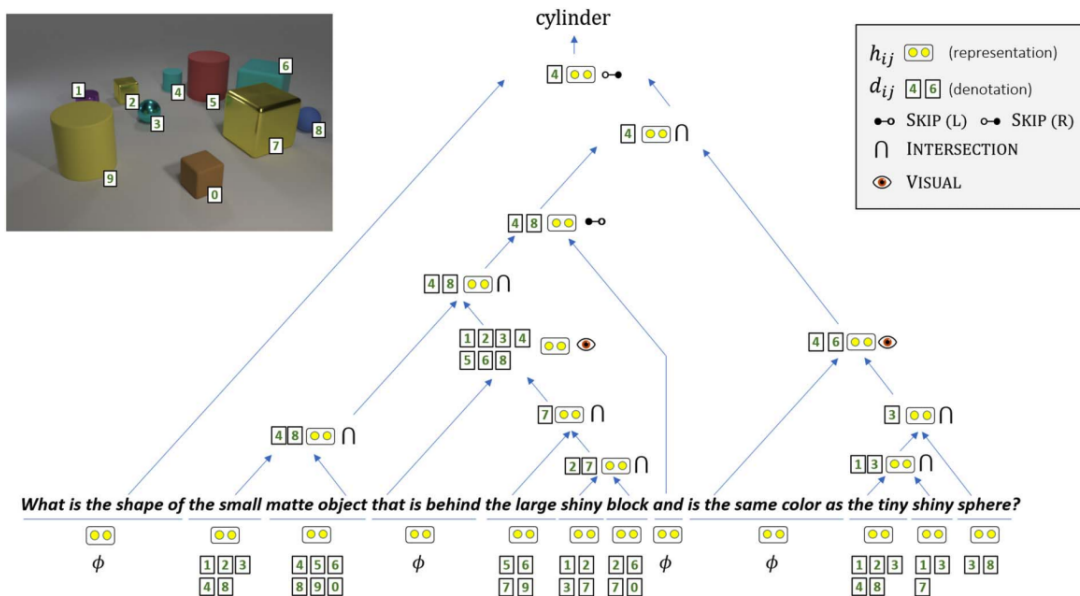


Figure 18: The Neural compositional module developed by (Bogin et al., 2021b). It uses a CKY-like algorithm to derive, in a bottom-up fashion, a compositional representation of the question that is then used, via its denotation, to generate the answer. Figure adapted from (Bogin et al., 2021b).

because the model does not focus on learning spurious correlations between the dataset but it learns to effectively use the provided “modules” to complete the task.

Relying on ad-hoc symbolic representations to facilitate the generation of compositional representations was proposed first by (Socher, Lin, Ng, & Manning, 2011) and then revisited by (Bogin et al., 2021b). As shown in Figure 18, (Bogin et al., 2021b) use a CKY-based algorithm (Kasami, 1966; Younger, 1967) that starts by grounding the single word embeddings associated to the words in the question. After this step, in a bottom-up fashion, compositional representations are induced from sub-span latent representations. They use a weighted sum of sub-span representations to make the computation differentiable. Each sub-span weight is derived as an attention score that represents how suitable that span is for a given context. Similar in spirit, a denotation for each span is used to derive the meaning of the sub-expression, and ultimately to derive the answer. Following the work by (Andreas, Rohrbach, Darrell, & Klein, 2016), they have several hand-crafted modules for executing operations on the object set that the model learns to execute. This model seems to do very well on the synthetic dataset CLEVR (Johnson et al., 2017) but, when tested on questions written by humans, it generalises marginally better than a fully neural counterpart (e.g., *MAC* (Hudson & Manning, 2018)). This might be because some operators (e.g., computing the maximum value in a set) cannot be obtained by composing elements in the question, but instead, require higher-order reasoning. The other problem is that more sophisticated grounding is required in some cases. Their model only grounds specific phrases to objects. However, in some cases, more fine-grained properties or numbers have to be grounded which is something that the model doesn’t support. The same *ground-compose*

approach was implemented in the work presented by (Zhang et al., 2021). They apply a syntactic parser on a caption to generate a symbolic representation that will guide the composition of grounded elements. First, specific noun phrases are grounded via dedicated modules to obtain a visually-grounded representation. Then, the symbolic representation is used to combine such grounded representations together.

All of the analysed papers rely on distributed representations for the words that completely ignore the extralinguistic information associated to a concept and only capture information about the use of the word in context (Westera & Boleda, 2019). In fact, these models do not have a representation for the concept of `cat` that defines a cat as an animal with 4 legs, a nose, and coloured fur, etc. Even multi-modal models can partially incorporate in their latent representation information related to the colour and to specific shapes that can be associated with specific concepts. However, these models do not seem to be able to take into account more complex attributes that people learn about concepts. In this sense, a compositional representation of `cat` should be derived by integrating specific sensory-motor properties (e.g., visual, auditory, etc.) as well as situational properties (e.g., affordances) that uniquely characterise a cat. Additionally, because such learned representations are typically encoded in a single vector representation, it is harder to understand what the relevant factors of variations are, and how they have been combined together. For instance, (Gardenfors, 2014) in its *Conceptual Spaces* theory defines an *object category* as composed of domains, a set of convex regions in such domains, and information about mereonomic (part-whole) relations as well as relations between dimensions across different domains. It is very possible that distributed representations derived from CNNs are able, for example, to extract information about object colours, which resemble one possible domain identified in (Gardenfors, 2014)’s theory. However, it is not clear to what extent distributed representations are expressed in terms of relational information. Such relational knowledge can be intended as the essence of compositionality because it would allow representations to be derived by composing simpler ones together.

4.2.1 ASSESSING COMPOSITIONALITY

Despite neural networks’ impressive performance on several benchmarks, many researchers have advocated that their representations and reasoning mechanisms are still very brittle (Marcus, 2018). As stated in (Linzen, 2020), current Deep Learning models are evaluated using a protocol that does not assess their ability to generalise to out-of-distribution data. This is a capability akin to systematic generalisation, and is therefore a core property of compositionality. Previous work on visual question answering (VQA) measures generalisation to novel question-answer pairs generated for natural (Agrawal et al., 2017; Whitehead, Wu, Ji, Feris, & Saenko, 2021), or synthetic images (Bahdanau, Murty, Noukhovitch, Nguyen, de Vries, & Courville, 2018b; Johnson et al., 2017). Similarly, (Suglia et al., 2020), proposes an evaluation framework that augments a goal-oriented evaluation metric with two auxiliary tasks aimed at assessing systematic generalisation. Prior work has examined systematic generalisation for image captioning (Atzmon, Berant, Kezami, Globerson, & Chechik, 2016; Nikolaus, Abdou, Lamm, Aralikkatte, & Elliott, 2019; Bugliarello & Elliott, 2021). (Ruis et al., 2020) examines compositionality in the context of language-guided action execution in grid world environments extending the framework developed by (Lake, 2019).

Systematic generalisation is not the only property that a compositional system can demonstrate. For instance, (Pantazopoulos, Suglia, & Eshghi, 2022) extends the multifaceted compositional evaluation framework presented by (Hupkes, Dankers, Mul, & Bruni, 2020) to image captioning. According to their definition, an image captioning model can be considered *compositional* if possesses 3 different properties: (1) *systematicity*: the ability to generalise to unseen combinations of concepts learned in isolation during training; (2) *productivity*: the capacity to extend predictions beyond the training observations; and (3) *substitutivity*: the robustness of predictions under synonym substitution.

By analysing a model under the above-mentioned compositional frameworks, we can assess its ability to be robust to unseen examples. Therefore, concepts such as compositionality and robustness, are connected to the concept of visual grounding. In the context of visual question answering, there have been several other works that have tested the robustness of V+L models. For instance, (Hudson & Manning, 2019) proposes a VQA benchmark that augments the VQA accuracy with robustness measures such as *plausibility*, *validity*, and *consistency*. The work presented in (Thrush et al., 2022) assesses the ability of the model to discriminate correct (image, caption) pairs from foil ones that are obtained by rearranging nouns and adjectives. CrossVQA (Akula, Changpinyo, Gong, Sharma, Zhu, & Soricut, 2021) presents a collection of test splits for assessing VQA generalisation in a variety of visual domains spanning from photographs (e.g., VQA v2) to real-world images taken by visually impaired users (e.g., VizWiz (Gurari, Li, Stangl, Guo, Lin, Grauman, Luo, & Bigham, 2018)).

4.3 The importance of Dialogue for Language Learning

We argue that Interactive Grounded Language Learning is a plausible paradigm for training embodied agents to understand and use Natural Language. However, few of the papers that we analysed try to solve a task involving a multi-turn conversation between two agents. Additionally, the analysed papers can be further categorised based on the degree of interaction that they implement.

4.3.1 SCRIPTED DIALOGUE POLICY

Dialogue can be *scripted* such as in the work by (Patki & Howard, 2018) and (Antunes et al., 2017). In this case, the dialogue policy follows a predefined script that allows the model to complete the reference task. A modification of this setup involves having dedicated dialogue managers *to learn when to learn* (Yu et al., 2017a; Thomason et al., 2019b). Specifically, the dialogue policy relies on specific intents to detect when the model should update its internal representation as a result of an error in classification.

4.3.2 LEARNED DIALOGUE POLICY

A more relevant setup for this work is one where the agent is supposed to learn *when* and *what* to generate directly from the data. For instance, (Agarwal et al., 2019) propose a collaborative game where two bots are present. One of them – A-bot – receives an image and a caption. The other – the Q-bot – receives only the caption. Given the caption, it has to generate the image after asking the other bot some questions about it. In this way, it can create a representation that is suitable for retrieving, ranking or generating that image.

In this way, no human supervision is required and a reward can be defined in terms of how close to the target the generated representation is (based on Euclidean distance). To tackle the *GuessWhat?!* dataset (see Section 3.3), (De Vries et al., 2017) propose to train separate models for the 3 different tasks that need to be solved namely question generation, target object guessing and answer generation. (Shekhar et al., 2019) propose an end-to-end neural model that uses a multi-task learning approach to jointly train the question generation and the guesser module. However, their model treats the image as a single feature vector generated by a pretrained CNN ignoring more fine-grained information present in real-world scenes.

An approach that tries to more directly integrate the image features in the process of discovering the target object was presented in (Deng et al., 2018). They propose a model that implements an iterative attention process that tries to co-ground the query feature to object-centric features associated with the image. However, they take into account only the guessing problem reducing its applicability to the full set of tasks that has to be solved in *GuessWhat?!*. A similar setup is also described in visual dialogue (e.g., (Kottur et al., 2019; Das, Kottur, Gupta, Singh, Yadav, Moura, Parikh, & Batra, 2017b)). For instance, CLEVR-Dialog (Kottur et al., 2019) has been defined to study visual coreference resolution in a multi-turn dialogue involving two agents. However, their main interest is to define models that are able to correctly track the state of the conversation rather than *actually generating* the questions/answers.

