MULTIMODAL ONE-SHOT LEARNING OF SPEECH AND IMAGES

Ryan Eloff Herman A. Engelbrecht

Herman Kamper

E&E Engineering, Stellenbosch University, South Africa

rpeloff@sun.ac.za, hebrecht@sun.ac.za, kamperh@sun.ac.za

ABSTRACT

Imagine a robot is shown new concepts visually together with spoken tags, e.g. "milk", "eggs", "butter". After seeing *one* paired audiovisual example per class, it is shown a new set of unseen instances of these objects, and asked to pick the "milk". Without receiving any hard labels, could it learn to match the new continuous speech input to the correct visual instance? Although unimodal one-shot learning has been studied, where one labelled example in a single modality is given per class, this example motivates *multimodal one-shot learning*. Our main contribution is to formally define this task, and to propose several baseline and advanced models. We use a dataset of paired spoken and visual digits to specifically investigate recent advances in Siamese convolutional neural networks. Our best Siamese model achieves twice the accuracy of a nearest neighbour model using pixel-distance over images and dynamic time warping over speech in 11-way cross-modal matching.

Index Terms— Multimodal modelling, one-shot learning, crossmodal matching, low-resource speech processing, word acquisition.

1. INTRODUCTION

Humans possess the remarkable ability to learn new words and object categories from only one or a few examples. For example, a child hearing the word "lego" for the first time in the context of receiving a new toy, can quickly learn to associate the spoken word "lego" to the new (visual) concept *lego*. Current state-of-the-art speech and vision processing algorithms require thousands of labelled examples to complete a similar task. This has lead to research in *one-shot learning* [1–3], where the task is acquisition of novel concepts from only one or a few labelled examples.

One-shot learning studies have focused on problems where novel concepts in a single modality are observed along with class labels. This is different from the example above: the child directly associates the spoken word "lego" to the visual signal of lego without any class labels, and can generalise this single example to other visual or spoken instances of lego. This motivates multimodal one-shot learning, a new task we formalise in this paper. Consider an agent such as a household robot that is shown a single visual example of milk, eggs, butter and a mug, novel objects each paired with a spoken description. During subsequent use, a speech query is given and the agent needs to identify which visual object the query refers to. This setting is relevant in modelling infant language acquisition, where models can be used to test particular cognitive hypotheses [4]; lowresource speech processing, where new concepts could be taught in an arbitrary language [5]; and robotics, where novel concepts must be acquired online from co-occurring multimodal sensory inputs [6-10].

Here we specifically consider multimodal one-shot learning on a dataset of isolated spoken digits paired with images. A model is shown a set of speech-image pairs, one for each of the 11 digit classes. We refer to this set, which is acquired before the model is applied, as the *support set*. During testing, the model is shown a new instance of a spoken digit and a new set of test images, called the *matching set*. It then needs to predict which test image in the matching set corresponds to the spoken input query. We tackle this problem by extending existing unimodal one-shot models to the multimodal case.

To do cross-modal test-time matching, we propose a framework relying on unimodal comparisons through the support set. Given an input speech query, we find the closest speech segment in the support set. We then take its paired support image, and find its closest image in the matching set. This image is predicted as the match. Metrics for speech-speech and image-image comparisons need to be defined, and this is where we take advantage of the large body of work in unimodal one-shot learning to investigate several options. One approach is to use labelled background training data not containing any of the classes under consideration. Using such speech and image background data, we specifically investigate Siamese neural networks [11, 12] as a way to explicitly train unimodal distance metrics. We also incorporate recently proposed advances [3, 13–18] for such networks.

We compare novel Siamese convolutional neural network (CNN) architectures to traditional direct feature matching models. We show that a single CNN with a triplet loss [13, 14] and online semi-hard mining [17] is more efficient and results in higher accuracies than the offline variant which uses shared weight networks, both approaches outperforming the direct feature matching baseline. Our main contribution is the formal definition of multimodal one-shot learning. We also develop a one-shot cross-modal matching dataset that may be used to benchmark other approaches. As an intermediate evaluation in our work, we also consider unimodal one-shot speech classification. Apart from [19], our paper is to our knowledge the only work that considers one-shot unimodal learning of spoken language. We present several new models and baselines not considered in [19].

