Who am I?

Rob van der Goot

Bla bla bla ...

Frisian

Dutch

Danish

Chinese

English

ROBustness in NLP

Lexical normalization

Multi-task learning

Is X solved?

Problem

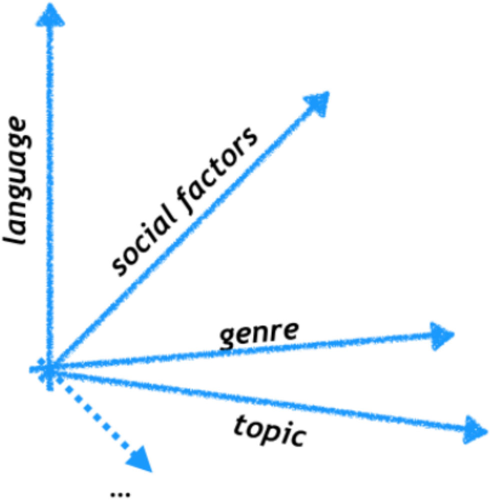

From: Barbara Plank. What to do about non-standard (or non-canonical) language in NLP.

Problem

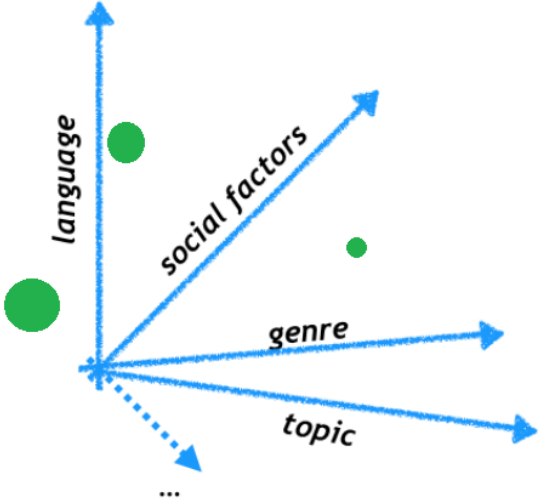

1. Lexical Normalization

u	hve	to	let	ppl	decide	what	dey	want	to	do
you	have	to	let	people	decide	what	they	want	to	do

Lexical Normalization

Situation in 2015:

- Some benchmarks for English: main one LexNorm
- Many models assume gold detection
- Some people working on their own languages
- Differences in models, task definitions and metrics

Lexical Normalization

Situation in 2015:

- Some benchmarks for English: main one LexNorm
- Many models assume gold detection
- Some people working on their own languages
- Differences in models, task definitions and metrics

MoNoise

- First multi-lingual normalization model
- SOTA wherever evaluated
- Outputs top-n; successfully integrated in syntactic parsers.

Parsing

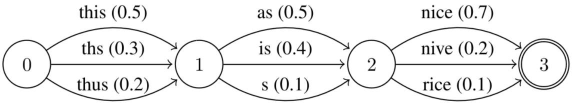

Figure 1: A possible output of the normalization model for the sentence 'ths s nice'.

MultiLexNorm: A Shared Task on Multilingual Lexical Normalization

Rob van der Goot, Alan Ramponi, Arkaitz Zubiaga, Barbara Plank, Benjamin Muller, Iñaki San Vicente Roncal, Nikola Ljubešić, Özlem Çetinoğlu, Rahmad Mahendra, Talha Çolakoğlu, Timothy Baldwin, Tommaso Caselli and Wladimir Sidorenko

