

Neural Semantic Encoders

Tsendsuren Munkhdalai and Hong Yu

University of Massachusetts, MA, USA

tsendsuren.munkhdalai@umassmed.edu, hong.yu@umassmed.edu

Abstract

We present a memory augmented neural network for natural language understanding: Neural Semantic Encoders. NSE is equipped with a novel memory update rule and has a variable sized encoding memory that evolves over time and maintains the understanding of input sequences through *read*, *compose* and *write* operations. NSE can also access¹ multiple and shared memories. In this paper, we demonstrated the effectiveness and the flexibility of NSE on five different natural language tasks: natural language inference, question answering, sentence classification, document sentiment analysis and machine translation where NSE achieved state-of-the-art performance when evaluated on publically available benchmarks. For example, our shared-memory model showed an encouraging result on neural machine translation, improving an attention-based baseline by approximately 1.0 BLEU.

1 Introduction

Recurrent neural networks (RNNs) have been successful for modeling sequences (Elman, 1990). Particularly, RNNs equipped with internal short memories, such as long short-term memories (LSTM) (Hochreiter and Schmidhuber, 1997) have achieved a notable success in sequential tasks (Cho et al., 2014; Vinyals et al., 2015). LSTM is powerful because it learns to control its short term memories. However, the short term memories in LSTM are a part of the training parameters. This imposes some practical difficulties in training and modeling long sequences with LSTM.

¹By access we mean changing the memory states by the *read*, *compose* and *write* operations.

Recently several studies have explored ways of extending the neural networks with an external memory (Graves et al., 2014; Weston et al., 2015; Grefenstette et al., 2015). Unlike LSTM, the short term memories and the training parameters of such a neural network are no longer coupled and can be adapted. In this paper we propose a novel class of memory augmented neural networks called Neural Semantic Encoders (NSE) for natural language understanding. NSE offers several desirable properties. NSE has a variable sized encoding memory which allows the model to access entire input sequence during the reading process; therefore efficiently delivering long-term dependencies over time. The encoding memory evolves over time and maintains the memory of the input sequence through *read*, *compose* and *write* operations. NSE sequentially processes the input and supports word compositionality inheriting both temporal and hierarchical nature of human language. NSE can read from and write to a set of relevant encoding memories simultaneously or multiple NSEs can access a shared encoding memory effectively supporting knowledge and representation sharing. NSE is flexible, robust and suitable for practical NLU tasks and can be trained easily by any gradient descent optimizer.

We evaluate NSE on five different real tasks. For four of them, our models set new state-of-the-art results. Our results suggest that a NN model with the shared memory between encoder and decoder is a promising approach for sequence transduction problems such as machine translation and abstractive summarization. In particular, we observe that the attention-based neural machine translation can be further improved by shared-memory models. We also analyze memory access pattern and compositionality in NSE and show that our model captures semantic and syntactic structures of input sentence.




Figure 1: High-level architectures of the Neural Semantic Encoders. NSE reads and writes its own encoding memory in each time step (a). MMA-NSE accesses multiple relevant memories simultaneously (b).

2 Related Work

One of the pioneering work that attempts to extend deep neural networks with an external memory is Neural Turing Machines (NTM) (Graves et al., 2014). NTM implements a centralized controller and a fixed-sized random access memory. The NTM memory is addressable by both content (i.e. soft attention) and location based access mechanisms. The authors evaluated NTM on algorithmic tasks such as copying and sorting sequences.

Comparison with Neural Turing Machines:

NSE addresses certain drawbacks of NTM. NTM has a single centralized controller, which is usually an MLP or RNN while NSE takes a modular approach. The main controller in NSE is decomposed into three separate modules, each of which performs for *read*, *compose* or *write* operation. In NSE, the *compose* module is introduced in addition to the standard memory update operations (i.e. read-write) in order to process the memory entries and input information.

The main advantage of NSE over NTM is in its memory update. Despite its sophisticated addressing mechanism, the NTM controller does not have mechanism to avoid *information collision* in the memory. Particularly the NTM controller emits two separate set of access weights (i.e. read weight and erase and write weights) that do not explicitly encode the knowledge about where information is read from and written to. Moreover the fixed-size memory in NTM has no memory allocation or de-allocation protocol. Therefore unless the controller is intelligent enough to track the previous read/write information, which is hard for an RNN when processing long sequences, the memory content is overlapped and information is over-

written throughout different time scales. We think that this is a potential reason that makes NTM hard to train and makes the training not stable. We also note that the effectiveness of the location based addressing introduced in NTM is unclear. In NSE, we introduce a novel and systematic memory update approach based on the soft attention mechanism. NSE writes new information to the most recently read memory locations. This is accomplished by sharing the same memory key vector between the read and write modules. The NSE memory update is scalable and potentially more robust to train. NSE is provided with a variable sized memory and thus unlike NTM, the size of the NSE memory is more relaxed. The novel memory update mechanism and the variable sized memory together prevent NSE from the *information collision* issue and avoid the need of the memory allocation and de-allocation protocols. Each memory location of the NSE memory stores a token representation in input sequence during encoding. This provides NSE with an anytime-access to the entire input sequence including the tokens from the future time scales, which is not permitted in NTM, RNN and attention-based encoders.

