

Story Cloze Ending Selection Baselines and Data Examination

Todor Mihaylov

Research Training Group AIPHES
Institute for Computational Linguistics
Heidelberg University

mihaylov@cl.uni-heidelberg.de

Anette Frank

Research Training Group AIPHES
Institute for Computational Linguistics
Heidelberg University

frank@cl.uni-heidelberg.de

Abstract

This paper describes two supervised baseline systems for the Story Cloze Test Shared Task (Mostafazadeh et al., 2016a). We first build a classifier using features based on word embeddings and semantic similarity computation. We further implement a neural LSTM system with different encoding strategies that try to model the relation between the story and the provided endings. Our experiments show that a model using representation features based on average word embedding vectors over the given story words and the candidate ending sentences words, joint with similarity features between the story and candidate ending representations performed better than the neural models. Our best model achieves an accuracy of 72.42, ranking 3rd in the official evaluation.

1 Introduction

Understanding common sense stories is an easy task for humans but represents a challenge for machines. A recent common sense story understanding task is the 'Story Cloze Test' (Mostafazadeh et al., 2016a), where a human or an AI system has to read a given four-sentence story and select the proper ending out of two proposed endings. While the majority class baseline performance on the given test set yields an accuracy of 51.3%, human performance achieves 100%. This makes the task a good challenge for an AI system. The Story Cloze Test task is proposed as a Shared Task for LSDSem 2017¹. 17 teams registered for the Shared Task and 10 teams submitted their results².

¹Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics 2017

²<https://competitions.codalab.org/competitions/15333> - Story Cloze Test at CodaLab

Our contribution is that we set a new baseline for the task, showing that a simple linear model based on distributed representations and semantic similarity features achieves state-of-the-art results. We also evaluate the ability of different embedding models to represent common knowledge required for this task. We present an LSTM-based classifier with different representation encodings that tries to model the relation between the story and alternative endings and argue about its inability to do so.

2 Task description and data construction

The Story Cloze Test is a natural language understanding task that consists in selecting the right ending for a given short story. The evaluation data consists of a *Dev set* and a *Test set*, each containing samples of four sentences of a story, followed by two alternative sentences, from which the system has to select the proper story ending. An example of such a story is presented in Table 1.

The instances in the *Dev* and *Test* gold data sets (1871 instances each) were crowd-sourced together with the related ROC Stories corpus (Mostafazadeh et al., 2016a). The ROC stories consists of around 100,000 crowd-sourced short five sentence stories ranging over various topics. These stories do not feature a wrong ending, but with appropriate extensions they can be deployed as training data for the Story Cloze task.

Task modeling. We approach the task as a supervised classification problem. For every classification instance (*Story*, *Ending1*, *Ending2*) we predict one of the labels in $\text{Label}=\{\textit{Good},\textit{Bad}\}$.

Obtaining a small training set from *Dev* set. We construct a (small) training data set from the *Dev* set by splitting it randomly into a *Dev-Train* and a *Dev-Dev* set containing 90% and 10% of the

Story context	Good Ending	Bad ending
Mary and Emma drove to the beach. They decided to swim in the ocean. Mary turned to talk to Emma. Emma said to watch out for the waves.	A big wave knocked Mary down.	The ocean was a calm as a bathtub.

Table 1: Example of a given story with a bad and a good ending.

original *Dev* set. From each instance in *Dev-Train* we generate 2 instances by swapping *Ending1* and *Ending2* and inverting the class label.

Generating training data from ROC stories.

We also make use of the ROC Stories corpus in order to generate a large training data set. We experiment with three methods:

(i.) **Random endings.** For each story we employ the first 4 sentences as the story context. We use the original ending as *Good* ending and define a *Bad* ending by randomly choosing some ending from an alternative story in the corpus. From each story with one *Good* ending we generate 10 *Bad* examples by selecting 10 random endings.

