

Multimodal Grounding for Language Processing

Lisa Beinborn^{◊*}

Teresa Botschen^{*△}

Iryna Gurevych [△]

◊ Language Technology Lab, University of Duisburg-Essen

△ Ubiquitous Knowledge Processing Lab (UKP) and Research Training Group AIPHES

Department of Computer Science, Technische Universität Darmstadt

www.ukp.tu-darmstadt.de

Abstract

This survey discusses how recent developments in multimodal processing facilitate conceptual grounding of language. We categorize the information flow in multimodal processing with respect to cognitive models of human information processing and analyze different methods for combining multimodal representations. Based on this methodological inventory, we discuss the benefit of multimodal grounding for a variety of language processing tasks and the challenges that arise. We particularly focus on multimodal grounding of verbs which play a crucial role for the compositional power of language.

Title and Abstract in German

Multimodale konzeptuelle Verankerung für die automatische Sprachverarbeitung

Dieser Überblick erörtert, wie aktuelle Entwicklungen in der automatischen Verarbeitung multimodaler Inhalte die konzeptuelle Verankerung sprachlicher Inhalte erleichtern können. Die automatischen Methoden zur Verarbeitung multimodaler Inhalte werden zunächst hinsichtlich der zugrundeliegenden kognitiven Modelle menschlicher Informationsverarbeitung kategorisiert. Daraus ergeben sich verschiedene Methoden um Repräsentationen unterschiedlicher Modalitäten miteinander zu kombinieren. Ausgehend von diesen methodischen Grundlagen wird diskutiert, wie verschiedene Forschungsprobleme in der automatischen Sprachverarbeitung von multimodaler Verankerung profitieren können und welche Herausforderungen sich dabei ergeben. Ein besonderer Schwerpunkt wird dabei auf die multimodale konzeptuelle Verankerung von Verben gelegt, da diese eine wichtige kompositorische Funktion erfüllen.

1 Introduction

Natural languages are continually developing constructs that include numerous variations and irregularities. Modeling the subtleties of language in a formal, processable way has driven computational linguistics for decades. In recent years, distributional approaches have become the most widely accepted solution to model the associative character of word meaning (Harris, 1954; Collobert et al., 2011; Mikolov et al., 2013; Pennington et al., 2014). These approaches learn word representations in a high-dimensional vector space based on context patterns in large text collections. Machine learning researchers aim at reducing external knowledge to an absolute minimum and simply interpret language as a continuous stream of characters. From an engineering perspective, these data-driven approaches are highly attractive because they reduce the need of domain experts.

From a cognitive perspective, processing language in isolation without information on situational context seems to be an overly artificial setup. Human acquisition of semantic representations does not occur based on pure language input. The term conceptual grounding refers to the idea that language is

[◊] The work by Lisa Beinborn has been carried out during her affiliation with Ubiquitous Knowledge Processing Lab (UKP) and Research Training Group AIPHES, Technische Universität Darmstadt.

^{*} The first and the second authors contributed equally to this work.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>




Figure 1: Information flow in multimodal tasks. Blue and yellow shapes refer to modality A and B .

grounded in perceptual experience and sensorimotor interactions with the environment (Barsalou, 2008). In its strictest interpretation, this embodied perspective implies that language production and language comprehension involve perceptual and motor simulations of the described situation (Goldman, 2006). An impressive number of recent neuroimaging studies indicate that processing a word activates areas in the brain that correspond to the associated sensory modality of its semantic category: action-related words like *kick* trigger activity in the motor cortex and object-related words like *cup* activate visual areas (Pulvermüller et al., 2005; Garagnani and Pulvermüller, 2016). While it remains a controversial question to what extent conceptual representations are actually shared across modalities (Louwerse, 2011; Leshinskaya and Caramazza, 2016), it has been widely accepted that conceptual and sensorimotor representations are tightly coupled and interact with each other. Cognitively plausible language processing should thus interpret language as one modality within a multimodal environment.

This survey discusses how recent developments in multimodal processing facilitate conceptual grounding of language. It intends to provide a bridge between the field of multimodal machine learning (Baltrušaitis et al., 2017) and the cognitive theories for grounding distributional semantics (Baroni, 2016). As this is a wide interdisciplinary topic which influences many subfields, we focus on multimodal grounding for computational linguistics. For a better understanding of the interaction between modalities, we categorize multimodal tasks according to the information flow between the modalities. In a second step, we analyze different methods for combining multimodal information. Based on this methodological inventory, we discuss the benefit of multimodal grounding for a variety of language processing tasks. In multimodal processing, grounding is usually limited to concrete concepts leading to a reduction of referential ambiguity. We provide a detailed analysis of the challenges that arise when multimodal grounding is extended to open-domain language settings. We particularly focus on multimodal grounding of verbs which is essential for the interpretation of sequences and the identification of relations between concepts.

2 Multimodal processing models

The term “multimodal” has been used in a broad range of different interpretations even in the computational linguistics literature alone. In the common interpretation, modalities refer to sensory input such as audio, vision, touch, smell, and taste. Other definitions stretch over different communicative channels such as language and gesture, or simply different “modes” of the same modality (e.g., day and night pictures). In this section, we analyze the flow of multimodal information in different multimodal tasks exemplified by three modalities: natural language encoded as texts, visual signals encoded as images or videos, and audio signals encoded as sound files. For an overview of the challenges and machine learning methods associated with each task, the interested reader is referred to Baltrušaitis et al. (2017). We propose a classification of multimodal tasks with respect to the information flow between modalities into cross-modal transfer, cross-modal interpretation, and joint multimodal processing. From a historical perspective, progress in multimodal processing can be aligned with cognitive theories of multimodal organization in the human brain.

2.1 Cross-modal transfer

In the 1980s and 90s, cognitive processing theories were heavily influenced by the theory of the modularity of mind (Fodor, 1985). It assumes that processing occurs in domain-specific encapsulated modules

that do not interact with each other. Earlier approaches to multimodal engineering have taken a similar modular perspective. They model the information flow in each modality separately and the final outcome is then transferred or aligned to another modality. We group tasks in which one modality serves as the interface to query or represent the content from another modality under the category of cross-modal transfer, see Figure 1a.