Lastly, NTM addresses small algorithmic problems while NSE focuses on a set of large-scale language understanding tasks.

The RNNSearch model proposed in (Bahdanau et al., 2015) can be seen as a variation of memory augmented networks due to its ability to read the historic output states of RNNs with soft attention. The work of Sukhbaatar et al. (2015) combines the soft attention with Memory Networks (MemNNs) (Weston et al., 2015). Similar to RNNSearch, MemNNs are designed with non-writable memories. It constructs layered memory representa-

tions and showed promising results on both artificial and real question answering tasks. We note that RNNSearch and MemNNs avoid the memory update and management overhead by simply using a non-writable memory storage. Another variation of MemNNs is Dynamic Memory Network (Kumar et al., 2016) that is equipped with an episodic memory and seems to be flexible in different settings.

Although NSE differs from other memory-augmented NN models in many aspects, they all use soft attention mechanism with a type of similarity measures to retrieve relevant information from the external memory. For example, NTM implements cosine similarity and MemNNs use vector dot product. NSE uses the vector dot product for the similarity measure in NSE because it is faster to compute.

Other related work includes Neural Program-Interpreters (Reed and de Freitas, 2016), which learns to run sub-programs and to compose them for high-level programs. It uses execution traces to provide the full supervision. Researchers have also explored ways to add unbounded memory to LSTM (Grefenstette et al., 2015) using a particular data structure. Although this type of architecture provides a flexible capacity to store information, the memory access is constrained by the data structure used for the memory bank, such as stack and queue.

Overall it is expensive to train and to scale the previously proposed memory-based models. Most models required a set of clever engineering tricks to work successfully. Most of the aforementioned memory augmented neural networks have been tested on synthetic tasks whereas in this paper we evaluated NSE on a wide range of real and large-scale natural language applications.

3 Proposed Approach

Our training set consists of N examples $\{X^i, Y^i\}_{i=1}^N$, where the input X^i is a sequence $w_1^i, w_2^i, \dots, w_{T_i}^i$ of tokens while the output Y^i can be either a single target or a sequence. We transform each input token w_t to its word embedding x_t .

Our Neural Semantic Encoders (NSE) model has four main components: read, compose and write modules and an encoding memory $M \in R^{k \times l}$ with a variable number of slots, where k is the embedding dimension and l is the length

of the input sequence. Each memory slot vector $m_t \in R^k$ corresponds to the vector representation of information about word w_t in memory. In particular, the memory is initialized by the embedding vectors $\{x_t\}_{t=1}^l$ and is evolved over time, through *read*, *compose* and *write* operations. Figure 1 (a) illustrates the architecture of NSE.

3.1 Read, Compose and Write

NSE performs three main operations in every time step. After initializing the memory slots with the corresponding input representations, NSE processes an embedding vector x_t and retrieves a memory slot $m_{r,t}$ that is expected to be associatively coherent (i.e. semantically associated) with the current input word w_t .² The slot location r (ranging from 1 to l) is defined by a key vector z_t which the read module emits by attending over the memory slots. The compose module implements a composition operation that combines the memory slot with the current input. The write module then transforms the composition output to the encoding memory space and writes the resulting new representation into the slot location of the memory. Instead of composing the raw embedding vector x_t , we use the hidden state o_t produced by the read module at time t

Concretely, let $e_l \in R^l$ and $e_k \in R^k$ be vectors of ones and given a read function f_r^{LSTM} , a composition f_c^{MLP} and a write f_w^{LSTM} NSE in Figure 1 (a) computes the key vector z_t , the output state h_t , and the encoding memory M_t in time step t as

$$o_t = f_r^{LSTM}(x_t) \quad (1)$$

$$z_t = \text{softmax}(o_t^\top M_{t-1}) \quad (2)$$

$$m_{r,t} = z_t^\top M_{t-1} \quad (3)$$

$$c_t = f_c^{MLP}(o_t, m_{r,t}) \quad (4)$$

$$h_t = f_w^{LSTM}(c_t) \quad (5)$$

$$M_t = M_{t-1}(\mathbf{1} - (z_t \otimes e_k)^\top) + (h_t \otimes e_l)(z_t \otimes e_k)^\top \quad (6)$$

where $\mathbf{1}$ is a matrix of ones, \otimes denotes the outer product which duplicates its left vector l or k times to form a matrix. The read function f_r^{LSTM} sequentially maps the word embeddings to the internal space of the memory M_{t-1} . Then Equation 2 looks for the slots related to the input by computing association degree between each memory slot

