
Latent Variable Algorithms for Multimodal Learning and Sensor Fusion

Lijiang Guo

Department of Intelligent Systems Engineering
Indiana University
Bloomington, IN 47405
lijguo@indiana.edu

Abstract

1 Multimodal learning has been lacking principled ways of combining information
2 from different modalities and learning a low-dimensional manifold of meaningful
3 representations. We study multimodal learning and sensor fusion from a latent
4 variable perspective. In the first part, we present a regularized recurrent attention
5 filter for sensor fusion. This algorithm can dynamically combine information from
6 different types of sensors in a sequential decision making task. Each sensor is
7 bonded with a modular neural network to maximize utility of its own information.
8 A gating modular neural network dynamically generates a set of mixing weights
9 for outputs from sensor networks by balancing utility of all sensors' information.
10 We design a co-learning mechanism to encourage co-adaption and independent
11 learning of each sensor at the same time, and propose a regularization based co-
12 learning method. In the second part, we focus on recovering the manifold of latent
13 representation. We propose a co-learning approach using probabilistic graphical
14 model which imposes a structural prior on the generative model: multimodal
15 variational RNN (MVRNN) model, and derive a variational lower bound for its
16 objective functions. In the third part, we extend the siamese structure to sensor
17 fusion for robust acoustic event detection. We perform experiments to investigate
18 the latent representations that are extracted; works will be done in the following
19 months. Our experiments show that the recurrent attention filter can dynamically
20 combine different sensor inputs according to the information carried in the inputs.
21 We consider MVRNN can identify latent representations that are useful for many
22 downstream tasks such as speech synthesis, activity recognition, and control and
23 planning. Both algorithms are general frameworks which can be applied to other
24 tasks where different types of sensors are jointly used for decision making. To our
25 knowledge, this algorithm is the first that addresses a online multimodal decision
26 making problem.

27	Contents	
28	1 Executive Summary	4
29	2 Introduction	4
30	3 Multimodal Learning and Sensor Fusion	6
31	3.1 An Example: Audio-Visual Learning	6
32	3.2 An Example: Load Balancing in Distributed System for Streaming Data	6
33	3.3 Sensor Fusion	7
34	4 Model Combination for Sensor Fusion	8
35	4.1 Motivation	8
36	4.2 Related Works	8
37	4.2.1 Temporal Attention	8
38	4.2.2 Spatial-Temporal Attention	10
39	4.2.3 Attention mechanism in audio-visual problems	11
40	4.3 Mixture of Sub-task Experts	12
41	4.4 Recurrent Attention Filter for Multimodal Sensor Fusion	13
42	4.4.1 Comments	14
43	4.5 Co-learning Latent Features	14
44	4.5.1 Distance based regularization	15
45	4.6 Recurrent Attention Filter for Audio-Visual Speech Separation	15
46	5 Model Based Sensor Fusion: Separating Modality Invariant and Modality Dependent Information	15
47	5.1 Motivation	16
48	5.2 Related Works	17
49	5.2.1 Inference in Graphical Models with Conjugate Prior	17
50	5.2.2 Variational Autoencoder	17
51	5.2.3 Recurrent Neural Network	18
52	5.2.4 Stochastic RNN	18
53	5.2.5 State Space Model: HMM and LDS	19
54	5.3 Multimodal HMM and LDS	20
55	5.4 Multimodal Variational RNN	21
56		
57	6 Learning Latent Variable from Similarity	23
58	6.1 Motivation	23
59	6.2 Global Embedding from Local Pairwise Similarity	23
60	6.3 Weakly Supervised Metric Learning without a Metric	24
61	6.4 A Example: Audio Event Detection in Noisy Environment	24
62	7 Experiments	26

63	7.1 Speech Activity Detection and Speech Separation	26
64	7.2 Data	28
65	7.3 Video Pre-processing and Pre-training	28
66	7.4 Audio Pre-processing and Pre-training	29
67	7.5 Training	31
68	7.6 Experiment Results	31
69	7.7 Spatial Attention Outputs	32
70	8 Conclusion and Future Works	33

1 Executive Summary

We present three multimodal learning and sensor fusion methods from a latent variable perspective. In [section 2](#) we briefly introduce the multimodal learning problem, [section 3](#) reviews background with concrete examples. In [section 4](#), [section 5](#), and [section 6](#) we introduce three different methods: a model-free sensor fusion approach, a model-based sensor fusion approach, and sensor fusion for weakly supervised embedding. In each of these sections we first discuss related works, then present our approach. Experiment results are discussed in [section 7](#). We discuss future works in [section 8](#).

2 Introduction

Multimodal learning has been studied in various forms for decades. Broadly speaking, multimodal input can be considered as streams of loosely synchronized multi-dimensional data, with each modality being one subset of dimensions. In simple cases when measurements are both taken in real domain, with certain metrics assumed, we often impose a model which leverages on the dependency across modalities to jointly solve a problem. Two commonly used dependency models are linear correlation and conditional probability.

Consider an example from [Shumway and Stoffer \[2011\]](#) where a patient’s biometric markers such as log(white blood count) [WBC], log(platelet) [PLT], and hematocrit [HCT] are used to predict probability of a patient’s long term survival using Bernoulli linear regression:

$$\mathbb{E}[Y|X] = f(\beta_0 + X_{\text{WBC}}\beta_1 + X_{\text{PLT}}\beta_2 + X_{\text{HCT}}\beta_3), \quad (1)$$

where f is a proper transformation function. We can consider X_{WBC} , X_{PLT} , X_{HCT} as different input modalities as they are essentially measuring different physical concepts. Classical statistical solutions focus on statistical inference in the sense of quantifying the error in estimating $\mathbb{E}[Y|X]$ using probability distributions. To this end, (1) f is often chosen to be functions which are friendly to manipulate, and (2) the relation between different modalities are chosen to be simple linear additive for the same reason. In this simple linear model for regression, we use linear combinations of the input variables. We can extend this model by considering linear combinations of fixed nonlinear functions of the input variables:

$$\mathbb{E}[Y|x] = f(\beta_0 + \phi_1(X_1)\beta_1 + \dots + \phi_k(X_K)\beta_K). \quad (2)$$

In this model, the relation between different input variables are still linear additive, but we allow for flexible representation, i.e. ϕ , of each variable, i.e. features, to be learned from training by optimizing some criterion.

Different from classical statistics, in machine learning we often face very high dimensional input data, and each input modality has essentially different representation forms; for example, video, audio, and text. The challenging questions in creating intelligent systems for processing audio or video inputs are

1. How to describe the complex relation between a single input with the target output, and
2. How to combine different input modalities when there are no straight forward physical model to describe them.

Let’s consider speech recognition as an example. Consider we have two input modalities, audio and video, and text as output. The first question corresponds to how are we going to describe the relation between audio and text, or video and text, respectively. The second question corresponds to how the audio input is related to video input in terms of text generation. In this research, we address the second question in a principled manner, and, equally important, create a multimodal algorithm that can outperform uni-modal algorithms.

We want to focus on multimodal time-series data as they are universal in speech, video, computer network, robot control, etc. Our objective is sequential inference under sequential multimodal input. Consider we have two input streams $\{X_t\}$ and $\{Y_t\}$, and an output stream $\{L_t\}$ where $t \in \{1, \dots, N\}$ are time indices. Our objective is to predict L_k from $\{(X_t, Y_t)\}_{t=1}^k$. The challenge is to model X_t given $(X_{t-1}, \dots, X_1, Y_t, Y_{t-1}, \dots, Y_1)$, and to model Y_t given $(Y_{t-1}, \dots, Y_1, X_t, X_{t-1}, \dots, X_1)$. This is particularly important when one input modality is corrupted, and the multimodal algorithm has to recognize or infer the missing information from other input modalities. In particular, we want a




Figure 1: Structural generative model.

119 algorithm which can estimate the reliability of each sensor input. We will discuss this in detail in
 120 [section 4](#).

121 From a generative model perspective, we assume that there are some latent random variable Z which
 122 generates the observed random variables X (see [Figure 1](#)). In multimodal data, X can be partitioned
 123 into disjoint sets of input signal where each $X^{(m)}$ corresponds to one modality. The observed speech
 124 activity state L is a function of X . The latent variables Z can also be partitioned into modality
 125 invariant and modality dependent subsets $\{Z^{(0)}, Z^{(1)}, \dots, Z^{(k)}\}$ where $Z^{(0)}$ is shared by all modality
 126 $\{X^{(m)}\}_{m=1}^k$, and each $Z^{(m)}$ is specific to $X^{(m)}$ for $m \in \{1, \dots, k\}$. Having a shared latent variable
 127 $Z^{(0)}$ enables the model to explain for the synchronization between different modalities. In [section 5](#)
 128 we discuss this in detail.

129 We will begin by reviewing a few related topics. Then we present a deep learning based regularized
 130 recurrent attention filter, which dynamically combines information from multiple sensor inputs in a
 131 sequential decision making task. We present a co-learning mechanism and discuss its advantages
 132 in learning robust latent features. Next we marry the strength of deep learning with probabilistic
 133 graphical models in a coherent way for multimodal learning. We present a multimodal nonlinear
 134 state-space model with a structural prior for separating modality dependent and modality invariant
 135 features. To demonstrate the proposed algorithms, we tackle the problem of speech activity detection
 136 and speech separation in corrupted audio-video streams.

137 This paper is a first step towards a faraway goal, thus we cannot solve the problem completely. Our
 138 model has the following contributions (no. 3, 4 and 5 are unique to our knowledge):

- 139 1. The model combines information collected from multiple sensors. Each sensor collects
 140 different types of data; for example, image, motion, audio, etc.
- 141 2. The model has independent processing module for each sensor. Each module is developed
 142 to maximize the utility of its sensor input in a decision making task.
- 143 3. The model addresses a sequential decision making task. It (1) takes new inputs at each time
 144 step, and (2) generates new attention over modalities by combining the new inputs and its
 145 memory, then (3) apply attention to modules to make a final decision for the current time
 146 step. The model does not need future information to make decision for current time step.
- 147 4. Each module is divided into two sub-components. One sub-component co-learns with other
 148 modules; this sub-component allows for simultaneous co-learning features that are shared
 149 by all sensors, which allows for robust detection of these features, even a subset of sensors'
 150 inputs are corrupted. The other sub-component learns independently from other modules;
 151 this allows each module to focus on its sensor's unique input features that are not shared by
 152 other sensors. This co-learning design prevents overfitting during training, and improves
 153 robustness of the modules.
- 154 5. The probabilistic approach to co-learning can be used to separate modality-dependent features
 155 from modality invariant features. It has wide applications such as speech recognition, speaker
 156 identity recognition, etc.

3 Multimodal Learning and Sensor Fusion

3.1 An Example: Audio-Visual Learning

Human senses the world with multiple modalities such as vision, sound, texture, etc. An AI agent, e.g. a robot, also relies on multiple sensors to collect data to perform tasks such as localization, path planning, control a robotic arm, etc. Baltrušaitis et al. [2018] refers a *sensory modality* as our primary channels of communication and sensation, such as vision or touch. A problem is characterized as *multimodal* when it includes multiple sensory modalities. In this work, as a concrete example, we focus on three modalities that are important to human speech: visual, audio, and motion. However, the method we discuss are general enough to tackle other sensor modalities on AI agents.

Many researches leverage natural synchrony between simultaneously recorded visual and audio signals to solve speech related problems. In recent works, lip-reading, the practice of using visual signals to understand speech, has been proven to be able to effectively extract useful information for speech recognition [Ephrat et al., 2017, Assael et al., 2016]. Hence it is natural to consider audio-visual approaches for many speech related tasks, including speech recognition [Ngiam et al., 2011, Mroueh et al., 2015, Feng et al., 2017], speech separation [Hou et al., 2017], and voice activity detection [Ariav et al., 2018].

Among these problems, speech separation is one of the fundamental problems in audio processing. Wang and Chen [2018] gave a overview of recent advances in deep-learning based audio-only speech separation system. Traditionally audio-only methods are used to separate a speech signal from other background signals. When multiple human speeches overlap each other, the speech separation problem is referred to as the *cocktail party problem*. Such a problem is especially challenging to solve because human speeches tend to have similar features than with non-speech sounds. It is also difficult to assign a speech to the corresponding speaker with audio signals only, often referred as *label permutation problem*. A few solutions have been proposed to address multi-speaker speech separation using single channel audio recording. For example, Hershey et al. [2016] proposed a method called deep clustering (DPCL) which uses discriminatively trained speech embeddings to cluster and separate speeches. They also proposed a permutation invariant loss function to solve the label permutation problem. Yu et al. [2017] and Isik et al. [2016] successfully use a permutation invariant loss function to train a DNN for multi-speaker speech separation.

In recent years, deep learning based audio-visual methods have been used for speech separation. Hou et al. [2017] proposed a CNN-based model which outputs a denoised speech spectrogram as well as a reconstruction of the input mouth region. Gabbay et al. [2017] trained a speech enhancement model on videos where other speech samples of the same speaker are used as background noise. Gabbay et al. [2018] use a video-to-sound synthesis method to filter noisy audio. While most audio-visual speech separation methods are speaker dependent, Ephrat et al. [2018] proposed a speaker-independent audio-visual model for separating single speech from a mixture of sounds such as multiple speakers and background noise. They use a pretrained face detection model to identify all speakers in a video, then outputs a complex spectrogram mask for each speaker. Alternatively, Owens and Efros [2018] train a deep neural network to predict whether audio and visual streams are temporally aligned. Learned features extracted from this self-supervised model are then used to condition an on/off screen speakers source separation model. Gao et al. [2018] and Zhao et al. [2018] addressed the closely related problem of separating the sound of multiple on-screen objects (e.g. musical instruments).

