

WIKIREADING: A Novel Large-scale Language Understanding Task over Wikipedia

Daniel Hewlett, Alexandre Lacoste, Llion Jones, Illia Polosukhin,
Andrew Fandrianto, Jay Han, Matthew Kelcey and David Berthelot

Google Research

{dhewlett, allac, llion, ipolosukhin, fto, hanjay, matkelcey, dberth}@google.com

Abstract

We present WIKIREADING, a large-scale natural language understanding task and publicly-available dataset with 18 million instances. The task is to predict textual values from the structured knowledge base Wikidata by reading the text of the corresponding Wikipedia articles. The task contains a rich variety of challenging classification and extraction sub-tasks, making it well-suited for end-to-end models such as deep neural networks (DNNs). We compare various state-of-the-art DNN-based architectures for document classification, information extraction, and question answering. We find that models supporting a rich answer space, such as word or character sequences, perform best. Our best-performing model, a word-level sequence to sequence model with a mechanism to copy out-of-vocabulary words, obtains an accuracy of 71.8%.

1 Introduction

A growing amount of research in natural language understanding (NLU) explores end-to-end deep neural network (DNN) architectures for tasks such as text classification (Zhang et al., 2015), relation extraction (Nguyen and Grishman, 2015), and question answering (Weston et al., 2015). These models offer the potential to remove the intermediate steps traditionally involved in processing natural language data by operating on increasingly raw forms of text input, even unprocessed character or byte sequences. Furthermore, while these tasks are often studied in isolation, DNNs have the potential to combine multiple forms of reasoning within a single model.

Supervised training of DNNs often requires a

large amount of high-quality training data. To this end, we introduce a novel prediction task and accompanying large-scale dataset with a range of sub-tasks combining text classification and information extraction. The dataset is made publicly-available at <http://goo.gl/wikireading>. The task, which we call WIKIREADING, is to predict textual values from the open knowledge base *Wikidata* (Vrandečić and Krötzsch, 2014) given text from the corresponding articles on *Wikipedia* (Ayers et al., 2008). Example instances are shown in Table 1, illustrating the variety of subject matter and sub-tasks. The dataset contains 18.58M instances across 884 sub-tasks, split roughly evenly between classification and extraction (see Section 2 for more details).

In addition to its diversity, the WIKIREADING dataset is also at least an order of magnitude larger than related NLU datasets. Many natural language datasets for question answering (QA), such as WIKIQA (Yang et al., 2015), have only thousands of examples and are thus too small for training end-to-end models. Hermann et al. (2015) proposed a task similar to QA, predicting entities in news summaries from the text of the original news articles, and generated a NEWS dataset with 1M instances. The bAbI dataset (Weston et al., 2015) requires multiple forms of reasoning, but is composed of synthetically generated documents. WIKIQA and NEWS only involve pointing to locations within the document, and text classification datasets often have small numbers of output classes. In contrast, WIKIREADING has a rich output space of millions of answers, making it a challenging benchmark for state-of-the-art DNN architectures for QA or text classification.

We implemented a large suite of recent models, and for the first time evaluate them on common grounds, placing the complexity of the task in context and illustrating the tradeoffs inherent in each

	Categorization		Extraction	
Document	Folkart Towers are twin skyscrapers in the Bayrakli district of the Turkish city of Izmir. Reaching a structural height of 200 m (656 ft) above ground level, they are the tallest ...	Angeles blancos is a Mexican telenovela produced by Carlos Sotomayor for Televisa in 1990. Jacqueline Andere, Rogelio Guerra and Alfonso Iturralde star as the main ...	Canada is a country in the northern part of North America. Its ten provinces and three territories extend from the Atlantic to the Pacific and northward into the Arctic Ocean , ...	Breaking Bad is an American crime drama television series created and produced by Vince Gilligan. The show originally aired on the AMC network for five seasons, from January 20, 2008 , to ...
Property	country	original language of work	located next to body of water	start time
Answer	Turkey	Spanish	Atlantic Ocean, Arctic Ocean, Pacific Ocean	20 January 2008

Table 1: Examples instances from WIKIREADING. The task is to predict the answer given the document and property. Answer tokens that can be extracted are shown in bold, the remaining instances require classification or another form of inference.

approach. The highest score of 71.8% is achieved by a sequence to sequence model (Kalchbrenner and Blunsom, 2013; Cho et al., 2014) operating on word-level input and output sequences, with special handling for out-of-vocabulary words.

