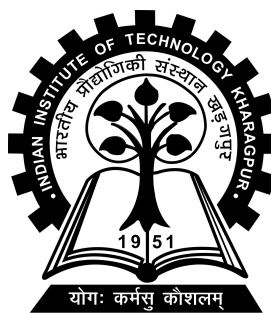


Model Blending for Text Classification

Project-II (MA57012) report submitted to
Indian Institute of Technology Kharagpur
in partial fulfilment for the award of the degree of
Master of Science
in
Mathematics and Computing

by
Ramit Pahwa
(14MA20029)

Under the supervision of
Prof. Jayanta Mukhopadhyaya, Dr Sunav Choudhary, Dr Vishwa Vinay



Department of Computer Science
Indian Institute of Technology Kharagpur

Autumn Semester, 2018-19

April 29, 2018

DECLARATION

I certify that

- (a) The work contained in this report has been done by me under the guidance of my supervisor.
- (b) The work has not been submitted to any other Institute for any degree or diploma.
- (c) I have conformed to the norms and guidelines given in the Ethical Code of Conduct of the Institute.
- (d) Whenever I have used materials (data, theoretical analysis, figures, and text) from other sources, I have given due credit to them by citing them in the text of the thesis and giving their details in the references. Further, I have taken permission from the copyright owners of the sources, whenever necessary.

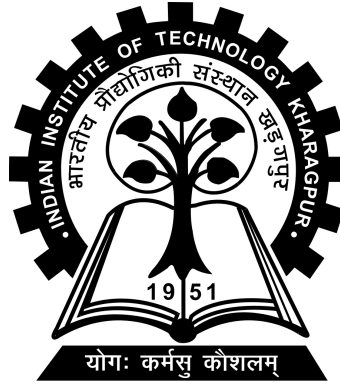
Date: April 29, 2018

Place: Kharagpur

(Ramit Pahwa)

(14MA20029)

DEPARTMENT OF COMPUTER SCIENCE
INDIAN INSTITUTE OF TECHNOLOGY KHARAGPUR
KHARAGPUR - 721302, INDIA



CERTIFICATE

This is to certify that the project report entitled “Model Blending for Text Classification” submitted by Ramit Pahwa (Roll No. 14MA20029) to Indian Institute of Technology Kharagpur towards partial fulfilment of requirements for the award of degree of Master of Science in Mathematics and Computing is a record of bona fide work carried out by him under my supervision and guidance during Autumn Semester, 2018-19.

Date: April 29, 2018
Place: Kharagpur

Prof. Jayanta Mukhopadhyaya, Dr Sunav
Choudhary, Dr Vishwa Vinay
Department of Computer Science
Indian Institute of Technology Kharagpur
Kharagpur - 721302, India

Acknowledgements

I express my deep sense of gratitude to my guide Prof. Jayanta Mukhopadhyaya for her valuable guidance and inspiration throughout the course of this work. I am thankful to her for her help, guidance, directions and support and his motivation towards independent thinking. It has been a great experience working under her in the cordial environment. I would also like to thank Prof. Dejeni Chakraborty, Department of Computer Science, IIT Kharagpur, Dr. Viswa Vinay and Dr. Sunav Choudhary, Adobe Research for their constant guidance and support.

Abstract

Name of the student: **Ramit Pahwa**

Roll No: **14MA20029**

Degree for which submitted: **Master of Science**

Department: **Department of Computer Science**

Thesis title: **Model Blending for Text Classification**

Thesis supervisor: **Prof. Jayanta Mukhopadhyaya, Dr Sunav Choudhary, Dr Vishwa Vinay**

Month and year of thesis submission: **April 29, 2018**

Deep neural networks (DNNs) have proven successful in a wide variety of applications such as speech recognition and synthesis, computer vision, machine translation, and game playing, to name but a few. However, existing deep neural network models are computationally expensive and memory intensive, hindering their deployment in devices with low memory resources or in applications with strict latency requirements. Therefore, a natural thought is to perform model compression and acceleration in deep networks without significantly decreasing the model performance, which is what we call **reducing the complexity**. In the following work, we try reducing the complexity of state of the art LSTM models for natural language tasks such as text classification, by distilling their knowledge to CNN based models, thus reducing the inference time(or latency) during testing.

Contents

Declaration	i
Certificate	ii
Acknowledgements	iii
Abstract	iv
Contents	v
List of Figures	vii
List of Tables	viii
Contents	ix
1 Introduction	1
1.1 Motivation	1
Reducing Complexity of Deep Neural Networks - Why ?	1
1.2 Arriving at the Problem	2
Related Work	2
Research gaps	3
1.3 Text Classification	4
1.4 Problem Statement	5
2 RNNs vs CNNs	6
2.1 RNNs vs CNNs	6
Advantages of using CNNs over RNNs	6
Disadvantages of using CNNs over RNNs	7
Conclusion	8
3 Methodology	9
3.1 Teacher: LSTM	10
3.2 Student: CNN	14

3.3	Blending	16
4	Experiments and Results	18
4.1	Dataset	18
4.1.1	Implementation Details	19
4.1.2	Baseline and Results	20
4.2	Main Results	21
5	Conclusion	22
6	Future Work	23
7	Appendix A	24

