

# Towards a general purpose machine translation system for Sranantongo

Just Zwennicker<sup>\*†</sup> David Stap  
Language Technology Lab  
University of Amsterdam

## Abstract

Machine translation for Sranantongo (Sranan, srn), a low-resource Creole language spoken predominantly in Surinam, is virgin territory. In this study we create a general purpose machine translation system for srn. In order to facilitate this research, we introduce the SRNcorpus, a collection of parallel Dutch (nl) to srn and monolingual srn data. We experiment with a wide range of proven machine translation methods. Our results demonstrate a strong baseline machine translation system for srn.

## 1 Introduction

The official language in Surinam is Dutch, however the language that you will most likely hear on the streets of Paramaribo, its capital city, is Sranan. It is a lingua franca that finds its origins in the 17th century when the transatlantic slave trade brought people of different cultures and languages together. This Creole language broke down the language barrier between the different social groups.

There are approximately 600k speakers of Sranan worldwide (Vissio and Zakharov, 2021), most of which live in Surinam, and in the Netherlands, their former colonizer. Around 1975, when Surinam became independent, a large group of people emigrated from Surinam to the Netherlands. That so called first generation is mostly fluent in Sranan, but this is often not the case for the second and third generation. A machine translation (MT) system for nl/srn could assist with (re)learning the language (Lent et al., 2022) and facilitate familiarizing themselves with their cultural background.

Although Sranan is the second largest language spoken in Surinam there are relatively few written sources available. This could in part be explained by the stigmatization of the language, as is the case with many Creole languages (Lent et al., 2022).

## 2 Method, setup and results

**SRNcorpus** We have collected data from various domains to create the SRNcorpus. It consists of

parallel nl-srn as well as monolingual srn data. The largest data source contains religious data, in particular the Jehova Witness bible translations (Željko Agic and Vulic, 2020)<sup>1</sup>. Furthermore we scraped around 3k non-religious, parallel sentences from an online Sranan dictionary<sup>2</sup>. For numerous words listed in the dictionary, it contains an example sentence showing how it could be used, along with the Dutch translation of that sentence. This website also contained links to various (children) stories<sup>3</sup> in Sranan, yielding over 6.5k monolingual sentences. Lastly the SRNcorpus contains 5 smaller parallel sources from various domains. See Table 1 for a detailed description of its contents.

While religious data is over-represented, our goal is to create a general purpose translation system. We therefore create non-religious validation and test sets, by sampling two times 256 parallel sentences from the non-religious sources within SRNcorpus to create test and validation sets.

**Other data** In addition to SRNcorpus, we used the following parallel data for our multilingual and transfer learning experiments: Wikimatrix (Schwenk et al., 2021) (nl-English (en) 3.3M; nl-Portugese (pt) 1.2M), TED2020 (Reimers and Gurevych, 2020) (nl-en 317K; nl-pt 260K) and Europarl (Koehn, 2005) (nl-en 2M; nl-pt 1.9M). We used pt since srn inherited features from this language (Seuren, 1981), and similarity between parent and child languages can stimulate transfer (Johnson et al., 2017).

**Domain temperature sampling** Temperature sampling is often used to overcome size differences between language pairs by oversampling lower resource pairs. In contrast, we applied temperature sampling with the goal of reducing domain imbalance.

<sup>1</sup>Unfortunately the JW300 became unavailable during our research, reportedly due to copyright issues. This limited the options for our multilingual setup, as we were not able to secure data for other language pairs containing Sranan

<sup>2</sup>SIL parallel sentences

<sup>3</sup>SIL stories

source	type	language(s)	Domain	#sentences
JW300	parallel	nl-srn	religious	307,866
SIL	parallel	nl-srn	general	2,927
Z-Library-1	parallel	nl-srn	stories	351
Z-Library-2	parallel	nl-srn	stories	163
Z-Library-3	parallel	nl-srn	stories	399
Naks Sranan fb	parallel	nl-srn	general	62
Dutch DOJ	parallel	nl-srn	legal	220
SIL	monolingual	srn	stories	6,572
<b>total:</b>				<b>318,560</b>

Table 1: Contents of SRNcorpus. The Summer Institute of Linguistics (SIL) sentences are scraped from an online Sranan dictionary. Z-library 1, 2 and 3 are authored by Tori di switi fu leisi / H.C. Tiendalli, Lafu tori / J. Redan and Dri Anansi tori / H.C. Tiendalli respectively. Naks Sranan is the facebook page of a cultural organisation and contains member introductions. Finally, the Dutch DOJ data contains warrants for arrest.

ance. We distinguish between religious vs non-religious data within SRNcorpus.

