Visually Grounded, Situated Learning in Neural Models

Alexander G. Ororbia, Ankur Mali, Matthew A. Kelly, and **David Reitter**
The Pennsylvania State University
University Park, PA 16802, USA
ago109, aam35, matthew.kelly, reitter @psu.edu

arXiv:1805.11546v1 [cs.CL] 29 May 2018

Abstract

The theory of situated cognition postulates that language is inseparable from its physical context–words, phrases, and sentences must be learned in the context of the objects or concepts to which they refer. Yet, statistical language models are trained on words alone. This makes it impossible for language models to connect to the real world–the world described in the sentences presented to the model. In this paper, we examine the generalization ability of neural language models trained with a visual context. A multimodal connectionist language architecture based on the Differential State Framework is proposed, which outperforms its equivalent trained on language alone, even when no visual context is available at test time. Superior performance for language models trained with a visual context is robust across different languages and models.

1 Introduction

The theory of situated cognition postulates that a person's knowledge is inseparable from the physical or social context in which it is learned and used (Greeno and Moore, 1993). Knowledge of language cannot be separated from its physical context, which allows words and sentences to be learned by grounding them in reference to objects or natural concepts on hand (see Roy and Reiter, 2005, for a review). Nor can knowledge of language be separated from its social context, where language is learned interactively through communicating with others to facilitate problem-solving. Simply put, language does not occur in a vacuum.

Yet, statistical language models, typically connectionist systems, are often trained in such a vacuum. Sequences of symbols, such as sentences or phrases composed of words in any language, such as English or German, are often fed into the model independently of any real-world context they might describe. In the classical language modeling framework, a neural model is tasked with a series of next-step prediction tasks, learning to predict a word based on a history of words it has seen so far. While these models learn a great deal of linguistic structure from these symbol sequences alone, acquiring the essence of basic syntax, it is highly unlikely that this approach can create models that acquire much in terms of semantics or pragmatics, which are integral to the human experience of language. How might one build neural language models that "understand" the semantic content held within the symbol sequences, of any language, presented to it?

In this paper, we take a small step towards a model that understands language by training a neural architecture jointly on corresponding linguistic and visual data. From an image-captioning dataset, we create a multi-lingual corpus where sentences are mapped to the real-world images they describe. We ask how adding such real-world context at training can improve the performance of language models. We extend the Δ-RNN (Ororbia II et al., 2017), the Long Short Term Memory (LSTM; Hochreiter and Schmidhuber, 1997) and the Gated Recurrent Unit (GRU; Cho et al., 2014) to incorporate visual context information, creating a unified multi-modal connectionist architecture. We find that the models acquire more knowledge of language than if they were trained without corresponding, real-world visual context.

2 Related Work

The Perceptual Symbol Systems theory holds that all of cognition, language, reasoning, and mem-

ory, is grounded in perceptual features (Barsalou, 1999). Both behavioral and neuroimaging studies have found considerable evidence for the contribution of perceptual information to linguistic tasks (Barsalou, 2008). Cognitive theory has long held that language is acquired jointly with perception through interaction with the environment (e.g. Frank et al., 2008). Cognitive models can account for bootstrapped learning of word meaning and syntax when language is paired with ambiguous and limited perceptual experience (Abend et al., 2017), and for the ability of children to rapidly acquire new words by inferring the referent from their physical environment (Alishahi et al., 2008).

A number of models of distributional semantics integrate word co-occurrence data extracted from a corpus with perceptual data, either to achieve a better model of language as it exists in the minds of humans (Kievit-Kylar and Jones, 2011; Johns and Jones, 2012) or to improve performance on machine learning tasks such as object recognition (Frome et al., 2013), image captioning (Kiros et al., 2014), or image search (Socher et al., 2014).

Integrating language and perception can facilitate language acquisition by allowing models to infer how a new word is used from the perceptual features of its referent (Johns and Jones, 2012). Likewise, this integration allows models to infer the perceptual features of an unobserved referent from how a word is used in language (Johns and Jones, 2012). As a result, language data can be used to improve object recognition by providing information about unobserved or infrequently observed objects (Frome et al., 2013).

By representing the referents of concrete nouns as arrangements of elementary visual features (Biederman, 1987), Kievit-Kylar and Jones (2011) find that the visual features of nouns capture semantic typicality effects, and that a combined representation, consisting of both visual features and word co-occurrence data, more strongly correlates with human judgments of semantic similarity than representations extracted from a corpus alone. While modeling similarity judgments is distinct from the problem of predictive language modeling, we take this finding as evidence that visual perception informs semantics, which suggests there are gains to be had integrating perception with predictive language models.