List of Figures

3.1	Blending Process	10
3.2	The repeating module in an LSTM contains four interacting layers. .	11
3.3	Student Network	15

List of Tables

4.1	Description of datasets	18
4.2	Comparison across datasets. one column for each dataset is accuracy	19
4.3	Comparison on the execution time of the methods is execution time compared to the our proposed CNN (15) on IMDB dataset.	19

Listings

3.1	Algorithm for Recurrent Neural Network LSTM applied to Text Data	11
3.2	Algorithm for Recurrent Neural Network AWD-LSTM applied to Text Data	12
3.3	Algorithm for Convolution Neural Network applied to Textual Data .	14

Chapter 1

Introduction

1.1 Motivation

Reducing Complexity of Deep Neural Networks - Why ? Deep neural networks (DNNs) have proven successful in a wide variety of applications such as speech recognition and synthesis, computer vision, machine translation, and game playing, to name but a few. However, existing deep neural network models are computationally expensive and memory intensive, hindering their deployment in devices with low memory resources or in applications with strict latency requirements. Therefore, a natural thought is to perform model compression and acceleration in deep networks without significantly decreasing the model performance, which is what we call **reducing the complexity**. As larger neural networks with more layers and nodes are considered, reducing their complexity becomes critical, especially for some real-time applications such as online learning and incremental learning. In addition, recent years witnessed significant progress in virtual reality, augmented

reality, and smart wearable devices, creating unprecedented opportunities for researchers to tackle fundamental challenges in deploying deep learning systems to portable devices with limited resources (e.g. memory, CPU, energy, bandwidth).

1.2 Arriving at the Problem

Related Work At the high level, there are two research directions for reducing model complexity: one is to introduce a completely new architecture with reduced complexity which mimics the original model and other is to make changes to the original model itself, without altering its overall structure. By not altering the structure, we mean applying techniques like quantization, pruning, huffman coding to the model we want to compress. Here again there are 2 directions: one is to compress while training, for example training a model with weights that are quantized, like in trained ternary quantization by (29), and other is compressing a pre-trained model, for example using quantization, pruning etc. like in deep compression by (10).

The second approach is to introduce a new architecture that mimics the original model. Here one idea is knowledge distillation due to (11), which is basically to make the student model learn the behaviour of the teacher model. The caveat here is that the architecture of the student model has to be hand designed. Due to this, there have been explorations on learning the ‘optimal’ architecture so to speak. There have been some very recent works on architecture search, which try to search the architecture from scratch via reinforcement learning or genetic algorithms, with the downside being it is computationally expensive. Another approach could be restricting the search space to the teacher model only. It relies on the premise that since the teacher model is able to achieve high accuracy on the dataset, it already contains the components required to solve the task and therefore is a suitable search

space for compressed architecture - an idea that has been used in N2N compression by (1). There also have been some very recent works that try to combine some of these approaches. For example, quantized distillation by (20) which is basically knowledge distillation but with the weights quantized while training to learn the behaviour of the teacher model.

Research gaps Much of the recent work on reducing model complexity is focused on learning new architectures from scratch or using a teacher model, via reinforcement learning, which is the second approach mentioned above. However, the reward functions that these reinforcement learning methods use are based on just accuracy and compression. None of these methods incorporate latency in their reward function, which can be important in latency strict applications. We first thought of working in this direction, but the reinforcement learning paradigm for architecture search requires a huge amount of computational resources, which we didn't have at our disposal. Thus we thought of reducing the scope of the problem, and during literature reading came across a paper by (2) that compared RNNs and CNNs for sequence modelling. It said that CNNs are advantageous over RNNs for sequence modelling in that parallel processing is possible for them, resulting in reduced latency during training and testing. Since almost all state of the art models for sequence modelling tasks like machine translation etc. are RNN based, we thought of working on distilling the knowledge from these architectures to CNN based models, a recent study proposes utilizing the outputs of LSTM while training the CNN which tries to incorporate 'dark knowledge' (11) for automatic speech recognition task (7). We propose blending for text classification task.

1.3 Text Classification

Text classification is fundamental problem in Natural Language Processing(NLP) and is an essential component in applications such as sentiment analysis (16), question classification (24) and topic classification (27).In recent year there has been a shift from traditional bag of words (BoW) model, which treats textual data as an unordered set of words (26) to more recent order sensitive models which have achieved exceptional performance across numerous NLP tasks. The challenge for textual modelling is to tackle how to capture features from various textual units such as phrases, sentences and documents. Benefitting from the sequential nature of the textual data, the recurrent neural network (RNN) in particular Long short term memory (LSTM) (28) units has been successful in modelling numerous classification tasks. But it comes with its limitations as it processes the data in a sequential manner, thereby increasing the execution latency and hindering its deployment in strict latency environments. The alternative deep learning architecture which in the recent years has gained popularity over LSTM is the convolution neural network (CNN) (15) primarily due to its parallelism, flexible receptive filter window (25) and low memory requirement while training in comparison to LSTM network. The inherent difference in the architectures of CNN and LSTM pose an interesting question of whether the functions they represent are different and if so, is there a method to distill knowledge from one of them to other. In this paper we apply **model blending**, a method to train very accurate CNN models by using LSTM priors to guide the training process, to text classification. The choice to distill the knowledge from LSTM to CNN and not vice-versa is motivated by the fact that CNN's are faster during inference time due to their parallelism. We thus arrive at the problem statement.