2. RELATED WORK

Most one-shot learning studies have been primarily interested in image classification. Apart from Siamese models [3], which we focus on here, other metric learning based approaches have been proposed [20–22] which build on advances in attention and memory mechanisms for neural networks. Along with the more recent meta-learning approaches [23–25], these have each produced improvements in one-shot image classification. However, only small improvements have been made over Siamese networks: Finn *et al.* [24] achieved state-of-the-art results with only 1.4% increase in accuracy over Siamese networks for a 5-way one-shot learning task.

A limited number of studies have considered other domains, such as robotics [6,26], video [27], and gesture recognition [8,28]. Lake *et al.* [19] investigated one-shot speech learning using a generative hierarchical hidden Markov model to recognise novel words from learned primitives. This Bayesian model is based on prior work [29]



Fig. 1. (a) Unimodal one-shot speech learning and classification and (b) multimodal one-shot learning and matching of speech and images.

in vision and has displayed strong results, although on visual tasks the Siamese model seems superior [3]. Our work is also related to learning multimodal representations from paired images and unlabelled speech [30–34]. We extend this research to the one-shot domain. shot learning, the support set $S = \{(\mathbf{x}_a^{(i)}, \mathbf{x}_v^{(i)})\}_{i=1}^{L \times K}$ consists of K speech-image example pairs for each of the L classes. This would occur, for instance, when a user teaches a robot speech-image correspondences by presenting it with multiple paired examples per class.

3. MULTIMODAL ONE-SHOT LEARNING

The goal of unimodal one-shot learning is to build a model that can acquire new concepts after observing only a single labelled example from each class. This model must then successfully generalise to new instances of those concepts in tasks such as classification or regression. Formally, a model is shown a support set S, containing one labelled example for each of L classes. From this set, it must learn a classifier C_S for unseen queries \hat{x} . This is illustrated in Figure 1(a) for five-way one-shot speech classification. In this case the support set S contains spoken utterances along with hard textual labels. The model uses this information to classify the spoken test query "two" as the concept label *two*. Note that the test-time query does not occur in the support set.

We now modify this scheme to fit multimodal one-shot learning. Instead of a labelled unimodal support set, we are now given features in multiple modalities with the only supervisory signal being that these features co-occur. In our case we consider speech and images as the two modalities, although this may be applied to any source of paired multi-sensory information. Formally, in an L-way one-shot problem we are given a multimodal support set $S = \{(\mathbf{x}_a^{(i)}, \mathbf{x}_v^{(i)})\}_{i=1}^L$, where each spoken caption $\mathbf{x}_a^{(i)} \in \mathcal{A}$ (audio space) is paired with an image $\mathbf{x}_{v}^{(i)} \in \mathcal{V}$ (vision space). During testtime, the modal is presented with a test query in one modality, and asked to determine the matching item in a test (or matching) set in the other modality. This is similar to the information retrieval task used by [32] to evaluate a multimodal network. Formally, we match query $\hat{\mathbf{x}}_a$ in one modality (speech) to a matching set $\mathcal{M}_v = \{(\hat{\mathbf{x}}_v)\}_{i=1}^N$ in the other modality (images) according to some metric $D_{\mathcal{S}}(\hat{\mathbf{x}}_a, \hat{\mathbf{x}}_v)$ learned from the support set S. Neither the query $\hat{\mathbf{x}}_a$ or the items in the match set \mathcal{M}_v occur exactly in the support set \mathcal{S} . We refer to this task as one-shot cross-modal matching. It is illustrated in Figure 1(b), where the spoken query "two" is most similar to the image of a two in the matching set according to the model trained on the support set.

