

2020

Universal Vector Neural Machine Translation with Effective Attention

Joshua Yi
SMU, joshuayi@smu.edu

Satish Mylapore
SMU, smylaporesaravanabha@mail.smu.edu

Ryan Paul
SMU, rqpaul@mail.smu.edu

Robert Slater
Advisor SMU, rslater@mail.smu.edu

Follow this and additional works at: <https://scholar.smu.edu/datasciencereview>



Part of the [Applied Statistics Commons](#), and the [Other Mathematics Commons](#)

Recommended Citation

Yi, Joshua; Mylapore, Satish; Paul, Ryan; and Slater, Robert (2020) "Universal Vector Neural Machine Translation with Effective Attention," *SMU Data Science Review*. Vol. 3 : No. 1 , Article 10.
Available at: <https://scholar.smu.edu/datasciencereview/vol3/iss1/10>

This Article is brought to you for free and open access by SMU Scholar. It has been accepted for inclusion in SMU Data Science Review by an authorized administrator of SMU Scholar. For more information, please visit <http://digitalrepository.smu.edu>.

Universal Vector Neural Machine Translation With Effective Attention

Satish Mylapore¹, Ryan Quincy Paul¹, Joshua Yi¹, and Robert D. Slater¹

Master of Science in Data Science, Southern Methodist University, Dallas TX 75275
USA {smylaporesaravanabhava, rqpaul, jyi, rslater}@smu.edu

Abstract. Neural Machine Translation (NMT) leverages one or more trained neural networks for the translation of phrases. Sutskever introduced a sequence to sequence based encoder-decoder model which became the standard for NMT based systems. Attention mechanisms were later introduced to address the issues with the translation of long sentences and improving overall accuracy. In this paper, we propose a singular model for Neural Machine Translation based on encoder-decoder models. Most translation models are trained as one model for one translation. We introduce a neutral/universal model representation that can be used to predict more than one language depending on the source and a provided target. Secondly, we introduce an attention model by adding an overall learning vector to the multiplicative model. With these two changes, by using the novel universal model the number of models needed for multiple language translation applications are reduced.

1 Introduction

Neural Machine Translation (NMT) [1] is a significant recent development in large scale translation [2, 3]. The traditional translation model introduced by Koehn et al. 2003 [4] was trained on a single large neural model with layers that are trained separately requiring many resources and effort. Today, most industry players have adopted a neural network based machine translation system derived from the Recurrent Neural Network (RNN) encoder-decoder model introduced by Cho et al. 2014 [5]. For machine translation, the encoder is used with the source language to encode the sentence input into a vector representation for the decoder. The decoder uses the encoded sequence to begin predicting the target sequence. There were several advancements to this model by the introduction of different types of RNNs such as LSTM (Long Short Term Memory) [6–9], GRU (Gated Recurrent Unit) [10], and Bi-RNN (Bidirectional RNN) [11] which was introduced to address the vanishing gradient problem [12] that was encountered during the training of the simple recurrent neural network.

Gated recurrent networks failed to fully resolve the problem of the encoder-decoder network [13] which is the ability to learn and maintain information of the encoder for longer sentences. This is where the attention mechanism was introduced by Graves et al. 2014 that is based on the cosine similarity of the sentences [14], Bahdanau et al. 2014 which concatenates the encoder and decoder

information [15], and Loung et al. 2015 that uses the dot product of the the encoder and decoder information to score the attention on the target sequence [3, 16]. The introduction of attention mechanisms increased the scalability of machine translation at the cost of performance during training.

The latest development in the machine translation space is the introduction of the Transformer model by Vaswani et al. 2017 [17]. The Transformer model focuses more on self-attention and fully leverages recurrent networks. It promotes self-attention in both the encoder and the decoder where the encoding of the source sequences are done in parallel. This reduces the training time significantly. The decoder prediction is auto-regressive which means it predicts each word at a time in a regressive state. Vaswani claims that the results of the transformer model has a significant improvement in prediction accuracy when compared to other recent models in the NMT space with the use of a German translation task [17].