While knowledge of concrete nouns benefits most directly from integrating perceptual data with language, verbs also benefit, as the perceptual features of verbs can be inferred from the features of the nouns they act upon (Johns and Jones, 2012), such that a model with access to perceptual features gains the ability to discriminate between actions afforded by a verb and actions that are not afforded by the verb (e.g., *hanging* a coat on a vacuum versus a cup).

Image Captioning (Kiros et al., 2014; Vinyals et al., 2015; Xu et al., 2015) systems have shown promising results in generating captions by mapping between vision and language. However such models are restricted to a single language and can introduce irreversible corruption to a vision signal if trained jointly, since randomly initialized language parameters generates Gaussian noise that can harm contextual interaction information. If a jointly trained vision and language model is trained on multiple languages then each language introduces language specific noise that would corrupt visual information.

In contrast to prior work in machine learning, our goal in integrating visual and linguistic data is not to accomplish a task such as image search or image captioning that inherently requires a mapping between these two modalities. Rather, our goal is to demonstrate that perceptual information is intrinsic to how humans process language, and as such, a language model that is trained on both visual and linguistic data will be a better model, consistently across languages, than a model trained on linguistic data alone.

Prior work in cognitive modeling has focused on models of distributional semantics that capture the similarity relations between words (e.g. Johns and Jones, 2012; Kievit-Kylar and Jones, 2011), whereas the model we propose here is a predictive language model.

Due to the ability of language models to probabilistically constrain input on the basis of preceding context and to classify linguistic material, these models play a central role in natural-language and speech processing applications. However, the psycholinguistic questions surrounding how people acquire and use linguistic knowledge are fundamentally different from the aims of machine learning. Using NLP-style language models to address psycholinguistic questions is a new approach that integrates well with the theory of predictive coding in cognitive psychology (Clark, 2013; Rao and Ballard, 1999). For

language processing this means that when reading text or comprehending speech, humans constantly anticipate what will be said next. This is a fast, implicit cognitive process that does not require symbol manipulation, but that can make use of the kind of sequence learning that recurrent neural models excel at. We do not propose such models as direct accounts of human language processing. Instead, our intent is to examine what can and cannot be learned with the addition of a non-linguistic modality (vision) at training time.

3 The Multimodal Neural Architecture

In designing our neural model, we start from the Differential State Framework (DSF, Ororbia II et al. (2017)), which unifies gated recurrent architectures under the general view that state memory is a simple parametrized mixture of "fast" and "slow" states. Our aim is to model sequences of symbols, such as the words that compose sentences, where at each time we process $\mathbf{x}_t$, or the one-hot encoding of a token.

One of the simplest models that can be derived from the DSF is the Δ-RNN, which has been shown to outperform most complex neural models in next-step symbol prediction tasks (Ororbia II et al., 2017). The model, with parameters $\Theta = \{W, U, V, \mathbf{b}, \mathbf{c}, \mathbf{b}_r, \beta_1, \beta_2, \alpha\}$, is defined as:

$$\mathbf{d}_t^{rec} = V\mathbf{h}_{t-1},\ \mathbf{d}_t^{dat} = W\mathbf{e}_{w,t}, \tag{1}$$

$$\mathbf{d}_t^1 = \alpha \otimes \mathbf{d}_t^{rec} \otimes\ \mathbf{d}_t^{dat} \tag{2}$$

$$\mathbf{d}_t^2 = \beta_1 \otimes \mathbf{d}_t^{rec} + \beta_2 \otimes \mathbf{d}_t^{dat}, \tag{3}$$

$$\mathbf{z}_t = \phi_{hid}(\mathbf{d}_t^1 + \mathbf{d}_t^2 + \mathbf{b}), \tag{4}$$

$$\mathbf{h}_t = \Phi((1 - \mathbf{r}) \otimes \mathbf{z}_t + \mathbf{r} \otimes \mathbf{h}_{t-1}), \text{ and,} \tag{5}$$

$$\mathbf{r} = 1/(1 + exp(-[\mathbf{d}_t^{dat} + \mathbf{b}_r])). \tag{6}$$

where $\mathbf{e}_{w,t}$ is the 1-of-k encoding of the word w at time t. Note that $\{\alpha, \beta_1, \beta_2\}$ are learnable bias vectors that modulate the internal multiplicative interactions and the rate gate $\mathbf{r}$ reuses the computed pre-activation term $\mathbf{d}_t^{dat}$. In contrast to the model originally trained in Ororbia II et al. (2017), the outer activation is the linear rectifier, $\Phi(v) = max(0, v)$, instead of the identity or hyperbolic tangent, because we found that it worked much better. We set the inner activation function $\phi_{hid}(v)$ to be $tanh(v) = \frac{(e^{(2v)}-1)}{(e^{(2v)}+1)}$.