2 WIKIREADING

We now provide background information relating to Wikidata, followed by a detailed description of the WIKIREADING prediction task and dataset.

2.1 Wikidata

Wikidata is a free collaborative knowledge base containing information about approximately 16M *items* (Vrandečić and Krötzsch, 2014). Knowledge related to each item is expressed in a set of *statements*, each consisting of a (property, value) tuple. For example, the item `Paris` might have associated statements asserting (instance of, city) or (country, France). Wikidata contains over 80M such statements across 884 properties. Items may be linked to articles on Wikipedia.

2.2 Dataset

We constructed the WIKIREADING dataset from Wikidata and Wikipedia as follows: We consolidated all Wikidata statements with the same item and property into a single (item, property, answer) triple, where answer is a set of values. Replacing each item with the text of the linked Wikipedia article (discarding unlinked items) yields a dataset of 18.58M (document, property, answer) instances. Importantly, all elements in each instance are human-readable strings, making the task entirely textual. The only modification we made to these strings was to

convert timestamps into a human-readable format (e.g., “4 July 1776”).

The WIKIREADING task, then, is to predict the answer string for each tuple given the document and property strings. This setup can be seen as similar to information extraction, or question answering where the property acts as a “question”. We assigned all instances for each document randomly to either training (12.97M instances), validation (1.88M), and test (3.73M) sets following a 70/10/20 distribution. This ensures that, during validation and testing, all documents are unseen.

2.3 Documents

The dataset contains 4.7M unique Wikipedia articles, meaning that roughly 80% of the English-language Wikipedia is represented. Multiple instances can share the same document, with a mean of 5.31 instances per article (median: 4, max: 879). The most common categories of documents are human, taxon, film, album, and human settlement, making up 48.8% of the documents and 9.1% of the instances. The mean and median document lengths are 489.2 and 203 words.

2.4 Properties

The dataset contains 884 unique properties, though the distribution of properties across instances is highly skewed: The top 20 properties cover 75% of the dataset, with 99% coverage achieved after 180 properties. We divide the properties broadly into two groups: *Categorical* properties, such as instance of, gender and country, require selecting between a relatively small number of possible answers, while *relational* properties, such as date of birth,

Property	Frequency	Entropy
instance of	2,574,038	0.431
sex or gender	941,200	0.189
country	803,252	0.536
date of birth	785,049	0.936
given name	767,916	0.763
occupation	716,176	0.589
country of citizenship	674,560	0.501
located in ... entity	478,372	0.802
place of birth	384,951	0.800
date of death	364,910	0.943

Table 2: Training set frequency and scaled answer entropy for the 10 most frequent properties.

parent, and capital, typically require extracting rare or totally unique answers from the document.

To quantify this difference, we compute the entropy of the answer distribution A for each property p , scaled to the $[0, 1]$ range by dividing by the entropy of a uniform distribution with the same number of values, i.e., $\hat{H}(p) = H(A_p) / \log |A_p|$. Properties that represent essentially one-to-one mappings score near 1.0, while a property with just a single answer would score 0.0. Table 2 lists entropy values for a subset of properties, showing that the dataset contains a spectrum of sub-tasks. We label properties with an entropy less than 0.7 as categorical, and those with a higher entropy as relational. Categorical properties cover 56.7% of the instances in the dataset, with the remaining 43.3% being relational.

2.5 Answers

The distribution of properties described above has implications for the answer distribution. There are a relatively small number of very high frequency “head” answers, mostly for categorical properties, and a vast number of very low frequency “tail” answers, such as names and dates. At the extremes, the most frequent answer `human` accounts for almost 7% of the dataset, while 54.7% of the answers in the dataset are unique. There are some special categories of answers which are systematically related, in particular *dates*, which comprise 8.9% of the dataset (with 7.2% being unique). This distribution means that methods focused on either head or tail answers can each perform moderately well, but only a method that handles both types of answers can achieve maximum performance. Another consequence of the long tail of answers is that many (30.0%) of the answers in the test set never appear in the training set, meaning they must be read out of the document. An answer

is present verbatim in the document for 45.6% of the instances.

