

# Symbolic inductive bias for visually grounded learning of spoken language

Grzegorz Chrupała

Tilburg University

g.chrupala@uvt.nl

## Abstract

A widespread approach to processing spoken language is to first automatically transcribe it into text. An alternative is to use an end-to-end approach: recent works have proposed to learn semantic embeddings of spoken language from images with spoken captions, without an intermediate transcription step. We propose to use multitask learning to exploit existing transcribed speech within the end-to-end setting. We describe a three-task architecture which combines the objectives of matching spoken captions with corresponding images, speech with text, and text with images. We show that the addition of the SPEECH/TEXT task leads to substantial performance improvements on image retrieval when compared to training the SPEECH/IMAGE task in isolation. We conjecture that this is due to a strong inductive bias transcribed speech provides to the model, and offer supporting evidence for this.

## 1 Introduction

Understanding spoken language is one of the key capabilities of intelligent systems which need to interact with humans. Applications include personal assistants, search engines, vehicle navigation systems and many others. The standard approach to understanding spoken language both in industry and in research has been to decompose the problem into two components arranged in a pipeline: Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU). The audio signal representing a spoken utterance is first transcribed into written text, which is subsequently processed to extract some semantic representation of the utterance. Recent works have proposed to learn semantic embeddings of spoken language by using photographic images of everyday situations matched with their spoken captions, without an intermediate transcription step

(Harwath et al., 2016; Chrupała et al., 2017). The weak and noisy supervision in these approaches is closer to how humans learn to understand speech by grounding it in perception and thus more useful as a cognitive model. It can also have some practical advantages: in certain circumstances it may be easier to find or collect speech associated with images rather than transcribed speech – for example when dealing with language whose speakers are illiterate, or for languages with no standard writing system (note that even some languages with many millions of speakers, like Cantonese, may not have a standardized writing system). On the other hand, the learning problem in this type of framework is less constrained, and harder, than standard ASR.

In order to alleviate this shortcoming, we propose to use multitask learning (MTL) and exploit transcribed speech within the end-to-end visually-grounded setting, and thus combine some features of both the pipeline and end-to-end approaches. Incorporating speech transcriptions into the end-to-end architecture via multi-task learning means that the amount of transcribed speech and its quality do not need to be as high as needed for training an ASR system within the pipeline architecture, since the role of this data is only to guide the end-to-end model via an auxiliary task.

We describe a three-task architecture which combines the main objective of matching speech with images with two auxiliary objectives: matching speech with text, and matching text with images.

The plain end-to-end SPEECH/IMAGE matching task, modeled via standard architectures such as recurrent neural networks, lacks a language-specific learning bias. This type of model may discover in the course of learning that speech can be represented as a sequence of symbols (such as for example phonemes or graphemes), but it is in no way predisposed to make this discovery. Hu-

man learners may be more efficient at least in part thanks to their innate inductive bias whereby they assume that language is symbolic. They arguably acquired such bias via the process of evolution by natural selection. In the context of machine learning, inductive bias can instead be injected via multi-task learning, where supervision from the secondary task guides the model towards appropriately biased representations.

Specifically, our motivation for the SPEECH/TEXT task is to encourage the model to learn speech representations which are correlated with the encoding of spoken language as a sequence of characters. Additionally, and for completeness, we also consider a second auxiliary task matching text to images.

Our contribution consists in formulating and answering the following questions:

- Do the auxiliary tasks improve the main SPEECH/IMAGE task? The SPEECH/TEXT task helps but we have no evidence of the TEXT/IMAGE task improving performance.
- If so, is this mainly because MTL allows us to exploit extra data, or because the additional task injects an appropriate inductive bias into the model? The inductive bias is key to the performance gains of MTL, while extra data makes no impact.
- Which parameters should be shared between tasks and which should be task specific? Best performance is achieved by sharing only the lower layers of the speech encoder.
- What are the specific effects of the symbolic inductive bias on the learned representations? SPEECH/TEXT contributes to make the encoded speech more speaker invariant, and more strongly correlated to the written or phonetically represented form of the utterances.

## 2 Related work

### 2.1 Visually grounded semantic embeddings of spoken language

The most relevant strand of related work is on visually-grounded learning of (spoken) language. It dates back at least to [Roy and Pentland \(2002\)](#), but has recently attracted further interest due to better-performing modeling tools based on neural networks.

