

DaNetQA: a yes/no Question Answering Dataset for the Russian Language

Taisia Glushkova^{*1[0000-0003-3521-3834]}, Alexey Machnev^{1[0000-0001-5828-9810]},
Alena Fenogenova^{2[0000-0003-3139-1668]}, Tatiana Shavrina^{2[0000-0002-6976-0185]},
Ekaterina Artemova^{1[0000-0003-4920-1623]}, and
Dmitry I. Ignatov^{1[0000-0002-6584-8534]}

¹ National Research University Higher School of Economics, Moscow, Russia
<https://cs.hse.ru/en/ai/computational-pragmatics/>

² Sberbank, Moscow, Russia

tolushkova@edu.hse.ru

{Fenogenova.A.S,Shavrina.T.O}@sberbank.ru

{echernyak,amachnev,dignatov}@hse.ru

Abstract. DaNetQA, a new question-answering corpus, follows BoolQ [1] design: it comprises natural yes/no questions. Each question is paired with a paragraph from Wikipedia and an answer, derived from the paragraph. The task is to take both the question and a paragraph as input and come up with a yes/no answer, i.e. to produce a binary output. In this paper we present a reproducible approach to DaNetQA creation and investigate transfer learning methods for task and language transferring. For task transferring we leverage three similar sentence modelling tasks: 1) a corpus of paraphrases, Paraphraser, 2) an NLI task, for which we use the Russian part of XNLI, 3) another question answering task, SberQUAD. For language transferring we use English to Russian translation together with multilingual language fine-tuning.

Keywords: question answering · transfer learning · transformers.

1 Introduction

Creation of new datasets, aimed at new, challenging tasks, describing complex phenomena, related to various aspects of language understanding and usage, is core to the current view of modern language technologies. However, the majority of the datasets, created and published at the best venues, target the English-language tasks and cultural aspects, related to English-speaking society.

In response to the bias towards English, new datasets and benchmarks are developed, that comprise multiple languages. One of the well-known examples of such multilingual datasets is XNLI [2], which is a natural language inference dataset. Although this dataset is developed for 15 languages, including low-resource ones, the approach to its creation still excessively utilizes English data:

* The first two authors have equal contribution.

the dataset entries are manually translated from English to other languages, without any specific adjustments. Although the translation-based approach is quick and dirty and allows to overcome lack of dataset for any language, other concerns arise. From a general point of view, we understand, that translating is different from natural every day language usage [3]. Thus translated datasets may have different statistics and word usage in comparison to text, composed from scratch. This may affect the quality of the models, trained further on the translated datasets, when applied to the real-life data.

Other approach to dataset creation involves collecting datasets following the guidelines and annotation schemes, designed for original datasets in English. RuRED [4] is a recent example of such dataset: it is created following TACRED [5] annotation scheme and collection pipeline. This approach however is criticized for the obvious lack of novelty.

Nevertheless in this paper we stick to the second approach to the dataset creation and more or less follow the pipeline, developed for BoolQ [1], to create a new dataset for binary questions in Russian, which we refer to as DaNetQA. We only deviate from the BoolQ pipeline, if we do not have access to proprietary data sources and use instead crowd sourcing. The motivation to re-create BoolQ lies, first, in the lack of question-answering dataset for Russian, and second, in the fact, that binary question answering appears to be a more challenging task, when compared to SQuAD-like and natural language inference tasks [6]. Thus we hope that DaNetQA may become of great use both for chat-bot technologies, which massively use question answering data, and for thorough evaluation on deep contextual encoders, such as BERT [7] or XLM-R [8].

As the annotations for DaNetQA are crowd-sourced and require payment, we explore different strategies that can help to increase the quality without the need to annotate larger amounts of data. This leads us to two strategies of transfer learning: transferring from tasks, which have similar setting, and transferring from English, keeping in mind, that DaNetQA recreates BoolQ, but in a different language. Our main contributions are the following:

1. We create and intent to publish in open access a new middle scale dataset for the Russian language, DaNetQA, which comprises yes/no questions, paired with paragraphs, providing enough information to answer the questions (Subsection 2.1) and report its statistics (Subsection 2.2);
2. We establish a simple baseline and a more challenging deep baseline for DaNetQA (Subsection 2.2);
3. We explore applicability of transfer learning techniques, of which some overcome the established baseline (Section 3).

