

Grounding ‘Grounding’ in NLP

Khyathi Raghavi Chandu, Yonatan Bisk, Alan W Black

Language Technologies Institute

Carnegie Mellon University

{kchandu, ybisk, awb}@cs.cmu.edu

Abstract

The NLP community has seen substantial recent interest in grounding to facilitate interaction between language technologies and the world. However, as a community, we use the term broadly to reference *any* linking of text to data or non-textual modality. In contrast, Cognitive Science more formally defines “grounding” as **the process of establishing what mutual information is required for successful communication between two interlocutors** – a definition which might implicitly capture the NLP usage but differs in intent and scope.

We investigate the gap between these definitions and seek answers to the following questions: (1) *What aspects of grounding are missing from NLP tasks?* Here we present the dimensions of coordination, purviews and constraints. (2) *How is the term “grounding” used in the current research?* We study the trends in datasets, domains, and tasks introduced in recent NLP conferences. And finally, (3) *How to advance our current definition to bridge the gap with Cognitive Science?* We present ways to both create new tasks or repurpose existing ones to make advancements towards achieving a more complete sense of grounding. github.com/khyathiraghavi/Grounding-Grounding

1 Introduction

We as humans communicate and interact for a variety of reasons with a goal. We use language to seek and share information, clarify misunderstandings that conflict with our prior knowledge and contextualize based on the medium of interaction to develop and maintain social relationships. However, language is only one of the enablers of this communication reliant on several auxiliary signals and sources such as documents, media, physical context etc., This linking of concepts to context is *grounding* and within NLP context is often a knowledge base, images or discourse.

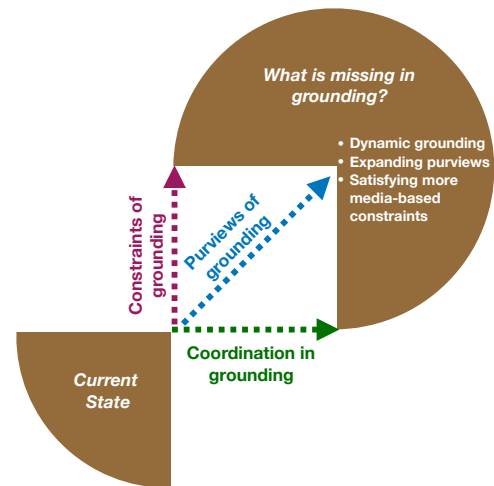


Figure 1: Dimensions of grounding – required to bridge the gap between current state of research and what is missing from a more complete sense of grounding.

In contrast, research in cognitive science defines grounding as the process of building a common ground based on shared mutual information in order to successfully communicate (Clark and Carlson, 1982; Krauss and Fussell, 1990; Clark and Brennan, 1991; Lewis, 2008). We argue that this definition subsumes NLP’s current working definition and provides concrete guidance on which phenomena are missing to ensure the naturalness and long term utility of our technologies.

In Section 2, we formalize 3 dimensions key to grounding: Coordination, Purviews and Constraints, to systematize our analysis of limitations in current work. Section 3 presents a comprehensive review of the current progress in the field including the interplay of different domains, modalities, and techniques. This analysis includes understanding when techniques have been specifically designed for a single modality, task, or form of grounding. Finally, Section 4 outlines strategies to repurpose existing datasets and tasks to align with the new

richer definition from cognitive science literature. These introspections, re-formulations, and concrete steps situate NLP ‘grounding’ in larger scientific discourse, to increase its relevance and promise.

2 Dimensions of grounding

Defining grounding loosely as *linking* or tethering concepts is insufficient to achieve a more realistic sense of grounding. Figure 1 presents the research dimensions missing from most current work.

2.1 Dimension 1: Coordination in grounding

The first and the most important dimension that bridges the gap between the two definitions of grounding is the aspect of coordination – alternatively viewed as the difference between *static* and *dynamic* grounding (Fig 2).

Static grounding is the most common type and assumes that the evidence for common ground or the gold truth for grounding is given or attained pseudo-automatically. This is demonstrated in Figure 2 (a). The sequence for this form of interaction includes: (1) human querying the agent, (2) agent querying the data or the knowledge it acquired, (3) agent retrieving and framing a response and (4) agent delivering it to the human. In this setting the common ground is the ground truth KB/data. The human and the agent have common ground by assuming its universality (i.e. no external references). Therefore, successfully grounding the query in this case relies solely on the agent being able to link the query to the data. For instance, in a scenario where a human wants to know the weather report, the accuracy of the database itself is axiomatic and we build a model for the agent to accurately retrieve the queried information in natural language.

Most current research assumes static grounding so progress is measured by the ability of the agent to *link* more concepts to more data. However, the axiomatic common ground often does not exist and needs to be established in real world scenarios.

Dynamic grounding posits that common ground is built via interactions and clarifications. The mutual information needed to communicate successfully is built via interactions including: Requesting and providing clarifications, Acknowledging or confirming the clarifications, Enacting or demonstrating to receive confirmations, and so forth. This dynamically-established-grounding guides the rest of the interaction by course-correcting any misun-




Figure 2: Coordination sequence in grounding

derstandings. The sequence of actions in dynamic grounding is demonstrated in Figure 2 (b). The steps for establishing grounding is a part of the interaction that includes: (1) The human querying the agent, (2) The agent requesting clarification or acknowledging, (3) The human clarifying or confirming. These three steps loop until a common ground is established. The remaining steps of (4) querying the data, (5) retrieving or framing a response, and (6) delivering the response, are same as that of static grounding. The agent and the human may not be on the same common ground but steps 2 and 3 loop as the conversation progresses to build this common ground. The process of successfully grounding the query not only relies on the ability of the agent to *link* the query but also to *construct the common ground from the mutually shared information* with respect to the human. Although there are efforts about clarification questioning (), the coverage of phenomena are still far from comprehensive (Benotti and Blackburn, 2021b).

Cognitive sciences in the perspective of language acquisition (Carpenter et al., 1998) present two ways of dynamic grounding via joint attention (Kolva et al., 2015; Tan et al., 2020): Dyadic joint attention and Triadic joint attention. In our case, dyadic attention describes the interaction between the human and the agent and any clarification or confirmation is done strictly between the both of them. Triadic attention also includes a tangible entity along with the human and the agent. The human can provide clarifications by gazing or pointing to this additional piece in the triad.

Summary: The community should prioritize dynamic grounding as it is more general and more accurately matches real experiences.

2.2 Dimension 2: Purviews of grounding

Next, we present the different stages behind reaching a common ground, known as purviews. Most

of the current approaches and tasks address these purviews individually and independently, while they are often co-dependent in real world scenarios.

Stage 1: Localization: The first stage is the localization of the concept either in the physical or mental contexts. This step is idiosyncratic and relates to the ability of the agent alone to localize the concept. These concepts often are also linked in a *compositional* form. For instance, consider a scenario in which the agent is to locate a ‘blue sweater’. The agent needs to understand each of the concepts of ‘blue’ and ‘sweater’ individually and then locate the composition of the whole unit. Clark and Krych (2004) from cognitive sciences demonstrate how *incremental* grounding (Schlangen and Skantze, 2009; DeVault and Traum, 2013; Eshghi et al., 2015) is performed with these compositions and show how recognition and interpretation of fragments help in this by breaking down instructions into simpler ones. This localization occurs at word, phrase and even sentence level in the language modality and pixel, object and scene level in the visual modality.

Stage 2: External Knowledge: After localizing the concept, the next step is to ensure consistency of the current context of the concept with existing knowledge. Often times, the references of grounding either match or contradict the references from our prior knowledge and external knowledge. This might lead to misunderstandings in the consequent rounds of communication. Hence, in addition to localizing the concept, it is also essential to make the concept and its attributes consistent with the available knowledge sources. Most of the current research is focused on localizing with few efforts towards extending it to maintain a consistency of the grounded concept with other knowledge sources.

Stage 3: Common sense: After establishing consistency of the concept, a human-like interaction additionally calls for grounding the common sense associated with the concept in that scenario. In addition to the basic level of practical knowledge that concerns with day to day scenarios Sap et al. (2020), the concept should also be reasoned based on that particular context. This contextual common sense moves the idiosyncratic sense towards a sense of collective understanding. For instance, if the human feels cold and asks the agent to get a blue coat, the agent needs to understand that the coat in this instance is a sweater coat and not a formal coat. This implicit common sense minimizes the effort

in building a common ground reducing articulation of meticulous details. Therefore it is essential to incorporate this explicitly in our modeling as well.

Stage 4: Personalized consensus: As a part of the evolving conversations, the references in the language evolve as well. The grounded term might have different meanings for the agent in the context with access to the history as opposed to a fresh agent without access to the history. This multi-instance multi-turn process to achieve consensus makes this collective or a shared stage continually adapting to personalization leading to better engagement (Bohus and Horvitz, 2014). In such settings, it is sufficient that the human and the agent are in consensus with the truth value of the grounded term, which need not be the same as the ground truth. This shift in the truth value of the meanings of the grounded terms often arise due to developing short-cuts for ease of communication and personalization, which is an acceptable shift as long as the communication is successful.

Summary: Common ground requires expanding to verticals of local, general, common-sense and personalized contextual knowledge.

2.3 Dimension 3: Constraints of grounding

The medium and mode of communication constrain communicative goals in practical scenarios. The number and availability of such media have increased and facilitated ubiquitous communication around the world, presenting a diversity in the mode of interaction. Motivated by this, we resurface and adapt the constraints of grounding with respect to media of interaction as defined by Clark and Brennan (1991). Here are the definitions of these constraints in the context of grounded language processing and the corresponding categorization of the majority of the representative domains in grounding satisfying different constraints.

