

# Combining Language and Vision with a Multimodal Skip-gram Model

Angeliki Lazaridou Nghia The Pham Marco Baroni

Center for Mind/Brain Sciences

University of Trento

{angeliki.lazaridou|thenghia.pham|marco.baroni}@unitn.it

## Abstract

We extend the SKIP-GRAM model of Mikolov et al. (2013a) by taking visual information into account. Like SKIP-GRAM, our *multimodal* models (MMSKIP-GRAM) build vector-based word representations by learning to predict linguistic contexts in text corpora. However, for a restricted set of words, the models are also exposed to visual representations of the objects they denote (extracted from natural images), and must predict linguistic and visual features jointly. The MMSKIP-GRAM models achieve good performance on a variety of semantic benchmarks. Moreover, since they propagate visual information to all words, we use them to improve image labeling and retrieval in the zero-shot setup, where the test concepts are never seen during model training. Finally, the MMSKIP-GRAM models discover intriguing visual properties of abstract words, paving the way to realistic implementations of embodied theories of meaning.

## 1 Introduction

Distributional semantic models (DSMs) derive vector-based representations of meaning from patterns of word co-occurrence in corpora. DSMs have been very effectively applied to a variety of semantic tasks (Clark, 2015; Mikolov et al., 2013b; Turney and Pantel, 2010). However, compared to human semantic knowledge, these purely textual models, just like traditional symbolic AI systems (Harnad, 1990; Searle, 1984), are severely impoverished, suffering of *lack of grounding* in extra-linguistic modalities (Glenberg and Robertson, 2000). This observa-

tion has led to the development of *multimodal distributional semantic models* (MDSMs) (Bruni et al., 2014; Feng and Lapata, 2010; Silberer and Lapata, 2014), that enrich linguistic vectors with perceptual information, most often in the form of visual features automatically induced from image collections.

MDSMs outperform state-of-the-art text-based approaches, not only in tasks that directly require access to visual knowledge (Bruni et al., 2012), but also on general semantic benchmarks (Bruni et al., 2014; Silberer and Lapata, 2014). However, current MDSMs still have a number of drawbacks. First, they are generally constructed by first separately building linguistic and visual representations of the same concepts, and then merging them. This is obviously very different from how humans learn about concepts, by hearing words in a situated perceptual context. Second, MDSMs assume that both linguistic and visual information is available for all words, with no generalization of knowledge across modalities. Third, because of this latter assumption of full linguistic and visual coverage, current MDSMs, paradoxically, cannot be applied to computer vision tasks such as image labeling or retrieval, since they do not generalize to images or words beyond their training set.

We introduce the *multimodal skip-gram* models, two new MDSMs that address all the issues above. The models build upon the very effective skip-gram approach of Mikolov et al. (2013a), that constructs vector representations by learning, incrementally, to predict the linguistic contexts in which target words occur in a corpus. In our extension, for a subset of the target words, relevant visual evidence from

natural images is presented together with the corpus contexts (just like humans hear words accompanied by concurrent perceptual stimuli). The model must learn to predict these visual representations jointly with the linguistic features. The joint objective encourages the propagation of visual information to representations of words for which no direct visual evidence was available in training. The resulting multimodally-enhanced vectors achieve remarkably good performance both on traditional semantic benchmarks, and in their new application to the “zero-shot” image labeling and retrieval scenario. Very interestingly, indirect visual evidence also affects the representation of abstract words, paving the way to ground-breaking cognitive studies and novel applications in computer vision.

## 2 Related Work

There is by now a large literature on multimodal distributional semantic models. We focus here on a few representative systems. Bruni et al. (2014) propose a straightforward approach to MDSM induction, where text- and image-based vectors for the same words are constructed independently, and then “mixed” by applying the Singular Value Decomposition to their concatenation. An empirically superior model has been proposed by Silberer and Lapata (2014), who use more advanced visual representations relying on images annotated with high-level “visual attributes”, and a multimodal fusion strategy based on stacked autoencoders. Kiela and Bottou (2014) adopt instead a simple concatenation strategy, but obtain empirical improvements by using state-of-the-art convolutional neural networks to extract visual features, and the skip-gram model for text. These and related systems take a two-stage approach to derive multimodal spaces (unimodal induction followed by fusion), and they are only tested on concepts for which both textual and visual labeled training data are available (the pioneering model of Feng and Lapata (2010) did learn from text and images jointly using Topic Models, but was shown to be empirically weak by Bruni et al. (2014)).

