

GIRNet: Interleaved Multi-Task Recurrent State Sequence Models

Divam Gupta,* Tanmoy Chakraborty,* Soumen Chakrabarti[†]

*IIT Delhi, India, [†]IIT Bombay, India

{divam14038, tanmoy}@iiitd.ac.in, soumen@cse.iitb.ac.in

Abstract

In several natural language tasks, labeled sequences are available in separate domains (say, languages), but the goal is to label sequences with mixed domain (such as code-switched text). Or, we may have available models for labeling whole passages (say, with sentiments), which we would like to exploit toward better position-specific label inference (say, target-dependent sentiment annotation). A key characteristic shared across such tasks is that different positions in a primary instance can benefit from different ‘experts’ trained from auxiliary data, but labeled primary instances are scarce, and labeling the best expert for each position entails unacceptable cognitive burden. We propose GIRNet, a unified position-sensitive multi-task recurrent neural network (RNN) architecture for such applications. Auxiliary and primary tasks need not share training instances. Auxiliary RNNs are trained over auxiliary instances. A primary instance is also submitted to each auxiliary RNN, but their state sequences are gated and merged into a novel composite state sequence tailored to the primary inference task. Our approach is in sharp contrast to recent multi-task networks like the cross-stitch and sluice networks, which do not control state transfer at such fine granularity. We demonstrate the superiority of GIRNet using three applications: sentiment classification of code-switched passages, part-of-speech tagging of code-switched text, and target position-sensitive annotation of sentiment in monolingual passages. In all cases, we establish new state-of-the-art performance beyond recent competitive baselines.

1 Introduction

Neural networks have shown outstanding results in many Natural Language Processing (NLP) tasks, particularly involving sequence labeling (Schmid 1994; Lample et al. 2016) and sequence-to-sequence translation (Bahdanau, Cho, and Bengio 2014). The success is generally attributed to their ability to learn good representations and recurrent models in an end-to-end manner. Most deep models require generous volumes of training data to adequately train their large number of parameters. Collecting sufficient labeled data for some tasks entails unacceptably high cognitive burden. To overcome this bottleneck, some form of transfer

learning, semi-supervised learning or multi-task learning (MTL) is used.

In the most common form of MTL, related tasks such as part-of-speech (POS) tagging and named entity recognition (NER) share representation close to the input, e.g., as character n-gram or word embeddings, followed by separate networks tailored to each prediction task (Søgaard and Goldberg 2016; Maurer, Pontil, and Romera-Paredes 2016). Care is needed to jointly train shared model parameters to prevent one task from hijacking the representation (Ruder 2017).

Sequence labeling brings a new dimension to MTL. In many NLP tasks, labeled sequences are readily available in separate single languages, but our goal may be to label sequences from code-switched multilingual text. Or, we may have trained models for labeling whole sentences or passages with overall sentiment, but the task at hand is to infer the sentiment expressed toward a specific entity mentioned in the text. Ideally, we would like to build a composite state sequence representation where each position draws state information from the best ‘expert’ auxiliary sequence, but instances are diverse in terms of where the experts switch, and annotating these transitions would entail prohibitive cognitive cost. We survey related work in Section 2, and explain why most of them do not satisfy our requirements.

By abstracting the above requirements into a unified framework, in Section 3 we present GIRNet (**G**ated **I**nterleaved **R**ecurrent **N**etwork)¹, a novel MTL network tailored for dynamically daisy-chaining at a *word level* the best experts along the input sequence to derive a composite state sequence that is ideally suited for the primary task. The whole network over all tasks is jointly trained in an end-to-end manner. GIRNet applies both when we have common instances for the tasks and when we have only disjoint instances over the tasks. Motivated by low-resource NLP challenges, we assume that one task is *primary* but with scant labeled data, whereas the *auxiliary* tasks have more labeled data available. We train standard LSTMs on the auxiliary models; but when a primary instance is run through the auxiliary networks, their states are interleaved using a novel, dynamic gating mechanism. Remarkably, without access to any segmentation of primary instances, GIRNet learns gate

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The code is available at https://github.com/divamgupta/mtl_girnet

values such that auxiliary states are daisy-chained to maximize primary task performance.

We can think of GIRNet as emulating an LSTM whose cell unit changes dynamically with input tokens. E.g., for each token in a code-switched sentence over two languages, GIRNet learns to choose the cell unit from an LSTM trained on text in the language in which that token is written. GIRNet achieves this with an end-to-end differentiable network that does not need supervision about which language is used for each token.

In Section 4, we instantiate the GIRNet template to three concrete applications: sentiment labeling of code-switched passages, POS tagging of code-switched passages, and target-dependent (monolingual) sentiment classification. Where the primary task is to tag code-switched multilingual sequences, the auxiliary tasks would be to tag the component monolingual texts. For target-dependent sentiment classification, the auxiliary task would be target-independent whole-passage sentiment classification. In all three applications, we consistently beat competitive baselines to establish new state-of-the-art performance. These experiments are described in Section 5.

Summarizing, our contributions are four-fold:

- GIRNet, a novel position- and context-sensitive dynamic gated recurrent network to compose state sequences from auxiliary recurrent units.
- Three applications that are readily expressed as concrete instantiations of the GIRNet framework.
- Superior performance in all three applications.
- Thorough diagnostic interpretation of GIRNet’s behavior, demonstrating the intended gating effect.