One-shot learning can be generalised to K-shot learning, where, in the unimodal case, a model is shown a support set containing Lnovel classes and K examples per class. In multimodal L-way K-

4. MULTIMODAL ONE-SHOT MODELLING

We now turn to the general framework we use to perform multimodal one-shot learning. Assume we have a method or model that can measure similarity within a modality. One-shot cross-modal matching is then accomplished by first comparing a query to all the items in the support set in the query modality (e.g. speech). The most similar (speech-image) support-set pair is retrieved. Finally, the retrieved support-set instance is used to determine the closest item in the matching set, based on comparisons in the matching-set modality (e.g. images). This approach thus defines a metric D_S as a mapping $\mathcal{A} \to \mathcal{V}$ which can match speech to images by unimodal comparisons through the multimodal support set S. As a concrete example, in Figure 1(b) the speech query in (b)-ii is compared to all the support set speech segments in (b)-i, and the closest speech item determined. The corresponding support-set image of this item is then taken and compared to all the images in the matching set \mathcal{M}_v in (b)-ii. The closest image is predicted as the match.

Several different methods or models can be used to determine within-modality similarity: we compare directly using the raw image pixels and extracted speech features ($\S4.1$) to feature embeddings learned by neural network functions ($\S4.2$ and $\S4.3$).

4.1. Direct feature matching

Our first approach consists of directly using image pixels and acoustic speech features. We specifically use cosine similarity between image pixels and dynamic time warping (DTW) [35] to measure similarity between speech segments. This is essentially our direct nearest neighbour baseline, as used in unimodal one-shot studies [3, 19, 20].

4.2. Neural network classifiers

Another method, also used in unimodal one-shot learning, is to train a supervised model on a large background dataset. This background dataset should not contain instances of the target one-shot classes. The idea is that features learned by such a model would still be useful for determining similarity on classes which it has not seen [20]. I.e., it follows the *transfer learning* principle: first train a classifier on a large labelled dataset, and then apply the learned representations to new tasks which are related but have too few training instances [36].

Here we specifically train both a feedforward (FFNN) and a convolutional neural network (CNN) classifier. We train such networks separately for both the spoken and visual modalities. We then take representations from the final hidden layer (before the softmax output) and apply these as the learned feature embeddings for nearest neighbour matching using cosine similarity.

4.3. Siamese neural networks

Networks can also be trained on background data to directly measure similarity between inputs instead of predicting a class label. Siamese neural networks have been used for this task [11, 12]. Two identical neural network functions with shared parameters (hence "Siamese") are trained to map input features to a target embedding space where the "semantic" relationship between input pairs may be captured based on proximity: inputs of the same type should ideally be mapped to similar embeddings, while inputs that are unrelated should be far apart. Early approaches [3, 11, 12] took in pairs of training examples and either maximised or minimised a distance based on whether the inputs were of same or different types. Recent studies [13–17] have argued that the *relative* rather than *absolute* distance between embeddings are meaningful, and we also follow this approach here.

Concretely, let \mathbf{x}_a and \mathbf{x}_p be inputs of the same class, while \mathbf{x}_a and \mathbf{x}_n are of different classes. The intuition is that we want to push the so-called anchor example \mathbf{x}_a and positive example \mathbf{x}_p together such that the distance between them is smaller (by some specified margin) than the distance between the anchor \mathbf{x}_a and negative example \mathbf{x}_n . Models using this approach are sometimes referred to as triplet models, since there are three tied networks for inputs ($\mathbf{x}_a, \mathbf{x}_p, \mathbf{x}_n$). We apply this approach where we learn embedding function $f(\cdot)$ as the final fully-connected layer of a CNN. We train models with a hinge loss for triplet pairs [13–15, 17], defined as:

$$l(\mathbf{x}_a, \mathbf{x}_p, \mathbf{x}_n) = \max\{0, m + D(\mathbf{x}_a, \mathbf{x}_p) - D(\mathbf{x}_a, \mathbf{x}_n)\}$$
(1)

where $D(\mathbf{x}_1, \mathbf{x}_2) = ||f(\mathbf{x}_1) - f(\mathbf{x}_2)||_2^2$ is the squared Euclidean distance and *m* is the margin between the pairs $(\mathbf{x}_a, \mathbf{x}_p)$ and $(\mathbf{x}_a, \mathbf{x}_n)$.