Howell et al. (2005) propose an incremental multimodal model based on simple recurrent networks (Elman, 1990), focusing on *grounding propagation*

from early-acquired concrete words to a larger vocabulary. However, they use subject-generated features as surrogate for realistic perceptual information, and only test the model in small-scale simulations of word learning. Hill and Korhonen (2014), whose evaluation focuses on how perceptual information affects different word classes more or less effectively, similarly to Howell et al., integrate perceptual information in the form of subject-generated features and text from image annotations into a skip-gram model. They inject perceptual information by merging words expressing perceptual features with corpus contexts, which amounts to linguistic-context re-weighting, thus making it impossible to separate linguistic and perceptual aspects of the induced representation, and to extend the model with non-linguistic features. We use instead authentic image analysis as proxy to perceptual information, and we design a robust way to incorporate it, easily extendible to other signals, such as feature norm or brain signal vectors (Fyshe et al., 2014).

The recent work on so-called *zero-shot learning* to address the annotation bottleneck in image labeling (Frome et al., 2013; Lazaridou et al., 2014; Socher et al., 2013) looks at image- and text-based vectors from a different perspective. Instead of combining visual and linguistic information in a common space, it aims at learning a mapping from image- to text-based vectors. The mapping, induced from annotated data, is then used to project images of objects that were not seen during training onto linguistic space, in order to retrieve the nearest word vectors as labels. Multimodal word vectors should be better-suited than purely text-based vectors for the task, as their similarity structure should be closer to that of images. However, traditional MDSMs cannot be used in this setting, because they do not cover words for which no manually annotated training images are available, thus defeating the generalizing purpose of zero-shot learning. We will show below that our multimodal vectors, that are not hampered by this restriction, do indeed bring a significant improvement over purely text-based linguistic representations in the zero-shot setup.

Multimodal language-vision spaces have also been developed with the goal of improving caption generation/retrieval and caption-based image retrieval (Karpathy et al., 2014; Kiros et al., 2014;

Mao et al., 2014; Socher et al., 2014). These methods rely on necessarily limited collections of captioned images as sources of multimodal evidence, whereas we automatically enrich a very large corpus with images to induce general-purpose multimodal word representations, that could be used as input embeddings in systems specifically tuned to caption processing. Thus, our work is complementary to this line of research.

### 3 Multimodal Skip-gram Architecture

#### 3.1 Skip-gram Model

We start by reviewing the standard SKIP-GRAM model of Mikolov et al. (2013a), in the version we use. Given a text corpus, SKIP-GRAM aims at inducing word representations that are good at predicting the *context* words surrounding a *target* word. Mathematically, it maximizes the objective function:

$$\frac{1}{T} \sum_{t=1}^T \left( \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t) \right) \quad (1)$$

where  $w_1, w_2, \dots, w_T$  are words in the training corpus and  $c$  is the size of the window around target  $w_t$ , determining the set of context words to be predicted by the induced representation of  $w_t$ . Following Mikolov et al., we implement a subsampling option randomly discarding context words as an inverse function of their frequency, controlled by hyperparameter  $t$ . The probability  $p(w_{t+j} | w_t)$ , the core part of the objective in Equation 1, is given by softmax:

$$p(w_{t+j} | w_t) = \frac{e^{u'_{w_{t+j}} \cdot u_{w_t}}}{\sum_{w'=1}^W e^{u'_{w'} \cdot u_{w_t}}} \quad (2)$$

where  $u_w$  and  $u'_w$  are the context and target vector representations of word  $w$  respectively, and  $W$  is the size of the vocabulary. Due to the normalization term, Equation 2 requires  $O(|W|)$  time complexity. A considerable speedup to  $O(\log |W|)$ , is achieved by using the hierarchical version of Equation 2 (Morin and Bengio, 2005), adopted here.

#### 3.2 Injecting visual knowledge

We now assume that word learning takes place in a *situated* context, in which, for a subset of the target words, the corpus contexts are accompanied by a




Figure 1: “Cartoon” of MMSKIP-GRAM-B. Linguistic context vectors are actually associated to classes of words in a tree, not single words. SKIP-GRAM is obtained by ignoring the visual objective, MMSKIP-GRAM-A by fixing  $M^{u \rightarrow v}$  to the identity matrix.

visual representation of the concepts they denote (just like in a conversation, where a linguistic utterance will often be produced in a visual scene including some of the word referents). The visual representation is also encoded in a vector (we describe in Section 4 below how we construct it). We thus make the skip-gram “multimodal” by adding a second, *visual* term to the original *linguistic* objective, that is, we extend Equation 1 as follow:

$$\frac{1}{T} \sum_{t=1}^T (\mathcal{L}_{ling}(w_t) + \mathcal{L}_{vision}(w_t)) \quad (3)$$

where  $\mathcal{L}_{ling}(w_t)$  is the text-based skip-gram objective  $\sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t)$ , whereas the  $\mathcal{L}_{vision}(w_t)$  term forces word representations to take visual information into account. Note that if a word  $w_t$  is not associated to visual information, as is systematically the case, e.g., for determiners and non-imageable nouns, but also more generally for any word for which no visual data are available,  $\mathcal{L}_{vision}(w_t)$  is set to 0.