- *Copresence:* Agent and human share the same physical environment of the data. Most of the current research in the category of embodied agents satisfy this constraint.
- *Visibility:* The data is visible to the agent and/or human. The domains of images, images & speech, videos, embodied agents satisfy this constraint.
- *Audibility:* Agent and human communicate by speaking about the data. Domains like speech, spoken image captions and videos satisfy this.
- *Cotemporality:* The agent/human receives at roughly the same time as the human/agent pro-

duces. The lag in the domains like conversations or interactive embodied agents is considered negligible and satisfy this constraint.

- *Simultaneity*: The agent and the human can send and receive at once simultaneously. Most media are cotemporal but do not engage in simultaneous interaction. This often disrupts the understanding of the current utterance and the participant may have to repeat it to avoid misunderstandings, which is commonly observed in real world scenarios.

- *Sequentiality*: The turn order of the agent and the human cannot get out of sequence. Face-to-face conversations usually follow this constraint but an email thread with active participants and the comments sections in online portals (such as Youtube, Twitch etc.,) do not necessarily follow a sequence. In such cases a reply to the message may be separated by arbitrary number of irrelevant messages. These categories are usually understudied but are commonly observed online.

- *Reviewability*: The agent reviews the common ground to the human to adapt to imperfect human memories. For instance, we reiterate full references instead of adapting to short cut references when the conversation resurfaces after a while. This is to develop a personalized adaptation between the interlocutors based on the media to enable ease of communication.

- *Revisability*: The interaction between the agent and the human indexes to a specific utterance in the conversation sequence and revise it, therefore changing the course of the interaction henceforth. Human errors are only natural in a conversation and the agent needs to be ready to rectify the previously grounded understanding.

There has been a good and continual effort in formulating tasks and datasets that satisfy the constraints of visibility, audibility and cotemporality. Contemporary efforts also see an increased interest in addressing copresence in grounded contexts. Very recently, (Benotti and Blackburn, 2021a) highlights the importance of recovering from mistakes while establishing the collaborative nature of grounding, contributing to the ability of revisability.

Summary: Key to progress is to focus on largely a blind spot in grounding: simultaneity, sequentiality & revisability to revive from mistakes.

3 Grounding ‘Grounding’

Having covered a more formal definition of grounding adapted to NLP, we turn our attention to cat-

aloging the precise usage of ‘grounding’ in our research community. We present an analysis on the various domains and techniques NLP has explored.

3.1 Data and Annotations

To this end, since our aim is to investigate how the community understands the loosely defined term ‘grounding’, we subselected all the papers that mention terms for ‘grounding’ in the title or abstract from the S2ORC data (Lo et al., 2020) between the years 1980-2020. In this way, we grounded the term ‘grounding’ in literature¹ to collect the relevant papers. We acknowledge that the papers analyzed here are not exhaustive with respect to concept of ‘grounding’.

Each of the paper is annotated with answers to the following questions: (i) is it introducing a new task? (ii) is it introducing a new dataset? (iii) what is the world scope (iv) is it working on multiple languages? (v) what are the grounding domains? (vi) what is the grounding task? (vii) what is the grounding technique?

3.2 Domains of grounding

Real world contexts we interact with are diverse and can be derived from different modalities such as textual or non-textual, each of which comprises of domains. Our categorization of these is inspired from the constraints of grounding as described in §2.3. Based on this, the modality based categorization include the following domains:

- *Textual modality comprising plain text, entities & events, knowledge bases and knowledge graphs.*
- *Non-textual modality comprising images, speech, images & speech and videos.*

Numerous other domains including *numbers and equations, colors, programs, tables, brain activity signals* etc., are studied in the context of grounding at relatively lower scale in comparison to the aforementioned ones. Each of these can further be interacted with along the variation in the coordination dimension of grounding from §2.1, that give rise to the following settings including *conversations, embodied agents and face-to-face interactions.*

3.3 Approaches to grounding

This section presents a list of approaches tailored to grounding. The obvious solution is to expand the datasets to promote a research platform. The

¹Please note that this is not an exhaustive list of papers working on grounding as there are several others that do mention this term and still work on some form of grounding.

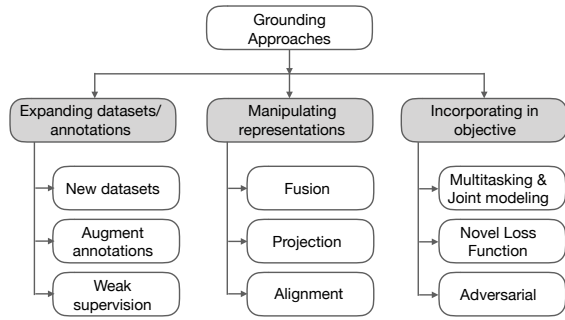


Figure 3: Categorical approaches to grounding

second is to manipulate different representations to *link* and bring them together. Finally the learning objective can leverage grounding. The sub-categories within each are presented in Figure 3.

1. Expanding datasets / annotations: The first step towards building an ecosystem for research in grounding is to curate the necessary datasets which is accomplished with expensive human efforts, augmenting existing annotations and automatically deriving annotations with weak supervision.

1a) New datasets: There has been an increase in efforts for curating new datasets with task specific annotations. These are briefly overlaid in Table 1 along with their modalities, domains and tasks.

1b) Augment annotations: These curated datasets can also be used subsequently to augment with task specific annotations instead of collecting the data from scratch, which might be more expensive.

- *Non-textual Modality:* Static grounding here includes using adversarial references to ground visual referring expressions (Akula et al., 2020), narration (Chandu et al., 2019b, 2020a), language learning (Suglia et al., 2020; Jin et al., 2020) etc.,

- *Textual Modality:* Static grounding includes entity slot filling (Bisk et al., 2016).

- *Interactive:* Though not fully dynamic grounding, some efforts here are amongst tasks like understanding spatial expressions (Udagawa et al., 2020), collaborative drawing (Kim et al., 2019) etc.,

1c) Weak supervision: While the above two are based on human efforts, we can also perform weak supervision to use a model trained to derive automatic soft annotations required for the task.

- *Non-Textual Modality:* In the visual modality, weak supervision is used in the contexts of automatic object proposals for different tasks like spoken image captioning (Srinivasan et al., 2020), visual semantic role labeling (Silberer and Pinkal, 2018), phrase grounding (Chen et al., 2019), loose

| Modality | Domain | Task | Work |
|-------------|--------|------------------------------|----------------------------------|
| Non-textual | Images | caption relevance | (Suhr et al., 2019) |
| | | multimodal MT | (Zhou et al., 2018c) |
| | | sports commentaries | (Koncel-Kedziorski et al., 2014) |
| | | semantic role labeling | (Silberer and Pinkal, 2018) |
| | | instruction following | (Han and Schlangen, 2017) |
| | | navigation | (Andreas and Klein, 2014) |
| | | causality | (Gao et al., 2016) |
| | | spatial expressions | (Kelleher et al., 2006) |
| | | spoken image captioning | (Alishahi et al., 2017) |
| | | entailment | (Vu et al., 2018) |
| Videos | | image search | (Kiros et al., 2018) |
| | | scene generation | (Chang et al., 2015) |
| | | action segmentation | (Regneri et al., 2013) |
| | | semantic parsing | (Ross et al., 2018) |
| Textual | Text | instruction following | (Liu et al., 2016) |
| | | question answering | (Lei et al., 2020) |
| | | content transfer | (Prabhumoye et al., 2019) |
| | | commonsense inference | (Zellers et al., 2018) |
| | | reference resolution | (Kennington and Schlangen, 2015) |
| | | symbol grounding | (Kameko et al., 2015) |
| | | bilingual lexicon extraction | (Laws et al., 2010) |
| | | POS tagging | (Cardenas et al., 2019) |
| | | negotiations | (Cadilhac et al., 2013) |
| | | documents | (Zhou et al., 2018b) |
| Interactive | Text | improvisation | (Cho and May, 2020) |
| | | referring expressions | (Habert et al., 2019) |
| | | emotions and styles | (Takmaz et al., 2020) |
| | Visual | media interviews | (Shuster et al., 2020) |
| | | spatial reasoning | (Majumder et al., 2020) |
| | | navigation | (Janner et al., 2018) |
| | | navigation | (Ku et al., 2020) |
| | Other | problem solving | (Li and Boyer, 2015) |

Table 1: Example datasets introduced for *grounding*.

temporal alignments between utterances and a set of events (Koncel-Kedziorski et al., 2014) etc.,

- *Textual Modality:* In the contexts of text, Tsai and Roth (2016a) work towards disambiguating concept mentions appearing in documents and grounding them in multiple KBs which is a step towards Stage 3 in §2.2. Poon (2013) perform question answering with a single database and (Parikh et al., 2015) with symbols.

Summary: While augmentation and weak supervision can be leveraged for dimensions of coordination and purviews, curating new datasets is the need of the hour to explore various constraints.

2. Manipulating representations: Grounding concepts often involves multiple modalities or representations that are *linked*. Three major methods to approach this are detailed here.

2a) Fusion and concatenation: Fusion is a very common technique in scenarios involving multiple modalities. In scenarios with a single modality, representations are often concatenated.

- *Non-textual modality:* Fusion is applied with images for tasks like referring expressions (Roy et al., 2019), SRL (Yang et al., 2016) etc., For videos, some tasks are grounding action descriptions (Regneri et al., 2013), spatio-temporal QA (Lei et al.,

2020), concept similarity (Kiela and Clark, 2015), mapping events (Fleischman and Roy, 2008) etc.,

- *Textual Modality*: With text, this is similar to concatenating context (Prabhumoye et al. (2019) perform content transfer by augmenting context).

- *Interactive*: In a conversational setting, work is explored in reference resolution (Takmaz et al., 2020; Haber et al., 2019), generating engaging response (Shuster et al., 2020), document grounded response generation Zhou et al. (2018b), etc.,

- *Others*: Nakano et al. (2003) study face-to-face grounding in instruction giving for agents.