To integrate visual context information into the Δ-RNN, we fuse the model with a neural vision system, motivated by promising recent work done in automated image captioning (Xu et al., 2015). We adopt a transfer learning approach and incorporate a state-of-the-art convolutional neural network into the Δ-RNN model, namely the Inception-v3 network (Szegedy et al., 2016).[1] The parameters of the vision network are fixed. As our focus is on language modeling and how the addition of visual context can improve neural network performance on the task, fixing the vision system prevents any noise from the language model from potentially corrupting the vision model and damaging its distributed representations. We leave learning the vision system jointly with the language model as future work.

To obtain a distributed representation of an image from the Inception-v3 network, we extract the vector produced from the final max-pooling layer, $\mathbf{c}$, after running an image through the model (note that this operation occurs right before the final, fully-connected processing layers which are usually task-specific parameters, such as in object classification). The Δ-RNN can make use of the information in this visual context vector if we modify its state computation in one of two ways. The first way would be to modify its inner state function to be a linear combination of the data-dependent pre-activation, the filtration, and a learned linear mapping of $\mathbf{c}$ as follows:

$$\mathbf{z}_t = \phi_{hid}(\mathbf{d}_t^1 + \mathbf{d}_t^2 + M\mathbf{c} + \mathbf{b}) \tag{7}$$

where M is a learnable synaptic connections that connect the visual context representation with the inner state. The second way to modify the Δ-RNN would be change its outer mixing function instead:

$$\mathbf{h}_t = \Phi([(1 - \mathbf{r}) \otimes \mathbf{z}_t + \mathbf{r} \otimes \mathbf{h}_{t-1}] \otimes (M\mathbf{c})) \tag{8}$$

Here we see the linearly-mapped visual context embedding interacts with the currently computation state through a multiplicative operation, allowing the visual-context to persist and work in a longer-term capacity. In either situation, using a parameter matrix M frees us from having to set the dimensionality of the hidden state to be the same as the context vector produced by the Inception-v3 network.

[1] In preliminary experiments, we also examined VGGNet and a few other variations, but found that the Inception worked the best when it came to acquiring somewhat more general distributed representations of natural images.

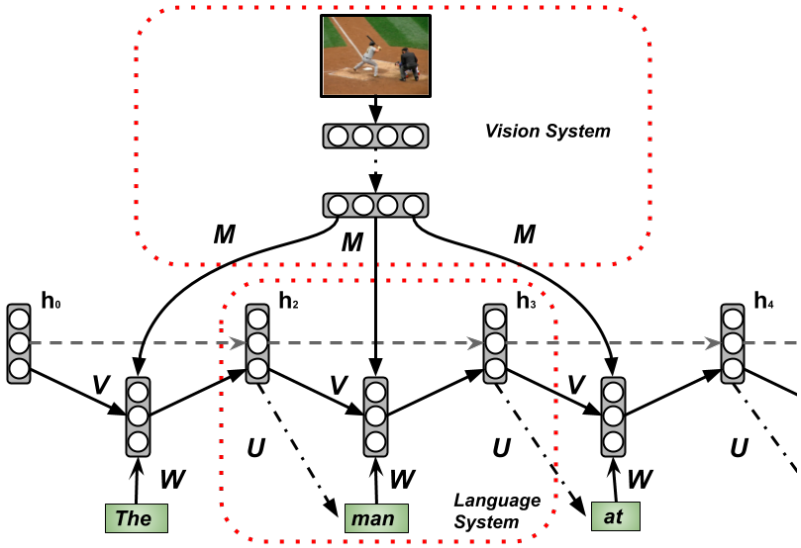

Figure 1: The multimodal Δ-RNN, unrolled over time. The gray-dashed connections represent the identity connections that carry over the slow-moving state while the dash-dotted black lines represent the next-step predictions made by the model. Solid black lines correspond to synaptic weight matrices (labeled accordingly).

We do not use regularization techniques with this model. The application of regularization techniques is, in principle, possible (and typically improves performance of the Δ-RNN), albeit it is inappropriate and indeed damaging to performance in this particular case, where an already compressed and regularized representation of the images from Inception-v3 serves as input to the multimodal language modeling network.

