

# Småprat: DialoGPT for Natural Language Generation of Swedish Dialogue by Transfer Learning

Tosin Adewumi<sup>\*1</sup>, Rickard Brännvall<sup>1,2</sup>, Nosheen Abid<sup>1</sup>, Maryam Pahlavan<sup>1</sup>,  
Sana Sabah Sabry<sup>1</sup>, Foteini Liwicki<sup>1</sup>, and Marcus Liwicki<sup>1</sup>

<sup>1</sup>ML Group, EISLAB, Luleå University of Technology, Sweden

<sup>2</sup>RISE Research Institutes of Sweden

firstname.lastname@ltu.se

## Abstract

Building open-domain conversational systems (or chatbots) that produce convincing responses is a recognized challenge. Recent state-of-the-art (SoTA) transformer-based models for the generation of natural language dialogue have demonstrated impressive performance in simulating human-like, single-turn conversations in English. This work investigates, by an empirical study, the potential for transfer learning of such models to Swedish language. DialoGPT, an English language pre-trained model, is adapted by training on three different Swedish language conversational datasets obtained from publicly available sources: Reddit, Familjeliv and the GDC. Perplexity score (an automated intrinsic metric) and surveys by human evaluation were used to assess the performances of the fine-tuned models. We also compare the DialoGPT experiments with an attention-mechanism-based seq2seq baseline model, trained on the GDC dataset. The results indicate that the capacity for transfer learning can be exploited with considerable success. Human evaluators asked to score the simulated dialogues judged over 57% of the chatbot responses to be human-like for the model trained on the largest (Swedish) dataset. The work agrees with the hypothesis that deep monolingual models learn some abstractions which generalize across languages. We contribute the codes, datasets and model checkpoints and host the demos on the HuggingFace platform.

---

<sup>\*</sup>Corresponding Author: tosin.adewumi@ltu.se

**Keywords:** Conversational Systems, Chatbots, Dialogue, DialoGPT, Swedish.

## 1 Introduction

The introduction of Eliza, the chatbot, in the 1960s marked an epoch in the area of conversational systems [34]. Since then, open-domain conversational systems have evolved [1, 2]. Advances in deep neural networks, such as the transformer-based architectures, have brought improvements to the field [9, 24, 14]. These models have demonstrated SoTA performances in natural language understanding (NLU) and Natural Language Generation (NLG) tasks [32, 13].

The advancements notwithstanding, challenges still exist with building conversational systems [18, 36]. These challenges include technical and ethical challenges[17]. This is more so that many of the models are originally pre-trained on English data [36], though researchers have recently been producing multilingual versions of some of the models [10, 8, 35]. Some of these multilingual models, however, have been shown to have poor performance compared to models trained completely on the target language data [31, 28].

In this work, we perform an empirical study of the performance of one of the recent SoTA models, DialoGPT (medium), on various Swedish datasets of different sizes. Their perplexity results are compared to the one we trained on the English MultiWOZ benchmark dataset. We further compare these experiments with an LSTM-based seq2seq baseline model with the attention mechanism and

trained on the GDC dataset. DialoGPT is an English pre-trained model for open-domain chatbots [36]. We are not familiar with any previous published work that investigates an English pre-trained dialogue model fine-tuned to produce a different target language model. However, there are multilingual models, which are pre-trained on unstructured text of several languages [10, 35]. We thereby investigate how the English pre-trained model performs in NLG (of dialogues) by fine-tuning on a foreign target language.

We contribute the codes<sup>1</sup>, datasets<sup>2</sup> and model checkpoints for public use and host the demos<sup>3</sup> on the HuggingFace platform. The Swedish models are fine-tuned on extracted/crawled datasets. The Swedish language is the official language of Sweden and is spoken by more than 8.5 million people [26]. We show that generation of dialogues is possible, with reasonable performance, for a foreign, target language though the pre-training was in English.