3 Methods

Recently, neural network architectures for NLU have been shown to meet or exceed the performance of traditional methods (Zhang et al., 2015; Dai and Le, 2015). The move to deep neural networks also allows for new ways of combining the property and document, inspired by recent research in the field of question answering (with the property serving as a question). In sequential models such as Recurrent Neural Networks (RNNs), the question could be prepended to the document, allowing the model to “read” the document differently for each question (Hermann et al., 2015). Alternatively, the question could be used to compute a form of *attention* (Bahdanau et al., 2014) over the document, to effectively focus the model on the most predictive words or phrases (Sukhbaatar et al., 2015; Hermann et al., 2015). As this is currently an ongoing field of research, we implemented a range of recent models and for the first time compare them on common grounds. We now describe these methods, grouping them into broad categories by general approach and noting necessary modifications. Later, we introduce some novel variations of these models.

3.1 Answer Classification

Perhaps the most straightforward approach to WIKIREADING is to consider it as a special case of document classification. To fit WIKIREADING into this framework, we consider each possible answer as a class label, and incorporate features based on the property so that the model can make different predictions for the same document. While the number of potential answers is too large to be practical (and unbounded in principle), a substantial portion of the dataset can be covered by a model with a tractable number of answers.

3.1.1 Baseline

The most common approach to document classification is to fit a linear model (e.g., Logistic Regression) over bag of words (BoW) features. To serve as a baseline for our task, the linear model needs to make different predictions for the same Wikipedia article depending on the property. We enable this behavior by computing two N_w element BoW vectors, one each for the document

and property, and concatenating them into a single $2N_w$ feature vector.

3.1.2 Neural Network Methods

All of the methods described in this section encode the property and document into a *joint representation* $\mathbf{y} \in \mathbb{R}^{d_{\text{out}}}$, which serves as input for a final softmax layer computing a probability distribution over the top N_{ans} answers. Namely, for each answer $i \in \{1, \dots, N_{\text{ans}}\}$, we have:

$$P(i|\mathbf{x}) = e^{\mathbf{y}^\top \mathbf{a}_i} / \sum_{j=1}^{N_{\text{ans}}} e^{\mathbf{y}^\top \mathbf{a}_j}, \quad (1)$$

where $\mathbf{a}_i \in \mathbb{R}^{d_{\text{out}}}$ corresponds to a learned vector associated with answer i . Thus, these models differ primarily in how they combine the property and document to produce the joint representation. For existing models from the literature, we provide a brief description and note any important differences in our implementation, but refer the reader to the original papers for further details.

Except for character-level models, documents and properties are tokenized into words. The N_w most frequent words are mapped to a vector in $\mathbb{R}^{d_{\text{in}}}$ using a learned embedding matrix¹. Other words are all mapped to a special out of vocabulary (OOV) token, which also has a learned embedding. d_{in} and d_{out} are hyperparameters for these models.

Averaged Embeddings (BoW): This is the neural network version of the baseline method described in Section 3.1.1. Embeddings for words in the document and property are separately averaged. The concatenation of the resulting vectors forms the joint representation of size $2d_{\text{in}}$.

Paragraph Vector: We explore a variant of the previous model where the document is encoded as a *paragraph vector* (Le and Mikolov, 2014). We apply the PV-DBOW variant that learns an embedding for a document by optimizing the prediction of its constituent words. These unsupervised document embeddings are treated as a fixed input to the supervised classifier, with no fine-tuning.

LSTM Reader: This model is a simplified version of the Deep LSTM Reader proposed by Hermann et al. (2015). In this model, an LSTM (Hochreiter and Schmidhuber, 1997) reads the property and document sequences word-by-word

¹Limited experimentation with initialization from publicly-available word2vec embeddings (Mikolov et al., 2013) yielded no improvement in performance.

and the final state is used as the joint representation. This is the simplest model that respects the order of the words in the document. In our implementation we use a single layer instead of two and a larger hidden size. More details on the architecture can be found in Section 4.1 and in Table 4.

Attentive Reader: This model, also presented in Hermann et al. (2015), uses an attention mechanism to better focus on the relevant part of the document for a given property. Specifically, Attentive Reader first generates a representation \mathbf{u} of the property using the final state of an LSTM while a second LSTM is used to read the document and generate a representation \mathbf{z}_t for each word. Then, conditioned on the property encoding \mathbf{u} , a normalized attention is computed over the document to produce a weighted average of the word representations \mathbf{z}_t , which is then used to generate the joint representation \mathbf{y} . More precisely:

$$\begin{aligned} \mathbf{m}_t &= \tanh(W_1 \text{concat}(\mathbf{z}_t, \mathbf{u})) \\ \alpha_t &= \exp(\mathbf{v}^\top \mathbf{m}_t) \\ \mathbf{r} &= \sum_t \frac{\alpha_t}{\sum_\tau \alpha_\tau} \mathbf{z}_t \\ \mathbf{y} &= \tanh(W_2 \text{concat}(\mathbf{r}, \mathbf{u})), \end{aligned}$$

where W_1 , W_2 , and \mathbf{v} are learned parameters.