A classical example for cross-modal transfer are search and retrieval tasks. The human user provides a natural language description to query an artifact (i.e., an image, video, or audio file) from a database (Atrey et al., 2010). The cross-modal alignment between the query and the artifact requires query expansion and disambiguation for referential indexing. In speech-related transfer tasks, textual content needs to be mapped to audio samples. Speech synthesis transforms text into artificially generated phonemes for users who cannot read (Zen et al., 2009). The reverse task of transcribing audio and video content is addressed by approaches for speech recognition (Juang and Rabiner, 2005) and subtitle generation (Daelemans et al., 2004). For lipreading tasks, mute video input of people speaking is transformed into text representing their utterances (Ngiam et al., 2011).

In these cross-modal transfer tasks, synchronous processing of the input in one modality is not directly influenced by information from the output modality. The main challenge lies in finding appropriate translations or alignments from one modality to the other. Information from the output modality is mainly used for reranking of input hypotheses. This view corresponds to mental models of a language hub in the brain that does not directly incorporate perceptual information (Chomsky, 1986).

2.2 Cross-modal interpretation

In order to explain how humans can select relevant information from perceptual input, the concept of attention has become very popular. Bridewell and Bello (2016) argue that attention serves as "a bottleneck for information flow in a cognitive system" that redirects mental resources. In multimodal processing, the concept of attention as a mediator between modalities is relevant for cross-modal interpretation. For these tasks, the goal is to obtain a compressed and structured intermediate representation of the input to generate a useful interpretation in the target modality. Attention mechanisms (Bahdanau et al., 2014) are used for the identification of relevant information, see Figure 1b.

A textual interpretation of a visually presented scene is generated in image captioning (Xu et al., 2015) and sketch recognition (Li et al., 2015). The goal is to identify relevant elements, group individual elements to semantic concepts, identify relations between concepts, and express these relations in natural language. The output sequence is generated while paying attention to different salient areas in the image. To our knowledge, a bidirectional information flow that includes cues from the language generation module in the image recognition process has not yet been implemented. However, semantic information could help to better direct the attention for image recognition, e.g., the generation of a verb like *eat* could constrain the visual recognition to edible objects as filler roles. Yatskar et al. (2016) propose the task of situation recognition to approximate this problem.

Complementary approaches attempt to generate visual representations to summarize documents and present the most relevant information in an intuitive way (Kucher and Kerren, 2015). The most popular approach are so-called word clouds which are frequency-based visualizations for topic modeling (Bateman et al., 2008). More recent approaches include semantic relations between words for a more conceptual-driven interpretation (Xu et al., 2016). Concept maps highlight structural relations between concepts in a graph-based visualization (Zubrinic et al., 2012). One key challenge for cross-modal interpretation tasks lies in the evaluation of the output because interpretations are by definition subjective and divergent solutions can be equally valid. Accumulations over various human ratings are currently considered to be better quality approximations than any automatic metrics (Vedantam et al., 2015).

2.3 Joint multimodal processing

Due to a wave of experimental findings that support the cognitive theory of embodied processing, the separating aspects between different modalities have become blurred (Pulvermüller et al., 2005). A similar development can be observed in multimodal machine learning. Tasks which explicitly require the combination of knowledge from different modalities gave rise to joint multimodal processing (Figure 1c). For

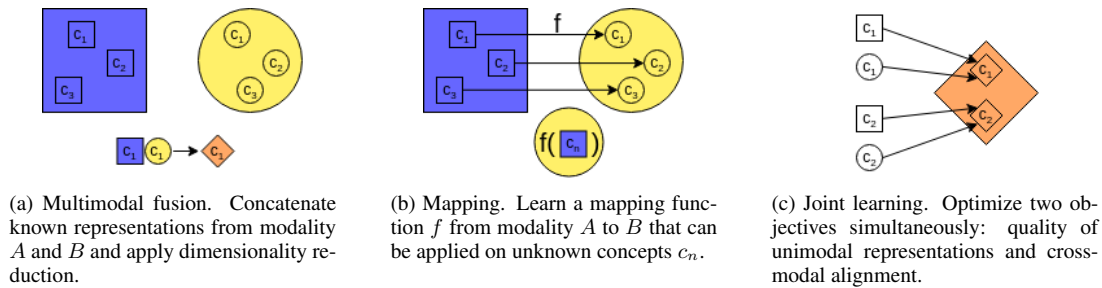


Figure 2: Methods for learning multimodal representations. Blue and yellow shapes indicate the representation space of modality A and B .

emotion recognition (Morency et al., 2011) or persuasiveness prediction (Santos et al., 2016), the actual content of an utterance and paraverbal cues (e.g., pitch, facial expression) need to be jointly evaluated. An ironic tone of voice might reverse the conceptual interpretation of the language content.

Recent work from the vision community goes one step further and tackles tasks that imperatively require an interactive flow of information. In visual question answering, a human user can ask questions about an image that the system should answer (Malinowski et al., 2015). This requires several steps: understanding the question, determining the salient elements in the image, interpreting the image with respect to the question, and generating a coherent natural language answer that matches the question. For this task, exchange of information between the modalities is crucial. In an overview, Wu et al. (2017) compare 29 approaches to visual question answering. 23 of these approaches use a joint representation of textual and visual information. The remaining 6 approaches organize the exchange either through a coordinated network architecture or through shared external knowledge bases. Novel interactive approaches make it possible to directly modulate the information flow in one modality by input from another modality (Vries et al., 2017) or by human feedback (Ling and Fidler, 2017).

The main challenge for joint processing lies in efficiently combining information from the modalities, so that redundant information is integrated without losing complementary cues. In human language understanding, this process seems to be performed in an effortless and highly accurate manner (Crocker et al., 2010). However, the underlying mechanisms of multimodal representations remain poorly understood. In Section 3, we discuss different methods for obtaining joint representations computationally.

3 Multimodal representation learning

Multimodal representations combine information from separate modalities. We discuss methods for representing known concepts, projecting information to represent unknown concepts, and for combining concept representations into compositional representations for sequences.