2b) Alignment: An alternative to combining representations is aligning them with one another.

- *Non-textual modality*: Wang et al. (2020) perform phrase localization in images and Hessel et al. (2020) study temporal alignment in videos.

- *Interactive*: Han and Schlangen (2017) align GUI actions to sub-utterances in conversations and Jänner et al. (2018) align local neighborhoods to the corresponding verbalizations.

2c) Projecting into a common space: A widely used approach is to also bring the different representations on to a joint common space.

- *Non-textual modality*: Projection to a joint semantic space is used in spoken image captioning (Chrupala et al., 2017; Alishahi et al., 2017; Havard et al., 2019), bicoding for learning image attributes (Silberer and Lapata, 2014), representation learning of images (Zarriß and Schlangen, 2017) and speech (Vijayakumar et al., 2017).

- *Textual modality*: Tsai and Roth (2016b) demonstrate cross-lingual NER and mention grounding model by activating corresponding language features. Yang et al. (2019) perform imputation of embeddings for rare and unseen words by projecting a graph to the pre-trained embeddings space.

Summary: Modeling different representations effectively aid in improving both consistency across purviews and media based constraints.

3. Learning Objective: Grounding is often performed to support a more defined end purpose task. We identified 3 ways that are broadly adopted to incorporate grounding in objective functions.

3a) Multitasking and Joint Modeling: The *linking formulation* of grounding is often used as an auxiliary or dependent to model another task.

- *Non-textual Modality*: Multitasking with images is used to perform spoken image captioning (Chrupala, 2019) and grammar induction (Zhao

and Titov, 2020). Joint modeling was used in multi-resolution language grounding Koncel-Kedziorski et al. (2014), identifying referring expressions Roy et al. (2019), multimodal MT (Zhou et al., 2018c), video parsing Ross et al. (2018), learning latent semantic annotations (Qin et al., 2018) etc.,

- *Interactive*: In a conversational setting, multitasking is used to compute concept similarity judgements (Silberer and Lapata, 2014), knowledge grounded response generation (Majumder et al., 2020), grounding language instructions Hu et al. (2019). Joint modeling is used by Li and Boyer (2015) to address dialog for complex problem solving in computer programs.

3b) Loss Function: It is crucial to utilize appropriate loss designed for the specific grounding task. The main difference between multitasking and a loss function adaptation is that while multitasking reweights combinations of existing loss functions, novel loss functions are informed by the data/task at hand, adapting to a novel use case.

- *Non-textual Modality*: Grujicic et al. (2020) design soft organ distance loss to model inter and intra organ interactions for relative grounding. Ilharco et al. (2019) improve diversity in spoken captions with a masked margin softmax loss.

3c) Adversarial: Leveraging deceptive grounded inputs in an attempt to fool the model is capable of making it robust to certain errors.

- *Non-textual Modality*: Chen et al. (2018); Akula et al. (2020) present an algorithm to craft visually-similar adversarial examples.

- *Textual Modality*: Zellers et al. (2018) perform adversarial filtering and constructs a de-biased dataset by iteratively training stylistic classifiers.

Summary: Manipulating learning objective is a modeling capability aiding as an additional component in bringing grounding adjunct to several other end tasks across all the dimensions.

3.4 Analysis of trends

Based on the categories of approaches and different datasets from §3.3, we presented a representative set of analyses that highlight the major avenues that addressing the key missing pieces of work on grounding to advance future research.

Figure 4 presents the trends in the development of grounding over the past decade including: specific approaches (a,b) that presents new tasks/challenges; world scopes (Bisk et al., 2020) (c) contributing to grounding language in different

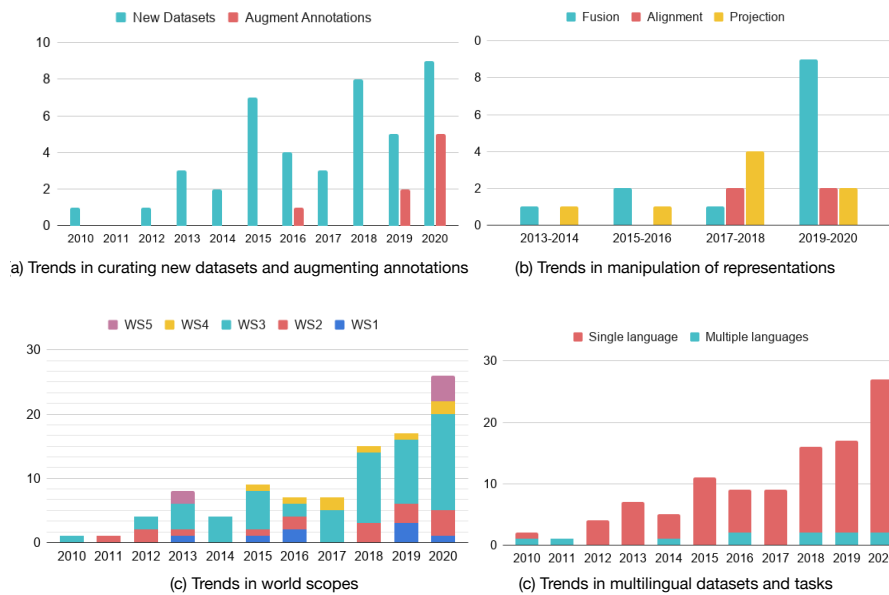


Figure 4: Analysis on the trends in grounding

data types; and multilinguality (d) contributing to a part of linguistic diversity. We also present hierarchical pie charts in Figure 5 and in Appendix to analyze the compositions of modalities and domains for these approaches. While we believe our analysis targets several of the most critical dimensions paving way for future research directions, it is not exhaustive and welcome suggestions from the community for additional analysis. For example, it is also interesting to study domain diversity, task formulation/usefulness, etc., in future.

Trends in datasets expansion: The introduction of new datasets has seen a rapid increase over the years, while there is also a subtle increasing trend in augmenting annotations to the existing datasets, as observed in Figure 4 (a). As we can see from Figure 5 (a), across all the domains, gathering new datasets seem to be prominent than augmenting them with additional annotations to repurpose the data for a new task. There seems to be a higher emphasis of expansion of datasets in the non-textual modalities, particularly in the domain of images. A similar rise is not observed in interactive settings including conversational data and interaction with embodied agents; which is the propitious way to bridge the gap towards real sense of grounding. It is indeed encouraging to see an increasing trend in the efforts for expanding datasets but *the need of the hour is to redirect some of these resources to address dynamic grounding in the coordination dimension which is scarcely studied in existing datasets.*

Trends in manipulating representations: From

Figure 4 (b), we note that the fusion technique has and is increasingly becoming popular in grounding through manipulating representations in comparison to alignment and projection. This is also observed in Figure 5 (b) with the dominance of non-textual modality. In the context of textual modality, this technique is equivalent to concatenation of the context or history in a conversation. Projecting onto a common space is the next popular technique in comparison to alignment. Similarly, we observe that the non-textual modality overwhelmingly occupies the space of manipulating representations with exceeding prominence of fusion. *Fusion and projecting onto common space currently are exceedingly used methodologies to ground within a single purview. They demonstrate a promising direction to manipulate representations across different stages to maintain consistency along the purviews.*

Trends in World Scopes: We also study the development of the field based on the definitions of the world scopes presented by Bisk et al. (2020). Based on this, last decade has seen an increasing dominance in research on world scope 3 (world of sights and sounds). However, this is limited to this scope and the same trend is not clear in world scope 4 (world of embodiment and action). An encouraging observation is the focus of the field in world scope 5 (social world) which is closer to real interactions in the last year. *We need to accelerate development of datasets and tasks in world scopes 4 and 5. It is highly recommended to take dynamic grounding scenario into account in the efforts for*

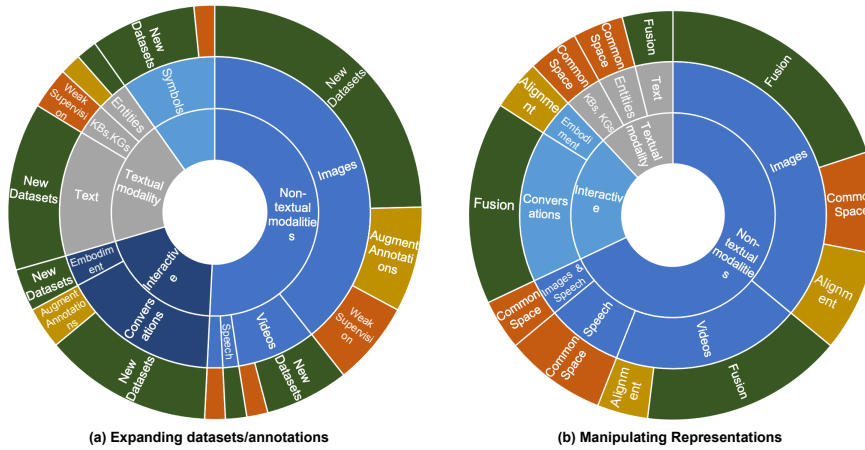


Figure 5: Analysis of Domains and Techniques

curating datasets in these scopes.

Inclusivity of multiple languages: Figure 4 (c) shows that research into grounding in multiple languages is still incredibly rare. As noted by [Bender \(2011\)](#), improvements in one language do not necessarily mandate comparable performances in other languages. *The norm for benchmarking large scale tasks still remains anglo-centric and we need serious efforts to drift this trend to identify challenges in grounding across languages. As a first step, a relatively less expensive way to navigate this dearth is to augment the annotations of existing datasets with other languages.*

4 Path Ahead: Towards New Tasks and Repurposing Existing Datasets

We presented the dimensions of grounding that require serious attention to bridge the gap between the definitions in cognitive sciences and language processing communities in §2. Based on this, we analyzed the language processing research to understand where we stand and where we fall short with the ongoing efforts in trends in grounding in §3. While we strongly advocate for efforts in building new datasets and tasks considering progress along these dimensions, we believe in a smoother transition towards this goal. Hence we present strategies to repurpose existing resources to maximum utility as we stride towards achieving grounding in real sense. In this section, we focus on concrete suggestions to improve along each of the dimensions.