Memory Network: Our implementation closely follows the End-to-End Memory Network proposed in Sukhbaatar et al. (2015). This model maps a property p and a list of sentences $\mathbf{x}_1, \dots, \mathbf{x}_n$ to a joint representation \mathbf{y} by attending over sentences in the document as follows: The input encoder I converts a sequence of words $\mathbf{x}_i = (x_{i1}, \dots, x_{iL_i})$ into a vector using an embedding matrix (equation 2), where L_i is the length of sentence i .² The property is encoded with the embedding matrix U (eqn. 3). Each sentence is encoded into two vectors, a memory vector (eqn. 4) and an output vector (eqn. 5), with embedding matrices M and C , respectively. The property encoding is used to compute a normalized attention vector over the memories (eqn. 6).³ The joint representation is the sum of the output vectors weighted

²Our final results use the position encoding method proposed by Sukhbaatar et al. (2015), which incorporates positional information in addition to word embeddings.

³Instead of the linearization method of Sukhbaatar et al. (2015), we applied an entropy regularizer for the softmax attention as described in Kurach et al. (2015).

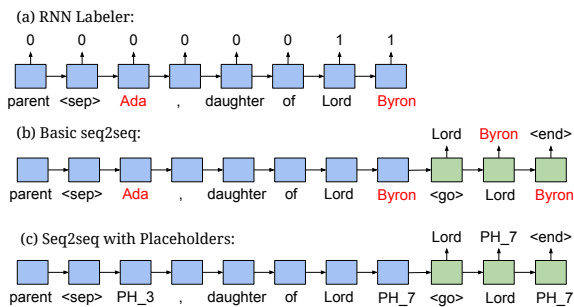


Figure 1: Illustration of RNN models. Blocks with same color share parameters. Red words are out of vocabulary and all share a common embedding.

by this attention (eqn. 7).

$$I(x_i, W) = \sum_j W x_{ij} \quad (2)$$

$$\mathbf{u} = I(p, U) \quad (3)$$

$$\mathbf{m}_i = I(x_i, M) \quad (4)$$

$$\mathbf{c}_i = I(x_i, C) \quad (5)$$

$$\mathbf{p}_i = \text{softmax}(\mathbf{q}^\top \mathbf{m}_i) \quad (6)$$

$$\mathbf{y} = \mathbf{u} + \sum_i \mathbf{p}_i \mathbf{c}_i \quad (7)$$

3.2 Answer Extraction

Relational properties involve mappings between arbitrary entities (e.g., `date of birth`, `mother`, and `author`) and thus are less amenable to document classification. For these, approaches from information extraction (especially relation extraction) are much more appropriate. In general, these methods seek to identify a word or phrase in the text that stands in a particular relation to a (possibly implicit) subject. Section 5 contains a discussion of prior work applying NLP techniques involving entity recognition and syntactic parsing to this problem.

RNNs provide a natural fit for extraction, as they can predict a value at every position in a sequence, conditioned on the entire previous sequence. The most straightforward application to WIKIREADING is to predict the probability that a word at a given location is part of an answer. We test this approach using an RNN that operates on the sequence of words. At each time step, we use a sigmoid activation for estimating whether the current word is part of the answer or not. We refer to this model as the *RNN Labeler* and present it graphically in Figure 1a.

For training, we label all locations where any answer appears in the document with a 1, and other positions with a 0 (similar to distant supervision (Mintz et al., 2009)). For multi-word an-

swers, the word sequences in the document and answer must fully match⁴. Instances where no answer appears in the document are discarded for training. The cost function is the average cross-entropy for the outputs across the sequence. When performing inference on the test set, sequences of consecutive locations scoring above a threshold are chunked together as a single answer, and the top-scoring answer is recorded for submission.⁵

3.3 Sequence to Sequence

Recently, sequence to sequence learning (or *seq2seq*) has shown promise for natural language tasks, especially machine translation (Cho et al., 2014). These models combine two RNNs: an *encoder*, which transforms the input sequence into a vector representation, and a *decoder*, which converts the encoder vector into a sequence of output tokens, one token at a time. This makes them capable, in principle, of approximating any function mapping sequential inputs to sequential outputs. Importantly, they are the first model we consider that can perform any combination of answer classification and extraction.