3.1 Concept representations

Even in unimodal tasks, researchers experiment with many different variations of representing concepts and their relations. Earlier work on multimodal representations used human-elicited visual features (Silberer and Lapata, 2012; Roller and Schulte Im Walde, 2013). Conveniently, integrating knowledge from different modalities has been facilitated due to the now common low-level representations of the input (also known as embeddings). Images are represented as groups of pixels, videos as series of image frames, audio data as windows of waveform samples, and language as sequences of distributional word representations. These values are then fed into a neural network that learns to compress and normalize the representation such that it better generalizes across input samples (Kielbaso and Bottou, 2014). A concept representation is usually obtained by averaging over many different samples for the concept (e.g., the concept *bird* is represented by averaging over the representation for n images showing a bird). The representations are expressed as high-dimensional matrices which can be projected into a common space. This approach facilitates a joint information flow between different modalities and has contributed to the growing success of multimodal processing.

Fusion The most intuitive approach is multimodal fusion (Figure 2a). Assuming that a unimodal vector representation v for the concept c and the modalities m_1 and m_2 exists, the multimodal representation v_{mm} consists of the concatenation $\hat{\cdot}$ of the two vectors weighted by a tunable parameter α :

$$v_{mm}(c) = \alpha \cdot v_{m_1}(c) \hat{\cdot} (1 - \alpha) \cdot v_{m_2}(c).$$

The concatenation occurs directly on the concept level and is thus called feature-level fusion or early fusion (Leong and Mihalcea, 2011; Bruni et al., 2011). In the case of pure concatenation, the unimodal representations reside in separate conceptual spaces. The concatenated representation for *cat* could give us the information that *cat* is visually similar to *panther* and textually similar to *dog*, but it is not possible to determine cross-modal similarity. In order to smooth the concatenated representations while maintaining multimodal correlations, dimensionality reduction techniques such as singular value decomposition (Bruni et al., 2014) or canonical correlation analysis (Silberer and Lapata, 2012) have been applied.

3.2 Projection

In practice, concepts that have a representation in one modality are not necessarily covered by representations in another modality. The projection of unseen concepts is known as zero shot learning. It can either be performed on a mapped or a joint representational space.

Mapping To overcome the lack of representations for one modality, several researchers proposed to map representations from one modality to the other (Figure 2b). The idea is to learn a mapping function f from m_1 to m_2 that maximizes the similarity between a known representation of c in m_2 and its projection from the representation in m_1 : $c_{m_2} \sim f(c_{m_1})$.

The choice of the similarity and the loss measures for learning the mapping function vary. A max-margin optimization which maximizes the similarity between true pairs of concept representations (c_{m_1}, c_{m_2}) and minimizes the similarity for pairs with random target representations $(c_{m_1}, random_{m_2})$ has been shown to be a good choice for image labeling (Frome et al., 2013). In this task, the mapping approach is applied in the image-to-text direction to classify unknown objects in images based on their semantic similarity to known objects (Socher et al., 2013). Lazaridou et al. (2014) and Collell et al. (2017) proceed in the reverse text-to-image direction to ground words in the visual world. Similar propagation approaches had already been examined by Johns and Jones (2012) and Hill and Korhonen (2014), but they used human-elicited perceptual features from the McRae dataset (McRae et al., 2005) instead of automatically derived image representations.

Joint learning The mapping approaches assume a directed transformation from one modality to the other. Joint estimation approaches aim to learn shared representations instead (Figure 2b). An approach inspired by topic modeling interprets aligned data as a multimodal document and uses Latent Dirichlet Allocation to derive multimodal topics (Andrews et al., 2009; Feng and Lapata, 2010; Silberer and Lapata, 2012; Roller and Schulte Im Walde, 2013). Unfortunately, this approach cannot be easily used for zero shot learning. Lazaridou et al. (2015) enrich the skip-gram model by Mikolov et al. (2013) with visual features. Their model optimizes two constraints: the representation of concepts c with respect to their textual contexts (unsupervised skip-gram objective in m_1) and the similarity between word representations and their visual counterparts (supervised max-margin objective for (c_{m_1}, c_{m_2})). In their approach, the visual representations remain fixed, but the textual representations are learned from scratch. Silberer and Lapata (2014) go one step further and use stacked multimodal autoencoders to simultaneously learn good representations for each modality (unsupervised reconstruction objective for m_1 and m_2) and their optimal multimodal combination (supervised classification objective for (c_{m_1}, c_{m_2})). Both approaches implicitly also learn a mapping between the two modalities and can be adjusted to induce a directional projection for zero shot learning. Joint learning of multimodal representations is very popular in the vision community (Karpathy et al., 2014; Srivastava and Salakhutdinov, 2012; Ngiam et al., 2011).

3.3 Compositional representations

For tasks that require representing longer sequences, a naïve approach is sequence-level fusion. In this setting, the unimodal sequence representation is obtained by performing an arithmetic operation (e.g.,

average, max) over the concept representations for each word in the sequence. Multimodal fusion is then performed on this averaged representation (Glavaš et al., 2017; Bruni et al., 2014). Shutova et al. (2016) work with short phrases consisting of two words and directly learn phrase representations. Missing concept representations in one modality can be obtained by mapping functions (Botschen et al., 2018).

For image captioning approaches, representations for a pair of an image and the corresponding caption are learned jointly (Kiros et al., 2014; Socher et al., 2014). Pre-trained unimodal representations are fed into a neural network which is trained with the max-margin objective to distinguish between true and false captions for an image. The multimodal sequence representation can be obtained from the last hidden layer of the network. The introduction of attention variables can function as a mediator between the visual and the textual modality (see Section 2.2). For a more detailed overview of multimodal sequence representations in the vision community, the interested reader is referred to Wu et al. (2017). The approaches for compositional representations have focused on enriching noun and adjective meaning multimodally. The multimodal interpretation of verbs as an integral part of compositional sequences has not yet been thoroughly examined.

4 Multimodal grounding for language processing

The progress in joint multimodal processing and the increasing availability of multimodal datasets and representations open up new possibilities for grounded approaches to language processing. We review recent works for grounding concepts, grounding phrases, and grounding interaction. The challenges that arise from these efforts are discussed in Section 5.