One problem with this approach is that the number of triplet pairs grows cubically with the dataset, and it may become infeasible to fit all possible triplets in memory. We follow the *online semi-hard mining* scheme [17], where all possible anchor-positive pairs in a mini-batch are used. For each positive pair, the most difficult negative example \mathbf{x}_n satisfying $D(\mathbf{x}_a, \mathbf{x}_p) < D(\mathbf{x}_a, \mathbf{x}_n)$ is then used, except if there is no such negative example in which case the one with the largest distance is used. According to [17], although it might seem natural to simply choose the hardest negative example in the mini-batch, this constraint is required for stability.

We also incorporate another recent advance, specifically to improve efficience: we simplify the three shared-parameter networks with a single neural network that embeds a mini-batch of examples and then samples triplet pairs online from these embeddings. This is done with an efficient implementation of pairwise distances, similar to [37]. We build this single network model with semi-hard triplet mining, and refer to it as *Siamese CNN (online)*. We also compare to using three shared-parameter networks with the same CNN architecture, where we generate triplets offline at each training step from the current mini-batch. We refer to this model as *Siamese CNN (offline)* in our experiments. Similar to our neural network classifier baselines above, these models are trained on triplets from a large disjoint labelled dataset which do not contain the target one-shot classes.

5. EXPERIMENTS

5.1. Experimental setup

We perform multimodal one-shot learning on a simple benchmark dataset: learning from examples of spoken digits paired with handwritten digit images. For speech we use the TIDigits corpus which contains spoken digit sequences from 326 different speakers [38], and for images we use the MNIST handwritten digits dataset which contains 28×28 grayscale images [39]. We use utterances from men, women, and children, and split digit sequences into isolated digits using forced alignments. Speech is parametrised as Mel-frequency cepstral coefficients with first and second order derivatives. We centre zero-pad or crop speech segments to 120 frames. Image pixels are normalised to the range [0, 1]. Each isolated spoken digit is then paired with an image of the same type. Unlike previous work which used the same dataset combination for learning multimodal representations [32, 33], we treat utterances labelled "oh" and "zero" as separate classes, resulting in 11 class labels.

Neural network models are trained on large labelled background datasets to obtain feature representations which may be applied to the one-shot problem on classes not occurring in the background data. We use the speech corpus of [40] and the Omniglot handwritten character dataset [41] as background data for the within-modality speech and vision models, respectively. Utterances in the audio corpus are split into isolated words using forced alignments, and features are extracted using the same process as for TIDigits. We ensure that none of the target digit classes occur in this audio data. Images in the Omniglot dataset are downsampled to 28×28 and pixel values are normalised and inverted in order to match MNIST. Again, none of the Omniglot classes overlap with digit classes.

Models are implemented in TensorFlow and trained using the Adam optimiser [42] with a learning rate of 10^{-3} which is step decayed by 0.96 at each new epoch. A batch size of 200 is used for the neural network classifiers ($\S4.2$). For the Siamese models ($\S4.3$) we follow an alternate approach where mini-batches are formed using the batch all strategy proposed in [18]. Specifically, we randomly sample p classes and k examples per class to produce balanced batches of pk examples each. This results in pk(pk - k)(k - 1) valid triplet combinations, maximising the number of triplets within a mini-batch. Our Siamese CNN (online) variant is capable of large pk combinations due to the efficient single network implementation and online triplet sampling scheme. We choose p = 128 and k = 8 (total of 7 282 688 triplets per mini-batch). For Siamese CNN (offline), trained in the standard way where three networks are explicitly tied [14, 16, 17], we use p = 32 and k = 2, giving 3 968 triplets per mini-batch which is the largest batch we could fit on a single NVidia Titan Xp GPU. Models are trained for a maximum of 100 epochs using early stopping based on one-shot validation error on the background data.