The transformer model is still in the incubation and adoption stage in current industry practice. This is due to its restricted context length during translation (fixed-length context). Furthermore, at present all RNN encoder-decoder based machine translation models still use a single model architecture for a translation job. For example, if a task requires translation from Spanish to English, one model will be trained. Another model would be trained to translate from English to Spanish. One model corresponds to one translation task, hence separate models are required. In this research, we seek to build a singular model to translate multiple languages. For the purpose of this research we have considered English-Spanish and Spanish-English translation using the same model.

All machine translation mechanisms to date use language specific encoders for each source language [18]. This paper will detail a novel method of hosting multiple neural machine translation tasks within the same model as follows. Section 2 will cover related works on the fundamental concepts of the sequence to sequence Recurrent Neural Network based Encoder-Decoder model, the additive attention model by Bahdanau, and wrap-up with the Dual Learning method introduced by Microsoft. Section 3 outlines the architecture for the universal vector model and discusses each layer. Section 4 discusses the training method for the universal model, while Section 5 explains the dataset and how it is used for training. Section 6 is an overview of the BLEU score. The translation results of the Universal Vector is explained in Section 7 then Section 8 presents the analysis of the BLEU score, loss results, and attention model performance. Section 9 goes over limitations and potential steps to take in the future with Section 10 discussing previously considered experiments. Finally, the paper is concluded with some parting thoughts on the development of this novel model in Section 11.

2 Related Work

This section will go over the associated work related to building the Universal Vector Neural Machine Translation model. First Recurrent Neural Network Based Encoder-Decoder Models proposed by Sutskever et al. and Cho et al. will

be discussed. Next, the attention mechanism first proposed by Bahdanau et al. will be detailed. Finally, the Dual Learning model training approach is explained.

2.1 Recurrent Neural Network Based Encoder-Decoder Models

Many NMTs are built upon the fundamental Recurrent Neural Network (RNN) based Encoder-Decoder model as proposed by Sutskever et al. (2014) and by Cho et al. (2014) [5, 10, 19, 20]. This model uses two networks, an encoder and a decoder, to learn sequences of information and make predictions. In this model a sequence of input x is provided to the encoder, an RNN. An RNN allows for outputs of iterations through a network to be passed on as input to future iterations [21–24]. x is processed word by word $(1, 2, \dots, t)$ over multiple iterations. Each iteration calculates a hidden state that is based on the current word in a phrase (x_t) and the hidden states of previous iterations (h_{t-1}) . This is represented at a high level in Equation 1 below with a non-linear equation q calculating hidden states at each position [15].

$$h_t = f(x_t, h_{t-1}) \quad (1)$$

Once all hidden states have been calculated, a function g will then return a single fixed length context vector c with each hidden state as inputs like in Equation 2 below. c represents the full summary of the output of the encoder network [15].

$$c = q(\{h_i, \dots, h_{T_x}\}) \quad (2)$$

The output of the encoder, c , is then fed into the decoder which is another trained RNN. The decoder emits the prediction for each input y at iteration t where these conditional outputs come together as a probability distribution like below in Equation 3 [15].

$$p(y_t | \{y_1, \dots, y_{t-1}\}, c) = g(y_{t-1}, s_t, c) \quad (3)$$

g is another non-linear function that takes in the previously predicted words (y_{t-1}) , the hidden state of the current iteration of the network (s_t) , and the context vector from above (c) . p represents a predicted target sequence of words for a given input sequence of words with conditional probability [25]. This is the basis of the Encoder-Decoder model that has been used heavily in neural machine translation.

2.2 Attention Mechanism

Attention mechanisms have gained visibility recently as they are able to improve the performance of translation by helping the encoder and decoder to align by providing guidance on what parts of a large sentence will be most useful in predicting the next word [15–17, 26]. In recent years many attention models have been introduced such as Bahdanau et al [15] which concatenates (referred

to as "concat" in Luong, et al., 2015 [16] and as "additive attention" in Vaswani, et al., 2017 [17]) forward and backward information from the source. This model changes the fundamental RNN Encoder-Decoder described above in a variety of ways.