## 2 Related Work

There are a number of pretrained models for open-domain conversational systems. Some of them include Texar [16], DLGnet [21], Meena [3] and BlenderBot [27]. These are pretrained on dialogue datasets. There exist, also, models pretrained on large text and adapted for conversational systems. Examples of such models include T5 [25] and BART [19]. Another pretrained model on conversational data, DialoGPT (dialogue generative pre-trained transformer), was trained on Reddit conversations of 147M exchanges [36]. In single-turn conversations, it achieved performance close to human in open-domain dialogues. DialoGPT is based on GPT-2 [24]. It is an autoregressive model, which achieved SoTA results in different NLP tasks [24].

In a recent work on cross-lingual transferability [5], Artetxe et al (2020) suggest that deep monolingual models learn abstractions that generalize across languages. This is in contrast to past hypothesis that attributes the generalization ability of multilingual models to the shared subword vo-

cabulary used across the languages and joint training, as demonstrated for mBERT [23]. The performance of such multilingual models on low-resource languages and unseen languages are known to be poor [22, 33].

In evaluating the performance of open-domain chatbots, it has been shown that automatic metrics, like the BLEU score, can be very poor but they are still used in some cases [20]. Conversation turns per session is another metric of interest [37]. Perplexity is widely used for intrinsic evaluation of language models in pilot experiments and its theoretical minimum, which is its best value, is 1 [3]. Probably the best evaluation is done by human evaluators (or annotators) but this can be subjective. The judgment of human evaluators is seen as very important, especially since humans are usually the end-users of such systems [36].

## 3 Methodology

We used the DialoGPT medium model with 345M parameters and 24 transformer layers in this work. We chose this because it was reported to have the best performance (compared to its small and big versions) across a set of related tasks [36]. The experiments were carried out on several Tesla V100 GPUs on an Nvidia DGX-1 server running Ubuntu 18. The datasets were split in the ratio 80:10:10 for training, dev and test sets. Multiple runs (5) per experiment were conducted and the average perplexity reported in section 4. Although one automatic metric (perplexity) was used to evaluate the models, it has been shown to correlate with another proposed human evaluation metric called Sensibleness and Specificity Average (SSA) [3]. The conversation context was set as 7 during training. Larger contexts bring memory challenges, hence 7 appears to be a good balance for training [3]. The finetuning process involved adjusting all the parameters of the pretrained model.

Furthermore, we compare the DialoGPT experiments with a reasonable baseline model: a seq2seq model trained on the Swedish GDC dataset. The seq2seq model is based on the LSTM architecture [15] and uses the attention mechanism [6]. The model has 6M trainable parameters. A batch size of 64 is used.

<sup>1</sup>[github.com/tosingithub/gemdesk](https://github.com/tosingithub/gemdesk)

<sup>2</sup>We may provide the datasets or the APIs for extracting them, where applicable.

<sup>3</sup>[huggingface.co/tosin/dialogpt\\_mwoz](https://huggingface.co/tosin/dialogpt_mwoz)  
[huggingface.co/tosin/dialogpt\\_sv](https://huggingface.co/tosin/dialogpt_sv)

### 3.1 Perplexity

Perplexity models the average predictability (i.e. minimizing the uncertainty of predicting the next token). The lower the perplexity, the better the model performs [3]. This is used often to evaluate the language models built with n-grams of text dataset [29]. Perplexity,  $PP$ , calculates the probability  $\rho$  of the test corpus, normalized by the total number of words,  $N$ , in the test corpus  $W_{test}$ . The normalization is done by taking the  $N$ th root of inverse of calculated probability (see Equation 1).

$$PP(W_{test}) = \sqrt[N]{\left(\frac{1}{\rho(W_{test})}\right)} \quad (1)$$

### 3.2 Human Evaluation

In addition to intrinsic evaluation, the ideal model in each category of datasets was evaluated on single-turn conversations by native/near-native Swedish speakers. The authors decided to test single-turn conversations as practiced in the original paper. This is because the lack of long-term contextual information is still an existing problem in conversational systems [36]. A scale of *clearly human-like (4.0)*, *somewhat human-like (3.0)*, *not very human-like (2.0)*, *clearly not human (1.0)* was provided. Similarly to the original DialoGPT work [36], we drew 30 input sentences randomly from the test set for each model and recorded their corresponding responses.

