

Multi-Label Transfer Learning for Multi-Relational Semantic Similarity

Li Zhang and Steven R. Wilson and Rada Mihalcea

Computer Science and Engineering

University of Michigan

{zharry, steverw, mihalcea}@umich.edu

Abstract

Multi-relational semantic similarity datasets define the semantic relations between two short texts in multiple ways, e.g., similarity, relatedness, and so on. Yet, all the systems to date designed to capture such relations target one relation at a time. We propose a multi-label transfer learning approach based on LSTM to make predictions for several relations simultaneously and aggregate the losses to update the parameters. This multi-label regression approach jointly learns the information provided by the multiple relations, rather than treating them as separate tasks. Not only does this approach outperform the single-task approach and the traditional multi-task learning approach, it also achieves state-of-the-art performance on all but one relation of the Human Activity Phrase dataset.

1 Introduction

Semantic similarity, or relating short texts or sentences¹ in a semantic space – be those phrases, sentences or short paragraphs – is a task that requires systems to determine the degree of equivalence between the underlying semantics of the two sentences. Although relatively easy for humans, this task remains one of the most difficult natural language understanding problems. The task has been receiving significant interest from the research community. For instance, from 2012 to 2017, the International Workshop on Semantic Evaluation (SemEval) has been holding the Semantic Textual Similarity (STS) shared tasks (Agirre et al., 2012, 2013b, 2015, 2016; Cer et al., 2017), dedicated to tackling this problem, with close to 100 team submissions each year.

In some semantic similarity datasets, an example consists of a sentence pair and a single annotated similarity score, while in others, each pair

¹In this work, we do not consider word level similarity.

comes with multiple annotations. We refer to the latter as *multi-relational semantic similarity* tasks. The inclusion of multiple annotations per example is motivated by the fact that there can be different relations, namely different types of similarity between two sentences. So far, these relations have been treated as separate tasks, where a model trains and tests on one relation at a time while ignoring the rest. However, we hypothesize that each relation may contain useful information about the others, and training on only one relation inevitably neglects some relevant information. Thus, training jointly on multiple relations may improve performance on one or more relations.

We propose a joint **multi-label** transfer learning setting based on LSTM, and show that it can be an effective solution for the multi-relational semantic similarity tasks. Due to the small size of multi-relational semantic similarity datasets and the recent success of LSTM-based sentence representations (Wieting and Gimpel, 2018; Conneau et al., 2017), the model is pre-trained on a large corpus and transfer learning is applied using fine-tuning. In our setting, the network is jointly trained on multiple relations by outputting multiple predictions (one for each relation) and aggregating the losses during back-propagation. This is different from the traditional multi-task learning setting where the model makes one prediction at a time, switching between the tasks. We treat the multi-task setting and the single-task setting (i.e., where a separate model is learned for each relation) as baselines, and show that the multi-label setting outperforms them in many cases, achieving state-of-the-art performance on all but one relation of the Human Activity Phrase dataset (Wilson and Mihalcea, 2017).

In addition to success on multi-relational semantic similarity tasks, the multi-label transfer learning setting that we propose can easily be