4.1 Grounding concepts

Multimodal concept representations are motivated by the idea that semantic relations between words are grounded in perception. Being able to assess semantic relations between concepts is an important prerequisite for modeling generalization capabilities in language processing. The combination of the textual and the visual modality has received most attention for conceptual grounding, but perceptual information from the auditory and the olfactory channel have also been used for dedicated tasks (Kiela et al., 2015; Kiela and Clark, 2017). In order to provide a more concrete discussion, we focus on the combination of textual and visual cues for the remainder of the survey.

The quality of concept representations is commonly evaluated by their ability to model semantic properties. Different approaches to learning conceptual models are compared by their performance on similarity datasets, e.g., *WordSim353* (Finkelstein et al., 2002), *SimLex-999* (Hill et al., 2015), *MEN* (Bruni et al., 2012), *SemSim*, and *VisSim* (Silberer and Lapata, 2014)). These datasets contain pairs of words that have been annotated with similarity scores for the two concepts. Several evaluations of semantic models have shown that multimodal concept representations outperform unimodal ones (Feng and Lapata, 2010; Silberer and Lapata, 2012; Bruni et al., 2014; Kiela et al., 2014). Kiela et al. (2016) perform a comparison of different image sources and architectures and their ability to model semantic similarity. Despite the advantages of multimodal models in capturing semantic relations, it remains an open question whether they contribute to a cognitively more plausible approximation of human conceptual grounding. Bulat et al. (2017b) and Anderson et al. (2017) conduct experiments to label brain activity scans by human subjects with the corresponding concepts that elicited the brain activity. They compare different distributional semantic models and obtain mixed results. Bulat et al. (2017b) find that visual information is beneficial for modeling concrete words, whereas Anderson et al. (2017) conclude that textual models sufficiently integrate visual properties. Further interdisciplinary research involving computer science, neuroscience, and psycholinguistics is required to obtain a deeper understanding of cognitively plausible language processing (Embick and Poeppel, 2015).

4.2 Grounding phrases

Most experiments for conceptual grounding indicate that providing a multimodal representation for abstract concepts is significantly more challenging due to the lack of perceptual patterns associated with abstract words (Hill et al., 2014). For grounding phrases, the meaning for concrete and abstract concepts

need to be combined (see Section 3.3). Bruni et al. (2012) examine the compositional meaning of color adjectives and find that multimodal representations are superior in modeling color. However, they fail to distinguish between literal and non-literal usage of color adjectives (e.g., *green cup* vs *green future*).

Vivid imagery and synaesthetic associations play an important role in the interpretation of figurative language. In their influential theory of metaphor, Lakoff and Johnson (1980) argue that abstract concepts can be grounded metaphorically in embodied and situated knowledge. They assume that metaphors can be understood as a mapping from a concrete source domain to a more abstract target domain. For example, *time* is often viewed as a *stream* that *flows* in a direction. Turney et al. (2011) operationalize this theoretical account by identifying metaphoric phrases based on the discrepancy in concreteness of source and target term. Shutova et al. (2016) and Bulat et al. (2017a) build on their approach and use multimodal models for identifying metaphoric word usage in adjective-noun combinations. They show that words used in a metaphorical combination (*dry wit*) exhibit less similarity than words in non-metaphorical phrases (*dry skin*). We strongly believe that progress in multimodal compositional grounding will pave the way for a more holistic understanding of figurative language processing. As a prerequisite, multimodal grounding needs to be examined beyond the representation of concrete objects (see Section 5.2). Representing verbs, compositional phrases, and even full sentences by means of multimodal information has not yet been sufficiently examined.

4.3 Grounding interaction

The origins of grounding theories were initiated to account for situational language use and interaction. We distinguish two main scenarios for interactive language use: language learning and situational grounding of action descriptions.

Grounded language learning Language learning is deeply rooted in social interaction and initially emerges with respect to a concrete referential context (Tomasello, 2010). Children acquire language in interaction with their parents and foreign language learning proceeds much faster in an environment that forces the learner to interact in the foreign language (Nation, 1990). Usage-based approaches to language learning that account for the frequency and the quality of the language stimulus have a long tradition (Dale and Chall, 1948). Brysbaert and New (2009) have shown that frequency information grounded in auditory and visual communicative cues can better model human processing effects than frequency information extracted from purely textual corpora. Lazaridou et al. (2016) show that a multimodal distributional approach better approximates word learning from interactive child-directed input than unimodal approaches. The same model can also convincingly simulate word meaning induction by adults (Lazaridou et al., 2017a). Psycholinguistic research indicates that conceptual mapping modulated by visual properties is not only relevant for first language acquisition, but is also used as a means to establish cross-lingual links in foreign language learning (Beinborn et al., 2014). Bergsma and Van Durme (2011) and Vulić et al. (2016) take advantage of this observation and use multimodal representations to induce multilingual representations.

Grounding sequences in actions Situational grounding of action descriptions requires the representation of sequences and their compositional interpretation. Regneri et al. (2013) build a corpus that grounds descriptions of actions in videos showing these actions. For the interpretation of sequences, evaluating verbs and their arguments plays a fundamental role. Yatskar et al. (2016) developed the *imSitu* dataset which consists of images depicting verbs and annotations which link the verb arguments to visual referents. This dataset can be used for the multimodal task of situation recognition (Mallya and Lazebnik, 2017; Zellers and Choi, 2017), and it serves as a multimodal resource for verb processing. Grounding verbs is particularly challenging because of the variety of their possible visual instantiations. For example, an image of an adult drinking beer has very little in common with a zebra drinking water.

Multimodal interpretation of sequences is highly relevant for robotics research (Chaplot et al., 2017). Mordatch and Abbeel (2017) examine the emergence of compositionality in grounded multi-agent communication. The language learned by artificial agents is not necessarily interpretable by humans. Lazaridou et al. (2017b) show that agents which develop their own language for representing concepts that are grounded in images infer similar taxonomic relations as humans. Their work suggests that the learned

concepts can even be mapped back into natural language. Agent-agent communication has already been examined in the talking head experiments, in which two agents learn to discriminate between objects and develop their own language of referring expressions (Steels and Vogt, 1997; Steels, 2002). Hermann et al. (2017) and Heinrich and Wermter (2018) explicitly focus on human-robot interaction and train their agent to associate natural language descriptions of actions with perceptual input from its sensors.