Seven human annotators then scored each conversation turn online, where the turns were assigned randomly from one of the three models or the human ground truth. The annotators were to use their best judgment to decide what they felt was human-like or otherwise on the four grade scale for each conversation. For some further analysis it was decided to also aggregate annotator scores into an aggregate binary score with classes *not human-like (0)* and *human-like (1)* such that inter-annotator agreement could be defined simply as the mean of class agreement for all annotator pairs that scored the same model-question combinations.

The English translation of part of the single-turn conversations of the familjeliv 1M+ model is available in the appendix. They were translated using Google translate and reviewed by a Swedish native

speaker. The original Swedish conversations are also available<sup>1</sup>.

### 3.3 Byte-Pair Encoding

A token vocabulary that includes all common words can become quite large and requires the use of an additional `<unk>` special token for unknown words. Byte-Pair Encoding (BPE) was introduced [30] to address both these shortcomings by first identifying a base vocabulary of tokens consisting of all symbols that occur in the text and then defining merge rules based on frequencies of compound symbols to form new tokens from two tokens of the base vocabulary. The merge process proceeds until the vocabulary has attained a desired fixed size.

GPT-2 uses bytes as the base vocabulary, which forces the base vocabulary to be of size 256 while allowing all unicode characters as well as higher level subword components, basic words and common compounds to be expressed in compressed form. With some additional rules to deal with punctuation, the tokenizer of GPT-2 can handle every text without the need for the `<unk>` symbol. For GPT-2 the merge set consists of 50,000 tokens which combined with the 256 bytes base tokens and a special end-of-text token yields a total vocabulary size of 50,257.

GPT-2's tokenizer trained on English text can thus also express words in Swedish (including words with special Swedish characters å, ä and ö) as it relies on BPE, however, one can expect the compression rate to be lower as the frequency and composition of subwords are different in the two languages. More often, it will have to rely on character level tokens. Indeed, a comparison of the merge vocabularies of the (English) GPT-2 tokenizer and one constructed<sup>4</sup> based on the Swedish language Gothenburg Dialogues Corpus (GDC) shows that only about 8,000 of the compound tokens are shared.

### 3.4 Datasets Used

The authors experimented with various Swedish datasets and the English MultiWOZ. The Swedish datasets are conversational data from Reddit (2 sizes), Familjeliv (3 sizes) and the GDC [4]. The extracted data were pre-processed by removing

<sup>4</sup>by using the BPE tokenizer of the HuggingFace library

emails, URLs, numbers and some special characters. Table 1 summarizes the datasets. Example conversation lines from all the Swedish datasets are available in appendix A.

Table 1: Summary of Datasets

Dataset	File Size	Conversation Lines
Reddit 4K	0.57M	4,300
Reddit 60K	10.4M	59,437
Familjeliv 70K	10.3M	71,470
Familjeliv 400K	45.3M	347,590
Familjeliv 1M+	200M	1,576,360
GDC	6.6M	108,571
MultiWoZ (English)	11M	143,048

### 3.4.1 Reddit

Reddit is a social discussion website with various communities or subreddits<sup>5</sup>. The discussions can be very informal and contain slangs, offensive text and emojis. The Reddit data was constructed with breadth-first search traversal, using the applicable API: PAWN. Two sizes of the data were created and experimented with: the hot 800 topics (with about 60K conversation lines) and hot 50 topics (with over 4K conversation lines) of the Swedish subreddits. The four Swedish subreddits from which data were extracted are *sweden*, *svenskpolitik*, *swedishproblems* and *stockholm*. The sentences in the conversation were given a maximum length of 500 characters.