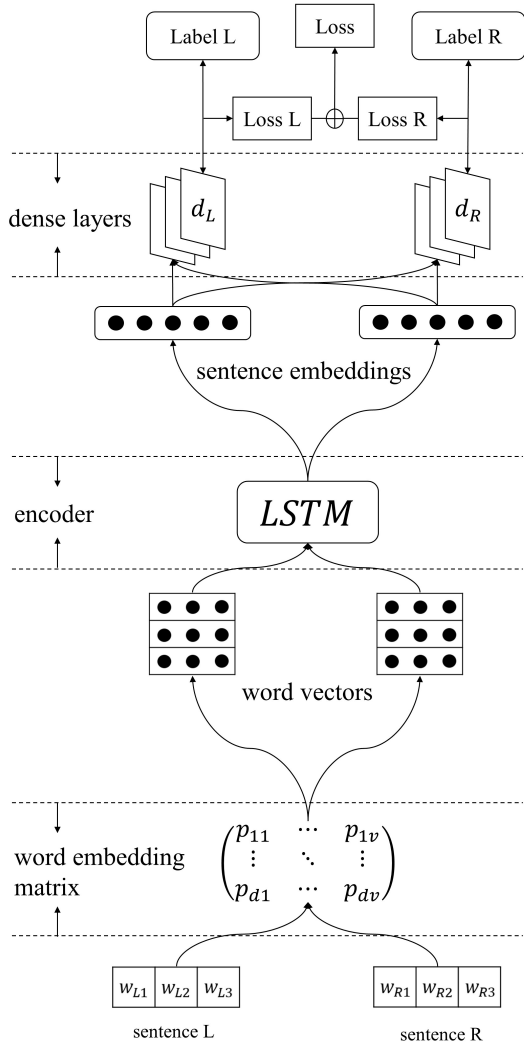


Figure 1: Overview of the multi-label architecture.

paired with other neural network architectures and applied to any dataset with multiple annotations available for each training instance.

2 Multi-Label Transfer Learning

We introduce a multi-label transfer learning setting by modifying the architecture of the LSTM-based sentence encoder, specifically designed for multi-relational semantic similarity tasks.

2.1 Architecture

We employ the “hard-parameter sharing” setting (Caruana, 1998), where some hidden layers are shared across multiple tasks while each task has its own specific output layer. As shown in Figure 1, using an example of a semantic similarity dataset with two relations, sentence L and sentence R in a

pair are first mapped to word vector sequences and then encoded as sentence embeddings. Up to this step, the choice of the word embedding matrix and sentence encoder is flexible, and we outline our choice in the sections to follow. For each relation that has been annotated with a ground-truth label, a dedicated output dense layer takes the two sentence embeddings as input and outputs a probability distribution across the range of possible scores. The output dense layers follow the methods of Tai et al. (2015).

With two such dense output layers, two losses are calculated, one for each relation. The total loss is calculated as the sum of the two losses for back-propagation which updates all parameters in the end-to-end network.

2.2 Model

We use InferSent (Conneau et al., 2017) as the sentence encoder due to its outstanding performances reported on various semantic similarity tasks.

Due to the small sizes of the evaluation datasets, we use the sentence encoder pre-trained on the Stanford Natural Language Inference corpus (Bowman et al., 2015) and Multi-Genre Natural Language Inference corpus (Williams et al., 2018), and transfer to the semantic similarity tasks using fine-tuning. In this process, the output layers for multi-label learning discussed above are stacked on top of the InferSent network, forming an end-to-end model for training and testing on semantic similarity tasks.

2.3 Comparison with Multi-Task Learning

Neither multi-task nor multi-label learning have been used for multi-relational semantic similarity datasets. For these datasets, either multi-task or multi-label learning can be achieved by treating each relation as a “task.” The key differences between the two are the relations involved in each forward-backward pass and the timing of the parameter updates.

Consider a training step in the two-relation example in Figure 1:

A **multi-task learning** model would pick a batch of sentences pairs, only consider *Label L*, only calculate *Loss L*, and all parameters except those of dense layer d_R are updated. Then, within the same batch,² the model would only consider

²In general multi-task learning, a new batch is picked after switching tasks. In multi-relational semantic similarity datasets, each task is a relational label, which shares the same

Label R, only calculate *Loss R*, and all parameters except those of dense layer d_L are updated.

A **multi-label learning** model (our model) would pick a batch of sentences pairs, consider both *Label L* and *Label R*, calculate *Loss L* and *Loss R*, aggregate them as the total loss, and update all parameters.

3 Experiments

To show the effectiveness of the multi-label transfer learning setting, we experiment on three semantic similarity datasets with multiple relations annotated, and use one LSTM-based sentence encoder that has been very successful in many downstream tasks.