We now propose two variants of the visual objective, resulting in two distinguished multi-modal versions of the skip-gram model.

#### 3.3 Multi-modal Skip-gram Model A

One way to force word embeddings to take visual representations into account is to try to directly increase the similarity (expressed, for example, by the *cosine*) between linguistic and visual rep-

representations, thus aligning the dimensions of the linguistic vector with those of the visual one (recall that we are inducing the first, while the second is fixed), and making the linguistic representation of a concept “move” closer to its visual representation. We maximize similarity through a max-margin framework commonly used in models connecting language and vision (Weston et al., 2010; Frome et al., 2013). More precisely, we formulate the visual objective  $\mathcal{L}_{vision}(w_t)$  as:

$$- \sum_{w' \sim P_n(w)} \max(0, \gamma - \cos(u_{w_t}, v_{w_t}) + \cos(u_{w_t}, v_{w'})) \quad (4)$$

where the minus sign turns a loss into a cost,  $\gamma$  is the margin,  $u_{w_t}$  is the target multimodally-enhanced word representation we aim to learn,  $v_{w_t}$  is the corresponding visual vector (fixed in advance) and  $v_{w'}$  ranges over visual representations of words (featured in our image dictionary) randomly sampled from distribution  $P_n(w_t)$ . These random visual representations act as “negative” samples, encouraging  $u_{w_t}$  to be more similar to its own visual representation than to that of other words. The sampling distribution is currently set to uniform, and the number of negative samples controlled by hyperparameter  $k$ .

### 3.4 Multi-modal Skip-gram Model B

The visual objective in MMSKIP-GRAM-A has the drawback of assuming a direct comparison of linguistic and visual representations, constraining them to be of equal size. MMSKIP-GRAM-B lifts this constraint by including an extra layer mediating between linguistic and visual representations (see Figure 1 for a sketch of MMSKIP-GRAM-B). Learning this layer is equivalent to estimating a cross-modal mapping matrix from linguistic onto visual representations, jointly induced with linguistic word embeddings. The extension is straightforwardly implemented by substituting, into Equation 4, the word representation  $u_{w_t}$  with  $z_{w_t} = M^{u \rightarrow v} u_{w_t}$ , where  $M^{u \rightarrow v}$  is the cross-modal mapping matrix to be induced. To avoid overfitting, we also add an L2 regularization term for  $M^{u \rightarrow v}$  to the overall objective (Equation 3), with its relative importance controlled by hyperparameter  $\lambda$ .

## 4 Experimental Setup

The parameters of all models are estimated by back-propagation of error via stochastic gradient descent.

Our text corpus is a Wikipedia 2009 dump comprising approximately 800M tokens.<sup>1</sup> To train the multimodal models, we add visual information for 5,100 words that have an entry in ImageNet (Deng et al., 2009), occur at least 500 times in the corpus and have concreteness score  $\geq 0.5$  according to Turney et al. (2011). On average, about 5% tokens in the text corpus are associated to a visual representation. To construct the visual representation of a word, we sample 100 pictures from its ImageNet entry, and extract a 4096-dimensional vector from each picture using the Caffe toolkit (Jia et al., 2014), together with the pre-trained convolutional neural network of Krizhevsky et al. (2012). The vector corresponds to activation in the top (FC7) layer of the network. Finally, we average the vectors of the 100 pictures associated to each word, deriving 5,100 aggregated visual representations.

**Hyperparameters** For both SKIP-GRAM and the MMSKIP-GRAM models, we fix hidden layer size to 300. To facilitate comparison between MMSKIP-GRAM-A and MMSKIP-GRAM-B, and since the former requires equal linguistic and visual dimensionality, we keep the first 300 dimensions of the visual vectors. For the linguistic objective, we use hierarchical softmax with a Huffman frequency-based encoding tree, setting frequency subsampling option  $t = 0.001$  and window size  $c = 5$ , without tuning. The following hyperparameters were tuned on the text9 corpus:<sup>2</sup> MMSKIP-GRAM-A:  $k=20$ ,  $\gamma=0.5$ ; MMSKIP-GRAM-B:  $k=5$ ,  $\gamma=0.5$ ,  $\lambda=0.0001$ .

## 5 Experiments

### 5.1 Approximating human judgments

**Benchmarks** A widely adopted way to test DSMs and their multimodal extensions is to measure how well model-generated scores approximate human similarity judgments about pairs of words. We put together various benchmarks covering diverse aspects of meaning, to gain insights on the effect of perceptual information on different similarity facets. Specifically, we test on general relatedness (*MEN*, Bruni et al. (2014), 3K pairs), e.g., pickles are related to hamburgers, semantic ( $\approx$  taxonomic) simi-

<sup>1</sup><http://wacky.sslmit.unibo.it>

<sup>2</sup><http://mattmahoney.net/dc/textdata.html>

larity (*Simlex-999*, Hill et al. (2014), 1K pairs; *SemSim*, Silberer and Lapata (2014), 7.5K pairs), e.g., pickles are similar to onions, as well as visual similarity (*VisSim*, Silberer and Lapata (2014), same pairs as *SemSim* with different human ratings), e.g., pickles look like zucchinis.