### 3.4.2 Familjeliv.se dataset

Familjeliv<sup>6</sup> is a popular website forum in Swedish. The name literally translates as family life and mainly targets adults with questions concerning pregnancy, parenthood, and domestic life hacks. It has more than a million daily visitors (2016) with lively user forum participation. The fora were scraped by automated software over the extent of a month to provide a dataset with over a million conversation turns.

### 3.4.3 Gothenburg Dialogue Corpus

The Gothenburg Dialogue Corpus (GDC) is comprised of 360 individual dialogues transcribed from

<sup>5</sup>reddit.com

<sup>6</sup>www.familjeliv.se

recordings of about 25 different social activities, including debates, academic seminars and situations [4]. The content is somewhat different from the other Swedish corpora used in this study as it originates from real speech conversations, and contains casual language, dialect, slang and sometimes broken grammar, typical of spoken language. It counts over 108K dialogue turns with over 1.3M tokens.

### 3.4.4 MultiWOZ

Among the many English conversational datasets available is the MultiWOZ [7]. It is a large, multi-domain and multi-task conversational dataset that has been extensively used since its creation. It consists of more than 10,000 dialogues distributed between 70% multi-domain and 30% single domain dialogues. It has been a standard benchmark for different dialogue problems. There are several versions of the dataset, with each new one bringing improvements [12].

## 4 Results & Discussion

Table 2 shows the mean perplexity results for the various datasets. We observe a trend of decreasing perplexity score with increased dataset - as expected. More data implies improved score. The best perplexity score on the test set is obtained with the MultiWOZ. This is followed by the Familjeliv size of over 1M turns while the Reddit 4K had the worst perplexity of the models trained with DialoGPT. The English MultiWOZ model has a better perplexity compared to any of the Swedish models, though some of the latter have more conversation turns. This observation is not surprising, as the pre-trained model was pretrained in English, though it used the Reddit-style conversation. The seq2seq model had the worst perplexity result overall. This should be partly because it was not pretrained. Its architecture is also not completely comparable to that of DialoGPT.

Table 3 shows results from human evaluation of the presented Swedish single-turn conversations for three of the models, with the last row including scores for the original dialogue (i.e. replies by real humans). The scale is as discussed in the previous section. The overarching target of the human evaluation was to collect a subjective opinion of how

Table 2: Mean perplexity results for the different datasets after training for 3 epochs

Dataset	Dev set	Test set
Reddit 4K	71.94	88.31
Reddit 60K	65.86	51.70
Familjeliv 70K	11.12	12.27
Familjeliv 400K	7.02	7.44
Familjeliv 1M+	7.150	7.148
GDC	29.17	23.95
Seq2seq-GDC	2,864	2,865
MultiWOZ (English)	6.41	6.21

Table 3: Human evaluation score per model (%)

annotator score	1.0	2.0	3.0	4.0
Familjeliv 1M+	0.233	0.194	0.252	0.320
Gothenburg DC	0.317	0.221	0.221	0.240
Reddit 60K	0.582	0.220	0.154	0.044
Human dialogue	0.000	0.079	0.281	0.640

natural the simulated conversations appear to humans. The aggregate *human-like* ratio across the survey is presented in Table 4 where we see that the model trained on the large familjeliv dataset was considered human-like in 57.3% of the assessments. This is when the scores on the four-level scale are collapsed into binary classes for simplicity: *human-like* (assigned value 1.0), which aggregates the *clearly human-like* and *somewhat human-like* assessments into one category, placing the other two assessments in the other class *not human-like* (value 0.0). The inter-annotator agreement scores in table 4 can then be calculated as the mean of binary class agreement across all annotator pairs that scored the same model-question combinations.