4.3.3 MULTI-TURN INTERACTION

Another important element is the relevance of multi-turn interaction and what is the granularity of the learning task. Specifically, some papers consider only single-turn interactions while others assume that the interaction can last up to N steps.

Papers studying Natural Language evolution (e.g., (Lazaridou et al., 2020, 2016; Choi et al., 2018; Havrylov & Titov, 2017)) can be considered as instances of a simplified dialogue namely a signalling game (Lewis, 1969). Here, they limit the interaction to single-turn exchanges where an agent (Speaker) generates a message and the other (Hearer) has to select the target image among a set of distractors. For instance, (Lazaridou et al., 2018) studies the communication protocol developed by agents equipped with different types of visual encoders. In these studies, they are interested in the type of language developed by the agents as a result of the task completion objective.

For instance, a system that tries to learn word meanings in an interactive fashion was presented by (Zhang et al., 2018). They present a synthetic task where an agent needs to learn the mapping between an image of an object, and a phrase, as shown in Figure 19. The images depict small-sized animals. They trained the model in a one-shot learning regime where the model learns to classify a test instance by relying on a single training instance of that class. To implement their model, they rely on a multimodal variant of the Neural Turing Machine read-write memory mechanism (Graves, Wayne, & Danihelka, 2014). Specifically, it learns to associate a specific word (e.g., “cat”) with a corresponding visual representation (e.g., “cat image”). In this paper, the dialogue involves learning multiple object categories at the same time. A single turn is required to learn the association between an image and a given category. Furthermore, the language used was synthetically generated by a

grammar and therefore lacks the diversity of Natural Language. Despite these limitations, they tackle an important problem in the Machine Learning community namely *Few-shot Learning* (Wang & Yao, 2019): the agent needs to learn to maximise the performance P for a task T using a limited amount of experience E about it. In addition, they use a training strategy for their agent that combines both Imitation Learning and Reinforcement Learning. In particular, they assume that early in the training the agent is incapable of generating meaningful responses so it imitates what the teacher says. Later on in the training, it starts generating more interesting responses so they use Reinforcement Learning to provide a more precise signal about the quality of the generated response. This will give the agent the chance to learn to “babble” and to correctly solve the conversational game.

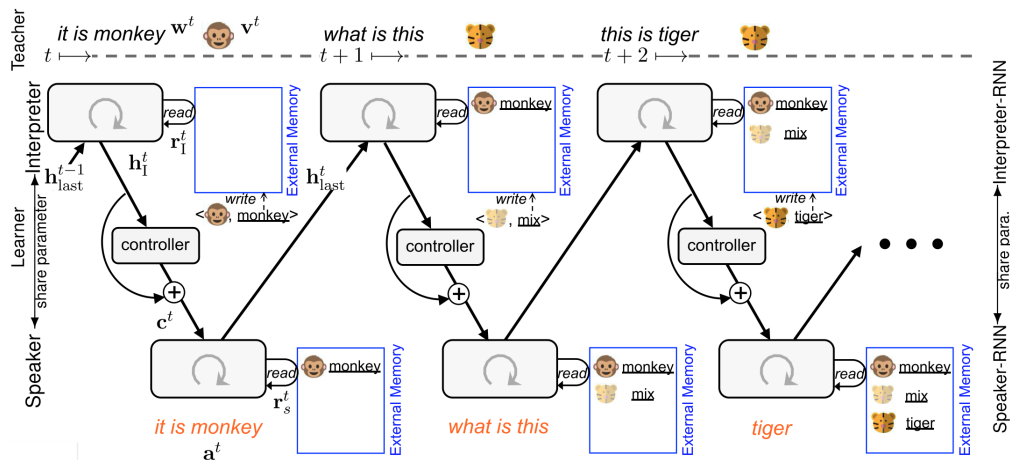


Figure 19: Model presented by (Zhang et al., 2018). The Speaker learns to mimic the Teacher by generating a statement associated with the images available. The Speaker has to learn animal concepts (and fruits) in a one-shot learning regime. After a single example of the class is shown to the Speaker by the Teacher, they are asked to classify an image associated with the same concept.

A model able to support a conversation composed of multiple turns can automatically generate extra training data from a very small bootstrap set. For instance, the work presented by (Misra et al., 2018b) describes a *learning by asking* paradigm where the model improves its VQA accuracy by learning to ask relevant questions about a given image. In this way, the model is able to retrieve data points that go beyond the training set and therefore learn better representations for the downstream task.

Depending on the task, it is debatable whether dialogue history plays an important role in the decision-making strategy of the learning agent. Particularly, tasks requiring single-turn interaction do not model the history in the first place. Additionally, even tasks based on multi-turn interaction do not necessarily require dialogue history. For instance, VisDial (Das et al., 2017b) is designed as a dialogue task but the questions that the humans develop can be answered by simply looking at the image. (Agarwal et al., 2020) show that models having sophisticated history encoding mechanisms bring very marginal improvements due to the fact that only a portion of the data *actually* requires the dialogue context. Therefore,

special attention should be paid to the task design to make sure that dialogue history actually matters.

4.4 Embodiment & Grounding

In the analysed papers a common trend emerges. In most of the studies, the reference for grounding is an image that represents the context of a *situated* interaction. As discussed in the previous section about action execution tasks (see Section 3.3), 3D simulators and virtual environments have been used to study the problem of grounding in visually rich contexts. (Yu et al., 2018a) propose *Guided Feature Transformation* which is a grounded language model extending classical CNN with an end-to-end differentiable component able to model the interactions between language and vision. They compute a recursive relationship where the input representation generated by the CNN is refined after J steps according to the words in the input command. They design several navigation tasks, as shown in Figure 20, in a 2D/3D environment: 1) *nav*: navigate to a specific object; 2) *nav_nr*: navigate to an object near the specified one; 3) *nav_bw*: navigate to a location between two objects; 4) *nav_avoid*: navigate to any object but the specified one; and 5) *nav_dir*: navigate to an object specified by a relative direction w.r.t. another object.




Figure 20: Examples of navigation tasks solved in the work presented by (Yu et al., 2018a).

(Chevalier-Boisvert et al., 2018) propose a first model for the BabyAI environment. They use a GRU module to model the language instruction and a FiLM (Perez, Strub, De Vries, Dumoulin, & Courville, 2018) for the vision module which learns to condition the visual representation on the question representation. The agent accuracy on these tasks seems to be very high (nearly 100%) posing some doubts related to the complexity of the world and of the language expressions used. Instead, (Bahdanau et al., 2018a) describes a learning framework for instruction-conditional RL agents having to execute instructions in a grid world. In their RL setup, the agent learns from a reward that is derived from a reward model trained from expert demonstrations. The proposed model is able to learn a reward model that performs well from a limited amount of data. Specifically, with a training set of only 8000 examples, the agent could reach a performance of 60% that can be improved to reach optimal performance when a set of 100,000 examples is used to train the model.

Motivated by the desire to create artificial agents that can follow instructions in domestic environments, there have also been several efforts transitioning from grid-world setups to benchmarks for language-guided navigation in indoor 3D environments with photo-realistic images. Several models have been proposed in the literature to solve this problem. For instance, several Transformer-based models have been designed for this task (e.g., (Majum-




Figure 21: Architecture of the Episodic Transformer model presented by (Pashevich et al., 2021). It jointly encodes language instruction, camera observations and previous actions using a dedicated multimodal encoder. The hidden state representation of the last state is then used to predict the action and the object class. Figure adapted from (Pashevich et al., 2021).

dar, Shrivastava, Lee, Anderson, Parikh, & Batra, 2020b; Magassouba et al., 2021; Hong, Wu, Qi, Rodriguez-Opazo, & Gould, 2020)). An important ingredient for all of these models is a careful pretraining design. Specifically, to deal with the complexity of the language instructions, language-only pretrained weights are used (e.g., BERT (Devlin, Chang, Lee, & Toutanova, 2019b)). Then, large-scale image-text pairs from the Web are leveraged to train the model. Using such additional pretraining data helps finetuning the cross-modal encoder that will be used for the downstream task evaluation in a photo-realistic environment such as Matterport3D (Chang et al., 2017).

Several projects have tried to investigate visual question-answering capabilities in 3D environments. For instance, the work presented by (Gordon et al., 2018) introduces the IQA dataset that requires the agent to move in the environment in order to answer the questions. It extends VQA in several ways: 1) navigation in the environment; 2) understanding of the environment (object, actions, affordances); and 3) executing actions conditioned on the question. The main question types are 1) counting; 2) existence; and 3) spatial relationships. They present *Hierarchical Interactive Memory Network (HIMN)*, a Deep Learning model that incorporates a hierarchy of modules: a *Planner* intended as the decision-making component responsible for activating several *task-specific modules*. The designed sub-modules are 1) navigator; 2) scanner; 3) detector; 4) manipulator; and 5) answerer. HIMN is unable to differentiate between an object being inside a container or on top of the container. From the language perspective, the questions are scripted and the language is relatively simple, a downside whose solution is considered as future work by the authors.