The encoder is built using a bi-directional recurrent neural network that contains two models. Each model will compute hidden states in either direction from a given input x_i . This will yield two hidden states, \vec{h}_i and \overleftarrow{h}_i . These two hidden states are concatenated together to form a vector h_i as below in Equation 4 that will represent the whole sentence emanating out from a given input word and are referred to as annotations [15].

$$h_i = \begin{bmatrix} \vec{h}_i \\ \overleftarrow{h}_i \end{bmatrix} \quad (4)$$

Due to RNNs tendency toward recency bias, the words immediately surrounding around a given input (x_i) will be better represented in the input word's annotation (h_i). This will be reflected when calculating attention which begins with a replacement to the fixed length context vector c mentioned in Section 2.1. A new context vector c_j is calculated for every output word y_j . This begins with a scoring function e_{ij} which will represent the importance of the hidden state output from the previous iteration of the decoder s_{j-1} to a given annotation h_i represented by Equation 5 below [15]. A higher score will represent higher importance.

$$e_{ij} = a(s_{j-1}, h_i) \quad (5)$$

e_{ij} is then fed into a softmax function α_{ij} found below in Equation 6 which will return a vector of numbers that all sum up to one that represents the weight of each annotation with respect to the given position of y_j [15].

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^{T_x} \exp(e_{jk})} \quad (6)$$

Finally, the context vector c_j unique to each word output by the decoder y_j is calculated with the summation found in Equation 7 below.

$$c_j = \sum_{i=1}^{T_x} \alpha_{ij} h_i \quad (7)$$

Vector c_j will be used in the calculation of hidden states in the decoder s_j found in Equation 8. s_j , the previously predicted words y_{j-1} , and the c_j will then be used in calculating the output of each iteration of the decoder at step j as in Equation 9 below. The output is a vector of probabilities of each possible word that could be predicted at y_j . The context vector c_i will weigh in words at input position i that scored a higher importance from e_{ij} in Equation 5 more than others which represents attention. This in contrast to taking the whole vector of input words into account at every j th position of y [15].

$$s_j = f(s_{j-1}, y_{j-1}, c_j) \quad (8)$$

$$p(y_j | \{y_1, \dots, y_{j-1}\}, x) = g(y_{j-1}, s_j, c_j) \quad (9)$$

This is an early implementation of attention proposed by Bahdanau et al. 2014 [15]. Many other forms of attention have been proposed since. Luong et al refers to Bahdanau’s attention mechanism as ”global attention.” In turn, Luong et al. proposed a ”local attention method” that focuses on smaller portions of context instead of applying attention weights on the entire source text [16]. The new attention mechanism proposed in this paper combines the two.

2.3 Dual Learning

In a paper proposed by Microsoft Research [27], the team considered a dual learning mechanism to handle the complexities in the training data labeling. The dual learning mechanism considers two agents, one agent for the forward translation model (source to target language) and the second agent is considered for the dual translation (target to source language). These models use two different corpora for training which are not parallel data sets. This enables reinforcement learning for the convergence of source and target language. The inputs considered on the Microsoft Research paper are ”Monolingual corpora D_A and D_B , initial translation models θ_{AB} and θ_{BA} , language models LMA and LMB, hyper-parameter α , beam search size K , learning rates $\gamma_{1,t}$, $\gamma_{2,t}$.” [27] The experiment used to test dual training uses two separate models, one for each translation direction.

In this paper, there are two contributions based on RNN Encoder-Decoder based machine translations. First, a neutral/universal vector representation for machine translation is introduced. Then a modified attention mechanism based on global attention mechanism proposed in Luong et al. [16] and Bahdanau et al. [15] is discussed. Finally, testing of the proposed neutral vector representation with modified attention mechanism are examined and the results are presented.

3 Model Architecture

The architecture of the model is built on top of the basic sequence to sequence model and modified to translate more than one language. A high level architecture diagram is found in Fig. 1 below. This model contains two networks, an encoder and a decoder, with embedded inputs and outputs for each. It also contains the modified attention mechanism and a Fully Connected Layer. In the current structure, the source text is inserted into the Input Embedding layer which contains the Encoder RNN. There are multiple Input Embedding Layers to handle different source texts such as Spanish, English, German, etc. From the Input Embedding layer, the results (context vectors) are fed into the Target Embedding layer which contains the Decoder RNN along with the modified Attention layer. As is the case with the Encoder portion of the system, there are

multiple Target Embedding layers for multiple target languages. Lastly, the output from the Target Embedding layer is passed into the Target Fully Connected layer. The result is a vector of probabilities for words in the target language. From this vector, the predicted phrase is converted from a numeric vector representation to words in a natural language.