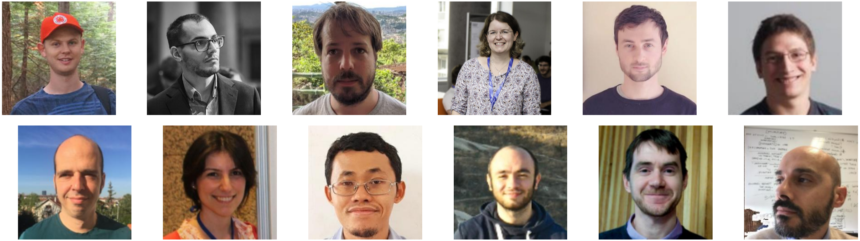

MultiLexNorm

Lang.	Language name	Normalization example
DA	Danish	De skarpe lamper gjorde destromindre ek bedre . De skarpe lamper gjorde destro mindre ikke bedre .
DE	German	ogäj isch hätts auch dwiddern könn Okay ich hätte es auch twittern können
EN	English	u hve to let ppl decide what dey want to do you have to let people decide what they want to do
ES	Spanish	@username cuuxamee sii peroo veen yaa eem @username escúchame sí pero ven ya eh
HR	Croatian	svi frendovi mi nešto rade , veceras san osta sam . svi frendovi mi nešto rade , večeras sam ostao sam .
ID-EN	Indonesian-English	pdhal not fully bcs those ppl jg sih . padahal not fully because those people juga sih .
IT	Italian	a Roma è cosí primavera che sembra gia giov a Roma è così primavera che sembra già giovedì
NL	Dutch	Kga me wss trg rolle vant lachn Ik ga me waarschijnlijk terug rollen van het lachen
SL	Slovenian	jst bi tud najdu kovanec vreden veliko denarja . jaz bi tudi našel kovanec vreden veliko denarja .
SR	Serbian	komunalci kace pocne kaznjavanje ? komunalci kad počne kažnjavanje ?
TR	Turkish	He o dediyin suala cvb verdim He o dediğin suale cevap verdim
TR-DE	Turkish-German	@username Yerimm senii , damkee schatzymm :-* @username Yerim seni , danke Schatzym :-*

MultiLexNorm

- ÚFAL: ByT5 for every word; synthetic data
- HEL-LJU: Pre-classify type of normalization (BERT) $\mapsto$ Char-SMT
- MoNoise: Feature-based, generate candidates and rank
- BLUE: NMT MBart-50
- CL-MoNise: Cross-lingual
- MaChAmp: Normalization as sequence labeling

Results

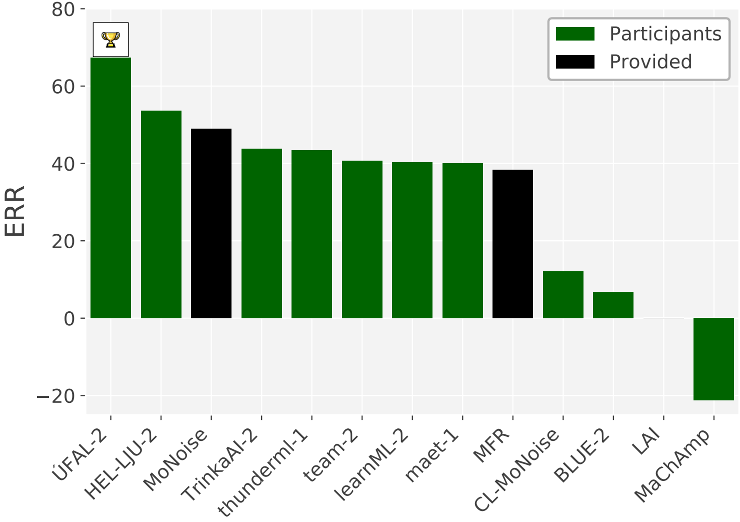

80
60
40
20
0
−20
ERR
Participants
Provided
ÚFAL-2
HEL-LJU-2
MoNoise
TrinkaAI-2
thunderml-1
team-2
learnML-2
maet-1
MFR
CL-MoNoise
BLUE-2
LAI
MaChAmp

Results

- Include detection in task (= the hardest part)
- Multi-lingual benchmark
- Wide variety of models
- Near-human performance for some datasets (in-lang/in-domain)

Open problems

- Cross-lingual/multi-lingual normalization
- Tokenization
- Limited downstream gains; lexical level might not be enough
- Bias in languages
- Bias in data source

MultiLexNorm 2

To be held at WNUT2025 (if accepted), including non Indo-European languages!