Let $w_1, \ldots, w_N$ be a variable-length sequence of N words corresponding to an image I. In general, the distribution over the variables follows the graphical model:

$$P_\theta(w_1, \ldots, w_T | I) = \prod_{t=1}^{T} P_\Theta(w_t | w_{<t}, I), \quad (9)$$

For all model variants the state $\mathbf{h}_t$ calculated at any time step is fed into a maximum-entropy classifier[2] defined as:

$$P(w, \mathbf{h}_t) = P_\Theta(w|\mathbf{h}_t) = \frac{\exp\left(w^{\mathrm{T}} U \mathbf{h}_t\right)}{\sum_{w'} \exp\left((w')^{\mathrm{T}} U \mathbf{h}_t\right)}, \quad (10)$$

The model parameters Θ optimized with respect to the sequence negative log likelihood:

$$\mathcal{L} = -\sum_{i=1}^{N} \sum_{t=1}^{T} \log P_\Theta(w_t|\mathbf{h}), \quad (11)$$

We employ back-propagation of errors, or differentiate with respect to the negative log likelihood objective function above, to calculate the gradients needed to update parameters.

4 Experiments

The experiments in this paper were conducted using the MS-COCO image-captioning dataset.[3] Images in the dataset contains significant amount of contextual information and also five human annotated captions per image.We extracted all the five sentences from the dataset and created 5 different ground truth splits. We translated ground truth splits into German and Spanish splits using state of the art Google Translation API. To our knowledge, this represents the first Multi-lingual MSCOCO dataset on situated learning. We process the corpus at the word-level and obtain a 16.6K vocabulary for English, 33.2K for German and 18.2k for Spanish.

Our primary concern is with the next-step prediction of words/tokens, which means the negative log likelihood and perplexity of the learned

[2]Bias term omitted for clarity.

[3]https://competitions.codalab.org/competitions/3221

Table 1: Generalization performance of language models trained and evaluated on linguistic data only (L), *full:* trained and evaluated on multimodal linguistic and visual data (LV), and, *blind*: trained on multimodal data (LV) but evaluated on language only (L).

	English		German		Spanish	
$ModelType$	Test-**NLL**	Test-**PPL**	Test-**NLL**	Test-**PPL**	Test-**NLL**	Test-**PPL**
Δ-RNN (L-L)	2.714	15.086	2.836	17.052	2.546	12.755
MM-Δ-RNN (full LV-LV)	2.645	14.086	2.777	16.082	2.405	11.082
MM-Δ-RNN (blind LV-L)	2.694	14.786	2.808	16.582	2.458	11.682
GRU (L-L)	2.764	15.871	2.854	17.369	2.554	12.866
MM-GRU (full LV-LV)	2.654	14.189	2.790	16.285	2.426	11.3089
MM-GRU (blind LV-L)	2.687	14.689	2.815	16.701	2.466	11.781
LSTM (L-L)	2.722	15.217	2.814	17.070	2.494	12.114
MM-LSTM (full LV-LV)	2.645	14.089	2.773	16.001	2.405	11.081
MM-LSTM (blind LV-L)	2.708	15.002	2.822	16.806	2.487	12.028

generative model is of high importance. This is different from the goals of machine translation or image captioning, which, in most cases, is concerned with a ranking of possible captions where one measures how similar the model's generated sequences are to ground-truth target phrases.

Baseline results were obtained with neural language models of text alone. For the Δ-RNN, this meant implementing a model using only Equations 1-7. To verify that the experiment generalizes beyond the specific architecture chosen, a Gated Recurrent Unit (GRU, Cho et al., 2014) and a Long Short Term Memory (LSTM, Hochreiter and Schmidhuber, 1997) were also trained. We compare these symbol-only baselines to the two variations of our proposed multimodal Δ-RNN, as described in the previous section. The multimodal variant of the GRU, where the context information is directly integrated into its inner function, is defined as follows:

$$\mathbf{d}_c = M\mathbf{c} \tag{12}$$

$$\mathbf{z}_t = \sigma(W_z\mathbf{x}_t + V_z\mathbf{h}_{t-1}) \tag{13}$$

$$\mathbf{r}_t = \sigma(W_r\mathbf{x}_t + V_r\mathbf{h}_{t-1}) \tag{14}$$

$$\widehat{\mathbf{h}}_t = \tanh(W_{\widehat{h}}\mathbf{x}_t + V_{\widehat{h}}(\mathbf{r}_t \otimes \mathbf{h}_{t-1})) \tag{15}$$

$$\mathbf{h}_t = [\mathbf{z}_t \otimes \mathbf{h}_{t-1} + (1 - \mathbf{z}_t) \otimes \widehat{\mathbf{h}}_t] \otimes \mathbf{d}_c \tag{16}$$