(ii.) **Coherent stories and endings with common participants and noun arguments.** Given that some random story endings are too clearly unconnected to the story, here we aim to select *Bad* candidate endings that are coherent with the story, yet still distinct from a *Good* ending. To this end, for each story in the ROC Stories corpus, we obtain the lemmas of all pronouns (tokens with part of speech tag starting with ‘PR’) and lemmas of all nouns (tokens with part of speech tag starting with ‘NN’) and select the top 10 endings from other stories that share most of these features as *Bad* endings.

(iii.) **Random coherent stories and endings.** We also modify (ii.) so that we select the nearest 500 endings to the story context and select 10 randomly.

3 A Baseline Method

For tackling the problem of right story ending selection we follow a feature-based classification approach that was previously applied to bi-clausal classification tasks in (Mihaylov and Frank, 2016; Mihaylov and Nakov, 2016). It uses features based on word embeddings to represent the clauses and semantic similarity measured between these representations for the clauses. Here, we adopt this approach to model parts of the story and the candidate endings. For the given *Story* and the given

candidate *Endings* we extract features based on word embeddings. An advantage of this approach is that it is fast for training and that it only requires pre-trained word embeddings as an input.

3.1 Features

In our models we only deploy features based on word embedding vectors. We are using two types of features: (i) **representation features** that model the semantics of parts of the story using word embedding vectors, and (ii) **similarity scores** that capture specific properties of the relation holding between the story and its candidate endings. For computing similarity between the embedding representations of the story components, we employ cosine similarity.

The different feature types are described below.

(i) **Embedding representations for *Story* and *Ending*.** For each *Story* (sentences 1 to 4) and story endings *Ending1* and *Ending2* we construct a centroid vector from the embedding vectors \vec{w}_i of all words w_i in their respective surface yield.

(ii.) **Story to Ending Semantic Vector Similarities.** We calculate various similarity features on the basis of the centroid word vectors for all or selected sentences in the given *Story* and the *Ending1* and *Ending2* sentences, as well as on parts of the these sentences:

Story to Ending similarity. We assume that a given *Story* and its *Good Ending* are connected by a specific semantic relation or some piece of common sense knowledge. Their representation vectors should thus stand in a specific similarity relation to each other. We use their cosine similarity as a feature. Similarity between the story sentences and a candidate ending has already been proposed as a baseline by Mostafazadeh et al. (2016b) but it does not perform well as a single feature.

Maximized similarity. This measure ranks each word in the *Story* according to its similarity with the centroid vector of *Ending*, and we compute the average similarity for the top-ranked N

words. We chose the similarity scores of the top 1,2,3 and 5 words as features. Our assumption is that the average similarity between the *Story* representation and the top N most similar words in the *Ending* might characterize the proper ending as the *Good* ending. We also extract **maximized aligned similarity**. For each word in *Story*, we choose the most similar word from the yield of *Ending* and take the average of all best word pair similarities, as suggested in Tran et al. (2015).

Part of speech (POS) based word vector similarities. For each sentence in the given four sentence story and the candidate endings we performed part of speech tagging using the Stanford CoreNLP (Manning et al., 2014) parser, and computed similarities between centroid vectors of words with a specific tag from *Story* and the centroid vector of *Ending*. Extracted features for POS similarities include symmetric and asymmetric combinations: for example we calculate the similarity between *Nouns* from *Story* with *Nouns* from *Ending* and similarity between *Nouns* from *Story* with *Verbs* from *Ending* and vice versa.

The assumption is that embeddings for some parts of speech between *Story* and *Ending* might be closer to those of other parts of speech for the *Good* ending of a given story.

3.2 Classifier settings

For our feature-based approach we concatenate the extracted representation and similarity features in a feature vector, scale their values to the 0 to 1 range, and feed the vectors to a classifier. We train and evaluate a L2-regularized Logistic Regression classifier with the LIBLINEAR (Fan et al., 2008) solver as implemented in *scikit-learn* (Pedregosa et al., 2011).

For each separate experiment we tune the regularization parameter C with 5 fold cross-validation on the *Dev* set and then train a new model on the entire *Dev* set in order to evaluate on the *Test* set.

4 Neural LSTM Baseline Method

We compare our feature-based linear classifier baseline to a neural approach. Our goal is to explore a simple neural method and to investigate how well it performs with the given small dataset. We implement a Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) recurrent neural network model.