2 Related work

The most common neural MTL architecture shares parameters in initial layers (near the inputs) and trains separate task-specific layers for per-task prediction (Caruana 1993; Søgaard and Goldberg 2016; Maurer, Pontil, and Romera-Paredes 2016). In applications where tasks are not closely related, finding a common useful representation for all tasks is hard. Moreover, jointly training shared model parameters, while preventing one task from hijacking the representation, may be challenging (Ruder 2017). In soft parameter sharing (Duong et al. 2015), each task has its separate set of parameters and the distance between the inter-task parameters is minimized by adding an additional loss while training.

In some cases, rather than just sharing the parameters (completely or partially), state sequence features extracted by the model for one task are fed into the model of another task. In shared-private MTL models, there is one common shared model over all tasks and also separate private models for each task. The following approaches combine information from shared and private LSTMs at the granularity of token positions. Liu, Qiu, and Huang (2017) and Chen et al. (2018b) use one shared LSTM and one private LSTM per task. Liu, Qiu, and Huang (2017) concatenate their outputs, whereas Chen et al. (2018b) concatenate the

shared LSTM state with the input embeddings. In the low-supervision MTL model (Søgaard and Goldberg 2016), auxiliary tasks are trained on the lower layers and the primary task is trained on the higher layer. None of these models control the amount of information shared between different tasks. To overcome this problem, various network architectures have been evolved to control more carefully the transfer across different tasks. Cross-stitch (Misra et al. 2016) and sluice (Ruder et al. 2017b) networks are two such frameworks. Chen et al. (2018a) have used reinforcement learning to search for the best patterns of sharing between tasks. However, transfer happens at the granularity of layers, and not recurrent positions. Also, transfer happens usually via vector concatenation. No gating information crosses RNN boundaries. Meta-MTL (Chen et al. 2018b) uses an LSTM shared across all tasks to control the parameters of the task-specific LSTMs. As will become clear, we gain more representational power by gating and interleaving of auxiliary states driven by the input. Further details of some of the above approaches are discussed along with experiments in Section 5.

3 Formulation and proposed architecture

Our abstract problem setting consists of a **primary** task and $m \geq 1$ **auxiliary** tasks. E.g., the primary task may be part-of-speech (POS) tagging of code-mixed (say, English and Spanish) text, and the two auxiliary tasks may be POS tagging of pure English and Spanish text. All tasks involve labeling sequences, but the labels may be per-token (e.g., POS tagging) or global (e.g., sentiment analysis). Labeled data for the primary task is generally scarce compared to the auxiliary tasks. Different tasks will generally have disjoint sets of instances. Our goal is to mitigate the paucity of primary labeled data by transferring knowledge from auxiliary labeled data. Our proposed architecture is particularly suitable when different parts or spans of a primary instance are related to different auxiliary tasks, and these primary spans are too expensive or impossible to identify during training.

In this section, we introduce GIRNet (Gated Interleaved Recurrent Network), our deep multi-task architecture which learns to skip or select spans in auxiliary RNNs based on its ability to assist the primary labeling task. We describe the system by using LSTMs, but similar systems can be built using other RNNs such as GRUs. Before going into details, we give a broad outline of our strategy. Auxiliary labeled instances are input to auxiliary LSTMs to reduce auxiliary losses (Section 3.2). Each primary instance is input to a ‘gating’ LSTM (Section 3.3) and a variation of each auxiliary LSTM, which run concurrently with a composite sequence assembled from the auxiliary state sequences using the scalar gate values (Section 3.4). In effect, our network learns to dynamically daisy-chain the best auxiliary experts word by word. Finally, the composite state sequence is combined with a primary LSTM state sequence to produce the primary prediction, with its corresponding primary loss (Section 3.5).

Auxiliary and primary losses are jointly optimized. Weights of each auxiliary LSTM are updated by its own loss function and the loss of the primary task, but not any

other auxiliary task. Remarkably, the scalar gating LSTM can learn *without supervision* how to assemble the composite state sequence. Qualitative inspection of the gate outputs show that the selection of the auxiliary states indeed gets tailored to the primary task. A high-level sketch of our architecture is shown in Figure 1 with two auxiliary tasks and one primary task.

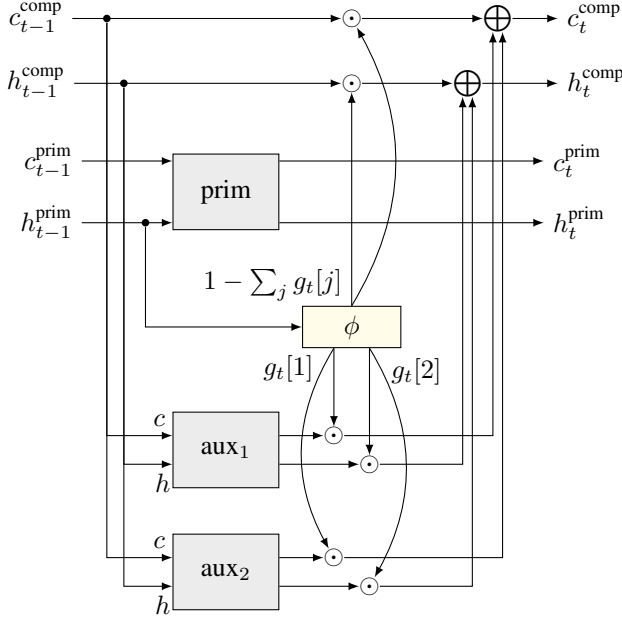


Figure 1: Part of GIRNet that processes primary input x_t^{prim} , which is provided to primary LSTM ‘prim’ and auxiliary LSTMs ‘aux $_j$ ’ for $j = 1, 2$ in this example, as well as gating logic ϕ . (x_t^{prim} has been elided to reduce clutter.) Training of auxiliary LSTMs on auxiliary inputs is standard, and has been omitted for clarity. h_t^{comp} is the gated composite state sequence. \oplus represents elementwise addition and \odot represents elementwise multiplication of the vector input with the scalar gate input.