1.4 Problem Statement

Distilling knowledge from state of the art RNN architectures () for natural language tasks such as text classification, machine translation etc. into CNN based architectures and comparing the latency and accuracy. As well as effectively employ a Computer Vision (CV) like transfer learning mechanism to Natural Language Processing (NLP) . In summary, the contributions to achieve this are :

- We propose blending into CNN using state of the art LSTM priors resulting in accurate CNN classifiers which are 15x more computationally efficient at test time and take a fraction of the training time as compared with the training of LSTM network.
- We empirically find that using a more accurate Teacher LSTM we can improve on the performance of the model.
- We extensively evaluate the blending process across text classification tasks such sentiment analysis (16), question classification (24) and topic classification and obtain competitive performance when compared with state of the art models.
- We provide vision style transfer learning by using recurrent language model trained on the corpus(12), instead of traditional embedding based transfer (19).

Chapter 2

RNNs vs CNNs

2.1 RNNs vs CNNs

Deep learning practitioners commonly regard recurrent architectures as the default starting point for sequence modeling tasks. On the other hand, recent research indicates that certain convolutional architectures can reach state-of-the-art accuracy in audio synthesis, word-level language modeling, and machine translation ((18); (14); (4); (6)). This raises the question of whether these successes of convolutional sequence modeling are confined to specific application domains or whether a broader reconsideration of the association between sequence processing and recurrent networks is in order. The paper by (2) addresses this question by conducting a systematic empirical evaluation of convolutional and recurrent architectures on a broad range of sequence modeling tasks. We provide a summary of the advantages and disadvantages of using CNNs over RNNs for sequence modelling, as presented in the paper.

Advantages of using CNNs over RNNs

- **Parallelism:** Unlike in RNNs where the predictions for later timesteps must wait for their predecessors to complete, convolutions can be done in parallel. Therefore, in both training and evaluation, a long input sequence can be processed as a whole in CNN, instead of sequentially as in RNN.
- **Flexible receptive field size:** A CNN can change its receptive field size in multiple ways. For instance, stacking more dilated (causal) convolutional layers, using larger dilation factors, or increasing the filter size are all viable options (with possibly different interpretations). CNNs thus afford better control of the model's memory size, and are easy to adapt to different domains.
- **Stable gradients:** Unlike recurrent architectures, CNN has a backpropagation path different from the temporal direction of the sequence. CNN thus avoids the problem for RNNs.
- **Low memory requirement for training:** Especially in the case of a long input sequence, RNNs and GRUs can easily use up a lot of memory to store the partial results for their multiple cell gates. However, in a CNNs the filters are shared across a layer, with the backpropagation path depending only on network depth. Therefore in practice, gated RNNs likely to use up to a multiplicative factor more memory than CNNs.
- **Variable length inputs:** Just like RNNs, which model inputs with variable lengths in a recurrent way, CNNs can also take in inputs of arbitrary lengths by sliding the 1D convolutional kernels. This means that CNNs can be adopted as drop-in replacements for RNNs for sequential data of arbitrary length.

Disadvantages of using CNNs over RNNs

- **Data storage during evaluation:** In evaluation/testing, RNNs only need to maintain a hidden state and take in a current input x_t in order to generate

a prediction. In other words, a “summary” of the entire history is provided by the fixed-length set of vectors h_t , and the actual observed sequence can be discarded. In contrast, CNNs need to take in the raw sequence up to the effective history length, thus possibly requiring more memory during evaluation.

- **Potential parameter change for a transfer of domain:** Different domains can have different requirements on the amount of history the model needs in order to predict. Therefore, when transferring a model from a domain where only little memory is needed to a domain where much longer memory is required, CNN may perform poorly for not having a sufficiently large receptive field.

Conclusion The comparison above demonstrates that CNNs may well be the best for sequence modelling in certain situations. The advantage that we are most concerned about is the reduced inference time due to parallelism. In the next chapter, we introduce a method for distilling the knowledge of RNN architectures into CNN architectures.

Chapter 3

Methodology

Although LSTMs yield state of the art performance on numerous tasks, they are slow to evaluate at test time, thus restricting their deployment in strict latency environments. CNNs on the other hand are much faster at test time on account of their parallelism, though subpar in performance as compared to LSTMs. The proposed system aims to improve the performance of the CNNs while reaping the benefits we get from it during deployment. The methodology, which we call *model blending*, is inspired by (7). Their objective is to improve the performance of the CNN model on automatic speech recognition task by mimicking the output of the LSTM model. The process is somewhat analogous to knowledge distillation (11), in which a bigger more powerful model referred as the teacher is used to train a compact model referred as the student. The teacher here is a variant of Recurrent Neural Network i.e Long Short-Term Memory (LSTM) which can capture capture long term dependencies and the student is a Convolutional Neural Network (CNN). The choice of having the CNN model as the student is what gives us benefits in terms of latency at test time. Thus, the blending process gives the benefits of ensembling the predictions of both the models (namely LSTM and CNN) while bypassing the

usual objection of an ensemble requiring too much computation at test time. An intuitive implementation is shown in figure 1.