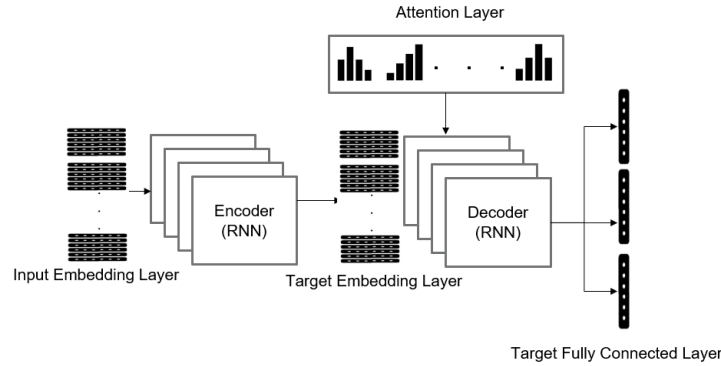


Fig. 1. Model architecture detailing the encoder and decoder networks and their inputs.

3.1 Embedding Layers

The model starts with an encoder layer to generate a vector that will be fed to the decoder to generate predictions in a target language. Embedding vectors for the encoder will be built as layers, which are considered to be the source input to the encoding layer as Equation 10 below.

$$es \sim es_1, es_2, \dots, es_n \tag{10}$$

Here, es is the embedding vector and each number represented by s represents a different language used as a source for translation. This will be used as the first layer in the encoder network. Similarly, an embedding layer for the decoder is also built as in Equation 11 below.

$$et \sim et_1, et_2, \dots, et_n \tag{11}$$

et is the embedding vector and each number represented by t represents a different language used as the target prediction.

3.2 Attention Layer

The modified attention mechanism considers a context vector c created by the encoder. This vector c is created based on all the hidden vectors of the hidden

states during the encoding phase. This vector has a representation of each word from the source. The attention score used to predict each target word is calculated by the dot product of the hidden value of each prediction and the encoded output [17]. This scoring mechanism is based on the global attention method proposed by Luong et al. [16]. Learning weights were introduced into the dot product score, which is calculated using Equation 12 below. The purpose of this is to learn the overall weights of the dot product score.

$$score(h_t, \bar{h}_s) = v_a^\top (h_t^\top W_a \bar{h}_s) \quad (12)$$

The context vector is computed by taking the dot product of the encoder output. This is done to add global alignment to the context vector, which will be used to estimate the score of the next prediction.

The attention mechanism will be used to align decoder predictions of the target vector. Attention weights for each target language is defined as Equation 13 below. Where a is the attention weight and t is the target language.

$$a_t \sim (a_{t1}, a_{t2}, \dots, a_{tn}) \quad (13)$$

3.3 Fully Connected Layer

The last layer is a fully connected layer, which will be attached to the size of the target language as seen in Equation 14 below where fc is a connected layer and t is each target language. The purpose of the fully connected layer is to act as a classifier for each targeted translated text.

$$fc \sim (fc_{t1}, fc_{t2}, \dots, fc_{tn}) \quad (14)$$

4 Model Training

The model training process considers training for each set of translations in a sequence. For this experiment, Spanish and English languages are considered for training with Gated Recurrent Units (GRUs) as the recurrent unit were considered to address the long term dependencies [6, 28–30] where $W1_{en}$, V_{en} , $W1_{sp}$, and V_{sp} all act as attention weight matrices, e_1 is the embedding layer for Spanish and e_2 is the embedding layer for English, d_1 is the fully connected layer for Spanish and d_2 is the fully connected layer for English. The weight matrices for each gate in the GRU are represented as Γ_{ue} and Γ_{ud} for the update gates, Γ_{re} and Γ_{rd} for relevance gates, c_e and c_d for the context gates, and $h_{(te)}$ and $h_{(td)}$ represent the hidden vectors. The weights will be initialized using the Glorot Uniform Initializer [31]. The Adam optimization algorithm has been used with a learning rate α and a decayed learning rate γ . Loss will be measured using the discrete classification methodology which leverages the sparse Softmax cross-entropy with logits loss. Spanish-English will be used as the parallel database where each example of language will be trained in parallel.

