

Generating Adequate Distractors for Multiple-Choice Questions

Cheng Zhang¹, Yicheng Sun¹, Hejia Chen² and Jie Wang¹

¹*Department of Computer Science, University of Massachusetts, Lowell, MA 01854, U.S.A.*

²*School of Computer Science and Technology, Xidian University, Xi'an 710126, P.R. China*

Keywords: Multiple-Choice Questions, Distractors, Word Embeddings, Word Edit Distance.

Abstract: This paper presents a novel approach to automatic generation of adequate distractors for a given question-answer pair (QAP) generated from a given article to form an adequate multiple-choice question (MCQ). Our method is a combination of part-of-speech tagging, named-entity tagging, semantic-role labeling, regular expressions, domain knowledge bases, word embeddings, word edit distance, WordNet, and other algorithms. We use the US SAT (Scholastic Assessment Test) practice reading tests as a dataset to produce QAPs and generate three distractors for each QAP to form an MCQ. We show that, via experiments and evaluations by human judges, each MCQ has at least one adequate distractor and 84% of MCQs have three adequate distractors.

1 INTRODUCTION

Generating MCQs on a given article is a quick and effective method for assessing the reader's comprehension of the article. An MCQ typically consists of a QAP and a few distractors. This paper is focused on generating adequate distractors for a given QAP in connection to the underlying article.

An adequate distractor must satisfy the following requirements: (1) it is an incorrect answer to the question; (2) it is grammatically correct and consistent with the underlying article; (3) it is semantically related to the correct answer; and (4) it provides distraction so that the correct answer could be identified only with some understanding of the underlying article.

Given a QAP for a given article, we study how to generate adequate distractors that are grammatically correct and semantically related to the correct answer in the sense that the distractors, while incorrect, look similar to the correct answer with a sufficient distracting effect—that is, it should be hard to distinguish distractors from the correct answer without some understanding of the underlying article. A distractors could be a single word, a phrase, a sentence segment, or a complete sentence.

In particular, we are to generate three adequate distractors for a QAP to form an MCQ. One way to generate a distractor is to substitute a word or a phrase contained in the answer with an appropriate word or

a phrase that maintains the original part of speech. Such a word or phrase could be an answer itself or contained in an answer sentence or sentence segment. For convenience, we refer to such a word or phrase as a target word.

If a target word is a number with an explicit or implicit quantifier, or anything that can be converted to a number, we call it a type-1 target. If a target word is a person, location, or organization, we call it a type-2 target. Other target words (nouns, phrasal nouns, adjectives, verbs, and adjectives) are referred to as type-3 targets. We use different methods to generate distractors for targets of different types.

Our distractor generation method is a combination of part-of-speech (POS) tagging (Toutanova et al., 2003), named-entity (NE) tagging (Nadeau and Sekine, 2007; Ali et al., 2010; Peters et al., 2017), semantic-role labeling (Martha et al., 2005; Shi and Lin, 2019), regular expressions, domain knowledge bases on people, locations, and organizations, word embeddings (such as Word2vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), Subwords (Bojanowski et al., 2017), and spherical text embedding (Meng1 et al., 2019)), word edit distance (Levenshtein, 1966), WordNet (<https://wordnet.princeton.edu>), and some other algorithms. We show that, via experiments, our method can generate adequate distractors for a QAP to form an MCQ with a high successful rate.

The rest of the paper is organized as follows: We

describe related work in Section 2 and present our distractor generation method in Section 3. We then show in Section 4 that, via experiments on the official SAT practice reading tests and evaluation by human judges, each MCQ has at least one adequate distractor and 84% of MCQs have three adequate distractors. We conclude the paper in Section 5.

2 RELATED WORK

Methods of generating adequate distractors for MCQs are typically following two directions: domain-specific knowledge bases and semantic similarity (Pho et al., 2014; Rao and Saha, 2018).