3.2 An Example: Load Balancing in Distributed System for Streaming Data

Suppose we have K streams of input data $x^{(1)}, \dots, x^{(K)}$, each are different but correlated. They are jointly used to solve a problem whose answer is a stream y . Consider we have many *workers* $\{w^{(1)}, \dots, w^{(N)}\}$ who are processing these inputs $\{x_k\}_1^K$. In a distributed system we need to partition the input streams to assign to different workers, and then gather them to solve for the answer. Let's assume for stream k , the information complexity fluctuates in time according to an unknown function $h^{(k)}(x)$, such that the processing time of a worker w is proportional to $t(k, x) = f(h^{(k)}(x))$. For example, a encoded video could take various amount of computations to decode.




Figure 2: Multimodal deep Boltzmann machine (reproduced from [Ngiam et al. \[2011\]](#)).

3.3 Sensor Fusion

Sensor fusion is one of the core problems in multimodal learning. Technically speaking, in multimodal sensor fusion, we need to integrate information from multiple modalities with the goal of fulfilling some tasks. Comparing with unimodal methods, [Baltrušaitis et al. \[2018\]](#) suggested multimodal sensor fusion has the benefit of (1) providing more robust predictions, (2) allowing the model to utilize complementary information from different modalities, and (3) imputing missing information for corrupted signals.

Sensor fusion algorithms can be divided into two big categories [[Baltrušaitis et al., 2018](#)]: model-agnostic and model-based. Model-agnostic sensor fusion does not depend on a specific machine learning method. It usually combines different sensors in a blind way and the relationship between sensors are not exploited. Model-based methods explicitly address fusion in their construction, e.g. using graphical models to impose a structural prior on modalities. In this work, we will propose both a model-agnostic approach and a model-based approach for sensor fusion.

The most common fusion approach is to concatenate the inputs from different sensors at certain modeling stage to solve a unimodal learning problem. [Ngiam et al. \[2011\]](#) proposed a multimodal autoencoder which first use deep Boltzmann machine (DBM) to learn each audio and video features independently, then concatenate the learned features to an multimodal autoencoder to learn a shared representation ([Figure 2](#)). To reproduce original inputs, the shared representation is mapped back to each modality. [Srivastava and Salakhutdinov \[2012\]](#) also use DBM to learn a generative model of multimodal input data. [Simonyan and Zisserman \[2014a\]](#) proposed a two-stream convolutional neural networks for action recognition in video. The two-stream model try to capture the complementary information from two modalities: still image and motion between frames. In their work, the last layers concatenates the two streams into a final classification layer. [Torfi et al. \[2017\]](#) proposed a coupled 3D-CNN to map multiple modalities into a representation space to evaluate the correspondence of audio-visual streams. In [Ephrat et al. \[2018\]](#) the audio and visual streams are combined by concatenating the feature maps of each stream, which are subsequently fed into a BLSTM followed by three FC layers.

Whether to combine information at early or late stage affects sensor fusion results. [Ngiam et al. \[2011\]](#) argued that naive concatenation of different sensors' inputs discourages co-learning since the within sensor correlation is much stronger than between sensor correlation, hence the model are biased towards learning dominant patterns in each sensor separately instead of learning patterns that occur simultaneously in multiple sensors. In this regard, they propose to learn joint representations that are shared across multiple modalities at the higher layer of the deep network after learning layers of modality-specific networks. As [Sohn et al. \[2014\]](#) commented, the rationale is that, they believe, the learned features may have less within-modality correlation than raw features, and this makes it easier to capture patterns across data modalities. However, in [Ngiam et al. \[2011\]](#), the middle layer is a simple fully connected layer, and we can only hope that it can always capture meaningful representations shared by all modalities.

One challenge of multimodal learning is how to deal with missing data in different modalities, and infer the missing data from other modalities — it is a unique advantage of multimodal learning

over unimodal learning as we can use the information extracted from one modality to improve the recognition ability of the other modality by complementing the missing information, i.e. given two input streams $\{X\}$ and $\{Y\}$, how to model X_t given $(X_{t-1}, \dots, X_1, Y_t, Y_{t-1}, \dots, Y_1)$, and Y_t given $(Y_{t-1}, \dots, Y_1, X_t, X_{t-1}, \dots, X_1)$. A naive approach is to use an autoencoder to map all modalities into a common latent space, and then map back to each modality, hoping that the latent representation can recover all modalities. For example, [Ngiam et al. \[2011\]](#) and [Ariav et al. \[2018\]](#) use a sparse autoencoder to map different modalities into the same space. [Ngiam et al. \[2011\]](#) requires each pre-trained modality-specific network is also jointly trained with other modality-specific networks for cross-modality learning. To infer missing data in one modality, they suppress other modalities with zero input during training.

Model-based sensor fusion explicitly incorporate the relation between sensors into the fusion stage. For example, we can combine multiple sensor signals into modality-dependent and modality-invariant latent random variables through a probabilistic graphical model. Model-based sensor fusion has the advantage of probabilistically detect and recover corrupted signal. We will discuss model-based sensor fusion in detail in [section 5](#).

In an interesting work, [Sohn et al. \[2014\]](#) proposed a multimodal representation learning framework which minimizes the information distance between data modalities through the shared latent representations. To this end, they use variation of information to measure the conditional probability of each modality given the other modalities, thereby determine the quality of shared representation across modalities. As a generative model it can predict missing data modalities given partial observation.

Despite the promise, there still remains missing a principled way of how to learn a good association between multiple data modalities that can effectively deal with missing data modalities in the testing time. Both [Ngiam et al. \[2011\]](#) and [Sohn et al. \[2014\]](#)'s approaches promote learning a shared representation by different modalities, and rely on the collaboration of different modalities. However, co-adaptation of different modalities may reduce the individual strength of each modality as we discuss in next sections. We propose two strategies to coordinate the collaboration between modalities. In the first approach we let each local expert focus on its own modality, and let a dedicated gate expert to coordinate experts. We will discuss this in [section 4](#). In the second approach, we propose a probabilistic graphical model which decompose latent representation into modality dependent and modality invariant features. We will discuss this in [section 5](#).

4 Model Combination for Sensor Fusion

4.1 Motivation

In this section we will develop a model-agnostic sensor fusion algorithm. We are motivated by model combination and aggregation as a way to use parallel models to jointly solve a common problem. Model combination algorithms such as *boosting* has been proven to be an effective approach in reducing prediction error and variance [\[Bishop et al., 2006\]](#). A natural way to approach multimodal learning is to divide a task into multiple parallel sub-tasks such that each modality uses a dedicated processing module to solve a sub-task. We consider a *conditional mixture model* [\[Bishop et al., 2006\]](#) to coordinate all processing modules towards a common learning goal. From the success of *attention mechanism* in sequence-to-sequence model, we observe that mixing weights can be generated in analogous to a spatial attention through a deep neural network. Therefore we combine model combination methods with recent development in deep learning to propose a recurrent attention filter for multimodal sensor fusion.

4.2 Related Works

4.2.1 Temporal Attention

In sequential prediction problems, sequence-to-sequence (seq2seq) model is a ground-breaking work [\[Sutskever et al., 2014, Cho et al., 2014a\]](#). One problem with seq2seq networks is that their performance will deteriorate rapidly as the length of input sequence increases [\[Cho et al., 2014b\]](#). RNN such as LSTM still suffers from difficulty in memorizing long-range relations, especially when a lot of information is cluttered in a sequence. The recently proposed attention mechanism solved this problem by providing a skip-connection and let the decoder focus on a sub-region of the

input [Bahdanau et al., 2014, Chorowski et al., 2015, Mnih et al., 2014, Ba et al., 2014]. Instead of memorizing the entire information of the past sequence, attention filter let the RNN memorize the location of the past information. This greatly reduces the information load of the RNN memory cell. Recently, convolutional and fully-attentional feed-forward architectures like the Transformer model [Vaswani et al., 2017] have emerged as a viable alternative to RNNs for a range of sequence modeling tasks.

In [Bahdanau et al., 2014], attention mechanism was introduced to help the decoder RNN to focus on relevant parts of the input sequence without relying on the RNN to encode all the information through repeated updating the hidden unit. Recall in a seq2seq model, the encoder takes a length M input sequence $\{x_t\}_{t=1}^M$ and generates an *encoder* hidden state sequence $\{h_t^e\}_{t=1}^M$. The last hidden state of encoder h_M^e is used as the initial hidden state for the decoder h_0^d to generate another state sequence $\{h_t^d\}_{t=1}^N$. During training, to making the decoder learn faster, the decoder also takes an auxiliary input which is the true label sequence, $\{y_t\}_{t=1}^N$; during testing, use decoder output prediction y_{t-1} as input to step t . In this simple encoder-decoder structure, suppose the last hidden node of the decoder, h_N^d is related to the first hidden node of the encoder, h_1^e , then the relevant information has to pass through $M + N$ RNN transitions. In general, the average distance between encoder and decoder hidden units is $O(M + N)$ — the number of operations required to relate signals from two arbitrary input or output positions grows linearly in the distance between positions. This long path has made it more difficult to learn dependencies between distant positions. RNN such as LSTM suffers from difficulty in memorizing long-range relations, especially when a lot of information is cluttered in a sequence. To solve this problem, Bahdanau et al. [2014] proposed attention mechanism which provides a skip-connection and let the decoder focus on a sub-region of the input. Instead of memorizing the entire information of the past sequence, attention filter let the RNN remember the location of the past information. This greatly reduce the information load of the RNN. With attention, the average distance between encoder and decoder hidden units is $O(1)$ — a constant number of operations.

The attention mechanism is similar to searching in a database and involves an interaction between query, key, and values. Consider the input to time step t of the decoder as a **query** (i.e. input is the previous hidden state h_{t-1}^d),

$$q_t = h_{t-1}^d. \quad (3)$$

The encoder hidden states $\{h_t^e\}_{t=1}^M$ are **keys** whose information are the **values**. For example, h_t^e is the key to t -th encoder state. To get the t -th decoder output y_t , we search for the relevant encoders values by comparing its key with the query,

$$p(y_t | \{x_t\}_{t=1}^M, \{y_t\}_1^{t-1}) = g(y_{t-1}, h_{t-1}^d, c_t) \quad (4)$$

where c_t (often named the context) is a result of the query procedure. Context vector can be seen as a continuous bag of weighted features of encoder hidden states \mathbf{h}^e [Chan et al., 2016].

Finding the appropriate query procedure is an area of current research. In a simple form, the query is performed in the following steps. First, we compute a similarity measure between q_t and each of h_k^e for $k \in \{1, \dots, M\}$, e.g. the kernel inner product $e_{t,k} = \langle q_t, h_k^e \rangle_{W^a} = q_t' W^a h_k^e$ where $'$ indicates transpose and W^a is a (square) matrix. We normalize the similarity over k by

$$w_{t,k} = \frac{\exp(e_{t,k})}{\sum_j \exp(e_{t,j})}. \quad (5)$$

The context is computed as the weighted linear combination of the encoder hidden states

$$c_t = \sum_{k=1}^M w_{t,k} h_k^e. \quad (6)$$

In the original paper that proposed attention, Bahdanau et al. [2014] commented that the approach of taking a weighted sum of all the annotations (e.g. $\{h_k^e\}_{k=1}^M$) is as if computing an expected annotation, i.e. $\mathbb{E}[h^e]$, over all the possible alignments. Here, alignment is defined by the weights $\{w_{t,k}\}_{k=1}^M$, which essentially aligns the t decoder with the entire encoder annotation sequence $\{h_k^e\}_{k=1}^M$. In this sense, $w_{t,k}$ is an estimated probability that the target word y_t is aligned to or translated from a source

work x_k . Thus the t -th context vector c_t is the expected annotation over all the annotations with probability $w_{t,k}$.

Two kinds of attention models have been proposed. The hard attention model uses the weights to randomly select one location [Mnih et al., 2014, Ba et al., 2014]; the soft attention model uses the weights to form a convex combination of many locations [Bahdanau et al., 2014, Ba et al., 2014]. In this work, we choose not to use stochastic “hard” attention for two reasons. First, stochastic attention requires sampling one input stream to make a decision. If both the visual and audio input are noisy, the stochastic attention model has to choose one that is more informative in making a decision. However, if the audio noise and visual noise are independent, the information from each stream may compensate the other. Therefore, we choose to use a deterministic “soft” attention model which dynamically combines multiple streams of data. Second, stochastic attention model involves an intractable objective function with a multinomial latent variable. Monte Carlo and REINFORCE [Williams, 1992] are two common methods for solving this problem. However, it is well known that the both estimation procedures are slow and unstable. On the contrary, soft attention model is fully differentiable and can be optimized using stochastic gradient descent.

4.2.2 Spatial-Temporal Attention

The temporal attention mechanism enables RNN to selectively focus on a subset of the input sequence. In multimodal learning, in addition to temporal attention, we aim to selectively focus on one (or a few) modalities of current time step and the past, which requires a spatial-temporal attention mechanism.

A relevant problem is image caption generation whose task is to transcribe an image into a word sentence that best describes the image. In Xu et al. [2015], a context vector is used to dynamically select relevant parts of a image to generate a word at time t . The context vector is a function of the RNN hidden states of time $t - 1$ and the image features. Xu et al. [2015] define a mechanism that generates a positive weight for each image feature location, which can be interpreted as the probability that location is the right place to focus for producing the next word in a sequence.

The problem we address here is different in a few places. First of all, in Xu et al. [2015] the image input to the decoder RNN is the same for all time steps. The attention model gives time dependent weight to each location of the image based on the decoder hidden state input and decoder model output of $t - 1$. We try to address a problem where at each time step the inputs (e.g. audio, motion, image) is different. This makes our model applicable to, e.g. online video description. Secondly, in Xu et al. [2015] the attention is imposed on the static extracted image features. In our case, because at each time step the inputs are different, we defer attention to each local expert’s decision. Third, although the image features are at different locations, they are same type of data and carry similar information. In our model, we have input features of totally different data types, e.g. image, motion and audio. Last, the image features are spatially dependent and the dependency does not change. In video, the spatial dependency of different input streams are dynamic and asynchronous. In summary, we use a dynamic-input cross-modality attention.