We also tune all models using unimodal one-shot learning on the validation sets of the background data. This gave the following architecture for speech CNNs: 39×9 convolution with 128 filters; ReLU; 1×3 max pooling; 1×10 convolution with 128 filters; ReLU; 1×28 max pooling over remaining units; 2048-unit fully-connected; ReLU. Vision CNNs have the architecture: 3×3 convolution with 32 filters; ReLU; 2×2 max pooling; 3×3 convolution with 64 filters; ReLU; 2×2 max pooling; 3×3 convolution with 128 filters; ReLU; 2048-unit fully-connected; ReLU; 1024-unit fully-connected. The speech and vision FFNNs have the same structure: 3 fully-connected layers with 512 units each. For the classifier networks (§4.2), the speech networks have an additional 5534-unit softmax output layer (the number of word types in the background data), while vision networks have a 964-unit softmax (background image classes).

 Table 1. 11-way one-shot and five-shot speech classification results

 on isolated spoken digits.

Model	Train	11-way Accuracy	
	time	one-shot	five-shot
DTW	_	$67.99\% \pm 0.29$	$91.30\% \pm 0.20$
FFNN Classifier	13.1m	$71.39\% \pm 0.81$	$89.49\% \pm 0.45$
CNN Classifier	60.6m	$82.07\% \pm 0.92$	$93.58\% \pm 0.98$
Siamese CNN (offline)	70.5m	$89.40\%\pm0.54$	$95.12\% \pm 0.37$
Siamese CNN (online)	15.0m	$\textbf{92.85\%} \pm \textbf{0.38}$	$97.65\% \pm 0.22$

We evaluate our models on the one-shot tasks according to the accuracy averaged over 400 test episodes. Each episode randomly samples a support set of isolated spoken digits paired with images for each of the L = 11 classes (the digits "oh" to "nine" and "zero"). For testing, a matching set is sampled, containing 10 digit images not in the support set. Finally, a random query is sampled, also not in the support set. The query then needs to be matched to the correct item in the matching set. The matching set only contains 10 items since there are 10 unique handwritten digit classes. Within an episode, 10 different query instances are also sampled while keeping the support and matching sets fixed. Results are averaged over 10 models trained with different seeds and we report average accuracies with 95% confidence intervals.

5.2. One-shot speech classification

We first consider unimodal one-shot speech classification, which has so far only been considered in [19]. Table 1 shows one-shot and five-shot (see end of §3) speech classification results. Average training time is also shown; all models trained within a few seconds of the average. Siamese models outperform the direct feature matching baseline using DTW, as well as the neural network classifiers. The *Siamese CNN (online)* model achieves best overall performance, outperforming the *Siamese CNN (offline)* variant, while training almost five times faster. The single network with online semi-hard mining is thus more efficient and accurate than the three shared-network approach. None of these Siamese models were considered in [19].

5.3. One-shot cross-modal matching of speech to images

We now turn to multimodal one-shot learning. Table 2 shows results for one- and five-shot cross-modal matching of speech to images. Here the Siamese models are again stronger overall compared to direct feature matching or transferring features from neural network classifiers. The *Siamese CNN* (*online*) model achieves our best results, with double the accuracy of direct feature matching using pixeldistance over images and DTW over speech. The *Siamese CNN* (*offline*) model follows closely, but is again slower to train.

While the Siamese models achieve promising results compared to the baselines here, our best one-shot multimodal accuracy is lower than the accuracy in unimodal one-shot speech classification (Table 1). Our best unimodal one-shot vision model achieves an accuracy of 74% (comparing favourably to the best result of 72% reported in [20]). The multimodal one-shot results here are therefore worse than both the individual unimodal matching results. This is due to compounding errors in our retrieval framework (\S 4): errors in comparisons with the support set affects comparisons in the subsequent matching step. This suggests investigating an end-to-end architecture which can directly compare test queries in one modality to the matching set items in the other modality, without doing explicit comparisons to the support set.

 Table 2. 11-way one- and five-shot cross-modal matching of spoken and visual digits.