Methods in the first direction are focused on lexical information have used part-of-speech (POS) tags, word frequency, WordNet, domain ontology, distributional hypothesis, and pattern matching, to find the target word’s synonym, hyponym and hypernym as distractor candidates. (Mitkov and Ha, 2003; Correia et al., 2012; Susanti et al., 2015)

Methods in the second direction analyze the semantic similarity of the target word using Word2Vec model for generating distractor candidates (Jiang and Lee, 2017; Susanti et al., 2018).

However, it is difficult to use Word2Vec or other word-embedding methods to find adequate distractors for polysemous answer words. Moreover, previous efforts have focused on finding some forms of distractors, instead of making them look more distracting. This paper is an attempt to tackle these issues.

3 DISTRACTOR GENERATION

Our method takes the original article and the answer as input, and generates distractors as output. Figure 1 depicts the data flow of our method.

Answers in QAPs are classified into two kinds. The first kind consists of just a single target word while the second kind consists of multiple target words. The latter is the case when the answer is a sentence or a sentence segment.

For an answer of the first kind, if it is a type-1 or type-2 target, we use the methods described in Section 3.1 to generate three distractors; if it is a type-3 target, we use the method described in Section 3.2 to generate distractor candidates. If there are at least three candidates, then select three candidates with the highest ranking scores.

For an answer of the second kind, for each type of a target word contained in it, we use the methods described in both Sections 3.1 and 3.2 to generate

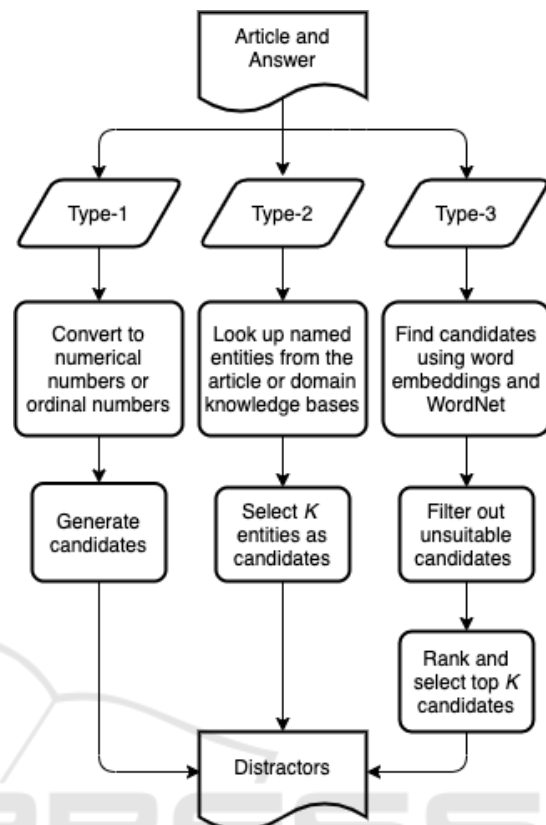


Figure 1: Distractor generation flowchart.

distractors for target words in a fixed ordered preference of subjects, objects, adjectives for subjects, adjectives for objects, predicates, adverbs, which can be obtained by semantic-role labeling. Target words are replaced according to the following preference: type-1 temporal, type-1 numerical value, type-2 person, type-2 location, type-2 organization, type-3 noun (phrasal noun), type-3 adjective, type-3 verb (phrasal verb), and type-3 adverb.

If the number of distractors for a given preference is less than three, then we generate extra distractors for a target word in the next preference. If all preference is gone through we still need more distractors, we could extend the selection threshold values to allow more candidates to be selected.

3.1 Distractors for Type-1 and Type-2 Targets

If a type-1 target is a point in time, a time range, a numerical number, an ordinal number, or anything that can be converted to a numerical number or an ordinal number (e.g. Friday may be converted to 5), which can be recognized by regular expressions based on a POS tagger, then we devise several algorithms to al-

ter time and number, and randomly select one of these algorithms when generating distractors. For example, we may increase or decrease the answer value by one or two units, change the answer value at random from a small range of values around the answer, or simply change the answer value at random. If a numerical value or an ordinal number is converted from a word, then the result is converted back to the same form. For example, suppose that the target word is “Friday”, which is converted to a number 5. If the distractor is a number 4, then it is converted to Thursday.