[Harwath and Glass \(2015\)](#) collect spoken descriptions for the Flickr8K captioned image dataset and present a model which is able to map pre-segmented spoken words to aspects of visual context. [Harwath et al. \(2016\)](#) describe a larger dataset of images paired with spoken captions (Places Audio Caption Corpus) and present an architecture that learns to project images and unsegmented spoken captions to the same embedding space. The sentence representation is obtained by feeding the spectrogram to a convolutional network. Further elaborations on this setting include [Harwath and Glass \(2017\)](#), which shows a clustering-based method to identify grounded words in the speech-image pairs, and [Harwath et al. \(2018b\)](#) which constructs a three-dimensional tensor encoding affinities between image regions and speech segments.

The work of [Chrupała et al. \(2017\)](#) is similar in that it exploits datasets of images with spoken captions, but their grounded speech model is based around multi-layer Recurrent Highway Networks, and focuses on quantitative analyses of the learned representations. They show that the encoding of meaning tends to become richer in higher layers, whereas encoding of form tends to initially increase and then stay constant or decrease. [Al-ishihi et al. \(2017a\)](#) further analyze the representations of the same model and show that phonological form is reliably encoded in the lower recurrent layers of the network but becomes substantially attenuated in the higher layers.

[Drexler and Glass \(2017\)](#) also analyze the representations of a visually grounded speech model with view of using such representations for unsupervised speech recognition, and show that they contain more linguistic and less speaker information than filterbank features.

[Kamper et al. \(2017\)](#) use images as a pivot to learn to associate textual labels with spoken utterances, by mapping utterances and images into joint semantic space. After labeling the images with an object classifier, these labels can be further associated with utterances, providing bag-of-words representation of spoken language which can be useful in speech retrieval.

### 2.2 Multi-task learning for speech and language

The concept of multi-task learning (MTL) was introduced by [Caruana \(1997\)](#). Neural architectures

widely used in the fields of speech and language processing make it easy to define parameter-sharing architectures and exploit MTL, and thus there has been a recent spurt of reports on its impact.

Within Natural Language Processing (NLP), [Luong et al. \(2016\)](#) explore sharing encoders and decoders in a sequence-to-sequence architecture for translation, syntactic parsing, and image captioning, and show gains on some configurations. [Bingel and Søgaard \(2017\)](#) investigate which particular pairs of NLP tasks lead to gains, concluding that learning curves and label entropy of the tasks may be used as predictors. [McCann et al. \(2018\)](#) propose a 10-task NLP challenge, and a single MTL model which performs reasonably well on all tasks.

[Søgaard and Goldberg \(2016\)](#) show that which parameters are shared in a multi-task architecture matters a lot: they find that when sharing parameters between syntactic chunking or supertagging and POS tagging as an auxiliary task, it was consistently better to only share the lower-layers of the model. Relatedly, [Hashimoto et al. \(2017\)](#) propose a method of training NLP tasks at multiple levels of complexity by growing the depth of the model to solve increasingly more difficult tasks. [Swayamdipta et al. \(2018\)](#) use similar ideas and show that syntactic information can be incorporated in a semantic task with MTL, using auxiliary syntactic tasks without building full-fledged syntactic structure at prediction time.

MTL can lead to a bewildering number of choices regarding which tasks to combine, which parameters to share and how to schedule and weight the tasks. Some recent works have suggested specific approaches to deal with this complexity: [Ruder et al. \(2017\)](#) propose to learn from data which parameters to share in MTL with sluice networks and show some gains on NLP tasks. [Kiperwasser and Ballesteros \(2018\)](#) investigate how to interleave learning syntax and translation and how to schedule these tasks.

Several works show that exploiting MTL via the use of multiple language versions of the same or comparable data leads to performance gains (e.g. [Lee et al., 2017](#); [Johnson et al., 2017](#); [de Lhoneux et al., 2018](#)). [Gella et al. \(2017\)](#) and [Kádár et al. \(2018\)](#) learn visual semantic embeddings from textual-visual datasets and show gains from additional languages which reuse the same encoder.

[Kádár et al. \(2018\)](#) additionally show that an extra objective linking the languages directly rather than only via the visual modality provides additional performance gains. In the context of audio-visual data, [Harwath et al. \(2018a\)](#) applies a type of MTL in the setting where there are images paired with descriptions in English and Hindi. They project the images, English speech and Hindi speech into a joint semantic space, and show that training on multiple tasks matching both languages to images works better compared to only using a single monolingual task.