²Such a coherence is calculated by a soft attention with dot product similarity.

and the hidden state o_t . We calculate the association degree by the dot product and transform this scores to the fuzzy key vector z_t by normalizing with *softmax* function. Since our key vector is fuzzy, the slot to be composed is retrieved by taking weighted sum of the all slots as in Equation 3. This process can also be seen as the soft attention mechanism (Bahdanau et al., 2015). In Equation 4 and 5, we compose and process the retrieved slot with the current hidden state and map the resulting vector to the encoder output space. Finally, we write the new representation to the memory location pointed by the key vector in Equation 6 where the key vector z_t emitted by the read module is reused to inform the write module of the most recently read slots. First the slot information that was retrieved is erased and then the new representation is located. NSE performs this iterative process until all words in the input sequence are read. The encoding memories $\{M\}_{t=1}^T$ and output states $\{h\}_{t=1}^T$ are further used for the tasks.

Although NSE reads a single word at a time, it has an anytime-access to the entire sequence stored in the encoding memory. With the encoding memory, NSE maintains a mental image of the input sequence. The memory is initialized with the raw embedding vector at time $t = 0$. We term such a freshly initialized memory a *baby* memory. As NSE reads more input content in time, the *baby* memory evolves and refines the encoded mental image.

The read f_r^{LSTM} , the composition f_c^{MLP} and the write f_w^{LSTM} functions are neural networks and are the training parameters in our NSE. As the name suggests, we use LSTM and multi-layer perceptron (MLP) in this paper. Since NSE is fully differentiable, it can be trained with any gradient descent optimizer.

3.2 Shared and Multiple Memory Accesses

For sequence to sequence transduction tasks like question answering, natural language inference and machine translation, it is beneficial to access other relevant memories in addition to its own one. The shared or the multiple memory access allows a set of NSEs to exchange knowledge representations and to communicate with each other to accomplish a particular task throughout the encoding memory.

NSE can be extended easily, so that it is able to read from and write to multiple memories si-

multaneously or multiple NSEs are able to access a shared memory. Figure 1 (b) depicts a high-level architectural diagram of a multiple memory access-NSE (MMA-NSE). The first memory (in green) is the shared memory accessed by more than one NSEs. Given a shared memory $M^n \in R^{k \times n}$ that has been encoded by processing a relevant sequence with length n , MMA-NSE with the access to one relevant memory is defined as

$$o_t = f_r^{LSTM}(x_t) \quad (7)$$

$$z_t = softmax(o_t^\top M_{t-1}) \quad (8)$$

$$m_{r,t} = z_t^\top M_{t-1} \quad (9)$$

$$z_t^n = softmax(o_t^\top M_{t-1}^n) \quad (10)$$

$$m_{r,t}^n = z_t^{n\top} M_{t-1}^n \quad (11)$$

$$c_t = f_c^{MLP}(o_t, m_{r,t}, m_{r,t}^n) \quad (12)$$

$$h_t = f_w^{LSTM}(c_t) \quad (13)$$

$$M_t = M_{t-1}(\mathbf{1} - (z_t \otimes e_k)^\top) + (h_t \otimes e_l)(z_t \otimes e_k)^\top \quad (14)$$

$$M_t^n = M_{t-1}^n(\mathbf{1} - (z_t^n \otimes e_k)^\top) + (h_t \otimes e_n)(z_t^n \otimes e_k)^\top \quad (15)$$

and this is almost the same as standard NSE. The read module now emits the additional key vector z_t^n for the shared memory and the composition function f_c^{MLP} combines more than one slots.

In MMA-NSE, the different memory slots are retrieved from the shared memories depending on their encoded semantic representations. They are then composed together with the current input and written back to their corresponding slots. Note that MMA-NSE is capable of accessing a variable number of relevant shared memories once a composition function that takes in dynamic inputs is chosen.

4 Experiments

We describe in this section experiments on five different tasks, in order to show that NSE can be effective and flexible in different settings.³ We report results on natural language inference, question answering (QA), sentence classification, document sentiment analysis and machine translation. All five tasks challenge a model in terms of language understanding and semantic reasoning.

The models are trained using Adam (Kingma and Ba, 2014) with hyperparameters selected on

³Code for the experiments and NSEs is available at <https://bitbucket.org/tsendeemts/nse>.

