

How to Encode Domain Information in Relation Classification

Elisa Bassignana^{1,2} Viggo Unmack Gascou² Frida Nøhr Laustsen²
Gustav Kristensen² Marie Haahr Petersen² Rob van der Goot^{2,3}
Barbara Plank^{2,3}

¹IT University of Copenhagen ²Pioneer Centre for AI ³LMU Munich
{elba, viga, fril, gukr, mhpe, robv}@itu.dk b.plank@lmu.de

Abstract

Current language models require a lot of training data to obtain high performance. For Relation Classification (RC), many datasets are domain-specific, so combining datasets to obtain better performance is non-trivial. We explore a multi-domain training setup for RC, and attempt to improve performance by encoding domain information. Our proposed models improve > 2 Macro-F1 against the baseline setup, and our analysis reveals that not all the labels benefit the same: The classes which occupy a similar space across domains (i.e., their interpretation is close across them, for example *physical*) benefit the least, while domain-dependent relations (e.g., *part-of*) improve the most when encoding domain information.

Keywords: Relation Classification, Domain, Multi-domain training, Robustness

1. Introduction

Relation Classification (RC) is the task of identifying the semantic relation between two given entities. The task is beneficial for many different downstream tasks which involve Natural Language Understanding. For example, question answering, knowledge base population, or summarization. In addition to the wide variety of downstream applications, as most information extraction tasks, RC is topic-specific, meaning that depending on the topic the information to extract can vary a lot. For example, in the music domain we may want to extract that a song is included in a musical album, while in the politics domain we may have a politician winning a political election. While current deep learning models require a lot of training data, collecting and annotating text from every domain is time-consuming and expensive.

In this project, we explore the critical setup of multi-domain training with the aim of identifying the best setup for maximizing the training data (by including data coming from different domains), without losing domain-specific information. To do so, we compare multiple ways of enriching the input instances with domain information (see Section 2).

Encoding information about where a certain utterance originates from has been previously explored in other Natural Language Processing fields. In the multi-lingual space, [Conneau and Lample \(2019\)](#) exploited language embeddings for multi-lingual model training. [Ammar et al. \(2016\)](#) first proposed to use language embeddings for training a multi-lingual syntactic parser for seven European languages, and showed improved performance. Later work also successfully trained parsers with the so called treebank embeddings for datasets within the

same language ([Stymne et al., 2018](#)) or language family ([Smith et al., 2018](#)). Other work have used special language ids to mark the language of each instance in the context of machine translation ([Liu et al., 2020](#)). To the best of our knowledge, these approaches have been exploited mostly in multi-lingual setups and syntactic tasks. In this work, we explore a gap and test their effectiveness for encoding domain information in a semantic setup: Relation Classification. We compare “dataset embeddings” and “domain markers” from previous work with a new approach exploiting domain-specific entity types. Our contributions are:

- CrossRE 2.0, an extension of the CrossRE dataset ([Bassignana and Plank, 2022](#)) with 3.3k new annotations in the news domain in order to balance data across domains;
- We propose the first multi-domain training baseline on CrossRE;
- We test previous work for encoding dataset information in RC, and compare it with a new RC-specific technique; We present an in-depth analysis of the results obtained.

2. Domain Encoding for Relation Classification

2.1. Dataset Embeddings

The dataset embedding model tries to encode information about the domain with ad-hoc embeddings on the encoder side. Dataset embeddings are vector representations learned at training time that aim at capturing distinctive properties of multiple data sources into a continuous vector, without losing

their heterogeneous characteristics. Originally they were often concatenated to the word embedding and then used in e.g., a Bi-LSTM (Stymne, 2020; Wagner et al., 2020; van der Goot et al., 2021). However, since large language models have become the standard, this has become trickier, as they have a pre-determined input size. To enable usage of dataset embeddings, van der Goot and de Lhoneux (2021) propose to sum them to the input representation. In our setup, we treat each domain as a separate data source.

2.2. Special Domain Markers

An intuitive and simple alternative way of encoding the domain is by using special tokens appended to the input text itself. This has been previously done in machine translation in order to mark the different languages (Liu et al., 2020). We concatenate a special token at the beginning of each instance containing the corresponding domain (e.g., [MUSIC] or [NEWS]). These domain markers are treated by the tokenizer as special tokens, i.e., they are not tokenized into subwords, so the model learns a representation for each of them during training.