MTL has also recently seen some success in speech processing. Similar to what we see in machine translation, in ASR parameter sharing between different languages is also beneficial ([Heigold et al., 2013](#)). More recently, [Dalmia et al. \(2018\)](#) show that exploiting this effect is especially useful for low-resource languages.

[Seltzer and Droppo \(2013\)](#) apply MTL for phone recognition with three lower-level auxiliary tasks and show noticeable reductions in error rates. [Toshniwal et al. \(2017\)](#) use MTL for conversational speech recognition with lower-level tasks (e.g. phoneme recognition) in an encoder-decoder model for direct character transcription. [Rao and Sak \(2017\)](#) learn to align utterances with phonetic transcriptions in a lower layer and graphemic transcriptions in the final layer, exploiting again the relation between task level of complexity and levels of neural architecture in a MTL setting. They also show a benefit of sharing model parameters between different varieties of the same language, specifically US, British, Indian and Australian English. [McMahan and Rao \(2017\)](#) demonstrate the effectiveness of transfer from generic audio classification to speech command recognition, which can also be considered a particular instance of MTL.

**How our work fits in.** The current paper uses an intuition also present in several of the works mentioned above: namely that an end-to-end model which needs to induce several levels of intermediate latent representations should be guided to find useful ones by including auxiliary prediction tasks at the intermediate layers. These auxiliary prediction tasks typically use lower-level linguistically-motivated structures such as phonemes for end-to-end ASR, or syntactic trees for semantic parsing.

The present study extends this setting to a full speech-to-semantics setup: the main task is to take

spoken language as input and learn a semantic representation based on feedback from the visual modality, while an ASR-like task (SPEECH/TEXT MATCHING) is merely auxiliary. The lower-level linguistic structures in our case are the sequences of phoneme-like units approximated by the written form of the language.

### 3 Methods

#### 3.1 Models

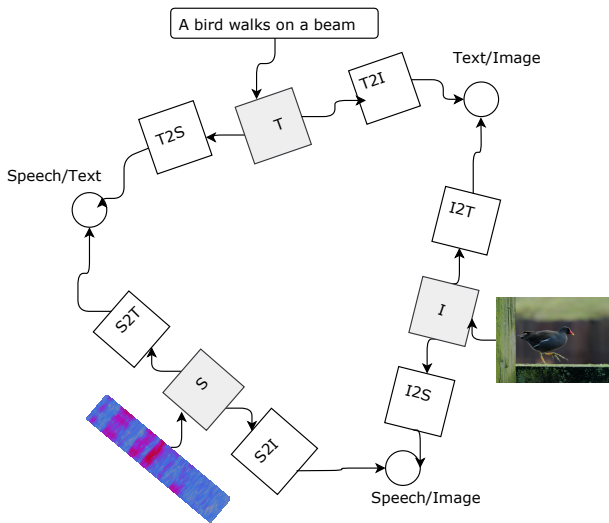


Figure 1: Overview of the task architecture. T: shared text encoder, S: shared speech encoder, I: shared image encoder. The notation X2Y stands for an encoder for input type X which is only used for the loss between encoded input types X and Y.

The modeling framework uses a multi-task setup. The core model is a three-task architecture depicted in Figure 1: there are three encoders, one for each modality: speech, image, and text. Each modality has a shared encoder which works directly on the input modality, and two specialized encoders which take as input the encoded data from the shared encoder. The three tasks correspond to three losses (depicted with circles in the figure): each loss works with a pair of modalities and attempts to minimize the distance between matching encoded items, while maximizing the distance between mismatching ones. For a pair of modalities with encoded objects  $\mathbf{u}$  and  $\mathbf{i}$ , the loss is defined as follows

$$\sum_{\mathbf{u}, \mathbf{i}} \left( \sum_{\mathbf{u}'} \max[0, \alpha + d(\mathbf{u}, \mathbf{i}) - d(\mathbf{u}', \mathbf{i})] + \sum_{\mathbf{i}'} \max[0, \alpha + d(\mathbf{u}, \mathbf{i}) - d(\mathbf{u}, \mathbf{i}')] \right) \quad (1)$$

where  $(\mathbf{u}, \mathbf{i})$  are matching objects (for example an utterance and a matching image), and  $(\mathbf{u}', \mathbf{i})$  and  $(\mathbf{u}, \mathbf{i}')$  are mismatched objects within a batch, while  $d(\cdot, \cdot)$  is the cosine distance between encoded objects.