4.1 Representations

We are using the raw LSTM output of the encoder. We also experiment with an encoder with attention to model the relation between a story and a candidate ending, following (Rocktäschel et al., 2015).

(i) Raw LSTM representations. For each given instance (*Story*, *Ending1*, *Ending2*) we first encode the *Story* token word vector representations using a recurrent neural network (RNN) with long short-term memory (LSTM) units. We use the last output \mathbf{h}_L^s and \mathbf{c}_L^s states of the *Story* to initialize the first LSTM cells for the respective encodings \mathbf{e}_1 and \mathbf{e}_2 of *Ending1* and *Ending2*, where L is the token length of the *Story* sequence.

We build the final representation \mathbf{o}_{se1e2} by concatenating the \mathbf{e}_1 and \mathbf{e}_2 representations. Finally, for classification we use a softmax layer over the output \mathbf{o}_{se1e2} by mapping it into the target space of the two classes (Good, Bad) using a parameter matrix \mathbf{M}_o and bias \mathbf{b}_o as given in (Eq.1). We train using the cross-entropy loss.

$$label_{prob} = softmax(W^o o_{se1e2} + b_o) \quad (1)$$

(ii) Attention-based representations We also model the relation h^* between the *Story* and each of the *Endings* using the attention-weighted representation \mathbf{r} between the last token output \mathbf{h}_N^e of the *Ending* representation and each of the token representations $[\mathbf{h}_1^s \dots \mathbf{h}_L^s]$ of the *Story*, strictly following the attention definition by Rocktäschel et al. (2015). The final representation for each ending is presented by Eq.2, where \mathbf{W}^p and \mathbf{W}^x are trained projection matrices.

$$h^* = tanh(W^p r + W^x h_N^e) \quad (2)$$

We then present the output representation \mathbf{o}_{se1e2} as a concatenation of the encoded *Ending1* and *Ending2* representations h_{e1}^* and h_{e2}^* and use Eq.1 to obtain the output label likelihood.

(iii) Combined raw LSTM output and attention representation We also perform experiments with combined LSTM outputs and representations. In this setting we present the output \mathbf{o}_{se1e2} as presented in Eq.3:

$$o_{se1e2} = concat(e_1, h_{e1}^*, e_2, h_{e2}^*) \quad (3)$$

4.2 Model Parameters and Tuning

We perform experiments with configurations of the model using grid search on the batch size (50,

System	Accuracy
Human	100.00
msap	75.20
cogcomp	74.39
Our features baseline	72.42
Our neural system	72.10
ukp	71.67
DSSM	58.50
Skip-thoughts sim	55.20
Word2Vec sim	53.90
Majority baseline	51.30

Table 2: Comparison of our models to shared task participants’ results and other baselines. *Word2Vec sim*, *Skip-thoughts sim* and *DSSM* are described in (Mostafazadeh et al., 2016b).

100, 200, 300, 400, 500) and LSTM output size (128, 256, 384, 512), by training a simple model with raw LSTM encoding on *Dev-Train* and evaluating on the *Dev-Dev*. For each configuration we train 5 models and take the parameters of the best. The best result on the *Dev-Dev* set is achieved for LSTM with output size 384 and batch Size 500 after 7 epochs and achieves accuracy of 72.10 on the official *Test*. For learning rate optimization we use Adam optimizer (Kingma and Ba, 2015) with initial learning rate 0.001.

Parameter initialization. We initialize the LSTM weights with Xavier initialization (Glorot and Bengio, 2010) and bias with a constant zero vector.

5 Experiments and Results

Overall results. In Table 2 we compare our best systems to existing baselines, Shared Task participant systems³ and human performance. *Our features baseline* system is our best feature-based system using embeddings and *word2vec* trained on *Dev* and tuned with cross-validation. *Our neural system* employs raw LSTM encodings as described in Section 4.1(i) and it is trained on the *Dev-Dev* dataset which consists of 90% of the *Dev* dataset selected randomly and tuned on the rest of *Dev*. The best result in the task is achieved by Schwartz et al. (2017) (*msap*) who employ stylistic features combined with RNN representations. We have no information about *cogcomp* and *ukp*.