3.3.1 Basic seq2seq

This model resembles LSTM Reader augmented with a second RNN to decode the answer as a sequence of words. The embedding matrix is shared across the two RNNs but their state to state transition matrices are different (Figure 1b). This method extends the set of possible answers to any sequence of words from the document vocabulary.

3.3.2 Placeholder seq2seq

While Basic seq2seq already expands the expressiveness of LSTM Reader, it still has a limited vocabulary and thus is unable to generate some answers. As mentioned in Section 3.2, RNN Labeler can extract any sequence of words present in the document, even if some are OOV. We extend the basic seq2seq model to handle OOV words by adding *placeholders* to our vocabulary, increasing the vocabulary size from N_w to $N_w + N_{doc}$. Then, when an OOV word occurs in the document, it is replaced at random (without replacement) by one of these placeholders. We also replace the corresponding OOV words in the target output se-

⁴Dates were matched semantically to increase recall.

⁵We chose an arbitrary threshold of 0.5 for chunking. The score of each chunk is obtained from the harmonic mean of the predicted probabilities of its elements.

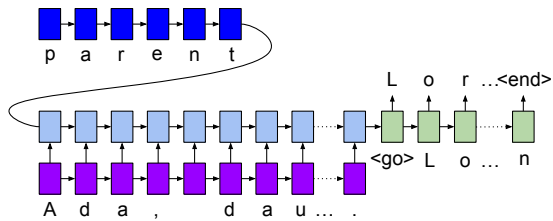


Figure 2: Character seq2seq model. Blocks with the same color share parameters. The same example as in Figure 1 is fed character by character.

quence by the same placeholder,⁶ as shown in Figure 1c. Luong et al. (2015) developed a similar procedure for dealing with rare words in machine translation, copying their locations into the output sequence for further processing.

This makes the input and output sequences a mixture of known words and placeholders, and allows the model to produce any answer the RNN Labeler can produce, in addition to the ones that the basic seq2seq model could already produce. This approach is comparable to *entity anonymization* used in Hermann et al. (2015), which replaces named entities with random ids, but simpler because we use word-level placeholders without entity recognition.

3.3.3 Basic Character seq2seq

Another way of handling rare words is to process the input and output text as sequences of characters or bytes. RNNs have shown some promise working with character-level input, including state-of-the-art performance on a Wikipedia text classification benchmark (Dai and Le, 2015). A model that outputs answers character by character can in principle generate any of the answers in the test set, a major advantage for WIKIREADING.

This model, shown in Figure 2, operates only on sequences of mixed-case characters. The property encoder RNN transforms the property, as a character sequence, into a fixed-length vector. This property encoding becomes the initial hidden state for the second layer of a two-layer document encoder RNN, which reads the document, again, character by character. Finally, the answer decoder RNN uses the final state of the previous RNN to decode the character sequence for the answer.

⁶The same OOV word may occur several times in the document. Our simplified approach will attribute a different placeholder for each of these and will use the first occurrence for the target answer.

3.3.4 Character seq2seq with Pretraining

Unfortunately, at the character level the length of all sequences (documents, properties, and answers) is greatly increased. This adds more sequential steps to the RNN, requiring gradients to propagate further, and increasing the chance of an error during decoding. To address this issue in a classification context, Dai and Le (2015) showed that initializing an LSTM classifier with weights from a language model (LM) improved its accuracy. Inspired by this result, we apply this principle to the character seq2seq model with a two-phase training process: In the first phase, we train a character-level LM on the input character sequences from the WIKIREADING training set (no new data is introduced). In the second phase, the weights from this LM are used to initialize the first layer of the encoder and the decoder (purple and green blocks in Figure 2). After initialization, training proceeds as in the basic character seq2seq model.

4 Experiments

We evaluated all methods from Section 3 on the full test set with a single scoring framework. An answer is correct when there is an exact string match between the predicted answer and the gold answer. However, as describe in Section 2.2, some answers are composed from a set of values (e.g. third example in Table 1). To handle this, we define the *Mean F1* score as follows: For each instance, we compute the F1-score (harmonic mean of precision and recall) as a measure of the degree of overlap between the predicted answer set and the gold set for a given instance. The resulting per-instance F1 scores are then averaged to produce a single dataset-level score. This allows a method to obtain partial credit for an instance when it answers with at least one value from the golden set. In this paper, we only consider methods for answering with a single value, and most answers in the dataset are also composed of a single value, so this Mean F1 metric is closely related to accuracy. More precisely, a method using a single value as answer is bounded by a Mean F1 of 0.963.