Model	11-way Accuracy		
	one-shot	five-shot	
DTW + Pixels	$34.92\%\pm0.42$	$44.46\% \pm 0.69$	
FFNN Classifier	$36.49\% \pm 0.41$	$44.29\% \pm 0.56$	
CNN Classifier	$56.47\% \pm 0.76$	$63.97\% \pm 0.91$	
Siamese CNN (offline)	$67.41\% \pm 0.56$	$70.92\% \pm 0.36$	
Siamese CNN (online)	$\textbf{70.12\%} \pm \textbf{0.68}$	$73.53\%\pm 0.52$	

Table 3. Speaker invariance tests for 11-way one-shot cross-modal speech-image digit matching. All support set items are from the same speaker as the query, except for the item actually matching the query.

Model	11-way Accuracy one-shot
DTW + Pixels FFNN Classifier CNN Classifier	$\begin{array}{c} 28.00\% \pm 1.86 \\ 34.95\% \pm 2.28 \\ 53.71\% \pm 2.20 \end{array}$
Siamese CNN (offline) Siamese CNN (online)	$\begin{array}{c} 66.70\% \pm 0.92 \\ \textbf{69.73\%} \pm \textbf{1.04} \end{array}$

5.4. Invariance to speakers

In all the experiments above we chose spoken queries such that the speaker uttering the query does not appear in the support set. This is representative of an extreme case where one user teaches an agent and another then uses the system. An even more extreme case could occur: the matching item in the support set could be the only item not coming from the query speaker. This is problematic since the same word uttered by different speakers might be acoustically more different than different words uttered by the same speaker. We test this worst-case setting: we sample a support set where all spoken digits are from the same speaker as the query, except for the one instance matching the query word which is produced by a different speaker. Cross-modal matching results for this case are shown in Table 3. All of the models experience a drop in accuracy compared to the results in Table 2 (first column). This decrease is smallest for the Siamese models, with the DTW + Pixels approach dropping most. This indicates that the neural models learn features from the background data which are more independent of speaker and can generalise to other speakers, whereas DTW over speech is affected more by speaker mismatch.

6. CONCLUSION

We introduced and formalised multimodal one-shot learning, specifically for learning from spoken and visual representations of digits. Observing just one paired example from each class, a model is asked to pick the correct digit image for an unseen spoken query. We proposed and evaluated several baseline and more advanced models. Although our Siamese convolutional approach outperforms a rawfeature nearest neighbour model, the performance in the cross-modal case is still worse than in unimodal one-shot learning. We argued that this is due to a compounding of errors in our framework, which relies on successive unimodal comparisons. In future work, we will therefore explore a method that can directly match one modality to another, particularly looking into recent meta-learning approaches.

We thank NVIDIA for sponsoring a Titan Xp GPU for this work.