3.1 Datasets

We study three semantic similarity datasets with multiple relations with texts of different lengths, spanning phrases, sentences, and short paragraphs.

Human Activity Phrase (Wilson and Mihalcea, 2017): a collection of pairs of phrases regarding human activities, annotated with the following four different relations.

- Similarity (SIM): The degree to which the two activity phrases describe the same thing, semantic similarity in a strict sense. Example of high similarity phrases: *to watch a film* and *to see a movie*.
- Relatedness (REL): The degree to which the activities are related to one another, a general semantic association between two phrases. Example of strongly related phrases: *to give a gift* and *to receive a present*.
- Motivational alignment (MA): The degree to which the activities are (typically) done with similar motivations. Example of phrases with potentially similar motivations: *to eat dinner with family members* and *to visit relatives*.
- Perceived actor congruence (PAC): The degree to which the activities are expected to be done by the same type of person. An example of a pair with a high PAC score: *to pack a suitcase* and *to travel to another state*.

The phrases are generated, paired and scored on Amazon Mechanical Turk.³ The annotated input.

³<https://www.mturk.com/>

scores range from 0 to 4 for SIM, REL and MA, and -2 to 2 for PAC. The evaluation is based on the Spearman's ρ correlation coefficient between the systems' predicted scores and the human annotations. There are 1,000 pairs in the dataset. We also use the supplemental 1,373 pairs from Zhang et al. (2018) in which 1,000 pairs are randomly selected for training and the rest are used for development. We then treat the original 1,000 pairs as a held-out test set so that our results are directly comparable with those previously reported.

SICK (Marelli et al., 2014b,a): the Sentences Involving Compositional Knowledge benchmark, which includes a large number of sentence pairs that are rich in the lexical, syntactic and semantic phenomena. Each pair of sentences is annotated in two dimensions: relatedness and entailment. The relatedness score ranges from 1 to 5, and Pearson's r is used for evaluation; the entailment relation is categorical, consisting of entailment, contradiction, and neutral. There are 4439 pairs in the train split, 495 in the trial split used for development and 4906 in the test split. The sentence pairs are generated from image and video caption datasets before being paired up using some algorithm. Due to the lack of human supervision in the process, some sentence pairs display minimal difference in semantic components, making the SICK tasks simpler than the others we study.

Typed-Similarity (Agirre et al., 2013b): a collection of meta-data describing books, paintings, films, museum objects and archival records taken from Europeana,⁴ presented as the pilot track in the SemEval 2013 STS shared task. Typically, the items consist of title, subject, description, and so on, describing a cultural heritage item and, sometimes, a thumbnail of the item itself. For the purpose of measuring semantic similarity, we concatenate all the textual entries such as title, creator, subject and description into a short paragraph that is used as input, although the annotations might be informed of the image aspects of the meta-data. Each pair of items is annotated on eight dimensions of similarity: general similarity, author, people involved, time, location, event or action involved, subject and description. There are 750

⁴<http://www.europeana.eu/>

pairs in the train split, of which we randomly sample 500 for training and 250 for development, and 721 in the test split. Pearson’s r is used for evaluation.

3.2 Baselines

We compare the multi-label setting with two baselines:

- **Single-task**, where each relation is treated as an individual task. For each relation, a model with only one output dense layer is trained and tested, ignoring the annotations of all other relations.
- **Multi-task**, where only one relation is involved during each round of feed-forward and back-propagation.

3.3 Experimental Details

In each experiment, we use stochastic gradient descent and a batch size of 16. We tune the learning rate over $\{0.1, 0.5, 1, 5\}$ and number of epochs over $\{10, 20\}$. For each dataset discussed above, we tune these hyperparameters on the development set. All other hyperparameters maintain their values from the original code.⁵ In the single-task setting, the model is trained and tested on each relation, ignoring the annotations of other relations. In the multi-task settings, the model is trained and tested on all the relations in a dataset. In the multi-task setting, relations are presented to the model in the order they are listed in the result tables within each batch.