If a type-2 target is a person, then we first look for different person names that appear in the article using an NE tagger to identify them, and then randomly choose a name as a distractor. If there are no other names in the article, then we use Synonyms (<http://www.synonyms.com>) or a domain knowledge base on notable people we constructed to find a distractor. If a type-2 target is a location or an organization, we find a distractor in the same way by first looking for other locations or organizations in the article, and then using Synonyms and domain knowledge bases to look for them if they cannot be found in the article. For example, If the target word is a city, then a distractor should also be city that is “closely” related to the target word. Distractors to the answer word “New York” should be cities in the same league, such as “Boston”, “Philadelphia”, and “Chicago”.

3.2 Distractors for Type-3 Targets

For a type-3 target, we find distractor candidates using word embeddings with similarity in a threshold interval (e.g., [0.6, 0.85]) so that a candidate is not too close nor too different from the correct answer and hypernyms using WordNet (Miller, 1995). Note that a similarity interval of [0.6, 0.85] for word embeddings often include antonyms of the target word, and we can use WordNet or an online dictionary to determine antonyms.

Not all distractor candidates are suitable. Thus, we first filter out unsuitable candidates as follows:

1. Remove distractor candidates that contain the target word, for it may be too close to the correct answer. For example, if “breaking news” is a generated distractor candidate for the target word “news”, then it is removed from the candidate list since it contains the target word.
2. Remove distractor candidates that have the same prefix of the target word with edit distance less than three, for such candidates may often be misspelled words from the target word. For example, suppose the target word is “knowledge”, then Word2vec may return a misspelled candidate

“knowldage” with a high similarity, which should be removed.

We then rank each remaining candidate using the following measure:

1. Compute the Word2vec cosine similarity score S_v for each distractor candidate w_c with the target word w_t . Namely,

$$S_v = \text{sim}(v(w_c), v(w_t)),$$

where $v(w)$ denotes a word embedding of w .

2. Compute the WordNet WUP score (Wu and Palmer, 1994) S_n for each distractor candidate with the target word. If the distractor candidate cannot be found in the WordNet dataset, set the WUP score to 0.1 for the following reason: If a word with a high score of word-embedding similarity to the target word but does not exist in the WordNet dataset, then it is highly likely that the word is misspelled, and so its ranking score should be reduced.
3. Compute the edit distance score S_d of each distractor candidate with target word by the following formula:

$$S_d = 1 - \frac{1}{1 + e^E},$$

where E is the edit distance. Thus, a larger edit distance E results in a smaller score S_d .

4. Compute the final ranking score R for each distractor candidate w_c with respect to the target word w_t by

$$R'(w_c, w_t) = \begin{cases} \frac{1}{4}(2S_v + S_n + S_d), & \text{if } w_c \text{ is an} \\ & \text{antonym of } w_t, \\ \frac{1}{3}(S_v + S_n + S_d), & \text{otherwise,} \end{cases}$$

$$R(w_c, w_t) = -R'(w_c, w_t) \log R'(w_c, w_t).$$

Note that S_v, S_n, S_d are each between 0 and 1, and so $R'(w_c, w_t)$ is between 0 and 1, which implies that $-\log R'(w_c, w_t) > 0$. Also note that we give more weight to antonyms.

4 EVALUATIONS

We implemented our method using the latest versions of POS tagger¹, NE tagger (Peters et al., 2017), semantic-role labeling (Shi and Lin, 2019), and fast-Text (Mikolov et al., 2018). We used the US SAT practice reading tests² as a dataset for evaluations.