For experiments on grounded language understanding, the situational environment is usually artificially restricted to a very small domain. This confined setting facilitates the analysis of compositional expressions and their referential interpretation as complex object descriptions or action sequences. In open-domain language understanding, semantic disambiguation is even more challenging. Approaches using multimodal information for the disambiguation of concepts (Xie et al., 2017), named entities (Moon et al., 2018), and sentences (Botschen et al., 2018; Shutova et al., 2017) show promising tendencies, but the underlying compositional principles are not yet understood.

5 Challenges for grounded language processing

Multimodal grounding of language has been a longstanding goal of language researchers. The discussion has gained new momentum due to the recent developments in learning distributed multimodal representations. Most evaluations indicate that multimodal representations are beneficial for a variety of tasks, but explanatory analyses of this effect are still in a developing phase. In this section, we discuss open challenges that arise from existing work. For future work, we propose to examine multimodal grounding beyond concrete nouns and adjectives. In order to do this, larger multimodal datasets encompassing a wider range of word classes need to be build. These datasets would enable us to analyze compositional representations in more detail and to develop more elaborate models of selective multimodal grounding.

5.1 Combining complementary information

Different modalities contribute qualitatively different conceptual information. Bruni et al. (2014) argue that highly relevant visual properties are often not represented by linguistic models because they are too obvious to be explicitly mentioned in text (e.g., birds have wings, violins are brown). Textual models, on the other hand, provide a better intuition of taxonomic and functional relations between concepts which cannot easily be derived from images (Collell and Moens, 2016). Ideally, multimodal representations should integrate the complementary perspectives for a more coherent grounded interpretation of language. From a more skeptical perspective, Louwerse (2011) states that perceptual information is already sufficiently encoded in textual cues. In this case, the superior performance of multimodal representations that has been established by several researchers would mainly be due to a more robust representation of highly redundant information. The results by Silberer and Lapata (2014) and Hill et al. (2014) support the intuitive assumption that textual representations better model textual similarity and visual representations better model visual similarity. As the multimodal models improve on both similarity tasks, the integration of complementary information seems to be successful. Interestingly, both evaluations show that simply concatenating the two modalities already yields a quite competitive model. The reported findings have been evaluated on models working with human-annotated perceptual features. These features inherently represent taxonomic knowledge that cannot be directly inferred from visual input. It remains an open question to which extent automatically derived image representations can contribute complementary information when combined with textual representations. Most multimodal research to date focuses on the representation of individual concepts (nouns) and their properties (adjectives). The benefit of multimodal representations for language tasks going beyond concept similarity needs to be examined in more detail from both, engineering and theoretical perspectives.

Multimodal grounding of verbs Verbs play a fundamental role for expressing relations between concepts and their situational functionality (Hartshorne et al., 2014). The dynamic nature of verbs poses a challenge for multimodal grounding. To our best knowledge, only Hill et al. (2014) and Collell et al. (2017) consider verbs in their evaluation. They report that results for verbs are significantly worse, but




Figure 3: Illustration for the quality of verb representations indicated as Spearman correlation between the cosine similarity of verbs and their corresponding similarity rating in the *SimVerb* dataset.

do not elaborate on this finding. We present first steps towards an investigation of verb grounding.¹ Figure 3 illustrate the quality of verb representations in the most common publicly available approaches for multimodal representations. In line with previous work, the quality of the representations is evaluated as the Spearman correlation between the cosine similarity of two verbs and their corresponding similarity rating in the *SimVerb* dataset (Gerz et al., 2016). We compare the quality of 3498 verb pairs² in textual Glove representations (Pennington et al., 2014) and two visual datasets: the Google dataset that performed best in Kiela et al. (2016) and has the highest coverage for the verb pairs (493 pairs, 14%)³ and the *imSitu* dataset which has been intentionally designed for verb identification (354 pairs, 10%). The results show that models which include visual information outperform purely textual representations for known concepts. However, the general quality of the verb representations is much lower than the quality reported for nouns. As a consequence, the mapping to unseen verb pairs yields unsatisfactory results for the full *SimVerb* dataset. Our encouraging results for the *imSitu* dataset indicate that it is recommended to directly obtain visual representations for verbs instead of projecting the meaning. Building larger multimodal datasets with a focus on verbs seems to be a promising strand of research for future work.

5.2 Imageability of abstract words

Conceptual grounding of language can be intuitively performed for concrete words that have direct referents in sensory experience. Bruni et al. (2014) and Hill and Korhonen (2014) show that multimodal representations are beneficial for evaluating concrete words, but have little to no impact on the evaluation of abstract words. Projecting unseen concepts into the representation space based on their relations to seen concepts in another modality provides an elegant method for zero shot learning, but it is questionable whether multimodal relations between concrete concepts are sufficient to infer relations between abstract concepts. Lazaridou et al. (2015) analyze projected abstract words by extracting the nearest visual neighbor from their multimodal representation. The neighbors were paired with random images and human raters judged how well each image represents the word. The hypothesis that concrete objects are more likely to be captured adequately by multimodal representations was confirmed. However, they also provide illustrating examples which represent abstract words like *together* or *theory* surprisingly well.

Embodiment of verbs From a multimodal perspective, verbs can be categorized according to their degree of embodiment. This measure indicates to which extent verb meanings involve bodily experience (Sidhu et al., 2014). We obtain embodiment ratings for 1163 pairs.⁴ The class *high embodiment* contains pairs like *fall-dive* in which the embodiment of both verbs can be found in the highest quartile (135 pairs), *low embodiment* contains pairs with embodiment ratings in the lowest quartile (81 pairs) like *know-decide*.⁵ Coherent with previous work on concrete and abstract nouns (Hill et al., 2014), it can be

¹The pre-trained embeddings and the script to reproduce our results are available for research purposes: <https://github.com/UKPLab/coling18-multimodalSurvey>.

²Two pairs had to be excluded because *misspend* was not covered in the textual representations.

³The coverage in *WN9-IMG* (Xie et al., 2017) and the dataset used by Collell et al. (2017) is lower.