2.3. Entity Type Information

The domain-specific entity types carry out information which can be relevant for identifying the correct relation label. Following Zhong and Chen (2021) we add entity type information in the representation of the input (see model description in Section 3.2). We test two different approaches to do this. First, we use the 39 fine-grained types proposed by Liu et al. (2021) including e.g., *musician* or *political party*, which are domain-specific. In the second setup we map these fine-grained types into five coarse-grained types. For example, *musician* and *political party* are mapped to *person* and *organization* respectively. While this last approach shades domain information, it guarantees a more condensed entity type distribution, and it can be combined with the other two setups.

3. Experimental Setup

3.1. Data

CrossRE (Bassignana and Plank, 2022),¹ is a manually-annotated dataset for multi-domain RC including 17 relation types spanning over six diverse text domains: news (📰), politics (🏛️), natural science (🔬), music (🎵), literature (📖), and artificial intelligence (🤖). The dataset was annotated on top of CrossNER (Liu et al., 2021), a Named Entity Recognition (NER) dataset. Table 1 reports

¹Released with a GNU General Public License v3.0.

	SENTENCES				RELATIONS			
	train	dev	test	tot.	train	dev	test	tot.
📰	217	1,320	3,053	4,590	156	1,043	2,115	3,314
🏛️	164	350	400	914	175	300	396	871
🔬	101	350	400	851	502	1,616	1,831	3,949
🎵	103	351	400	854	355	1,340	1,393	3,088
📖	100	350	399	849	496	1,861	2,333	4,690
🤖	100	400	416	916	397	1,539	1,591	3,527
🎮	100	350	431	881	350	1,006	1,127	2,483
tot.	885	3,471	5,499	9,855	2,431	8,705	10,786	21,922

Table 1: **CrossRE 2.0 Statistics.** Number of sentences and number of relations of the news extension (first row), and statistics of the original domains of CrossRE (below).

the statistics of CrossRE. While the train, dev, and test splits include similar amounts of sentences across the six domains, the number of annotated relations varies over a wider range. The reason for this includes different average sentence lengths, and different relation densities across the domains. In the original dataset, the news domain is particularly small. This domain comes from a different source with respect to the other five—the CoNLL-2003 dataset (Tjong Kim Sang and De Meulder, 2003) and Wikipedia (Liu et al., 2021) respectively.

3.1.1. CrossRE 2.0

With the aim of mitigating the effect of dataset size on the model performance, which influences the comparison of results across domains, we expand the news domain of CrossRE. We follow the annotation guidelines by Bassignana and Plank (2022)² and manually annotate more than 4.5k sentences from the CoNLL-2003 dataset (Tjong Kim Sang and De Meulder, 2003)—the original data source of this domain. The data is annotated by a hired linguist compensated fairly according to national salary scales, who got extensive training for the task. We refer to Bassignana and Plank (2022) for the discussion on the annotation agreement because for consistency we employed the same annotator who annotated the original version of CrossRE. Table 1 reports the statistics of our extension, with over 3k annotated relations and an overall total in news (including the original dataset) of 4.1k, which is in line with the other domains. The dataset extension is public available in the CrossRE repository.³ We train the model in a multi-domain setup, i.e., mixing the six training sets of CrossRE.

²https://github.com/mainlp/CrossRE/blob/main/crossre_annotation/CrossRE-annotation-guidelines.pdf

³<https://github.com/mainlp/CrossRE/>

3.2. Model Architecture

We use the baseline model of the original CrossRE paper.⁴ Following the model architecture first proposed by Baldini Soares et al. (2019), the implementation by Bassignana and Plank (2022) augments the sentence with four entity markers $e_1^{start}, e_1^{end}, e_2^{start}, e_2^{end}$ surrounding the two entities. When exploiting the entity types, the information is injected in the entity markers (e.g., [E1:person]) The augmented sentence is then passed through a pre-trained encoder, and the classification is made by a linear layer over the concatenation of the start markers $[\hat{s}_{e_1^{start}}, \hat{s}_{e_2^{start}}]$. We run our experiments over five random seeds and report the average. All hyperparameters follow Bassignana and Plank (2022). Our code is available on GitHub.⁵

4. Results

Table 2 reports the Macro-F1 results of our experiments. The dataset embeddings setup fails to beat the baseline. The main reason for this is the limited amount of training data in our setup, which challenges the model in learning them.⁶ The dataset embeddings are summed to the word, segment, and position embeddings, which are then updated all at once in the forward pass. Additionally, in settings where they are successful, these embeddings are usually used to disambiguate datasets coming from different data sources or languages. Here instead we are at a more fine-grained level, trying to model different topics, with data extracted from the same source (except for news).