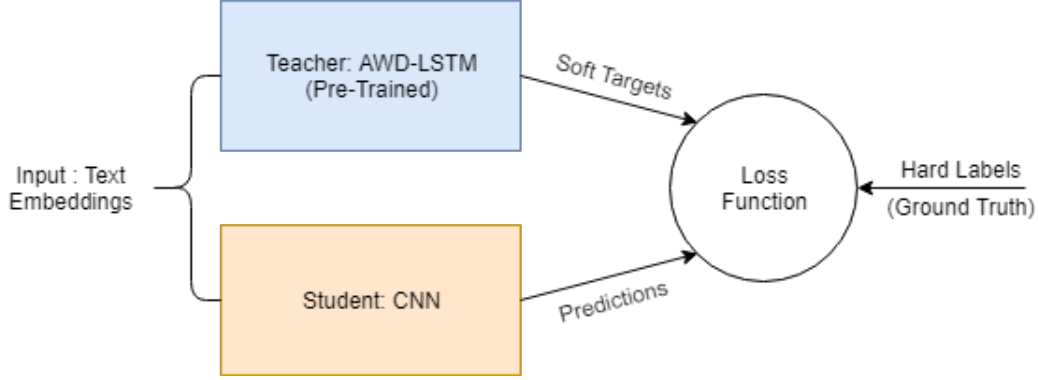


FIGURE 3.1: Blending Process

3.1 Teacher: LSTM

LSTMs exhibit superior performance not only in classification (17; 12) and speech , but also in handwriting recognition (9; 8) and parsing (23), this is primarily due their ability to capture long range interaction which is an essential component in understanding the semantics of the sentence. The Figure shows the LSTM unit. The governing equations are as follows:

$$f_t = \sigma_g(W_f x_t + U_f h_{t-1} + b_f) \quad (3.1)$$

$$i_t = \sigma_g(W_i x_t + U_i h_{t-1} + b_i) \quad (3.2)$$

$$o_t = \sigma_g(W_o x_t + U_o h_{t-1} + b_o) \quad (3.3)$$

$$c_t = f_t \circ c_{t-1} + i_t \circ \sigma_c(W_c x_t + U_c h_{t-1} + b_c) \quad (3.4)$$

$$h_t = o_t \circ \sigma_h(c_t) \quad (3.5)$$

For classification tasks, we train two networks, the first one is a Bi-Directional

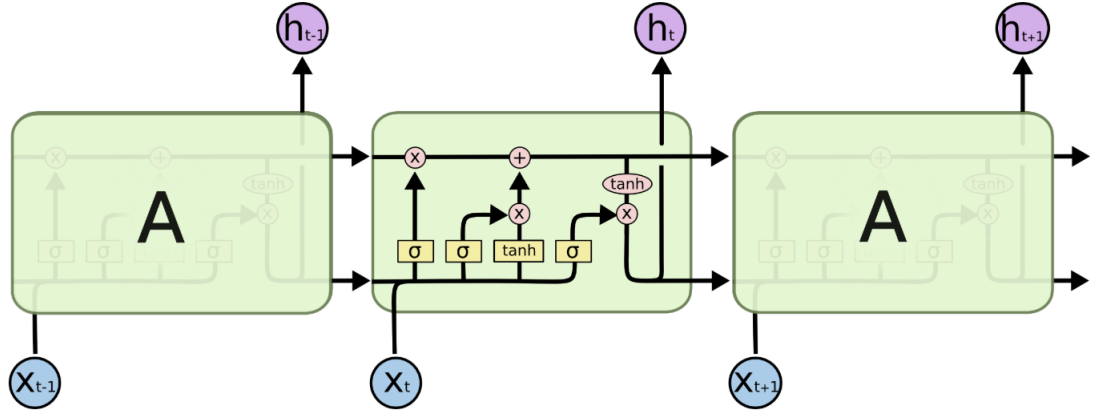


FIGURE 3.2: The repeating module in an LSTM contains four interacting layers.