**Alternative Multimodal Models** We compare our models against several recent alternatives. We test the vectors made available by Kiela and Bottou (2014). Similarly to us, they derive textual features with the skip-gram model (from a portion of the Wikipedia and the British National Corpus) and use visual representations extracted from the ESP dataset (von Ahn and Dabbish, 2004) through a convolutional neural network (Oquab et al., 2014). They concatenate textual and visual features after normalizing to unit length and centering to zero mean. We also test the vectors that performed best in the evaluation of Bruni et al. (2014), based on textual features extracted from a 3B-token corpus and SIFT-based Bag-of-Visual-Words visual features (Sivic and Zisserman, 2003) extracted from the ESP collection. Bruni and colleagues fuse a weighted concatenation of the two components through SVD. We further reimplement both methods with our own textual and visual embeddings as CONCATENATION and SVD (with target dimensionality 300, picked without tuning). Finally, we present for comparison the results on *SemSim* and *VisSim* reported by Silberer and Lapata (2014), obtained with a stacked-autoencoders architecture run on textual features extracted from Wikipedia with the Strudel algorithm (Baroni et al., 2010) and attribute-based visual features (Farhadi et al., 2009) extracted from ImageNet.

All benchmarks contain a fair amount of words for which we did not use direct visual evidence. We are interested in assessing the models both in terms of how they fuse linguistic and visual evidence when they are both available, and for their robustness in lack of full visual coverage. We thus evaluate them in two settings. The visual-coverage columns of Table 1 (those on the right) report results on the subsets for which all compared models have access to direct visual information for both words. We further report results on the full sets (“100%” columns of Table 1) for models that can propagate visual information and that, consequently, can meaningfully be tested

on words without direct visual representations.

**Results** The state-of-the-art visual CNN FEATURES alone perform remarkably well, outperforming the purely textual model (SKIP-GRAM) in two tasks, and achieving the best absolute performance on the visual-coverage subset of *Simlex-999*. Regarding multimodal *fusion* (that is, focusing on the visual-coverage subsets), both MMSKIP-GRAM models perform very well, at the top or just below it on all tasks, with comparable results for the two variants. Their performance is also good on the full data sets, where they consistently outperform SKIP-GRAM and SVD (that is much more strongly affected by lack of complete visual information). They’re just a few points below the state-of-the-art MEN correlation (0.8), achieved by Baroni et al. (2014) with a corpus 3 larger than ours and extensive tuning. MMSKIP-GRAM-B is close to the state of the art for *Simlex-999*, reported by the resource creators to be at 0.41 (Hill et al., 2014). Most impressively, MMSKIP-GRAM-A reaches the performance level of the Silberer and Lapata (2014) model on their *SemSim* and *VisSim* data sets, despite the fact that the latter has full visual-data coverage and uses attribute-based image representations, requiring supervised learning of attribute classifiers, that achieve performance in the semantic tasks comparable or higher than that of our CNN features (see Table 3 in Silberer and Lapata (2014)). Finally, if the multimodal models (unsurprisingly) bring about a large performance gain over the purely linguistic model on visual similarity, the improvement is consistently large also for the other benchmarks, confirming that multimodality leads to better semantic models in general, that can help in capturing different types of similarity (general relatedness, strictly taxonomic, perceptual).

While we defer to further work a better understanding of the relation between multimodal grounding and different similarity relations, Table 2 provides qualitative insights on how injecting visual information changes the structure of semantic space. The top SKIP-GRAM neighbours of *donuts* are places where you might encounter them, whereas the multimodal models relate them to other take-away food, ranking visually-similar pizzas at the top. The *owl* example shows how multimodal

Model	MEN		Simlex-999		SemSim		VisSim	
	100%	42%	100%	29%	100%	85%	100%	85%
KIELA AND BOTTOU	-	0.74	-	0.33	-	0.60	-	0.50
BRUNI ET AL.	-	0.77	-	0.44	-	0.69	-	0.56
SILBERER AND LAPATA	-	-	-	-	0.70	-	0.64	-
CNN FEATURES	-	0.62	-	0.54	-	0.55	-	0.56
SKIP-GRAM	0.70	0.68	0.33	0.29	0.62	0.62	0.48	0.48
CONCATENATION	-	0.74	-	0.46	-	0.68	-	0.60
SVD	0.61	0.74	0.28	0.46	0.65	0.68	0.58	0.60
MMSKIP-GRAM-A	0.75	0.74	0.37	0.50	0.72	0.72	0.63	0.63
MMSKIP-GRAM-B	0.74	0.76	0.40	0.53	0.66	0.68	0.60	0.60

Table 1: Spearman correlation between model-generated similarities and human judgments. Right columns report correlation on visual-coverage subsets (percentage of original benchmark covered by subsets on first row of respective columns). First block reports results for out-of-the-box models; second block for visual and textual representations alone; third block for our implementation of multimodal models.