A pseudo algorithm of the training process is given below.

Algorithm 1 Model Training Process

Require: Parallel Dataset D with phrases in both $Lang_A$ and $Lang_B$

- 1: **repeat**
 - 2: **for all** Phrases p made up of $Lang_A, Lang_B$ in D **do**
 - 3: **Block 1:** Encode $Lang_A$ example and compute encoder GRU layer $\Gamma_{ue}, \Gamma_{re}, c_e, h_{(te)}$
 - 4: Decode to predict $Lang_B$ (English) using the encoder output of $c_e, h_{(te)}$
 - 5: Compute Compute $W1_{en}$ and V_{en} for alignment model
 - 6: Compute $\Gamma_{ud}, \Gamma_{rd}, c_d, h_{(td)}$ for each prediction.
 - 7: Compute d_1
 - 8: Compute loss L_1 using the sparse softmax cross entropy with logits loss
 - 9: **Block 2:** Encode $Lang_B$ example and compute encoder GRU layer $\Gamma_{ue}, \Gamma_{re}, c_e, h_{(te)}$
 - 10: Decode to predict $Lang_A$ (Spanish) using the encoder output of $c_e, h_{(te)}$
 - 11: Compute $W1_{sp}$ and V_{sp} for alignment model
 - 12: Compute $\Gamma_{ud}, \Gamma_{rd}, c_d, h_{(td)}$ for each prediction.
 - 13: Compute d_2
 - 14: Compute loss L_2 using the sparse softmax cross entropy with logits loss
 - 15: Compute total loss $L = L_1 + L_2$
 - 16: Optimize using Adam optimization with learning rate α and decayed learning rate γ .
 - 17: **end for**
 - 18: **until** All Phrases $Lang_A, Lang_B$ have been processed
-

This process is repeated for all the examples. Here, we keep the encoder and decoder same for all the languages that are trained for prediction. If another translation is added, then the blocks are repeated for each language.

5 Dataset

Parallel datasets for Spanish and English are used for training of the Universal Vector model. Data is taken from Many Things, an online resource for English as a Second Language Students¹ [32]. It contains 122,936 pairs of phrases in English and a corresponding Spanish translation.

5.1 Training

The universal vector model is trained using a modified version of the Dual Training method proposed by Xia et al. [27]. This model is trained with a sequence for each training data, first with Spanish to English and then English to Spanish for every iteration of the dataset mentioned above. Sample phrases in both

¹ The primary source of the dataset used in this study as well as many more language pairings can be found at <http://www.manythings.org/anki/>. We used a copy hosted by the TensorFlow team at <http://storage.googleapis.com/download.tensorflow.org/data/spa-eng.zip>

English and Spanish were used to test the predictive ability of the network into both languages. The model has been trained at 20, 30, and 40 epochs to see the effectiveness of the model as the amount of training increases.

6 BLEU Score

A Bilingual Evaluation Understudy (BLEU) score was used as a metric to determine the effectiveness of our NMT. BLEU was developed as a replacement for human-based validation of machine based translation that was becoming an expensive bottleneck due to the need for language expertise. The formula to calculate the score is language independent, does not need to be trained, and is able to mimic human evaluation. The function takes in the translated sentence and one or more reference sentences that it will be compared with. Groups of words, or n -grams, in the translated sentence to be evaluated are matched with n -grams in the reference sentences.