3.1 Input embedding layer

Each input instance is a sentence- or passage-like sequence of words or tokens that are mapped to integer IDs. An embedding layer maps each ID to a d -dimensional vector. The resulting sequence of embedding vectors for a sentence will be called x . We use a common embedding matrix over all auxiliary and primary tasks. The embedding matrix is initialized randomly or using pre-trained embeddings (Pennington, Socher, and Manning 2014; Mikolov et al. 2013), and then trained along with the rest of our network. Labeled instances will be accompanied by a suitable label y .

3.2 Auxiliary LSTMs run on auxiliary input

For each auxiliary task $j \in [1, m]$, we train a separate auxiliary LSTM using a separate set of labeled instances, which we call $x^{\text{auxaux}_j} = (x_t^{\text{auxaux}_j} : t = 1, \dots, T)$, where t indexes positions, and the ground-truth is y^{auxaux_j} . The signature and

dimension of y^{auxaux_j} can vary with the nature of the auxiliary task (classification, sequence labeling, etc.)

For auxiliary task j , we define the LSTM state at each position t : input gate $i_t^{\text{auxaux}_j}$, output gate $o_t^{\text{auxaux}_j}$, forget gate $f_t^{\text{auxaux}_j}$, memory state $c_t^{\text{auxaux}_j}$ and hidden state $h_t^{\text{auxaux}_j}$, all vectors in \mathbb{R}^d . For all auxiliary task models, we use the same number of hidden units.

$$\begin{bmatrix} \tilde{c}_t^{\text{auxaux}_j} \\ o_t^{\text{auxaux}_j} \\ i_t^{\text{auxaux}_j} \\ f_t^{\text{auxaux}_j} \end{bmatrix} = \begin{bmatrix} \tanh \\ \sigma \\ \sigma \\ \sigma \end{bmatrix} \left(W^{\text{aux}_j} \begin{bmatrix} x_t^{\text{auxaux}_j} \\ h_{t-1}^{\text{auxaux}_j} \end{bmatrix} \right) \quad (1)$$

$$c_t^{\text{auxaux}_j} = \tilde{c}_t^{\text{auxaux}_j} \odot i_t^{\text{auxaux}_j} + c_{t-1}^{\text{auxaux}_j} \odot f_t^{\text{auxaux}_j} \quad (2)$$

$$h_t^{\text{auxaux}_j} = o_t^{\text{auxaux}_j} \odot \tanh c_t^{\text{auxaux}_j} \quad (3)$$

(For compact notation we have written the operators to be applied in a column vector.) Using the hidden states of the LSTM for the auxiliary task, we get the desired output using another model M^{auxaux_j} . E.g., we can use a fully connected layer over the last hidden states or pooled hidden states for the whole-sequence classification. We can use a fully connected layer on each hidden state for sequence labeling. Using generic notation,

$$\text{out}^{\text{auxaux}_j} = M^{\text{auxaux}_j}(h_1^{\text{auxaux}_j}, h_2^{\text{auxaux}_j}, \dots, h_n^{\text{auxaux}_j}) \quad (4)$$

For each auxiliary task j , a separate loss

$$\text{loss}^{\text{auxaux}_j}(\text{out}^{\text{auxaux}_j}, y^{\text{auxaux}_j}) \quad (5)$$

is computed using the model output and the ground-truth. The losses of all auxiliary tasks are added to the final loss.

3.3 LSTM on primary instance to produce gating signal

A primary input instance is written as $x^{\text{prim}} = (x_1^{\text{prim}}, \dots, x_n^{\text{prim}})$ and the ground-truth label as y^{prim} . At each token position t , the gating LSTM has internal representations as follows: input gate i_t^{prim} , output gate o_t^{prim} , forget gate f_t^{prim} , memory state c_t^{prim} and hidden state h_t^{prim} . All these vectors are in \mathbb{R}^d , where d is the number of hidden units in the gating LSTM.

$$\begin{bmatrix} \tilde{c}_t^{\text{prim}} \\ o_t^{\text{prim}} \\ i_t^{\text{prim}} \\ f_t^{\text{prim}} \end{bmatrix} = \begin{bmatrix} \tanh \\ \sigma \\ \sigma \\ \sigma \end{bmatrix} \left(W^{\text{prim}} \begin{bmatrix} x_t^{\text{prim}} \\ h_{t-1}^{\text{prim}} \end{bmatrix} \right) \quad (6)$$

$$c_t^{\text{prim}} = \tilde{c}_t^{\text{prim}} \odot i_t^{\text{prim}} + c_{t-1}^{\text{prim}} \odot f_t^{\text{prim}} \quad (7)$$

$$h_t^{\text{prim}} = o_t^{\text{prim}} \odot \tanh c_t^{\text{prim}} \quad (8)$$