The paper is organised as follows. Section 2 describes the data collection process and provides the reader with their basic statistics. In Section 3, all the conducted experiments are described. Section 4 discusses the obtained results. In Section 5, related work is summarised. Section 6 concludes the paper.

2 Dataset

2.1 Collection

Our approach on DaNetQA dataset creation is inspired by the work of [1], where the pipeline from NQ [9] is used as a base.

Questions are generated on a crowdsourcing platform Yandex.Toloka³, which has proven to be a good source for slavic languages data generation compared to other language groups. Questions are created by crowd workers following an instruction, that suggests phrases which can be used to start a phrase – a broad list of templates (more than 50 examples).

Generated yes/no questions are then treated as queries in order to retrieve relevant Wikipedia pages with the use of Google API. Questions are only kept if 3 Wikipedia pages could be returned, in which case the question and the text of an article are passed forward to querying a pre-trained BERT-based model for SQuAD to extract relevant paragraphs.

Finally, crowd workers label question and paragraph pairs with “yes” or “no” answers, and the questions that could not be answered based on the information from the paragraph are marked as “not answerable”. Annotating data in such manner is quite expensive, since not only crowd workers read through and label thousands of pairs manually, but there is also used a high overlap of votes to ensure the high quality of the dataset. In addition, Google API also requires some payments to process high numbers of queries.

The final labels for each pair are picked based on the majority of votes. In case of uncertainty when there is no label picked by majority, the pair is being checked and labeled manually by the authors.

Aside from the overlap used on the final step, we use a number of gold-standard control questions that are randomly mixed into the tasks to make sure that crowd workers do the annotation responsibly. Also, each annotator goes through a small set of learning tasks before starting labelling the actual pairs.

2.2 Statistics

In the following section we analyze our corpus to better understand the nature of the collected questions and paragraphs, and compute all kinds of descriptive statistics. Statistics presented below, include the minimum, average and maximum length of questions and paragraphs in tokens (Table 1) and the distribution of text lengths in the entire dataset as a whole. Total number of yes/no questions in Dev/Test/Train sets can be found in Table 2.

We also provide a brief overview of categories covered by our question and paragraph pairs (Table 3) and a t-SNE visualization of a LDA model for 15 topics, trained on concatenated questions and paragraphs (Figure 1). Each topic is labelled with top-3 tokens that represent a topic summary.

Since the crowd workers were provided (but not restricted) with question starter templates, in order to generate yes/no questions, we decided to check

³ <https://toloka.yandex.ru>

Table 1: Descriptive statistics of collected question/paragraph pairs.

	Count	Mean	Min	25%	50%	75%	Max
Questions	2691	5	2	5	5	6	14
Paragraphs	2691	90	37	72	88	106	206

Table 2: Number of yes/no questions in Dev, Test and Train sets.

	Size	#Yes	Yes%	#No	No%
Dev set	821	672	81.8	149	18.1
Test set	805	620	77	185	22.9
Train set	1065	805	75.5	260	24.4

Table 3: Question categorization of DaNetQA. Top-10 question topics.

Category	Example	Percent	Yes%
Война (War/Military)	Был ли взят калинин немцами?	10.51	54.7
История (History)	Были ли фамилии у крепостных крестьян?	9.55	87.1
Космос (Space/Galaxy)	Есть ли жизнь на других планетах солнечной системы?	9.10	79.5
СССР (USSR)	Была ли конституция в СССР?	7.61	83.9
Здоровье (Health)	Передаётся ли чумка от кошки к человеку?	7.17	84.9
Семья (Family)	Был ли у гагарина брат?	6.76	81.3
Искусство (Art/Literature)	Существует ли жанр эклектика в живописи?	6.54	75.5
Евросоюз (EU)	Входит ли чехия в евросоюз?	6.28	86.9
География (Geography)	Была ли албания в составе югославии?	6.05	76.6
Связи (Communications)	Был ли лермонтов знаком с пушкиным?	5.42	82.8

the frequency of the used bigrams that the generated questions begin with (see Table 4).

Note that this list of starter phrases is consistent with the one in the [1] pipeline. Indeed, in the paper, authors select questions that begin with a certain set of words (“did”, “do”, “does”, “is”, “are”, “was”, “were”, “have”, “has”, “can”, “could”, “will”, “would”). Our list and the list used in the creation of BoolQ are somewhat similar, if we omit the particle “ли”, which is added in Russian when creating interrogative sentences. The majority of the questions in the corpus are starting with these words.