Coordination: This is based on simulating interaction for dynamic grounding. As establishing a common ground is not integrated within datasets, we propose an iterative paradigm to explicitly settle on a common ground based on our priors.

The first family of methods to perform this is human-in-the-loop interactions. The traditional methods of data collection do not cater to human feedback or generation. Some recent approaches to incorporate human feedback are during data collection ([Wallace et al., 2019](#)), training ([Stiennon et al., 2020](#)), inference ([Hancock et al., 2019](#)). While the feedback in a human in the loop setting can be via scores, we argue for natural language feedback ([Wallace et al., 2019](#)) loop, which resembles human-human grounding via communication.

The second family of methods are inspired from the theory of mind ([Gopnik and Wellman, 1992](#)) to iteratively or progressively ask and clarify to establish a common ground ([Roman et al., 2020](#)). [de Vries et al. \(2017\)](#); [Suglia et al. \(2020\)](#) disambiguate or clarify the referenced object through a series of questions in a guessing game. This iterative paradigm can be related to work by [Shwartz et al. \(2020\)](#) that generates clarification questions and answers to incorporate in the task of question answering. This loop of semi-automatic generation of clarifications establishes a common ground. This is also in spirit similar to generating an explanation or a hypothesis for question answering ([Latcinnik and Berant, 2020](#)). The process of generating an acceptable explanation to human before acts as establishing a common ground.

We believe that datasets and tasks along the following 3 directions encourage dynamic grounding: (1) conversational language learning ([Chevalier-Boisvert et al., 2019](#)) or acquisition, and (2) clarification questioning and ambiguity resolution ([Shwartz et al., 2020](#)) (3) mixed initiative for grounding in conversations ([Morbini et al., 2012](#)). The need of the hour that can revolutionize this

paradigm is the *development of evaluation strategies to monitor evolution of the common ground*. This dynamic grounding data helps improve performance/robustness and encourages human’s trust while using these interactive systems.

Purviews: This is based on establishing consistency across stages of grounding with an incremental paradigm. A simple solution is a modular approach where the purviews flow into the next stage after reasonably satisfying the previous stage. The current benchmarking approaches are mostly lateral i.e., our current strategies collate multiple datasets of a single task to benchmark. This approach implicitly establishes boundaries between the purviews. In contrast, *we advocate for a longitudinal approach for benchmarking* i.e in addition to collating different datasets for a task, we also extend the purviews of the task such that the output from the previous purview flows into the next purview. An example of establishing a longitudinal benchmark for visual dialog. The tasks flow from object detection (stage 1: localization) to knowledge graphs (stage 2: external knowledge) to common sense understanding (stage 3: common sense) to empathetic dialogue (stage 4: personalization) for the *same* dataset. This helps us dissect which aspect of grounding is the model good and bad at to understand the weak areas.

Constraints: With media imposed constraints, there is a need for paradigm shift in the way these datasets are curated. The optimal way to navigate this problem is curating new datasets to *specifically focus on the less studied constraints of simultaneity, sequentiality and revisability*. At the heart of revisability in a collaborative dialog is clarification questioning and resolving ambiguities (Boni and Manandhar, 2003; Rao and III, 2018; Braslavski et al., 2017; Kumar and Black, 2020; Aliannejadi et al., 2020; Benotti and Blackburn, 2021b) However, they are rarely explored and are not systematically standardized across modalities. Transferring knowledge for shared constraints across tasks is a promising way to leverage the existing datasets.

Augment with multilingual annotations: Different languages also bring novel challenges to each of these issues (e.g. pronoun drop dialogue in Japanese, morphological alignments, etc). However, as observed in §3.4, the increase in expanding datasets is not proportionally reflected to include multiple languages. We recommend a relatively less expensive process of translating the datasets

for grounding into other languages to kick start this inclusion. The research community has already seen such efforts in image captioning with human annotated German captions in Multi30k (Elliott et al., 2016) extended from Flickr30k (Plummer et al., 2015) and Japanese captions in STAIR (Yoshikawa et al., 2017) based on MS-COCO images (Lin et al., 2014). Instead of using human annotations, some efforts have also been made to use automatic translations such as the work by Thapliyal and Soricut (2020) and denoising (Chandu et al., 2020b) extending from (Sharma et al., 2018). Not just augmentation, but there are also ongoing efforts in gathering datasets in multiple languages (Ku et al., 2020) extending (Anderson et al., 2018).

5 Conclusions

We discussed the missing pieces and dimensions that bridge the gap between the definitions of grounding in Cognitive Sciences and NLP communities. Thereby, we chart out executable actions in steering existing resources along 3 dimensions to achieve a more realistic sense of grounding. Specifically: (1) Static grounding still remains the central tenet for existing tasks and datasets. However, dynamic grounding is key moving forward. (2) Current benchmarking strategies evaluate model generalization. In tandem, we also need to steer towards longitudinal benchmarking to naturally proliferate across purviews of grounding that is closer to human interactions. (3) Constraints imposed by the interaction medium present nuanced categories of communicative goals. While discerning learning from shared constraints, we also urge the community to invest resources on revisability as a way to recover from contextually mistaken groundings. While ruminating on the above phenomena, the challenge of expanding them to multiple languages and domains still persists. We also recommend systematic evaluation of grounding along these dimensions in addition to the existing linking capabilities.

Ethical Considerations

The analytical and ontological discussion here focuses exclusively on the question of *grounding* and common ground and does not address the harmful biases inherent in these datasets. Further, the common ground for which we are advocating is culturally specific and future work that introduces tasks and data for these purposes must be explicit about who they serve (culturally and linguistically).

References

- Arjun R. Akula, Spandana Gella, Yaser Al-Onaizan, Song-Chun Zhu, and Siva Reddy. 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, ACL 2020, Online, July 5-10, 2020*, pages 6555–6565. Association for Computational Linguistics.
- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail S. Burtsev. 2020. [Convai3: Generating clarifying questions for open-domain dialogue systems \(clariq\)](#). *CoRR*, abs/2009.11352.
- Afra Alishahi, Marie Barking, and Grzegorz Chrupala. 2017. [Encoding of phonology in a recurrent neural model of grounded speech](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017), Vancouver, Canada, August 3-4, 2017*, pages 368–378. Association for Computational Linguistics.
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian D. Reid, Stephen Gould, and Anton van den Hengel. 2018. [Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3674–3683. IEEE Computer Society.
- Jacob Andreas and Dan Klein. 2014. [Grounding language with points and paths in continuous spaces](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning, CoNLL 2014, Baltimore, Maryland, USA, June 26-27, 2014*, pages 58–67. ACL.
- Leonor Becerra-Bonache, Henning Christiansen, and M Dolores Jiménez-López. 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*, pages 1–9.
- Emily M Bender. 2011. On achieving and evaluating language-independence in nlp. *Linguistic Issues in Language Technology*, 6(3):1–26.
- Luciana Benotti and Patrick Blackburn. 2021a. [Grounding as a collaborative process](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 515–531. Association for Computational Linguistics.
- Luciana Benotti and Patrick Blackburn. 2021b. [A recipe for annotating grounded clarifications](#). *CoRR*, abs/2104.08964.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph P. Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8718–8735. Association for Computational Linguistics.
- Yonatan Bisk, Siva Reddy, John Blitzer, Julia Hockenmaier, and Mark Steedman. 2016. [Evaluating induced CCG parsers on grounded semantic parsing](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2022–2027. The Association for Computational Linguistics.
- Dan Bohus and Eric Horvitz. 2014. [Managing human-robot engagement with forecasts and... um... hesitations](#). In *Proceedings of the 16th International Conference on Multimodal Interaction, ICMI 2014, Istanbul, Turkey, November 12-16, 2014*, pages 2–9. ACM.
- Marco De Boni and Suresh Manandhar. 2003. [An analysis of clarification dialogue for question answering](#). In *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, HLT-NAACL 2003, Edmonton, Canada, May 27 - June 1, 2003*. The Association for Computational Linguistics.
- Benjamin Börschinger, Bevan K. Jones, and Mark Johnson. 2011. [Reducing grounded learning tasks to grammatical inference](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1416–1425. ACL.
- Pavel Braslavski, Denis Savenkov, Eugene Agichtein, and Alina Dubatovka. 2017. [What do you mean exactly?: Analyzing clarification questions in CQA](#). In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, CHIIR 2017, Oslo, Norway, March 7-11, 2017*, pages 345–348. ACM.
- Anaïs Cadilhac, Nicholas Asher, Farah Benamara, and Alex Lascarides. 2013. [Grounding strategic conversation: Using negotiation dialogues to predict trades in a win-lose game](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 357–368. ACL.
- Ronald Cardenas, Ying Lin, Heng Ji, and Jonathan May. 2019. [A grounded unsupervised universal part-of-speech tagger for low-resource languages](#). In *Proceedings of the 2019 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2428–2439. Association for Computational Linguistics.
- Malinda Carpenter, Katherine Nagell, Michael Tomasello, George Butterworth, and Chris Moore. 1998. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the society for research in child development*, pages i–174.
- Khyathi Chandu, Shrimai Prabhumoye, Ruslan Salakhutdinov, and Alan W Black. 2019a. “my way of telling a story”: Persona based grounded story generation. In *Proceedings of the Second Workshop on Storytelling*, pages 11–21.
- Khyathi Raghavi Chandu and Alan W. Black. 2020. [Style variation as a vantage point for code-switching](#). In *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pages 4761–4765. ISCA.
- Khyathi Raghavi Chandu, Ruo-Ping Dong, and Alan W. Black. 2020a. [Reading between the lines: Exploring infilling in visual narratives](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1220–1229. Association for Computational Linguistics.
- Khyathi Raghavi Chandu, Eric Nyberg, and Alan W. Black. 2019b. [Storyboarding of recipes: Grounded contextual generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6040–6046. Association for Computational Linguistics.
- Khyathi Raghavi Chandu, Piyush Sharma, Soravit Changpinyo, Ashish Thapliyal, and Radu Soricut. 2020b. [Weakly supervised content selection for improved image captioning](#). *arXiv preprint arXiv:2009.05175*.
- Angel X. Chang, Will Monroe, Manolis Savva, Christopher Potts, and Christopher D. Manning. 2015. [Text to 3d scene generation with rich lexical grounding](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 53–62. The Association for Computer Linguistics.
- David L. Chen. 2012. [Fast online lexicon learning for grounded language acquisition](#). In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8-14, 2012, Jeju Island, Korea - Volume 1: Long Papers*, pages 430–439. The Association for Computer Linguistics.
- Hongge Chen, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, and Cho-Jui Hsieh. 2018. [Attacking visual language grounding with adversarial examples: A case study on neural image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2587–2597. Association for Computational Linguistics.
- Zhenfang Chen, Lin Ma, Wenhan Luo, and Kwan-Yee Kenneth Wong. 2019. [Weakly-supervised spatio-temporally grounding natural sentence in video](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1884–1894. Association for Computational Linguistics.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. 2019. [Babyai: A platform to study the sample efficiency of grounded language learning](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Hyundong Cho and Jonathan May. 2020. [Grounding conversations with improvised dialogues](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2398–2413. Association for Computational Linguistics.
- Grzegorz Chrupala. 2019. [Symbolic inductive bias for visually grounded learning of spoken language](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6452–6462. Association for Computational Linguistics.
- Grzegorz Chrupala, Lieke Gelderloos, and Afra Alishahi. 2017. [Representations of language in a model of visually grounded speech signal](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 613–622. Association for Computational Linguistics.
- Chenhui Chu, Mayu Otani, and Yuta Nakashima. 2018. [iparaphrasing: Extracting visually grounded paraphrases via an image](#). In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 3479–3492. Association for Computational Linguistics.
- Herbert H. Clark and Susan E. Brennan. 1991. [Grounding in communication](#). In Lauren B. Resnick,