Table 4: Aggregate human-likeness scores (%) and annotator agreement for each model

human-likeness	mean score	agreement
Familjeliv 1M+	0.573	0.762
Gothenburg DC	0.462	0.731
Reddit 60K	0.198	0.814
Human dialogue	0.921	0.860

To understand better the variation in the human evaluation scores, a statistical bootstrap [11] exercise (i.e. sampling with replacement) was carried

out and reported in table 5, where we note that the bootstrap mean agree well with the simple survey means for the aggregate (in table 4) with relatively narrow standard deviations and confidence intervals (last column of table 5). The confidence intervals for the GDC and Familjeliv1M+ overlap somewhat, while the model trained on the Reddit60K data shows considerably worse performance. This dataset had a higher occurrence of English dialogues, which may have counteracted the transfer learning into Swedish.

Table 5: Bootstrap means, standard deviations and 95% confidence bands for the human-likeness scores

human-likeness	mean	std	conf95
Familjeliv 1M+	0.571	0.049	0.47-0.66
Gothenburg DC	0.463	0.050	0.37-0.56
Reddit 60K	0.198	0.042	0.12-0.28
Human dialogue	0.921	0.028	0.86-0.97

We also note that the lower percentile for the human dialogue scores are above the higher percentile for the best performing chatbot (also visualised in Figure 1). We must therefore conclude that none of the models achieved human performance. However, this investigation seemingly agrees with the hypothesis that deep monolingual models learn abstractions that generalize across languages, as demonstrated also by [5], even though their experiments are different from those carried out in this work. Indeed, in this work, less computational effort was needed to demonstrate this.

Demos of the English and Swedish chatbots are hosted on the HuggingFace platform<sup>3</sup>, including the model checkpoints. It should be pointed out that there are risks with using the models, such as producing dialogue texts that contain unwanted bias, stereotypes or offensive language. This is a well-known, but hard to avoid problem due to the difficulty of removing such material from the underlying sources of data which the models are trained on [36].

## 5 Conclusion

In this work, we show through an empirical study, using the recent SoTA model DialoGPT-medium [36], that a model pre-trained in English can be




Figure 1: Bootstrap histograms for human-likeness scores

adapted by fine-tuning on a foreign language corpus such that the performance of NLG of dialogues for the target language can be good. The foreign, target language in this study is Swedish. The investigation was carried out using 3 differently-sourced Swedish datasets. The human evaluators asked to score the simulated single-turn conversations judged that about 57% of the chatbot responses are human-like for the model trained on the largest Swedish dataset.

Despite the advancements in deep neural networks and the demonstration of SoTA performances in NLG tasks [32, 13] by a model such as DialoGPT, challenges still exist [36]. Hence, future work may involve attempts to understand the mathematics of languages/linguistics and their relatedness, comparison of the performance observed here with a model pretrained on Swedish data and the possibility of using a Swedish tokenizer on the model pretrained on Swedish data. It may also involve finding better decoding schemes than existing ones.

## References

- [1] T. P. Adewumi, F. Liwicki, and M. Liwicki. Conversational systems in machine learning from the point of view of the philosophy of science—using alime chat and related studies. *Philosophies*, 4(3):41, 2019. doi: 10.3390/philosophies4030041.
- [2] T. P. Adewumi, F. Liwicki, and M. Liwicki. Vector representations of idioms in chatbots. In *32nd annual workshop of the Swedish Artificial Intelligence Society (SAIS 2020), 16-17 June, 2020, Gothenburg, Sweden (Online)*. Chalmers University of Technology, 2020.
- [3] D. Adiwardana, M.-T. Luong, D. R. So, J. Hall, N. Fiedel, R. Thoppilan, Z. Yang, A. Kulshreshtha, G. Nemade, Y. Lu, et al. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*, 2020. doi: 10.48550/arXiv.2001.09977.
- [4] J. Allwood, L. Grönqvist, E. Ahlsén, and M. Gunnarsson. Annotations and tools for an activity based spoken language corpus. In *Current and new directions in discourse and dialogue*, pages 1–18. Springer, 2003. doi: 10.1007/978-94-010-0019-2\_1.
- [5] M. Artetxe, S. Ruder, and D. Yogatama. On the cross-lingual transferability of monolingual representations. In *Proceedings of the*