Another relevant work is dual-stage attention RNN (DA-RNN) [Qin et al., 2017]. Dual-stage attention solves two unique problems different from classical attentions. First, an input attention is used to attentively extract relevant information from multiple parallel exogenous input sequences instead of one input sequence. Then a temporal attention mechanism is used to select relevant encoder hidden states generated by input attention across all time steps. The decoder outputs a time-series prediction instead of a classification. Experiments show that DA-RNN is effective in time-series prediction and robust to noisy inputs.

RNN performance will deteriorate rapidly as the length of input sequence increases, Qin et al. [2017] used a temporal attention in the decoder to adaptively select relevant encoder hidden states across all time steps. They use the concatenate score function in [Luong et al., 2015].

$$l_t^i = \mathbf{v}_d^T \tanh(\mathbf{W}_d[\mathbf{d}_{t-1}; \mathbf{s}_{t-1}'] + \mathbf{U}_d \mathbf{h} + i), \quad 1 \leq i \leq T \quad (7)$$

and

$$\beta_t^i = \frac{\exp(l_t^i)}{\sum_{j=1}^T \exp(l_t^j)}, \quad (8)$$

In the encoder, $\mathbf{x}_t \in \mathbb{R}^n$ is the observation of n exogenous driving input sequences at time t . In a normal RNN, the input sequences $\{\mathbf{x}_t\}_{t=1}^T$ is fed to a hidden unit such as LSTM which is updated by

$$\mathbf{h}_t = f(\mathbf{h}_{t-1}, \mathbf{x}_t).$$

Qin et al. [2017] proposed to use \mathbf{h}_{t-1} and \mathbf{x}_t to generate a set of n weights for x_t :

$$e_t^k = \mathbf{v}_e^T \tanh(\mathbf{W}_e[\mathbf{h}_{t-1}; \mathbf{s}_{t-1}] + \mathbf{U}_e \mathbf{x}^k) \quad (9)$$

and

$$\alpha_t^k = \frac{\exp(e_t^k)}{\sum_{i=1}^n \exp(e_t^i)} \quad (10)$$

where $\mathbf{v}_e \in \mathbb{R}^T$, $\mathbf{W}_e \in \mathbb{R}^{T \times 2m}$ and $\mathbf{U}_e \in \mathbb{R}^{T \times T}$ are parameters to learn. Using the weights α_t , a new input is

$$\tilde{\mathbf{x}}_t = \mathbf{x}_t \odot \alpha_t. \quad (11)$$

Then the hidden state at time t is updated as

$$\mathbf{h}_t = f(\mathbf{h}_{t-1}, \tilde{\mathbf{x}}_t). \quad (12)$$

With the proposed input attention mechanism, the encoder can selectively focus on certain driving series rather than treating all the input driving series equally.

Note in [Qin et al., 2017] the prediction is on $\mathbf{y}_t \in \mathbb{R}$. But in denoising, the prediction is on $\mathbf{y}_t \in \mathbb{R}^D$ where D is the number of frequency bins.

4.2.3 Attention mechanism in audio-visual problems

Temporal attention mechanism has been used to improve RNN performance for speech related problems. Chan et al. [2016] proposed Listen, Attend and Spell (LAS), which is an audio-only speech recognition model based on sequence-to-sequence learning framework with attention. The model learns to transcribe an audio sequence signal to a word sequence, one character at a time. It consists of an encoder RNN, and a decoder RNN. The encoder RNN is a pyramidal RNN which converts low level speech signals into higher level features. The decoder is an RNN that converts high level features into output utterances by specifying a probability distribution over sequences of characters using the attention mechanism. The encoder and decoder are trained jointly. The key contribution of Chan et al. [2016] is the pyramidal RNN model for the encoder, which reduces the number of time steps that the attention model has to extract relevant information from.

RNN has been well known for its power in modeling sequential relation. It is also hard to train and can easily overfit. The attention mechanism has been used as a way to reduce overfitting. Chan et al. [2016] commented that without the attention mechanism, the model overfits the training data significantly and memorizes the training transcripts without paying attention to the acoustics.

Chorowski et al. [2015] proposed an attention based audio-video speech recognition model. Their model has a two-stream network with temporal attention in each stream. Each stream's encoder consists of a feature extractor, and a RNN encoder. The RNN encoder takes per-frame input from feature extractor in reverse time order. In the end, the two streams' RNN sequence outputs and final state are fed into a RNN decoder to generate texts. They found the attention mechanism is critical for the speech recognition system to work. Without attention, the model appears to forget the input signal, and produce output sequence that correlates very little to the input. Chorowski et al. [2015] use separate temporal attention to each modality in the decoder. However, the decoder gives equal weights to each modality when generating outputs. That is, each modality receives a sum of attention equals 1. Suppose one modality is corrupted, the weight on that modality is not adjusted. In this regard, Chorowski et al. [2015] does not consider the variation of information in modalities.

Here we present a spatial attention based sensor fusion model as a principled way of combining different modality of signals (i.e. audio and video) for e.g. speech enhancement. The model is trained end-to-end, and simultaneously learns spatiotemporal audio-visual features and a sequence model.




Figure 3: Mixture of sub-task experts (left) and mixture of experts (right). In mixture of experts, each expert receives the identical task. In mixture of sub-task expert, each expert receives a distinct subset of the task.

4.3 Mixture of Sub-task Experts

Dynamically combining multiple functions in supervised learning can be traced back to *mixture of experts* model [Jacobs et al., 1991]. A mixture of experts model is driven by the assumption that a set of training cases may be naturally divided into subsets that correspond to distinct tasks. However, the interference between different subsets of tasks would lead to slow learning and poor generalization. Such interference can be reduced by using a system composed of several different *expert functions* where each one is trained for a subset of tasks. A *gating function* is trained to decide which of the experts should be used for each training case. Instead of a hard decision such as decision tree, the gating function makes probabilistic decision by assigning mixing weights to the experts. Neural networks can be used for both the expert functions and gating function, which is known as *mixture density network* (MDN) [Bishop, 1994]. One advantage of MDN is it can approximate a flexible family of distributions, including distribution of multiple modes.

Jacobs et al. [1991] want a system to learn how to allocate cases to experts. They designed a loss function such that the gating network allocates a new case to one or a few experts, and if the output is incorrect, the weight changes are localized to these experts and the gating network. The experts are therefore local in the sense that the weights in one expert are decoupled from the weights in other experts. In addition they will often be local in the sense that each expert will be allocated to only a small local region of the space of possible input vectors.

Different from dividing all training cases into subsets of tasks, we observe that a training case can be divided into *parallel sub-tasks* (Figure 3). For example, in using two hands to open a jar, the parallel sub-tasks are left-hand movement and right-hand movement. In speech, the lip movement and audio are data of two sub-tasks, where speech is the super-task.

Many recent researches have implicitly utilized such a parallel sub-task design. In activity recognition, the two-stream network [Simonyan and Zisserman, 2014a] has a sub-task design for extracting motion and image features. The shared decoder works as a fusion function which takes concatenated inputs from sub-networks. In audio-visual speech recognition, Chorowski et al. [2015] use two encoder neural networks for processing audio and video signals. The two also share one decoder which takes concatenated inputs from the two encoders. In these works, the gating function is implicit in the decoder. Such a blackbox methods does not explicitly assign sub-tasks to experts, and the error backpropagation may cause global changes to expert functions.

We make three extensions to MDN for multimodal sensor fusion. First, we separate expert networks into different functional components by dividing a training case into parallel sub-tasks, one for each sensor. Sub-tasks could have same objective or partially overlapping objectives. Each expert network receives input from only one unique sensor, hence is forced to solve only one sub-task. The gating network receives all sensor inputs to decide a soft mixing weight. This contrasts a boosting algorithm, e.g. mixture of expert, where each expert receives the complete and identical task.

Second, in the mixture of density network, one function (i.e. neural network) is used to predict the parameters of all of the component densities as well as the mixing coefficients, so the hidden units between input and output layers are shared among the input-dependent functions. We further divide

the functions into different modular neural networks, such that each expert learns a unique function, e.g. neural network for the assigned sub-task.

A major concern of mixture of experts model is co-adaption of experts. The third contribution of our work is to re-address this issue by promoting multiple experts to learn similar features, and also learning dissimilar features at the same time; we propose a mechanism to partition latent variables into co-adaption set and independent set; see [subsection 4.5](#).

4.4 Recurrent Attention Filter for Multimodal Sensor Fusion

In this section we introduce recurrent attention filter for sensor fusion. We try to marry the strength of spatial-temporal attention and mixture of sub-task experts. For demonstration, we use speech activity detection where the inputs are audio, image, and motion, and outputs are binary labels. This model can be converted to a, e.g. speech and video enhancement system where outputs are, e.g. spectrograms and images.

Let $\{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ be a sequence of input signals. Partition \mathbf{x}_t into $\{\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \mathbf{x}_t^{(3)}\}$ which are audio, image, and motion features, respectively. These features are functions, e.g. neural networks, of input signals. Let $\mathbf{y} = \{y_1, \dots, y_T\}$ denote the sequence of speech activity labels, i.e. $\{0, 1\}$. Let $\mathbf{x}_{\leq t}^{(m)}$ and $\mathbf{y}_{\leq t}$ denote the subsequences $\{\mathbf{x}_1^{(m)}, \dots, \mathbf{x}_t^{(m)}\}$ and $\{y_1, \dots, y_t\}$. We model y_t as a Bernoulli random variable with conditional distribution:

$$\mathbb{P}(y_t = 1 | \mathbf{x}_{\leq t}) = \langle g(\mathbf{x}_{\leq t}), [f_1(\mathbf{x}_{\leq t}^{(1)}), f_2(\mathbf{x}_{\leq t}^{(2)}), f_3(\mathbf{x}_{\leq t}^{(3)})] \rangle, \quad (13)$$

where $\langle \cdot, \cdot \rangle$ is Euclidean inner product. Each *expert* function $f_m(\cdot)$ is a neural network which predicts the conditional distribution of y_t given input from modality m only:

$$f_m(\mathbf{x}_{\leq t}^{(m)}) = \mathbb{P}(y_t = 1 | \mathbf{x}_{\leq t}^{(m)}). \quad (14)$$

The *spatial attention* function $g : \mathbb{R}^n \mapsto \mathbb{R}^m$ is also a neural network which specifies a mixing weights for each f_m . Input to g are features of all modalities, as g has to output weights on all modalities. The output weights are non-negative and sum to 1. Therefore, $y_t | \mathbf{x}_{\leq t}$ has a mixture of Bernoulli distributions, where each component is a conditional distribution:

$$p(y_t | \mathbf{x}_{\leq t}) = \sum_m p(m | \mathbf{x}_{\leq t}) p(y_t | \mathbf{x}_{\leq t}^{(m)}). \quad (15)$$

In the *recurrent attention mixture model* ([Equation 15](#)), for each expert function we apply temporal attention on the input $\mathbf{x}_{\leq t}^{(m)}$, such that $f_m(\mathbf{x}_{\leq t}^{(m)})$ can look back in time and focus on relevant time segments. The recurrent attention mixture model can be solved using stochastic gradient descent. RNN with attention is often slow to train and has high memory consumption. Therefore the RNN with temporal attention can be replaced with 1-D dilated convolutional neural network.

A simplified version is akin to Markov mixture of expert [[Meila and Jordan, 1996](#)] where current prediction is independent of the past given a latent random variable:

$$p(y_t | \mathbf{x}_{\leq t}) = \sum_m p(m | \mathbf{x}_t, \mathbf{h}'_{t-1}) p(y_t | \mathbf{x}_t^{(m)}, \mathbf{h}_{t-1}^{(m)}) \quad (16)$$

where \mathbf{h}' and $\mathbf{h}^{(m)}$ are RNN hidden units in attention and expert networks. The *Markov attention mixture model* ([Equation 16](#)) is shown in [Figure 4](#). Regularization is discussed in [subsection 4.5](#).

We can further simplify the model by assuming $\mathbf{y}_t \perp \mathbf{x}_{< t} | \mathbf{x}_t$ such that

$$p(y_t | \mathbf{x}_{\leq t}) = \sum_m p(m | \mathbf{x}_t) p(y_t | \mathbf{x}_t^{(m)}). \quad (17)$$

This simplified model ([Equation 17](#)) is a *conditional attention mixture model* — each time step is a independently distributed mixture of conditional Bernoulli distributions. We can follow the standard EM receipt [[Bishop et al., 2006](#)] to solve the Markov attention mixture model and conditional attention mixture model (see [Appendix A](#)).

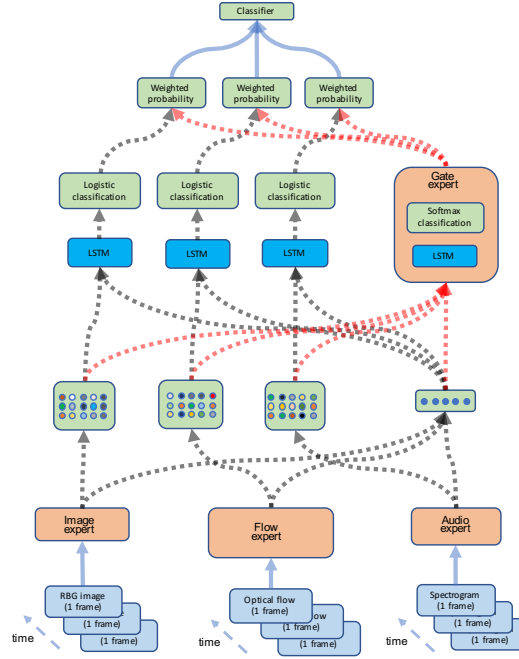


Figure 4: Regularized recurrent attention multimodal sensor fusion model.