³<https://competitions.codalab.org/competitions/15333> - The Story Cloze Test Shared Task home page

Model variations and experiments. The Story Cloze Test is a story understanding problem. However, the given stories are very short and they require background knowledge about relations between the given entities, entity types and events defining the story and their endings, as well as relations between these events. We first train our feature-based model with alternative embedding representations in order to select the best source of knowledge for further experiments.

We experiment with different word embedding models pre-trained on a large number of tokens including *word2vec*⁴ (Mikolov et al., 2013), *GloVe* (Pennington et al., 2014) and *ConceptNet Numberbatch* (Speer and Chin, 2016). Results on training the feature-based model with different word embeddings are shown in Table 3. The results indicate how well the vector representation models perform in terms of encoding common sense stories. We present the performance of the embedding models depending on the defined features. We perform feature ablation experiments to determine the features which contribute most to the overall score for the different models. Using *All* features defined in Section 3.1, the *word2vec* vectors, trained on Google News 100B corpus perform best followed by *ConceptNet* enriched embeddings and *Glove* trained on Common Crawl 840B. The *word2vec* model suffers most when similarity features are excluded. We note that the *ConceptNet* embeddings do not decrease performance when similarity features are excluded, unlike all other models. We also see that the *POS similarities* are more important than the *MaxSim* and the *Sim (cosine between all words in Story and Ending)* as they yield worse results, for almost all models, when excluded from *All* features.

In column *WE E1, E2* we report results on features based only on *Ending1* and *Ending 2*. We note that the overall results are still very good. From this we can derive that the difference of *Good vs. Bad* endings is not only defined in the story context but it is also characterized by the information present in these sentences in isolation. This could be due to a reporting bias (Gordon and Van Durme, 2013) employed by the crowdworkers in the corpus construction process.

The last column *Sims only* shows results with features based only on similarity features. It in-

⁴<https://code.google.com/archive/p/word2vec/> - Pre-trained embeddings on Google News dataset (100B words).

Model	All	All wo POS sim	All wo MaxSim	All wo Sim	WE S, E1, E2+Sim	WE E1, E2	Sims only
Word2Vec GN 100B 300d	72.42	71.41	71.94	72.10	71.51	70.71	58.15
Concepnet 300d	72.05	72.05	72.05	72.05	71.83	71.67	61.67
Glove 840B 300d	71.41	71.09	71.89	72.26	70.82	70.71	60.28
Glove 6B 200d	69.43	69.75	68.31	69.64	68.04	68.68	62.37
Glove 6B 300d	68.84	69.32	69.21	69.05	68.79	68.89	61.19
Glove 6B 100d	68.84	68.09	67.93	68.41	67.66	67.56	60.82
Glove 6B 50d	64.89	66.01	64.19	64.67	64.78	64.83	58.57

Table 3: Experiments using linear classifier with features based on word embeddings. Trained on *Dev* (tuned with cross-validation) and evaluated on *Test*.

Model	Epoch	Dev-Dev	Test
LSTM Raw	7	77.12	72.10
LSTM Raw + Att	2	79.25	68.30
Attention	9	72.79	63.22

Table 4: Comparison between LSTM representation strategies.

cludes all story-to-ending semantic vector similarities described in Section 3.1.

We also perform experiments with the neural LSTM model. In Table 4 we compare results of the LSTM representation models that we examined for the task. We trained the models on the *Dev-Train* for 10 epochs and take the best performing model on the *Dev-Dev* dataset.

Our best LSTM model uses only raw LSTM encodings of the *Story* and the candidate *Endings*, without using attention. Here the *Attention* representation is intended to capture semantic relations between the *Story* context and the candidate *Endings*, similar to the *Similarities only* setup examined with the feature-based approach. Considering the low performance of *Attention*, the poor results of the *semantic similarity features* and the high performance of the feature-based model with *Ending only* features we hypothesize that the reason for this unexpected result is that the background knowledge presented in the training data is not enough to learn strong relations between the story context and the endings.