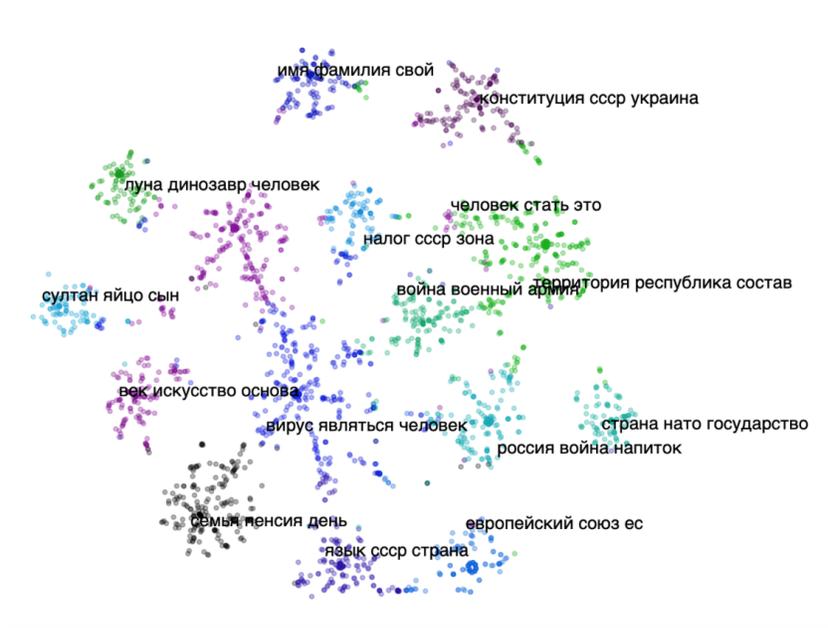


Fig. 1: t-SNE visualization of LDA model trained on concatenated questions and paragraphs, 15 topics.

Table 4: Top-15 frequent bigrams, that start the questions.

Indicator words	Frequency
был ли (was (<i>3rd pers., sing., masc.</i>) there)	575
есть ли (is (<i>3rd pers., sing., masc.</i>) there)	362
была ли (was (<i>3rd pers., sing., fem.</i>) there)	317
были ли (was (<i>3rd pers., plur.</i>) there)	302
входит ли (is included (<i>3rd pers., sing.</i>))	201
едят ли (do [they] eat (<i>3rd pers., plur.</i>))	130
правда ли (is [it] true)	130
разрешено ли (is [it] allowed (<i>3rd pers., sing., neutr.</i>))	74
передаётся ли (is [it] transmitted (<i>3rd pers., sing.</i>))	70
состоит ли (does [it] consist (<i>3rd pers., sing.</i>) in)	68
бывает ли (is (<i>3rd pers., sing., fem.</i>) there)	55
вреден ли (is [it] harmful (<i>3rd pers., sing., masc.</i>))	55
вредна ли (is [it] harmful (<i>3rd pers., sing., fem.</i>))	50
существует ли (does [it] exist (<i>3rd pers., sing.</i>))	41
а была (was (<i>3rd pers., sing., fem.</i>) there)	30

3 Experiments

The DaNetQA task is formulated as a binary classification task. The input to the model consists of the question and paragraph pair. The task is to return

either 0 or 1, so that the positive answer indicates that the answer to the question, based on the information from the paragraph is “yes”. Otherwise, 0 stands for the negative answer. We evaluated two baseline algorithms, to which multiple transferring techniques are compared. As the source for the transfer, we used either other tasks, which have similar setting, or sources in other languages.

3.1 Baseline

We exploit two baseline approaches:

1. FastText for binary classification [10];
2. A deep learning approach: fine-tuning RuBERT model on DaNetQA.

FastText. The FastText classifier trained on the concatenated vectors of questions and paragraphs with the use of pretrained vector representations (taken from the official website) showed slightly higher accuracy, than the one trained only on DaNetQA dataset (Table 5).

BERT. Fine-tuning DeepPavlov RuBERT [11] on the imbalanced DaNetQA, leads to resulting model being uncalibrated. Therefore the predicted probability for the positive class is higher in most cases. To tackle this problem, we aim to find an optimal predicted probability threshold for validation dataset. This is done in two steps:

1. For each threshold t in some discrete subset of range $[0, 1]$, “yes” answer $precision(t)$ and $recall(t)$ calculated;
2. $F_1(t) = \frac{2 \cdot precision(t) \cdot recall(t)}{precision(t) + recall(t)}$ is calculated;
3. Threshold t for the largest $F_1(t)$ is selected.