The first step in the scoring process is to calculate a precision score by taking the number of matching n -grams between the evaluated sentence and the reference sentences. This number is then divided by the total count of the n -grams in both the references sentences and the evaluated translation. This equation can be found below in Equation 15 [33]. Another consideration when determining a score for a translation is the length of the output. There are many ways to say the same thing in most languages, but using too many words can introduce ambiguity and using too few words may not provide enough nuance.

$$p_n = \frac{\sum_{C \in \{Candidates\}} \sum_{n-gram \in C} Count_{clip}(n-gram)}{\sum_{C \in \{Candidates\}} \sum_{n-gram \in C} Count(n-gram)} \quad (15)$$

Penalties are in place to ensure sentence of proper length score better. The precision score equation has a built in penalty for candidate sentences that are too long as more n -grams will increase the denominator and lead to a smaller score. For translations that are too short, a penalty is introduced in the form of a Brevity Penalty (BP) as in Equation 16 [33] below. r is the count of the words in the reference sentence that is closest to the translated sentence being evaluated. c is the length of the candidate sentence. If there is a match, the BP is 1 and there is no penalty assessed. If there is not an exact match in length, then a penalty is assessed according to an exponentiation of e .

$$BP = \begin{cases} 1, & \text{if } c > r \\ e^{(1-r/c)} & \text{if } c \leq r \end{cases} \quad (16)$$

The overall BLEU score for a candidate sentence is the product of the brevity penalty and the exponential sum of the product of the log of the precision score multiplied by a positive weight w_n . This weight is based on the number of n -grams N such that $w_n = 1/N$. The overall score is found by using Equation 17 below. Equation 18 below is a form of the equation that is used to provide values

that are more able to be ranked among other candidate translated sentences by applying a log to the whole sentence.

$$BLEU = BP \cdot \exp\left(\sum_{n=1}^N w_n \log p_n\right) \quad (17)$$

$$\log BLEU = \min\left(1 - \frac{r}{c}, 0\right) + \sum_{n=1}^N w_n \log p_n \quad (18)$$

The NLTK BLUE Score package is used for evaluation of the model².

7 Results

The following section will cover the translation results obtained from the Universal Vector model. It will discuss the translations from English to Spanish and Spanish to English.

7.1 Translations

Example phrases in each language were fed to the model. Two example pairs of phrases are found in Table 1 below.

Table 1. Example phrases used for testing

English	Spanish
They abandoned their country	Ellos abandonaron su país
This is my life	Esta es mi vida

The results of the English to Spanish task can be found in Table 2 below. In the case of our model, a BLEU score cannot capture the accuracy since it is based on matching n -grams. The sentences were too short to have anything larger than matching bigrams which are too small for the scoring algorithm. The result of the first phrase perfectly matched the reference sentence found in Table 1. The output of the second phrase switched the gender of the word for "this" in English from "esto" to "esta". Without more context before a phrase, the model is not able to consistently determine the genders of specific words.

When English and Spanish are flipped, the model provided similar results. The resulting English outputs can be found below in Table 3. Small differences are present, again due to small gender differences that small sentences will be expected to yield without proper context for pronouns.

² documentation for the BLEU score functionality can be found https://www.nltk.org/_modules/nltk/translate/bleu_score.html

Table 2. English input and Spanish output

English Input	Spanish Output
They abandoned their country	Ellos abandonaron su país
This is my life	Esto es mi vida

Table 3. Spanish input and English output

Spanish Input	English Output
Ellos abandonaron su país	They abandoned his country
Esta es mi vida	This is my life

Sentences longer than four or five words yielded very poor results. This is due to the small dataset and low number of training iterations when compared with other papers in the NMT space such as most of those cited in this paper. With a larger dataset and more training time the model will better handle longer phrases.

8 Model Analysis

The following section covers the analysis and each subsection that follows is the analysis discussion for the BLEU score, Loss analysis and Attention model.

8.1 BLEU Score Analysis

Applying the BLEU score to the Universal Vector Model resulted in unfavorable scores. Table 4 shows the results of the BLEU score from Spanish to English and English to Spanish. BLEU score calculations are provided as part of this work to show the minimum capability of this model to translate more than one language using a single universal model. The score from this work should not be compared with other translation models like Bert and other Transformer based models [17, 34]. There are two main reasons behind this. First, the tested sentence words were short in length. Second, the short sentences did not meet the minimal n-gram length of 2 for proper scoring. The use of longer sentences could have solved these issues, however the model had difficulty translating longer sentences at the level of training we were able to accomplish in the time given (60 epochs).