LSTM network (28) which is a tough baseline to break for most NLP tasks. The sentence $S = (x_1, x_2, \dots, x_m)$, where m is the length of the sentence is taken as input and then we use glove embedding (19) of size 300 to embed the words of the sentence. The above obtained sentence matrix S has a dimension of $N \times d$, here N is the batch size and d is the fixed embedding dimension of each word. This matrix is given as input to the Bi-direction LSTM network, the last hidden state \vec{h}_t of the forward network along with the last hidden state of the backward network \overleftarrow{h}_t is passed into the fully connected network to predict the outputs. This forms a baseline RNN teacher for our experiments. We further utilize the state of the art architecture i.e AWD-LSTM (17) architecture and employ the same training regime mentioned in (12) namely discriminative fine-tuning, gradual unfreezing of layers and slanted triangular learning rates to achieve state of the art performance for these tasks. The Pseudo code is provided below to define the RNN network in pytorch.

```
import torch.nn as nn

class RNN(nn.Module):
    def __init__(self, v_s, emb_dim, h_dim, o_dim, n, bi, d):
        super().__init__()

        self.embedding = nn.Embedding(v_s, emb_dim)
        self.rnn = nn.LSTM(emb_dim, h_dim, n, bi, dropout=d)
```

```

self.fc = nn.Linear(h_dim*2, o_dim)
self.dropout = nn.Dropout(d)

def forward(self, x):
    #x = [sent len, batch size]
    embedded = self.dropout(self.embedding(x))
    #embedded = [sent len, batch size, emb dim]

    output, (hidden, cell) = self.rnn(embedded)

    #output = [sent len, batch size, hid dim * num directions]
    #hidden = [num layers * num directions, batch size, hid dim]
    #cell = [num layers * num directions, batch size, hid dim]

    #concat the final forward (hidden[-2,:,:])
    # and backward (hidden[-1,:,:]) hidden layers
    #and apply dropout

    hidden = self.dropout(torch.cat((hidden[-2,:,:], hidden[-1,:,:]), dim=1))

    #hidden = [batch size, hid dim * num directions]

    return self.fc(hidden.squeeze(0))

```

LISTING 3.1: Algorithm for Recurrent Neural Network LSTM applied to Text Data

```

from torch import *
class RNN_Encoder(nn.Module):

    """A custom RNN encoder network
    """

    initrange=0.1
    def __init__(self, ntoken, emb_sz, nhid, nlayers, pad_token, bidir=False,
                  dropout=0.3, dropouti=0.65, dropoute=0.1, wdrop=0.5):
        """ Default constructor for the RNN_Encoder class
            Returns:
                None
        """

        super().__init__()
        self.ndir = 2 if bidir else 1

```

```

self.bs = 1
self.encoder = nn.Embedding(ntoken, emb_sz, padding_idx=pad_token)
self.encoder_with_dropout = EmbeddingDropout(self.encoder)
self.rnn = [nn.LSTM(emb_sz if l == 0 else nhid, self.nhid, 1, 1) for l in range(nlayers)]
if wdrop: self.rnn = [WeightDrop(rnn, wdrop) for rnn in self.rnn]
self.rnn = torch.nn.ModuleList(self.rnn)
self.encoder.weight.data.uniform_(-self.initrage, self.initrage)

self.emb_sz, self.nhid, self.nlayers, self.dropoute = emb_sz, nhid, nlayers, dropoute
self.dropouti = LockedDropout(dropoute)
self.dropouths = nn.ModuleList([LockedDropout(dropouth) for l in range(nlayers)])

def forward(self, input):
    """ Invoked during the forward propagation of the RNN_Encoder module.
    Args:
        input (Tensor): input of shape (sentence length x batch_size)

    Returns:
        raw_outputs (tuple(list (Tensor), list(Tensor))): (no dropout, dropout)
    """
    sl, bs = input.size()
    if bs != self.bs:
        self.bs = bs
        self.reset()

    emb = self.encoder_with_dropout(input, dropout=self.dropoute if self.training else 0)
    emb = self.dropouti(emb)

    raw_output = emb
    new_hidden, raw_outputs, outputs = [], [], []
    for l, (rnn, drop) in enumerate(zip(self.rnn, self.dropouths)):
        current_input = raw_output
        with warnings.catch_warnings():
            warnings.simplefilter("ignore")
            raw_output, new_h = rnn(raw_output, self.hidden[l])
        new_hidden.append(new_h)
        raw_outputs.append(raw_output)
        if l != self.nlayers - 1: raw_output = drop(raw_output)
        outputs.append(raw_output)

    self.hidden = repackage_var(new_hidden)
    return raw_outputs, outputs

```

LISTING 3.2: Algorithm for Recurrent Neural Network AWD-LSTM applied to
Text Data

3.2 Student: CNN

Convolution Neural Networks have gained popularity as they perform at par with state of the art LSTM's networks and have the added advantage of being faster at test time. The CNN network for text classification differs from the convolutional networks used for computer vision (22; 21). The architecture for classification uses only two or three convolutional layers with large filters followed by more fully connected layers and only use convolution or pooling over one dimension, either time or frequency (15). For the purpose of our experiment we operate over time. The embedding matrix is taken as input, keeping the embedding dimensions and the initialization same as that used by the teacher network. Inspired by this we construct our CNN model (referred to as a-CNN) which will act as student learner. We also use a deep CNN network (3) (referred as b-CNN), which works at the character level and performs small convolution and pooling operations as state of the art baseline of CNN architecture. The architecture of a-CNN is shown in figure . The Pseudo code of the class is provided below.

```
import torch.nn as nn
import torch.nn.functional as F

class CNN(nn.Module):
    def __init__(self, v, emb_dim, n, f_sizes, o_dim, d):
        super().__init__()
        self.embedding_dim = emb_dim
        self.embedding = nn.Embedding(v, self.embedding_dim)
        self.convs = nn.ModuleList([nn.Conv2d(in_channels=1, out_channels=n,
        kernel_size=(fs, embedding_dim)) for fs in filter_sizes])
```

```

self.fc = nn.Linear(len(filter_sizes)*n_filters, output_dim)
self.dropout = nn.Dropout(dropout)

def forward(self, x):
    #x = [sent len, batch size]
    x = x.permute(1, 0)
    #x = [batch size, sent len]

    embedded = self.embedding(x)
    #embedded = [batch size, sent len, emb dim]

    embedded = embedded.unsqueeze(1)
    #embedded = [batch size, 1, sent len, emb dim]

    convded = [F.relu(conv(embedded)).squeeze(3) for conv in self.convs]
    #conv_n = [batch size, n_filters, sent len - filter_sizes[n]]

    pooled = [F.max_pool1d(conv, conv.shape[2]).squeeze(2) for conv in convded]
    #pooled_n = [batch size, n_filters]

    cat = self.dropout(torch.cat(pooled, dim=1))
    #cat = [batch size, n_filters * len(filter_sizes)]

    return self.fc(cat)