We train a model for 40 epochs and then find the best model according to $F1$ metric value. The best model is selected by precision-recall-AUC metric value for the “No” answer. For this task, we set learning rate to $3e-5$, and linearly decrease it up to zero to the end of training.

Table 5: Accuracy for binary classification with FastText.

	FastText	pre-trained FastText	RuBERT
ACC	0.81366	0.80745	0.7975

FastText can be seen as a strong baseline, as it manages to achieve results comparable to larger RuBERT model. Further we will refer to RuBERT baseline to compare with transfer learning techniques.

3.2 Task Transferring

To transfer learning from other tasks, we make two steps:

1. Fine-tuning the pretrained transformer model on a similar task;
2. Fine-tuning the model on DaNetQA.

The tasks, used for the first step, are: Paraphraser, the Russian part of XNLI, and the Task A of the SberQUAD. For each task, we initialize model weights with DeepPavlov RuBERT. All three datasets are used independently.

Paraphraser [12] is a dataset for the paraphrasing task: it consists of sentence pairs, each of which is labeled as paraphrase, not paraphrase or maybe paraphrase. This task is close to DaNetQA as the model is required to detect linkage between sentences. We transform this task to binary classification problem, where label is true only for definitely paraphrase.

XNLI [2] is the dataset for the language inference task, translated to 15 languages, including Russian. The task is to predict if one sentence is entailment or contradiction of another, or two sentences are neutral. Each pair has initial annotation and five other crowd workers annotations. For our task, we kept only Russian part of dataset and only sentences where majority label is the same as initial. As for Paraphraser, we transformed the task into binary classification problem, and only entailment is treated as positive class.

SberQUAD task A is the part of Sberbank question answer challenge. In task A, model needs to predict if given text contains answer to given question.

Since each task in the first step is, as DaNetQA, a binary classification problem, we use the same approach with training and selecting the best model, as for fine-tuning BERT on DaNetQA in baseline approach. Best model selection and evaluation on further DaNetQA fine-tuning is also the same as in baseline solution.

The model is fine-tuned on first-step task with initial learning rate $3e-5$ and then fine-tuned on DaNetQA with initial $2e-5$ learning rate, linearly decreasing to zero in each case.

3.3 Language Transferring

Similarly to task transferring, for language transferring we fine-tune transformer on a similar dataset and then further fine-tune it on DaNetQA. The difference is that dataset in the first-step task is not in Russian, so we should train model on it's machine translation or use multilingual pre-trained model. We tested the following configurations:

- Finetune RuBERT on translated BoolQ and test on DaNetQA;
- Finetune RuBERT on translated BoolQ and then finetune and test on DaNetQA;
- Finetune XLM-R on BoolQ and then finetune on DaNetQA.

Hyper parameters and algorithm of the best model selection and evaluation are the same as in baseline approach for fine-tuning RuBERT on DaNetQA.

Intuitively, when pre-trained on BoolQ, which inspired DaNetQA, the model should have relevant knowledge to solve our task. As the BoolQ dataset is in English, we should use some kind of multilingual technique to enable transfer learning. The core methods in this case are:

1. To use machine translation models to translate either BoolQ to Russian, or DaNetQA to English, and fine-tune a monolingual model on BoolQ and further use it for DaNetQA;
2. To use a multi-lingual model and BoolQ for pre-training in a straightforward way.

To translate BoolQ into Russian we used the Helsinki-NLP/opus-mt-en-ru machine translation model from HuggingFace Transformers⁴ library.

XLN-R [8] is a transformer-based model, trained on different languages. This model should capture semantic information in higher layers, so fine-tuning on English data could be good initialization to process further DaNetQA.

4 Results

To evaluate the transfer learning approaches and compare against baselines, we calculated accuracy and F1 scores for “Yes” answer. However, since “Yes” is the most common answer, we should pay more attention to “No” answer, which is less frequent. Therefore we also calculated F1 score for “No” answer. Moreover, we need to understand how sensitive is the model. For this reason, we calculated ROC-AUC, and as the models are trained on imbalanced data, we also calculated precision-recall AUC for both classes. The result values of each configuration are described in Table 6.