⁴<https://psychology.ucalgary.ca/languageprocessing/node/22>. We only include a pair, if an embodiment rating is available for both verbs.

⁵It should be noted that not all instances of the two classes are covered by the visual representations. The small number of instances might have an impact on the correlation values.

seen that visual representations better capture the similarity of verbs with a high level of embodiment. The mapped representations maintain this sense of embodiment, whereas the concatenated and fused representations better capture the similarity for verbs referring to more conceptual actions. This finding indicates that multimodal information is not equally beneficial for all words.

5.3 Selective multimodal grounding

The expressive power of language is essentially due to its combinatorial capabilities. Understanding how to combine concept representations to represent multi-word expressions or even full sentences has been a question of ongoing research in computational linguistics for decades. The inclusion of additional modalities further complicates this debate. Glavaš et al. (2017) and Botschen et al. (2018) obtain multimodal sentence representations by averaging over the multimodal representations for each word. They report improved results for the tasks of sentence similarity and frame identification. Our comparison above indicates that this superior performance is mainly due to a better representation of concepts. This raises the assumption that multimodal grounding should only be performed on selected words. Glavaš et al. (2017) propose to condition the inclusion of visual information on the prototypicality of a concept as measured by the image dispersion score (Kiela et al., 2014). This measure calculates the average pairwise cosine distance in a set of images to model the assumption that an image collection for an abstract concept like *happiness* is more diverse than for a concrete concept like *ladder*. Lazaridou et al. (2015) and Hessel et al. (2018) propose alternative concreteness measures based on the same idea. Unfortunately, these measures are highly dependent on the image retrieval algorithm which might be optimized towards obtaining a diverse range of images. Nevertheless, we assume that selective multimodal grounding constitutes a more plausible approach to sentence processing. Some functional words (e.g., locative expressions) might benefit from multimodal information, but it currently remains unclear how words with syntactic functions (e.g., coordinating expressions) should be represented visually.

6 Conclusion

We analyzed how multimodal processing has developed from transfer between encapsulated modalities to interactive processing over joint multimodal representations. These developments contribute to new avenues of research for grounded language processing. We strongly believe that the integration of multimodal information will improve our understanding of conceptual semantic models, figurative language processing, language learning, and situated interaction. Image datasets are often optimized towards providing a variety of visual instantiations. Developing algorithms for determining more prototypical visual representations could contribute to better grounding of verbs and might also serve as a criterion for selective multimodal grounding.

Acknowledgements

This work has been supported by the DFG-funded research training group “Adaptive Preparation of Information from Heterogeneous Sources” (AIPHES, GRK 1994/1) at Technische Universität Darmstadt. We thank Faraz Saeedan for his assistance with the computation of the visual embeddings for the imSitu images. We thank the anonymous reviewers for their insightful comments.

References

- Andrew J Anderson, Douwe Kiela, Stephen Clark, and Massimo Poesio. 2017. Visually grounded and textual semantic models differentially decode brain activity associated with concrete and abstract nouns. *Transactions of the Association for Computational Linguistics (TACL)*, 5(1):17–30.
- Mark Andrews, Gabriella Vigliocco, and David Vinson. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological review*, 116(3):463–498.
- Pradeep K Atrey, M Anwar Hossain, Abdulmotaleb El Saddik, and Mohan S Kankanhalli. 2010. Multimodal fusion for multimedia analysis: a survey. *Multimedia systems*, 16(6):345–379.

- Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2017. Multimodal machine learning: A survey and taxonomy. *arXiv preprint arXiv:1705.09406*.
- Marco Baroni. 2016. Grounding distributional semantics in the visual world. *Language and Linguistics Compass*, 10(1):3–13.
- Lawrence W Barsalou. 2008. Grounded cognition. *Annual Review of Psychology*, 59:617–645.
- Scott Bateman, Carl Gutwin, and Miguel Nacenta. 2008. Seeing things in the clouds: the effect of visual features on tag cloud selections. In *Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*, pages 193–202. ACM.
- Lisa Beinborn, Torsten Zesch, and Iryna Gurevych. 2014. Readability for foreign language learning: The importance of cognates. *International Journal of Applied Linguistics*, 165(2):136–162.
- Shane Bergsma and Benjamin Van Durme. 2011. Learning bilingual lexicons using the visual similarity of labeled web images. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 22, pages 1764–1769.
- Teresa Botschen, Iryna Gurevych, Jan-Christoph Klie, Hatem Mousselly Sergieh, and Stefan Roth. 2018. Multimodal frame identification with multilingual evaluation. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 1481–1491. Association for Computational Linguistics.
- Will Bridewell and Paul F Bello. 2016. A theory of attention for cognitive systems. *Advances in Cognitive Systems*, 4:1–16.
- Elia Bruni, Giang Binh Tran, and Marco Baroni. 2011. Distributional semantics from text and images. In *Proceedings of the GEMS workshop on geometrical models of natural language semantics*, pages 22–32. Association for Computational Linguistics.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proceedings of the 50th Annual Conference of the Association for Computational Linguistics (ACL)*, pages 136–145. Association for Computational Linguistics.
- Elia Bruni, Nam-Khanh Tram, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Marc Brysbaert and Boris New. 2009. Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior research methods*, 41(4):977–990.
- Luana Bulat, Stephen Clark, and Ekaterina Shutova. 2017a. Modelling metaphor with attribute-based semantics. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 523–528. Association for Computational Linguistics.
- Luana Bulat, Stephen Clark, and Ekaterina Shutova. 2017b. Speaking, seeing, understanding: Correlating semantic models with conceptual representation in the brain. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1081–1091.
- Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. 2017. Gated-attention architectures for task-oriented language grounding. *arXiv preprint arXiv:1706.07230*.
- Noam Chomsky. 1986. *Knowledge of language: Its nature, origin, and use*. Greenwood Publishing Group.
- Guillem Collell and Marie-Francine Moens. 2016. Is an image worth more than a thousand words? On the fine-grain semantic differences between visual and linguistic representations. In *Proceedings of the 26th International Conference on Computational Linguistics (COLING)*, pages 2807–2817.
- Guillem Collell, Ted Zhang, and Marie-Francine Moens. 2017. Imagined visual representations as multimodal embeddings. In *Proceedings of the Thirty-First Conference on Artificial Intelligence (AAAI)*, pages 4378–4384.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537.