Concatenating a special domain marker at the beginning of the sentence results in the best performance (36.90 Macro-F1), with the highest improvement in the music domain (+2.96) and sometimes small yet consistent improvements across all domains. The fine-grained entity types lead to decreased performance, because their distribution is very sparse across the six domains. For example, out of the 39, the news domain from CoNLL 2003 only includes *person, location, organization* and *miscellaneous*, resulting in a performance decrease of -1.78. Using the coarse-grained entity types—shared across all the domains—results in a slightly better average Macro-F1 (34.99) than employing the fine-grained ones (34.10), but it does not improve over the baseline either. As this setup lacks domain information, we try combining the coarse-grained entity representation with the special domain markers. Within this setup results are

⁴<https://github.com/mainlp/CrossRE>

⁵Project repository <https://github.com/viggo-gascou/multi-domain-rc>

⁶We manually inspected the dataset embeddings before and after training.

		📄	🏠	🍀	🎵	📺	📰	avg.
dev	BASELINE	25.45	31.35	39.46	39.69	38.84	38.09	35.48
	DATASET EMB.	15.38	22.22	24.77	32.64	30.95	29.80	25.96
	DOMAIN MARK.	26.36	32.77	40.31	42.65	40.59	38.71	36.90
	FINE-GRAIN.	23.67	32.67	35.35	38.76	38.23	35.94	34.10
	COARSE-GRAIN.	24.46	31.56	38.59	39.33	38.09	37.90	34.99
	DOM. + COARSE	24.52	32.02	39.63	42.19	40.01	37.17	35.92
test	BASELINE	24.73	34.12	39.67	39.96	44.64	35.71	36.47
	DOMAIN MARK.	26.72	37.62	43.57	41.48	44.88	37.69	38.66

Table 2: **Performance Scores.** Macro-F1 scores of the explored setups. DOM. + COARSE indicates the combination of special domain markers with the coarse-grained entity types.




Figure 1: **Domain Representation.** PCA plot of the untrained embeddings of the instances in the development set, colored by domain.

mixed across the domains: While most of them (except AI) improve over the coarse-grained entity type (without domain information), only politics, science, music and literature overcome the baseline. The overall average across the domains results in a minor improvement of +0.44.

We evaluate the best setup (the special domain markers) on the test set in order to confirm our findings. Following the trend on the development set, the improvement over the baseline is +2.19 Macro-F1. The lower performance range of news over the other domains (both in dev and test) indicates that the different data source has a high impact even with the more uniform data distribution across domains proposed with our dataset extension.

5. Analysis

Domain Representation To inspect how much domain information the out-of-the-box embeddings contain, in Figure 1 we plot the PCA representations of the untrained embeddings (with *bert-base-cased*, the encoder used by the RC model) of the instances in the development set. We do not include the news domain in this plot because given its different data source (consisting of shorter sentences, typically news headlines), the news instances are very distant from the other domains, resulting in an high overlap of the latter. The current




Figure 2: **Relation Representation.** PCA plot of the trained embeddings of the most frequent relation labels in the development set, colored by relation labels and shaped by domain.

setup allows us to shed light on the remaining five domains, besides news which we already know is very distinctive. The domains are already relatively distinguishable with the untrained encoder. The two technical domains, science and AI, marginally overlap; politics is completely detached; and only music and literature overlap significantly. Our intuition is that encoding additional domain information (see Section 2) may not be particularly relevant.

Relation Representation We dive deeper into analyzing the relations and explore whether in the baseline setup (i.e., without additional domain information) the representation of instances coming from different domains, but belonging to the same class, are close to each other. In Figure 2 we plot the PCA representation of the trained baseline model of the instances with the most frequent relation labels in the development set, separated by class and domain. The main finding from this plot is that most of the classes are quite clustered, independently from the domain they belong to. For example, the *physical* relation on the right side has instances from all the domains. Similarly, the *arti-*

fact and the *role* labels towards the bottom-left corner of the plot. Interestingly, the *general-affiliation* relation presents clustered representations of the instances in the literature and music domains, but it still dominates the upper left side of the plot. Less surprisingly, the *related-to* label, listed as None-Of-The-Above (NOTA) in the guidelines, has a more sparse distribution across the plot.