```

LISTING 3.3: Algorithm for Convolution Neural Network applied to Textual Data

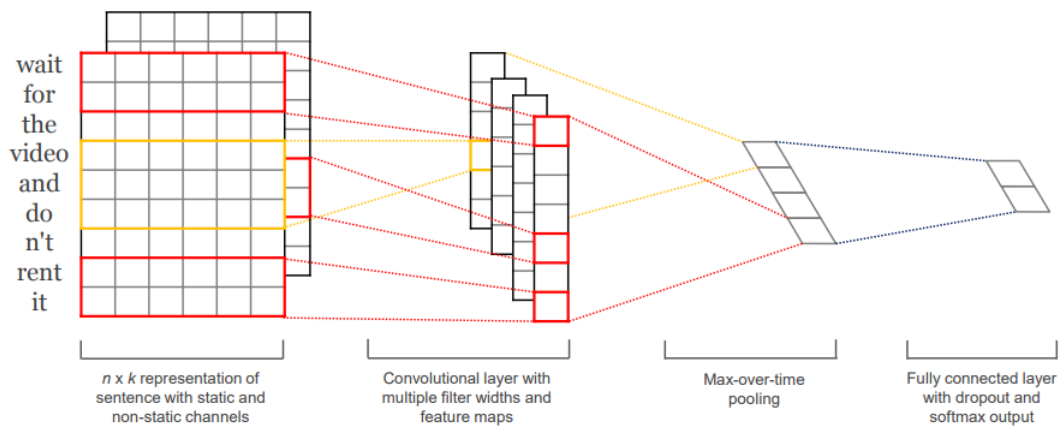


FIGURE 3.3: Student Network

3.3 Blending

Both RNNs and CNNs are powerful models, but the mechanisms that guide their learning are quite different. That creates an opportunity to combine their predictions, implicitly averaging their inductive biases. A classic way to perform this is ensembling, that is, to mix posterior predictions of the two models in the following manner :

$$p_{ensemble}(y|s) = \gamma p_{LSTM}(y|s) + (1 - \gamma) p_{CNN}(y|s) \quad (3.6)$$

where $\gamma \in [0, 1]$ and $s_i = (x_1, x_2, \dots, x_j, \dots, x_m)$. Also, $p_{LSTM}(y|s)$ and $p_{CNN}(y|s)$ are the probability of y output class given a feature vector $x_j \in^d$ or sentence in this case. This combination is in accordance with conditions proposed by (5) i.e. *“a necessary and sufficient condition for an ensemble of classifiers to be more accurate than any of its individual members is if the classifiers are accurate and diverse”* to form an ensemble of classifiers, but is inefficient at test time as it requires the posterior prediction of both the models at test time to compute the final predictions. The method we propose is inspired by the process of knowledge distillation (11) which aims to distill knowledge from complex teacher models to much simpler student models by matching output logits, as well as model blending (7) which uses the posterior predictions of LSTM to train vision style CNN architectures for automatic speech recognition task. We use a different class of CNN which utilizes large filter sizes and are shallow in comparison to the vision style CNN. We train our CNN architecture for text classification task namely sentiment analysis, question classification and topic classification. To the best of our knowledge, this is first time CNN architectures are trained in the presence of LSTM priors for these classification tasks. To incorporate the LSTM information during training of the CNN we modify the loss function to combine loss from hard labels from the training data with a loss function which penalises deviation from predictions of the LSTM teacher. The

training objective is to minimize the weighted sum of the cross-entropy loss :

$$\mathcal{L} = \lambda \cdot \mathcal{L}_{soft} + 1 - \lambda \cdot \mathcal{L}_{hard} \quad (3.7)$$

$$\begin{aligned} \mathcal{L}_{soft} &= H(p_{LSTM}(c|s_i), q_{CNN}(c|s_i)) \\ &= - \sum_i p_{LSTM}(c|s_i) \cdot \log q_{CNN}(c|s_i) \end{aligned} \quad (3.8)$$

$$\begin{aligned} \mathcal{L}_{hard} &= H(q_{CNN}(y_i|s_i), y_i^{true}) \\ &= - \sum_i q_{CNN}(y_i|s_i) \cdot \log y_i^{true} \end{aligned} \quad (3.9)$$

where $p_{LSTM}(c|s_i)$ is the probability of class c for the the given training example s_i predicted by the teacher LSTM network, $q_{CNN}(c|s_i)$ is the probability of class c for the the given training example s_i predicted by the student CNN and y_i^{true} is the true class for the training example s_i . Here $\lambda \in [0, 1]$, helps us control the relative contribution of the two losses. This modified loss function helps us to incorporate information learnt by LSTM network while training the student CNN network, this method give similar performance to the ensemble of the two models, but it is computationally 17x less expensive at test time than the ensemble. Thus blending results in more accurate and efficient classifiers at test time.