Here we describe the RNN that produces the gating signal as a uni-directional LSTM, but, depending on the application, we could use bi-directional LSTMs. Using the hidden state of the gating LSTM and its input at token position t , we compute gate vector $g_t \in \mathbb{R}^m$, where m is the number of auxiliary tasks. I.e., for each auxiliary task, we predict a scalar gate value.

$$g_t = \phi \left(W^{\text{gator}} \begin{bmatrix} x_t^{\text{prim}} \\ h_{t-1}^{\text{prim}} \end{bmatrix} \right) \quad (9)$$

Here ϕ is an activation function which should ensure $\sum_{j=1}^m g_t[j] \leq 1$. The rationale for this stipulation will become clear shortly from Equations (10) and (11). We implement ϕ by including a fully-connected layer that generates $m+1$ scalar values, followed by applying soft-max, and discarding the last value.

3.4 Auxiliary and gated composite LSTMs run on primary input

For prediction on a primary instance to benefit from auxiliary models, the primary instance is also processed by a variant of each auxiliary LSTM. The model weights W^{aux_j} of each auxiliary LSTM will be borrowed from the corresponding auxiliary task, but the input is x_t^{prim} and the states will be composite states over all auxiliary LSTMs. Therefore the internal cell variables will have different values, which we therefore give different names: input gate $i_t^{\text{primaux}_j}$, output gate $o_t^{\text{primaux}_j}$, forget gate $f_t^{\text{primaux}_j}$, memory state $c_t^{\text{primaux}_j}$ and hidden state $h_t^{\text{primaux}_j}$. All these vectors are in \mathbb{R}^d .

A key innovation in our architecture is that, along with these auxiliary states, we will compute, position-by-position, a *gated, composite* state sequence comprised of c_t^{comp} and h_t^{comp} . The idea is to draw upon that auxiliary task, if any, that is best qualified to lend representation to the primary task, position by position. Recall from Section 3.3 that $g_t[j]$ denotes the relevance of the auxiliary model j at token position t of the primary instance. If $g_t[j] \rightarrow 1$, then the state of auxiliary model j strongly influences h_t^{comp} in the next step. If, for all j , $g_t[j] \rightarrow 0$, then no auxiliary model is helpful at position t . Therefore, the previous composite state is passed to the next time step as is. The skip possibility also makes training easier by countering vanishing gradients.

$$h_t^{\text{comp}} = \sum_{j=1}^m h_t^{\text{primaux}_j} g_t[j] + \left(1 - \sum_{j=1}^m g_t[j]\right) h_{t-1}^{\text{comp}} \quad (10)$$

$$c_t^{\text{comp}} = \sum_{j=1}^m c_t^{\text{primaux}_j} g_t[j] + \left(1 - \sum_{j=1}^m g_t[j]\right) c_{t-1}^{\text{comp}} \quad (11)$$

E.g., in target-dependent sentiment analysis, the auxiliary LSTM may identify sentiment-bearing words irrespective of target, and the composite state sequence prepares the stage for detecting the polarity of sentiment at a designated target.

$$\begin{bmatrix} \tilde{c}_t^{\text{primaux}_j} \\ o_t^{\text{primaux}_j} \\ i_t^{\text{primaux}_j} \\ f_t^{\text{primaux}_j} \end{bmatrix} = \begin{bmatrix} \tanh \\ \sigma \\ \sigma \\ \sigma \end{bmatrix} \left(W^{\text{aux}_j} \begin{bmatrix} x_t^{\text{prim}} \\ h_{t-1}^{\text{comp}} \end{bmatrix} \right) \quad (12)$$

$$c_t^{\text{primaux}_j} = \tilde{c}_t^{\text{primaux}_j} \odot i_t^{\text{primaux}_j} + c_{t-1}^{\text{comp}} \odot f_t^{\text{primaux}_j} \quad (13)$$

$$h_t^{\text{primaux}_j} = o_t^{\text{primaux}_j} \odot \tanh c_t^{\text{primaux}_j} \quad (14)$$

The overall flow of information between many embedding variables is complicated. Therefore, we have sketched it for clarity in Figure 2. To elaborate further, at step $t-1$, we




Figure 2: Variables defined using other variables and observed constants. Auxiliary and primary tasks are tied via W^{aux} . Self-loops indicate recurrence. Only states h shown for simplicity; c assumed to accompany them.

compute h_{t-1}^{comp} by combining all h^{primaux} using gate signals. The composite state is then fed into each auxiliary cell in the current step t .

3.5 Combining gating and composite states for primary prediction

Using the hidden states from the auxiliary LSTMs and the primary LSTM, we get the desired output using another model M^{prim} . For example, we can use a fully connected layer over the last hidden states or pooled hidden states for classification. We can use a fully connected layer on all hidden states for sequence labeling tasks.

$$\text{out}^{\text{prim}} = M^{\text{prim}}(h_1^{\text{prim}}, \dots, h_n^{\text{prim}}; h_1^{\text{comp}}, \dots, h_n^{\text{comp}}) \quad (15)$$

The loss of the primary task,

$$\text{loss}^{\text{prim}}(\text{out}^{\text{prim}}, y^{\text{prim}}), \quad (16)$$

is computed using the model output and the ground truth of the primary task data point.