- John M. Levine, and Stephanie D. Teasley, editors, *Perspectives on socially shared cognition*, pages 127–149. American Psychological Association.
- Herbert H Clark and Thomas B Carlson. 1982. Hearers and speech acts. *Language*, pages 332–373.
- Herbert H Clark and Meredyth A Krych. 2004. Speaking while monitoring addressees for understanding. *Journal of memory and language*, 50(1):62–81.
- David DeVault and David R. Traum. 2013. [A method for the approximation of incremental understanding of explicit utterance meaning using predictive models in finite domains](#). In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pages 1092–1099. The Association for Computational Linguistics.
- Ruo-Ping Dong, Khyathi Raghavi Chandu, and Alan W. Black. 2019. [Induction and reference of entities in a visual story](#). *CoRR*, abs/1909.09699.
- Gabriel Doyle and Michael C. Frank. 2015. [Shared common ground influences information density in microblog texts](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1587–1596. The Association for Computational Linguistics.
- Judith Eckle-Kohler. 2016. [Verbs taking clausal and non-finite arguments as signals of modality - revisiting the issue of meaning grounded in syntax](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Desmond Elliott, Stella Frank, Khalil Sima'an, and Lucia Specia. 2016. [Multi30k: Multilingual english-german image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language, hosted by the 54th Annual Meeting of the Association for Computational Linguistics, VL@ACL 2016, August 12, Berlin, Germany*. The Association for Computer Linguistics.
- Arash Eshghi, Christine Howes, Eleni Gregoromichelaki, Julian Hough, and Matthew Purver. 2015. [Feedback in conversation as incremental semantic update](#). In *Proceedings of the 11th International Conference on Computational Semantics, IWCS 2015, 15-17 April, 2015, Queen Mary University of London, London, UK*, pages 261–271. The Association for Computer Linguistics.
- Zhihao Fan, Zhongyu Wei, Siyuan Wang, and Xuanjing Huang. 2019. [Bridging by word: Image grounded vocabulary construction for visual captioning](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6514–6524. Association for Computational Linguistics.
- Michael Fleischman and Deb Roy. 2008. [Grounded language modeling for automatic speech recognition of sports video](#). In *ACL 2008, Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, June 15-20, 2008, Columbus, Ohio, USA*, pages 121–129. The Association for Computer Linguistics.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. [Multimodal compact bilinear pooling for visual question answering and visual grounding](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 457–468. The Association for Computational Linguistics.
- Qiaozhi Gao, Malcolm Doering, Shaohua Yang, and Joyce Yue Chai. 2016. [Physical causality of action verbs in grounded language understanding](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Alison Gopnik and Henry M Wellman. 1992. Why the child's theory of mind really is a theory.
- Dusan Grujicic, Gorjan Radevski, Tinne Tuytelaars, and Matthew B. Blaschko. 2020. [Learning to ground medical text in a 3d human atlas](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning, CoNLL 2020, Online, November 19-20, 2020*, pages 302–312. Association for Computational Linguistics.
- Janosch Haber, Tim Baumgärtner, Ece Takmaz, Lieke Gelderloos, Elia Bruni, and Raquel Fernández. 2019. [The photobook dataset: Building common ground through visually-grounded dialogue](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1895–1910. Association for Computational Linguistics.
- Ting Han and David Schlangen. 2017. [Grounding language by continuous observation of instruction following](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*, pages 491–496. Association for Computational Linguistics.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazaré, and Jason Weston. 2019. [Learning from dialogue after deployment: Feed yourself, chatbot!](#) In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1895–1910. Association for Computational Linguistics.

- 1: *Long Papers*, pages 3667–3684. Association for Computational Linguistics.
- William Havard, Laurent Besacier, and Jean-Pierre Chevrot. 2020. [Catplayinginthesnow: Impact of prior segmentation on a model of visually grounded speech](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning, CoNLL 2020, Online, November 19-20, 2020*, pages 291–301. Association for Computational Linguistics.
- William N. Havard, Jean-Pierre Chevrot, and Laurent Besacier. 2019. [Word recognition, competition, and activation in a model of visually grounded speech](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL 2019, Hong Kong, China, November 3-4, 2019*, pages 339–348. Association for Computational Linguistics.
- Jack Hessel, Zhenhai Zhu, Bo Pang, and Radu Soricut. 2020. [Beyond instructional videos: Probing for more diverse visual-textual grounding on youtube](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8812–8822. Association for Computational Linguistics.
- Ronghang Hu, Daniel Fried, Anna Rohrbach, Dan Klein, Trevor Darrell, and Kate Saenko. 2019. [Are you looking? grounding to multiple modalities in vision-and-language navigation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6551–6557. Association for Computational Linguistics.
- Pingping Huang, Jianhui Huang, Yuqing Guo, Min Qiao, and Yong Zhu. 2019. [Multi-grained attention with object-level grounding for visual question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3595–3600. Association for Computational Linguistics.
- Gabriel Ilharco, Yuan Zhang, and Jason Baldridge. 2019. [Large-scale representation learning from visually grounded untranscribed speech](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL 2019, Hong Kong, China, November 3-4, 2019*, pages 55–65. Association for Computational Linguistics.
- Michaela Jänner, Karthik Narasimhan, and Regina Barzilay. 2018. [Representation learning for grounded spatial reasoning](#). *Trans. Assoc. Comput. Linguistics*, 6:49–61.
- Sujay Kumar Jauhar, Chris Dyer, and Eduard H. Hovy. 2015. [Ontologically grounded multi-sense representation learning for semantic vector space models](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 683–693. The Association for Computational Linguistics.
- Xisen Jin, Junyi Du, Arka Sadhu, Ram Nevatia, and Xiang Ren. 2020. [Visually grounded continual learning of compositional phrases](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2018–2029. Association for Computational Linguistics.
- Mark Johnson, Katherine Demuth, and Michael C. Frank. 2012. [Exploiting social information in grounded language learning via grammatical reduction](#). In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8-14, 2012, Jeju Island, Korea - Volume 1: Long Papers*, pages 883–891. The Association for Computer Linguistics.
- Hirotaaka Kameko, Shinsuke Mori, and Yoshimasa Tsuruoka. 2015. [Can symbol grounding improve low-level nlp? word segmentation as a case study](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 2298–2303. The Association for Computational Linguistics.
- Kazuya Kawakami, Chris Dyer, and Phil Blunsom. 2019. [Learning to discover, ground and use words with segmental neural language models](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6429–6441. Association for Computational Linguistics.
- John D. Kelleher, Geert-Jan M. Kruijff, and Fintan J. Costello. 2006. [Proximity in context: An empirically grounded computational model of proximity for processing topological spatial expressions](#). In *ACL 2006, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, Sydney, Australia, 17-21 July 2006*. The Association for Computer Linguistics.
- Casey Kennington and David Schlangen. 2015. [Simple learning and compositional application of perceptually grounded word meanings for incremental reference resolution](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 292–301. The Association for Computer Linguistics.
- Douwe Kiela, Luana Bulat, and Stephen Clark. 2015. [Grounding semantics in olfactory perception](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language*

- Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 2: Short Papers*, pages 231–236. The Association for Computer Linguistics.
- Douwe Kiela and Stephen Clark. 2015. [Multi- and cross-modal semantics beyond vision: Grounding in auditory perception](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 2461–2470. The Association for Computational Linguistics.
- Jin-Hwa Kim, Nikita Kitaev, Xinlei Chen, Marcus Rohrbach, Byoung-Tak Zhang, Yuandong Tian, Dhruv Batra, and Devi Parikh. 2019. [Codraw: Collaborative drawing as a testbed for grounded goal-driven communication](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6495–6513. Association for Computational Linguistics.
- Jamie Ryan Kiros, William Chan, and Geoffrey E. Hinton. 2018. [Illustrative language understanding: Large-scale visual grounding with image search](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 922–933. Association for Computational Linguistics.
- Nikolina Koleva, Martin Villalba, Maria Staudte, and Alexander Koller. 2015. [The impact of listener gaze on predicting reference resolution](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 2: Short Papers*, pages 812–817. The Association for Computer Linguistics.
- Rik Koncel-Kedziorski, Hannaneh Hajishirzi, and Ali Farhadi. 2014. [Multi-resolution language grounding with weak supervision](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 386–396. ACL.
- Robert M Krauss and Susan R Fussell. 1990. Mutual knowledge and communicative effectiveness. *Intellectual teamwork: Social and technological foundations of cooperative work*, pages 111–146.
- Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. 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 2020, Online, November 16-20, 2020*, pages 4392–4412. Association for Computational Linguistics.
- Vaibhav Kumar and Alan W. Black. 2020. [Clarq: A large-scale and diverse dataset for clarification question generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7296–7301. Association for Computational Linguistics.
- Veronica Latcinnik and Jonathan Berant. 2020. [Explaining question answering models through text generation](#). *CoRR*, abs/2004.05569.
- Florian Laws, Lukas Michelbacher, Beate Dorow, Christian Scheible, Ulrich Heid, and Hinrich Schütze. 2010. [A linguistically grounded graph model for bilingual lexicon extraction](#). In *COLING 2010, 23rd International Conference on Computational Linguistics, Posters Volume, 23-27 August 2010, Beijing, China*, pages 614–622. Chinese Information Processing Society of China.
- Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. 2020. [TVQA+: spatio-temporal grounding for video question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8211–8225. Association for Computational Linguistics.
- David Lewis. 2008. *Convention: A philosophical study*. John Wiley & Sons.
- Xiaolong Li and Kristy Boyer. 2015. [Semantic grounding in dialogue for complex problem solving](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 841–850. The Association for Computational Linguistics.
- Zekang Li, Cheng Niu, Fandong Meng, Yang Feng, Qian Li, and Jie Zhou. 2019. [Incremental transformer with deliberation decoder for document grounded conversations](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 12–21. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft COCO: common objects in context](#). In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer.
- Changsong Liu, Lanbo She, Rui Fang, and Joyce Y. Chai. 2014. [Probabilistic labeling for efficient referential grounding based on collaborative discourse](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL*

- 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 2: Short Papers, pages 13–18. The Association for Computer Linguistics.
- Changsong Liu, Shaohua Yang, Sari Saba-Sadiya, Nishant Shukla, Yunzhong He, Song-Chun Zhu, and Joyce Yue Chai. 2016. [Jointly learning grounded task structures from language instruction and visual demonstration](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1482–1492. The Association for Computational Linguistics.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. 2020. [S2ORC: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4969–4983, Online. Association for Computational Linguistics.
- Minh-Thang Luong, Michael C. Frank, and Mark Johnson. 2013. [Parsing entire discourses as very long strings: Capturing topic continuity in grounded language learning](#). *Trans. Assoc. Comput. Linguistics*, 1:315–326.
- Bodhisattwa Prasad Majumder, Shuyang Li, Jianmo Ni, and Julian J. McAuley. 2020. [Interview: Large-scale modeling of media dialog with discourse patterns and knowledge grounding](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8129–8141. Association for Computational Linguistics.
- Alexandre Blondin Massé, Guillaume Chicoisne, Yasmine Gargouri, Stevan Harnad, Odile Marcotte, and Olivier Picard. 2008. [How is meaning grounded in dictionary definitions?](#) In *Coling 2008: Proceedings of the 3rd Textgraphs workshop on Graph-based Algorithms for Natural Language Processing*, pages 17–24.
- Brian McMahan and Matthew Stone. 2015. [A bayesian model of grounded color semantics](#). *Trans. Assoc. Comput. Linguistics*, 3:103–115.
- Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. [Colors in context: A pragmatic neural model for grounded language understanding](#). *Trans. Assoc. Comput. Linguistics*, 5:325–338.
- Fabrizio Morbini, Eric Forbell, David DeVault, Kenji Sagae, David R. Traum, and Albert A. Rizzo. 2012. [A mixed-initiative conversational dialogue system for healthcare](#). In *Proceedings of the SIGDIAL 2012 Conference, The 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 5-6 July 2012, Seoul National University, Seoul, South Korea*, pages 137–139. The Association for Computer Linguistics.
- Yukiko I. Nakano, Gabe Reinstein, Tom Stocky, and Justine Cassell. 2003. [Towards a model of face-to-face grounding](#). In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics, 7-12 July 2003, Sapporo Convention Center, Sapporo, Japan*, pages 553–561. ACL.
- Sushobhan Nayak and Amitabha Mukerjee. 2012. [Grounded language acquisition: A minimal commitment approach](#). In *COLING 2012, 24th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, 8-15 December 2012, Mumbai, India*, pages 2059–2076. Indian Institute of Technology Bombay.
- Joel Nothman, Matthew Honnibal, Ben Hachey, and James R. Curran. 2012. [Event linking: Grounding event reference in a news archive](#). In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8-14, 2012, Jeju Island, Korea - Volume 2: Short Papers*, pages 228–232. The Association for Computer Linguistics.
- Tim Oates. 2003. [Grounding word meanings in sensor data: Dealing with referential uncertainty](#). In *Proceedings of the HLT-NAACL 2003 workshop on Learning word meaning from non-linguistic data*, pages 62–69.
- Brian E. Pangburn, S. Sitharama Iyengar, Robert C. Mathews, and Jonathan P. Ayo. 2003. [EBLA: A perceptually grounded model of language acquisition](#). In *Proceedings of the HLT-NAACL 2003 Workshop on Learning Word Meaning from Non-Linguistic Data*, pages 46–53.
- Nikolaos Pappas, Phoebe Mulcaire, and Noah A. Smith. 2020. [Grounded compositional outputs for adaptive language modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1252–1267. Association for Computational Linguistics.
- Ankur P. Parikh, Hoifung Poon, and Kristina Toutanova. 2015. [Grounded semantic parsing for complex knowledge extraction](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 756–766. The Association for Computational Linguistics.
- Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. [Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2641–2649. IEEE Computer Society.
- Hoifung Poon. 2013. [Grounded unsupervised semantic parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*,

- ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers, pages 933–943. The Association for Computer Linguistics.
- Shrimai Prabhunoye, Chris Quirk, and Michel Galley. 2019. [Towards content transfer through grounded text generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2622–2632. Association for Computational Linguistics.
- Guanghui Qin, Jin-Ge Yao, Xuening Wang, Jinpeng Wang, and Chin-Yew Lin. 2018. [Learning latent semantic annotations for grounding natural language to structured data](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3761–3771. Association for Computational Linguistics.
- Sudha Rao and Hal Daumé III. 2018. [Learning to ask good questions: Ranking clarification questions using neural expected value of perfect information](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2737–2746. Association for Computational Linguistics.
- Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. [Grounding action descriptions in videos](#). *Trans. Assoc. Comput. Linguistics*, 1:25–36.
- Homero Roman Roman, Yonatan Bisk, Jesse Thomason, Asli Celikyilmaz, and Jianfeng Gao. 2020. [Rmm: A recursive mental model for dialog navigation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1732–1745.
- Candace Ross, Andrei Barbu, Yevgeni Berzak, Batushig Myanganbayar, and Boris Katz. 2018. [Grounding language acquisition by training semantic parsers using captioned videos](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2647–2656. Association for Computational Linguistics.
- Deb Roy, Kai-Yuh Hsiao, and Nikolaos Mavridis. 2003. [Conversational robots: building blocks for grounding word meaning](#). In *Proceedings of the HLT-NAACL 2003 workshop on Learning word meaning from non-linguistic data*, pages 70–77.
- Subhro Roy, Michael Noseworthy, Rohan Paul, Daehyung Park, and Nicholas Roy. 2019. [Leveraging past references for robust language grounding](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL 2019, Hong Kong, China, November 3-4, 2019*, pages 430–440. Association for Computational Linguistics.
- Subhro Roy, Shyam Upadhyay, and Dan Roth. 2016. [Equation parsing : Mapping sentences to grounded equations](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1088–1097. The Association for Computational Linguistics.
- Maarten Sap, Vered Shwartz, Antoine Bosselut, Yejin Choi, and Dan Roth. 2020. [Commonsense reasoning for natural language processing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts, ACL 2020, Online, July 5, 2020*, pages 27–33. Association for Computational Linguistics.
- David Schlangen and Gabriel Skantze. 2009. [A general, abstract model of incremental dialogue processing](#). In *EACL 2009, 12th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference, Athens, Greece, March 30 - April 3, 2009*, pages 710–718. The Association for Computer Linguistics.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2556–2565. Association for Computational Linguistics.
- Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. 2019. [Visually grounded neural syntax acquisition](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1842–1861. Association for Computational Linguistics.
- Robik Shrestha, Kushal Kafle, and Christopher Kanan. 2020. [A negative case analysis of visual grounding methods for VQA](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8172–8181. Association for Computational Linguistics.
- Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. 2020. [Image-chat: Engaging grounded conversations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2414–2429. Association for Computational Linguistics.
- Ekaterina Shutova, Niket Tandon, and Gerard de Melo. 2015. [Perceptually grounded selectional preferences](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of*

- Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 950–960. The Association for Computational Linguistics.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Unsupervised commonsense question answering with self-talk](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4615–4629. Association for Computational Linguistics.
- Carina Silberer and Mirella Lapata. 2014. [Learning grounded meaning representations with autoencoders](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 721–732. The Association for Computational Linguistics.
- Carina Silberer and Manfred Pinkal. 2018. [Grounding semantic roles in images](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2616–2626. Association for Computational Linguistics.
- Georgios P. Spithourakis, Isabelle Augenstein, and Sebastian Riedel. 2016. [Numerically grounded language models for semantic error correction](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 987–992. The Association for Computational Linguistics.
- Tejas Srinivasan, Ramon Sanabria, Florian Metze, and Desmond Elliott. 2020. [Fine-grained grounding for multimodal speech recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 2667–2677. Association for Computational Linguistics.
- Luc Steels. 2004. [Constructivist development of grounded construction grammar](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics, 21-26 July, 2004, Barcelona, Spain*, pages 9–16. ACL.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. 2020. [Learning to summarize from human feedback](#). *CoRR*, abs/2009.01325.
- Michael Strube and Udo Hahn. 1999. Functional centering - grounding referential coherence in information structure. *Comput. Linguistics*, 25(3):309–344.
- Alessandro Suglia, Ioannis Konstas, Andrea Vanzo, Emanuele Bastianelli, Desmond Elliott, Stella Frank, and Oliver Lemon. 2020. [Compguesswhat?!: A multi-task evaluation framework for grounded language learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7625–7641. Association for Computational Linguistics.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. [A corpus for reasoning about natural language grounded in photographs](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6418–6428. Association for Computational Linguistics.
- Ece Takmaz, Mario Giulianelli, Sandro Pezzelle, Arabella Sinclair, and Raquel Fernández. 2020. Refer, reuse, reduce: Grounding subsequent references in visual and conversational contexts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4350–4368.
- Xiang Zhi Tan, Sean Andrist, Dan Bohus, and Eric Horvitz. 2020. [Now, over here: Leveraging extended attentional capabilities in human-robot interaction](#). In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2020, Cambridge, UK, March 23-26, 2020*, pages 468–470. ACM.
- Ashish V. Thapliyal and Radu Soriccut. 2020. [Cross-modal language generation using pivot stabilization for web-scale language coverage](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 160–170. Association for Computational Linguistics.
- Chen-Tse Tsai and Dan Roth. 2016a. [Concept grounding to multiple knowledge bases via indirect supervision](#). *Trans. Assoc. Comput. Linguistics*, 4:141–154.
- Chen-Tse Tsai and Dan Roth. 2016b. [Illinois cross-lingual wikifier: Grounding entities in many languages to the english wikipedia](#). In *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference System Demonstrations, December 11-16, 2016, Osaka, Japan*, pages 146–150. ACL.
- Takuma Udagawa, Takato Yamazaki, and Akiko Aizawa. 2020. [A linguistic analysis of visually grounded dialogues based on spatial expressions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 750–765. Association for Computational Linguistics.
- Ashwin K. Vijayakumar, Ramakrishna Vedantam, and Devi Parikh. 2017. [Sound-word2vec: Learning word representations grounded in sounds](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*,

- pages 920–925. Association for Computational Linguistics.
- Harm de Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron C. Courville. 2017. [Guesswhat?! visual object discovery through multi-modal dialogue](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4466–4475. IEEE Computer Society.
- Hoa Trong Vu, Claudio Greco, Aliia Erofeeva, Somayeh Jafaritazehjan, Guido Linders, Marc Tanti, Alberto Testoni, Raffaella Bernardi, and Albert Gatt. 2018. [Grounded textual entailment](#). In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 2354–2368. Association for Computational Linguistics.
- Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. *Transactions of the Association for Computational Linguistics*, 7:387–401.
- Qinxin Wang, Hao Tan, Sheng Shen, Michael W. Mahoney, and Zhewei Yao. 2020. [MAF: multimodal alignment framework for weakly-supervised phrase grounding](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2030–2038. Association for Computational Linguistics.
- Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. [Conversational graph grounded policy learning for open-domain conversation generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1835–1845. Association for Computational Linguistics.
- Shaohua Yang, Qiaozhi Gao, Changsong Liu, Caiming Xiong, Song-Chun Zhu, and Joyce Y. Chai. 2016. [Grounded semantic role labeling](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 149–159. The Association for Computational Linguistics.
- Tsung-Yen Yang, Andrew S. Lan, and Karthik Narasimhan. 2020. [Robust and interpretable grounding of spatial references with relation networks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 1908–1923. Association for Computational Linguistics.
- Ziyi Yang, Chenguang Zhu, Vin Sachidananda, and Eric Darve. 2019. [Embedding imputation with grounded language information](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3356–3361. Association for Computational Linguistics.
- Yuya Yoshikawa, Yutaro Shigeto, and Akikazu Takeuchi. 2017. [STAIR captions: Constructing a large-scale japanese image caption dataset](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 417–421. Association for Computational Linguistics.
- Sina Zarrieß and David Schlangen. 2017. [Deriving continuous grounded meaning representations from referentially structured multimodal contexts](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 959–965. Association for Computational Linguistics.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 93–104. Association for Computational Linguistics.
- Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2020. [Grounded conversation generation as guided traverses in commonsense knowledge graphs](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2031–2043. Association for Computational Linguistics.
- Yanpeng Zhao and Ivan Titov. 2020. [Visually grounded compound pcfgs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4369–4379. Association for Computational Linguistics.
- Victor Zhong, Mike Lewis, Sida I. Wang, and Luke Zettlemoyer. 2020. [Grounded adaptation for zero-shot executable semantic parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6869–6882. Association for Computational Linguistics.
- Ben Zhou, Daniel Khashabi, Chen-Tse Tsai, and Dan Roth. 2018a. [Zero-shot open entity typing as type-compatible grounding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2065–2076. Association for Computational Linguistics.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W. Black. 2018b. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 708–713. Association for Computational Linguistics.

Mingyang Zhou, Runxiang Cheng, Yong Jae Lee, and Zhou Yu. 2018c. [A visual attention grounding neural model for multimodal machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3643–3653. Association for Computational Linguistics.

A Examples for dimensions of grounding

Static Grounding: In static grounding, when you ask an agent “*Can you place the dragon fruit on the rack?*”, the agent links the entities and places the dragon fruit on the rack. The challenge here is mainly the linking part which is crucial to ensure it accurately understood the instruction.

Dynamic Grounding: The same is not true for dynamic grounding. There are primarily 2 ways to materialize this. First, with respect to language learning: What if the agent does not know dragon fruit? The agent needs to first ask “*What is a dragon fruit?*”, and the human provides an answer. Lets say the human responded by describing the physical attributes such as *reddish pink fruit* and/or a spatial reference by referring to it as *the fruit on the bottom left*. The important aspect here is that the agent asks and learns what a dragon fruit is and use this knowledge later.

The second is ambiguity resolution. Consider a scenario where there are multiple racks. It is very natural for a human to ask the agent which rack to resolve ambiguity. We expect the same from the agent to ask a clarifying question to resolve ambiguity and then place it on the second rack.

Purviews - Localization: Consider this example of a conversation between an agent and a human.

Human: *What is the name of the role Robert Downey Jr played in Avengers?*

Agent: *He played the role of Tony Stark, and sometimes is also referred to as Iron Man.*

The agent begins by localizing and linking Robert Downey Jr to Tony Stark and Iron Man to provide the appropriate answer to the query.

Purviews - External Knowledge: However, natural conversations also extend beyond the purview of localization to discuss a broadened scope involving external knowledge of the context including entities, actions etc., For example, consider this conversation which seems to be a natural continuation to the earlier one.

Human: *Is he the head of SHIELD?*

Agent: *Tony Stark has never been the head of SHIELD in the movies but has been the acting head upon Maria Hill’s suggestion in the Comics.*

Once we localized Tony Stark, asking additional information like whether he is the head of SHIELD is natural in conversations; However, access to required external knowledge is rarely present in the datasets as well as evaluated. Here, we need to refer to external sources spanning from movies to comics to conclude that he has been the acting head in the comics but was never in the movies.

Purviews - Common sense: One of the branches of natural progression to this context can extend to the following turns:

Human: *How long was the contract between Tony Stark and Marvel?*

Agent: *Tony Stark is the name of the character in Marvel. Would you like to know the contract length for Robert Downey Jr who played the role?*

Here, the agent needs to understand that Tony Stark is not a real person, but is a character in Marvel. Hence, any contract is with the actor but not the character who played the role. The agent needs to have the common sense to understand this and clarify the question.

Purviews - Personalization: Upon a continuous exchange regarding this topic (and perhaps a few other times earlier), the agent needs to adapt and personalize to the interacting human over time.

Human: *Can you give me any movie suggestions?*

Agent: *Yes, since you like Disney movies and seem interested in Robert Downey Jr, would you like to watch “Dolittle”?*

Having discussed about Robert Downey Jr in prior contexts and retaining from the prior interactions that the human likes Disney movies, when the human asks about a movie recommendation, the agent continually learns and contextually suggests Robert Downey Jr’s Disney movie “Dolittle” as a recommendation.

Constraints - Copresence: Modality is an important medium that affects communicative goals and the nature of interaction. Here is an example in a copresent environment.

Human: *I want to play with my cat. Can you get me the ball on your right?*

In the above example, the human and the agent

| | Modality | | | Cue | | | | |
|-------------------------|------------|------------|------------|---------------|--------------|---------------|---------------|--------------|
| | Copresence | Visibility | Audibility | Cotemporality | Simultaneity | Sequentiality | Reviewability | Revisability |
| Face-to-face | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | | |
| Telephone | | | ✓ | ✓ | ✓ | ✓ | | |
| Video Teleconference | | ✓ | ✓ | ✓ | ✓ | ✓ | | |
| Terminal Teleconference | | | | ✓ | | ✓ | ✓ | |
| Answering Machines | | | ✓ | | | | ✓ | |
| E-mail | | | | | | | ✓ | ✓ |
| Letters | | | | | | | ✓ | ✓ |

Table 2: Constraints of grounding along with their medium of communication (Clark and Brennan, 1991)

are copresent in the same environment. The above utterance for instance, includes executable actions in the environment along with references being either person-centric or agent-centric.