With the aim of solving the ALFRED benchmark (Shridhar et al., 2020), several models have been proposed. The proposed solutions are very specific to the task at hand. In the work presented by (Pashevich et al., 2021), they propose a Transformer-based architecture

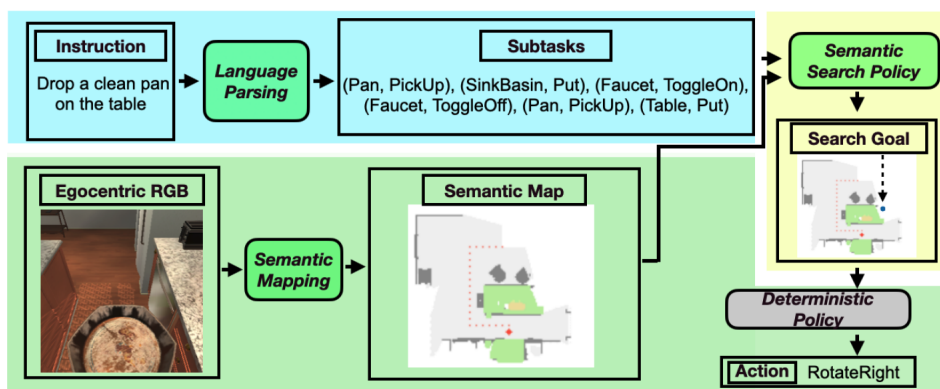


Figure 22: Overview of the FILM architecture for language-guided task completion in ALFRED. They developed a modular architecture composed of several pretrained components that are connected together to solve specific tasks. The goal instruction is the only linguistic information that is available. From it, an execution plan is derived where a semantic search policy looks for the relevant objects in the scene, while a deterministic policy actually executes the relevant actions. Figure adapted from (Min et al., 2021).

that encodes the CNN representation associated with each timestep, together with the language embeddings as well as the previous actions (see Figure 21). In this way, they can use the hidden state representations to generate actions conditioned on the context. Specifically, they use the hidden state of the last trajectory timestep, to predict the action to be executed as well as a possible object class that the agent should manipulate. The limitation of such an approach is that their architecture is specifically designed and trained only for ALFRED. For instance, for manipulation actions, they assume that only a single instance of a given object class is available in the manipulation area (e.g., the counter will contain only one “apple” so no ambiguity is present). Therefore, it is possible to use an external object detector to detect the objects and then just select the one that we want to manipulate based on the class predicted by the model. This forces the model to reason only in terms of which actions can manipulate certain classes but it doesn’t allow the model to have a good understanding of the actual objects that the agent has in its front view because its visual representation is not fine-grained enough. A more generic approach was proposed by (Suglia, Gao, Thomason, Thattai, & Sukhatme, 2021) who designed a multi-modal Transformer able to model the fine-grained, and multi-modal ALFRED trajectories represented in terms of detected objects in the visual scene. As highlighted by (Storks, Gao, Thattai, & Tur, 2021), navigating towards the target object is the hardest task in ALFRED because it can easily put the agent in a situation that the agent cannot recover from. Additionally, the navigation problem is exacerbated by the training regime used to train such models. Their policy is trained using supervised learning, mimicking human trajectories. However, at training time, agents do not have the ability to explore the environment to gain a better understanding of the objects and their location in the environment which might be important to successfully solve the tasks.

Another possible solution to this problem is to avoid solving the problem in an end-to-end fashion, and instead divide it into multiple modules. For instance, as shown in Figure 22,

modular architectures such as FILM (Min et al., 2021) can achieve better performance for this task. First, a dedicated module derives a semantic representation from the goal instruction. This will represent a high-level plan that the agent will have to execute. Specific navigation modules, paired with a map of the environment, are then used to navigate towards target objects. Despite its effectiveness compared to concurrent approaches ², its utility is debatable for models that are able to learn grounded representations that are useful for multiple tasks as well as domains. Indeed, such a model relies on specific algorithms that are used at inference time to complete the task. At the same time, it is important to understand that future models should have a way to derive a map of the environment in which they are located if they want to solve increasingly complex tasks with complete autonomy.

5. Discussion

In this section, we provide a summary of our study of the literature on visually grounded language learning. As a result of our analysis, we derive several research questions that will be detailed in the following subsections. To answer these questions, novel evaluation benchmarks that are more suitable for grounded language learning as well as novel modelling solutions will be required. This survey represents an important reference for studying how to extend both models and related benchmarks to push the boundaries of grounded language learning. Finally, we compare our analysis with previous surveys in the field of V+L, and describe the importance of our categorisation to further study the field of V+L in the future.

5.1 How does this survey relate to other V+L surveys?

The main focus of this survey is to provide a systematic categorisation of several tasks that have been explored in the V+L community. In particular, we provide a categorisation based on 3 broad categories namely *discriminative games*, *generative games*, and *interactive games*. This aims to provide a higher-level perspective of the different tasks that we believe will help researchers in sharing knowledge and in generalising results that have been produced by a specific task but that could be applied to another one as well. For instance, (Ferraro, Mostafazadeh, Huang, Vanderwende, Devlin, Galley, & Mitchell, 2015) presents a survey of image captioning datasets that dates back only to 2015 which can be categorised in the *generative games* category of our survey. The work from (Li, Zhang, Zhang, Liu, Guo, Ni, Zhang, & Zhang, 2022) presents a very detailed analysis of recent Transformer-based architectures that focuses on *discriminative* and *generative games* only ignoring the importance of interactive and embodied games for grounded language learning. A recent survey presented by (Gu et al., 2022) underlines the importance of embodiment and interactive learning. They analyse the field of Vision+Language navigation in great detail. However, as specified in our description of *embodied interactive games*, visual navigation ignores an important part of learning which is object interaction. (Mogadala, Kalimuthu, & Klakow, 2021) presents a very broad survey of general vision and language research up to the year 2019, whereas this paper provides a more up-to-date survey and analysis focused on Interactive Grounded Language Learning tasks. We consider Interactive Grounded Language

2. FILM was 2nd in the ALFRED leaderboard as of 5th May 2022.

Learning a new avenue for multimodal machine learning research. In particular, we argue for the importance of bringing together two relevant problems such as “symbol grounding” and “conversational grounding”, which we believe should be studied together to build flexible learning agents that can smoothly adapt to different speakers and domains using Natural Language.

Finally, we discuss the relationship between our survey and surveys which provide a broader view of artificial agents able to learn language. For instance, (Schlangen, 2019b) uses the notion of language games for describing Natural Language Processing tasks. This survey identifies the task of “unrestricted situated language interaction” as “a natural supremum” in NLP tasks which is akin to the *interactive* language games that we consider in this work. In our survey, as first defined in (Bisk, Holtzman, Thomason, Andreas, Bengio, Chai, Lapata, Lazaridou, May, Nisnevich, et al., 2020b), we underline the importance of considering NLP tasks that require that the agent can “experience” the result of its actions in the world as well as communicate and collaborate with other agents uncovering the social role of Natural Language.

5.2 How can Language Games be used for Visually Grounded Representation Learning?

In this survey, we have analysed several tasks as well as associated models that have been proposed to solve them. Each of them has a specific focus of interest in modelling visual scenes as well as language. Depending on the way the output was generated, we divided them into three types: discriminative, generative, and interactive tasks. We argue that interactive tasks are the most important for future work because they are closer to the language games scenario that we have described. Among them, we consider visual guessing games an important task for learning grounded representations of objects. Objects in an environment, and their attributes, are the first step in learning grounded representations. Other interactive language games are embodied ones. Particularly, we consider language-conditioned action execution in a 3D world as the most relevant benchmark for learning other important elements of grounded representations, such as actions and events. In particular, it will be important to consider new modelling techniques that facilitate the development of agents that can negotiate and coordinate plans of actions in the world.

5.3 What Type of Representations can be Learned via Language Games?

When investigating the quality of distributed representations, it is possible to define several desiderata for distributed representations (Bengio, Courville, & Vincent, 2013). This is an important problem, especially for symbol grounding where we want to learn representations that are compositions of elements that are in turn grounded in the real world (Harnad, 1990). As previously stated, if we focus only on increasing the performance on the downstream task which is evaluated based on a held-out test set, we might have an opaque understanding of the real system capabilities. Most datasets and models assume this scenario where the model is never confronted with novel object categories. Its generalisation abilities are assessed only via novel instances of known categories. Therefore, we want to make sure to investigate language games that are complex enough to let the agent learn grounded representations. A representation can be considered grounded when it is a composition of

other grounded elements. Therefore, to assess this ability of the model, we should consider tasks involving zero-shot learning scenarios where novel categories are involved. In this case, the only way for the agent to solve the task is to use what it has learned during training in novel ways, therefore demonstrating its ability to systematically generalise. To mitigate this issue, (Suglia et al., 2020) describes a novel multi-task evaluation framework that extends the traditional evaluation framework for grounded language learning to give value to the quality of the learned representations. Thanks to this evaluation framework, the authors have shown several limitations of current multimodal models used for situated visual guessing games. Therefore, it is very important that future benchmarks address this issue by designing evaluation setups that assess the model generalisation capabilities in a more systematic way (Linzen, 2020).

5.4 Are Learned Multimodal Representations generic enough to be transferred to other tasks?

From the works described in Section 4, a common trend in the V+L literature is to propose a model architecture having inductive biases that can be applied to a specific task only. However, it is important to underline that another property of high-quality distributed representations is their suitability for other tasks different from the one they have been originally trained on. Concretely, this implies that language games should favour the emergence of important properties of the real world that can be captured by the model. If such knowledge is generic enough, the agent should be able to reuse it for solving other tasks as well. For instance, (Suglia, Bisk, Konstas, Vergari, Bastianelli, Vanzo, & Lemon, 2021) argue that visual guessing games can be developed as a generic transfer learning procedure for grounded language learning. This stresses the idea that more general architectures able to solve multiple tasks are required (Brown et al., 2020).

When considering the generalisation power of the developed models, it is also important to study the ability of these multimodal models to be applied to scenarios where not all the modalities are involved. For instance, it is still not clear whether current instances of multimodal architectures are able to bring an improvement in text-only scenarios such as Question Answering, Natural Language Inference, and so on. Several recent works have tried to address this issue (e.g., Vokenization (Tan & Bansal, 2020)) but they are not generic enough to be easily transferred to other tasks because they rely on additional components that are not part of the same trained architecture. It is important that future work investigates *truly* multimodal models that can rely on multimodal modalities to create richer semantic representations that can be reused in scenarios that are different from the one they were originally trained on.

5.5 What Representations are useful for Action-driven and Event-driven Language Games?

Among the grounded language learning tasks in the literature, visual guessing games represent the first step in a curriculum for learning object categories and their attributes. However, objects are only one part of the spectrum required for understanding. As argued in (McClelland, Hill, Rudolph, Baldrige, & Schütze, 2020), the main target of language is the understanding of situations (Barwise & Perry, 1981). Situations are highly relational

constructs composed of entities and relationships between them. Quoting (McClelland et al., 2020): “To understand a situation is to construct a representation of it that captures aspects of the participating objects, their properties, relationships and interactions, and resulting outcomes”. This resonates with the idea that even the understanding of nouns (i.e., entities) requires composing affordances of them (Glenberg & Kaschak, 2002). Following the idea of grounded cognition (Barsalou, 2008), embodied experiences of concepts and events are required for the development of meanings. Thanks to the development of 3D environments, it is now possible to study embodied tasks and their relationship with grounded cognition.