3.6 Training and regularization

The system is trained jointly across auxiliary and primary tasks. We sample $\{x^{\text{prim}}, y^{\text{prim}}\}$ from the primary task dataset and $\{x^{\text{auxaux}_1}, y^{\text{auxaux}_1}\}, \{x^{\text{auxaux}_2}, y^{\text{auxaux}_2}\}, \dots, \{x^{\text{auxaux}_m}, y^{\text{auxaux}_m}\}$ from the auxiliary task datasets respectively. We compute the loss for the primary task and the auxiliary tasks as defined above. The total loss is the weighted sum of the loss of the primary task and the auxiliary tasks, where α_j is the weight of the auxiliary task (a tuned hyperparameter).

$$\text{loss}^{\text{all}} = \text{loss}^{\text{prim}} + \sum_{j=1}^m \alpha_j \text{loss}^{\text{auxaux}_j} \quad (17)$$

Optionally, we can use activity regularization on g_t , such that it is either close to 1 or close to 0. The regularization loss is:

$$\text{loss}^{\text{reg}} = \lambda \sum_{t=1}^n \|\min(g_t, 1 - g_t)\|_1 \quad (18)$$

which discourages g from “sitting on the fence” near 0.5. Regularization loss is added to loss^{all} .

4 Concrete instantiations of GIRNet

In this section, we present three concrete instantiations of GIRNet, and their detailed architecture.

4.1 Sentiment classification of code-switched passages

In code-switched text, words from two languages are used. Code switching is common in informal conversations and in social media where participants are multilingual users. Usually, ‘imported’ tokens are transliterated into a character set commonly used in the ‘host’ language, making span language identification (the ‘language identification task’) non-trivial. Let the languages be A and B with vocabularies D_A and D_B , respectively. We are given a word sequence (w_1, w_2, \dots, w_N) , where each word $w_i \in D_A \cup D_B$. Our goal is to infer a sentiment label from $\{-1, 0, +1\}$ for the whole sequence. As stipulated in Section 3, labeled instances and classification techniques are readily available for text in single languages (auxiliary task) (Tang et al. 2015a; Wang et al. 2017), but rare for code-switched text (primary task). Recent sequence models for sentiment classification in one language essentially recognize the sentiments in words or short spans. For such monolingual models to work well for code-switched text, the words and spans must be labeled with their languages, which is difficult and error-prone. Combining signals from auxiliary models is not trivial because some words in language A (e.g., ‘not’ in English) could modify sentiments expressed by other words in language B .

Datasets: For the primary task we use the sentiment classification dataset of English-Spanish code-switched sentences (Vilares, Alonso, and Gómez-Rodríguez 2015). Each sentence has a human labeled sentiment class in $\{-1, 0, 1\}$. The training and test sets contain 2,449 and 613 instances, respectively. We use two disjoint auxiliary task datasets. For sentiment classification of English sentences, we use the Twitter dataset provided by Sentistrength², which has 7,217 labeled instances. For sentiment classification of Spanish sentences, we use the Twitter dataset by Villena Roman et al. (2015), containing 4,241 labeled instances.

Model description: We use two LTSMs with 64 hidden units for the auxiliary task of English and Spanish sentiment classification. The output of only the last step is fed into a fully connected layer with 3 units and softmax activation. We use three separate fully connected layers for English, Spanish and English-Spanish tasks. For the primary RNN which produces the gating signal, we use a bi-directional LSTM with 32 units. The gating signal is produced by adding a fully connected layer of 3 units with softmax activation on each step of the primary RNN. The word embeddings are initialized randomly and trained along with the model.

4.2 POS tagging of code-switched sentences

Our second application is part-of-speech (POS) tagging of code-switched sentences. Unlike sentiment classification,

here a label is associated with each token in the input sequence. As in sentiment classification, we can use auxiliary models built from learning to tag monolingual sentences. A word in language A may have multiple POS tags depending on context, but context information may be provided by neighboring words in language B . E.g., in the sentence “*mein iss car ko like karta hu*” (meaning “I like this car”), the POS tag of ‘like’ depends on the Hindi text around it. The input sentence is again denoted by (w_1, w_2, \dots, w_N) , with each word $w_i \in D_A \cup D_B$. The goal is to infer a label sequence (y_1, y_2, \dots, y_N) , where y_i comes from the set of all POS tags over the languages. To predict the POS tag at each word, we apply a fully connected layer to the composite hidden state.

Datasets: For the primary task we use a Hindi-English code-switch dataset provided in a shared task of ICON’16 (Patra, Das, and Das 2018). It contains sentences from Facebook, Twitter and Whatsapp. It has 19 POS tags, and the training and test sets have 2102 and 528 instances, respectively. For the auxiliary dataset of Hindi POS tagging, we use the data released by (Sachdeva et al. 2014), containing 14,084 instances with 25 POS tags. For the auxiliary dataset of English POS tagging, we use the data released in a shared task of CoNLL 2000 (Tjong Kim Sang and Buchholz 2000), containing 8,936 instances with 45 POS tags.

Model description: We use two LTSMs with 64 hidden units for the auxiliary task of English and Hindi POS tagging. For the primary task, h_t^{comp} at each position is fed into a fully connected layer with 19 units with softmax activation. For the two auxiliary tasks, $h_t^{\text{auxaux}_1}$ and $h_t^{\text{auxaux}_2}$ are fed into the two separate fully connected layers of 25 and 45 units at each word for Hindi and English tags, respectively. For the primary RNN which produces the gating signal, we use a bi-directional LSTM with 32 units. The gating signal is produced by adding a fully connected layer of 3 units with softmax activation in each step of the primary RNN. The first two elements are the English and Hindi gates. The word embeddings are initialized randomly and trained along with the model.