Model	d	$ \theta _M$	Train	Test
Classifier with handcrafted features (Bowman et al., 2015)	-	-	99.7	78.2
LSTM encoders (Bowman et al., 2015)	300	3.0M	83.9	80.6
Dependency Tree CNN encoders (Mou et al., 2016)	300	3.5M	83.3	82.1
SPINN-PI encoders (Bowman et al., 2016)	300	3.7M	89.2	83.2
NSE	300	3.4M	86.2	84.6
MMA-NSE	300	6.3M	87.1	84.8
LSTM attention (Rocktäschel et al., 2016)	100	242K	85.4	82.3
LSTM word-by-word attention (Rocktäschel et al., 2016)	100	252K	85.3	83.5
MMA-NSE attention	300	6.5M	86.9	85.4
mLSTM word-by-word attention (Wang and Jiang, 2015)	300	1.9M	92.0	86.1
LSTMN with deep attention fusion (Cheng et al., 2016)	450	3.4M	89.5	86.3
Decomposable attention model (Parikh et al., 2016)	200	582K	90.5	86.8
Full tree matching NTI-SLSTM-LSTM global attention (Munkhdalai and Yu, 2017)	300	3.2M	88.5	87.3

Table 1: Training and test accuracy on natural language inference task. d is the word embedding size and $|\theta|_M$ the number of model parameters.

development set. We chose two one-layer LSTM for read/write modules on the tasks other than QA on which we used two-layer LSTM. The pre-trained 300-D Glove 840B vectors and 100-D Glove 6B vectors (Pennington et al., 2014) were obtained for the word embeddings.⁴ The word embeddings are fixed during training. The embeddings for out-of-vocabulary words were set to zero vector. We crop or pad the input sequence to a fixed length. A padding vector was inserted when padding. The models were regularized by using dropouts and an l_2 weight decay.⁵

4.1 Natural Language Inference

The natural language inference is one of the main tasks in language understanding. This task tests the ability of a model to reason about the semantic relationship between two sentences. In order to perform well on the task, NSE should be able to capture sentence semantics and be able to reason the relation between a sentence pair, i.e., whether a premise-hypothesis pair is entailing, contradictory or neutral. We conducted experiments on the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015), which consists of 549,367/9,842/9,824 premise-hypothesis pairs for train/dev/test sets and target label indicating their relation.

Following the setting in (Mou et al., 2016; Bowman et al., 2016) the NSE output for each sentence was the input to a MLP, where the input layer computes the concatenation $[h_i^p; h_i^h]$, absolute difference $h_i^p - h_i^h$ and elementwise product $h_i^p \cdot h_i^h$ of the two sentence representations. In addition, the MLP has a hidden layer with 1024 units with

ReLU activation and a *softmax* layer. We set the batch size to 128, the initial learning rate to $3e-4$ and l_2 regularizer strength to $3e-5$, and train each model for 40 epochs. The write/read neural nets and the last linear layer were regularized by using 30% dropouts.

We evaluated three different variations of NSE show in Table 1. The NSE model encodes each sentence simultaneously by using a separate memory for each sentence. The second model - MMA-NSE first encodes the premise and then the hypothesis sentence by sharing the premise encoded memory in addition to the hypothesis memory. For the third model, we use inter-sentence attention which selectively reconstructs the premise representation.

Table 1 shows the results of our models along with the results of published methods for the task. The classifier with handcrafted features extracts a set of lexical features. The next group of models are based on sentence encoding. While most of the sentence encoder models rely solely on word embeddings, the dependency tree CNN and the SPINN-PI models make use of sentence parser output. The SPINN-PI model is similar to NSE in spirit that it also explicitly computes word composition. However, the composition in the SPINN-PI is guided by supervisions from a dependency parser. NSE outperformed the previous sentence encoders on this task. The MMA-SNE further slightly improved the result, indicating that reading the premise memory is helpful while encoding the hypothesis.

The last set of methods designs inter-sentence relation with parameterized soft attention (Bahdanau et al., 2015). Our MMA-NSE attention model is similar to the LSTM attention model.

⁴<http://nlp.stanford.edu/projects/glove/>

⁵More detail on hyper-parameters can be found in code.

Particularly, it attends over the premise encoder outputs $\{h^p\}_{t=1}^T$ in respect to the final hypothesis representation h_t^h and constructs an attentively blended vector of the premise. This model obtained 85.4% accuracy score. The best performing model for this task performs tree matching with attention mechanism and LSTM.

4.2 Answer Sentence Selection

Answer sentence selection is an integral part of the open-domain question answering. For this task, a model is trained to identify the correct sentences that answer a factual question, from a set of candidate sentences. We experiment on WikiQA dataset constructed from Wikipedia (Yang et al., 2015). The dataset contains 20,360/2,733/6,165 QA pairs for train/dev/test sets.