We have analysed several interactive tasks that study the problem of completing a task guided by Natural Language. First, language-guided navigation was proposed as a benchmark for embodied AI. Equipping the agent only with navigation actions is limiting its ability to learn grounded meanings. Therefore, several other benchmarks have been proposed to study more complex task-completion tasks involving manipulating objects as well. These environments represent an important avenue for embodied AI. Particularly, a general embodied AI architecture should be able to encode arbitrary long action sequences and make decisions upon them. However, simulated environments are limited by the capabilities of the physics engine, the object repertoire and the number of scenes available. Another interesting perspective for a learning agent is represented by videos in the third-person view (i.e., exocentric view). Learning from such multimodal inputs can expose the learning agents to many more situations that are useful in real-world contexts.

5.6 How does Symbol Grounding relate to Conversational Grounding?

As we highlighted in the introduction, interaction with the world and with other agents are key ingredients for language learning. When playing language games, speakers coordinate among themselves to refer to things in the world and achieve goals. Additionally, the ability to engage in conversation allows agents to be more robust in situations of uncertainty. Many recent works have indicated the importance of *conversational grounding* (Brennan & Clark, 1996) as foundational for the future of conversational systems. More importantly, many recent position papers have convincingly argued the importance of connecting symbols and conversational grounding to create truly collaborative systems (Liu, Fang, She, & Chai, 2013; Chai, Fang, Liu, & She, 2016; Schlangen, 2019a; Benotti & Blackburn, 2021; Chandu et al., 2021). It is therefore important to define new tasks where symbol grounding and conversational grounding together have the role of first-class citizens, and where communicating agents have divergent information which they need to coordinate (Lemon, 2022). The need for divergence is essential because it is what motivates agents to communicate to achieve specific goals. Additionally, this idea of divergence is related to the concept of Theory of Mind (Astington & Baird, 2005) which is used in this context to refer to social understanding. Following (Tomasello, 2010), social understanding interacts with language communication in two ways depending on the type of social understanding desired. In particular, understanding the goals and perceptions of an agent will engage different language activities than the ones demonstrated when trying to understand one’s desires and beliefs.

Studying the interplay between symbol and conversational grounding in a systematic way will require dedicated data collection processes aimed at collecting data in realistic

embodied environments from which it is possible to elicit both symbol and communicative grounding phenomena that a model can learn from. These phenomena require additional methods to model both the conversational context as well as to derive relevant signals from the language feedback provided by other agents. It is important to consider a learning scenario that is closer to a real-world scenario involving tasks providing an ecologically valid scenario (De Vries et al., 2020) for studying multimodal conversational grounding.

6. Future Research Directions

In this section, we will describe potential interesting research directions that we argue are worth exploring as a result of this literature review. These research directions are important for the development of truly generalist agents that can *perceive*, *reason*, and *act* in novel scenarios.

6.1 Perceptual Symbols

Language is a means of communication, which has evolved by transmitting symbols grounded in perceptual experience (e.g., gestures, sounds, etc.). However, as described in Section 4, current AI models learn a language only from symbolic representations derived from a tokenizer that are: 1) specific for a given language and are costly to transfer to low-resource languages, so limiting people’s to access technology; 2) sensitive to noise (e.g., spelling mistakes); 3) hand-crafted because they do not represent language input in a multimodal way, as humans do.

Thanks to recent innovations in Computer Vision, researchers have attempted to render language data as images, and use new Deep Learning architectures for Computer Vision to attempt this task. This solution was first proposed for Machine Translation by (Salesky, Etter, & Post, 2021) who use a CNN architecture to derive visual representations of the input text that was rendered as an image. From these representations, the authors were able to translate to a target language without the need for an input tokenizer. (Rust, Lotz, Bugliarello, Salesky, de Lhoneux, & Elliott, 2022) generalised this approach by proposing *PIXEL*, the first language model pre-trained with rendered images only. This model is pretrained using a self-supervised reconstruction loss that uses the image patches to learn latent representations of the language. In their paper, the authors demonstrate comparable performance with BERT on several tasks such as POS tagging and Natural Language inference. Most importantly, they show superior performance for very low-resource languages as well as languages that do not have a Latin script.

Deriving language representations from the audio signal is another recent trend that goes under the name of *textless NLP* (Kharitonov, Copet, Lakhota, Nguyen, Tomasello, Lee, Elkahky, Hsu, Mohamed, Dupoux, et al., 2022). For instance, the work presented in (Harwath & Glass, 2017) demonstrates how to derive word-like representation units directly from visual and audio features. Additionally, there is a very interesting line of work interested in learning representations from vision, language, and audio data (e.g., (Zellers, Lu, Lu, Yu, Zhao, Salehi, Kusupati, Hessel, Farhadi, & Choi, 2022)). However, this approach is still constrained by the use of a tokenizer which limits applicability to specific language scripts.

Deriving language representations from multimodal inputs is essential to achieve more realistic NLP systems that are more widely applicable such as Spoken Language Understanding systems. Second, spoken language carries many nuances (prosody, irony, anger, etc.) and expressive vocalisations (laughter, yawning, etc) that are not captured by text. Additionally, modelling language via the visual modality can allow the implementation of more sophisticated agents that are able to understand the content of websites, forms, and other media that contain multimodal content. In general, modelling language directly from other perceptual modalities has the potential to make AI applications more natural and expressive. To advance the state of the art in Multimodal Representation Learning, we should aim to design *perceptual symbol systems* whose aim is to derive language representations from multimodal perceptual experiences in a process akin to the seminal work from (Barsalou, 1999).

6.2 Multimodal Robustness

To date, many large-scale V+L pretrained models have been proposed, and they are typically only used for V+L tasks. However, there are very few works (e.g., (Tan & Bansal, 2020; Singh et al., 2022)) that demonstrate the ability of such models to bring a performance gain when it comes to single-modality tasks (e.g., vision-only or language-only). Therefore, we believe that the new generation of V+L models should be able to effectively encode the input modalities as well as be robust to *perturbations* of them (i.e., cases where such are potentially missing or are noisy). For instance, the work by (Ma, Ren, Zhao, Testuggine, & Peng, 2022) shows how multimodal models receiving noisy or missing modalities achieve performance that is inferior or just on par with unimodal models. Therefore, another important ingredient for the implementation of a sound and effective *perceptual symbol systems* is the ability to encode and *reconstruct* modalities when these are missing or noisy.

Crucially, it is important to investigate the robustness of multimodal representations in general. Much recent literature has tackled this problem from the point of view of domain generalisation (e.g., (Akula et al., 2021)) as well as the effect of spurious correlations and noise (e.g., (Tu, Lalwani, Gella, & He, 2020)). This calls for a more systematic empirical investigation of the effect of such shifts in distribution and how they affect models across multiple tasks. For this reason, it is important to develop experimental frameworks that diagnose what matters in learning representations for V+L tasks (e.g., (Zhu, Qi, Narayana, Sone, Basu, Wang, Wu, Eckstein, & Wang, 2022; Akula, Gella, Al-Onaizan, Zhu, & Reddy, 2020)).

6.3 Interaction with the World and with other Agents

As we described in Section 5.6, future V+L agents need to develop the ability to interact with the world around them. This is important because the ultimate goal is to create AI systems that can communicate with humans in Natural Language and develop a seamless symbiotic relationship with them. The current wave of Generative AI models has demonstrated that having conversational skills is essential for user-facing systems such as ChatGPT (OpenAI, 2022).

In this scenario, we will assume that the agent is embodied in an environment (either simulated or realistic) and that there are one or multiple agents in it. The agent, motivated

by a goal g , will have to produce a sequence of actions that will achieve their goal. There could be two types of actions that the agent can perform: *physical actions* and *verbal actions*. Physical actions involve a change in the agent’s position or a change in the environment around them. In both cases, the agent will be able to *observe* the results of their actions and derive signals from these changes that can be used for learning purposes. Verbal actions instead are intended as utterances that are formulated when engaging in conversations with other agents. This is required to acquire information and resolve possible ambiguities arising in the process of conversational grounding. Additionally, we also envision other types of verbal actions that involve *inner monologue* (Vygotsky, 2012). This form of *self-referential dialogue* could be very effective if intended as a form of planning that allows the agent to foresee the outcome of certain actions, and recover from possible problems. This approach has already been implemented in some form in text-only language models via *chain of thought* (Wei, Wang, Schuurmans, Bosma, Chi, Le, & Zhou, 2022; Yao, Zhao, Yu, Du, Shafran, Narasimhan, & Cao, 2022) as well as using text-only models for planning in 3D simulated environments (e.g., (Huang, Xia, Xiao, Chan, Liang, Florence, Zeng, Tompson, Mordatch, Chebotar, et al., 2022)).

7. Conclusions

In this work, we surveyed datasets and tasks proposed to study the problem of visual symbol grounding in multimodal machine learning models. Following the language games proposed in the Talking Heads experiments (Steels, 2015), we proposed a categorisation of these tasks into 3 main categories: discriminative, generative, and interactive language games. We argue that *interactive games* represent the most relevant class of language games to study the problem of grounded language learning. Additionally, we studied the properties and the skills required for solving the proposed language games. We find that language games situated in 3D environments enabling navigation and object manipulation represent an important step forward. In this way, artificial agents can learn object reference in the world (i.e., required when navigating in an environment), as well as the effects of certain actions on certain objects (i.e., observed when manipulating objects) which are essential for the development of rich grounded meaning representations.

After categorising all the available tasks, we analysed the multimodal machine learning models that have been proposed to solve them. We completed an analysis based on several key characteristics that these models have. We find that in order to support interactive language games, we need more sophisticated ways of encoding rich multimodal context, as well as supporting conversations aimed at collecting important information for the task at hand. In addition, we highlighted the need for multimodal models able to generate representations that are meaningful even in scenarios where some modalities are missing (i.e., text-only models), or when some of the modalities are noisy. This is an essential ingredient for developing systems that are resilient to adverse scenarios that are very common in the real world.

As a result of this survey, we have identified several weaknesses primarily concerned with the datasets used to train such models. The development of sophisticated 3D engines has sparked research into embodied multimodal models able to solve tasks in simulated environments. This represents an important direction for future work with the aim of de-

veloping more ecologically valid studies (De Vries et al., 2020) of how grounded meanings can be developed. Additionally, to adequately study language learning, we stressed the idea that we require datasets that require multi-agent communication skills. These are essential to account for recovery from possible breakdown and uncertainty during task solving. Moreover, it will facilitate agents to learn new tasks by learning from interaction (Gluck et al., 2018).

Acknowledgments

We would like to thank Arash Eshghi and Raquel Fernandez for their feedback on a preliminary version of this manuscript which was part of the first author’s PhD thesis (Suglia et al., 2022).

