

Grapheme to Phoneme models for Norwegian Bokmål

Version: 2.0 (2024-02-09)

This repo contains G2P models for Norwegian bokmål^[^1], which produce phonemic transcriptions for *close-to-spoken* pronunciations (such as in spontaneous conversations: **spoken**) and *close-to-written* pronunciations (such as when reading text aloud: **written**) for 5 different dialect areas:

1. East Norwegian (**e**)
2. South West Norwegian (**sw**)
3. West Norwegian (**w**)
4. Central Norwegian (Trøndersk) (**t**)
5. North Norwegian (**n**)

[^1]: Bokmål is the most widely used written standard for Norwegian. The other written standard is Nynorsk. Read more on [Wikipedia](#).

Setup

Follow installation instructions from [Phonetisaurus](#). You only need the steps "Next grab and install OpenFst-1.7.2" and "Checkout the latest Phonetisaurus from master and compile without bindings".

Data

The pronunciation lexica that were used to train the G2P-models are free to download and use from Språkbanken's resource catalogue: [NB Uttale](#)

For more information about the lexica, see the Github repo: [Sprakbanken/nb_uttale](#)

Content

- **models/**: contains the models, as well as auxiliary files used by Phonetisaurus
 - **nb_*.fst**: model files to use with **phonetisaurus-apply**. The expansion of ***** is a string of a dialect and pronunciation style, e.g. **e_spoken** or **t_written**.
 - **nb_*.o8.arpa**: 8-gram-models for phoneme sequences that Phonetisaurus uses during training.
 - **nb_*.corpus**: aligned graphemes and phonemes from the lexica.
- **data/**: contains various lexica used for training and testing, including predictions from the models on the test set
 - **NB-uttale*_train.dict**: training data for **models/nb_*.fst**. Each file contains 543 495 word-transcription pairs (WTP), and makes up 80% of all unique WTPs in the lexicon.
 - **NB-uttale*_test.dict**: test data for **models/nb_*.fst**. Each file contains the remaining 20% of the WTPs in the lexicon, i.e. 135 787 WTPs.
 - **predicted_nb_*.dict**: The words from the testdata with the model's predicted transcriptions.

- `wordlist_test.txt`: The orthographic words from the test data, which the models predict transcriptions for.
- `evaluate.py`: script to evaluate the models. The method for calculating WER and PER were re-implemented.
- `g2p_stats.py`: script to evaluate the models from V1.0, which can be used to compare results between these models and the NDT models (with and without tone and stress markers) from version 1.
- **LICENSE**: The license text for CC0, which this resource is distributed with.

Usage

```
phonetisaurus-apply --model models/nb_e_spoken.fst --word_list
data/wordlist_test.txt -n 1 -v > output.txt
```

- Input data (`--word_list`) should be a list of newline-delimited words. See the file `data/wordlist_test.txt` for an example.
- The trained G2P-models are `.fst` files located in the `models` folder. The same folder also contains aligned `.corpus` files and phoneme 8-gram models (`.arpa` files), also from the **Phonetisaurus** training process.
- `-n` lets you adjust the number of most probable predictions.

Evaluation

There are 2 scripts to calculate WER and PER statistics, which give slightly different results.

`evaluate.py`

Calculates stats for all the provided models by default. You can give a pronunciation variant (e.g. `-l e_spoken`) to calculate stats for specific models.

- The WER score is calculated as the count of all mismatching transcriptions (1 error = 1 mismatching word) divided by the count of all words in the reference, i.e. a `*_test.dict` file.
- PER is calculated as the count of all errors (1 error = a mismatching phoneme) divided by the total count of phonemes in the reference file.

```
python evaluate.py
```

Model	Word Error Rate	Phoneme Error Rate
<code>nb_e_written.fst</code>	13.661238654564869	1.9681178920293207
<code>nb_e_spoken.fst</code>	13.72501038144391	1.9832518286152074
<code>nb_sw_written.fst</code>	13.240048644480037	1.8396612197218096
<code>nb_sw_spoken.fst</code>	16.422702734768936	2.426312206336983

Model	Word Error Rate	Phoneme Error Rate
<i>nb_w_written.fst</i>	13.240048644480037	1.8396612197218096
<i>nb_w_spoken.fst</i>	16.892833837574894	2.5064155890730686
<i>nb_t_written.fst</i>	13.736133357062347	1.98774986044724
<i>nb_t_spoken.fst</i>	16.47992288013051	2.5809178688066843
<i>nb_n_written.fst</i>	13.736133357062347	1.98774986044724
<i>nb_n_spoken.fst</i>	17.22590930999963	2.8209779677747715

g2p_stats.py

Calculates WER and PER for two input files.