4.4.1 Comments

In multimodal sensor fusion, each expert function would process one of the parallel sensor input, rather than a subset of all the training cases as in mixture of experts (Figure 3). The gating function is replaced with a spatial attention function, which at each time step, generates soft attention over sensors.

Baltrušaitis et al. [2018] suggested that late fusion ignores the low level interaction between the modalities. However, Ngiam et al. [2011] found that it is difficult to capture cross-modality relation with low level features, because within-modality correlation is much stronger than between-modality correlation. They suggested that fusion at higher level would encourage learning cross-modality features, because higher level features may have less within-modality correlation than raw features.

In our model, the gating network could take input features from any level of the expert networks. It is possible to apply a hybrid of late and early fusion by taking both high and low level features. Hence our model can be considered as a model agnostic hybrid fusion. These features are used to create dynamic mixing weights, similar to a spatial attention mechanism.

The spatial attention function g is similar to soft attention [Xu et al., 2015]. But there are two key differences. First, $g(t)$ does not have input from $\mathbf{x}_{>t}$; this is critical because g does not rely on future information to make current decision, which allows for online decision making. Second, for $t + 1$, g takes new input \mathbf{x}_{t+1} in addition to previous inputs $\mathbf{x}_{\leq t}$, hence a filter.

The recurrent attention mixture model is similar to mixture density network [Bishop, 1994]. The major difference is that we use dedicated neural networks to predict each component density (i.e. expert) and mixing coefficients (i.e. gate), thus maximizing the utility of each sensor's input. To recognize the covariance between different sensors, we design a co-learning mechanism as explained in next section.

4.5 Co-learning Latent Features

Each expert function f_m is trained to maximize the utility of its sensor input. Although sensor inputs are different, they are likely to have some similar latent features. For example, when uttering some words, part of the lip movements (i.e. visemes) are coordinated with the sound (i.e. phoneme). In this sense, two input modalities are not independent. We can make an assumption that for the speech

activity detection task, for each experts, there are some common features that “represent” some states; for example, a fixed set of nodes in both experts’ hidden layer are activated when speech is found. If we can apply some prior information which encourage latent variable sharing, we may extract more robust features for the classification task.

To this end, we consider two techniques from statistical learning. One is to apply a penalty term to part of hidden unit outputs to encourage the output of different experts be similar to each other. For example, we can add a term to the loss function which is proportional to the sum of squared distances between outputs and average of outputs. Alternatively, instead of tying hidden unit outputs we can put penalty on weights to encourage some weights are tied between different expert networks.

A second approach is to define a structure prior of the generative model via probabilistic graphical model. We will discuss this approach extensively in [section 5](#).

4.5.1 Distance based regularization

For each expert network, choose one hidden layer. Let $\mathbf{z}_1, \dots, \mathbf{z}_M$ denote the hidden unit outputs of these layers of the M experts. Let \mathbf{z}_m^* denote the first n hidden nodes of \mathbf{z}_m . We assume that

$$\mathbb{E}[\mathbf{z}_m^*] = \zeta. \quad (18)$$

That is, for all input sensors, \mathbf{z}_m^* has the same expectation. Thus

$$\hat{\zeta} = \frac{1}{K} \sum_{k=1}^K \mathbf{z}_k^* \quad (19)$$

is an unbiased estimator of ζ , where \mathbf{z}_k^* is a sample of \mathbf{z}_k^* . Define the co-learning loss as

$$\mathcal{L}_{co} = \sum_{m=1}^M \lambda_m \|\mathbf{z}_m^* - \hat{\zeta}\|_2^2. \quad (20)$$

\mathcal{L}_{co} is a sum of squared $L2$ norms; λ_k is a sensor specific penalty parameter selected by cross-validation. Adding \mathcal{L}_{co} to the loss function has the effect of shrinking \mathbf{z}_k^* towards ζ , as in Tikhonov regularization, to prevent overfitting and stabilize parameter estimation.

We want to point out that if we assume \mathbf{z}^* are Gaussian random variables, then \mathcal{L}_{co} is equivalent to impose a Gaussian prior distribution over ζ . In this sense, the \mathcal{L}_{co} estimate of \mathbf{z}^* is a MAP estimate. The advantage of formulating \mathcal{L}_{co} as Tikhonov regularization is that we can replace $\hat{\zeta}$ with a moving average of the current mini-batch of samples in stochastic gradient descent.

4.6 Recurrent Attention Filter for Audio-Visual Speech Separation

It is easy to modify our model from speech activity detection to speech separation. For speech separation, a common technique is to use a ideal ratio mask (or a ideal binary mask), which is a element-wise ratio (or binarized ratio) between clean and noisy spectrogram. This mask is then multiplied with the input spectrogram to get a denoised spectrogram. The audio input could be either complex spectrogram or magnitude spectrogram. In the case of complex spectrogram, two masks will be generated for real and imaginary component respectively. When using magnitude spectrogram, one mask will be generated, and the phase of noisy input will be used to as denoised phase. The speech activity detection model discussed above can be seamlessly transformed into a speech separation model.

5 Model Based Sensor Fusion: Separating Modality Invariant and Modality Dependent Information

In the previous section, we discussed a deep learning model for multimodal sensor fusion. Towards the end, we propose a co-learning idea which encourages co-adaption and independent learning of each sensor at the same time. In the following sections, we propose probabilistic models which combines multiple sensors’ signals into separate modality-dependent and modality-invariant features. We begin by introduce a variational RNN (VRNN) model where the transition between hidden units

are stochastic. The VRNN model is a non-Markovian model in the sense the conditional independence assumption is broken in both transition and emission. We discuss this property by comparing VRNN with hidden Markov model and linear dynamic system. Using VRNN, we build a multimodal VRNN which imposes a structural prior on the generative model. As a result, the modality dependent and modality invariant factors are encouraged to separate into different latent variables. We emphasize that as oppose to PCA or VAE where latent factors are not known a priori without looking at the posterior distribution, our model explicitly matches latent variables with concepts. This is a result of the explicit graphical model structure.

The non-Markovian model brings a price. One direct outcome is that the latent state no longer contains *all the information* for generation or transition. In this sense, it is not a state-space model in the classical sense. As the latent state does not contain all the information, it is questionable to use the latent state for many downstream tasks; for example, control and planning. Another outcome is that the time derivative information is not captured by the VRNN model as the latent variable only focuses on recovery of the current observation.

5.1 Motivation

Our motivation is driven by a few observations. First, we observe from human representation learning that feature representations are either modality-invariant or modality-dependent: when you listen to a person, your ears hear the speech, and your eyes watch the person’s face, then you jointly use the visual and audio signals to decode the content of the speech which is embedded in both audio and visual signals, and other information such as speaker’s visual identity and vocal accent which are uniquely embedded in visual or audio signal. With multiple input channels, we can infer the speech content with more precision than with a single input. For this reason, if we consider audio and visual signals are generated by some latent explanatory factors, it is natural to partition the latent factors into modality-dependent and shared subsets.

A second observation is many signals have a temporal structure. If we naively separate each time step as a independent unit for analysis, the data does not carry the same information. A typical approach is to formulate a sequence of latent random variables with Markov property to generate the observed sequence. Two common models are hidden Markov model and linear Gaussian models. However, neither of which are well-suited to modeling long-term dependencies and complex probability distributions over high-dimensional sequences. Neural network models such as RNN have seen many successes in modeling complex sequences with long dependency. One limitation with RNN is structural variations are captured by deterministic transition. While a RNN can increase its memory capacity by increasing the size of its hidden unit, it is more likely to overfit to training data. There is recent evidence that when complex sequences such as speech and music are modeled, the performances of RNNs can be dramatically improved when uncertainty is included in their hidden state. In this work, we combine stochastic latent variable with RNN to take advantage of both methods in a coherent way.

Considering learning a generative model for video of speech. Learning interpretable representations for such data, and comparing them as the speech content or the speaker identity are changed, gives useful high-level tools for speech recognition and speech synthesis. Even though each image is encoded by thousands of pixels and each audio segment is represented as hundreds of frequency bins, the data lie near a low-dimensional nonlinear manifold. A useful model must not only learn this manifold but also provide an interpretable representation of the speech dynamics. A natural representation from speech is that the human speech audio is divided into *phonemes*, and speech video is divided into *visemes* [Bear and Harvey, 2017]. Therefore an appropriate model might switch between discrete states with each state representing the dynamics of a particular action. These two learning tasks, identifying a speech manifold and a structured dynamics model, are complementary. We want to learn the speech manifold in terms of coordinates in which the structured dynamics fit well.

In our preliminary experiments, we used variational autoencoder (VAE) to map images of digit according to content and style by learning from the MNIST dataset. In Figure 5a we project the latent space to a 2 dimension space, and we can see clear clustering of digits, i.e. contents. The row dimension and column dimension show the variation of style of each digit. We also plot the hidden layer activation map in Figure 5b, and we can observe meaningful patterns corresponding to digits.




Figure 5: Separating content and style from MNIST digits.

5.2 Related Works

5.2.1 Inference in Graphical Models with Conjugate Prior

In Gaussian linear dynamic system model with linear Gaussian observation, because the observation model $p(y|x)$ is conjugate to the latent variable model $p(x)$, e.g. Gaussian, then the optimal approximate distribution is naturally Gaussian, and we can use efficient message passing algorithms to perform exact inference. However, when the observation model is not conjugate to the latent variable model, these algorithmic structure break down. A solution is to use general variational inference such as mean-field method. Mean-field method makes strong independence assumption on the latent distribution for ease of computation, which lead to greater gap between the marginal likelihood and variational lower bound.

5.2.2 Variational Autoencoder

Two common approaches for solving latent variable models are variational inference and Markov Chain Monte Carlo methods. Recently, stochastic variational inference [Hoffman et al., 2013] has been successfully combined with deep neural networks into a class of deep Gaussian models named Variational Autoencoder [Kingma and Welling, 2013, Rezende et al., 2014]. Different from a vanilla autoencoder, a VAE introduces a set of latent random variables \mathbf{z} to explain the variations in the observed variables \mathbf{x} . Their joint distribution is defined as:

$$p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}). \quad (21)$$

The VAE typically parameterizes $p(\mathbf{x}|\mathbf{z})$ with a highly flexible function approximator such as a neural network. The neural network allows for highly non-linear mapping from \mathbf{z} to \mathbf{x} which is a powerful and unique feature of VAE. However, introducing a highly non-linear mapping from \mathbf{z} to \mathbf{x} results in intractable inference of the posterior $p(\mathbf{z}|\mathbf{x})$. VAE uses a variational approach to approximate the posterior distribution $p(\mathbf{z}|\mathbf{x})$ while incrementally raises the evident lower bound:

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \log p(\mathbf{x}|\mathbf{z}) - KL(q(\mathbf{z}|\mathbf{x})|p(\mathbf{z})), \quad (22)$$

where $q(\mathbf{z}|\mathbf{x})$ is a distribution modeled using a inference neural network.

The generative model $p(\mathbf{x}|\mathbf{z})$ and the inference model $p(\mathbf{z}|\mathbf{x})$ are jointed optimized by maximizing the evidence lower bound. The expectation with respect to $q(\mathbf{z}|\mathbf{x})$ is approximated stochastically in VAE.

In VAE we estimate the distribution of latent variables \mathbf{z} and use this information to reconstruct original signal \mathbf{x} . Because \mathbf{z} is random rather than deterministic, VAE allows for reconstructing different \mathbf{x} from different modes of \mathbf{z} rather than a single point estimate as in regular auto-encoder.

5.2.3 Recurrent Neural Network

An RNN is a autoregressive model which takes a sequence input $\{\mathbf{x}_t\}_{t=1}^T$ to predict a sequence output $\{\mathbf{y}_t\}_{t=1}^T$. At each time step t , the RNN reads the input \mathbf{x}_t and updates its hidden state \mathbf{h}_t by

$$\mathbf{h}_t = f(\mathbf{x}_t, \mathbf{h}_{t-1}; \theta), \quad (23)$$

where f is a deterministic non-linear function with parameters θ . Common choices of function f can are gated activation functions such as LSTM or GRU.

While the function f is deterministic, we can add randomness to RNN by letting the output be a distribution. For example, assuming a sequence of lag 1 autoregressive relation, i.e. \mathbf{x}_t depends on $\{\mathbf{x}_k\}_{k=1}^{t-1}$. Conceptually, RNN models this sequence by parameterizing a factorization of the joint sequence probability distribution as a product of conditional probabilities such that

$$p(\mathbf{x}_1, \dots, \mathbf{x}_T) = \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{x}_{<t}) \quad (24)$$

$$= \prod_{t=1}^T g(\mathbf{h}_{t-1}; \tau) \quad (25)$$

where g is a function which maps the hidden state \mathbf{h}_{t-1} to the conditional probability distribution $p(\mathbf{x}_t | \{\mathbf{x}_k\}_{k=1}^{t-1})$.

We can model the output function g as being composed of two parts. The first part φ_τ is a function that returns the parameter set ϕ_t given the hidden state \mathbf{h}_{t-1}

$$\phi_t = \varphi_\tau(\mathbf{h}_{t-1}). \quad (26)$$

The second part of g returns the conditional distribution of \mathbf{x}_t as $p(\mathbf{x}_t | \{\mathbf{x}_k\}_{k=1}^{t-1}) = p(\mathbf{x}_t | \phi_t)$.

Gaussian mixture model (GMM) is a common choice for modeling a high-dimensional and real-valued distribution, especially for modeling structured output density model. For GMM, φ_τ returns a set of parameters including the mixture coefficients and the means and covariance matrices of each Gaussian components.

However, given f is a deterministic function, a RNN model does not have the stochastic transition in a HMM. The lack of a mechanism to model the structural variation imposes a restriction on the RNN, as when it attempts to encode sufficient input variability to capture the signal variations, it inevitably overfits noise variations. In order to prevent overfitting, we must limit the capacity of the RNN and find another mechanism to model signal variations.