Table 6: Results of the task solving approaches.

Transfer dataset	Model	Finetune on DaNetQA	ACC	ROC - AUC	“No” precision -recall AUC
-	RuBERT	Yes	82.46 ± 0.48	75.92 ± 0.38	57.16 ± 2.38
-	XLN-R	Yes	81.44 ± 1.28	67.76 ± 2.53	49.41 ± 3.43
Paraphraser	RuBERT	Yes	82.48 ± 0.53	75.02 ± 0.48	55.22 ± 1.97
XNLI (ru)	RuBERT	Yes	81.89 ± 0.76	78.56 ± 1.99	57.15 ± 2.47
SberQUAD Task A	RuBERT	Yes	82.63 ± 0.88	80.74 ± 1.16	63.02 ± 1.54
BoolQ translated	RuBERT	No	74.04 ± 1.88	62.74 ± 2.61	32.64 ± 2.86
	RuBERT	Yes	82.56 ± 0.29	80.81 ± 1.50	62.49 ± 1.84
BoolQ	XLN-R	No	78.68 ± 0.97	66.53 ± 3.21	42.25 ± 4.65
	XLN-R	Yes	83.20 ± 0.30	79.55 ± 1.63	61.84 ± 1.00

⁴ <https://huggingface.co/transformers/>

Three Russian datasets in the task transferring approach help the model to achieve higher quality than without fine-tuning. Nevertheless, the highest improvement is achieved when transferring from XNLI in line with similar experiments of [1]. As the tasks differ from Yes/No question answering, results are still lower than those, achieved by the language transferring approach.

Our results show, that fine-tuning the XLM-R model on BoolQ (both translated and not) and using the model for further evaluation on DaNetQA, fails in comparison to other techniques. A possible explanation to this is the difference in the topics, covered by both datasets. However, if the model is further fine-tuned on DaNetQA, the superior results are achieved due to the cross-lingual nature of the model.

5 Related work

Binary question answering is a significant part of machine reading comprehension problem. Binary questions are present in the following datasets in English : CoQA [13], QuAC [14], HotPotQA [15] and ShARC [16]. Some of these datasets, in particular, HotPotQA require multi-hop reasoning, for which answering binary questions is crucial. However BoolQ [1] to the best of our knowledge is the only dataset, devoted to binary questions exclusively.

Transfer learning is one of the leading paradigms in natural language processing. It leverages pre-training on a large-scale datasets as a preliminary step before fine-tuning a contextualized encoder for the task under consideration. Applications of transfer learning range from part-of-speech tagging [17] to machine translation [18]. As the language modelling can be seen as independent tasks, some researchers see pre-training with language modelling objective as a part of transfer learning paradigm [19, 20]. Pre-training on natural language understanding tasks, in particular, on sentence modelling tasks, help not only to improve the quality of the task under consideration [1, 21, 22], but also to derive semantically meaningful sentence embeddings that can be compared using cosine-similarity [23].

New datasets for the Russian language are rarely published. Although Russian is one of the most widely spoken languages, the amount of datasets, suitable for NLP studies is quite limited. Dialogue, AIST and AINL conferences are main venues for the Russian datasets to appear. Among recent datasets are RuRED [4], SentiRusColl [24] and a corpus of questions, annotated according to their types.

6 Conclusion

In this paper a new question answering dataset, DaNetQA, is presented. It comprises binary yes/no questions, paired with paragraphs, which should be used to answer the questions. The overall collection procedure follows the design of BoolQ dataset, which is a magnitude larger in size than DaNetQA, partially due to the use of proprietary sources. We establish a straightforward baseline,

exploiting FastText and RuBERT models and experiment with multiple transfer learning settings. Our results show, that on the one hand, the English dataset can be leveraged to improve the results for the Russian one. However, we can not confirm, that we can re-use BoolQ for training the model, while keeping DaNetQA for evaluation only. This brings us to the following conclusion: although the re-creation of English datasets in other languages may seem as a redundant and secondary activity, the current state of the cross-lingual models does not allow of perfect language transfer. It is not enough to train the model on the English data. It seems impossible to gain high quality results, if the model is not trained in the target language. This highlights the need for future development: the development of more advanced cross-lingual contextualized encoders as well as more sophisticated datasets to evaluate cross-lingual tasks. As for DaNetQA development, we plan to enlarge the dataset with more question-paragraph pair and to extend the dataset with unanswerable question, affecting though the task setting.