4.3 Target-dependent sentiment classification

Our third application is target-dependent sentiment classification (TDSC), where we are given a sentence or short (assumed monolingual here) passage (w_1, w_2, \dots, w_N) with a designated span $[i, j]$ that mentions a named-entity target. The primary task is to infer the polarity of sentiment expressed toward the target. Our auxiliary task is whole-passage sentiment classification, for which collecting labeled instances is easier (Go, Bhayani, and Huang 2009; Sanders 2011). The passage-level task has a label associated with the whole passage, rather than a specific target position. E.g., in the tweet “*I absolutely love listening to electronic music, however artists like Avici & Tiesto copy it from others*”, the overall sentiment is positive, but the sentiments associated with both targets ‘Avici’ and ‘Tiesto’ are negative.

Datasets: For the primary task i.e., target-dependent sentiment classification, we use dataset of the SemEval 2014

²<http://sentistrength.wlv.ac.uk>

task 4 (Pontiki et al. 2014). It contains reviews from two domains – Laptops (2,328 training and 638 testing instances) and Restaurants (3,608 training and 1120 testing instances). The dataset for its corresponding auxiliary task is Yelp2014, consisting of similar types of 183,019 reviews. All of these datasets have three classes $\{+1, 0, -1\}$.

Model description: This application is another useful instantiation of GIRNet, because (i) by training on whole-passage instances, the auxiliary RNN learns to identify the polarity of words and their local combinations (“not very good”), (ii) the primary RNN learns to focus on the span related to the target entity.

In order to infer the sentiment of the target entity, the model would have to find the regions which are related to the given target entity and use the information learned from the auxiliary task. For qualitative analysis of our model, we will visualize the values of the scalar gates and find that words related to the target entity are passed through the auxiliary RNN and others are skipped/blocked. Here we implement GIRNet on top of TD-LSTM (Tang et al. 2015b). We can think of it as two separate instances, for the left and right context spans of the target.

For the primary task i.e., target-dependent sentiment classification, we use two separate RNNs to produce the gating signals and capture primary features which are not captured by the auxiliary RNN. Since there is only one auxiliary task, the controller either skips the RNN at a particular word or it may pass it to the auxiliary RNN. Given an input sentence and a target entity, we split the sentence at the position of the target entity. We input the left half of the sentence to the left primary and auxiliary RNN and the right half to the right auxiliary and primary RNN. Rather than pooling the states, we do a weighted pool where we sum the hidden states after multiplying with the gate values. Similar to TD-LSTM, The pooling of left and right is done separately and then concatenated.

For the auxiliary task i.e., sentiment classification of the complete sentence, we use a left RNN and a right RNN so that we can couple them with the two primary RNNs. We run both the left and right auxiliary RNNs on the input text and the reversed input text respectively. At each token of the input sentence we concatenate the hidden state of the left RNN and the right RNN. For sentiment classification of the complete sentence, we take an average pool the features of all the token positions and pass it to a fully connected layer with softmax activation for classification of the sentiment score.

5 Comparative Evaluation

In this section, we first describe baseline approaches, followed by a comparison between them and GIRNet over the three concrete tasks we set up in Section 4.

5.1 Baseline Methods

Here we briefly describes several MTL-based baselines (along with their variants) used in all three applications mentioned in Section 4. A few other application-specific baselines are mentioned in Section 5.2.

Model	Accuracy	Macro F ₁	Precision	Recall
LSTM (No MTL)	59.22	58.34	58.86	58.01
HardShare 1L	60.36	59.65	60.28	59.23
HardShare 2L	57.10	55.21	56.59	54.88
LowSupShare	55.95	54.42	55.42	54.23
LowSupConcat	56.61	55.84	57.22	55.25
XStitch 1L	59.71	58.59	59.94	58.08
XStitch 2L	56.44	48.31	58.76	51.35
Sluice	58.56	57.71	57.61	58.18
PSP-MTL	59.22	59.23	59.04	59.62
SSP-MTL	59.71	59.75	59.45	60.33
Meta-MTL	57.59	54.56	58.39	54.24
Coupled RNN	56.77	55.69	55.85	55.70
GIRNet 1L	63.30	62.36	63.35	61.83
GIRNet 2L	61.99	61.23	61.79	61.13

Table 1: Accuracy of the competing models for the sentiment classification of code-switched passages. F₁ score, precision and recall are macro-averaged.

LSTM (no MTL): LSTM trained on just the primary task.

Hard Parameter Sharing (HardShare): This model (Caruana 1993) uses the same LSTM for all the tasks along with the separate task-specific classifiers. We show results for both 1-layer LSTM (**HardShare 1L**) and 2-layer LSTM (**HardShare 2L**).

Low Supervision (LowSup): This model (Søgaard and Goldberg 2016) uses a 2-layer LSTM, where the auxiliary tasks are trained on the lower layer LSTM, and the primary task is trained on the higher layer LSTM. Here we show results on two schemes: 1) **LowSupShare:** same LSTM is used for all auxiliary tasks, 2) **LowSupConcat:** outputs of separate layer-1 LSTM concatenated.