5.2.4 Stochastic RNN

Modeling sequential data is a domain of interest to representation learning. In a temporal-aligned RNN, given a deterministic transition function, the only source of variability is the output distribution $p(y_t | x_{\leq t}) = p(y_t | x_t, h_{t-1})$. On the other hand, a limitation of standard HMM is that it is poor at capturing long-range correlation between the observed variables [Bishop et al., 2006]. Recently there is a trend in combining probabilistic models models such as state space model with deep neural networks [Chung et al., 2015]. The key is to incorporate some stochastic hidden states to RNN. Chung et al. [2015] introduced a sequence of latent random variables whose prior distribution at time step t is dependent on all the preceding inputs via a RNN hidden state \mathbf{h}_{t-1} . A similar model is proposed by Fraccaro et al. [2016]. These are stochastic RNN models, not state-space models in the strict sense as they have broken the conditional independence assumption in the emission model:

$$p(\mathbf{x}_{1:T} | \mathbf{z}_{q:T}) \neq \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{z}_t). \quad (27)$$

SRNN can be solved using variational inference algorithm. VAE at heart is a simple probabilistic graphic model of joint distribution of two time independent random variables \mathbf{x} and \mathbf{z} . When there is temporal relation, the assumption is broken. The recognition model would only try to encode the current observation into a latent variable and the generative model would decode the latent variable. Thus VAE would overlook the transition between latent states. In this sense, the VAE model overfits to the training data and discards the temporal information in the training data.




Figure 6: Graphical model of unimodal HMM.

5.2.5 State Space Model: HMM and LDS

A Hidden Markov Model (HMM) (Figure 6) is a doubly embedded random sequence whose underlying Markov chain is not directly observable, hence a hidden sequence. HMM model has three sets of parameters. The first set is the initial state distribution of the Markov chain, $p(\mathbf{z}_0)$. The second set is the transition distribution of the Markov chain, i.e. the conditional distribution $p(\mathbf{z}_t|\mathbf{z}_{t-1})$. The last set is the emission distribution $p(\mathbf{x}_t|\mathbf{z}_t)$. When \mathbf{z} has finite discrete state space, it is a Hidden Markov Model; when \mathbf{z} has continuous state space, it is a dynamical system. For a HMM model, we can choose arbitrary distributions, the efficient forward-backward algorithm can estimate the posterior distribution of the parameters, then use EM algorithm to improve the solution. For a review of HMM see [Rabiner, 1989]. A HMM requires a discrete state space with a known size, a dynamical system extends from finite state space to a continuous state space. A linear dynamical system simplifies a dynamical system by requiring $\mathbb{E}[\mathbf{z}_t|\mathbf{z}_{t-1}]$ to be a linear function of \mathbf{z}_{t-1} . Two common forms of LDS are the Kalman filter and Kalman Smoother. To get a linear time algorithm, in a Kalman filter, we take advantage of the conjugacy of exponential family. That is,

$$p(\mathbf{z}_t, \{\mathbf{x}_k\}_{k=0}^t) = p(\mathbf{x}_t|\mathbf{z}_t) \int_{\mathbf{z}_{t-1}} p(\mathbf{z}_t|\mathbf{z}_{t-1}) p(\mathbf{z}_{t-1}, \{\mathbf{x}_k\}_{k=0}^{t-1}) \quad (28)$$

which is equivalent to

$$p(\mathbf{z}_t|\{\mathbf{x}_k\}_{k=0}^t) = c_t p(\mathbf{x}_t|\mathbf{z}_t) \int_{\mathbf{z}_{t-1}} p(\mathbf{z}_t|\mathbf{z}_{t-1}) p(\mathbf{z}_{t-1}|\{\mathbf{x}_k\}_{k=0}^{t-1}) \quad (29)$$

where $c_t = p(x_t|x_1, \dots, x_{t-1})$ is a scaling factor.

The linear Gaussian restriction is that

$$p(\mathbf{z}_t|\mathbf{z}_{t-1}) \sim \mathcal{N}(\mathbf{A}\mathbf{z}_{t-1}, \Gamma) \quad (30)$$

$$p(\mathbf{x}_t|\mathbf{z}_t) \sim \mathcal{N}(\mathbf{C}\mathbf{z}_t, \Sigma). \quad (31)$$

The model can be written in the conventional Kalman filter form as

$$\mathbf{z}_t = \mathbf{A}\mathbf{z}_{t-1} + \mathbf{w}_t \quad (32)$$

$$\mathbf{x}_t = \mathbf{C}\mathbf{z}_t + \mathbf{v}_t \quad (33)$$

$$\mathbf{w}_t \sim \mathcal{N}(0, \Gamma) \quad (34)$$

$$\mathbf{v}_t \sim \mathcal{N}(0, \Sigma). \quad (35)$$

Note here we have omitted the control input \mathbf{u} which is assumed to be known in Kalman filter.

By recognizing the probability density function, we get the posterior distribution $p(\mathbf{z}_t|\{\mathbf{x}_k\}_{k=0}^t)$.

For Kalman smoother, we have the backward form

$$p(\mathbf{x}_t, \dots, \mathbf{x}_T|\mathbf{z}_{t-1}) = \int_{\mathbf{z}_t} p(\mathbf{x}_t, \dots, \mathbf{x}_T, \mathbf{z}_t|\mathbf{z}_{t-1}) \quad (36)$$

$$= \int_{\mathbf{z}_t} p(\mathbf{z}_t|\mathbf{z}_{t-1}) p(\mathbf{x}_t|\mathbf{z}_t) p(\mathbf{x}_{t+1}, \dots, \mathbf{x}_T|\mathbf{z}_t). \quad (37)$$

With forward-backward filters, we can get all the statistics we need for inference and learning. For example, we can easily compute the posterior probability of \mathbf{z}_t given the observed sequence (which is used in inference and learning)

$$p(\mathbf{z}_t|\mathbf{x}_1, \dots, \mathbf{x}_T) \propto p(\mathbf{z}_t, \mathbf{x}_1, \dots, \mathbf{x}_T)$$




Figure 7: Graphical model of multimodal HMM. Each X^m is a modality for $m \in \{1, \dots, M\}$. All modalities have the same shared latent explanatory variable Z^s .

$$= p(\mathbf{z}_t, \mathbf{x}_1, \dots, \mathbf{x}_t) p(\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | \mathbf{z}_t).$$

711 In speech recognition, in order to infer which word generated the sound, we need $p(\{\mathbf{x}_t\}_{t=1}^T | W)$
 712 where W is the word decides the probability distribution. With the forward-backward filter, we have

$$p(\{\mathbf{x}_t\}_{t=1}^T | W) = \int_{\mathbf{z}_t} p(\mathbf{z}_t, \mathbf{x}_1, \dots, \mathbf{x}_t) p(\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | \mathbf{z}_t) \quad (38)$$

$$= p(\mathbf{z}_T, \mathbf{x}_1, \dots, \mathbf{x}_T) \quad (39)$$

713 where \mathbf{z}_T is the final non-emitting state.

714 5.3 Multimodal HMM and LDS

715 First let's consider a multimodal Hidden Markov Model (HMM). Let $\{x_t^m\}_{t=1}^T$ denote the sequence of
 716 observed data of modality m , $\{z_t^m\}_{t=1}^T$ denote the sequence of latent explanatory factor of modality m ,
 717 $\{z_t^s\}_{t=1}^T$ denote the sequence of latent explanatory factor shared by all modalities. The shared latent
 718 factor is a key difference between multimodal HMM and regular HMM, as it enables the observed
 719 data to be correlated between modalities. When a event (e.g. speech) is described by multiple signal
 720 sequences, the sequence of shared latent variables contains modality invariant information such as
 721 the semantic content (e.g. speech content), whereas the modality-specific latent variables contain
 722 modality-dependent information such as styles pertains only to that modality (e.g. voice timbre, face
 723 contour). We define *style* latent variables to be modality-dependent, and *content* latent variables to be
 724 modality-invariant. We assume that at every time step t , a modality m has observed variable x_t^m that
 725 is generated by latent variables $\{z_t^m, z_t^s\}$. A graphical model of the sequences is shown in Figure 7.
 726 The latent variables $\{z_t^s\}$ and $\{z_t^m\}$ have Markov property such that $p(z_k^m | \{z_t^m\}_{t=1}^{k-1}) = p(z_k^m | z_{k-1}^m)$
 727 and $p(z_k^s | \{z_t^s\}_{t=1}^{k-1}) = p(z_k^s | z_{k-1}^s)$. Hence at time t , x_t^m is conditionally independent of the rest
 728 of the graph given $\{z_t^m, z_t^s\}$. Notice that the sequence $\{x_t^m\}_{t=1}^T$ is not temporally independent
 729 due to temporal dependency in $\{z_t^s\}_{t=1}^T$ and $\{z_t^m\}_{t=1}^T$. $\{x_t^m\}_{t=1}^T$ is also not spatially independent
 730 across m at fixed t due to the shared latent variable z_t^s . We assume that modality dependent latent
 731 variables are independent, i.e. z_k^m is independent of $z_k^{m'}$ if $m \neq m'$. The assumption of having
 732 orthogonal modality dependent latent factors is consistent with the objective to have disentangled
 733 latent dimensions [Bengio et al., 2013].

734 Using the Markov property, we can factorize the joint likelihood in a favorable form. To simplify
 735 notation, let's denote all latent variables at time step k by $\mathbf{z}_k = \{\{z_k^m\}_{m=1}^M, z_k^s\}$, and denote all
 736 observed variables at time step k by $\mathbf{x}_k = \{x_k^m\}_{m=1}^M$. For a sequence of T steps, the joint distribution
 737 of observed variables and latent variables is

$$\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_T, \mathbf{z}_1, \dots, \mathbf{z}_T) = \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t}) p(\mathbf{z}_t | \mathbf{z}_{< t}, \mathbf{x}_{< t}) \quad (40)$$

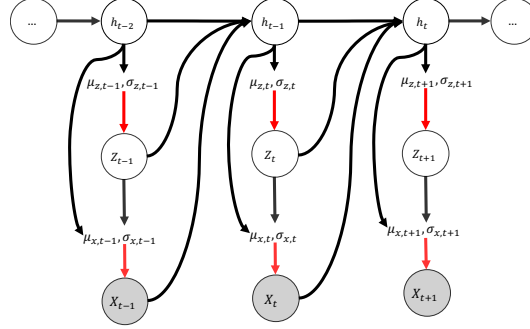


Figure 8: Stochastic RNN. Red line is stochastic relation, black line is deterministic relation. h are hidden states of a recurrent neural network, z are latent explanatory variable, x are observed variables, (μ, σ) are parameters of Gaussian distributions.

$$= \prod_{t=1}^T \left[\left(\prod_{m=1}^M p(x_t^m | z_t^m, z_t^s) \right) \left(p(z_t^s | z_{t-1}^s) \prod_{m=1}^M p(z_t^m | z_{t-1}^m) \right) \right] \quad (41)$$

where the second equality follows from Markov property.

Instead of solving this multimodal HMM model, next we will present a multimodal nonlinear dynamical system, and combine it with recurrent neural network, and solve it using stochastic variational inference.

5.4 Multimodal Variational RNN

Recurrent neural networks are able to represent long-term dependencies in sequential data by encoding inputs to update a deterministic hidden state. Recent works found that when complex sequences such as speech and music are modeled, the performance of RNN can be improved by including uncertainty in hidden states [Chung et al., 2015]. The reason is that when there is strong structural variation (i.e. high signal-to-noise ratio), it is mixed with random noise in inputs and outputs. While a deterministic RNN can increase its memory capacity by increasing the size of neural network, it also could bring over-fitting to training data, hence not a good solution.

The first order Markov property in HMM model is an effort to have a compromise between computation complexity and model complexity. However, the marginalization step in finding posterior distribution of z is often intractable. Only a few exact inference algorithms are available (i.e. HMM and Kalman filter).

Having $p(z_t | z_{t-1}) = p(z_t | z_{t < t}, x_{< t})$ would be ideal since it allows for maximum information passing. However, if we consider having a RNN layer on top of the latent variables as shown in Figure 8, the RNN hidden state variable h_t is going to pass all the previous $\mathbf{z}_{\leq t}$ and $\mathbf{x}_{\leq t}$ to timestep t , which breaks down the Markov property. In that case we still can factorize the joint distribution as

$$\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_T, \mathbf{z}_1, \dots, \mathbf{z}_T) = \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t}) p(\mathbf{z}_t | \mathbf{z}_{< t}, \mathbf{x}_{< t}). \quad (42)$$

This long dependency is challenge to work with in dynamic Bayesian networks. Chung et al. [2015] instead used neural networks to replace the transition probability matrix and emission probability, and solve the estimation problem using stochastic variational Bayes [Kingma and Welling, 2013, Rezende et al., 2014].

In this work, we propose a multimodal stochastic RNN model (Figure 9). In the generative model, there are $M + 1$ RNN sequences $\{\{h^m\}_{m=1}^M, h^s\}$, connecting to $M + 1$ latent explanatory variable $\{\{Z^m\}_{m=1}^M, Z^s\}$. M of the latent variables $\{Z^m\}_{m=1}^M$ are each associated with one modality, and Z^s is shared by all modalities. As we explained earlier, each modality's output X^m is a function of Z^m and Z^s . The RNN are deterministic and the latent explanatory variables are stochastic. By doing so, we let the latent variables model the structural variation in signals, and let the RNN pass information from the past to present. The observed variables stochastically depend on latent variables,




Figure 9: Generative model of multimodal stochastic RNN (Figure only shows one modality m). Red line is stochastic relation, black line is deterministic relation. h are hidden states of a recurrent neural network, Z are latent explanatory variable, X are observed variables, (μ, σ) are parameters of Gaussian distributions.