2. Multi-task learning

Massive Choice, Ample Tasks (MaChAmp):
A Toolkit for Multi-task Learning in NLP

Rob van der Goot **Ahmet Üstün** **Alan Ramponi** **Ibrahim Sharaf**
Barbara Plank
IT University of Copenhagen University of Groningen University of Trento
Fondazione the Microsoft Research - University of Trento COSBI Factmata
robv@itu.dk, a.ustun@rug.nl, alan.ramponi@unitn.it
ibrahim.sharaf@factmata.com, bapl@itu.dk

MaChAmp

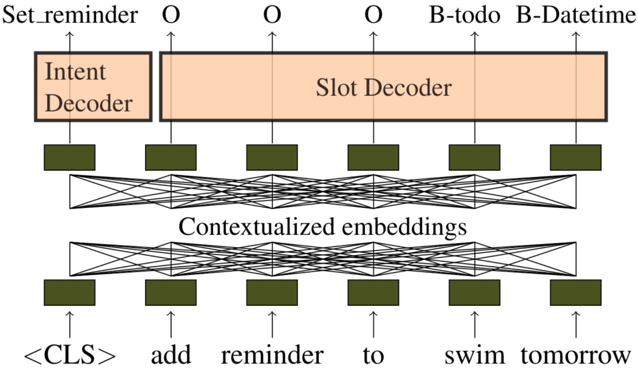
Set_reminder
O
O
O
B-todo
B-Datetime
Intent Decoder
Slot Decoder
Contextualized embeddings
<CLS>
add
reminder
to
swim
tomorrow

MaChAmp

- Supports (almost?) all huggingface models
- 10 task types
- Still SOTA on some tasks
- Many multi-task learning options

xSID: Cross-lingal Slot and Intent Detection

I'd like to see the showtimes for Silly Movie 2.0 at the movie house

xSID: Cross-lingal Slot and Intent Detection

أود أن أرى مواعيد عرض فيلم Silly Movie 2.0 في دار السينما
Jeg vil gerne se spilletiderne for Silly Movie 2.0 i biografen
Ich würde gerne den Vorstellungsbeginn für Silly Movie 2.0 im Kino sehen
I mecht es Programm fir Silly Movie 2.0 in Film Haus sechn
I'd like to see the showtimes for Silly Movie 2.0 at the movie house
Saya ingin melihat jam tayang untuk Silly Movie 2.0 di gedung bioskop
Mi piacerebbe vedere gli orari degli spettacoli per Silly Movie 2.0 al cinema
映画館のSilly Movie 2.0の上映時間を見せて。
Мен Silly Movie 2.0 бағдарламасының кинотеатрда көрсетілім уақытын
Ik wil graag de speeltijden van Silly Movie 2.0 in het filmhuis zien
Želela bih da vidim raspored prikazivanja za Silly Movie 2.0 u bioskopu
Silly Movie 2.0'ın sinema salonundaki seanslarını görmek istiyorum
我想看Silly Movie 2.0 在影院 的放映
I möcht d Spiuzyte für Silly Movie 2.0 im Chino gseh .
Mi piacesse 'e verè ll' orari d' 'e spettacoli pe' Silly Movie 2.0 'o cinema
Norėčiau pamatyti Silly Movie 2.0 seansų laiką kino teatre
I dad gern de Showzeitn füa Silly Movie 2.0 im Filmhaus seng

xSID: Cross-lingal Slot and Intent Detection

Test effect of:

- Auxiliary tasks
- (dialectal) varietion
- Translated vs native text
- Low resource dialects

NorSID:

- 14 different “dialects” of Norwegian
- dialect classification, slot detection, intent detection

https://sites.google.com/view/vardial-2025/shared-tasks?authuser=0

Multi-task learning

MaChAmp at SemEval-2022 Tasks 2, 3, 4, 6, 10, 11, and 12: Multi-task Multi-lingual Learning for a Pre-selected Set of Semantic Datasets

Rob van der Goot
IT University of Copenhagen
robv@itu.dk

MaChAmp at SemEval-2023 tasks 2, 3, 4, 5, 7, 8, 9, 10, 11, and 12: On the Effectiveness of Intermediate Training on an Uncurated Collection of Datasets.

Rob van der Goot
IT University of Copenhagen
robv@itu.dk

3. Future

To what extent are these tasks solved? what are the remaining issues?:

- Tokenization
- Language identification
- Cultural awareness in LLMs

Tokenization

The problem of finding/segmenting tokens (UD):

Input:
If␣momma␣ain't␣happy,␣nobody␣ain't␣happy.