The SPEECH/IMAGE part of the architecture is based on the grounded speech model from Chrupała et al. (2017), with the main difference being that these authors used Recurrent Highway Networks (Zilly et al., 2017) for the recurrent layers, while we chose the simpler Gated Recurrent Unit networks (Chung et al., 2014), because they have optimized low-level CUDA support which makes them much faster to run and enables us to carry out an at least somewhat comprehensive set of experiments.

#### 3.2 Image Encoders

The shared image encoder I is a pretrained, fixed Convolutional Neural Network which outputs a vector with image features; specifically, the activations of the pre-classification layer. The modality-specific encoders I2S and I2T are linear mappings which take the output of I.

#### 3.3 Speech Encoders

The shared encoder S consists of a 1-dimensional convolutional layer which subsamples the input, followed by a stack of recurrent layers. The modality specific encoders S2T and S2I consist of a stack of recurrent layers, followed by an attention operator. The encoder S is defined as follows:

$$S(\mathbf{x}) = \text{GRU}_\ell(\text{Conv}_{s,d,z}(\mathbf{x})) \quad (2)$$

where Conv is a convolutional layer with kernel size  $s$ ,  $d$  channels, and stride  $z$ , and  $\text{GRU}_\ell$  is a stack of  $\ell$  GRU layers. An encoder of modality X is defined as

$$S2X(\mathbf{x}) = \text{unit}(\text{Attn}(\text{GRU}_\ell(\mathbf{x}))) \quad (3)$$

where Attn is the attention operator and unit is L2-normalization. Note that for the case  $\ell = 0$

$\text{GRU}_\ell$  is simply the identity function. The attention operator computes a weighted sum of the RNN activations at all timesteps:

$$\text{Attn}(\mathbf{x}) = \sum_t \alpha_t \mathbf{x}_t \quad (4)$$

where the weights  $\alpha_t$  are determined by an MLP with learned parameters  $\mathbf{U}$  and  $\mathbf{W}$ , and passed through the timewise softmax function:

$$\alpha_t = \frac{\exp(\mathbf{U} \tanh(\mathbf{W} \mathbf{x}_t))}{\sum_{t'} \exp(\mathbf{U} \tanh(\mathbf{W} \mathbf{x}_{t'}))} \quad (5)$$

### 3.4 Text Encoders

The text encoders are defined in the same way as the speech encoders, with the only difference being that the convolutional layer is replaced by an embedding layer, i.e. a lookup table mapping characters to embedding vectors.

### 3.5 Multi-tasking

The model is trained by alternating between the tasks, and updating the parameters of each task in turn. Note that the input data for the three tasks can be the same, but can also be partly or completely disjoint. We report two conditions

- **ALIGNED:** all tasks use the same parallel data;
- **NON-ALIGNED:** the data for the SPEECH/TEXT task is disjoint from the data for the other two tasks.

We consider the NON-ALIGNED condition somewhat more realistic, in that it is easier to find separate datasets for each pair of modalities than it is to find a single dataset with all three modalities. However the main reason to including both conditions is that it allows us to disentangle via which mechanism MTL contributes: by enabling the use of extra data, or by enforcing an inductive bias.

### 3.6 Architecture variants

There is a multitude of ways in which the details of the core architecture can be varied. In order to reduce them to a manageable number we made the following choices:

- Keep the image encoder simple and fixed.
- Keep the architecture of the encoders fixed, and only vary encoder depth and the degree of sharing.

In addition to variants of the full three-task model, we also have single-task and two-task baselines which are the three-task model with the SPEECH/TEXT and TEXT/IMAGE tasks completely ablated, or with only the TEXT/IMAGE task ablated. Note that we do not include a condition with only the SPEECH/TEXT task ablated, as the two remaining tasks do not share any learnable parameters (since  $I$  is fixed).

### 3.7 Evaluation metrics

Below we introduce metrics evaluating performance on the image retrieval task, as well as additional analytical metrics which quantify some aspects of the internal representation learned by the encoders.

**Evaluating image retrieval** In order to evaluate how well the main SPEECH/IMAGE task performs we report the recall at 10 ( $R@10$ ) and median rank ( $\text{Medr}$ ) for the SPEECH/IMAGE task: utterances in the development set are encoded via S2I and images via I2S. For each utterance the images are ranked in order of cosine distance;  $R@10$  counts the mean proportion of correct images among top 10 ranked images, while  $\text{Medr}$  gives the median of the ranks of the correct image (where correct image counts as image originally paired with the utterance).