4.1 Training Details

We implemented all models in a single framework based on TensorFlow (Abadi et al., 2015) with shared pre-processing and comparable hyper-parameters whenever possible. All documents are

Method	Mean F1	Bound	Categorical	Relational	Date	Params
Answer Classifier						
Sparse BoW Baseline	0.438	0.831	0.725	0.063	0.004	500.5M
Averaged Embeddings	0.583		0.849	0.234	0.080	120M
Paragraph Vector	0.552		0.787	0.227	0.033	30M
LSTM Reader	0.680		0.880	0.421	0.311	45M
Attentive Reader	0.693		0.886	0.441	0.337	56M
Memory Network	0.612		0.861	0.288	0.055	90.1M
Answer Extraction						
RNN Labeler	0.357	0.471	0.240	0.536	0.626	41M
Sequence to Sequence						
Basic seq2seq	0.708	0.925	0.844	0.530	0.738	32M
Placeholder seq2seq	0.718	0.948	0.835	0.565	0.730	32M
Character seq2seq	0.677	0.963	0.841	0.462	0.731	4.1M
Character seq2seq (LM)	0.699	0.963	0.851	0.501	0.733	4.1M

Table 3: Results for all methods described in Section 3 on the test set. *F1* is the Mean F1 score described in 4. *Bound* is the upper bound on Mean F1 imposed by constraints in the method (see text for details). The remaining columns provide score breakdowns by property type and the number of model parameters.

truncated to the first 300 words except for Character seq2seq, which uses 400 characters. The embedding matrix used to encode words in the document uses $d_{in} = 300$ dimensions for the $N_w = 100,000$ most frequent words. Similarly, answer classification over the $N_{ans} = 50,000$ most frequent answers is performed using an answer representation of size $d_{out} = 300$.⁷ The first 10 words of the properties are embedded using the document embedding matrix. Following Cho et al. (2014), RNNs in seq2seq models use a GRU cell with a hidden state size of 1024. More details on parameters are reported in Table 4.

Method	Emb. Dims	Doc. Length	Property Length	Doc. Vocab. Size
Sparse BoW Baseline	N/A	300 words	10 words	50K words
Paragraph Vector	N/A	N/A	10 words	N/A
Character seq2seq	30	400 chars	20 chars	76 chars
All others	300	300 words	10 words	100K words

Table 4: Structural model parameters. Note that the Paragraph Vector method uses the output from a separate, unsupervised model as a document encoding, which is not counted in these parameters.

Optimization was performed with the Adam stochastic optimizer⁸ (Kingma and Adam, 2015) over mini-batches of 128 samples. Gradient clipping⁹ (Graves, 2013) is used to prevent instability in training RNNs. We performed a search over

⁷For models like Averaged Embedding and Paragraph Vector, the concatenation imposes a greater d_{out} .

⁸Using $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$.

⁹When the norm of gradient \mathbf{g} exceeds a threshold C , it is

50 randomly-sampled hyperparameter configurations for the learning rate and gradient clip threshold, selecting the one with the highest Mean F1 on the validation set. Learning rate and clipping threshold are sampled uniformly, on a logarithmic scale, over the range $[10^{-5}, 10^{-2}]$ and $[10^{-3}, 10^1]$ respectively.

4.2 Results and Discussion

Results for all models on the held-out set of test instances are presented in Table 3. In addition to the overall Mean F1 scores, the model families differ significantly in Mean F1 upper bound, and their relative performance on the relational and categorical properties defined in Section 2.4. We also report scores for properties containing dates, a subset of relational properties, as a separate column since they have a distinct format and organization. For examples of model performance on individual properties, see Table 5.

As expected, all classifier models perform well for categorical properties, with more sophisticated classifiers generally outperforming simpler ones. The difference in precision reading ability between models that use broad document statistics, like Averaged Embeddings and Paragraph Vectors, and the RNN-based classifiers is revealed in the scores for relational and especially date properties. As shown in Table 5, this difference is magnified in situations that are more difficult for a classifier, such as relational properties or properties with fewer training examples, where Attentive Reader outperforms Averaged Embeddings by a wide margin. This model family also has a high

scaled down i.e. $\mathbf{g} \leftarrow \mathbf{g} \cdot \min\left(1, \frac{C}{\|\mathbf{g}\|}\right)$.