**Experimental setup** First we experiment with bilingual models (+ backtranslation). Then we investigate transfer learning, where we first train a parent model and then finetune on a child model. Finally we train multilingual models by sharing all parameters between languages and prepending a target token (Johnson et al., 2017). For all experiments we used the Transformer base architecture (Vaswani et al., 2017). We implemented all our models using JoeyNMT (Kreutzer et al., 2019). For evaluation we calculated BLEU scores using sacrebleu (Post, 2018).

**Results** See Table 2 for an overview of our results.

**Bilingual** We found that for both nl-srn and srn-nl, using SRNcorpus instead of JW300 increased BLEU scores by 10+ points. 1000 BPE merge operations produced best results, which is in line with (Gowda and May, 2020) who studied the effect of the number of BPE merge operations in relation to the parallel corpus size on translation performance. For nl-srn best results were obtained with  $T = 2$ , whereas for srn-nl  $T = 3$  worked best.

**Backtranslation** For our backtranslation experiments we used the best performing bilingual srn-nl model (BPE=1000;  $T = 3$ ) to translate the monolingual data from SRNcorpus into nl. We then trained models in both directions using this synthetic generated data on top of our SRNcorpus. For BPE and  $T$  we used the best values according to the bilingual models. We found a +1.4 BLEU increase for nl-srn, scoring highest overall for this direction, while our srn-nl model decreased slightly. In addition we experimented with applying domain temperature sampling to the synthetic parallel data as well. We found this to hurt performance for all translation directions.

type	$T$	BPE	nl-srn	srn-nl
b1	-	1000	22.06	15.04
SRN	3	3000	-	28.00
SRN	2	4000	35.48	-
SRN	2	3000	36.78	-
SRN	2	2000	36.86	-
SRN	2	1000	37.48	27.47
SRN	2	500	36.60	-
SRN	1	1000	32.92	24.93
SRN	3	1000	36.99	28.47
bt	2	1000	<b>38.88</b>	-
bt*	2	1000	37.29	-
bt	3	1000	-	27.41
bt*	3	1000	-	27.87
t1	2	32000	37.33	-
t1	2	16000	37.49	-
t1	2	8000	38.85	-
t1	3	32000	-	<b>32.02</b>
t1	3	16000	-	30.78
t1	3	8000	-	29.70
m1	2	32000	30.04	-
m1	2	16000	29.88	23.02
m1	2	8000	29.00	-
m1	3	32000	-	21.71

Table 2: BLEU scores reported on the non-religious test set of the SRNcorpus.  $T$  is domain sampling temperature, BPE is the number of merge ops. b1 is baseline trained using only religious JW300 data. SRN are bilingual models trained on SRNcorpus. bt are backtranslation experiments (\* indicates inclusion of synthetic data). t1 are transfer learning experiments. m1 are multilingual experiments.

**Transfer Learning** Since Sranan is an English based Creole language, the language pair nl-en is a natural choice as a parent model. After training the parent model until convergence, we initialized our child models with the resulting parameters. We increased the number of merge ops to accommodate for the increased vocabulary. For nl-srn, BPE=8k resulted in best translation quality, almost on par with the backtranslation experiment. For srn-nl, best results were achieved with BPE=32k, resulting in the highest BLEU score for this direction.

**Multilingual** For our multilingual models we used our SRNcorpus plus other data as described in the previous Section. We report scores for nl-srn (which are obtained by training a one-to-many model on nl-en, nl-pt and nl-srn) and srn-nl (which are obtained by training a many-to-one model on the reverse directions). Note that we applied language temperature sampling ( $T = 5$ ) to oversample nl-srn and nl-pt. We found that resulting models perform substantially worse compared to our other models (except the baseline).

### 3 Conclusion

In this study we have put NMT for Sranan on the map. We introduced the SRNcorpus and used it for various experiments in search for a performant gen-

eral purpose machine translation for  $nl \rightarrow srn$  and  $srn \rightarrow nl$ . Our results demonstrate a strong baseline machine translation system for Sranantongo, which future work can build on.

Željko Agić and Ivan Vulić. 2020. [Jw300: A wide-coverage parallel corpus for low-resource languages.](#)

## References

- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation.](#)
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation.](#) *Transactions of the Association for Computational Linguistics*, 5.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation.](#) In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Julia Kreutzer, Joost Bastings, and Stefan Riezler. 2019. [Joey nmt: A minimalist nmt toolkit for novices.](#)
- Heather Lent, Kelechi Ogueji, Miryam de Lhoneux, Orevaoghene Ahia, and Anders Søgaard. 2022. [What a creole wants, what a creole needs.](#)
- Matt Post. 2018. [A call for clarity in reporting BLEU scores.](#) In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia.](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Pieter A.M. Seuren. 1981. [Tense and aspect in sranan.](#) *Linguistics*, 19.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. volume 2017-December.
- Nicolás Cortegoso Vissio and Viktor Zakharov. 2021. [Towards a part -of-speech tagger for sranan tongo.](#) *International Journal of Open Information Technologies ISSN: 2307-8162*, 9.