Constraints - Visibility: Certain communications like in the cases of visual question answering or visual dialog only presents a visible medium to interact *about*. The interaction requires information from an image or a video, but does not necessarily include executable actions or cater to external knowledge of the information. For example, with an access to an image a human can ask a question like the following:

Human: How many peaks are there in those mountain ranges?

Constraints - Audibility: This modality constrains the information scope to be within speech signals that are only heard and do not contain any visual or copresent information.

Table 2 presents the constraint of grounding.

B Further survey and categories

Here is a brief elaboration of the datasets presented in Table 1.

New datasets: The first solution to curate the entire dataset with annotations designed for the task.

- *Non-textual Modality:* For images, new datasets are curated for a variety of tasks including caption relevance (Suhr et al., 2019), multimodal MT (Zhou et al., 2018c), soccer commentaries (Koncel-Kedziorski et al., 2014) semantic role labeling (Silberer and Pinkal, 2018), instruction following (Han and Schlangen, 2017), navigation (Andreas and Klein, 2014), understanding physical causality of actions (Gao et al., 2016), understanding topological spatial expressions (Kelleher et al., 2006), spoken image captioning (Alishahi et al., 2017), entail-

ment (Vu et al., 2018), image search (Kiros et al., 2018), scene generation (Chang et al., 2015), etc., Coming to videos, datasets have become popular for several tasks like identifying action segments (Regneri et al., 2013), semantic parsing (Ross et al., 2018), instruction following from visual demonstration (Liu et al., 2016), spatio-temporal question answering (Lei et al., 2020), etc.,

- *Textual Modality:* Within text, there are several datasets for tasks like content transfer (Prabhumoye et al., 2019), commonsense inference (Zellers et al., 2018), reference resolution (Kennington and Schlangen, 2015), symbol grounding (Kameko et al., 2015), studying linguistic and non-linguistic contexts in microblogs (Doyle and Frank, 2015), bilingual lexicon extraction (Laws et al., 2010), universal part-of-speech tagging for low resource languages (Cardenas et al., 2019), entity linking and reference (Nothman et al., 2012) etc.,

- *Other:* More static grounding datasets correspond to tasks like identifying phrases representing variables (Roy et al., 2016), conceptual similarity in olfactory data (Kiela et al., 2015), identifying colors from descriptions (Monroe et al., 2017), correcting numbers (Spithourakis et al., 2016) etc.,

- *Interactive:* Coming to an interactive setting, the datasets span tasks like conversations based on negotiations (Cadilhac et al., 2013), referring expressions from images (Haber et al., 2019; Takmaz et al., 2020), emotions and styles (Shuster et al., 2020), media interviews (Majumder et al., 2020), documents (Zhou et al., 2018b), improvisation (Cho and May, 2020), problem solving (Li and Boyer, 2015), spatial reasoning in a simulated environment (Jänner et al., 2018), navigation (Ku et al., 2020) etc.,

In addition, there are several other techniques used to ground phenomenon in real world contexts.

In addition to the techniques discussed in the paper, we also studied the categorization based on stratification, which is explained here.

Stratification: The stratification technique characterizes the input or the model to explicitly cater to the compositionality property. This can be done by either breaking down the input to meaningful compositions or building the model to compose the representations. Utilizing grammatical rules need not necessarily lead to compositions, although there is an overlap between these two techniques.

A common strategy when language is involved is leveraging **syntax and parsing**. In the domain of images, [Udagawa et al. \(2020\)](#) design an annotation protocol to capture important linguistic structures based on predicate-argument structure, modification and ellipsis to utilize linguistic structures based on spatial expressions. [Becerra-Bonache et al. \(2018\)](#) study linguistic complexity from a developmental point of view by using syntactic rules to provide data to a learner, that identifies the underlying language from this data. [Shi et al. \(2019\)](#) use image-caption pairs to extract constituents from text, based on the assumption that similar spans should be matched to similar visual objects and these concrete spans form constituents. [Kelleher et al. \(2006\)](#) use combinatory categorial grammar (CCG) to build a psycholinguistic based model to predict absolute proximity ratings to identify spatial proximity between objects in a natural scene. [Ross et al. \(2018\)](#) employ CCG-based parsing to a fixed set of unary and binary derivation rules to generate semantic parses for videos.

- *Textual Modality:* [Johnson et al. \(2012\)](#) study the modeling the task of inferring the referred objects using social cues and grammatical reduction strategies in language acquisition. [Eckle-Kohler \(2016\)](#) attempt to understand meaning in syntax by a multi-perspective semantic characterization of the inferred classes in multiple lexicons. [Chen \(2012\)](#) develop a context-free grammar to understand formal navigation instructions that correspond better with words or phrases in natural language. [Börschinger et al. \(2011\)](#) study the probabilistic context-free grammar learning task using the inside-out algorithm in game commentaries. CCG parsers are also used to perform entity slot filling task ([Bisk et al., 2016](#)). When applied to question answering over a database, dependency rules are used to model the edge states as well as transitions such as the work done by using a treeHMM ([Poon, 2013](#)).

- *Other:* [Roy et al. \(2016\)](#) perform equation parsing that identifies noun phrases in a given sentence representing variables using high precision mathe-

matical lexicon to generate the correct relations in the equations. [Parikh et al. \(2015\)](#) perform prototype driven learning to learn a semantic parser in tables of nested events and unannotated text.

- *Interactive:* [Luong et al. \(2013\)](#) use parsing and grammar induction to produce a parser capable of representing full discourses and dialogs. [Steels \(2004\)](#) study games and embodied agents by modeling a constructivist approach based on invention, abduction and induction to language development.

Another frequently used technique when language is involved is by leveraging the principle of **compositionality**. This implies that the meaning of a complex expression is determined by the meanings of its constituents and how they interact with one another.

- *Non-textual Modality:* In the domain of images, [Suhr et al. \(2019\)](#) present a new dataset to understand challenges in language grounding including compositionality, semantic diversity and visual reasoning. [Shi et al. \(2019\)](#), discussed earlier also use grammar rules to compose the inputs. [Koncel-Kedziorski et al. \(2014\)](#) leverage the compositional nature of language to understand professional soccer commentaries. In the domain of videos, [Nayak and Mukerjee \(2012\)](#) study language acquisition by segmenting the world to obtain a meaning space and combining them to get a linguistic pattern.

- *Textual Modality:* With ontologies, [Pappas et al. \(2020\)](#) perform adaptive language modeling to other domains to get a fully compositional output embedding layer which is further grounded in information from a structured lexicon.

- *Interactive:* [Roy et al. \(2003\)](#) work on grounding word meanings for robots by composing perceptual, procedural, and affordance representations.

Hierarchical modeling is also applied to show effect of introducing phone, syllable, or word boundaries in spoken captions ([Havard et al., 2020](#)) and with a compact bilinear pooling in visual question answering ([Fukui et al., 2016](#)).

There is some work that presents a bayesian probabilistic formulation to learn referential grounding in dialog ([Liu et al., 2014](#)), user preferences ([Cadilhac et al., 2013](#)), color descriptions ([McMahan and Stone, 2015](#); [Andreas and Klein, 2014](#)).

A huge chunk of work also focus on leveraging attention mechanism for grounding multimodal phenomenon in images ([Srinivasan et al., 2020](#); [Chu et al., 2018](#); [Huang et al., 2019](#); [Fan et al., 2019](#); [Vu et al., 2018](#); [Kawakami et al., 2019](#); [Dong et al.,](#)

2019), videos (Lei et al., 2020; Chen et al., 2019) and navigation of embodied agents (Yang et al., 2020), etc.,

Some approach this using data structures such as graphs in the domains of grounding images (Chang et al., 2015; Liu et al., 2014), videos (Liu et al., 2016), text (Laws et al., 2010; Chen, 2012; Massé et al., 2008), entities (Zhou et al., 2018a), knowledge graphs and ontologies (Jauhar et al., 2015; Zhang et al., 2020) and interactive settings Jauhar et al. (2015); Xu et al. (2020).

Here is the technique wise representation of these categories of models in the literature.




Figure 6: Papers addressing stratification in grounding

C Prevalance of modalities and constraints

Here is the distribution of the papers studying various tasks based on the constraints imposed by the medium.




Figure 7: Papers addressing different constraints of grounding

As we can see, a major concentration of these efforts lie in grounding visual and textual media,

while a few cater to audibility i.e speech signals. Papers studying dialog are the main representatives of the constraints for sequentiality and co-temporality.

D Nuanced modeling variations for grounding

Here is a more nuanced and finer grained categorization of the various modeling techniques used in literature for grounding. Figure 8 presents these categories in depth.




Figure 8: Modeling variations in papers studying grounding

As discussed in the paper, most of the literature is focused on grounding in static visual modality. Attention based methods dominate the rest of the methods in both textual and non-textual modalities closely followed by graph based methods as observed in these trends.

This is not an exhaustive study of all the techniques that present grounding, but are some of the representative categories. Here are more studies that perform grounding with various techniques such as clustering (Shutova et al., 2015; Cardenas et al., 2019) regularization (Shrestha et al., 2020), CRFs (Gao et al., 2016), classification (Pangburn et al., 2003; Monroe et al., 2017), linguistic theories (Strube and Hahn, 1999), iterative refinement (Li et al., 2019; Chandu and Black, 2020), language modeling (Spithourakis et al., 2016; Cho and May, 2020), nearest neighbors (Kiela et al., 2015), contextual fusion (Chandu et al., 2019a), mutual information (Oates, 2003), cycle consistency (Zhong et al., 2020) etc.,