Target	SKIP-GRAM	MMSKIP-GRAM-A	MMSKIP-GRAM-B
donut	fridge, diner, candy	pizza, sushi, sandwich	pizza, sushi, sandwich
owl	pheasant, woodpecker, squirrel	eagle, woodpecker, falcon	eagle, falcon, hawk
mural	sculpture, painting, portrait	painting, portrait, sculpture	painting, portrait, sculpture
tobacco	coffee, cigarette, corn	cigarette, cigar, corn	cigarette, cigar, smoking
depth	size, bottom, meter	sea, underwater, level	sea, size, underwater
chaos	anarchy, despair, demon	demon, anarchy, destruction	demon, anarchy, shadow

Table 2: Ordered top 3 neighbours of example words in purely textual and multimodal spaces. Only *donut* and *owl* were trained with direct visual information.

models pick taxonomically closer neighbours of concrete objects, since often closely related things also look similar (Bruni et al., 2014). In particular, both multimodal models get rid of squirrels and offer other birds of prey as nearest neighbours. No direct visual evidence was used to induce the embeddings of the remaining words in the table, that are thus influenced by vision only by propagation. The subtler but systematic changes we observe in such cases suggest that this indirect propagation is not only non-damaging with respect to purely linguistic representations, but actually beneficial. For the concrete *mural* concept, both multimodal models rank paintings and portraits above less closely related sculptures (they are not a form of painting). For *tobacco*, both models rank cigarettes and cigar over coffee, and MMSKIP-GRAM-B avoids the arguably less common “crop” sense cued by corn. The last two examples show how the multimodal models turn up the embodiment level in their representation of abstract words. For *depth*, their neighbours suggest a concrete marine setup

over the more abstract measurement sense picked by the MMSKIP-GRAM neighbours. For *chaos*, they rank a demon, that is, a concrete agent of chaos at the top, and replace the more abstract notion of despair with equally gloomy but more imageable shadows and destruction (more on abstract words below).

## 5.2 Zero-shot image labeling and retrieval

The multimodal representations induced by our models should be better suited than purely text-based vectors to label or retrieve images. In particular, given that the quantitative and qualitative results collected so far suggest that the models propagate visual information across words, we apply them to image labeling and retrieval in the challenging zero-shot setup (see Section 2 above).<sup>3</sup>

<sup>3</sup>We will refer here, for conciseness’ sake, to *image* labeling/retrieval, but, as our visual vectors are aggregated representations of images, the tasks we’re modeling consist, more precisely, in labeling a set of pictures denoting the same object and retrieving the corresponding set given the name of the object.

**Setup** We take out as test set 25% of the 5.1K words we have visual vectors for. The multimodal models are re-trained without visual vectors for these words, using the same hyperparameters as above. For both tasks, the search for the correct word label/image is conducted on the whole set of 5.1K word/visual vectors.

In the image labeling task, given a visual vector representing an image, we map it onto word space, and label the image with the word corresponding to the nearest vector. To perform the vision-to-language mapping, we train a Ridge regression by 5-fold cross-validation on the test set (for SKIP-GRAM only, we also add the remaining 75% of word-image vector pairs used in estimating the multimodal models to the Ridge training data).<sup>4</sup>

In the image retrieval task, given a linguistic/multimodal vector, we map it onto visual space, and retrieve the nearest image. For SKIP-GRAM, we use Ridge regression with the same training regime as for the labeling task. For the multimodal models, since maximizing similarity to visual representations is already part of their training objective, we do not fit an extra mapping function. For MMSKIP-GRAM-A, we directly look for nearest neighbours of the learned embeddings in visual space. For MMSKIP-GRAM-B, we use the  $M^{u \rightarrow v}$  mapping function induced while learning word embeddings.

**Results** In image labeling (Table 3) SKIP-GRAM is outperformed by both multimodal models, confirming that these models produce vectors that are directly applicable to vision tasks thanks to visual propagation. The most interesting results however are achieved in image retrieval (Table 4), which is essentially the task the multimodal models have been implicitly optimized for, so that they could be applied to it without any specific training. The strategy of directly querying for the nearest visual vectors of the MMSKIP-GRAM-A word embeddings works remarkably well, outperforming on the higher ranks SKIP-GRAM, which requires an ad-hoc mapping function. This suggests that the multimodal

<sup>4</sup>We use one fold to tune Ridge  $\lambda$ , three to estimate the mapping matrix and test in the last fold. To enforce strict zero-shot conditions, we exclude from the test fold labels occurring in the LSVRC2012 set that was employed to train the CNN of Krizhevsky et al. (2012), that we use to extract visual features.