**Invariance to speaker** We measure how invariant the utterance encoding is to the identity of the speaker; in principle it is expected and desirable that the utterance encoding captures the meaning of the spoken language rather than other aspects of it such as who spoke it. To quantify this invariance we report the accuracy of an L2-penalized logistic regression model on the task of decoding the identity of the speaker from the output of the S2I encoder. The logistic model is trained on  $\frac{2}{3}$  of the development data and tested on the remaining  $\frac{1}{3}$ .

**Representational similarity** Representational Similarity Analysis (Kriegeskorte et al., 2008) gauges the correlation between two sets of pairwise similarity measurements. Here we use it to quantify the correlation of the learned representation space with the written text space and with the image space. For the encoder representations, the pairwise similarities between utterances are given by the cosine similarities. For the written form, the similarities are the inverse of the normalized Levenshtein distance between

the character sequences encoding each pair of utterances:

$$\text{sim}_{\text{text}}(a, b) = 1 - \frac{D(a, b)}{\max(|a|, |b|)} \quad (6)$$

where  $D(a, b)$  is the Levenshtein distance and  $|\cdot|$  is string length. We compute the Pearson correlation coefficient between two similarity matrices on the upper triangulars of the each matrix, excluding the diagonal.

**Phoneme decoding** A direct way of measuring whether neural representations of speech are biased towards encoding symbols is to try to decode the phonemes from the activation patterns aligned with a phonetic transcription of the utterance. We follow the methodology of [Alishahi et al. \(2017b\)](#) and train an L2-penalized logistic regression model on the output of the S encoders for phonemes from 2,500 utterances and report classification accuracies on data from 2,500 heldout utterances.

### 3.8 Experimental settings

**Data** The SPEECH/IMAGE and TEXT/IMAGE tasks are always trained on the Flickr8K Audio Caption Corpus ([Harwath et al., 2016](#)), which is based on the original Flickr8K dataset ([Hodosh et al., 2013](#)). Flickr8K consists of 8,000 photographic images depicting everyday situations. Each image is accompanied by five brief English descriptions produced by crowd workers. Flickr8K Audio Caption Corpus enriches this data with spoken versions of these descriptions, read aloud and recorded by crowd workers. The total amount of speech in this dataset is approximately 34 hours. One thousand images are held out for validation, and another one thousand for the test set, using the splits provided by [Karpathy and Fei-Fei \(2015\)](#). In the ALIGNED condition the SPEECH/TEXT task is also trained on this data. In the NON-ALIGNED condition, we train the SPEECH/TEXT task on the Libri dataset ([Panayotov et al., 2015](#)) which consists of approximately 1,000 hours of read English speech, derived from read audiobooks. There are 291,630 sentences in the corpus, of which 1,000 are held out for validation.

We preprocess the audio by extracting 12-dimensional mel-frequency cepstral coefficients (MFCC) plus log of the total energy. We use 25 millisecond windows, sampled every 10 milliseconds. The shared image encoder is fixed and

consists of 4096 dimensional activations of the pre-classification layer of VGG-16 ([Simonyan and Zisserman, 2014](#)) pre-trained on Imagenet ([Russakovsky et al., 2015](#)).

**Hyperparameters** Most of the hyperparameters are based from especially [Chrupala et al. \(2017\)](#). The models are trained for a maximum of 25 epochs with Adam, with learning rate 0.0002, and gradient clipping at 2.0. The loss function’s margin parameter is  $\alpha = 0.2$ . The GRUs have 1024 dimensions. The convolutional layer has 64 channels, kernel size of 6 and stride 2. The hidden layer of the attention MLP is 128. The linear mappings I2S and I2T project 4096 dimensions down to 1024. We apply early stopping and pick the results of each run after the epoch for which it scored best on R@10. We run three random initializations of each configuration.

**Multi-task training** We use a simple round-robin training scheme: we alternate between tasks, and for each task update the parameters of that task as well as the shared parameters based on supervision from one batch of data. The data ordering for each task is independent, both in the ALIGNED and NON-ALIGNED condition: for each epoch we reshuffle the dataset associated to each task and iterate through the batches until the smallest dataset runs out. This procedure makes sure that the only difference between the ALIGNED and NON-ALIGNED conditions is the actual data and not other aspects of training.

**Repository** The code needed to reproduce our results and analyses is available at <https://github.com/gchrupala/symbolic-bias>.

## 4 Results

Table 1 shows the evaluation results on the validation data, on the image retrieval task of 13 configurations of the model, including three versions with one or two tasks ablated.