Property	Test Instances	Mean F1						
		Averaged Embeddings	Attentive Reader	Memory Network	Basic seq2seq	Placeholder seq2seq	Character seq2seq	Character seq2seq (LM)
Categorical Properties								
instance of	734187	0.8545	0.8978	0.8720	0.8877	0.8775	0.8548	0.8659
sex or gender	267896	0.9917	0.9966	0.9936	0.9968	0.9952	0.9943	0.9941
genre	32531	0.5320	0.6225	0.5625	0.5511	0.5260	0.5096	0.5283
instrument	3665	0.7621	0.8415	0.7886	0.8377	0.8172	0.7529	0.7832
Relational Properties								
given name	218625	0.4973	0.8486	0.7206	0.8669	0.8868	0.8606	0.8729
located in	137253	0.4140	0.6195	0.4832	0.5484	0.6978	0.5496	0.6365
parent taxon	62685	0.1990	0.3467	0.2077	0.2044	0.7997	0.4979	0.5748
author	9517	0.0309	0.2088	0.1050	0.6094	0.6572	0.1403	0.3748
Date Properties								
date of birth	223864	0.0626	0.3677	0.0016	0.8306	0.8259	0.8294	0.8303
date of death	103507	0.0417	0.2949	0.0506	0.7974	0.7874	0.7897	0.7924
publication date	31253	0.3909	0.5549	0.4851	0.5988	0.5902	0.5903	0.5943
date of official opening	1119	0.1510	0.3047	0.1725	0.3333	0.3012	0.1457	0.1635

Table 5: Property-level Mean F1 scores on the test set for selected methods and properties. For each property type, the two most frequent properties are shown followed by two less frequent properties to illustrate long-tail behavior.

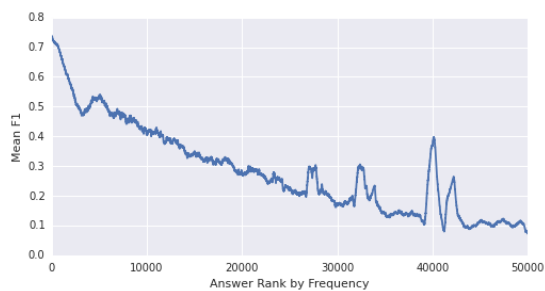


Figure 3: Per-answer Mean F1 scores for Attentive Reader (moving average of 1000), illustrating the decline in prediction quality as the number of training examples per answer decreases.

upper bound, as perfect classification across the 50,000 most frequent answers would yield a Mean F1 of 0.831. However, none of them approaches this limit. Part of the reason is that their accuracy for a given answer decreases quickly as the frequency of the answer in the training set decreases, as illustrated in Figure 3. As these models have to learn a separate weight vector for each answer as part of the softmax layer (see Section 3.1), this may suggest that they fail to generalize across answers effectively and thus require significant number of training examples per answer.

The only answer extraction model evaluated,

RNN Labeler, shows a complementary set of strengths, performing better on relational properties than categorical ones. While the Mean F1 upper bound for this model is just 0.434 because it can only produce answers that are present verbatim in the document text, it manages to achieve most of this potential. The improvement on date properties over the classifier models demonstrates its ability to identify answers that are typically present in the document. We suspect that answer extraction may be simpler than answer classification because the model can learn robust patterns that indicate a location without needing to learn about each answer, as the classifier models must.

The sequence to sequence models show a greater degree of balance between relational and categorical properties, reaching performance consistent with classifiers on the categorical questions and with RNN Labeler on relational questions. Placeholder seq2seq can in principle produce any answer that RNN Labeler can, and the performance on relational properties is indeed similar. As shown in Table 5, Placeholder seq2seq performs especially well for properties where the answer typically contains rare words such as the name of a place or person. When the set of possible answer tokens is more constrained, such

as in categorical or date properties, the Basic seq2seq often performs slightly better. Character seq2seq has the highest upper bound, limited to 0.963 only because it cannot produce an answer set with multiple elements. LM pretraining consistently improves the performance of the Character seq2seq model, especially for relational properties as shown in Table 5. The performance of the Character seq2seq, especially with LM pretraining, is a surprising result: It performs comparably to the word-level seq2seq models even though it must copy long character strings when doing extraction and has access to a smaller portion of the document. We found the character based models to be particularly sensitive to hyperparameters. However, using a pretrained language model reduced this issue and significantly accelerated training while improving the final score. We believe that further research on pretraining for character based models could improve this result.