	P@1	P@2	P@10	P@20	P@50
SKIP-GRAM	1.5	2.6	14.2	23.5	36.1
MMSKIP-GRAM-A	2.1	3.7	16.7	24.6	37.6
MMSKIP-GRAM-B	2.2	5.1	20.2	28.5	43.5

Table 3: Percentage precision@ $k$  results in the zero-shot image labeling task.

	P@1	P@2	P@10	P@20	P@50
SKIP-GRAM	1.9	3.3	11.5	18.5	30.4
MMSKIP-GRAM-A	1.9	3.2	13.9	20.2	33.6
MMSKIP-GRAM-B	1.9	3.8	13.2	22.5	38.3

Table 4: Percentage precision@ $k$  results in the zero-shot image retrieval task.

embeddings we are inducing, while general enough to achieve good performance in the semantic tasks discussed above, encode sufficient visual information for direct application to image analysis tasks. This is especially remarkable because the word vectors we are testing were not matched with visual representations at model training time, and are thus multimodal only by propagation. The best performance is achieved by MMSKIP-GRAM-B, confirming our claim that its  $M^{u \rightarrow v}$  matrix acts as a multimodal mapping function.

### 5.3 Abstract words

We have already seen, through the *depth* and *chaos* examples of Table 2, that the indirect influence of visual information has interesting effects on the representation of abstract terms. The latter have received little attention in multimodal semantics, with Hill and Korhonen (2014) concluding that abstract nouns, in particular, do not benefit from propagated perceptual information, and their representation is even harmed when such information is forced on them (see Figure 4 of their paper). Still, embodied theories of cognition have provided considerable evidence that abstract concepts are also grounded in the senses (Barsalou, 2008; Lakoff and Johnson, 1999). Since the word representations produced by MMSKIP-GRAM-A, including those pertaining to abstract concepts, can be directly used to search for near images in visual space, we decided to verify, experimentally, if these near images (of concrete things) are relevant not only for concrete words, as

expected, but also for abstract ones, as predicted by embodied views of meaning.

More precisely, we focused on the set of 200 words that were sampled across the USF norms concreteness spectrum by Kiela et al. (2014) (2 words had to be excluded for technical reasons). This set includes not only concrete (*meat*) and abstract (*thought*) nouns, but also adjectives (*boring*), verbs (*teach*), and even grammatical terms (*how*). Some words in the set have relatively high concreteness ratings, but are not particularly imageable, e.g.: *hot*, *smell*, *pain*, *sweet*. For each word in the set, we extracted the nearest neighbour picture of its MMSKIP-GRAM-A representation, and matched it with a random picture. The pictures were selected from a set of 5,100, all labeled with distinct words (the picture set includes, for each of the words associated to visual information as described in Section 4, the nearest picture to its aggregated visual representation). Since it is much more common for concrete than abstract words to be directly represented by an image in the picture set, when searching for the nearest neighbour we excluded the picture labeled with the word of interest, if present (e.g., we excluded the picture labeled *tree* when picking the nearest neighbour of the word *tree*). We ran a CrowdFlower<sup>5</sup> survey in which we presented each test word with the two associated images (randomizing presentation order of nearest and random picture), and asked subjects which of the two pictures they found more closely related to the word. We collected minimally 20 judgments per word. Subjects showed large agreement (median proportion of majority choice at 90%), confirming that they understood the task and behaved consistently.

We quantify performance in terms of proportion of words for which the number of votes for the nearest neighbour picture is significantly above chance according to a two-tailed binomial test. We set significance at  $p < 0.05$  after adjusting all p-values with the Holm correction for running 198 statistical tests. The results in Table 5 indicate that, in about half the cases, the nearest picture to a word MMSKIP-GRAM-A representation is meaningfully related to the word. As expected, this is more often the case for concrete than abstract words. Still, we also observe a

<sup>5</sup><http://www.crowdfunder.com>

	<i>global</i>	<i> words </i>	<i>unseen</i>	<i> words </i>
all	48%	198	30%	127
concrete	73%	99	53%	30
abstract	23%	99	23%	97

Table 5: Subjects’ preference for nearest visual neighbour of words in Kiela et al. (2014) vs. random pictures. Figure of merit is percentage proportion of significant results in favor of nearest neighbour across words. Results are reported for the whole set, as well as for words above (*concrete*) and below (*abstract*) the concreteness rating median. The *unseen* column reports results when words exposed to direct visual evidence during training are discarded. The *words* columns report set cardinality.