Figure 10: Inference Model of multimodal stochastic RNN (Figure only shows one modality m). Red line is stochastic relation, black line is deterministic relation. h are hidden states of a recurrent neural network, z are latent explanatory variable, x are observed variables, (μ, σ) are parameters of Gaussian distributions.

769 which allows for noise variations. By have two separate stochastic component, we separate structural
770 variation from random noise variation.

771 For inference, we consider the following q function as shown in Figure 10 to approximate the posterior
772 distribution $p(\mathbf{z}_1, \dots, \mathbf{z}_T | \mathbf{x}_1, \dots, \mathbf{x}_T)$

$$\begin{aligned} q(\mathbf{z}_1, \dots, \mathbf{z}_T) &= \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{x}_t, \mathbf{h}_{t-1}) \\ &= \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{z}_{<t}, \mathbf{x}_{\leq t}). \end{aligned} \quad (43)$$

773 The last equality holds because in the generative model (Figure 9), \mathbf{h}_{t-1} depends on all the \mathbf{x}_k and
774 \mathbf{z}_k up to time $k = t - 1$. Hence our inference model (Figure 10) implies (43).

775 Thus we can derive the variational lower bound as

$$\log p(\{\mathbf{x}_t\}_{t=1}^T) \geq \mathbb{E}_{q(\{\mathbf{z}_t\}_{t=1}^T)} \log \left[\frac{p(\{\mathbf{x}_t, \mathbf{z}_t\}_{t=1}^T)}{q(\{\mathbf{z}_t\}_{t=1}^T)} \right] \quad (44)$$

$$= \sum_{t=1}^T \sum_{m=1}^M \mathbb{E}_{q(\{\mathbf{z}_k\}_{k=1}^t)} \log p(x_t^m | z_{\leq t}^m, z_{\leq t}^s, x_{<t}^m) \quad (45)$$

$$- \sum_{t=1}^T \sum_{m=1}^M KL(q(z_t^m | \mathbf{z}_{<t}, \mathbf{x}_{\leq t}) || p(z_t^m | z_{<t}^m, x_{<t}^m)) \quad (46)$$

$$- \sum_{t=1}^T \sum_{m=1}^M KL(q(z_t^s | \mathbf{z}_{<t}, \mathbf{x}_{\leq t}) || p(z_t^s | z_{<t}^s, \{x_{<t}^m\}_{m=1}^M)). \quad (47)$$

776 Here, the first term is the reconstruction loss. The first set of KL terms is the modality specific latent
777 variable posterior approximation error, and the second set of KL terms is the shared latent variable

posterior approximation error. A key difference between RNN-VAE and regular VAE is the prior distribution for a latent variable Z_t is not the same Gaussian($0, I$). Instead, the prior depends on the previous states $Z_{<t}$. Here, we define the prior distribution to be a Gaussian distribution whose parameters are given by a neural network with input h_{t-1} (Figure 9).

6 Learning Latent Variable from Similarity

6.1 Motivation

Measuring the distance between objects has been useful in many clustering or classification tasks. People normally first find a mapping, called a embedding, of a object into a latent space, e.g. a Euclidean space. The mapping is normally hand-crafted. Here we discuss methods which learn a explicit embedding function from only knowing local similarity between sample points, not necessary distance measure, but only label of categories. Such a function can be updated as new data are collected, hence reducing the algorithm complexity. We propose a model which combines embedding and sensor fusion and discuss applications of such approach for tasks such as robust sound event detection and classification.

6.2 Global Embedding from Local Pairwise Similarity

Most methods for learning a latent embedding involve inverting a matrix of pairwise similarity between sample points, e.g. Locally Linear Embedding [Roweis and Saul, 2000] and ISOMAP [Tenenbaum et al., 2000]. In a high-dimensional input space, the cost of inverting such a matrix is $O(n^3)$, and often the matrix is sparse. While there are iterative methods for approximating a solution, there are other methods motivated on a objective function which can be naturally optimized using gradient descent, e.g. *Stochastic Neighborhood Embedding* (SNE) [Hinton and Roweis, 2002] and t-SNE [Van der Maaten and Hinton, 2008]. While LLE and ISOMAP lack a explicit embedding function, SNE and t-SNE models the similarity between two sample points using conditional probability, e.g. Gaussian. Suppose x_i would pick all its neighbors in proportion to their probability density under a Gaussian centered at x_i , the similarity of sample point x_i to x_j is the conditional probability $p_{j|i}$:

$$p_{j|i} = \frac{\exp(-||x_i - x_j||^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-||x_i - x_k||^2)} \quad (48)$$

Here σ_i is the variance of the Gaussian that is centered on x_i . In the embedding space (i.e. latent space), a lower dimensional counterpart z_i has a similar conditional probability, which is denoted by $q_{j|i}$:

$$q_{j|i} = \frac{\exp(-||z_i - z_j||^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-||z_i - z_k||^2)} \quad (49)$$

If the latent points z_i and z_j correctly model the similarity between the high-dimensional sample points x_i and x_j , the conditional probability $p_{j|i}$ and $q_{j|i}$ will be equal. Motivated by this observation, SNE finds a latent representation which minimizes the difference between $p_{j|i}$ and $q_{j|i}$ for all i, j . SNE minimizes the sum of Kullback-Leibler divergences over all sample points using a gradient descent method:

$$C = \sum_i KL(P_i || Q_i) = \sum_i \sum_j p_{j|i} \log \frac{p_{j|i}}{q_{j|i}} \quad (50)$$

One of the advantages of t-SNE is when new sample points are available, current embedding can be updated with low cost. For example, for a new point x^* , we can initialize its latent code z^* by the mean of sample latent code, and perform gradient descent for a few iterations.

While SNE is a unsupervised learning approach, its result can be used for classification. After a embedding function f is trained to convergence, it can be applied on new data for tasks involving similarity search. For example, Van der Maaten and Hinton [2008] trained a embedding function of MNIST dataset whose latent space shows good clustering results of digits. A new digit image can be send to the embedding function, and we can perform a k-nearest neighbor search to find its more probable label. Such a search is often cheaper since the embedding space is low-dimensional.

6.3 Weakly Supervised Metric Learning without a Metric

A drawback of SNE and t-SNE is the similarity metric is imposed with a structural form, e.g. Gaussian kernel, on both the input and latent space. Hadsell et al. [2006] proposed a *siamese structure* to learn a globally coherent non-linear function that maps the data evenly in a latent space. The learning relies solely on neighborhood relationships and does not require any distance measure in the input space.

Siamese structure is a weakly supervised model in the sense that the training data only use binary labels to learn a continuous embedding. Recognizing that a meaningful mapping maps similar input vectors to nearby points on the output space and dissimilar vectors to distant points, siamese structure constructs a contractive loss function whose minimization can produce such a embedding function. The loss function takes pairs of samples x_i and x_j , with a binary label $y = 1$ indicating they are similar and $y = 0$ indicating they are dissimilar. In classification, samples from same classes are considered as similar, and from different classes are considered as dissimilar. The output of the function would naturally allocate each sample according to its representativeness in each class without any distance measure. Define the distance function D_W as euclidean distance between the embedding output $G_W(x_1)$ and $G_W(x_2)$ where G_W is the embedding function to be learned:

$$D_W(x_1, x_2) = \|G_W(x_1) - G_W(x_2)\|. \quad (51)$$

The contrastive loss function \mathcal{L} is defined as

$$L(W, (x_1, x_2, y)) = (1 - y)L_p(D_W) + yL_n(D_W), \quad (52)$$

where

$$\mathcal{L}(W) = \sum_{i=1}^N L(W, (x_1, x_2, y)^{(i)}) \quad (53)$$

where $(x_1, x_2, y)^{(i)}$ is the i -th sample pair, L_p is the loss function for positive pairs and L_n is the loss function for negative pairs. Here L_p and L_n are designed to reflect similarity in the embedding space; for example, Hadsell et al. [2006] suggested using $L_p = D_W(x_1, x_2)^2$ and $L_n = \max\{0, (m - D_W(x_1, x_2))^2\}$ where $m > 0$ is some hyper-parameter represents a margin for dissimilar pairs.

Bell and Bala [2015] proposed a *siamese network* which utilizes deep neural network for G_W . The siamese network showed good result in embedding images of objects into a latent space according to their visual similarity represented by class labels.

6.4 A Example: Audio Event Detection in Noisy Environment

Analysis of environmental sound has the potential to be used in many applications, such as surveillance and smart homes. The Detection and Classification of Acoustic Scene and Events (DCASE) challenge is a venue for researchers to propose new methods for audio classification. Several tasks has been defined for audio classification including acoustic scene classification, sound event detection, and audio tagging. Recently, Google released an ontology and human-labeled large scale data set for audio events, i.e. Audio Set [Gemmeke et al., 2017], which consists of 527 classes and over 2 million human-labeled 10-second long sound clips drawn from YouTube videos. Audio Set is defined for tasks such as audio tagging. The objective of audio tagging is to perform multi-label classification on fixed-length audio chunks without predicting the precise boundaries of acoustic events. Recently, deep learning methods have been successfully applied to audio tagging.

Siamese network has been used for sound event detection [Zhang and Duan, 2016, 2017]. Here we explore sound event detection in a noisy environment by combining denoising autoencoder, siamese network, and conditional attention mixture model (Figure 11). We use siamese network to train a embedding encoder for a map of clean sound categories in e.g. Audio Set based on category label. Because we are not performing speech separation, we do not have to stick to STFT features as input to the network. Instead, we use pretrained features such as Audio Set VGGish features which are proved to be more effective than STFT features [Hershey et al., 2017]. We also compare other audio features such as mel-frequency cepstral coefficient (MFCC) features and SoundNet features [Aytar et al., 2016]. We extract VGGish features from audios for downstream tasks as described below.

In a noisy environment, the clean sound is mixed with noise, such that the trained siamese encoder would not embed noisy sound to the exact embedding location of the clean sound. To this end, we

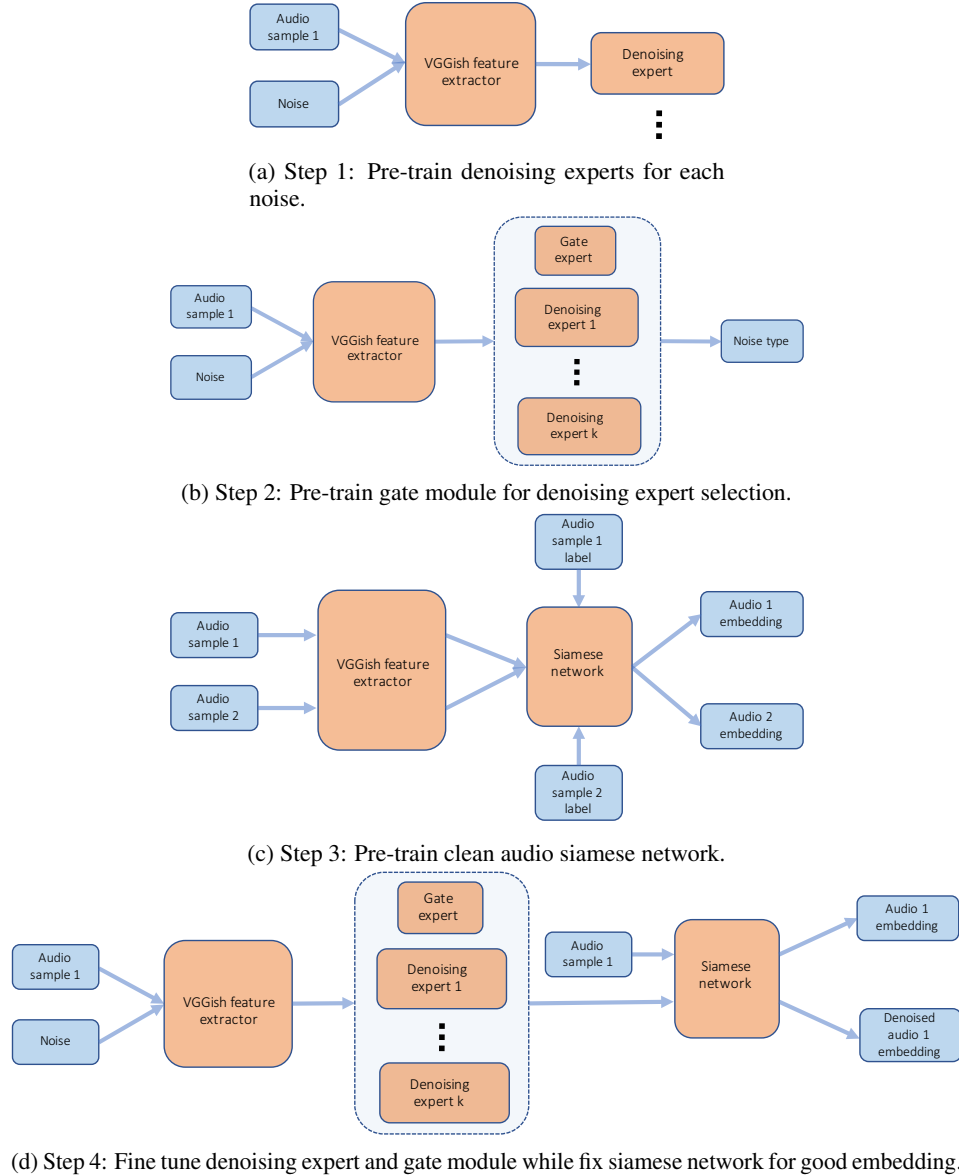


Figure 11: Robust sound event detection model.

867 train a denoising autoencoder (DAE) [Vincent et al., 2010] such that the denoised sound when passed
 868 through the trained siamese encoder would map to the original clean sound embedding location as
 869 close as possible. That is, the output of the DAE is connected to the input of the siamese network to
 870 get a embedding. The loss is the difference between clean-sound embedding and denoised-sound
 871 embedding. During back-propagation, we can (1) only update the weights for the denoising network,
 872 while fixing the siamese network, or (2) only update the weights for the siamese network, while fixing
 873 the denoising network, or (3) update both. Using the siamese embedding of the denoised sound, we
 874 can classify the class of the sound event by nearest neighbor search using euclidean distance in the
 875 embedding space. This has the advantage of searching by *audio similarity* [Zhang and Duan, 2017].