Experiments with generated data. We also try to employ the data from the ROC Stories corpus by generating labeled datasets following all approaches described in Section 2. Training our best neural model using all of the generated datasets separately without any further data selection yields results close to the random baseline of the ending selection task. We also try to filter the generated data by training several feature-based

and neural models with our best configurations and evaluating the generated data. We take only instances that have been classified correctly by all models. The idea here was to generate much more data (with richer vocabulary) that performs at least as good as the *Dev* data as training. However the results of the models trained on these datasets were not better than the one trained on *Dev* and *Dev-Dev* (for the neural models).

6 Conclusion and Future work

In this work we built two strong supervised baseline systems for the Story Cloze task: one based on semantic features based on word embedding representations and bi-clausal similarity features obtained from them, and one on based on a neural network LSTM-based encoder model. The neural network approach trained on a small dataset performs worse than the feature-based classifier by a small margin only, and our best model ranks 3rd according to the shared task web page.

In terms of data, it seems that the most important features are coming from word representations trained on large text corpora rather than relations between the data. Also we can train a model that performs well only on the given endings, without a given context which could mean that there is a bias in the annotation process. However, this requires more insights and analysis.

In future work we plan improve the current results on this (or a revised) dataset by collecting more external knowledge and obtaining more or different training data from the original ROC Stories corpus.

Acknowledgments. This work is supported by the German Research Foundation as part of the Research Training Group “Adaptive Preparation of Information from Heterogeneous Sources” (AIPHES) under grant No. GRK 1994/1.

We would like to thank the reviewers for their valuable and constructive comments.

References

- Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. Liblinear: A library for large linear classification. *J. Mach. Learn. Res.*, 9:1871–1874, June.
2010. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 9:249–256.
- Jonathan Gordon and Benjamin Van Durme. 2013. Reporting bias and knowledge acquisition. *Proc. 2013 Work. Autom. Knowl. base Constr. - AKBC '13*, (circa):25–30.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: a Method for Stochastic Optimization. *Int. Conf. Learn. Represent. 2015*, pages 1–15.
- Christopher D Manning, John Bauer, Jenny Finkel, Steven J Bethard, Mihai Surdeanu, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60.
- Todor Mihaylov and Anette Frank. 2016. Discourse Relation Sense Classification Using Cross-argument Semantic Similarity Based on Word Embeddings. In *Proceedings of the Twentieth Conference on Computational Natural Language Learning - Shared Task*, pages 100–107, Berlin, Germany. Association for Computational Linguistics.
- Todor Mihaylov and Preslav Nakov. 2016. SemanticZ at SemEval-2016 Task 3: Ranking relevant answers in community question answering using semantic similarity based on fine-tuned word embeddings. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval '16*, San Diego, California, USA.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT '13*, pages 746–751, Atlanta, Georgia, USA.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016a. A Corpus and Evaluation Framework for Deeper Understanding of Commonsense Stories. *Naacl*.
- Nasrin Mostafazadeh, Lucy Vanderwende, Wen-tau Yih, Pushmeet Kohli, and James Allen. 2016b. Story cloze evaluator: Vector space representation evaluation by predicting what happens next. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 24–29, Berlin, Germany, August. Association for Computational Linguistics.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global Vectors for Word Representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2015. Reasoning about Entailment with Neural Attention. *Unpublished*, (2015):1–9, sep.
- Roy Schwartz, Maarten Sap, Ioannis Konstas, Leila Zilles, Yejin Choi, and Noah A Smith. 2017. The Effect of Different Writing Tasks on Linguistic Style: A Case Study of the ROC Story Cloze Task. *Proc. Link. Model. Lexical, Sentential Discourse-level Semant. Shar. Task*, feb.
- Robert Speer and Joshua Chin. 2016. An ensemble method to produce high-quality word embeddings. *arXiv preprint arXiv:1604.01692*.
- Quan Hung Tran, Vu Tran, Tu Vu, Minh Nguyen, and Son Bao Pham. 2015. JAIST: Combining multiple features for answer selection in community question answering. In *Proceedings of the 9th International Workshop on Semantic Evaluation, SemEval '15*, pages 215–219, Denver, Colorado, USA.