Figure 2: Examples of nearest visual neighbours of some abstract words: on the left, cases where subjects preferred the neighbour to the random foil; on the right, cases where they did not.

significant preference for the model-predicted nearest picture for about one fourth of the abstract terms. Whether a word was exposed to direct visual evidence during training is of course making a big difference, and this factor interacts with concreteness, as only two abstract words were matched with images during training.<sup>6</sup> When we limit evaluation to word representations that were not exposed to pictures during training, the difference between concrete and abstract terms, while still large, becomes less dramatic than if all words are considered.

Figure 2 shows four cases in which subjects expressed a strong preference for the nearest visual neighbour of a word. *Freedom*, *god* and *theory* are strikingly in agreement with the view, from embodied theories, that abstract words are grounded in rel-

<sup>6</sup>In both cases, the images actually depict concrete senses of the words: a memory board for *memory* and a stop sign for *stop*.



evant concrete scenes and situations. The *together* example illustrates how visual data might ground abstract notions in surprising ways. For all these cases, we can borrow what Howell et al. (2005) say about visual propagation to abstract words (p. 260):

Intuitively, this is something like trying to explain an abstract concept like *love* to a child by using concrete examples of scenes or situations that are associated with love. The abstract concept is never fully grounded in external reality, but it does inherit some meaning from the more concrete concepts to which it is related.

Of course, not all examples are good: the last column of Figure 2 shows cases with no obvious relation between words and visual neighbours (subjects preferred the random images by a large margin).

The multimodal vectors we induce also display an interesting intrinsic property related to the hypothesis that grounded representations of abstract words are more complex than for concrete ones, since abstract concepts relate to varied and composite situations (Barsalou and Wiemer-Hastings, 2005). A natural corollary of this idea is that visually-grounded representations of abstract concepts should be more diverse: If you think of dogs, very similar images of specific dogs will come to mind. You can also imagine the abstract notion of freedom, but the nature of the related imagery will be much more varied. Recently, Kiela et al. (2014) have proposed to measure abstractness by exploiting this very same intuition. However, they rely on manual annotation of pictures via Google Images and define an *ad-hoc* measure of *image dispersion*. We conjecture that the representations naturally induced by our models display a similar property. In particular, the *entropy* of our multimodal vectors, being an expression of how varied the information they encode is, should correlate with the degree of abstractness of the corresponding words. As Figure 3(a) shows, there is indeed a difference in entropy between the most concrete (*meat*) and most abstract (*hope*) words in the Kiela et al. set.

To test the hypothesis quantitatively, we measure the correlation of entropy and concreteness on the 200 words in the Kiela et al. (2014) set.<sup>7</sup> Figure 3(b) shows that the entropies of both the

<sup>7</sup>Since the vector dimensions range over the real number line, we calculate entropy on vectors that are unit-normed after adding a small constant insuring all values are positive.




Figure 3: (a) Distribution of MMSKIP-GRAM-A vector activation for *meat* (blue) and *hope* (red). (b) Spearman  $\rho$  between concreteness and various measures on the Kiela et al. (2014) set.

MMSKIP-GRAM-A representations and those generated by mapping MMSKIP-GRAM-B vectors onto visual space (MMSKIP-GRAM-B\*) achieve very high correlation (but, interestingly, not MMSKIP-GRAM-B). This is further evidence that multimodal learning is grounding the representations of both concrete and abstract words in meaningful ways.

## 6 Conclusion

We introduced two multimodal extensions of SKIP-GRAM. MMSKIP-GRAM-A is trained by directly optimizing the similarity of words with their visual representations, thus forcing maximum interaction between the two modalities. MMSKIP-GRAM-B includes an extra mediating layer, acting as a cross-modal mapping component. The ability of the models to *integrate* and *propagate* visual information resulted in word representations that performed well in both semantic and vision tasks, and could be used as input in systems benefiting from prior visual knowledge (e.g., caption generation). Our results with abstract words suggest the models might also help in tasks such as metaphor detection, or even retrieving/generating pictures of abstract concepts. Their incremental nature makes them well-suited for cognitive simulations of grounded language acquisition, an avenue of research we plan to explore further.

## Acknowledgments

We thank Adam Liska, Tomas Mikolov, the reviewers and the NIPS 2014 Learning Semantics audience. We were supported by ERC 2011 Starting Independent Research Grant n. 283554 (COMPOSES).