where we note the parameter matrix M that maps the visual context c into the GRU state effectively gates the outer function.[4] The multimodal variant of the LSTM (with peephole connections) is defined as follows:

$$\mathbf{d}_c = M\mathbf{c} \tag{17}$$

$$\mathbf{h}_t = [\mathbf{r}_t \otimes \Phi(\mathbf{c}_t)] \otimes \mathbf{d}_c, \text{ where,} \tag{18}$$

$$\mathbf{r}_t = \sigma(W_r\mathbf{x}_t + V_r\mathbf{h}_{t-1} + U_r\mathbf{c}_t + \mathbf{b}_r) \tag{19}$$

$$\mathbf{c}_t = \mathbf{f}_t \otimes \mathbf{c}_{t-1} + \mathbf{i}_t \otimes \mathbf{z}_t, \text{ where,} \tag{20}$$

$$\mathbf{z}_t = \Phi(W_z\mathbf{x}_t + V_z\mathbf{h}_{t-1} + \mathbf{b}_z), \tag{21}$$

$$\mathbf{i}_t = \sigma(W_i\mathbf{x}_t + V_i\mathbf{h}_{t-1} + U_i\mathbf{c}_{t-1} + \mathbf{b}_i), \tag{22}$$

$$\mathbf{f}_t = \sigma(W_f\mathbf{x}_t + V_f\mathbf{h}_{t-1} + U_f\mathbf{c}_{t-1} + \mathbf{b}_f). \tag{23}$$

All models were trained to minimize the sequence loss of the sentences in the training split. The weight matrices of all models were initialized from uniform distribution, $U(-0.1, 0.1)$, biases were initialized from zero, and the Δ-RNN-specific biases $\{\alpha, \beta_1, \beta_2\}$ were all initialized to one. Parameter updates calculated through back-propagation through time required unrolling the model over 49 steps in time. All symbol sequences were zero-padded and appropriately masked to ensure efficient mini-batching. Gradients were hard-clipped at a magnitude bound of $l = 2.0$. Over mini-batches of 32 samples, model parameters were optimized using simple stochastic gradient descent with a learning rate that starts at $\lambda = 1.0$ and is halved if the perplexity, measured at the end of each epoch, goes up three or more times.

[4] In preliminary experiments, we tried both ways of integrating the visual context information as proposed before, Equations 7 and 8. We ultimately found the second formulation to give better performance.

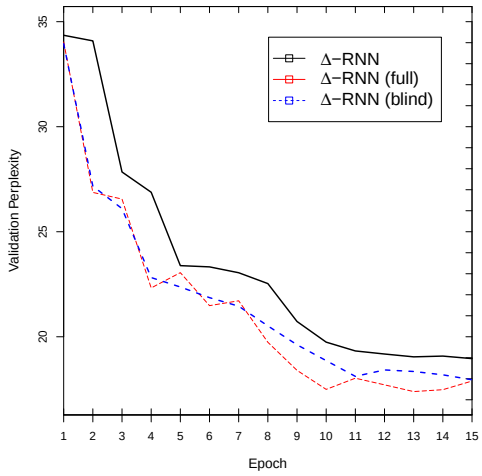

(a) English Δ-RNNs.

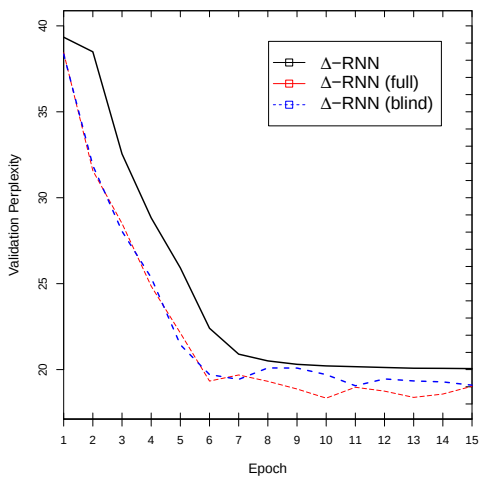

(b) German Δ-RNNs.

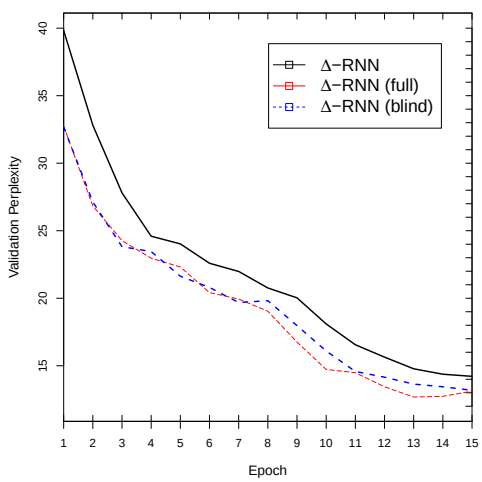

(c) Spanish Δ-RNNs.