The MLP setup used in the language inference task is kept same, except that we now replace the *softmax* layer with a *sigmoid* layer and model the following conditional probability distribution.

$$p_{\theta}(y = 1|h_t^q, h_t^a) = \text{sigmoid}(o^{QA}) \quad (16)$$

where h_t^q and h_t^a are the question and the answer encoded vectors and o^{QA} denotes the output of the hidden layer of the MLP. We trained the MMA-NSE attention model to minimize the *sigmoid* cross entropy loss. MMA-NSE first encodes the answers and then the questions by accessing its own and the answer encoding memories. In our preliminary experiment, we found that the multiple memory access and the attention over answer encoder outputs $\{h^a\}_{t=1}^T$ are crucial to this problem. Following previous work, we adopt MAP and MRR as the evaluation metrics for this task.⁶

We set the batch size to 4 and the initial learning rate to 1e-5, and train the model for 10 epochs. We used 40% dropouts after word embeddings and no

⁶We used *trac_eval* script to calculate the evaluation metrics

Model	MAP	MRR
Classifier with features (2013)	0.5993	0.6068
Paragraph Vector (2014)	0.5110	0.5160
Bigram-CNN (2014)	0.6190	0.6281
3-layer LSTM (2016)	0.6552	0.6747
3-layer LSTM attention (2016)	0.6639	0.6828
NASM (2016)	0.6705	0.6914
MMA-NSE attention	0.6811	0.6993

Table 2: Experiment results on answer sentence selection.

l_2 weight decay. The word embeddings are pre-trained 300-D Glove 840B vectors. For this task, a linear mapping layer transforms the 300-D word embeddings to the 512-D LSTM inputs.

Table 2 presents the results of our model and the previous models for the task.⁷ The classifier with handcrafted features is a SVM model trained with a set of features. The Bigram-CNN model is a simple convolutional neural net. While the LSTM and LSTM attention models outperform the previous best result by nearly 5-6% by implementing deep LSTM with three hidden layers, NASM improves it further and sets a strong baseline by combining variational auto-encoder (Kingma and Welling, 2014) with the soft attention. Our MMA-NSE attention model exceeds the NASM by approximately 1% on MAP and 0.8% on MRR for this task.

4.3 Sentence Classification

We evaluated NSE on the Stanford Sentiment Treebank (SST) (Socher et al., 2013). This dataset comes with standard train/dev/test sets and two subtasks: binary sentence classification or fine-grained classification of five classes. We trained our model on the text spans corresponding to labeled phrases in the training set and evaluated the model on the full sentences.

The sentence representations were passed to a two-layer MLP for classification. The first layer of the MLP has *ReLU* activation and 1024 or 300 units for binary or fine-grained setting. The second layer is a *softmax* layer. The read/write modules are two one-layer LSTM with 300 hidden units and the word embeddings are the pre-trained 300-D Glove 840B vectors. We set the batch size to 64, the initial learning rate to 3e-4 and l_2 regularizer strength to 3e-5, and train each model for 25 epochs. The write/read neural nets and the last linear layer were regularized by 50% dropouts.

Table 3 compares the result of our model with the state-of-the-art methods on the two subtasks. Most best performing methods exploited the parse tree provided in the treebank on this task with the exception of the DMN. The Dynamic Memory Network (DMN) model is a memory-augmented network. Our model outperformed the DMN and set the state-of-the-art results on both subtasks.

⁷Inclusion of simple word count feature improves the performance by around 0.15-0.3 across the board

Model	Bin	FG
RNTN (Socher et al., 2013)	85.4	45.7
Paragraph Vector (Le and Mikolov, 2014)	87.8	48.7
CNN-MC (Kim, 2014)	88.1	47.4
DRNN (Irsoy and Cardie, 2015)	86.6	49.8
2-layer LSTM(Tai et al., 2015)	86.3	46.0
Bi-LSTM(Tai et al., 2015)	87.5	49.1
CT-LSTM(Tai et al., 2015)	88.0	51.0
DMN (Kumar et al., 2016)	88.6	52.1
NSE	89.7	52.8

Table 3: Test accuracy for sentence classification. Bin: Binary, FG: fine-grained 5 classes.

4.4 Document Sentiment Analysis

We evaluated our models for document-level sentiment analysis on two publically available large-scale datasets: the IMDB consisting of 335,018 movie reviews and 10 different classes and Yelp 13 consisting of 348,415 restaurant reviews and 5 different classes. Each document in the datasets is associated with human ratings and we used these ratings as gold labels for sentiment classification. Particularly, we used the pre-split datasets of (Tang et al., 2015).