Acknowledgements. This paper was prepared in the Laboratory for Model and Methods of Computational Pragmatics within the framework of the HSE University Basic Research Program and funded by the Russian Academic Excellence Project '5-100'. The first and the last authors were also supported by Russian Science Foundation under grant XXX at YYY.

References

- [1] Christopher Clark et al. “BoolQ: Exploring the surprising difficulty of natural yes/no questions”. In: *arXiv preprint arXiv:1905.10044* (2019).
- [2] Alexis Conneau et al. “XNLI: Evaluating Cross-lingual Sentence Representations”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2018, pp. 2475–2485.
- [3] Leo Hickey. *The pragmatics of translation*. Vol. 12. Multilingual matters, 1998.
- [4] D. Gordeev et al. “Relation Extraction Dataset for the Russian”. In: *Proceedings of Dialogue*. 2020.
- [5] Yuhao Zhang et al. “Position-aware Attention and Supervised Data Improve Slot Filling”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*. 2017, pp. 35–45. URL: <https://nlp.stanford.edu/pubs/zhang2017tacred.pdf>.
- [6] Alex Wang et al. “SuperGlue: A stickier benchmark for general-purpose language understanding systems”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 3266–3280.
- [7] Jacob Devlin et al. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 2019, pp. 4171–4186.

- [8] Alexis Conneau et al. “Unsupervised cross-lingual representation learning at scale”. In: *arXiv preprint arXiv:1911.02116* (2019).
- [9] Tom Kwiatkowski et al. “Natural questions: a benchmark for question answering research”. In: *Transactions of the Association for Computational Linguistics* 7 (2019), pp. 453–466.
- [10] Armand Joulin et al. “Bag of Tricks for Efficient Text Classification”. In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. 2017, pp. 427–431.
- [11] Yuri Kuratov and Mikhail Arhipov. “Adaptation of deep bidirectional multilingual transformers for russian language”. In: *arXiv preprint arXiv:1905.07213* (2019).
- [12] Ekaterina Pronoza, Elena Yagunova, and Anton Pronoza. “Construction of a Russian paraphrase corpus: unsupervised paraphrase extraction”. In: *Russian Summer School in Information Retrieval*. Springer. 2015, pp. 146–157.
- [13] Siva Reddy, Danqi Chen, and Christopher D Manning. “CoQA: A Conversational Question Answering Challenge”. In: *Transactions of the Association for Computational Linguistics* 7 (2019), pp. 249–266.
- [14] Eunsol Choi et al. “QuAC: Question Answering in Context”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2018, pp. 2174–2184.
- [15] Zhilin Yang et al. “HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering”. In: *EMNLP*. 2018.
- [16] Marzieh Saeidi et al. “Interpretation of Natural Language Rules in Conversational Machine Reading”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2018, pp. 2087–2097.
- [17] Joo-Kyung Kim et al. “Cross-lingual transfer learning for pos tagging without cross-lingual resources”. In: *Proceedings of the 2017 conference on empirical methods in natural language processing*. 2017, pp. 2832–2838.
- [18] Baijun Ji et al. “Cross-Lingual Pre-Training Based Transfer for Zero-Shot Neural Machine Translation”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 01. 2020, pp. 115–122.
- [19] Jeremy Howard and Sebastian Ruder. “Universal Language Model Fine-tuning for Text Classification”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2018, pp. 328–339.
- [20] Sergey Golovanov et al. “Large-scale transfer learning for natural language generation”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2019, pp. 6053–6058.
- [21] Mingyue Shang et al. “Find a Reasonable Ending for Stories: Does Logic Relation Help the Story Cloze Test?” In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 2019, pp. 10031–10032.
- [22] Sanjay Kamath, Brigitte Grau, and Yue Ma. “How to pre-train your model? Comparison of different pre-training models for biomedical ques-

- tion answering”. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2019, pp. 646–660.
- [23] Nils Reimers and Iryna Gurevych. “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 2019, pp. 3973–3983.
- [24] Anastasia Kotelnikova and Evgeny Kotelnikov. “SentiRusColl: Russian Collocation Lexicon for Sentiment Analysis”. In: *Conference on Artificial Intelligence and Natural Language*. Springer. 2019, pp. 18–32.