Figure 2: Comparison of learning curves for the Δ-RNNs in each language (English, German, Spanish).

To determine if our multimodal language model actually captures knowledge that is different from a text-only language model, we evaluate each model twice. First, we compute the model perplexity on the test set using the sentences' visual context vectors. Next, we compute the model perplexity on the test sentences by feeding in a null-vector to the multimodal model as the visual context. If the model did truly pick up some semantic knowledge that is not exclusively dependent on the conditioned context vector, its perplexity in the second setting, while naturally worse than the first setting, should still outperform the text-only baselines.

In Table 3, we report the model negative log likelihood (NLL) and per-word perplexity (PPL). PPL is a function of NLL, and is simply calculated using the measure:

$$PPL = \exp\Big[-(1/N)\sum_{i=1}^{N}\sum_{t=1}^{T}\log P_\Theta(w_t|\mathbf{h})\Big] \tag{24}$$

We observe that in all cases the multimodal models outperform their respective text-only baselines. More importantly, **the multimodal models, when evaluated without the Inception-v3 representations on held-out samples, still perform better than the text-only baselines**. This improvement in generalization can be attributed to the visual context information given to the model in the training data, enriching its distributed representations over word sequences with knowledge of actual objects as provided by the Inception-v3 vision system. Figure 2 shows the validation perplexity of the various Δ-RNN on each language as a function of the first 15 epochs of learning. We observe that throughout the learning process, the improvement in generalization afforded by the visual context $\mathbf{c}$ is persistent. Validation performance was also tracked for the various GRU and LSTM models, where the same trend was also observed. We provide the plots for those models in the appendix.

4.1 Model Analysis

To further probe the differences between the text-only and multimodal models, we analyze the decoders of each. Specifically, we examine the parameter matrix U, which is directly involved in calculating the logits of the underlying generative model. U can be essentially thought of as "transposed embeddings", an idea that has also been ex-

Table 2: Decoder analysis: Word query similarity test.

Ocean		**Kite**		**Subway**		**Racket**	
Δ-RNN	MM-**Δ-RNN**	**Δ-RNN**	MM-**Δ-RNN**	**Δ-RNN**	MM-**Δ-RNN**	**Δ-RNN**	MM-**Δ-RNN**
surfing	boats	plane	kites	train	railroad	bat	bat
sandy	beach	kites	airplane	passenger	train	batter	players
filled	pier	airplane	plane	railroad	locomotive	catcher	batter
beach	wetsuit	surfboard	airplanes	trains	trains	skateboard	swing
market	cloth	planes	planes	gas	steam	umpire	catcher
crowded	surfing	airplanes	airliner	commuter	gas	soccer	hitter
topped	windsurfing	boats	helicopter	trolley	commuter	women	ball
plays	boardwalk	jet	jets	locomotive	passenger	pedestrians	umpire
cross	flying	aircraft	biplane	steam	crowded	players	tennis
snowy	biplane	jets	jet	it's	trolley	uniform	tatoos

ploited to introduce further regularization into the neural language model learning process (Press and Wolf, 2016; Inan et al., 2016). If we treat each row of this matrix (since we assume column-major orientation in implementation) as the learned embedding for a particular word, we can calculate its similarity to other column embeddings using cosine similarity.

In Table 2, we examine the top ten highest ranked words given several query terms, using the decoder parameter matrix. By observing the different sets of nearest-neighbors produced by the Δ-RNN and the MM-Δ-RNN, we can see that MM-Δ-RNN appears to have learned to combine the visual context information with the token sequence information in its distributed representations. For example, in the case of "ocean", we see that while the Δ-RNN does associate some relevant terms, such as "surfing" and "beach", nearly all of the terms the MM-Δ-RNN associates are relevant to the query. The same situation is observed for "kite" and "subway". In the case of "racket", while the text-only baseline does mostly seem to associate sports terms, especially sports equipment like "bat", the MM-Δ-RNN is actually able to relate the query correctly to the correct sport, "tennis".