We stack a NSE or LSTM on the top of another NSE for document modeling. The first NSE encodes the sentences and the second NSE or LSTM takes sentence encoded outputs and constructs document representations. The document representation is given to a output *softmax* layer. The whole network is trained jointly by backpropagating the cross entropy loss. We used one-layer LSTM with 100 hidden units for the read/write modules and the pre-trained 100-D Glove 6B vectors for this task. We set the batch size to 32, the initial learning rate to $3e-4$ and l_2 regularizer strength to $1e-5$, and trained each model for 50 epochs. The write/read neural nets and the document-level NSE/LSTM were regularized by 15% dropouts and the softmax layer by 20% dropouts. In order to speedup the training, we created document buckets by considering the number of sentences per document, i.e., documents with the same number of sentences were put together in the same bucket. The buckets were shuffled and updated per epoch. We did not use curriculum scheduling (Bengio et al., 2009), although it is observed to help sequence training.

Table 4 shows our results. We report two performance metrics: accuracy and MSE. The best results on the task were previously obtained by Conv-GRNN and LSTM-GRNN, which are also

Model	Yelp 13		IMDB	
	Acc	MSE	Acc	MSE
Classifier (2015)	59.8	0.68	40.5	3.56
PV (2015)	57.7	0.86	34.1	4.69
CNN (2015)	59.7	0.76	37.6	3.30
Conv-GRNN (2015)	63.7	0.56	42.5	2.71
LSTM-GRNN (2015)	65.1	0.50	45.3	3.00
NSE-NSE	66.6	0.48	48.3	1.94
NSE-LSTM	67.0	0.47	48.1	1.98

Table 4: Results of document-level sentiment classification. PV: paragraph vector, Acc: accuracy, and MSE: mean squared error.

Model	Train	Dev	Test
Baseline LSTM-LSTM	28.06	17.96	17.02
NSE-LSTM	28.73	17.67	17.13
NSE-NSE	29.89	18.53	17.93

Table 5: BLEU scores for English-German translation task.

stacked models. These models first learn the sentence representations with a CNN or LSTM and then combine them for document representation using a gated recurrent neural network (GRNN). Our NSE models outperformed the previous state-of-the-art models in terms of both accuracy and MSE, by approximately 2-3%. On the other hand, all systems tend to show poor results on the IMDB dataset. That is, the IMDB dataset contains longer documents than the Yelp 13 and it has 10 classes while the Yelp 13 dataset has five classes to distinguish.⁸ The stacked NSEs (NSE-NSE) performed slightly better than the NSE-LSTM on the IMDB dataset. This is possibly due to the encoding memory of the document level NSE that preserves the long dependency in documents with a large number of sentences.

4.5 Machine Translation

Lastly, we conducted an experiment on neural machine translation (NMT). The NMT problem is mostly defined within the encoder-decoder framework (Kalchbrenner and Blunsom, 2013; Cho et al., 2014; Sutskever et al., 2014). The encoder provides the semantic and syntactic information about the source sentences to the decoder and the decoder generates the target sentences by conditioning on this information and its partially produced translation. For an efficient encoding, the attention-based NTM was introduced (Bahdanau et al., 2015).

⁸The average number of sentences and words in a document for IMDB: 14, 152 and Yelp 13: 9, 326

For NTM, we implemented three different models. The first model is a baseline model and is similar to the one proposed in (Bahdanau et al., 2015) (RNNSearch). This model (LSTM-LSTM) has two LSTM for the encoder/decoder and has the soft attention neural net, which attends over the source sentence and constructs a focused encoding vector for each target word. The second model is an NSE-LSTM encoder-decoder which encodes the source sentence with NSE and generates the targets with the LSTM network by using the NSE output states and the attention network. The last model is an NSE-NSE setup, where the encoding part is the same as the NSE-LSTM while the decoder NSE now uses the output state and has an access to the encoder memory, i.e., the encoder and the decoder NSEs access a shared memory. The memory is encoded by the first NSEs and then read/written by the decoder NSEs. We used the English-German translation corpus from the IWSLT 2014 evaluation campaign (Cettolo et al., 2012). The corpus consists of sentence-aligned translation of TED talks. The data was pre-processed and lowercased with the Moses toolkit.⁹ We merged the dev2010 and dev2012 sets for development and the tst2010, tst2011 and tst2012 sets for test data¹⁰. Sentence pairs with length longer than 25 words were filtered out. This resulted in 110,439/4,998/4,793 pairs for train/dev/test sets. We kept the most frequent 25,000 words for the German dictionary. The English dictionary has 51,821 words. The 300-D Glove 840B vectors were used for embedding the words in the source sentence whereas a lookup embedding layer was used for the target German words. Note that the word embeddings are usually optimized along with the NMT models. However, for the evaluation purpose we in this experiment do not optimize the English word embeddings. Besides, we do not use a beam search to generate the target sentences.