1. The reference file (e.g. `data/NB-uttale_e_spoken_test.dict`)
 2. The model prediction file (e.g. `output.txt` from the command in [Usage](#), or `data/predicted_nb_e_spoken.dict`).
- The WER score is calculated as the count of errors (1 error = 1 mismatching word) divided by the count of all words in the predictions, i.e. a `predicted_*.dict` file.
 - PER is calculated as the sum of phone error rates for each transcription, divided by the total count of words in the predictions.

NOTE: This method doesn't take transcription lengths into account, so a transcription with 2 phonemes where 1 is wrong has a 0.5 PER while a word with length 10 with 1 error has a 0.1 PER, and the average score for the two words would be 0.35.

```
python g2p_stats.py data/NB-uttale_e_spoken_test.dict
data/predicted_nb_e_spoken.dict
# WER: 14.049209423582523
# PER: 2.714882650391985
```

Model	Word Error Rate	Phoneme Error Rate
<i>nb_e_written.fst</i>	13.97114598599277	2.7038190765903214
<i>nb_e_spoken.fst</i>	14.049209423582523	2.714882650391985
<i>nb_sw_written.fst</i>	13.541060631724685	2.5423757844377284
<i>nb_sw_spoken.fst</i>	16.729141964989285	3.34063477772742
<i>nb_w_written.fst</i>	13.541060631724685	2.5423757844377284
<i>nb_w_spoken.fst</i>	17.186475877661337	3.4137304874392114
<i>nb_t_written.fst</i>	14.059519688924565	2.7190289235234104
<i>nb_t_spoken.fst</i>	--	--

Model	Word Error Rate	Phoneme Error Rate
<i>nb_n_written.fst</i>	14.059519688924565	2.7190289235234104
<i>nb_n_spoken.fst</i>	--	--

NOTE: *The t_spoken and n_spoken model predictions are not the same length as the reference file, which causes the script to exit.*

Transcription standard

The G2P models have been trained on the NoFAbet transcription standard which is easier to read by humans than X-SAMPA. NoFAbet is in part based on [2-letter ARPAbet](#) and is made by [Nate Young](#) for the National Library of Norway in connection with the development of [NoFA](#), a forced aligner for Norwegian. The equivalence table below contains X-SAMPA, IPA and NoFAbet notations.

X-SAMPA-IPA-NoFAbet equivalence table

X-SAMPA	IPA	NoFAbet	Example
A:	ɑ:	AA0	bad
{:	æ:	AE0	vær
{	æ	AEH0	vært
{*I	æɪ	AEJ0	sei
E*u0	æu	AEW0	sau
A	ɑ	AH0	hatt
A*I	ɑɪ	AJ0	kai
@	ə	AX0	behage
b	b	B	bil
d	d	D	dag
e:	e:	EE0	lek
E	ɛ	EH0	penn
f	f	F	fin
g	g	G	gul
h	h	H	hes
I	ɪ	IH0	sitt
i:	i:	I0	vin
j	j	J	ja
k	k	K	kost

X-SAMPA	IPA	NoFAbet	Example
C	ç	KJ	kino
l	l	L	land
l=	l̥	LX0	
m	m	M	man
n	n	N	nord
N	ŋ	NG	eng
n=	ŋ̥	NX0	
o:	o:	OA0	rå
O	ɔ	OAH0	gått
2:	ø:	OE0	løk
9	œ	OEH0	høst
9*Y	œy	OEJ0	køye
U	u	OH0	f*ort
O*Y	ɔy	OJ0	konvoy
u:	u:	OO0	bod
@U	oʊ	OU0	show
p	p	P	pil
r	r	R	rose
d`	d̥	RD	rekord
l`	l̥	RL	perle
l`=	l̥̥	RLX0	
n`	ŋ̥	RN	barn
n`=	ŋ̥̥	RNX0	
s`	ʃ̥	RS	pers
t`	t̥	RT	stort
s	s	S	sil
S	ʃ	SJ	sju
t	t	T	tid
u0	ʊ	UH0	russ
}:	ʊ:	UU0	hus

X-SAMPA	IPA	NoFAbet	Example
v	ʊ	V	v ase
w	w	W	W ashington
Y	y	YH0	ny tt
y:	y:	YY0	ny

Unstressed syllables are marked with a 0 after the vowel or syllabic consonant. The nucleus is marked with a 1 for tone 1 and a 2 for tone 2. Secondary stress is marked with 3.

License

These models are shared with a [Creative Commons-ZERO \(CC-ZERO\)](#) license, and so are the lexica they are trained on. The models can be used for any purpose, as long as it is compliant with Phonetisaurus' license.