References

- Abramson, J., Ahuja, A., Barr, I., Brussee, A., Carnevale, F., Cassin, M., Chhaparia, R., Clark, S., Damoc, B., Dudzik, A., et al. (2020). Imitating interactive intelligence..
- Agarwal, A., Gurumurthy, S., Sharma, V., Lewis, M., & Sycara, K. (2019). Community Regularization of Visually-Grounded Dialog. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1042–1050. International Foundation for Autonomous Agents and Multiagent Systems.
- Agarwal, S., Bui, T., Lee, J.-Y., Konstas, I., & Rieser, V. (2020). History for Visual Dialog: Do we really need it?. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8182–8197.
- Agrawal, A., Kembhavi, A., Batra, D., & Parikh, D. (2017). C-vqa: A compositional split of the visual question answering (vqa) v1. 0 dataset..
- Akula, A., Changpinyo, S., Gong, B., Sharma, P., Zhu, S.-C., & Soricut, R. (2021). Crossvqa: Scalably generating benchmarks for systematically testing vqa generalization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 2148–2166.
- Akula, A., Gella, S., Al-Onaizan, Y., Zhu, S.-C., & Reddy, S. (2020). Words Aren’t Enough, Their Order Matters: On the Robustness of Grounding Visual Referring Expressions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6555–6565 Online. Association for Computational Linguistics.
- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., & Zhang, L. (2018a). Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6077–6086.
- Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., & Van Den Hengel, A. (2018b). Vision-and-language navigation: Interpreting

- visually-grounded navigation instructions in real environments. In *Computer Vision and Pattern Recognition (CVPR)*.
- Andreas, J., Rohrbach, M., Darrell, T., & Klein, D. (2016). Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 39–48.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., & Parikh, D. (2015). Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433.
- Antunes, A., Saponaro, G., Morse, A., Jamone, L., Santos-Victor, J., & Cangelosi, A. (2017). Learn, plan, remember: A developmental robot architecture for task solving. In *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pp. 283–289. IEEE.
- Astington, J. W., & Baird, J. A. (2005). *Why language matters for theory of mind*. Oxford University Press.
- Atzmon, Y., Berant, J., Kezami, V., Globerson, A., & Chechik, G. (2016). Learning to generalize to new compositions in image understanding..
- Austin, J. L. (1975). *How to do things with words*. Oxford university press.
- Bahdanau, D., Hill, F., Leike, J., Hughes, E., Hosseini, A., Kohli, P., & Grefenstette, E. (2018a). Learning to understand goal specifications by modelling reward..
- Bahdanau, D., Murty, S., Noukhovitch, M., Nguyen, T. H., de Vries, H., & Courville, A. (2018b). Systematic generalization: what is required and can it be learned?..
- Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2018). Multimodal machine learning: A survey and taxonomy.. Vol. 41, pp. 423–443. IEEE.
- Barsalou, L. W. (1999). Perceptual symbol systems.. Vol. 22, pp. 577–660. Cambridge University Press.
- Barsalou, L. W. (2008). Grounded cognition.. Vol. 59, pp. 617–645. Annual Reviews.
- Barwise, J., & Perry, J. (1981). Situations and attitudes.. Vol. 78, pp. 668–691. JSTOR.
- Batali, J. (1998). Computational simulations of the emergence of grammar.. pp. 405–426. Cambridge University Press.
- Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. (2016). Deepmind lab..
- Becerra-Bonache, L., Christiansen, H., & Jiménez-López, M. D. (2018). A Gold Standard to Measure Relative Linguistic Complexity with a Grounded Language Learning Model. In *Proceedings of the Workshop on Linguistic Complexity and Natural Language Processing*, pp. 1–9.

- Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5185–5198.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives.. Vol. 35, pp. 1798–1828. IEEE.
- Benotti, L., & Blackburn, P. (2021). Grounding as a Collaborative Process. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 515–531 Online. Association for Computational Linguistics.
- Bisk, Y., Holtzman, A., Thomason, J., Andreas, J., Bengio, Y., Chai, J., Lapata, M., Lazaridou, A., May, J., Nisnevich, A., et al. (2020a). Experience Grounds Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8718–8735.
- Bisk, Y., Holtzman, A., Thomason, J., Andreas, J., Bengio, Y., Chai, J., Lapata, M., Lazaridou, A., May, J., Nisnevich, A., et al. (2020b). Experience Grounds Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8718–8735.
- Bisk, Y., Yuret, D., & Marcu, D. (2016). Natural language communication with robots. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 751–761.
- Bogin, B., Gupta, S., Gardner, M., & Berant, J. (2021a). COVR: A Test-Bed for Visually Grounded Compositional Generalization with Real Images. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 9824–9846.
- Bogin, B., Subramanian, S., Gardner, M., & Berant, J. (2021b). Latent Compositional Representations Improve Systematic Generalization in Grounded Question Answering.. Vol. 9, pp. 195–210. MIT Press.
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching Word Vectors with Subword Information.. Vol. 5, pp. 135–146.
- Bouchacourt, D., & Baroni, M. (2018). How agents see things: On visual representations in an emergent language game..
- Bowman, S., Angeli, G., Potts, C., & Manning, C. D. (2015). A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 632–642.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation.. Vol. 22, pp. 482–493.
- Brodeur, S., Perez, E., Anand, A., Golemo, F., Celotti, L., Strub, F., Rouat, J., Larochelle, H., & Courville, A. (2017). HoME: A household multimodal environment..

- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners.. Vol. 33, pp. 1877–1901.
- Bugliarello, E., Cotterell, R., Okazaki, N., & Elliott, D. (2021). Multimodal pretraining unmasked: A meta-analysis and a unified framework of vision-and-language BERTs.. Vol. 9, pp. 978–994. MIT Press.
- Bugliarello, E., & Elliott, D. (2021). The Role of Syntactic Planning in Compositional Image Captioning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 593–607.
- Bugliarello, E., Liu, F., Pfeiffer, J., Reddy, S., Elliott, D., Ponti, E. M., & Vulić, I. (2022). IGLUE: A Benchmark for Transfer Learning across Modalities, Tasks, and Languages..
- Chai, J. Y., Fang, R., Liu, C., & She, L. (2016). Collaborative language grounding toward situated human-robot dialogue.. Vol. 37, pp. 32–45.
- Chandu, K. R., Bisk, Y., & Black, A. W. (2021). Grounding ‘Grounding’ in NLP. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4283–4305 Online. Association for Computational Linguistics.
- Chang, A., Dai, A., Funkhouser, T., Halber, M., Niebner, M., Savva, M., Song, S., Zeng, A., & Zhang, Y. (2017). Matterport3D: Learning from RGB-D Data in Indoor Environments. In *2017 International Conference on 3D Vision (3DV)*, pp. 667–676. IEEE.
- Chen, H., Suhr, A., Misra, D., Snavely, N., & Artzi, Y. (2019). Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12538–12547.
- Chen, Y.-C., Li, L., Yu, L., Kholy, A. E., Ahmed, F., Gan, Z., Cheng, Y., & Liu, J. (2020). Uniter: Universal image-text representation learning. In *European Conference on Computer Vision (ECCV)*.
- Chevalier-Boisvert, M., Bahdanau, D., Lahlou, S., Willems, L., Saharia, C., Nguyen, T. H., & Bengio, Y. (2018). BabyAI: First Steps Towards Grounded Language Learning With a Human In the Loop..
- Choi, E., Lazaridou, A., & de Freitas, N. (2018). Compositional oververter communication learning from raw visual input..
- Chou, S.-H., Chen, Y.-C., Zeng, K.-H., Hu, H.-N., Fu, J., & Sun, M. (2018). Self-view grounding given a narrated 360 video. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Chrupała, G. (2019). Symbolic Inductive Bias for Visually Grounded Learning of Spoken Language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6452–6462.

- Clark, C., Salvador, J., Schwenk, D., Bonafilia, D., Yatskar, M., Kolve, E., Herrasti, A., Choi, J., Mehta, S., Skjonsberg, S., et al. (2021). Iconary: A Pictionary-Based Game for Testing Multimodal Communication with Drawings and Text. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1864–1886.
- Clark, H. H. (1996). *Using language*. Cambridge university press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication.. American Psychological Association.
- Co-Reyes, J. D., Gupta, A., Sanjeev, S., Altieri, N., DeNero, J., Abbeel, P., & Levine, S. (2018). Guiding Policies with Language via Meta-Learning..
- Csurka, G., Dance, C., Fan, L., Willamowski, J., & Bray, C. (2004). Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, Vol. 1, pp. 1–2. Prague.
- Da, J., Forbes, M., Zellers, R., Zheng, A., Hwang, J. D., Bosselut, A., & Choi, Y. (2021). Edited Media Understanding Frames: Reasoning About the Intent and Implications of Visual Misinformation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 2026–2039.
- Das, A., Datta, S., Gkioxari, G., Lee, S., Parikh, D., & Batra, D. (2018). Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 2054–2063.
- Das, A., Kottur, S., Gupta, K., Singh, A., Yadav, D., Moura, J. M., Parikh, D., & Batra, D. (2017a). Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 326–335.
- Das, A., Kottur, S., Gupta, K., Singh, A., Yadav, D., Moura, J. M., Parikh, D., & Batra, D. (2017b). Visual Dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- De Vries, H., Bahdanau, D., & Manning, C. (2020). Towards ecologically valid research on language user interfaces..
- De Vries, H., Shuster, K., Batra, D., Parikh, D., Weston, J., & Kiela, D. (2018). Talk the walk: Navigating new york city through grounded dialogue..
- De Vries, H., Strub, F., Chandar, S., Pietquin, O., Larochelle, H., & Courville, A. (2017). Guesswhat?! visual object discovery through multi-modal dialogue. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5503–5512.
- Deitke, M., Han, W., Herrasti, A., Kembhavi, A., Kolve, E., Mottaghi, R., Salvador, J., Schwenk, D., VanderBilt, E., Wallingford, M., et al. (2020). Robothor: An open simulation-to-real embodied ai platform. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3164–3174.

- Deng, C., Wu, Q., Wu, Q., Hu, F., Lyu, F., & Tan, M. (2018). Visual grounding via accumulated attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7746–7755.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding..
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019a). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019b). 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)*, pp. 4171–4186.
- Dobnik, S., Kelleher, J. D., & Howes, C. (2020). Local alignment of frame of reference assignment in English and Swedish dialogue. In *German Conference on Spatial Cognition*, pp. 251–267. Springer.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale..
- Elliott, D., Frank, S., Sima'an, K., & Specia, L. (2016). Multi30K: Multilingual English-German Image Descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pp. 70–74.
- Elman, J. L. (1990). Finding structure in time.. Vol. 14, pp. 179–211. Wiley Online Library.
- Ferraro, F., Mostafazadeh, N., Huang, T.-H., Vanderwende, L., Devlin, J., Galley, M., & Mitchell, M. (2015). A Survey of Current Datasets for Vision and Language Research. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 207–213.
- Fodor, J. A. (1983). *The modularity of mind*. MIT press.
- Fodor, J. A., & Lepore, E. (2002). *The compositionality papers*. Oxford University Press.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis.. Vol. 28, pp. 3–71. Elsevier.
- Frege, G. (1892). Über sinn und bedeutung.. Vol. 100, pp. 25–50.