The LSTM encoder/decoders have two layers with 300 units. The NSE read/write modules are two one-layer LSTM with the same number of units as the LSTM encoder/decoders. This ensures that the number of parameters of the models is roughly the equal. The models were trained to minimize word-level cross entropy loss and were regularized by 20% input dropouts and the

⁹<https://github.com/moses-smt/mosesdecoder>

¹⁰We modified `prepareData.sh` script: <https://github.com/facebookresearch/MIXER>

30% output dropouts. We set the batch size to 128, the initial learning rate to 1e-3 for LSTM-LSTM and 3e-4 for the other models and l_2 regularizer strength to 3e-5, and train each model for 40 epochs. We report BLEU score for each models.¹¹

Table 5 reports our results. The baseline LSTM-LSTM encoder-decoder (with attention) obtained 17.02 BLEU on the test set. The NSE-LSTM improved the baseline slightly. Given this very small improvement of the NSE-LSTM, it is unclear whether the NSE encoder is helpful in NMT. However, if we replace the LSTM decoder with another NSE and introduce the shared memory access to the encoder-decoder model (NSE-NSE), we improve the baseline result by almost 1.0 BLEU. The NSE-NSE model also yields an increasing BLEU score on dev set. The result demonstrates that the attention-based NMT systems can be improved by a shared-memory encoder-decoder model. In addition, memory-based NMT systems should perform well on translation of long sequences by preserving long term dependencies.

5 Qualitative Analysis

5.1 Memory Access and Compositionality

NSE is capable of performing multiscale composition by retrieving associative slots for a particular input at a time step. We analyzed the memory access order and the compositionality of memory slot and the input word in the NSE model trained on the SNLI data.

Figure 2 shows the word association graphs for the two sentence picked from SNLI test set. The association graph was constructed by inspecting the key vector z . For an input word, we connect it to the most active slot pointed by z ¹².

Note the graph components clustered around the semantically rich words: "sits", "wall" and "autumn" (a) and "Three", "puppies", "tub" and "vet" (b). The memory slots corresponding to words that are semantically rich in the current context are the most frequently accessed. The graph is able to capture certain syntactic structures including phrases (e.g., "hand built rock wall") and modifier relations (between "sits" and "quietly" and

¹¹We computed the BLEU score with `multi-bleu.perl` script of the Moses toolkit

¹²Since z is fuzzy, we visualize the highest scoring slot. For a few inputs, z pointed to a slot corresponding to the same word. In this case, we masked out those slots and showed the second best scoring slot.




Figure 2: Word association or composition graphs produced by NSE memory access. The directed arcs connect the words that are composed via *compose* module. The source nodes are input words and the destination nodes (pointed by the arrows) correspond to the accessed memory slots. $\langle S \rangle$ denotes the beginning of sequence.

between "tub" and "sprayed with water"). Another interesting property is that the model tends to perform sensible compositions while processing the input sentence. For example, NSE retrieved the memory slot corresponding to "wall" or "Three" when reading the input "rock" or "are".

In Appendix A, we show a step-by-step visualization of NSE memory states for the first sentence. Note how the encoding memory is evolved over time. In time step four ($t = 4$), the memory slot for "quietly" encodes information about "quiet(ly) little child". When $t = 6$, the model forms another composition involving "quietly", "quietly sits". In the last time step, we are able to find the most or the least frequently accessed slots in the memory. The least accessed slots correspond to function words while the frequently accessed slots are content words and tend to carry out rich semantics and intrinsic compositions found in the input sentence. Overall the model is less constrained and is able to compose multiword expressions.

6 Conclusion

Our proposed memory augmented neural networks have achieved the state-of-the-art results when evaluated on five representative NLP tasks. NSE is capable of building an efficient architecture of the single, shared and multiple memory accesses for a specific NLP task. For example, for the NLI task NSE accesses premise encoded memory when processing hypothesis. For the QA task, NSE accesses answer encoded memory when reading question for QA. In machine translation, NSE shares a single encoded memory between encoder and decoder. Such flexibility in the architectural choice of the NSE memory access allows for

the robust models for a better performance.

The initial state of the NSE memory stores information about each word in the input sequence. We in this paper used word embeddings to represent the words in the memory. Different variations of word representations such as character-based models are left to be evaluated for memory initialization in the future. We plan to extend NSE so that it learns to select and access a relevant subset from a memory set. One could also explore unsupervised variations of NSE, for example, to train them to produce encoding memory and representation vector of entire sentences or documents using either new or existing models such as the skip-gram model (Mikolov et al., 2013).

Acknowledgments

We would like to thank Abhyuday Jagannatha and the anonymous reviewers for their insightful comments and suggestions. This work was supported in part by the grant HL125089 from the National Institutes of Health (NIH). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.