Chapter 4

Experiments and Results

4.1 Dataset

We evaluate our models on four public dataset namely IMDB dataset (16) for sentiment classification, TREC-6 (24) for question classification, AG’s corpus of news articles for topic modeling and DBpedia by (27). The description of the datasets are tabulated below:

TABLE 4.1: Description of datasets

Dataset	IMDB	TREC-6	AG	DBPedia
Training	25k	5452	120k	560k
Testing	25k	500	7600	70k
Classes	2	6	4	14
Avg-words	234	-	45	55

TABLE 4.2: Comparision across datasets. one column for each dataset is accuracy

Model		IMDB	TREC-6	AG	DBPedia
Linear	Fasttext (13)	85.22	88.2	92.6	98.5
RNN	LSTM	82.77	90.2	82.77	-
	Bi-LSTM (28)	84.11	93.0	93.0	99.12
	AWD-LSTM (17)	95.4	96.4	95.0	99.20
CNN	a-CNN (15)	88.22	92.0	91.6	98.6
	b-CNN (3)	-	93.0	91.3	98.4
	c-CNN (27)	-	93.2	87.2 (90.49)	98.3
Ensemble	CNN + LSTM	88.63	93.15	92.0	99.0
	CNN + AWD-LSTM	95.4	96.7	95.0	99.23
Blended*	CNN + LSTM	87.23	91.08	92.3	99.14
	CNN + AWD-LSTM	93.2	92.34	92.5	99.1

TABLE 4.3: Comparison on the execution time of the methods is execution time compared to the our proposed CNN (15) on IMDB dataset.

Model		Execution Time
Linear	Fasttext	1.1x
RNN	LSTM	5.55x
	Bi-LSTM	5.64x
	AWD-LSTM	15.5x
CNN	a-CNN	1.0x
	b-CNN	4.3x
Ensemble	a-CNN + LSTM	6.2x
	a-CNN + AWD-LSTM	16.7x
Blended*	a-CNN + LSTM	1.0x
	a-CNN + AWD-LSTM	1.0x

4.1.1 Implementation Details

Input. We adopted the word vectors from (19).The word embedding used were of size 300 for our LSTM network.

Training Setting. For training the state of the art AWD-LSTM (17) Teacher we follow similar training regime as (12) with an embedding size of 400, 4 layers,1150 hidden activations per layer, and a BPTT batch size of 70. We apply different

dropout to different types as suggested by and (17) The student CNN network (15; 27) has filter sizes of 3,4 and 5. The model were implemented in pytorch and were adapted from prior works (12; 15). The CNN trained using the proposed method outperforms the CNN architecture and the teacher network and is 15x more computationally efficient from its teacher. The experiments were first performed on IMDB dataset for sentiment analysis to tune γ and λ hyper-parameter to their optimal values of 0.4 and 0.5 we then validate these parameter using the TREC-6 data set for question classification. We use Adam with default setting $\beta_1 = 0.9$ and $\beta_2 = 0.99$ with a minibatch of 64. The practises otherwise are same as used by (12; 17)

4.1.2 Baseline and Results

We compare our method and utilize several of the popular models in our experiments : linear model fasttext (13), we use CNN architecture inspired by (15) and train it using numerous teacher (17; 28) which are state of the art for text classification. We compare the model with character level CNN (27) referred to as c-CNN and very Deep CNN network (3). We also benchmark it against (25) which proposes dense connection between convolutional layer and multi-scale feature attention for text classification. They also experiment with depth of the network an found that shallow networks perform comparable to the deep models with only a slight drop in performance. Shallowness of the model is advantageous for deployment to low - resource devices. Our proposed training regime improves the performance of shallow CNN and performs comparable to other popular networks of the type.

4.2 Main Results

The CNN trained through the proposed method is comparable to its Teacher LSTM network, due to the shallowness of the models it is more interpretable and is faster at the test time and hence the model can be deployed to mobile devices. The provides evidence towards using a trained complex network transferring the 'dark knowledge' (11) to a much simpler model which is computationally efficient at inference time.

Chapter 5

Conclusion

We see that blending the state of the art RNN into the CNN model, results in minimal drop in accuracy . But the inference time is reduced greatly, by about 18 times compared to the ensemble. This shows that exploring blending of state of the art RNN models into CNN models is promising, and much more can be achieved if we sacrifice a bit on the inference time (by making the CNN model more deep, for example). **This also increase deployability of these models to mobile devices i.e. Smartphones, Cameras etc.**

Chapter 6

Future Work

The current CNN architecture we used had only a single convolutional layer. We can try blending into CNN architectures with more than one layer. This would increase the inference time a bit, but the accuracy would also improve. Basically, we try experimenting with the CNN architecture until we reach a sweet spot with respect to both inference time and accuracy. Moreover, our current work deals with only sentence classification. We can try to apply this technique to other natural language tasks which require sequence modelling such as machine translation, in which the state of the art models are RNN based.