876 Because noise type can vary, we apply conditional attention mixture model for dynamic selection
 877 of DAE. The system is developed in four stages (Figure 11). We first pick a few kinds of noises,
 878 then pre-train one DAE for each of them (Figure 11a). On top of pre-trained DAE's, we add the
 879 conditional attention mixture model (Equation 17) where each DAE is a expert network, and let the
 880 gate module dynamically decide which kind of noise is most likely. This is supervised training since
 881 we know which DAE is correct during training time (Figure 11b). Meanwhile in parallel, we train a

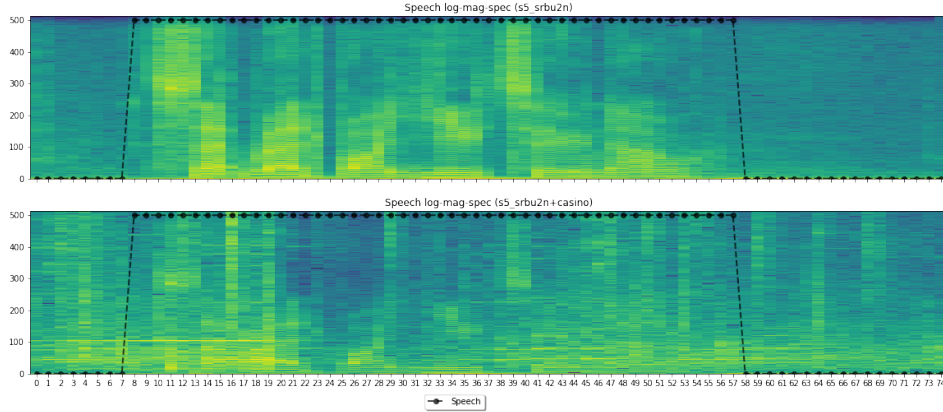


Figure 12: Log-mag-STFT of speech. Top figure shows speech, bottom figure shows speech plus casino noise with SNR = -5DB. Dotted black line shows speech activity.

882 siamese network on clean audio (Figure 11c). After both the gate module and siamese network are
 883 trained, we connect the denoising network with the siamese network (Figure 11d). We further fine
 884 tune the DAE's based on the distance of embedded pairs while holding siamese network as ground
 885 truth mapping. Recall the conditional attention mixture model applies a probability to each expert
 886 module, hence the choice of denoising network could be combined, and consequently the embedding
 887 into latent space could be a set of points of different weights. During test time, we can use a weighted
 888 mean of these points as the center for nearest neighbor search.

889 7 Experiments

890 We will first give a brief overview of speech activity detection and speech separation. After that
 891 we will introduce the experiment dataset and video and audio preprocessing protocol. We discuss
 892 some relevant existing deep learning techniques that can be used to improve model performance,
 893 e.g. pre-trained feature extractor. Then we explain the experiment set-up and results. Finally we
 894 discuss the experiment results and propose follow-up works to be completed in our follow up works.
 895 Due to time constraints, we are not able to complete all the works we hoped at this time. We plan to
 896 complete them as part of the dissertation thesis.

897 7.1 Speech Activity Detection and Speech Separation

898 Acoustic event detection (AED) has draw many attentions due to its wide application in intelligent
 899 systems such as robots and smart homes [Gemmeke et al., 2017, Hershey et al., 2017]. With new tools
 900 such as deep learning, we have seen significant improvement in the results in recent DCASE Acoustic
 901 Scene Classification (ADC) task. As one kind of AED tasks, Speech activity detection (SAD) is a
 902 classification problem of a given sequence of audio frames into speech active and non-active states.
 903 Most SAD models rely on audio signal, for example Li et al. [2017] used grid LSTM to detect
 904 speech endpoints. Although speech is audio signal, video signal has shown value to detection and
 905 understanding of speech. Sadly, most SAD systems developed so far either entirely rely on audio or
 906 video alone, only a limited number of systems utilize both audio and video signal, e.g. [Ariav et al.,
 907 2018]. In addition, it is not a straightforward task to effectively combine audio and video signals due
 908 to the natural differences between audio and video signals.

909 A unimodal system that relies solely on the audio signals can fail to do this job due to the common
 910 artifacts found in the real-world speech signals such as additive noise and reverberation. To give a
 911 concrete example, one speech audio is displayed as log transformed magnitude Short Time Fourier
 912 Transform (log-mag-STFT) in Figure 12. Notice although the speech only appears between 8th
 913 frame and 57th frame, it is difficult to determine where are true speech activities due to the noise we
 914 manually injected. In this scenario, video provides a opportunity to improve the accuracy of speech
 915 activity detection.




Figure 13: Images of speech (3 seconds, 25fps). Speech activity is from frame 8 to frame 57.

While video signal is invariant to acoustic environments, the relation between speech and face movement is dynamic and usually asynchronous. For example, when uttering some words, part of the lip movements (i.e. visemes) are loosely coordinated with the sound (i.e. phoneme). In this sense, two input modalities are not independent. Visually speaking, speech is closely related to mouth movements. Hence the visual speech activity detection problem can be easily contaminated by non-speech mouth movement such as breathing, eating, etc. For example, in Figure 13, speech starts at 8th frame, but there is noticeable mouth movement from frame 0 to 7. Phonetically speaking, vowel is a speech sound tends to require relatively open mouth, while a consonant is a sound made with mouth relatively closed. Hence speech does not necessarily activates expressive mouth movements, which puts challenge on video-only speech activity detection. More importantly, visual signals do not always provide a reliable cue due to variations such as head rotations, illuminations, different view points, etc.

We want to comment that video based Automatic Speech Recognition (ASR) has been addressed in a few previous researches; for example Assael et al. [2016]. These models directly classify video segments into words or phonemes, and hence can be used for SAD. We approach the SAD problem from a different perspective. We consider ASR as a multi-stage system, and SAD is a key step in front end processing. Instead of using video to predict text, we use visual cue to assist audio cue to classify noise audio from noisy speech audio. When noisy speech audio can be located (e.g. black line in Figure 12), we can apply source separation methods to extract clean speech from noisy speech. The denoised speech is subsequently fed into a audio based ASR system. To this end, video based ASR can work with audio based ASR in our model jointly in a multimodal ASR task.

A SAD system can be applied to convert a unsupervised speech separation systems into a supervised one. Consider a noisy acoustic environment, if the type of noise is unknown, we have a unsupervised speech separation problem. If we know the type of noise, we have a supervised speech separation problem, which is significantly less challenging than the unsupervised case. With a SAD system, we can detect the noisy period immediately before speech. If we assume that the same type of noise will continue during speech, then we can use dictionary based speech separation algorithms such as [Smaragdis et al., 2007] and [Guo and Kim, 2018] with a known dictionary, hence a supervised speech separation system. In a multi-person speech separation scenario, SAD system can assist a face recognition algorithm to determine which person is speaking, and use beam-forming to enhance that speaker’s speech.

We emphasize again it is easy to modify our model for speech activity detection to speech separation. For speech separation, a common technique is to use a ideal ratio mask (or a ideal binary mask), which is a element-wise ratio (or binarized ratio) between clean and noisy spectrogram. This mask is then multiplied with the input spectrogram to get a denoised spectrogram. The audio input could be either complex spectrogram or magnitude spectrogram. In the case of complex spectrogram, two masks will be generated for real and imaginary component respectively. When using magnitude spectrogram, one mask will be generated, and the phase of noisy input will be used to as denoised phase. The speech activity detection model discussed above can be seamlessly transformed into a speech separation model.

7.2 Data

In this experiment we use the GRID corpus data [Cooke et al., 2006] which contains 34 speakers. Each speaker has 1000 short speeches, each about 3 seconds long, recorded in a controlled environment with limited ambient noise. The speech is fully annotated, with phonemes mapped to time. Figure 14 shows a sample image of one video.

Due to limitation in computing resources, in this pilot study we randomly choose a subset of subjects to develop our model. For the experiment results presented in the following sections, we use 4 subjects, 2 males and 2 females, to train and test our model. For each subject, we randomly partition the video clips by 7 : 2 : 1 into training, validation, and in-sample testing sets. We have a 5th subject served as out-of-sample testing set. We are planning on extend to all subjects in following works.

7.3 Video Pre-processing and Pre-training

Each video clip is 3 seconds long with 25 fps. We pre-process the video data using opencv-python. However, due to ffmpeg issue, some unpacked videos have less than 75 frames. To align the video




Figure 14: Video image of GRID.

with audio, we pad the first and last frames to the front and end of video to increase the number of frames to 75.

In visual activity detection, still image and motion between frames have been shown to complement each other [Simonyan and Zisserman, 2014a]. Therefore, we extract two kinds of information from video: image and motion. We first extract a face from each video frame. Because the video is relatively clean with speaker's front face at the center the video, we found a Viola-Jones type face detector works well. We used opencv-python Harr cascade front face feature to detect faces. We are able to detect high quality face bounding box because the camera is directly pointing at face of the subject, and both camera and subject are stationary. However, there is still movements of subjects between frames, such that each bounding box shifts a little. Because the subject in each video is always stationary and the movement is not dramatic, we choose to use the bounding box in first frame for the 75 frames. Each face image is resized to 224×224 pixels in order to feed into a InceptionV3 feature extraction.

To create image features, pre-trained feature extractors can be used such as AlexNet [Krizhevsky et al., 2012], VGG [Simonyan and Zisserman, 2014b], InceptionV3 [Szegedy et al., 2016] and ResNet-50 [He et al., 2016]. In this experiment, we use InceptionV3 pre-trained on ImageNet without fine-tuning. The extracted InceptionV3 feature has dimension of 2048 for each $224 \times 224 \times 3$ image input. Alternatively, with enough data and computing resource, we could also train the feature extractor from scratch on GRID images or a general human face dataset, and fine-tune with the rest of the model. We will pursue these works in future.

Convolutional network trained on multi-frame dense optical flow has been shown to achieve good performance on capturing human motion [Simonyan and Zisserman, 2014a]. We compute dense optical flow using Farnback algorithm provided in opencv-python from the extracted face images. Optical flow is $224 \times 224 \times 2$ for each frame.

To create a corrupted video signal, we randomly apply a square masking patch, and randomly adjust the brightness of that patch area. The location of the patch is random for each video, and the size of the patch is randomly chosen between 50 and 150. For each video, we randomly pick a segment of 20 to 30 frames, and apply the patch to all frames in that segment. This mimics the effect of having a shadow or spot-light casting on the subject's face during a segment of a video clip. Figure 15 shows a example of corrupted image and corrupted optical flow of the same video.

7.4 Audio Pre-processing and Pre-training

For audio data, we first pre-process the input audio to synchronize with images extracted from video by matching number of audio frames with the number of image frames. We choose the STFT window size and re-sample audio such that there are 25 frames per second for audio. We first down-sampled the speech audio from 50kHz to 16kHz, then perform Short-Time-Fourier-Transform with a window size of 1024 and a stride of 760 which transforms each audio into 75×513 representation.




Figure 15: Corrupted images and optical of speech (subject 11, speech srbt7s, 3 seconds, 25fps). Speech activity between 9th and 52nd frames, video noise between 22nd and 49th frames.




Figure 16: A example of corrupted speech. Top clean speech, bottom corrupted speech (subject 11, speech srbt7s, -5db SNR with Casino noise). Black line marks speech activity, red line marks visual noise (audio noise is presented in all frames).

1005 Notice that 25 frames per second may not be sufficient to fully capture the dynamics of speech. Our
 1006 method can be easily extended to extract 50 or 75 audio frames per second.

1007 Similar to image feature extractor, we consider audio features such as mel-frequency cepstral
 1008 coefficient (MFCC) features, SoundNet features [Aytar et al., 2016] and Audio Set VGGish [Hershey
 1009 et al., 2017]. Because MFCC and VGGish features are not invertible, in speech separation, we use
 1010 STFT features as input to the network. In acoustic event detection, MFCC, SoundNet and VGGish
 1011 are proved to be more effective than STFT features [Hershey et al., 2017]. We will pursue these
 1012 works in future.

1013 To create a corrupted speech signal, we injected two kinds of noise to the audio. The first kind
 1014 is a casino noise proposed by [Duan et al., 2012]. The casino noise is challenging due to being a
 1015 non-stationary, wide-band noise which span from low to high frequency domain. Figure 16 shows a
 1016 example.

1017 We construct a second kind of noise using a mixture of three person from TIMIT training data
 1018 (TIMIT-3P). TIMIT training set contains 136 female and 326 male speakers, while the testing set
 1019 contains 56 female and 112 male speakers, which are from eight dialect regions in the US. Each
 1020 TIMIT speaker has 10 short utterances. For each GRID speech, we randomly choose 3 speeches from
 1021 TIMIT, then choose a random segment of 3 seconds from each chosen speech. We mix the 3 speeches
 1022 with equal magnitude. Finally, we mix the 3-person mixture with GRID speech with SNR equal 0
 1023 DB.

1024 7.5 Training

1025 During each training epoc, a minibatch of size 16 is used. Same minibatch size is used for evaluating
 1026 validation and testing set. For RNN we used a 5 frames for truncated backpropagation through time.
 1027 For DNN model, we feed 5 audio frames per sample to provide a context, which is a common practice
 1028 in audio based speech recognition system. All models are trained on a server with 4 Nvidia Titan X
 1029 GPUs with 12GB RAM per GPU. The multimodal RNN model takes about 1-4 hours ¹ per training
 1030 epoc, and we trained it for three days of 30 epocs. The recurrent attention mixture model takes about
 1031 4 hours per epoc to train .