Cross-stitch Network (XStitch): This model (Misra et al. 2016) has separate LSTM for each task. The amount of information shared to the next layer is controlled by trainable scalar parameters. We show results for both 1-layer LSTM (**XStitch 1L**) and 2-layer LSTM (**XStitch 2L**).

Sluice Network (Sluice): This is an extension (Ruder et al. 2017a) of cross-stitch network, where the outputs of the intermediate layers are fed to the final classifier. Here, rather than applying the stitch module on the LSTM outputs at a layer, they split the channels of each LSTM and apply the stitch module.

Shared-private sharing scheme: In this architecture there is a common shared LSTM over all tasks and separate LSTMs for each task. We show results on two schemes 1) parallel shared-private sharing scheme (**PSP-MTL**) as described by (Liu, Qiu, and Huang 2017), where the outputs of the private LSTM and shared LSTM are concatenated; and 2) stacked shared-private sharing scheme (**SSP-MTL**) as described by (Chen et al. 2018b), where output of the shared LSTM is concatenated with the sentence which is fed to the private LSTM.

Meta multi-task learning (Meta-MTL): In this model (Chen et al. 2018b), the weight of the task specific LSTM is a function of a vector produced by the Meta LSTM which is shared across all the tasks.

Model	Accuracy	Macro F ₁	Precision	Recall
LSTM (no MTL)	61.46	40.55	47.18	40.84
HardShare 1L	62.19	40.55	45.00	39.18
HardShare 2L	62.10	42.33	47.77	40.87
LowSupShare	61.99	43.62	46.88	42.53
LowSupConcat	62.66	43.82	50.17	41.75
XStitch 1L	63.28	44.39	54.51	42.11
XStitch 2L	62.88	41.78	48.03	39.75
Sluice	60.90	40.81	43.95	40.35
PSP-MTL	62.93	41.88	47.85	39.98
SSP-MTL	62.90	37.63	46.36	35.49
Meta-MTL	62.25	41.33	47.4	39.77
Coupled RNN	62.44	41.91	52.73	40.20
GIRNet 1L	64.29	47.58	53.51	45.48
GIRNet 2L	63.13	45.75	51.34	43.65

Table 2: Accuracy of the competing models for the POS tagging of code-switched sentences. F₁ score, Precision and Recall are macro averaged.

Coupled RNN: This model (Liu, Qiu, and Huang 2016) has an LSTM for each task which uses the information of the other task.

5.2 Experimental Results

Tables 1, 2 and 3 show the results for GIRNet and other MTL baselines for three applications mentioned in Section 4. We observe a significant improvement of GIRNet over all other models. For a fair comparison with multi-layer LSTM MTL models, we show results of GIRNet with two LSTM layers. This is because some models like Sluice and LowSup can only be implemented with > 1 LSTM layers. We see that for three applications, all the 2-layer LSTM based models have worse performance compared to the single layer LSTM models. However, GIRNet with two layers outperforms all other 2-layer LSTM models. The Coupled RNN and the Meta-MTL are defeated by GIRNet. It uses scalar gates and has fewer degrees of freedom, which helps in learning with less data. GIRNet beats XStitch and Sluice as they do not have any sharing of information at a granularity of words.

We also compare GIRNet with some single-task baselines in Table 3. These methods are designed particularly for target-specific sentiment classification. TD-LSTM (Tang et al. 2015b) is the no-MTL baseline for the task of TDSC as in this case, GIRNet is implemented on top of TDSC. In the TD-LSTM + Attention model (Wang et al. 2016), attention score at each token is computed which is used to do a weighted pooling of the hidden states. In Memory network (MemNet) (Tang, Qin, and Liu 2016), multiple modules of memory are stacked and the initial key is the target entity.

5.3 Visualization

To get insight into GIRNet’s success, we studied the scalar gate values at each word of input sentences. Table 4 shows scalar gate values for a few TDSC instances. We see that words associated with the target entity get larger gate values, which is particularly beneficial for multiple entities and diverse sentiments.

Model	Laptop		Restaurant	
	Accuracy	F ₁	Accuracy	F ₁
TD-LSTM	71.38	68.42	78.00	66.73
TD-LSTM + Att.	72.14	67.45	78.89	69.01
MemNet	70.33	64.09	78.16	65.83
HardShare 1L	72.27	66.71	78.66	66.08
HardShare 2L	70.72	63.70	78.13	66.16
LowSupShare	71.65	65.74	79.64	68.46
XStitch 1L	71.81	65.5	79.02	68.55
XStitch 2L	73.05	67.63	78.93	68.24
Sluice	71.50	66.10	78.84	69.62
PSP-MTL	71.65	65.45	79.55	68.75
SSP-MTL	70.87	65.93	79.11	69.32
Meta-MTL	71.34	66.19	78.66	68.17
Coupled RNN	71.34	64.68	79.19	65.98
GIRNet 1L	74.92	69.67	82.41	74.35
GIRNet 2L	75.86	71.39	80.18	69.14

Table 3: Accuracy and macro F1 of the competing models for the target-dependent sentiment classification.

Target	Gate heatmap
Soup	the service is great , my <u>soup</u> always ar- rives nice and hot .
Appetizers	<u>appetizers</u> are ok , but the service is slow.
Service	appetizers are ok , but <u>the service</u> is slow .

Table 4: Gating heatmaps of Target-dependent sentiment classification. In each case, words associated with the target get the largest gate values.