Table 2 shows the results on the test set with the 1-task baseline model and the best performing configuration compared to previously reported results on this dataset. As can be seen the baseline model is a bit worse than the best reported result on this data, while the 3-task model is much better.

## 5 Discussion

Below we discuss and interpret the patterns in performance on image retrieval as measured by Re-

	Data	Tasks	s	T	s2I	s2T	T2S	T2I	R@10	Medr
1	NA	1	2	.	2	.	.	.	0.218	63.8
2	Aligned	2	2	1	2	0	0	.	0.279	42.3
3	Non-aligned	2	2	1	2	0	0	.	0.280	41.3
4			2	1	1	0	0	1	0.280	43.0
5			2	1	1	1	1	1	0.266	44.3
6	Aligned	3	2	1	2	0	0	1	<b>0.281</b>	<b>39.7</b>
7			2	1	2	1	1	1	0.270	44.3
8			4	1	0	0	0	0	0.255	48.3
9			2	1	1	0	0	1	0.275	42.8
10			2	1	1	1	1	1	0.257	49.8
11	Non-aligned	3	2	1	2	0	0	1	0.280	41.7
12			2	1	2	1	1	1	0.252	50.7
13			4	1	0	0	0	0	0.223	59.3

Table 1: Results on the validation set with varying model configuration. R@10 is recall at 10 for the Speech/Image task, Medr is the median rank for the same task. All scores are averages over 3 runs with different random initializations; models were run for 25 epochs with early stopping with R@10 as a criterion. The numbers (1, 2) in the columns corresponding to encoders specify the number of RNN layers in each encoder; zero (0) indicates the encoder only consists of the self-attention with no RNN layers; dot (.) indicates the whole task in which the encoder participates is ablated.

Data	Tasks	s	T	s2I	s2T	T2S	T2I	R@10	Medr
NA	1	Harwath and Glass (2015)						0.179	-
NA	1	Chrupała et al. (2017)						0.253	48
NA	1	2	.	2	.	.	.	0.244	51
Aligned	3	2	1	2	0	0	1	0.296	34

Table 2: Results on the test set, obtained by using the best run/epoch determined on the validation data. The first two rows show the numbers reported in previous work.

call@10 and median rank.

**Impact of tasks** The most striking result is the large gap in performance between the 1-task condition (row 1) and most of the other rows. Comparing row 1 versus rows 2 and 3 we see that adding the SPEECH/TEXT task leads to a substantial improvement. However, comparing rows 2 and 3 versus rows 6 and 11, it seems that the addition of TEXT/IMAGE task does not seem to have a major impact on performance, at least to the extent that can be gleaned from the experiments we carried out. It is possible that with more effort put into engineering this component of the model we would see a better result.

**Role of data vs inductive bias** The other major finding is that whether we use the same or different data to train the main and auxiliary task has overall little impact: this is indicated by relatively small differences between configurations in the

ALIGNED vs NON-ALIGNED condition. The differences that are there tend to favor the ALIGNED setting. This lends supports to the conclusion that the SPEECH/TEXT auxiliary task contributes to improved performance on the main task via a strong inductive bias rather than merely via enabling the use of extra data. This is in contrast to many other applications of MTL.

**Impact of parameter sharing design** The third important effect is about how parameters between the tasks are shared, specifically how the shared and task-specific parts of the speech encoder are apportioned. The configuration with maximum sharing of parameters among the tasks (rows 8 and 13) performs poorly compared to sharing only the lower layers of the encoders for speech and text (i.e. rows 6 and 11). Additionally, we see that the inclusion of a text-specific speech encoder s2T degrades performance: compare for exam-

ple row 6 to 7, and row 11 to 12. Thus it is best to have a shared speech encoder whose output is directly used by the SPEECH/TEXT task, while the SPEECH/IMAGE task carries out further transformations of the input via an image-specific speech encoder S2I. We can interpret this as the MTL emulating a pipeline architecture to some extent: direct connection of the SPEECH/TEXT task to the shared encoder forces it to come up with a representation closely correlated with a written transcription, and then the image-specific speech encoder takes this as input and maps it to a more meaning-related representation, useful for the SPEECH/IMAGE task.

In addition to the above patterns of performance on image retrieval we now address our further research questions by investigating selected aspects of the encoder representations.

**Speaker invariance** Table 3 shows the accuracy of speaker identification from the activation patterns of the output of encoder S2I for the single task model, the 2-task model, and for the 3-task model which achieved the highest recall@10. The accuracy of the 2 task model is almost three times worse than for the single task model, indicating that the inclusion of SPEECH/TEXT strongly drives the learned representations towards speaker invariance. The TEXT/IMAGE task has only a minor impact.