- Gao, Q., Thattai, G., Gao, X., Shakiah, S., Pansare, S., Sharma, V., Sukhatme, G., Shi, H., Yang, B., Zheng, D., et al. (2023). Alexa arena: A user-centric interactive platform for embodied ai..
- Gao, X., Gao, Q., Gong, R., Lin, K., Thattai, G., & Sukhatme, G. S. (2022). Dialfred: Dialogue-enabled agents for embodied instruction following.. Vol. 7, pp. 10049–10056. IEEE.
- Gardenfors, P. (2014). *The geometry of meaning: Semantics based on conceptual spaces*. MIT press.
- Gibson, J. J. (1977). The theory of affordances. *Hilldale, USA*, 1(2), 67–82.
- Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action.. Vol. 9, pp. 558–565. Springer.
- Gluck, K. A., Laird, J., & Lupp, J. (2018). *Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks Through Natural Interactions*. MIT Press.
- Gordon, D., Kembhavi, A., Rastegari, M., Redmon, J., Fox, D., & Farhadi, A. (2018). Iqa: Visual question answering in interactive environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4089–4098.
- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., & Parikh, D. (2017a). Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6904–6913.
- Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., & Parikh, D. (2017b). Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6904–6913.
- Grauman, K., & Darrell, T. (2005). The pyramid match kernel: Discriminative classification with sets of image features. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, Vol. 2, pp. 1458–1465. IEEE.
- Graves, A., Wayne, G., & Danihelka, I. (2014). Neural turing machines..
- Grubinger, M., Clough, P., Müller, H., & Deselaers, T. (2006). The iapr tc-12 benchmark: A new evaluation resource for visual information systems. In *International workshop ontoImage*, Vol. 2.
- Gu, J., Stefani, E., Wu, Q., Thomason, J., & Wang, X. (2022). Vision-and-Language Navigation: A Survey of Tasks, Methods, and Future Directions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7606–7623.

- Gurari, D., Li, Q., Stangl, A. J., Guo, A., Lin, C., Grauman, K., Luo, J., & Bigham, J. P. (2018). Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3608–3617.
- Haber, J., Baumgärtner, T., Takmaz, E., Gelderloos, L., Bruni, E., & Fernández, R. (2019). The PhotoBook Dataset: Building Common Ground through Visually-Grounded Dialogue. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1895–1910.
- Harnad, S. (1990). The symbol grounding problem.. Vol. 42, pp. 335–346. Elsevier.
- Harwath, D., & Glass, J. (2017). Learning Word-Like Units from Joint Audio-Visual Analysis. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 506–517.
- Havrylov, S., & Titov, I. (2017). Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Advances in neural information processing systems*, pp. 2149–2159.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep residual learning for image recognition. eprint..
- Hermann, K. M., Hill, F., Green, S., Wang, F., Faulkner, R., Soyer, H., Szepesvari, D., Czarnecki, W. M., Jaderberg, M., Teplyashin, D., et al. (2017). Grounded language learning in a simulated 3d world..
- Hong, Y., Wu, Q., Qi, Y., Rodriguez-Opazo, C., & Gould, S. (2020). A Recurrent Vision-and-Language BERT for Navigation..
- Huang, T.-H., Ferraro, F., Mostafazadeh, N., Misra, I., Agrawal, A., Devlin, J., Girshick, R., He, X., Kohli, P., Batra, D., et al. (2016). Visual storytelling. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1233–1239.
- Huang, W., Xia, F., Xiao, T., Chan, H., Liang, J., Florence, P., Zeng, A., Tompson, J., Mordatch, I., Chebotar, Y., et al. (2022). Inner monologue: Embodied reasoning through planning with language models..
- Hudson, D. A., & Manning, C. D. (2018). Compositional Attention Networks for Machine Reasoning. In *International Conference on Learning Representations*.
- Hudson, D. A., & Manning, C. D. (2019). Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6700–6709.
- Hupkes, D., Dankers, V., Mul, M., & Bruni, E. (2020). Compositionality decomposed: how do neural networks generalise?.. Vol. 67, pp. 757–795.

- Ilinykh, N., Zarrieß, S., & Schlangen, D. (2019a). Meetup! a corpus of joint activity dialogues in a visual environment..
- Ilinykh, N., Zarrieß, S., & Schlangen, D. (2019b). Tell me more: A dataset of visual scene description sequences. In *Proceedings of the 12th international conference on natural language generation*, pp. 152–157.
- Jain, U., Weihs, L., Kolve, E., Rastegari, M., Lazebnik, S., Farhadi, A., Schwing, A. G., & Kembhavi, A. (2019). Two body problem: Collaborative visual task completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6689–6699.
- Jiang, H., Misra, I., Rohrbach, M., Learned-Miller, E., & Chen, X. (2020). In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10267–10276.
- Jiang, Y., Gu, S. S., Murphy, K. P., & Finn, C. (2019). Language as an Abstraction for Hierarchical Deep Reinforcement Learning.. Vol. 32, pp. 9419–9431.
- Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., & Girshick, R. (2017). Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2901–2910.
- Johnson, J., Karpathy, A., & Fei-Fei, L. (2016). Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4565–4574.
- Kafle, K., & Kanan, C. (2017). An analysis of visual question answering algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1965–1973.
- Kalyan, A., Lee, S., Kannan, A., & Batra, D. (2018). Learn from Your Neighbor: Learning Multi-modal Mappings from Sparse Annotations. In *International Conference on Machine Learning*, pp. 2454–2463.
- Kasami, T. (1966). An efficient recognition and syntax-analysis algorithm for context-free languages.. Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.
- Kazemzadeh, S., Ordonez, V., Matten, M., & Berg, T. (2014). Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 787–798.
- Kharitonov, E., Copet, J., Lakhotia, K., Nguyen, T. A., Tomasello, P., Lee, A., Elkahky, A., Hsu, W.-N., Mohamed, A., Dupoux, E., et al. (2022). textless-lib: A library for textless spoken language processing..
- Kiela, D. (2017). Deep embodiment: grounding semantics in perceptual modalities. Tech. rep. UCAM-CL-TR-899, University of Cambridge, Computer Laboratory.

- Kiela, D., Conneau, A., Jabri, A., & Nickel, M. (2018). Learning Visually Grounded Sentence 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)*, pp. 408–418.
- Kiela, D., Firooz, H., Mohan, A., Goswami, V., Singh, A., Ringshia, P., & Testuggine, D. (2020). The hateful memes challenge: Detecting hate speech in multimodal memes..
- Kim, J.-H., Kitaev, N., Chen, X., Rohrbach, M., Zhang, B.-T., Tian, Y., Batra, D., & Parikh, D. (2019). CoDraw: Collaborative Drawing as a Testbed for Grounded Goal-driven Communication. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6495–6513.
- Kim, W., Son, B., & Kim, I. (2021). Vilt: Vision-and-language transformer without convolution or region supervision..
- Kiros, J., Chan, W., & Hinton, G. (2018). Illustrative Language Understanding: Large-Scale Visual Grounding with Image Search. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 922–933.
- Kiseleva, J., Li, Z., Aliannejadi, M., Mohanty, S., ter Hoeve, M., Burtsev, M., Skrynnik, A., Zholus, A., Panov, A., Srinet, K., et al. (2022). Interactive Grounded Language Understanding in a Collaborative Environment: IGLU 2021..
- Kolve, E., Mottaghi, R., Han, W., VanderBilt, E., Weihs, L., Herrasti, A., Gordon, D., Zhu, Y., Gupta, A., & Farhadi, A. (2017). AI2-THOR: An Interactive 3D Environment for Visual AI..
- Kottur, S., Moon, S., Geramifard, A., & Damavandi, B. (2021). SIMMC 2.0: A Task-oriented Dialog Dataset for Immersive Multimodal Conversations..
- Kottur, S., Moura, J., Lee, S., & Batra, D. (2017). Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2962–2967.
- Kottur, S., Moura, J. M., Parikh, D., Batra, D., & Rohrbach, M. (2018). Visual coreference resolution in visual dialog using neural module networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 153–169.
- Kottur, S., Moura, J. M., Parikh, D., Batra, D., & Rohrbach, M. (2019). CLEVR-Dialog: A Diagnostic Dataset for Multi-Round Reasoning in Visual Dialog..
- Krishna, R., Hata, K., Ren, F., Fei-Fei, L., & Carlos Nibbles, J. (2017a). Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pp. 706–715.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. (2017b). Visual genome: Connecting language and vision using crowdsourced dense image annotations.. Vol. 123, pp. 32–73. Springer.

- Krojer, B., Adlakha, V., Vineet, V., Goyal, Y., Ponti, E., & Reddy, S. (2022). Image Retrieval from Contextual Descriptions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3426–3440.
- Ku, A., Anderson, P., Patel, R., Ie, E., & 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)*, pp. 4392–4412.
- Kuo, Y.-L., Katz, B., & Barbu, A. (2020). Compositional Networks Enable Systematic Generalization for Grounded Language Understanding..
- Küttler, H., Nardelli, N., Miller, A. H., Raileanu, R., Selvatici, M., Grefenstette, E., & Rocktäschel, T. (2020). The nethack learning environment..
- Lake, B. M. (2019). Compositional generalization through meta sequence-to-sequence learning..
- Lake, B. M., & Murphy, G. L. (2021). Word meaning in minds and machines... American Psychological Association.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people.. Vol. 40. Cambridge University Press.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge... Vol. 104, p. 211. American Psychological Association.
- Lazaridou, A., Chrupała, G., Fernández, R., & Baroni, M. (2016). Multimodal Semantic Learning from Child-Directed Input. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 387–392.
- Lazaridou, A., Hermann, K. M., Tuyls, K., & Clark, S. (2018). Emergence of linguistic communication from referential games with symbolic and pixel input..
- Lazaridou, A., Peysakhovich, A., & Baroni, M. (2016). Multi-agent cooperation and the emergence of (natural) language..
- Lazaridou, A., Potapenko, A., & Tieleman, O. (2020). Multi-agent communication meets natural language: Synergies between functional and structural language learning..
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning.. Vol. 521, pp. 436–444. Nature Publishing Group.
- Lei, J., Yu, L., Berg, T. L., & Bansal, M. (2020). Tvr: A large-scale dataset for video-subtitle moment retrieval. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pp. 447–463. Springer.