1032 7.6 Experiment Results

1033 Our first set of experiments is on the recurrent attention filter. We implemented four versions of the
 1034 model. The first one is the conditional attention mixture model (Equation 17). The second one is the

¹It seems the training speed is fluctuating on the server for some unknown reason.

Table 1: Experiment Results

	Train(*)	Test (+)
Uni-modal input		
Image	0.12	80.8
Optical flow	0.03	90.5
Audio (with casino noise)	0.06	92.5
Audio (with 3 persons noise)	0.06	91.1
Multimodal input without co-learning		
Conditional Attention Mixture (with casino noise)	0.03	96.38
Conditional Attention Mixture (with TIMIT-3P noise)	0.04	96.25
Recurrent Attention Mixture (with casino noise)	0.03	93.75
Recurrent Attention Mixture (with TIMIT-3P noise)	0.1	85.7
Multimodal input with distance based co-learning		
Conditional Attention Mixture (with casino noise)	0.03	91.74
Conditional Attention Mixture (with TIMIT-3P noise)	0.1	84.56
Recurrent Attention Mixture (with casino noise)	0.03	89.37
Recurrent Attention Mixture (with TIMIT-3P noise)	0.1	76.23
*: MSE, +: % correct.		

recurrent attention mixture model (Equation 15). For each of these two models, we construct two versions, i.e. with and without co-learning (distance based regularizer). For comparison, we also implemented three uni-modal DNN models (audio, optical flow, and image). The results are shown in Table 1.

First we notice that multimodal input model does outperform unimodal input model. While the no-audio models (i.e. image and optical flow) have lower accuracies than the audio-only model, most multimodal input models outperform audio-only model, which suggests image and motion provides addition information not seen in audio for speech activity detection. This validates our hypothesis that multimodal input is useful to speech related tasks.


Second, we notice that the RNN model easily overfits the training data. While the recurrent attention model has lower training error than the conditional attention model, during test time the result is reversed. This suggests we may need (1) to regularize RNN such using e.g. dropout or (2) stop RNN training earlier.

Third, we notice that a simple distance based regularizer for co-learning does not improve prediction accuracy. On the contrary, the un-regularized models have higher prediction accuracy, very likely due to the lack of restriction of the model. Recall our motivation for the model-based regularization is that co-learning helps to separate modality invariant features from modality-dependent features, however, such features may not be useful for speech activity detection. This suggests we may need to carefully design the co-learning structure in order to extract features that are useful to a task.


7.7 Spatial Attention Outputs

We would like to discuss the spatial attention generated by the gate module. Recall the spatial attention gives the mixing weights to the experts. When a sensor's signal is corrupted, we expect the mixing weight for that sensor would decrease, and the mixing weight for a robust sensor would increase, such that the overall decision is more likely to be correct. In Figure 18 we show the attention weights for a speech sample. Notice that speech begins at the 9th frame (black line), and the model prediction lags only one frame. The interesting phenomenon is that at the 22nd frame, the video stream changes to corrupted (red line), causing the mixing weight for audio expert to increase, while the mixing weights for image and optical flow decrease. This suggests that the attention module has learned to identify which sensor inputs are more reliable and assigns mixing weights accordingly. At the 51st frame, the video stream changes to uncorrupted, and we see the mixing weight for image increases to the largest. Recall we have overlap speech audio with noise audio of the entire speech duration, hence video signal is more reliable when uncorrupted.

We notice throughout the entire speech, the attention module prefers to choose only one of the three inputs as the dominant input. It is likely due to the nature of the relationship between the three inputs we have in this particular experiment. Another possibility is the softmax function we used for generating mixing weights may prefer sparse outputs. In order to investigate this property, we can (1) use a different set of sensor inputs, and (2) modify or replace the softmax function to encourage non-sparse outputs.




(a) Expert network outputs.




(b) Attention network outputs.

Figure 17: Model outputs (subject 11, speech srbt7s, mixed with casino noise at -5db.



(a) Expert network outputs.



(b) Attention network outputs.

Figure 18: Model outputs (subject 11, speech srbt7s, mixed with 3 random person noise at 0db.)

1073 8 Conclusion and Future Works

1074 In our speech activity detection experiment, we found that the best multimodal model has 95%
 1075 accuracy on testing set, whereas for flow model and audio model, the accuracy is 90 – 92%, and
 1076 85% for image model. The multimodal model prediction is more accurate and stable than unimodal
 1077 models. This suggests that multimodal model has successfully combined different sensor data for a
 1078 common task. We also found the attention function dynamically estimates how reliable each expert is
 1079 and assigns weights accordingly. Hence the gate module successfully opened the black-box of sensor

fusion, which provides insight to the relation between signal and sensor fusion process. The result shows a small step towards a structured sensor fusion method. Speech activity detection involves a sequential binary classification problem. As we discussed, another interesting problem is speech separation (e.g. denoising). We are currently working on this problem.

We found that co-learning using distance-based regularizer decreases prediction accuracy. This is possibly due to co-learning separates modality invariant features and is not designed for speech activity detection. We also proposed to combine probabilistic graphical model with deep neural network to construct a model-based sensor fusion model. The model has a co-learning design which tries to separate modality-invariant and modality-dependent features. We are pursuing this direction at the moment. We would like to investigate if model-free and model-based co-learning can successfully separate modality invariant features from modality dependent features.

The combination of denoising autoencoder, siamese network, and conditional attention mixture model is another work we are currently working on as part of the author's dissertation. Although we choose audio event detection as an application, this framework is general enough to solve other problems, e.g. speaker identification and video event detection in multimodal setting. We will explore these possibilities in future.

Acknowledgement

Part of this work is based on Lijiang Guo's PhD qualify exam paper. We would like to thank Dr. Geoffrey Fox, Dr. Minje Kim, Dr. Francesco Nesta, Dr. Michael Ryoo and Dr. Lantao Liu for helpful discussions.

References

- Robert H Shumway and David S Stoffer. *Time series analysis and its applications: with R examples*. Springer, 2011.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- Ariel Ephrat, Tavi Halperin, and Shmuel Peleg. Improved speech reconstruction from silent video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 455–462, 2017.
- Yannis M Assael, Brendan Shillingford, Shimon Whiteson, and Nando de Freitas. Lipnet: Sentence-level lipreading. *arXiv preprint arXiv:1611.01599*, 2016.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 689–696, 2011.
- Youssef Mroueh, Etienne Marcheret, and Vaibhava Goel. Deep multimodal learning for audio-visual speech recognition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2130–2134. IEEE, 2015.
- Weijiang Feng, Naiyang Guan, Yuan Li, Xiang Zhang, and Zhigang Luo. Audio visual speech recognition with multimodal recurrent neural networks. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 681–688. IEEE, 2017.
- Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang. Audio-visual speech enhancement based on multimodal deep convolutional neural network. *arXiv preprint arXiv:1709.00944*, 2017.
- Ido Ariav, David Dov, and Israel Cohen. A deep architecture for audio-visual voice activity detection in the presence of transients. *Signal Processing*, 142:69–74, 2018.
- DeLiang Wang and Jitong Chen. Supervised speech separation based on deep learning: An overview. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(10):1702–1726, 2018.

- 1125 John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. Deep clustering: Discriminative
1126 embeddings for segmentation and separation. In *Acoustics, Speech and Signal Processing (ICASSP),
1127 2016 IEEE International Conference on*, pages 31–35. IEEE, 2016.
- 1128 Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen. Permutation invariant training of
1129 deep models for speaker-independent multi-talker speech separation. In *2017 IEEE International
1130 Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 241–245. IEEE, 2017.
- 1131 Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R Hershey. Single-channel
1132 multi-speaker separation using deep clustering. *arXiv preprint arXiv:1607.02173*, 2016.
- 1133 Aviv Gabbay, Asaph Shamir, and Shmuel Peleg. Visual speech enhancement. *arXiv preprint
1134 arXiv:1711.08789*, 2017.
- 1135 Aviv Gabbay, Ariel Ephrat, Tavi Halperin, and Shmuel Peleg. Seeing through noise: Visually driven
1136 speaker separation and enhancement. In *2018 IEEE International Conference on Acoustics, Speech
1137 and Signal Processing (ICASSP)*, pages 3051–3055. IEEE, 2018.
- 1138 Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T
1139 Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: A speaker-independent
1140 audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*, 2018.
- 1141 Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory
1142 features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648,
1143 2018.
- 1144 Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching
1145 unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages
1146 35–53, 2018.
- 1147 Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio
1148 Torralba. The sound of pixels. In *Proceedings of the European Conference on Computer Vision
1149 (ECCV)*, pages 570–586, 2018.
- 1150 Nitish Srivastava and Ruslan R Salakhutdinov. Multimodal learning with deep boltzmann machines.
1151 In *Advances in neural information processing systems*, pages 2222–2230, 2012.
- 1152 Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition
1153 in videos. In *Advances in neural information processing systems*, pages 568–576, 2014a.
- 1154 Amirsina Torfi, Seyed Mehdi Iranmanesh, Nasser Nasrabadi, and Jeremy Dawson. 3d convolutional
1155 neural networks for cross audio-visual matching recognition. *IEEE Access*, 5:22081–22091, 2017.
- 1156 Kihyuk Sohn, Wenling Shang, and Honglak Lee. Improved multimodal deep learning with variation
1157 of information. In *Advances in Neural Information Processing Systems*, pages 2141–2149, 2014.
- 1158 Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New
1159 York, 2006.
- 1160 Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks.
1161 In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- 1162 Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger
1163 Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for
1164 statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014a.
- 1165 Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of
1166 neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014b.
- 1167 Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly
1168 learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- 1169 Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio.
1170 Attention-based models for speech recognition. In *Advances in neural information processing
1171 systems*, pages 577–585, 2015.

- Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pages 2204–2212, 2014.
- Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 4960–4964. IEEE, 2016.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pages 2048–2057, 2015.
- Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. A dual-stage attention-based recurrent neural network for time series prediction. *arXiv preprint arXiv:1704.02971*, 2017.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- Christopher M Bishop. Mixture density networks. Technical report, Citeseer, 1994.
- Marina Meila and Michael I Jordan. Learning fine motion by markov mixtures of experts. In *Advances in neural information processing systems*, pages 1003–1009, 1996.
- Helen L Bear and Richard Harvey. Phoneme-to-viseme mappings: the good, the bad, and the ugly. *Speech Communication*, 95:40–67, 2017.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988, 2015.
- Marco Fraccaro, Søren Kaae Sønderby, Ulrich Paquet, and Ole Winther. Sequential neural models with stochastic layers. In *Advances in neural information processing systems*, pages 2199–2207, 2016.
- Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Yoshua Bengio, Aaron Courville, and Pierre Vincent. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1798–1828, 2013.
- Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

- 1218 Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for
1219 nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- 1220 Geoffrey E Hinton and Sam T Roweis. Stochastic neighbor embedding. In *Advances in neural*
1221 *information processing systems*, pages 833–840, 2002.
- 1222 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine*
1223 *Learning Research*, 9(2579-2605):85, 2008.
- 1224 Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant
1225 mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*,
1226 volume 2, pages 1735–1742. IEEE, 2006.
- 1227 Sean Bell and Kavita Bala. Learning visual similarity for product design with convolutional neural
1228 networks. *ACM Transactions on Graphics (TOG)*, 34(4):98, 2015.
- 1229 Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing
1230 Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for
1231 audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*
1232 *(ICASSP)*, pages 776–780. IEEE, 2017.
- 1233 Yichi Zhang and Zhiyao Duan. Imisound: An unsupervised system for sound query by vocal imitation.
1234 In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*,
1235 pages 2269–2273. IEEE, 2016.
- 1236 Yichi Zhang and Zhiyao Duan. Iminet: Convolutional semi-siamese networks for sound search by
1237 vocal imitation. 2017.
- 1238 Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing
1239 Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for
1240 large-scale audio classification. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE*
1241 *International Conference on*, pages 131–135. IEEE, 2017.
- 1242 Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from
1243 unlabeled video. In *Advances in neural information processing systems*, pages 892–900, 2016.
- 1244 Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol.
1245 Stacked denoising autoencoders: Learning useful representations in a deep network with a local
1246 denoising criterion. *Journal of Machine Learning Research*, 11(Dec):3371–3408, 2010.
- 1247 Bo Li, Carolina Parada, Gabor Simko, Shuo-yiin Chang, and Tara Sainath. Endpoint detection using
1248 grid long short-term memory networks for streaming speech recognition. 2017.
- 1249 Paris Smaragdis, Bhiksha Raj, and Madhusudana Shashanka. Supervised and semi-supervised
1250 separation of sounds from single-channel mixtures. *Independent Component Analysis and Signal*
1251 *Separation*, pages 414–421, 2007.
- 1252 Lijiang Guo and Minje Kim. Bitwise source separation on hashed spectra: An efficient posterior
1253 estimation scheme using partial rank order metrics. In *2018 IEEE International Conference on*
1254 *Acoustics, Speech and Signal Processing (ICASSP)*, pages 761–765. IEEE, 2018.
- 1255 Martin Cooke, Jon Barker, Stuart Cunningham, and Xu Shao. An audio-visual corpus for speech
1256 perception and automatic speech recognition. *The Journal of the Acoustical Society of America*,
1257 120(5):2421–2424, 2006.
- 1258 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolu-
1259 tional neural networks. In *Advances in neural information processing systems*, pages 1097–1105,
1260 2012.
- 1261 Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image
1262 recognition. *arXiv preprint arXiv:1409.1556*, 2014b.
- 1263 Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the
1264 inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer*
1265 *Vision and Pattern Recognition*, pages 2818–2826, 2016.

- 1266 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
1267 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
1268 pages 770–778, 2016.
- 1269 Zhiyao Duan, Gautham J Mysore, and Paris Smaragdis. Speech enhancement by online non-negative
1270 spectrogram decomposition in nonstationary noise environments. In *Thirteenth Annual Conference*
1271 *of the International Speech Communication Association*, 2012.