¹<https://nlp.stanford.edu/software/tagger.shtml>

²<https://collegereadiness.collegeboard.org/sat/practice/full-length-practice-tests>

There are a total of eight SAT practice reading tests, each consisting of five articles for a total of 40 articles. Each article in the SAT practice reading tests consists of around 25 sentences and we generated about 10 QAPs from each article. To evaluate our distractor generation algorithm, we selected independently at random slightly over 100 QAPs. After removing a smaller number of QAPs with pronouns as target words, we have a total of 101 QAPs for evaluations.

We generated 3 distractors for each QAP for a total of 303 distractors, and evaluated distractors based on the following criteria:

1. A distractor is *adequate* if it is grammatically correct and relevant to the question with distracting effects.
2. An MCQ is *adequate* if each of the three distractors is adequate.
3. An MCQ is *acceptable* if one or two distractors are adequate.

We define two levels of distracting effects: (1) *sufficient distraction*: It requires an understanding of the underlying article to choose the correct answer; (2) *distraction*: It only requires an understanding of the underlying question to choose the correct answer. A distractor has no distracting effect if it can be determined wrong by just looking at the distractor itself.

Evaluations were carried out by humans and the results are listed below:

1. All distractors generated by our method are grammatically correct.
2. 98% distractors (296 out of 303) are relevant to the QAP with distraction.
3. 96% distractors (291 out of 303) provide sufficient distraction.
4. 84% MCQs are adequate.
5. All MCQs are acceptable (i.e., with at least one adequate distractor).

Given below are a few adequate MCQs with automatically generated distractors by our method:

Example 1

Question: What does no man like to acknowledge? (SAT practice test 2 article 1)

Correct answer: that he has made a mistake in the choice of his profession.

Distractors:

1. that he has made a mistake in the choice of his association.
2. that he has made a mistake in the choice of his engineering.

3. that he has made a mistake in the way of his profession.

Example 2

When should ethics apply? (SAT practice test 2 article 2)

Correct answer: when someone makes an economic decision.

Distractors:

1. when someone makes an economic request.
2. when someone makes an economic proposition.
3. when someone makes a political decision.

Example 3

Question: What did Chie hear? (SAT practice test 1 article 1)

Correct answer: her soft scuttling footsteps, the creak of the door.

Distractors:

1. her soft scuttling footsteps, the creak of the driveway.
2. her soft scuttling footsteps, the creak of the stairwell.
3. her soft scuttling footsteps, the knock of the door.

Example 4

Question: Who might duplicate itself? (SAT practice test 1 article 3)

Correct answer: the deoxyribonucleic acid molecule.

Distractors:

1. the deoxyribonucleic acid coenzyme.
2. the deoxyribonucleic acid polymer.
3. the deoxyribonucleic acid trimer

Example 5

Question: When does Deep Space Industries of Virginia hope to be harvesting metals from asteroids? (SAT practice test 1 article 5)

Correct answer: by 2020.

Distractors:

1. by 2021.
2. by 2030.
3. by 2019.

Example 6

Question: What did a British study of the way women search for medical information online indicate? (SAT practice test 2 article 3)

Correct answer: An experienced Internet user can, at least in some cases, assess the trustworthiness and probable value of a Web page in a matter of seconds.

Distractors:

1. An experienced Supernet user can, at least in some cases, assess the trustworthiness and probable value of a Web page in a matter of seconds.
2. An experienced CogNet user can, at least in some cases, assess the trustworthiness and probable value of a Web page in a matter of seconds.
3. An inexperienced Internet user can, at least in some cases, assess the trustworthiness and probable value of a Web page in a matter of seconds.

Example 7

What does a woman know better than a man? (SAT test 2 article 4)

Correct answer: the cost of life.

Distractors:

1. the cost of happiness.
2. the cost of experience.
3. the risk of life.

Example 8

This example presents a distractor without sufficient distraction.

Question: What are subject to egocentrism, social projection, and multiple attribution errors?

Correct answer: their insights.