4 Evaluation

The results are shown in Tables 1, 2 and 3. For every experiment (represented by a cell in the tables), 30 runs with different random seeds are recorded and the average is reported. For each relation (each column in the tables), let the true mean performance of multi-label learning, single-task baseline and multi-task baseline be μ_{MLL} , μ_{single} , μ_{MTL} , respectively. Two one-sided Student’s t-tests are conducted to test if multi-label learning outperforms the baselines for that relation. The significance level is chosen to be 0.05. A down-arrow \downarrow indicates that our proposed multi-label learning underperforms a baseline, while an up-arrow \uparrow indicates that our proposed multi-label learning outperforms a baseline.

⁵<https://github.com/facebookresearch/InferSent>

5 Discussion

5.1 Results

For the Human Activity Phrase dataset, the single-task setting already achieves state-of-the-art performances on SIM, REL and PAC relations, surpassing the previous best results reported by Zhang et al. (2018), which achieved Spearman’s correlation coefficient of .710 in SIM, .715 in REL, .690 in MA and .549 in PAC. This approach is based on fine-tuning a bi-directional LSTM with average-pooling pre-trained on translated texts (Wieting and Gimpel, 2018). Using multi-label learning, our model is able to gain a statistically significant improvement in the performance of REL compared to the single-task setting, while maintaining performance for the other relations. The traditional multi-task setting, however, performs significantly worse than the other settings.

For the entailment task on the SICK dataset, our multi-label setting outperforms the single-task baseline and the previous best results of InferSent. These best results consisted of an accuracy of 86.3% achieved using a logistic regression classifier and sentence embeddings generated by pre-trained InferSent as features (Conneau et al., 2017). In the relatedness task, this setting achieved a Pearson’s correlation coefficient of .885, which even our multi-label setting is unable to beat. However, the multi-label setting does have a statistically significant performance gain compared to the single-task setting in the relatedness task, while the traditional multi-task setting underperforms the other settings.

For the Typed-Similarity dataset, the previous best results are achieved using rich feature engineering without the use of sentence embeddings, with a different scoring scheme for each relation (Agirre et al., 2013a). While this method yielded better results than all of the transfer learning approaches we compare, it should be noted that this approach is specific to tackling this dataset, unlike the transfer learning settings that are generalizable to other scenarios. One potential reason for the discrepancy in performance is that some relations such as time, people involved, or events may be easily or sometimes trivially captured by information retrieval techniques such as named entity recognition. Using sentence embeddings and transfer learning for all the relations, though simpler, may face greater challenge in the rela-

| | general | author | people | time | location | event | subject | description |
|--------|----------------|---------------|---------------|-------------|-----------------|--------------|----------------|--------------------|
| MLL | .744 | .721 | .640 | .713 | .751 | .611 | .697 | .737 |
| Single | .750↓ | .690↑ | .619↑ | .712 | .744↑ | .606↑ | .694↑ | .718↑ |
| MTL | .718↑ | .689↑ | .611↑ | .697↑ | .723↑ | .579↑ | .669↑ | .714↑ |

Table 1: The performance in Pearson’s r on the Typed-Similarity dataset, in accordance with the specification of the dataset to allow for direct comparison with previous results. The results of single task and multi-task learning (MTL) are followed by \uparrow if it is statically significantly lower than those of multi-label learning (MLL), and they are followed by \downarrow otherwise.

| | SIM | REL | MA | PAC |
|--------|------------|------------|-----------|------------|
| MLL | .720 | .721 | .682 | .557 |
| Single | .719 | .717↑ | .682 | .555 |
| MTL | .683↑ | .686↑ | .651↑ | .515↑ |

Table 2: The performance in Spearman’s ρ on the Human Activity Phrase dataset.

| | Relatedness | Entailment |
|--------|--------------------|-------------------|
| MLL | .882 | 86.7 |
| Single | .874↑ | 86.4↑ |
| MTL | .871↑ | 86.2↑ |

Table 3: The performance in Pearson’s r on the SICK dataset, in accordance with the specification of the dataset to allow for direct comparison with previous results.

tions mentioned above. Among the three transfer learning approaches, our multi-label setting is still superior, outperforming the single-task setting in over half of the relations, and outperforming the multi-task setting in all relations.