4.2 Conditional Sampling

Another interesting way to see how visual context information influences the neural language architecture is to sample from the learned conditional generative model. While image-captioning generally focuses on ranking appropriate caption candidates, we intend to use the model to generate sentences using only the image for guidance. Sampling the learned generative model will allow us to gauge if the system can "explain", in some fashion, what it sees. Table 4.1 lists examples generated by the trained English model. Another sampling approach we implemented is beam search, where, iteratively, m best sentences are picked at time t from a set of generated sentences of length $t+1$. We experimented with a beam of size 13 and Table 34.1 shows the generated captions using this specific beam-search.

5 Discussion and Conclusions

We find that multi-modal neural models trained with a perceptual context are better at modeling language than models trained on language alone. Specifically, we find that augmenting a predictive language model with images that illustrate the sentences being learned enhances the ability of the model to make next-word predictions. This performance improvement persists even in situations devoid of visual representations, when the model is being used as a pure language model.

This research is a step towards taking neural language models more seriously as cognitive and psycholinguistic models of the non-symbolic, implicit aspects of language representation. There's a great deal of evidence that something like a predictive language model exists in the human mind. *Surprisal* is a concept in psycholinguistics that refers to the degree of mismatch between what a human listener expected to be said next and what is actually said, such as when a garden path sentence forces the listener to abandon a partial, incremental parse (Hale, 2001). More generally,

Table 3: Some captions generated by the Multimodal Δ-RNN in English.

	a skateboarder and person in front of skyscrapers. a person with skateboarder on air. a person doing a trick with skateboarder . a person with camera with blue background.
	a food bowl on the table a bowl full of food on the table a green and red bowl on the table a salad bowl with chicken
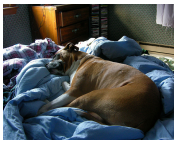	a dog on blue bed with blanket. a dog sleeps near wooden table. a dog sleeps on a bed. a dog on some blue blankets.

the idea of *predictive coding* holds that the mind forms expectations before perception occurs (see Clark, 2013, for review). How these predictions are formed is unclear. Predictive language models trained with a generic neural architecture, without specific linguistic universals, are a reasonable candidate for a model of predictive coding in language. This does not imply neuropsychological realism of the low-level representations or learning algorithms, and we cannot advocate for a specific neural architecture as being most plausible. However, we can show that an architecture that predicts linguistic input well learns better when its input mimics that of a human language learner.

In our (cognitive) view of language processing, we distinguish between symbolic language knowledge and processes that implement compositionality to produce semantics on the one hand, and implicit processes that leverage sequences and associations to produce expectations. With respect to acquiring the latter model, we note that children are exposed to a rich sensory environment, and a more detailed one than the visual environment provided to our language model here. If even static visual input alone improves language acquisition, then what could a sensorily rich environment achieve? When a multimodal learner is considered, then, perhaps, the language acquisition stimulus that has been famously labeled to be rather poor (Chomsky, 1959; Berwick et al., 2013), isn't so poor after all.

One direction for future work is to learn the visual architecture jointly with the language model. Error signals from the language model's back-propagation pathway can prove useful in tuning the multimodal model's ability to fuse information from the linguistic context and the image context. While our current architecture allows us to explore the visual grounding of human language, an architecture trained jointly on vision and language would allow us to also examine the theoretical influence of language on human visual perception.

References

Omri Abend, Tom Kwiatkowski, Nathaniel J. Smith, Sharon Goldwater, and Mark Steedman. 2017. Bootstrapping language acquisition. *Cognition* 164:116 – 143. https://doi.org/10.1016/j.cognition.2017.02.009.

Afra Alishahi, Afsaneh Fazly, and Suzanne Stevenson. 2008. Fast mapping in word learning: What probabilities tell us. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning*. Association for Computational Linguistics, pages 57–64.

Lawrence W Barsalou. 1999. Perceptions of perceptual symbols. *Behavioral and Brain Sciences* 22(4):637–660.

Lawrence W Barsalou. 2008. Grounded cognition. *Annual Review of Psychology* 59:617–645.

Robert C Berwick, Noam Chomsky, and Massimo Piattelli-Palmarini. 2013. Poverty of the stimulus stands: Why recent challenges fail. *Rich languages from poor inputs* pages 19–42.

Irving Biederman. 1987. Recognition-by-components: a theory of human image understanding. *Psychological Review* 94(2):115.

Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* .

Noam Chomsky. 1959. A review of bf skinner's verbal behavior. *Language* 35(1):26–58.

Andy Clark. 2013. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences* 36(3):181–204.

Michael C Frank, Noah D Goodman, and Joshua B Tenenbaum. 2008. A bayesian framework for cross-situational word-learning. In *Advances in neural information processing systems*. pages 457–464.

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*. pages 2121–2129.