- Matthew W Crocker, Pia Knoeferle, and Marshall R Mayberry. 2010. Situated sentence processing: The coordinated interplay account and a neurobehavioral model. *Brain and language*, 112(3):189–201.
- Walter Daelemans, Anja Höthker, and Erik F Tjong Kim Sang. 2004. Automatic sentence simplification for subtitling in dutch and english. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC)*, pages 1045–1048.
- Edgar Dale and Jeanne S Chall. 1948. A formula for predicting readability: Instructions. *Educational Research Bulletin*, 27(2):37–54.
- David Embick and David Poeppel. 2015. Towards a computational(ist) neurobiology of language: correlational, integrated and explanatory neurolinguistics. *Language, Cognition and Neuroscience*, 30(4):357–366.
- Yansong Feng and Mirella Lapata. 2010. Visual information in semantic representation. In *Proceedings of the 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 91–99, Stroudsburg, USA. Association for Computational Linguistics.
- Lev Finkelstein, Evgeniy Gabilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2002. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM.
- Jerry A Fodor. 1985. Precis of the modularity of mind. *Behavioral and brain sciences*, 8(1):1–5.
- Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. 2013. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2121–2129.
- Max Garagnani and Friedemann Pulvermüller. 2016. Conceptual grounding of language in action and perception: a neurocomputational model of the emergence of category specificity and semantic hubs. *European Journal of Neuroscience*, 43(6):721–737.
- Daniela Gerz, Ivan Vulić, Felix Hill, Roi Reichart, and Anna Korhonen. 2016. SimVerb-3500: A Large-Scale Evaluation Set of Verb Similarity. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2173–2182.
- Goran Glavaš, Ivan Vulić, and Simone Paolo Ponzetto. 2017. If sentences could see: Investigating visual information for semantic textual similarity. In *Proceedings of the 12th International Conference on Computational Semantics (IWCS)*, Montpellier, France.
- Alvin I Goldman. 2006. *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University Press.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Joshua K Hartshorne, Claire Bonial, and Martha Palmer. 2014. The verbcorner project: Findings from phase 1 of crowd-sourcing a semantic decomposition of verbs. In *Proceedings of the 52nd Annual Conference of the Association for Computational Linguistics (ACL)*, pages 397–402. Association for Computational Linguistics.
- Stefan Heinrich and Stefan Wermter. 2018. Interactive natural language acquisition in a multi-modal recurrent neural architecture. *Connection Science*, 30(1):99–133.
- Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojtek Czarnecki, Max Jaderberg, Denis Teplyashin, et al. 2017. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*.
- Jack Hessel, David Mimno, and Lillian Lee. 2018. Quantifying the visual concreteness of words and topics in multimodal datasets. *arXiv preprint arXiv:1804.06786*.
- Felix Hill and Anna Korhonen. 2014. Learning abstract concept embeddings from multi-modal data: Since you probably can’t see what i mean. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 255–265.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2014. Multi-modal models for concrete and abstract concept meaning. *Transactions of the Association for Computational Linguistics (TACL)*, 2:285–296.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.

- Brendan T Johns and Michael N Jones. 2012. Perceptual inference through global lexical similarity. *Topics in Cognitive Science*, 4(1):103–120.
- Biing-Hwang Juang and Lawrence R Rabiner. 2005. Automatic speech recognition – a brief history of the technology development. *Georgia Institute of Technology. Atlanta Rutgers University and the University of California. Santa Barbara*, 1:67.
- Andrej Karpathy, Armand Joulin, and Li F Fei-Fei. 2014. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1889–1897.
- Douwe Kiela and Léon Bottou. 2014. Learning Image Embeddings using Convolutional Neural Networks for Improved Multi-Modal Semantics. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Douwe Kiela and Stephen Clark. 2017. Learning neural audio embeddings for grounding semantics in auditory perception. *Journal of Artificial Intelligence Research*, 60:1003–1030.
- Douwe Kiela, Felix Hill, Anna Korhonen, and Stephen Clark. 2014. Improving multi-modal representations using image dispersion: Why less is sometimes more. In *Proceedings of the 52nd Annual Conference of the Association for Computational Linguistics (ACL)*, volume 2, pages 835–841. Association for Computational Linguistics.
- Douwe Kiela, Luana Bulat, and Stephen Clark. 2015. Grounding semantics in olfactory perception. In *Proceedings of the 53rd Annual Conference of the Association for Computational Linguistics (ACL)*, pages 231–236, Beijing, China, July. Association for Computational Linguistics.
- Douwe Kiela, Anita Veró, and Stephen Christopher Clark. 2016. Comparing Data Sources and Architectures for Deep Visual Representation Learning in Semantics. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 447–456.
- Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Kostiantyn Kucher and Andreas Kerren. 2015. Text visualization techniques: Taxonomy, visual survey, and community insights. In *Pacific Visualization Symposium (PacificVis)*, pages 117–121. IEEE.
- George Lakoff and Mark Johnson. 1980. *Metaphors We Live By*. University of Chicago press.
- Angeliki Lazaridou, Elia Bruni, and Marco Baroni. 2014. Is this a wampimuk? cross-modal mapping between distributional semantics and the visual world. In *Proceedings of the 52nd Annual Conference of the Association for Computational Linguistics (ACL)*, volume 1, pages 1403–1414. Association for Computational Linguistics.
- Angeliki Lazaridou, Nghia The Pham, and Marco Baroni. 2015. Combining language and vision with a multi-modal skip-gram model. In *Proceedings of the 14th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 153–163. Association for Computational Linguistics.
- Angeliki Lazaridou, Grzegorz Chrupała, Raquel Fernández, and Marco Baroni. 2016. Multimodal semantic learning from child-directed input. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 387–392. Association for Computational Linguistics.
- Angeliki Lazaridou, Marco Marelli, and Marco Baroni. 2017a. Multimodal word meaning induction from minimal exposure to natural text. *Cognitive science*, 41(S4):677–705.
- Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017b. Multi-agent cooperation and the emergence of (natural) language. *arXiv preprint arXiv:1612.07182*.
- Chee Wee Leong and Rada Mihalcea. 2011. Going beyond text: A hybrid image-text approach for measuring word relatedness. In *Proceedings of the 5th International Joint Conference on Natural Language Processing (IJCNLP)*, pages 1403–1407.
- Anna Leshinskaya and Alfonso Caramazza. 2016. For a cognitive neuroscience of concepts: Moving beyond the grounding issue. *Psychonomic bulletin & review*, 23(4):991–1001.
- Yi Li, Timothy M. Hospedales, Yi-Zhe Song, and Shaogang Gong. 2015. Free-hand sketch recognition by multi-kernel feature learning. *Computer Vision and Image Understanding*, 137:1–11.