Table 4. BLEU Score Results

Sentence	Direction	BLEU Score
esto es mi vida.	English to Spanish	8.3882e-155
this is my life.	Spanish to English	6.8681e-78

8.2 Loss Analysis




Fig. 2. Loss analysis by epoch during training of the model.

Since the BLEU score could not properly capture model accuracy for testing, more attention has been placed into minimizing loss. The loss explains how well the model is performing by minimizing error. A lower number for loss correlates to a better performing model. Figure 2 shows the performance of the universal vector model by loss and training time over each epoch. The results of the model are shown between 40-60 Epochs to show where the loss curve flattens. The figure shows that the loss gradually declined between 41 to 48 epochs and began to stabilize at about .6 – .7.

At the 60th epoch, a loss of 0.6963 was obtained which was sufficient to translate short sentences. The model struggled with training performance with respect to time between 40-50 epochs. The exponential jump in time could potentially be due to the model struggling to get to the local minima point during optimization. Overall the loss obtained is sufficient to translate short sentences and shows that the universal vector model can translate words with minimal error.

8.3 Attention Model Analysis

Heat maps were created to visualize how the attention mechanism directed the focus of the decoder when predicting the corresponding text in a translation. The diagram has each word in the source language on the top and each word in the predicted sentence in the target language on the left axis. Fig. 3 below was generated when the Spanish phrase "Esta es mi vida" was fed into the model. On the left is the output of the model which is a prediction of the English Translation. As a visual reminder, the heat map does not necessarily show how words are correlated from source to target. Instead the visual representation

of the heat map gives insight into the parts of input that the attention model focuses on when translating. For example, the yellow box in the upper left shows heavy focus on the Spanish word "esta" when the model predicts the English word "this". From there, the heavy areas of focus follow a diagonal line down and to the left. This means as the decoder moves on to predict words later in the sentence, the focus is directed to later parts of the source sentence which is generally good. Longer sentences would show more defined and more varied areas of heat as they get more complicated. Overall, the maps generated from the small sentence sizes that the model can handle, show potential that the modified attention mechanism is working as intended.




Fig. 3. Heat map showing areas of focus from Spanish to English.

9 Limitations and Future Expansion

For further model experimentation on translation of more than two languages, a parallel dataset containing a triad of language phrases is required. While the architecture and model as part of this experiment is created to handle more than two languages, we only consider using a single model for two languages. As of today, most parallel datasets available are bilingual. In the future, a parallel dataset with three or more languages will be used to train and modify the current universal vector representation model. Furthermore, larger datasets will be used

with more training iterations akin to other papers in the NMT space. A more standardized test such as those provided by the annual Workshop on Machine Translation can be used on the model translated text.

10 Previously Considered Experiments

Connected Learning was the first attempt at a novel proposal. During the time of initial research, there were no other papers proposing the methods that made up this new idea. This method would allow the weights to learn source and target language as a Z format. First the model is trained in the direction of Source \rightarrow Target, then immediately trained again with the direction of Target \rightarrow Source, and finally the weights are retrained from Source \rightarrow Target.

In connected learning, training is done on the source sequence of vectors as $x = (x_1, \dots, x_{T_x})$ and target sequence of vectors as $y = (y_1, \dots, y_{T_x})$. For each of the sequence pairs of vectors, the source and target are swapped twice by utilizing the hidden output as the input when swapped. For example, if vector sequence x represents Spanish and vector sequence y is English, the model would first generate $h_{(t)} = f(h_{(t-1)}, x_t)$ and $c = q(h_1, \dots, h_{T_x})$ and use the context vector "c" when combining the hidden state of the recurrent network and provide vector sequence y as the source and generate values for x .

The belief was that the weights in the contextual information would have all the target information, however the model could not converge to a local optimum point where it was aligned to both source and target information.

11 Conclusion

In this paper, the idea of a "Universal Vector" is proposed as a new facet of NMT that can be used to translate between multiple languages in the same vector space. Models are usually built to translate in one direction. There exists some work that has been done in using both directions between a source and target language for reinforcement learning of training sets. However, the "Universal Vector" model is a singular model that can be trained in both directions (source to target and target to source) for more than one pair of languages.