5.2 Empirical Recommendation

While our results above show that multi-label learning is almost always the most effective way to transfer sentence embeddings in multi-relational semantic similarity tasks, in some situations simply training with one relation might yield better performance (such as the general similarity relation in the Typed-Similarity dataset). This suggests that the choice of multi-label learning or single-task learning can be tuned as a hyperparameter empirically for the optimal performance on a task.

5.3 Other Considerations and Discussions

In the multi-label setting, we calculate the total loss by summing the loss from each dimension. We also explore weighting the loss from each di-

mension by factors of 2, 5 and 10, but doing so hurts the performance for all dimensions.

In the multi-task setting, we attempt different ordering of the dimensions when presenting them to the model within a batch of examples, but the difference in performance is not statistically significant. Furthermore, the multi-task setting takes about n times longer to train than the multi-label setting, where n is number of dimensions of annotations.

6 Conclusions

We introduced a multi-label transfer learning setting designed specifically for semantic similarity tasks with multiple relations annotations. By experimenting with a variety of relations in three datasets, we showed that the multi-label setting can outperform single-task and traditional multi-task settings in many cases.

Future work includes exploring the performance of this setting with other sentence encoders, as well as multi-label datasets outside of the domain of semantic similarity. This may include NLP datasets annotated with author information for multiple dimensions, or computer vision datasets with multiple annotations for scenes.

Acknowledgments

This material is based in part upon work supported by the Michigan Institute for Data Science, by the John Templeton Foundation (grant #61156), by the National Science Foundation (grant #1815291), and by DARPA (grant #HR001117S0026-AIDA-FP-045). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the Michigan Institute for Data Science, the John Templeton Foundation, the National Science Foundation, or DARPA.

References

- Eneko Agirre, Nikolaos Aletras, Aitor Gonzalez-Agirre, German Rigau, and Mark Stevenson. 2013a. [Ubc_uos-typed: Regression for typed-similarity](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 132–137, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Claire Cardie, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Weiwei Guo, Inigo Lopez-Gazpio, Montse Maritxalar, Rada Mihalcea, German Rigau, Larraitz Uria, and Janyce Wiebe. 2015. [Semeval-2015 task 2: Semantic textual similarity, english, spanish and pilot on interpretability](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 252–263, Denver, Colorado. Association for Computational Linguistics.
- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. [Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511, San Diego, California. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. [Semeval-2012 task 6: A pilot on semantic textual similarity](#). In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393, Montréal, Canada. Association for Computational Linguistics.
- Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. 2013b. [*sem 2013 shared task: Semantic textual similarity](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 32–43, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Rich Caruana. 1998. Multitask learning. In *Learning to learn*, pages 95–133. Springer.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. [Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Marco Marelli, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli. 2014a. [Semeval-2014 task 1: Evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 1–8, Dublin, Ireland. Association for Computational Linguistics.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014b. [A sick cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. [Improved semantic representations from tree-structured long short-term memory networks](#). 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 1556–1566, Beijing, China. Association for Computational Linguistics.
- John Wieting and Kevin Gimpel. 2018. [Paranmt-50m: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Steven Wilson and Rada Mihalcea. 2017. [Measuring semantic relations between human activities](#). In

Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 664–673, Taipei, Taiwan. Asian Federation of Natural Language Processing.

Li Zhang, Steven R Wilson, and Rada Mihalcea. 2018. Sequential network transfer: Adapting sentence embeddings to human activities and beyond. *arXiv preprint arXiv:1804.07835*.