Distractors:

1. their perspectives.
2. their findings.
3. their valuables.

The last distractor can be spotted wrong by just looking at the question: It is easy to tell that it is out of place without the need to read the article.

5 CONCLUSIONS AND FINAL REMARKS

We presented a novel method using various NLP tools for generating adequate distractors for a QAP to form

an adequate MCQ on a given article. This is an interesting area with important applications. Experiments and evaluations on MCQs generated from the SAT practice reading tests indicate that our approach is promising.

A number of improvements can be explored. For example, we may improve the ranking measure to help select a better distractor for a target word from a list of candidates. Another direction is explore how to produce generative distractors using neural networks, instead of just replacing a few target words in a given answer.

ACKNOWLEDGMENT

This work was supported in part by funding from Eola Solutions, Inc. We thank Hao Zhang and Changfeng Yu for discussions.

REFERENCES

- Ali, H., Chali, Y., and Hasan, S. A. (2010). Automation of question generation from sentences. In *Proceedings of QG2010: The Third Workshop on Question Generation*, pages 58–67.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Correia, R., Baptista, J., Eskenazi, M., and Mamede, N. (2012). Automatic generation of cloze question stems. In Caseli, H., Villavicencio, A., Teixeira, A., and Perdigão, F., editors, *Computational Processing of the Portuguese Language*, pages 168–178, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Jiang, S. and Lee, J. (2017). Distractor generation for chinese fill-in-the-blank items. In *BEA@EMNLP*.
- Levenshtein, V. I. (1966). Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady*, 10:707.
- Martha, P., Dan, G., and Paul, K. (2005). The proposition bank: a corpus annotated with semantic roles. *Computational Linguistics Journal*, 31(1):10–1162.
- Meng1, Y., Huang, J., Wang, G., Zhang, C., Zhuang, H., Kaplan, L., and Han1, J. (2019). Spherical text embedding. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*.
- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., and Joulin, A. (2018). Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of*

- the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'13, pages 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Miller, G. A. (1995). Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41.
- Mitkov, R. and Ha, L. A. (2003). Computer-aided generation of multiple-choice tests. In *Proceedings of the HLT-NAACL 03 Workshop on Building Educational Applications Using Natural Language Processing - Volume 2*, HLT-NAACL-EDUC '03, page 17–22, USA. Association for Computational Linguistics.
- Nadeau, D. and Sekine, S. (2007). A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1):3–26.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Peters, M., Ammar, W., Bhagavatula, C., and Power, R. (2017). Semi-supervised sequence tagging with bidirectional language models. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1756–1765, Vancouver, Canada. Association for Computational Linguistics.
- Pho, V.-M., André, T., Ligozat, A.-L., Grau, B., Illouz, G., and François, T. (2014). Multiple choice question corpus analysis for distractor characterization. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4284–4291, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Rao, D. and Saha, S. K. (2018). Automatic multiple choice question generation from text : A survey. *IEEE Transactions on Learning Technologies*, pages 14–25.
- Shi, P. and Lin, J. (2019). Simple BERT Models for Relation Extraction and Semantic Role Labeling. *arXiv e-prints*, page arXiv:1904.05255.
- Susanti, Y., Iida, R., and Tokunaga, T. (2015). Automatic generation of english vocabulary tests. In *CSEUDU*.
- Susanti, Y., Tokunaga, T., Nishikawa, H., and Obari, H. (2018). Automatic distractor generation for multiple-choice english vocabulary questions. In *RPTTEL 13*.
- Toutanova, K., Klein, D., Manning, C. D., and Singer, Y. (2003). Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the 2003 conference of the North American chapter of the association for computational linguistics on human language technology-volume 1*, pages 173–180. Association for Computational Linguistics.
- Wu, Z. and Palmer, M. (1994). Verbs semantics and lexical selection. In *Proceedings of the 32nd Annual Meeting on Association for Computational Linguistics, ACL '94*, page 133–138, USA. Association for Computational Linguistics.