Model	Accuracy
Model 1, S2I	0.297
Model 2, S2I	0.101
Model 6, S2I	0.085

Table 3: Speaker identification accuracy for three model configurations. Model numbers refer to rows in Table 1.

### RSA with regard to textual and visual spaces

Table 4 shows the RSA scores between the encoder representations of utterances and their representations in the spoken, written and visual modalities. Comparing the RSA scores between the S2I encoder of model 1 (single task) and model 6 (3 task) we see that the correlations with the textual modality and the visual modality are enhanced while the correlation with the input audio modality drops. This can be interpreted as the SPEECH/TEXT task nudging the model to align more closely with the text, which also ends up

	MFCC	Text	Image
Model 1, S2I	0.043	0.194	0.187
Model 6, S2I	0.030	0.212	0.222
Model 6, S2T	0.099	0.243	0.105
Image	0.008	0.083	1.000

Table 4: Pearson correlation between pairwise utterance similarity matrices, for utterances represented by Mean MFCC features, written text, three encoders, and the features of the image corresponding to the utterance. Model numbers refer to rows in Table 1. Analysis carried out on the single best seed/epoch for each configuration, according to Recall@10.

contributing to the correlation with the image space. For model 6 but using the output of the S2T encoder, we see the correlation with the text space is even higher while the correlation with the image space is low. These patterns are what we would expect if SPEECH/TEXT does indeed inject a symbolic inductive bias to the model. Finally, while the RSA score between the textual and visual modalities is low (0.083), nevertheless model 6’s encoder S2I is moderately correlated with both of these (0.212 and 0.222 respectively).

**Phoneme decoding** Table 5 shows how well phonemes can be decoded from time-aligned slices of four types of representations: input MFCC features, the activation patterns of a randomly initialized s encoder, and the activations of the s encoder for two trained models (1-task and 3-task). Phonemes are most decodable from the 3-task activation patterns, corroborating that the SPEECH/TEXT task biases the representations towards a symbolic encoding of speech.

Representation	Accuracy
MFCC	0.284
Random init, s	0.486
Model 1, s	0.528
Model 6, s	0.578

Table 5: Phoneme decoding accuracy for the four representations. Model numbers refer to rows in Table 1.

## 6 Conclusion

We show that the SPEECH/TEXT task leads to substantial performance improvements when compared to training the SPEECH/IMAGE task in isolation. Via controlled experiments and analyses we



show evidence that this is due to the role of inductive bias on the learned encoder representations.

**Limitations and future work** Our current model does not include an explicit speech-to-text decoder, which limits the types of analyses we can perform. For one, it makes it infeasible to carry out an apples-to-apples comparison with a pipeline architecture. Going forward we would like to go beyond matching tasks and evaluate the impact of an explicit speech-to-text decoder as an auxiliary task.

We are also planning to investigate how sensitive our approach is to amount of data for the auxiliary task. This would be especially interesting given that one motivation for a visually-supervised end-to-end approach is the un-availability of large amounts of transcribed speech in certain circumstances.

## Acknowledgements

I would like to thank Afra Alishahi, Lieke Gelderloos and Ákos Kádár, as well as several anonymous reviewers, for helpful comments and discussion about this work.