The labels which present a less defined cluster (i.e., the ones whose meaning shifts the most across domains) are the ones which benefit the most from the special domain markers. For example, *related-to* improves from a baseline value of 20.99 F1 up to 24.21 in the special domain markers setup; *named* goes from 68.25 to 71.30 F1; and *part-of* improves from 32.79 to 35.54 F1. In contrary, the relation labels which present a better defined cluster already within the baseline (see Figure 2) do not benefit much from the additional encoding of domain information. For example, the per-label F1 scores of the *physical*, *general-affiliation*, and *role* relations in the baseline and special domain markers setups are respectively 77.16 and 77.51, 54.09 and 54.46, 65.60 and 65.11.

6. Conclusion

We explore how to encode domain information in a multi-domain training setup for the domain-specific task of RC. We propose CrossRE 2.0, a dataset extension of CrossRE (Bassignana and Plank, 2022) for balancing the amount of data across the six domains included in it. We manage to improve the multi-domain training baseline by > 2 Macro-F1 with a simple, but effective technique which encodes domain information in special domain markers concatenated at the beginning of each input. Our analysis reveals that not all of the relation labels benefit the same from the domain encoding: The most generic, with a diverse interpretation across domains (e.g., *part-of*) are the ones which gain the most in terms of per-label F1.

7. Ethics Statement

We do not foresee any potential risk related to this work. The data we use is published freely by Liu et al. (2021) and Bassignana and Plank (2022).

For the dataset extension, we hired an expert with a linguistics degree employed following national salary rates. The cost of the annotation process amounts to ≈ 1 \$ per annotated relation.

Acknowledgments

We thank our annotator for the great job and substantial help given to this project. We thank the NLPnorth group at ITU and the MaiNLP group at LMU for feedback on an earlier version of this paper. Elisa Bassignana and Barbara Plank are supported by the Independent Research Fund Denmark (Danmarks Frie Forskningsfond; DFF) Sapere Aude grant 9063-00077B. Barbara Plank is in parts supported by the European Research Council (ERC) grant agreement No. 101043235.

8. Bibliographical References

Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. [Many languages, one parser](#). *Transactions of the Association for Computational Linguistics*, 4:431–444.

Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. [Matching the blanks: Distributional similarity for relation learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905, Florence, Italy. Association for Computational Linguistics.

Elisa Bassignana and Barbara Plank. 2022. [CrossRE: A cross-domain dataset for relation extraction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3592–3604, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.

Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2021. [Crossner: Evaluating cross-domain named entity recognition](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13452–13460.

Aaron Smith, Bernd Bohnet, Miryam de Lhoneux, Joakim Nivre, Yan Shao, and Sara Stymne. 2018. [82 treebanks, 34 models: Universal Dependency parsing with multi-treebank models](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.

Sara Stymne. 2020. [Cross-lingual domain adaptation for dependency parsing](#). In *Proceedings of the 19th International Workshop on Treebanks and Linguistic Theories*, pages 62–69, Düsseldorf, Germany. Association for Computational Linguistics.

Sara Stymne, Miryam de Lhoneux, Aaron Smith, and Joakim Nivre. 2018. [Parser training with heterogeneous treebanks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 619–625, Melbourne, Australia. Association for Computational Linguistics.

Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.

Rob van der Goot and Miryam de Lhoneux. 2021. [Parsing with pretrained language models, multiple datasets, and dataset embeddings](#). In *Proceedings of the 20th International Workshop on*

Trebanks and Linguistic Theories (TLT, SyntaxFest 2021), pages 96–104, Sofia, Bulgaria. Association for Computational Linguistics.

Rob van der Goot, Ahmet Üstün, and Barbara Plank. 2021. [On the effectiveness of dataset embeddings in mono-lingual, multi-lingual and zero-shot conditions](#). In *Proceedings of the Second Workshop on Domain Adaptation for NLP*, pages 183–194, Kyiv, Ukraine. Association for Computational Linguistics.

Joachim Wagner, James Barry, and Jennifer Foster. 2020. [Treebank embedding vectors for out-of-domain dependency parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8812–8818, Online. Association for Computational Linguistics.

Zexuan Zhong and Danqi Chen. 2021. [A frustratingly easy approach for entity and relation extraction](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 50–61, Online. Association for Computational Linguistics.