- Lemon, O. (2022). Conversational grounding in emergent communication – data and divergence. In *Emergent Communication (EmeComm) Workshop at ICLR*.
- Levesque, H., Davis, E., & Morgenstern, L. (2012). The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.
- Lewis, D. K. (1969). *Convention: A Philosophical Study*. Wiley-Blackwell, Cambridge, MA, USA.
- Li, F., Zhang, H., Zhang, Y.-F., Liu, S., Guo, J., Ni, L. M., Zhang, P., & Zhang, L. (2022). Vision-Language Intelligence: Tasks, Representation Learning, and Large Models..
- Li, J., Ataman, D., & Sennrich, R. (2021). Vision Matters When It Should: Sanity Checking Multimodal Machine Translation Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 8556–8562.
- Li, L. H., Yatskar, M., Yin, D., Hsieh, C.-J., & Chang, K.-W. (2019). VisualBERT: A simple and performant baseline for vision and language..
- Li, X., Yin, X., Li, C., Hu, X., Zhang, P., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., Choi, Y., & Gao, J. (2020). Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. In *European Conference on Computer Vision (ECCV)*.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer.
- Linzen, T. (2020). How Can We Accelerate Progress Towards Human-like Linguistic Generalization?. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5210–5217.
- Liu, C., Fang, R., She, L., & Chai, J. (2013). Modeling collaborative referring for situated referential grounding. In *Proceedings of the SIGDIAL 2013 Conference*, pp. 78–86.
- Liu, F., Bugliarello, E., Ponti, E. M., Reddy, S., Collier, N., & Elliott, D. (2021). Visually Grounded Reasoning across Languages and Cultures. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10467–10485.
- Lu, J., Batra, D., Parikh, D., & Lee, S. (2019). ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Neural Information Processing Systems (NeurIPS)*.
- Lu, J., Yang, J., Batra, D., & Parikh, D. (2018). Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7219–7228.
- Ma, M., Ren, J., Zhao, L., Testuggine, D., & Peng, X. (2022). Are Multimodal Transformers Robust to Missing Modality?. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18177–18186.

- Magassouba, A., Sugiura, K., & Kawai, H. (2021). CrossMap Transformer: A Crossmodal Masked Path Transformer Using Double Back-Translation for Vision-and-Language Navigation..
- Majumdar, A., Shrivastava, A., Lee, S., Anderson, P., Parikh, D., & Batra, D. (2020a). Improving vision-and-language navigation with image-text pairs from the web. In *European Conference on Computer Vision*, pp. 259–274. Springer.
- Majumdar, A., Shrivastava, A., Lee, S., Anderson, P., Parikh, D., & Batra, D. (2020b). Improving Vision-and-Language Navigation with Image-Text Pairs from the Web. In *European Conference on Computer Vision (ECCV)*.
- Marcus, G. (2018). Deep learning: A critical appraisal..
- Massiceti, D., Dokania, P. K., Siddharth, N., & Torr, P. H. (2018). Visual Dialogue without Vision or Dialogue..
- McClelland, J. L., Hill, F., Rudolph, M., Baldridge, J., & Schütze, H. (2020). Placing language in an integrated understanding system: Next steps toward human-level performance in neural language models.. Vol. 117, pp. 25966–25974. National Acad Sciences.
- Mikolov, T., Yih, W.-t., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pp. 746–751.
- Miller, G. A. (1995). WordNet: a lexical database for English.. Vol. 38, pp. 39–41. ACM New York, NY, USA.
- Min, S. Y., Chaplot, D. S., Ravikumar, P., Bisk, Y., & Salakhutdinov, R. (2021). FILM: Following Instructions in Language with Modular Methods..
- Mirowski, P., Banki-Horvath, A., Anderson, K., Teplyashin, D., Hermann, K. M., Malinowski, M., Grimes, M. K., Simonyan, K., Kavukcuoglu, K., Zisserman, A., et al. (2019). The StreetLearn Environment and Dataset..
- Misra, D., Bennett, A., Blukis, V., Niklasson, E., Shatkhin, M., & Artzi, Y. (2018a). Mapping Instructions to Actions in 3D Environments with Visual Goal Prediction. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Misra, I., Girshick, R., Fergus, R., Hebert, M., Gupta, A., & van der Maaten, L. (2018b). Learning by asking questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 11–20.
- Mogadala, A., Kalimuthu, M., & Klakow, D. (2021). Trends in integration of vision and language research: A survey of tasks, datasets, and methods.. Vol. 71, pp. 1183–1317.
- Montague, R. (1970). Universal grammar.. Vol. 36, pp. 373–398. Wiley Online Library.

- Mordatch, I., & Abbeel, P. (2018). Emergence of grounded compositional language in multi-agent populations. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Narayan-Chen, A., Jayannavar, P., & Hockenmaier, J. (2019). Collaborative dialogue in Minecraft. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5405–5415.
- Navigli, R., & Ponzetto, S. P. (2010). BabelNet: Building a very large multilingual semantic network. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pp. 216–225.
- Nenkova, A., & McKeown, K. (2011). *Automatic summarization*. Now Publishers Inc.
- Newell, A., Simon, H. A., et al. (1972). *Human problem solving*, Vol. 104. Prentice-hall Englewood Cliffs, NJ.
- Nikolaus, M., Abdou, M., Lamm, M., Aralikkatte, R., & Elliott, D. (2019). Compositional Generalization in Image Captioning. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pp. 87–98.
- Ogden, C. K., & Richards, I. A. (1923). *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*, Vol. 29. K. Paul, Trench, Trubner & Company, Limited.
- OpenAI (2022). Introducing ChatGPT..
- Padmakumar, A., Thomason, J., Shrivastava, A., Lange, P., Narayan-Chen, A., Gella, S., Piramuthu, R., Tur, G., & Hakkani-Tur, D. (2021). TEACH: Task-driven Embodied Agents that Chat..
- Pantazopoulos, G., Suglia, A., & Eshghi, A. (2022). Combine to Describe: Evaluating Compositional Generalization in Image Captioning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pp. 115–131.
- Papalampidi, P., Keller, F., Frermann, L., & Lapata, M. (2020). Screenplay Summarization Using Latent Narrative Structure. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1920–1933.
- Park, J. S., Bhagavatula, C., Mottaghi, R., Farhadi, A., & Choi, Y. (2020). Visualcomet: Reasoning about the dynamic context of a still image. In *European Conference on Computer Vision*, pp. 508–524. Springer.
- Pashevich, A., Schmid, C., & Sun, C. (2021). Episodic transformer for vision-and-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15942–15952.
- Patki, S., & Howard, T. (2018). Language-Guided Adaptive Perception for Efficient Grounded Communication with Robotic Manipulators in Cluttered Environments. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pp. 151–160.

- Peirce, C. S. (1902). *Logic as semiotic: The theory of signs.* p. 100.
- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., & Courville, A. (2018). Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Pezzelle, S., & Fernández, R. (2019). Is the red square big? malevic: Modeling adjectives leveraging visual contexts..
- Pezzelle, S., Greco, C., Gandolfi, G., Gualdoni, E., & Bernardi, R. (2020). Be different to be better! A benchmark to leverage the complementarity of language and vision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pp. 2751–2767.
- Pillai, N., & Matuszek, C. (2018). Unsupervised selection of negative examples for grounded language learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Puig, X., Shu, T., Li, S., Wang, Z., Liao, Y.-H., Tenenbaum, J. B., Fidler, S., & Torralba, A. (2020). Watch-and-help: A challenge for social perception and human-ai collaboration..
- Pylyshyn, Z. W. (1986). *Computation and cognition: Toward a foundation for cognitive science*. The MIT Press.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. (2021). Learning transferable visual models from natural language supervision..
- Radford, A., Wu, J., Amodei, D., Amodei, D., Clark, J., Brundage, M., & Sutskever, I. (2019). Better language models and their implications.. Vol. 1, p. 2.
- Ramakrishnan, S. K., Gokaslan, A., Wijmans, E., Maksymets, O., Clegg, A., Turner, J., Undersander, E., Galuba, W., Westbury, A., Chang, A. X., et al. (2021). Habitat-Matterport 3D Dataset (HM3D): 1000 Large-scale 3D Environments for Embodied AI..
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., & Sutskever, I. (2021). Zero-shot text-to-image generation..
- Reiter, E., & Dale, R. (2000). *Building Natural Language Generation Systems*. Studies in Natural Language Processing. Cambridge University Press.
- Rosenthal, S., Biswas, J., & Veloso, M. M. (2010). An effective personal mobile robot agent through symbiotic human-robot interaction.. In *AAMAS*, Vol. 10, pp. 915–922.
- Ruis, L., Andreas, J., Baroni, M., Bouchacourt, D., & Lake, B. M. (2020). A Benchmark for Systematic Generalization in Grounded Language Understanding.. Vol. 33.

- Rust, P., Lotz, J. F., Bugliarello, E., Salesky, E., de Lhoneux, M., & Elliott, D. (2022). Language Modelling with Pixels..
- Salesky, E., Etter, D., & Post, M. (2021). Robust Open-Vocabulary Translation from Visual Text Representations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7235–7252 Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., et al. (2019). Habitat: A platform for embodied ai research..
- Schlangen, D. (2019a). Grounded Agreement Games: Emphasizing Conversational Grounding in Visual Dialogue Settings..
- Schlangen, D. (2019b). Language tasks and language games: On methodology in current natural language processing research..
- Schuster, M., & Nakajima, K. (2012). Japanese and korean voice search. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5149–5152. IEEE.
- Shah, P., Fiser, M., Faust, A., Kew, J. C., & Hakkani-Tur, D. (2018). Follownet: Robot navigation by following natural language directions with deep reinforcement learning..
- Sharma, P., Ding, N., Goodman, S., & 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)*, pp. 2556–2565.
- Shekhar, R., Pezzelle, S., Klimovich, Y., Herbelot, A., Nabi, M., Sangineto, E., & Bernardi, R. (2017). FOIL it! Find One mismatch between Image and Language caption. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 255–265.
- Shekhar, R., Venkatesh, A., Baumgärtner, T., Bruni, E., Plank, B., Bernardi, R., & Fernández, R. (2019). Beyond task success: A closer look at jointly learning to see, ask, and GuessWhat. 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)*, pp. 2578–2587.
- Shi, H., Ball, L., Thattai, G., Zhang, D., Hu, L., Gao, Q. Q., Shakiah, S., Gao, X., Padmakumar, A., Yang, B., et al. (2023). Alexa, play with robot: Introducing the first Alexa Prize SimBot Challenge on embodied AI..
- Shridhar, M., Thomason, J., Gordon, D., Bisk, Y., Han, W., Mottaghi, R., Zettlemoyer, L., & Fox, D. (2020). ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *Computer Vision and Pattern Recognition (CVPR)*.