## References

- Afra Alishahi, Marie Barking, and Grzegorz Chrupała. 2017a. [Encoding of phonology in a recurrent neural model of grounded speech](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 368–378. Association for Computational Linguistics.
- Afra Alishahi, Marie Barking, and Grzegorz Chrupała. 2017b. [Encoding of phonology in a recurrent neural model of grounded speech](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 368–378.
- Joachim Bingel and Anders Søgaard. 2017. [Identifying beneficial task relations for multi-task learning in deep neural networks](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 164–169. Association for Computational Linguistics.
- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Grzegorz Chrupała, Lieke Gelderloos, and Afra Alishahi. 2017. [Representations of language in a model of visually grounded speech signal](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. [Empirical evaluation of gated recurrent neural networks on sequence modeling](#). In *NIPS 2014 Deep Learning and Representation Learning Workshop*.
- Siddharth Dalmia, Ramon Sanabria, Florian Metze, and Alan W. Black. 2018. [Sequence-based multilingual low resource speech recognition](#). *CoRR*, abs/1802.07420.
- Jennifer Drexler and James Glass. 2017. [Analysis of audio-visual features for unsupervised speech recognition](#). In *Proceedings of Grounded Language Understanding Workshop*.
- Spandana Gella, Rico Sennrich, Frank Keller, and Mirella Lapata. 2017. [Image pivoting for learning multilingual multimodal representations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2839–2845.
- David Harwath, Galen Chuang, and James Glass. 2018a. [Vision as an interlingua: Learning multilingual semantic embeddings of untranscribed speech](#). In *Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- David Harwath and James Glass. 2015. [Deep multimodal semantic embeddings for speech and images](#). In *IEEE Automatic Speech Recognition and Understanding Workshop*.
- David Harwath and James Glass. 2017. [Learning word-like units from joint audio-visual analysis](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 506–517. Association for Computational Linguistics.
- David Harwath, Adrià Recasens, Dídac Surís, Galen Chuang, Antonio Torralba, and James Glass. 2018b. [Jointly discovering visual objects and spoken words from raw sensory input](#). *arXiv preprint arXiv:1804.01452*.
- David Harwath, Antonio Torralba, and James Glass. 2016. [Unsupervised learning of spoken language with visual context](#). In *Advances in Neural Information Processing Systems*, pages 1858–1866.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. [A joint many-task model: Growing a neural network for multiple nlp tasks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1923–1933. Association for Computational Linguistics.
- Georg Heigold, Vincent Vanhoucke, Andrew Senior, Patrick Nguyen, Marc’Aurelio Ranzato, Matthieu Devin, and Jeff Dean. 2013. [Multilingual acoustic models using distributed deep neural networks](#). In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8619–8623.

- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Ákos Kádár, Desmond Elliott, Marc-Alexandre Côté, Grzegorz Chrupała, and Afra Alishahi. 2018. Lessons learned in multilingual grounded language learning. In *Proceedings of the 22nd Conference on Computational Natural Language Learning (CoNLL 2018)*.
- Herman Kamper, Shane Settle, Gregory Shakhnarovich, and Karen Livescu. 2017. [Visually grounded learning of keyword prediction from untranscribed speech](#). In *Proc. Interspeech 2017*, pages 3677–3681.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137.
- Eliyahu Kiperwasser and Miguel Ballesteros. 2018. Scheduled multi-task learning: From syntax to translation. *Transactions of the Association of Computational Linguistics*, 6:225–240.
- Simon Kirby, Mike Dowman, and Thomas L. Griffiths. 2007. [Innateness and culture in the evolution of language](#). *Proceedings of the National Academy of Sciences*, 104(12):5241–5245.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Biedtner. 2008. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. Fully character-level neural machine translation without explicit segmentation. *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Miryam de Lhoneux, Johannes Bjerva, Isabelle Augenstein, and Anders Søgaard. 2018. Parameter sharing between dependency parsers for related languages. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task sequence to sequence learning. In *International Conference on Learning Representations*.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. The natural language Decathlon: Multitask learning as question answering. *arXiv preprint arXiv:1806.08730*.
- Brian McMahan and Delip Rao. 2017. [Listening to the world improves speech command recognition](#). *CoRR*, abs/1710.08377.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An ASR corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Kanishka Rao and Haşim Sak. 2017. [Multi-accent speech recognition with hierarchical grapheme based models](#). In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4815–4819.
- Deb K Roy and Alex P Pentland. 2002. [Learning words from sights and sounds: a computational model](#). *Cognitive Science*, 26(1):113 – 146.
- Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. 2017. Learning what to share between loosely related tasks. *arXiv preprint arXiv:1705.08142*.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. [Imagenet large scale visual recognition challenge](#). *International Journal of Computer Vision*, 115(3):211–252.
- Michael L. Seltzer and Jasha Droppo. 2013. [Multi-task learning in deep neural networks for improved phoneme recognition](#). In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6965–6969.
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556.
- Anders Søgaard and Yoav Goldberg. 2016. [Deep multi-task learning with low level tasks supervised at lower layers](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 231–235. Association for Computational Linguistics.
- Swabha Swayamdipta, Sam Thomson, Kenton Lee, Luke S. Zettlemoyer, Chris Dyer, and Noah A. Smith. 2018. Syntactic scaffolds for semantic structures. *CoRR*, abs/1808.10485.
- Shubham Toshniwal, Hao Tang, Liang Lu, and Karen Livescu. 2017. Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition. In *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, pages 3532–3536.

Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. 2017. [Recurrent highway networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 4189–4198, International Convention Centre, Sydney, Australia. PMLR.