- Samuel R. Bowman, Jon Gauthier, Abhinav Rashtogi, Raghav Gupta, Christopher D. Manning, and Christopher Potts. 2016. A fast unified model for parsing and sentence understanding. *CoRR*, abs/1603.06021.
- Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. Wit³: Web inventory of transcribed and translated talks. In *EAMT*, Trento, Italy.
- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. *CoRR*, abs/1601.06733.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Jeffrey L. Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural turing machines. *arXiv preprint arXiv:1410.5401*.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. Learning to transduce with unbounded memory. In *NIPS 2015*, pages 1819–1827.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ozan Irsoy and Claire Cardie. 2015. Modeling compositionality with multiplicative recurrent neural networks. In *ICLR 2015*.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *EMNLP*, volume 3, page 413.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *ICLR*.
- Diederik P. Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR 2014*.
- Ankit Kumar, Ozan Irsoy, Jonathan Su, James Bradbury, Robert English, Brian Pierce, Peter Ondruska, Ishaan Gulrajani, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. *CoRR*, abs/1506.07285.
- Quoc V. Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *ICML 2014*.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural variational inference for text processing. In *ICLR 2016*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. 2016. Recognizing entailment and contradiction by tree-based convolution. In *ACL 2016*.
- Tsendsuren Munkhdalai and Hong Yu. 2017. Neural tree indexers for text understanding.
- Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. *arXiv preprint arXiv:1606.01933*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543.
- Scott Reed and Nando de Freitas. 2016. Neural programmer-interpreters. In *ICLR 2016*.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. In *ICLR 2016*.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP 2013*.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *NIPS 2015*, pages 2431–2439.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *ACL 2015*.
- Duyu Tang, Bing Qin, and Ting Liu. 2015. Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP*.
- Wen tau Yih, Ming-Wei Chang, Christopher Meek, and Andrzej Pastusiak. 2013. Question answering using enhanced lexical semantic models. In *ACL 2013*.
- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. Grammar as a foreign language. In *NIPS*.
- Shuohang Wang and Jing Jiang. 2015. Learning natural language inference with LSTM. *CoRR*, abs/1512.08849.

Jason Weston, Sumit Chopra, and Antoine Bordes. 2015. Memory networks. In *ICML 2015*.

Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *EMNLP 2015*.

Lei Yu, Karl Moritz Hermann, Phil Blunsom, and Stephen Pulman. 2014. Deep learning for answer sentence selection. In *NIPS Deep Learning Workshop 2014*.

A Step-by-step visualization of memory states in NSE

Each small table represents the memory state at a single time step. The current time step and input token are listed on the top of the table. The memory slots pointed by the query vector is highlighted in red color. The brackets represent the word composition order in each slot.

<i>t=0</i> <i>input:</i> <S> <S> A little child sits quietly on a hand built rock wall in autumn	<i>t=1</i> <i>input:</i> <S> <S> (<S>A) little child sits quietly on a hand built rock wall in autumn	<i>t=2</i> <i>input:</i> A (A <S>) (<S>A) little child sits quietly on a hand built rock wall in autumn	<i>t=3</i> <i>input:</i> little (A <S>) (<S>A) little child sits (little quietly) on a hand built rock wall in autumn	<i>t=4</i> <i>input:</i> child (A <S>) (<S>A) little child sits (child (little quietly)) on a hand built rock wall in autumn	<i>t=5</i> <i>input:</i> sits (A <S>) (<S>A) little child sits (child (little quietly)) on a hand built rock wall in (sits autumn)
<i>t=6</i> <i>input:</i> quietly (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand built rock wall in (sits autumn)	<i>t=7</i> <i>input:</i> on (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand built rock wall in (on (sits autumn))	<i>t=8</i> <i>input:</i> a (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand (a built) rock wall in (on (sits autumn))	<i>t=9</i> <i>input:</i> hand (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand (a built) rock (hand wall) in (on (sits autumn))	<i>t=10</i> <i>input:</i> built (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand (a built) rock (built (hand wall)) in (on (sits autumn))	<i>t=11</i> <i>input:</i> rock (A <S>) (<S>A) little child (quietly sits) (child (little quietly)) on a hand (a built) rock (rock (built (hand wall))) in (on (sits autumn))
<i>t=12</i> <i>input:</i> wall (A <S>) (<S>A) little child (wall (quietly sits)) (child (little quietly)) on a hand (a built) rock (rock (built (hand wall))) in (on (sits autumn))	<i>t=13</i> <i>input:</i> in (A <S>) (<S>A) little child (wall (quietly sits)) (child (little quietly)) on a hand (a built) (in rock) (rock (built (hand wall))) in (on (sits autumn))	<i>t=14</i> <i>input:</i> autumn (A <S>) (<S>A) little child (wall (quietly sits)) (child (little quietly)) on a hand (autumn (a built)) (in rock) (rock (built (hand wall))) in (on (sits autumn))			