Chapter 7

Appendix A

This contains basic implementation of the above mentioned algorithm in python it has the following dependencies which are mentioned below:

- *Python* ≥ 3.0
- *fastai* $== 0.7.0$
- *torchtext* $== 0.2.3$

The experiments were performed on **Google COLAB**, which is free jupyter notebook service which has GPU capabilities. Experiments were performed on **NVIDIA K80 GPU**, and testing of the models were done on **MI A2 mobile device** by converting the **Pytorch model** to **Tensorflow model** using **ONYX**.

Bibliography

- [1] Ashok, A., Rhinehart, N., Beainy, F., Kitani, K.M.: N2n learning: network to network compression via policy gradient reinforcement learning. arXiv preprint arXiv:1709.06030 (2017)
- [2] Bai, S., Kolter, J.Z., Koltun, V.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271 (2018)
- [3] Conneau, A., Schwenk, H., Barrault, L., Lecun, Y.: Very deep convolutional networks for text classification. arXiv preprint arXiv:1606.01781 (2016)
- [4] Dauphin, Y.N., Fan, A., Auli, M., Grangier, D.: Language modeling with gated convolutional networks. arXiv preprint arXiv:1612.08083 (2016)
- [5] Dietterich, T.G.: Ensemble methods in machine learning. In: International workshop on multiple classifier systems. pp. 1–15. Springer (2000)
- [6] Gehring, J., Auli, M., Grangier, D., Dauphin, Y.N.: A convolutional encoder model for neural machine translation. arXiv preprint arXiv:1611.02344 (2016)
- [7] Geras, K.J., Mohamed, A.r., Caruana, R., Urban, G., Wang, S., Aslan, O., Philipose, M., Richardson, M., Sutton, C.: Blending lstms into cnns. arXiv preprint arXiv:1511.06433 (2015)

-
- [8] Graves, A.: Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850 (2013)
 - [9] Graves, A., Schmidhuber, J.: Offline handwriting recognition with multidimensional recurrent neural networks. In: Advances in neural information processing systems. pp. 545–552 (2009)
 - [10] Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149 (2015)
 - [11] Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
 - [12] Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). vol. 1, pp. 328–339 (2018)
 - [13] Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759 (2016)
 - [14] Kalchbrenner, N., Espeholt, L., Simonyan, K., Oord, A.v.d., Graves, A., Kavukcuoglu, K.: Neural machine translation in linear time. arXiv preprint arXiv:1610.10099 (2016)
 - [15] Kim, Y.: Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882 (2014)
 - [16] Maas, A.L., Daly, R.E., Pham, P.T., Huang, D., Ng, A.Y., Potts, C.: Learning word vectors for sentiment analysis. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies. pp. 142–150. Association for Computational Linguistics, Portland, Oregon, USA (June 2011), <http://www.aclweb.org/anthology/P11-1015>

-
- [17] Merity, S., Keskar, N.S., Socher, R.: Regularizing and optimizing lstm language models. arXiv preprint arXiv:1708.02182 (2017)
 - [18] Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499 (2016)
 - [19] Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Empirical Methods in Natural Language Processing (EMNLP). pp. 1532–1543 (2014), <http://www.aclweb.org/anthology/D14-1162>
 - [20] Polino, A., Pascanu, R., Alistarh, D.: Model compression via distillation and quantization. arXiv preprint arXiv:1802.05668 (2018)
 - [21] Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
 - [22] Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: AAAI. vol. 4, p. 12 (2017)
 - [23] Vinyals, O., Kaiser, L., Koo, T., Petrov, S., Sutskever, I., Hinton, G.: Grammar as a foreign language. In: Advances in Neural Information Processing Systems. pp. 2773–2781 (2015)
 - [24] Voorhees, E.M., Tice, D.M.: The trec-8 question answering track evaluation. In: TREC. vol. 1999, p. 82 (1999)
 - [25] Wang, S., Huang, M., Deng, Z.: Densely connected cnn with multi-scale feature attention for text classification (2018)
 - [26] Wang, S., Manning, C.D.: Baselines and bigrams: Simple, good sentiment and topic classification. In: Proceedings of the 50th Annual Meeting of the

- Association for Computational Linguistics: Short Papers-Volume 2. pp. 90–94.
Association for Computational Linguistics (2012)
- [27] Zhang, X., Zhao, J., LeCun, Y.: Character-level convolutional networks for text classification. In: Advances in neural information processing systems. pp. 649–657 (2015)
- [28] Zhou, P., Qi, Z., Zheng, S., Xu, J., Bao, H., Xu, B.: Text classification improved by integrating bidirectional lstm with two-dimensional max pooling. arXiv preprint arXiv:1611.06639 (2016)
- [29] Zhu, C., Han, S., Mao, H., Dally, W.J.: Trained ternary quantization. arXiv preprint arXiv:1612.01064 (2016)