- 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.421. URL <https://aclanthology.org/2020.acl-main.421>.
- [6] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations, ICLR 2015*, 2015. doi: 10.48550/arXiv.1409.0473. URL <https://arxiv.org/pdf/1409.0473.pdf>.
- [7] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, and M. Gašić. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium, Oct.–Nov. 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1547. URL <https://aclanthology.org/D18-1547>.
- [8] A. CONNEAU and G. Lample. Cross-lingual language model pretraining. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. doi: 10.48550/arXiv.1901.07291. URL <https://proceedings.neurips.cc/paper/2019/file/c04c19c2c2474dbf5f7ac4372c5b9af1-Paper.pdf>.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. doi: 10.48550/arXiv.1810.04805.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Multilingual bert. 2018.
- [11] B. Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1–26, 1979. doi: 10.1214/aos/1176344552.
- [12] M. Eric, R. Goel, S. Paul, A. Sethi, S. Agarwal, S. Gao, A. Kumar, A. Goyal, P. Ku, and D. Hakkani-Tur. Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 422–428, Marseille, France, May 2020. European Language Resources Association. URL <https://www.aclweb.org/anthology/2020.lrec-1.53>.
- [13] S. Gehrmann, T. Adewumi, K. Aggarwal, P. S. Ammanamanchi, A. Anuoluwapo, A. Bosse-lut, K. R. Chandu, M. Clinciu, D. Das, K. D. Dhole, et al. The gem benchmark: Natural language generation, its evaluation and metrics. pages 96–120, Aug. 2021. doi: 10.18653/v1/2021.gem-1.10. URL <https://aclanthology.org/2021.gem-1.10>.
- [14] P. He, X. Liu, J. Gao, and W. Chen. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*, 2020. doi: 10.48550/arXiv.2006.03654.
- [15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.
- [16] Z. Hu, H. Shi, B. Tan, W. Wang, Z. Yang, T. Zhao, J. He, L. Qin, D. Wang, X. Ma, et al. Texar: A modularized, versatile, and extensible toolkit for text generation. *arXiv preprint arXiv:1809.00794*, 2018. doi: 10.48550/arXiv.1809.00794.
- [17] S. Javed, T. P. Adewumi, F. S. Liwicki, and M. Liwicki. Understanding the role of objectivity in machine learning and research evaluation. *Philosophies*, 6(1):22, 2021. doi: 10.3390/philosophies6010022.
- [18] D. Jurafsky and J. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Dorling Kindersley Pvt, Limited, 2020. ISBN 9789332518414. URL <https://books.google.se/books?id=ZalcjwEACAAJ>.
- [19] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language

- generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703. URL <https://aclanthology.org/2020.acl-main.703>.
- [20] M. Lundell Vinkler and P. Yu. Conversational chatbots with memory-based question and answer generation, 2020.
- [21] O. Olabiyi and E. T. Mueller. Multiturn dialogue response generation with autoregressive transformer models. *arXiv preprint arXiv:1908.01841*, 2019.
- [22] J. Pfeiffer, A. Rücklé, C. Poth, A. Kamath, I. Vulić, S. Ruder, K. Cho, and I. Gurevych. Adapterhub: A framework for adapting transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020): Systems Demonstrations*, pages 46–54, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.7. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.7>.
- [23] T. Pires, E. Schlinger, and D. Garrette. How multilingual is multilingual BERT? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1493. URL <https://aclanthology.org/P19-1493>.
- [24] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- [25] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- [26] M. Reuter. Swedish as a pluricentric language. *Pluricentric languages. Differing norms in different nations*, pages 101–116, 1992.
- [27] S. Roller, E. Dinan, N. Goyal, D. Ju, M. Williamson, Y. Liu, J. Xu, M. Ott, K. Shuster, E. M. Smith, et al. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*, 2020. doi: 10.48550/arXiv.2004.13637.
- [28] S. Rönqvist, J. Kanerva, T. Salakoski, and F. Ginter. Is multilingual bert fluent in language generation? *arXiv preprint arXiv:1910.03806*, 2019. doi: 10.48550/arXiv.1910.03806.
- [29] R. Sennrich. Perplexity minimization for translation model domain adaptation in statistical machine translation. Association For Computational Linguistics, 2012. doi: 10.5167/uzh-61712.
- [30] R. Sennrich, B. Haddow, and A. Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2016. doi: 10.18653/v1/p16-1162.
- [31] A. Virtanen, J. Kanerva, R. Ilo, J. Luoma, J. Luotolahti, T. Salakoski, F. Ginter, and S. Pyysalo. Multilingual is not enough: Bert for finnish. *arXiv preprint arXiv:1912.07076*, 2019. doi: 10.48550/arXiv.1912.07076.
- [32] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *arXiv preprint arXiv:1905.00537*, 2019.
- [33] Z. Wang, A. W. Yu, O. Firat, and Y. Cao. Towards zero-label language learning. *arXiv preprint arXiv:2109.09193*, 2021. doi: 10.48550/arXiv.2109.09193.
- [34] J. Weizenbaum. A computer program for the study of natural language. *Fonte: Stanford: http://web.stanford.edu/class/linguist238/p36*, 1969.
- [35] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, and



- C. Raffel. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.41. URL <https://aclanthology.org/2021.naacl-main.41>.
- [36] Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and B. Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, 2020. doi: 10.48550/arXiv.1911.00536.
- [37] L. Zhou, J. Gao, D. Li, and H.-Y. Shum. The design and implementation of xiaoice, an empathetic social chatbot. *Computational Linguistics*, 46(1):53–93, 2020. doi: 10.1162/coli.a\_00368.

## Acronyms

**GDC** Gothenburg Dialogue Corpus. 4, 10

**NLG** Natural Language Generation. 1, 2, 6

**NLP** Natural Language Processing. 2

**NLU** natural language understanding. 1

**SoTA** state-of-the-art. 1, 2, 5, 6

## Appendix

### A - Examples from the various datasets

#### Reddit examples

”Om man tycker att munskydd är fånigt, tycker man då att ansiktsmask är ansikt-slarv?”

”Du har en liten poäng, kylkedjan gör att

det kräver mer planering och det går lite långsammare. Men det stora problemet med dessa vaccin är inte här i Sverige utan i länder som är fattiga och inte har”

”\*Det kan jämföras med andra europeiska länder som övervakat förekomsten och som haft en tidigare spridning av denna variant där förekomsten utgör upp till procent av fallen.\* Det kan ju också jämf”

”Det är nu en kamp mellan vaccineringen och brittisk covid, får vi en tredje våg eller inte.”

”Mycket lovande siffror, ska bli intressant att se SÅBO siffror om ett par veckor, det borde vara mkt låga antal då.”

#### 5.0.1 Familjeliv examples on the proper use of the word mother’s day.

”Varför? Det heter MORS DAG och FARS DAG, inte morsdag och farsdag. Det uttalas isär, inte ihopsatta som torsdag . Så varför detta envisa ihopskrivande? Folk särskriver till höger och vänster men skriver ihop sånt som inte ska skrivas ihop. Varför?!Leo”,

”Nja, man får ju skapa nya sammansatta ord på svenska så morsdag är en acceptabel konstruktion. Det är alltså inte dialekt utan ett sammansatt ord som verkar vara vanligare i vissa delar av landet.”,

”Ja. Jag är också för dialekter. Men jag är inte för att skriva på dialekt. Skriftspråket är likadant var än i Sverige du befinner dig.”,

”Ja, precis. Inläggen från Butterscotch och Pandamma79 var faktiskt ganska dryga. Då kan jag lika gärna hävda att de som säger mors dag säger fel . Det var inte meningen att vara dryg - jag menade bara att det faktiskt är bra att man får lära sig nya saker ibland.”,