5 Related Work

The goal of automatically extracting structured information from unstructured Wikipedia text was first advanced by Wu and Weld (2007). As Wikidata did not exist at that time, the authors relied on the structured *infoboxes* included in some Wikipedia articles for a relational representation of Wikipedia content. Wikidata is a cleaner data source, as the infobox data contains many slight variations in schema related to page formatting. Partially to get around this issue, the authors restrict their prediction model Kylin to 4 specific infobox classes, and only common attributes within each class.

A substantial body of work in relation extraction (RE) follows the distant supervision paradigm (Craven and Kumlien, 1999), where sentences containing both arguments of a knowledge base (KB) triple are assumed to express the triple’s relation. Broadly, these models use these distant labels to identify syntactic features relating the subject and object entities in text that are indicative of the relation. Mintz et al. (2009) apply distant supervision to extracting Freebase triples (Bollacker et al., 2008) from Wikipedia text, analogous to the relational part of WIKIREADING. Extensions to distant supervision include explicitly modelling whether the relation is actually expressed in the sentence (Riedel et al., 2010), and jointly reasoning over larger sets of sentences and relations (Sur-

deanu et al., 2012). Recently, Rocktäschel et al. (2015) developed methods for reducing the number of distant supervision examples required by sharing information between relations.

6 Conclusion

We have demonstrated the complexity of the WIKIREADING task and its suitability as a benchmark to guide future development of DNN models for natural language understanding. After comparing a diverse array of models spanning classification and extraction, we conclude that end-to-end sequence to sequence models are the most promising. These models simultaneously learned to classify documents and copy arbitrary strings from them. In light of this finding, we suggest some focus areas for future research.

Our character-level model improved substantially after language model pretraining, suggesting that further training optimizations may yield continued gains. Document length poses a problem for RNN-based models, which might be addressed with convolutional neural networks that are easier to parallelize. Finally, we note that these models are not intrinsically limited to English, as they rely on little or no pre-processing with traditional NLP systems. This means that they should generalize effectively to other languages, which could be demonstrated by a multilingual version of WIKIREADING.

Acknowledgments

We thank Jonathan Berant for many helpful comments on early drafts of the paper, and Catherine Finegan-Dollak for an early implementation of RNN Labeler.

References

- Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2015. Tensorflow: Large-scale machine learning on heterogeneous systems. *Software available from tensorflow.org*.
- Phoebe Ayers, Charles Matthews, and Ben Yates. 2008. *How Wikipedia works: And how you can be a part of it*. No Starch Press.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations (ICLR)*.

- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 1247–1250, New York, NY, USA. ACM.
- Kyunghyun Cho, Bart Van Merriënboer, Çalar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October. Association for Computational Linguistics.
- Mark Craven and Johan Kumlien. 1999. Constructing biological knowledge bases by extracting information from text sources. In *Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology*, pages 77–86. AAAI Press.
- Andrew M Dai and Quoc V Le. 2015. Semi-supervised sequence learning. In *Advances in Neural Information Processing Systems*, pages 3061–3069.
- Alex Graves. 2013. Generating sequences with recurrent neural networks. *CoRR*, abs/1308.0850.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1684–1692.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. In *Proceedings of the CVSC Workshop*, Sofia, Bulgaria. Association of Computational Linguistics.
- Diederik P Kingma and Jimmy Ba Adam. 2015. A method for stochastic optimization. In *International Conference on Learning Representation*.
- Karol Kurach, Marcin Andrychowicz, and Ilya Sutskever. 2015. Neural random-access machines. In *International Conference on Learning Representations (ICLR)*.
- Quoc V. Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of The 31st International Conference on Machine Learning*, pp. , 2014, pages 1188?–1196.
- Thang Luong, Ilya Sutskever, Quoc Le, Oriol Vinyals, and Wojciech Zaremba. 2015. Addressing the rare word problem in neural machine translation. In *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)*, pages 11–19, Beijing, China, July. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2 - Volume 2*, ACL '09, pages 1003–1011, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2015. Relation extraction: Perspective from convolutional neural networks. In *Proceedings of NAACL-HLT*, pages 39–48.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer.
- Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. 2015. Injecting Logical Background Knowledge into Embeddings for Relation Extraction. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in Neural Information Processing Systems*, pages 2431–2439.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 455–465. Association for Computational Linguistics.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A free collaborative knowledgebase. *Commun. ACM*, 57:78–85.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. May.

- Fei Wu and Daniel S Weld. 2007. Autonomously semantifying wikipedia. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 41–50. ACM.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, pages 649–657.