James G Greeno and Joyce L Moore. 1993. Situativity and symbols: Response to Vera and Simon. *Cognitive Science* 17(1):49–59.

John Hale. 2001. A probabilistic earley parser as a psycholinguistic model. In *Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies*. Pittsburgh, PA, pages 1–8.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9(8):1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735.

Hakan Inan, Khashayar Khosravi, and Richard Socher. 2016. Tying word vectors and word classifiers: A loss framework for language modeling. *arXiv preprint arXiv:1611.01462* .

Brendan T Johns and Michael N Jones. 2012. Perceptual inference through global lexical similarity. *Topics in Cognitive Science* 4(1):103–120.

Brent Kievit-Kylar and Michael Jones. 2011. The semantic pictionary project. In *Proceedings of the Annual Meeting of the Cognitive Science Society*. volume 33.

Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539* .

Alexander G Ororbia II, Tomas Mikolov, and David Reitter. 2017. Learning simpler language models with the differential state framework. *Neural Computation* .

Ofir Press and Lior Wolf. 2016. Using the output embedding to improve language models. *arXiv preprint arXiv:1608.05859* .

Rajesh PN Rao and Dana H Ballard. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2(1):79.

Deb Roy and Ehud Reiter. 2005. Connecting language to the world. *Artificial Intelligence* 167(1-2):1–12.

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 of Computational Linguistics* 2(1):207–218.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 2818–2826.

Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*. IEEE, pages 3156–3164.

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.

arXiv:1805.11546v1 [cs.CL] 29 May 2018

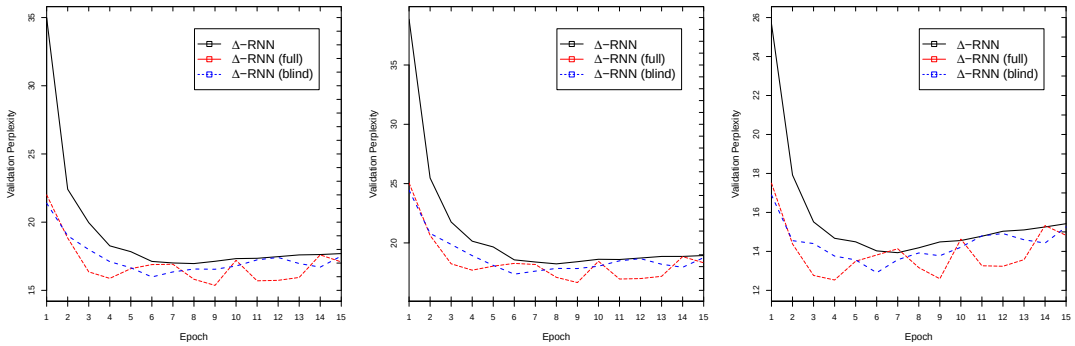

(a) English GRUs. (b) German GRUs. (c) Spanish GRUs.

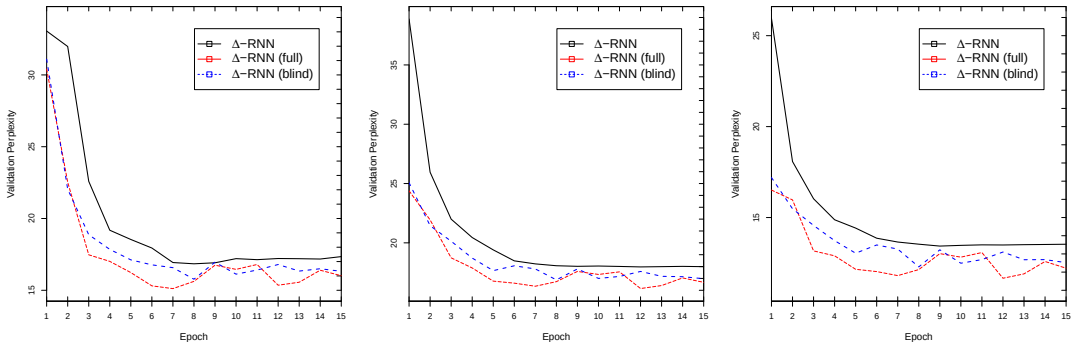

(d) English LSTMs. (e) German LSTMs. (f) Spanish LSTMs.

Figure 1: Appendix: Comparison of learning curves for the GRUs and LSTMs in each language (English, German, Spanish). To augment Figure 2 in the main paper, we also show the learning curves for all models experimented with in this paper beyond the Δ-RNN. Validation learning curves are provided for the GRU and LSTM language models, both multimodal and unimodal variations.