”Morsdag kanske kan bli ett nytt ord, men med någon annan innebörd då? Mors dag finns ju som etablerat begrepp (hur man uttalar det lägger jag mig inte i, men i skrift tycker jag att det är rimligt att alla skriver likadant).”,

”Ok, tack ska du ha, jag ber om ursäkt för missförståndet och att jag blev lite putt i så fall. Men nu är ju morsdag inte ett nytt ord, det har använts mycket länge. Jag tycker att det inte spelar så stor roll om namnen inte är identiska i olika delar av landet.”,

”Jag önskar alltså inte att alla i hela Sverige ska säga och skriva morsdag, men där begreppet är etablerat och till och med självklart finns det ingen anledning att ändra. Dessutom: För mig är morsdag INTE mors dag .”,

### 5.0.2 GDC examples

”ja men vi hade haft fest så vi var väl lite i i gasen så ja det ska jag göra sade jag”

”sade du det”

”... ja det sade jag så sen sen satte jag på den där skivan det var ett sådant där ackompanjemang till den där vi talade om den där biten när man har ... san francisco”

”...”

”ja”

”apropå san francisco så fick vi fick vi den hör skivan så gick vi upp och så sjöng jag till så så spelade vi upp den sen ja det lät lite knasigt men det var så löjligt men eh det är en tjusig melodi”

”ja”

”ja men det är en tjusig stad också”

”ja det är nog det va”

”ja”

”då har jag eh tror jag äntligen fått det hela och fungera”

”nejdå det var trevligt”

## B - Conversations with the Familjeliv 1M+ model

Prompt: We contacted the principal and she said they didn't do that. Although our principal doesn't seem very competent, surely... Maybe we should contact the school counselor ourselves...

BOT: I've spoken to the headmaster's boss, but she doesn't have any free hours at the school. I'll try to reach her on Monday.

Prompt: It turns out that the girl smeared powder into her son's hair and painted him in the face with mascara.

BOT: Then he got angry and shouted that he was going to report this to the police.

Prompt: All rear-facing chairs are safer than turning one over for a young child. And then they've gone through the same tests to be sold.

BOT: the only thing they do anyway is drive a chair. Is it really worth it?

Prompt: Probably wouldn't have been able to keep your mouth shut in front of them either. It's okay to get angry with both the parents and their children when she hurts your child.

BOT: I love my children and do everything I can to make them feel good and have a good time.

Prompt: So I think merino wool can work. I think the usual wool itched a little, because the baby wasn't completely satisfied before she fell asleep...

BOT: I'm also afraid it's going to be too much mess on the clothes, having a hard time buying second-hand and can't really afford to castrate them.

Prompt: I have worked in healthcare for 10 years and have worked 7 Christmases so far. So now it's my turn to be free.

BOT: The first year I have only been on sick leave for 3 weeks and it was changed to 3 days a week. But I don't have a job to go back to.

Prompt: I think it's perfectly OK to let your child become a child model provided the child enjoys the camera and thinks it's fun.

BOT: We've had cameras that have been known to be good, and it's worked great.

Prompt: Double-check with the social insurance fund so you don't get caught for benefit fraud.

BOT: Thanks for the tip! I have already been in contact with fk and they would send me forms that I will submit.

Prompt: I'm afraid if he goes over, he's going to be that flamboyant again who's going to assert himself with new classmates, and then he's safe in the class he's in. Anyone have any wise thoughts on this?

BOT: He loves being in kindergarten and therefore enjoys it. However, he eats poorly and does not want to eat anything other than water and juice.

Prompt: In many countries, babies are given to e.g. day mother, grandmother, aunt... because parental insurance is not as generous as ours.

BOT: Then even if it were still good ownership, it increases the cost of a single person to make ends meet. YOU: so we can go through it together ... BOT: but I still want to look around a little before I get into it, to see what to expect.

Prompt: there are no parking facilities where as much as I know

BOT: You have no obligation to report anything. It is not illegal to park in a family car park.