- Shridhar, M., Yuan, X., Côté, M.-A., Bisk, Y., Trischler, A., & Hausknecht, M. (2021). ALFWorld: Aligning Text and Embodied Environments for Interactive Learning. In *International Conference on Learning Representations (ICLR)*.
- Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., & Kiela, D. (2022). Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15638–15650.
- Singh, G., Deng, F., & Ahn, S. (2021). Illiterate DALL-E Learns to Compose..
- Sivic, J., Russell, B. C., Efros, A. A., Zisserman, A., & Freeman, W. T. (2005). Discovering objects and their location in images. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, Vol. 1, pp. 370–377. IEEE.
- Socher, R., Lin, C. C.-Y., Ng, A. Y., & Manning, C. D. (2011). Parsing natural scenes and natural language with recursive neural networks. In *ICML*.
- Steels, L. (2015). *The Talking Heads experiment: Origins of words and meanings*, Vol. 1. Language Science Press.
- Storks, S., Gao, Q., Thattai, G., & Tur, G. (2021). Are We There Yet? Learning to Localize in Embodied Instruction Following. In *HAI @ AAI 2021*.
- Suglia, A., Bisk, Y., Konstas, I., Vergari, A., Bastianelli, E., Vanzo, A., & Lemon, O. (2021). An Empirical Study on the Generalization Power of Neural Representations Learned via Visual Guessing Games. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2135–2144.
- Suglia, A., et al. (2022). *Visually grounded representation learning using language games for embodied AI*. Ph.D. thesis, Heriot-Watt University.
- Suglia, A., Gao, Q., Thomason, J., Thattai, G., & Sukhatme, G. (2021). Embodied bert: A transformer model for embodied, language-guided visual task completion..
- Suglia, A., Konstas, I., Vanzo, A., Bastianelli, E., Elliott, D., Frank, S., & Lemon, O. (2020). CompGuessWhat?!: A Multi-task Evaluation Framework for Grounded Language Learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7625–7641 Online. Association for Computational Linguistics.
- Suhr, A., Lewis, M., Yeh, J., & Artzi, Y. (2017). A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 217–223.
- Suhr, A., Yan, C., Schluger, J., Yu, S., Khader, H., Mouallem, M., Zhang, I., & Artzi, Y. (2019a). 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), pp. 2119–2130.

- Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., & Artzi, Y. (2019b). A Corpus for Reasoning about Natural Language Grounded in Photographs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6418–6428 Florence, Italy. Association for Computational Linguistics.
- Tan, H., & Bansal, M. (2019a). LXMERT: Learning cross-modality encoder representations from transformers. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Tan, H., & Bansal, M. (2019b). 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)*, pp. 5103–5114.
- Tan, H., & Bansal, M. (2020). Vokenization: Improving Language Understanding via Contextualized, Visually-Grounded Supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2066–2080.
- Thomason, J., Murray, M., Cakmak, M., & Zettlemoyer, L. (2019a). Vision-and-Dialog Navigation. In *Conference on Robot Learning (CoRL)*.
- Thomason, J., Padmakumar, A., Sinapov, J., Walker, N., Jiang, Y., Yedidsion, H., Hart, J., Stone, P., & Mooney, R. J. (2019b). Improving Grounded Natural Language Understanding through Human-Robot Dialog. In *International Conference on Robotics and Automation (ICRA)*.
- Thrush, T., Jiang, R., Bartolo, M., Singh, A., Williams, A., Kiela, D., & Ross, C. (2022). Winoground: Probing Vision and Language Models for Visio-Linguistic Compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5238–5248.
- Tokunaga, T., Iida, R., Terai, A., & Kuriyama, N. (2012). The REX corpora: A collection of multimodal corpora of referring expressions in collaborative problem solving dialogues. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pp. 422–429.
- Tomasello, M. (2010). *Origins of human communication*. MIT press.
- Tu, L., Lalwani, G., Gella, S., & He, H. (2020). An empirical study on robustness to spurious correlations using pre-trained language models.. Vol. 8, pp. 621–633. MIT Press.
- Vygotsky, L. S. (2012). *Thought and language*. MIT press.
- Wang, J., Madhyastha, P., Figueiredo, J., Lala, C., & Specia, L. (2021). Multisubs: A large-scale multimodal and multilingual dataset..

- Wang, X., Wu, J., Chen, J., Li, L., Wang, Y.-F., & Wang, W. Y. (2019). VateX: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4581–4591.
- Wang, Y., & Yao, Q. (2019). Few-shot Learning: A Survey.. Vol. abs/1904.05046.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., & Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models..
- Westera, M., & Boleda, G. (2019). Don't Blame Distributional Semantics if it can't do Entailment..
- Whitehead, S., Wu, H., Ji, H., Feris, R., & Saenko, K. (2021). Separating skills and concepts for novel visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5632–5641.
- Williams, E. C., Gopalan, N., Rhee, M., & Tellex, S. (2018). Learning to parse natural language to grounded reward functions with weak supervision. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–7. IEEE.
- Winograd, T. (1971). Procedures as a representation for data in a computer program for understanding natural language. Tech. rep., MASSACHUSETTS INST OF TECH CAMBRIDGE PROJECT MAC.
- Wittgenstein, L., Anscombe, G. E. M., & Rhees, R. (1953). *Philosophische Untersuchungen. (Philosophical investigations)*. Basil Blackwell.
- Wortsman, M., Ehsani, K., Rastegari, M., Farhadi, A., & Mottaghi, R. (2019). Learning to learn how to learn: Self-adaptive visual navigation using meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6750–6759.
- Xie, N., Lai, F., Doran, D., & Kadav, A. (2018). Visual Entailment Task for Visually-Grounded Language Learning..
- Yan, C., Misra, D., Bennett, A., Walsman, A., Bisk, Y., & Artzi, Y. (2018). CHALET: Cornell house agent learning environment..
- Yang, Y., Panagopoulou, A., Lyu, Q., Zhang, L., Yatskar, M., & Callison-Burch, C. (2021). Visual Goal-Step Inference using wikiHow..
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). React: Synergizing reasoning and acting in language models..
- Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., & Tenenbaum, J. (2018). Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, pp. 1031–1042.
- Young, P., Lai, A., Hodosh, M., & Hockenmaier, J. (2014). From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions.. Vol. 2, pp. 67–78. MIT Press.

- Younger, D. H. (1967). Recognition and parsing of context-free languages in time n3.. Vol. 10, pp. 189–208. Elsevier.
- Yu, H., Lian, X., Zhang, H., & Xu, W. (2018a). Guided Feature Transformation (GFT): A Neural Language Grounding Module for Embodied Agents. In *Conference on Robot Learning*, pp. 81–98.
- Yu, H., Zhang, H., & Xu, W. (2018b). Interactive grounded language acquisition and generalization in a 2d world..
- Yu, Y., Eshghi, A., & Lemon, O. (2017a). Learning how to Learn: An Adaptive Dialogue Agent for Incrementally Learning Visually Grounded Word Meanings. In *Proceedings of the First Workshop on Language Grounding for Robotics*, pp. 10–19.
- Yu, Y., Eshghi, A., Mills, G., & Lemon, O. (2017b). The BURCHAK corpus: a Challenge Data Set for Interactive Learning of Visually Grounded Word Meanings. In *Proceedings of the Sixth Workshop on Vision and Language*, pp. 1–10.
- Zarri , S., Hough, J., Kennington, C., Manuvinakurike, R., DeVault, D., Fern andez, R., & Schlangen, D. (2016). Pentoref: A corpus of spoken references in task-oriented dialogues. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pp. 125–131.
- Zellers, R., Bisk, Y., Farhadi, A., & Choi, Y. (2019). From recognition to cognition: Visual commonsense reasoning. In *Proceedings of IEEE CVPR*.
- Zellers, R., Holtzman, A., Peters, M., Mottaghi, R., Kembhavi, A., Farhadi, A., & Choi, Y. (2021). PIGLeT: Language Grounding Through Neuro-Symbolic Interaction in a 3D World..
- Zellers, R., Lu, J., Lu, X., Yu, Y., Zhao, Y., Salehi, M., Kusupati, A., Hessel, J., Farhadi, A., & Choi, Y. (2022). Merlot reserve: Neural script knowledge through vision and language and sound. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16375–16387.
- Zellers, R., Lu, X., Hessel, J., Yu, Y., Park, J. S., Cao, J., Farhadi, A., & Choi, Y. (2021). MERLOT: Multimodal Neural Script Knowledge Models..
- Zhang, B., Hu, H., Qiu, L., Shaw, P., & Sha, F. (2021). Visually Grounded Concept Composition..
- Zhang, H., Yu, H., & Xu, W. (2018). Interactive Language Acquisition with One-shot Visual Concept Learning through a Conversational Game. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vol. 1, pp. 2609–2619.
- Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., & Gao, J. (2021). Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5579–5588.

- Zheng, D., Meng, F., Si, Q., Fan, H., Xu, Z., Zhou, J., Feng, F., & Wang, X. (2022). Spot the Difference: A Cooperative Object-Referring Game in Non-Perfectly Co-Observable Scene..
- Zhong, V., Hanjie, A. W., Wang, S., Narasimhan, K., & Zettlemoyer, L. (2021). SILG: The Multi-domain Symbolic Interactive Language Grounding Benchmark.. Vol. 34, pp. 21505–21519.
- Zhou, L., Palangi, H., Zhang, L., Hu, H., Corso, J. J., & Gao, J. (2020). Unified Vision-Language Pre-Training for Image Captioning and VQA.. In *AAAI*, pp. 13041–13049.
- Zhu, H., Neubig, G., & Bisk, Y. (2021). Few-shot language coordination by modeling theory of mind. In *International Conference on Machine Learning*, pp. 12901–12911. PMLR.
- Zhu, W., Qi, Y., Narayana, P., Sone, K., Basu, S., Wang, X., Wu, Q., Eckstein, M., & Wang, W. Y. (2022). Diagnosing Vision-and-Language Navigation: What Really Matters. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5981–5993 Seattle, United States. Association for Computational Linguistics.