The "Universal Vector" model detailed in this paper was built to test the proposition by modifying an RNN based Encoder-Decoder model. Existing attention mechanisms were also modified and used to create context vectors that increased performance in predicting the next translated text for overall target phrase translation. Multiple fully connected layers are added, one for each target language, to facilitate translations into multiple target languages.

The model is trained with parallel English and Spanish datasets. Phrases from both languages are trained from English to Spanish and Spanish to English within a recurrent network using Dual Training based methods. It was tested with many examples of both Spanish and English phrases. The attention mechanism was evaluated by viewing heat maps of where the model selectively focused on input text for its corresponding translated text.

While the results are promising, with more time and resources the experiment would provide better results. With more computing power the model can be trained using more words with more languages in a reasonable amount of time. In the future, better accepted benchmarks in translation such as those provided by the annual Workshop on Machine Translation can be used. While limited in scope, these results point to potential for greater accuracy on using a singular model for translating between multiple languages.

References

1. Nal Kalchbrenner and Phil Blunsom. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709, Seattle, Washington, USA, October 2013. Association for Computational Linguistics.
2. Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. On using very large target vocabulary for neural machine translation. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015.
3. Thang Luong, Ilya Sutskever, Quoc V. Le, Oriol Vinyals, and Wojciech Zaremba. Addressing the rare word problem in neural machine translation. *CoRR*, abs/1410.8206, 2014.
4. Philipp Koehn, Franz J. Och, and Daniel Marcu. Statistical phrase-based translation. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 127–133, 2003.
5. Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
6. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997.
7. David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. *Learning Representations by Back-Propagating Errors*, page 696–699. MIT Press, Cambridge, MA, USA, 1988.
8. H. Schwenk. Cslm - a modular open-source continuous space language modeling toolkit. *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, pages 1198–1202, 01 2013.
9. Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. Lstm neural networks for language modeling. 09 2012.
10. Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, 2014.
11. M. Schuster and K. K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.

12. Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, page III–1310–III–1318. JMLR.org, 2013.
13. Jean Pouget-Abadie, Dzmitry Bahdanau, Bart van Merriënboer, Kyunghyun Cho, and Yoshua Bengio. Overcoming the curse of sentence length for neural machine translation using automatic segmentation. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 78–85, Doha, Qatar, October 2014. Association for Computational Linguistics.
14. Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines. *CoRR*, abs/1410.5401, 2014.
15. Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014.
16. Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015.
17. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
18. Karl Moritz Hermann and Phil Blunsom. Multilingual Distributed Representations without Word Alignment. In *Proceedings of ICLR*, apr 2014.
19. Mikel L. Forcada and Ramón P. Neco. Recursive hetero-associative memories for translation. In *Proceedings of the International Work-Conference on Artificial and Natural Neural Networks: Biological and Artificial Computation: From Neuroscience to Technology*, IWANN '97, page 453–462, Berlin, Heidelberg, 1997. Springer-Verlag.
20. Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
21. F. A. Gers, J. Schmidhuber, and F. Cummins. Learning to forget: continual prediction with lstm. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*, volume 2, pages 850–855 vol.2, 1999.
22. Zachary Chase Lipton. A critical review of recurrent neural networks for sequence learning. *CoRR*, abs/1506.00019, 2015.
23. Hojjat Salehinejad, Julianne Baarbe, Sharan Sankar, Joseph Barfett, Errol Colak, and Shahrokh Valaee. Recent advances in recurrent neural networks. *CoRR*, abs/1801.01078, 2018.
24. Alex Sherstinsky. Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. *CoRR*, abs/1808.03314, 2018.
25. Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. How to construct deep recurrent neural networks. In *Proceedings of the Second International Conference on Learning Representations (ICLR 2014)*, 2014.
26. Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google's neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144, 2016.

27. Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. Dual learning for machine translation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 820–828, USA, 2016. Curran Associates Inc.
28. Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
29. Sepp Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. 04 1991.
30. Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. 2001.
31. Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR.
32. Charles Kelly. Tab-delimited Bilingual Sentence Pairs, March 2020.
33. Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics.
34. Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. Incorporating bert into neural machine translation, 2020.