## References

- Marco Baroni, Eduard Barbu, Brian Murphy, and Massimo Poesio. 2010. Strudel: A distributional semantic model based on properties and types. *Cognitive Science*, 34(2):222–254.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of ACL*, pages 238–247, Baltimore, MD.
- Lawrence Barsalou and Katja Wiemer-Hastings. 2005. Situating abstract concepts. In D. Pecher and R. Zwaan, editors, *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thought*, pages 129–163. Cambridge University Press, Cambridge, UK.
- Lawrence Barsalou. 2008. Grounded cognition. *Annual Review of Psychology*, 59:617–645.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam Khanh Tran. 2012. Distributional semantics in Technicolor. In *Proceedings of ACL*, pages 136–145, Jeju Island, Korea.
- Elia Bruni, Nam Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Stephen Clark. 2015. Vector space models of lexical meaning. In Shalom Lappin and Chris Fox, editors, *Handbook of Contemporary Semantics*, 2nd ed. Blackwell, Malden, MA. In press; [http://www.cl.cam.ac.uk/~sc609/pubs/sem\\_handbook.pdf](http://www.cl.cam.ac.uk/~sc609/pubs/sem_handbook.pdf).
- Jia Deng, Wei Dong, Richard Socher, Lia-Ji Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of CVPR*, pages 248–255, Miami Beach, FL.
- Jeffrey L Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. 2009. Describing objects by their attributes. In *Proceedings of CVPR*, pages 1778–1785, Miami Beach, FL.
- Yansong Feng and Mirella Lapata. 2010. Visual information in semantic representation. In *Proceedings of HLT-NAACL*, pages 91–99, Los Angeles, CA.
- Andrea Frome, Greg Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. 2013. DeViSE: A deep visual-semantic embedding model. In *Proceedings of NIPS*, pages 2121–2129, Lake Tahoe, NV.
- Alona Fyshe, Partha P Talukdar, Brian Murphy, and Tom M Mitchell. 2014. Interpretable semantic vectors from a joint model of brain-and text-based meaning. In *Proceedings of ACL*, pages 489–499.
- Arthur Glenberg and David Robertson. 2000. Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, 3(43):379–401.
- Stevan Harnad. 1990. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346.
- 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 EMNLP*, pages 255–265, Doha, Qatar.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2014. SimLex-999: Evaluating semantic models with (genuine) similarity estimation. <http://arxiv.org/abs/arXiv:1408.3456>.
- Steve Howell, Damian Jankowicz, and Suzanna Becker. 2005. A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language*, 53:258–276.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.
- Andrej Karpathy, Armand Joulin, and Li Fei-Fei. 2014. Deep fragment embeddings for bidirectional image sentence mapping. In *Proceedings of NIPS*, pages 1097–1105, Montreal, Canada.
- Douwe Kiela and Léon Bottou. 2014. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of EMNLP*, pages 36–45, Doha, Qatar.
- 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 ACL*, pages 835–841, Baltimore, MD.
- Ryan Kiros, Ruslan Salakhutdinov, and Richard Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. In *Proceedings of the NIPS Deep Learning and Representation Learning Workshop*, Montreal, Canada. Published online: <http://www.dlworkshop.org/accepted-papers>.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Proceedings of NIPS*, pages 1097–1105, Lake Tahoe, Nevada.
- George Lakoff and Mark Johnson. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Basic Books, New York.
- 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 ACL*, pages 1403–1414, Baltimore, MD.
- Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, and Alan Yuille. 2014. Explain images with multimodal recurrent neural networks. In *Proceedings of the NIPS Deep Learning and Representation Learning Workshop*, Montreal, Canada. Published online: <http://www.dlworkshop.org/accepted-papers>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. 2013a. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*, pages 3111–3119, Lake Tahoe, NV.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. In *Proceedings of NAACL*, pages 746–751, Atlanta, Georgia.
- Frederic Morin and Yoshua Bengio. 2005. Hierarchical probabilistic neural network language model. In *Proceedings of AISTATS*, pages 246–252, Barbados.
- Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of CVPR*.
- John Searle. 1984. *Minds, Brains and Science*. Harvard University Press, Cambridge, MA.
- Carina Silberer and Mirella Lapata. 2014. Learning grounded meaning representations with autoencoders. In *Proceedings of ACL*, pages 721–732, Baltimore, Maryland.
- Josef Sivic and Andrew Zisserman. 2003. Video Google: A text retrieval approach to object matching in videos. In *Proceedings of ICCV*, pages 1470–1477, Nice, France.
- Richard Socher, Milind Ganjoo, Christopher Manning, and Andrew Ng. 2013. Zero-shot learning through cross-modal transfer. In *Proceedings of NIPS*, pages 935–943, Lake Tahoe, NV.
- Richard Socher, Quoc Le, Christopher Manning, and Andrew Ng. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218.
- Peter Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188.
- Peter Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. Literal and metaphorical sense identification through concrete and abstract context. In *Proceedings of EMNLP*, pages 680–690, Edinburgh, UK.
- Luis von Ahn and Laura Dabbish. 2004. Labeling images with a computer game. In *Proceedings of CHI*, pages 319–326, Vienna, Austria.
- Jason Weston, Samy Bengio, and Nicolas Usunier. 2010. Large scale image annotation: learning to rank with joint word-image embeddings. *Machine learning*, 81(1):21–35.