7. REFERENCES

- L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Trans. PAMI*, vol. 28, no. 4, pp. 594– 611, 2006.
- [2] B. M. Lake, R. Salakhutdinov, J. Gross, and J. B. Tenenbaum, "One shot learning of simple visual concepts," in *Proc. CogSci*, 2011.
- [3] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *Proc. ICML*, 2015.
- [4] O. Räsänen and H. Rasilo, "A joint model of word segmentation and meaning acquisition through cross-situational learning," *Psychol. Rev.*, vol. 122, no. 4, pp. 792–829, 2015.
- [5] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech Commun.*, vol. 56, pp. 85–100, 2014.
- [6] M. R. Walter, Y. Friedman, M. Anton, and S. Teller, "One-shot visual appearance learning for mobile manipulation," *IJRR*, vol. 31, no. 4, pp. 554–567, 2012.
- [7] T. Taniguchi et al., "Symbol emergence in robotics: A survey," Adv. Robot., vol. 30, no. 11–12, pp. 706–728, 2016.
- [8] W. Thomason and R. A. Knepper, "Recognizing unfamiliar gestures for human-robot interaction through zero-shot learning," in *Proc. ISER*, 2016.
- [9] V. Renkens and H. Van hamme, "Weakly supervised learning of hidden Markov models for spoken language acquisition," *IEEE/ACM Trans. ASLP*, vol. 25, no. 2, pp. 285–295, 2017.
- [10] V. Renkens and H. Van hamme, "Capsule networks for low resource spoken language understanding," in *Proc. Interspeech*, 2018.
- [11] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah, "Signature verification using a "Siamese" time delay neural network," in *Proc. NIPS*. 1994.
- [12] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. CVPR*, 2005.
- [13] G. Chechik, V. Sharma, U. Shalit, and S. Bengio, "Large scale online learning of image similarity through ranking," *JMLR*, vol. 11, no. 11, pp. 1109–1135, 2010.
- [14] J. Wang et al., "Learning fine-grained image similarity with deep ranking," in *Proc. CVPR*, 2014.
- [15] K. M. Hermann and P. Blunsom, "Multilingual distributed representations without word alignment," in *Proc. ICLR*, 2014.
- [16] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *Proc. SIMBAD*, 2015.
- [17] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proc. CVPR*, 2015.
- [18] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *arXiv preprint arXiv:1703.07737*, 2017.
- [19] B. M. Lake, C.-Y. Lee, J. R. Glass, and J. B. Tenenbaum, "Oneshot learning of generative speech concepts," in *Proc. CogSci*, 2014.
- [20] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, "Matching networks for one shot learning," in *Proc. NIPS*. 2016.

- [21] P. Shyam, S. Gupta, and A. Dukkipati, "Attentive recurrent comparators," in *Proc. ICML*, 2017.
- [22] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. NIPS*. 2017.
- [23] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "Meta-learning with memory-augmented neural networks," in *Proc. ICML*, 2016.
- [24] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic metalearning for fast adaptation of deep networks," in *Proc. ICML*, 2017.
- [25] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, "A simple neural attentive meta-learner," in *Proc. ICLR*, 2018.
- [26] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," in *Proc. Robot Learn.*, 2017.
- [27] T. Stafylakis and G. Tzimiropoulos, "Zero-shot keyword spotting for visual speech recognition in-the-wild," in *Proc. ECCV*, 2018.
- [28] D. Wu, F. Zhu, and L. Shao, "One shot learning gesture recognition from rgbd images," in *Proc. CVPR*, 2012.
- [29] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "One-shot learning by inverting a compositional causal process," in *Proc. NIPS*. 2013.
- [30] J. Ngiam et al., "Multimodal deep learning," in *Proc. ICML*, 2011.
- [31] D. Harwath, A. Torralba, and J. R. Glass, "Unsupervised learning of spoken language with visual context," in *Proc. NIPS*. 2016.
- [32] K. Kashyap, "Learning digits via joint audio-visual representations," M.S. thesis, MIT, 2017.
- [33] K. Leidal, D. Harwath, and J. R. Glass, "Learning modalityinvariant representations for speech and images," in *Proc. ASRU*, 2017.
- [34] H. Kamper, G. Shakhnarovich, and K. Livescu, "Semantic speech retrieval with a visually grounded model of untranscribed speech," *IEEE/ACM Trans. ASLP*, vol. 27, no. 1, pp. 89–98, 2019.
- [35] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Trans. ASSP*, vol. 26, no. 1, pp. 43–49, 1978.
- [36] J. Donahue et al., "Decaf: A deep convolutional activation feature for generic visual recognition," in *Proc. ICML*, 2014.
- [37] H. O. Song, Y. Xiang, S. Jegelka, and S. Savarese, "Deep metric learning via lifted structured feature embedding," in *Proc. CVPR*, 2016.
- [38] R. G. Leonard and G. Doddington, "TIDIGITS LDC93S10," Philadelphia: Linguistic Data Consortium, 1993.
- [39] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradientbased learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [40] D. Harwath and J. R. Glass, "Deep multimodal semantic embeddings for speech and images," in *Proc. ASRU*, 2015.
- [41] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, "Humanlevel concept learning through probabilistic program induction," *Science*, vol. 350, no. 6266, pp. 1332–1338, 2015.
- [42] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.