6 Conclusion

Sequence labeling tasks are often applied to multi-domain (such as code-switched) sequences. But labeled multi-domain sequences are more difficult to collect compared to single-domain (such as monolingual) sequences. We therefore need *sequence MTL*, which can train auxiliary sequence models on single-domain instances, and learn, in an unsupervised manner, how to interleave composite sequences by drawing on the best auxiliary sequence model cell at each token position. We tested our model on three concrete applications and obtained larger accuracy gains compared to other MTL architectures.

Acknowledgement: Partly supported by Microsoft Research India travel grant, IBM, Early Career Research Award (SERB, India), and the Center for AI, IIT Delhi, India.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *CoRR* abs/1409.0473.
- Caruana, R. 1993. Multitask learning: A knowledge-based source of inductive bias. In *ICML*, 41–48.
- Chen, J.; Chen, K.; Chen, X.; Qiu, X.; and Huang, X. 2018a. Exploring shared structures and hierarchies for multiple nlp tasks. *arXiv preprint arXiv:1808.07658*.

- Chen, J.; Qiu, X.; Liu, P.; and Huang, X. 2018b. Meta multi-task learning for sequence modeling. *arXiv preprint arXiv:1802.08969*.
- Duong, L.; Cohn, T.; Bird, S.; and Cook, P. 2015. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In *ACL/IJCNLP*, volume 2, 845–850.
- Go, A.; Bhayani, R.; and Huang, L. 2009. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford* 1(2009):12.
- Lample, G.; Ballesteros, M.; Subramanian, S.; Kawakami, K.; and Dyer, C. 2016. Neural architectures for named entity recognition. *arXiv preprint arXiv:1603.01360*.
- Liu, P.; Qiu, X.; and Huang, X. 2016. Recurrent neural network for text classification with multi-task learning. *arXiv preprint arXiv:1605.05101*.
- Liu, P.; Qiu, X.; and Huang, X. 2017. Adversarial multi-task learning for text classification. *arXiv preprint arXiv:1704.05742*.
- Maurer, A.; Pontil, M.; and Romera-Paredes, B. 2016. The benefit of multitask representation learning. *Journal of Machine Learning Research* 17(81):1–32.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS Conference*, 3111–3119.
- Misra, I.; Shrivastava, A.; Gupta, A.; and Hebert, M. 2016. Cross-stitch networks for multi-task learning. In *IEEE CVPR*, 3994–4003.
- Patra, B. G.; Das, D.; and Das, A. 2018. Sentiment analysis of code-mixed indian languages: An overview of sail_code-mixed shared task@ icon-2017. *arXiv preprint arXiv:1803.06745*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global vectors for word representation. In *EMNLP Conference*, volume 14, 1532–1543.
- Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; and Manandhar, S. 2014. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 27–35. Dublin, Ireland: ACL.
- Ruder, S.; Bingel, J.; Augenstein, I.; and Søgaard, A. 2017a. Learning what to share between loosely related tasks. *arXiv preprint arXiv:1705.08142*.
- Ruder, S.; Bingel, J.; Augenstein, I.; and Søgaard, A. 2017b. Sluice networks: Learning what to share between loosely related tasks. *arXiv preprint arXiv:1705.08142*.
- Ruder, S. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Sachdeva, K.; Srivastava, R.; Jain, S.; and Sharma, D. M. 2014. Hindi to english machine translation: Using effective selection in multi-model SMT. In *LREC*, 1807–1811.
- Sanders, N. 2011. Twitter sentiment corpus.
- Schmid, H. 1994. Part-of-speech tagging with neural networks. In *Computational Linguistics Conference*, 172–176.
- Søgaard, A., and Goldberg, Y. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *ACL Conference*, 231–235.
- Tang, D.; Qin, B.; Feng, X.; and Liu, T. 2015a. Effective LSTMs for target-dependent sentiment classification. *arXiv preprint arXiv:1512.01100*.
- Tang, D.; Qin, B.; Feng, X.; and Liu, T. 2015b. Effective lstms for target-dependent sentiment classification. *arXiv preprint arXiv:1512.01100*.
- Tang, D.; Qin, B.; and Liu, T. 2016. Aspect level sentiment classification with deep memory network. *arXiv preprint arXiv:1605.08900*.
- Tjong Kim Sang, E. F., and Buchholz, S. 2000. Introduction to the conll-2000 shared task: Chunking. In *Proceedings of the 2nd workshop on Learning language in logic and the 4th conference on Computational natural language learning-Volume 7*, 127–132.
- Vilares, D.; Alonso, M. A.; and Gómez-Rodríguez, C. 2015. Sentiment analysis on monolingual, multilingual and code-switching twitter corpora. In *6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 2–8.
- Villena Roman, J.; Garcia Morera, J.; Martinez Carmara, E.; and Jimenez Zafra, S. M. 2015. Tass 2014 - the challenge of aspect-based sentiment analysis. *Procesamiento del Lenguaje Natural* 54:61–68.
- Wang, Y.; Huang, M.; Zhao, L.; et al. 2016. Attention-based LSTM for aspect-level sentiment classification. In *EMNLP Conference*, 606–615.
- Wang, B.; Liakata, M.; Zubiaga, A.; and Procter, R. 2017. TDParse: Multi-target-specific sentiment recognition on Twitter. In *EACL*, 483–493.