- Huan Ling and Sanja Fidler. 2017. Teaching machines to describe images via natural language feedback. In *Advances in Neural Information Processing Systems (NIPS)*.
- Max M. Louwerse. 2011. Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3(2):273–302.
- Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December.
- Arun Mallya and Svetlana Lazebnik. 2017. Recurrent models for situation recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 455–463.
- Ken McRae, George S Cree, Mark S Seidenberg, and Chris McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, 37(4):547–559.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *International Conference on Learning Representations Workshop (ICLR)*.
- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. 2018. Multimodal named entity recognition for short social media posts. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 852–860. Association for Computational Linguistics.
- Igor Mordatch and Pieter Abbeel. 2017. Emergence of grounded compositional language in multi-agent populations. *arXiv preprint arXiv:1703.04908*.
- Louis-Philippe Morency, Rada Mihalcea, and Payal Doshi. 2011. Towards multimodal sentiment analysis: Harvesting opinions from the web. In *Proceedings of the 13th international conference on multimodal interfaces*, pages 169–176. ACM.
- Ian Stephen Paul Nation. 1990. *Teaching and Learning Vocabulary*. Teaching Methods. Heinle & Heinle.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 689–696.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Friedemann Pulvermüller, Olaf Hauk, Vadim V. Nikulin, and Risto J. Ilmoniemi. 2005. Functional links between motor and language systems. *European Journal of Neuroscience*, 21(3):793–797.
- Michaela Regneri, Marcus Rohrbach, Dominikus Wetzels, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. Grounding action descriptions in videos. *Transactions of the Association of Computational Linguistics (TACL)*, 1:25–36.
- Stephen Roller and Sabine Schulte Im Walde. 2013. A multimodal lda model integrating textual, cognitive and visual modalities. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1146–1157.
- Pedro Bispo Santos, Lisa Beinborn, and Iryna Gurevych. 2016. A domain-agnostic approach for opinion prediction on speech. In Malvina Nissim, Viviana Patti, and Barbara Plank, editors, *Proceedings of the Workshop on Computational Modeling of Peoples Opinions, Personality, and Emotions in Social Media (PEOPLES)*, pages 163–172.
- Ekaterina Shutova, Douwe Kiela, and Jean Maillard. 2016. Black holes and white rabbits: Metaphor identification with visual features. In *Proceedings of the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 160–170. Association for Computational Linguistics.
- Ekaterina Shutova, Andreas Wundsam, and Helen Yannakoudakis. 2017. Semantic frames and visual scenes: Learning semantic role inventories from image and video descriptions. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM)*, pages 149–154.
- David M Sidhu, Rachel Kwan, Penny M Pexman, and Paul D Siakaluk. 2014. Effects of relative embodiment in lexical and semantic processing of verbs. *Acta psychologica*, 149:32–39.

- Carina Silberer and Mirella Lapata. 2012. Grounded models of semantic representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1423–1433. Association for Computational Linguistics.
- Carina Silberer and Mirella Lapata. 2014. Learning grounded meaning representations with autoencoders. In *Proceedings of the 52nd Annual Conference of the Association for Computational Linguistics (ACL)*, volume 1, pages 721–732. Association for Computational Linguistics.
- Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-shot learning through cross-modal transfer. In *Advances in Neural Information Processing Systems (NIPS)*, pages 935–943.
- Richard Socher, Andrej Karpathy, Quoc V Le, Christopher D Manning, and Andrew Y Ng. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics (TACL)*, 2(1):207–218.
- Nitish Srivastava and Ruslan Salakhutdinov. 2012. Learning representations for multimodal data with deep belief nets. In *International conference on machine learning workshop*, volume 79.
- Luc Steels and Paul Vogt. 1997. Grounding adaptive language games in robotic agents. In *Proceedings of the 4th european conference on artificial life*, volume 97.
- Luc Steels. 2002. Grounding symbols through evolutionary language games. In *Simulating the evolution of language*, pages 211–226. Springer.
- Michael Tomasello. 2010. *Origins of human communication*. MIT press.
- Peter D Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. Literal and metaphorical sense identification through concrete and abstract context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 680–690. Association for Computational Linguistics.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4566–4575.
- Harm De Vries, Florian Strub, J  r  mie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C. Courville. 2017. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6597–6607.
- Ivan Vuli  , Douwe Kiela, Stephen Clark, and Marie-Francine Moens. 2016. Multi-modal representations for improved bilingual lexicon learning. In *Proceedings of the 54th Annual Conference of the Association for Computational Linguistics (ACL)*, pages 188–194. Association for Computational Linguistics.
- Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2017. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 163:21–40.
- Ruobing Xie, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2017. Image-embodied knowledge representation learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3140–3146.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pages 2048–2057.
- Jin Xu, Yubo Tao, and Hai Lin. 2016. Semantic word cloud generation based on word embeddings. In *Pacific Visualization Symposium (PacificVis)*, pages 239–243. IEEE.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5534–5542.
- Rowan Zellers and Yejin Choi. 2017. Zero-shot activity recognition with verb attribute induction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 946–958.
- Heiga Zen, Keiichi Tokuda, and Alan W Black. 2009. Statistical parametric speech synthesis. *Speech Communication*, 51(11):1039–1064.
- Krunoslav Zubrinic, Damir Kalpic, and Mario Milicevic. 2012. The automatic creation of concept maps from documents written using morphologically rich languages. *Expert systems with applications*, 39(16):12709–12718.