Tokenization:
If␣momma␣ain't␣happy␣,␣nobody␣ain't␣happy␣.

Multi-word expansions:
If␣momma␣is␣not␣happy,␣nobody␣is␣not␣happy.

Subword segmentation:
If␣mo␣##mma␣ai␣##n␣'␣t␣happy␣,␣no␣##body␣ai␣##n␣'␣t␣happy␣.

Methods

1)	Dr. Dron is his backup.
2)	s=[·][.][])} > ”′]∗*$=\1 \2\3 =g
3)	biiobiiiobiobiiobiiiiib

4)

Dr	.	Dro	##n	is	his	backup	.
b	i	b	i	b	b	b	b

LM for tokenization?

- Finetuning a language model for this task might be overkill
- Multi-task learning can make it efficient: add a decoder head for tokenization
- Adapters used before (costly to train)

Settings

- RB: Rule Based
- ST: Single Task: just tokenization
- MT: Multi-task: UPOS, morph. tagging, lemmatization, dep. parsing
- ML+MT: Multi-lingual Multi-task model

In treebank results

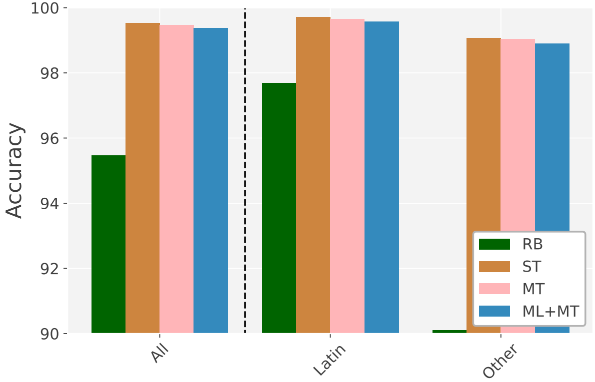

Figure: Results of tokenizers on Latin vs non-Latin languages. RB=RuleBased, ST=SingleTask, MT=MultiTask, ML+MT=Multi-Lingual+MultiTask

Cross-treebank results

setting	F1 tok.	# treebanks
all	93.23	90
in-language	95.11	34
in-script	94.16	84
new-script	80.11	6

Table: Results on test-only treebanks

More analysis

EACL 2024 findings

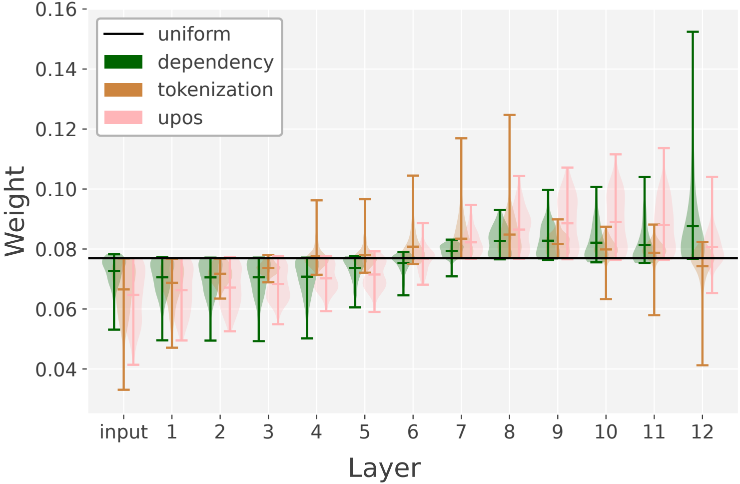

Open challenges in language identification

- Many tools/benchmarks available
- When to use which?:

Open challenges in language identification

- Many tools/benchmarks available
- When to use which?:
  - # languages
  - input size
  - # training instances per language
  - scripts
  - language families
  - domains

Data

Dataset	langs	scripts	fams	domains
OpenLID	139	25	16	literature, news, wiki, social, grammar, subtitles, spoken
UDHR	397	38	61	rights
LTI LangID	2,110	47	139	wiki, political, religious, grammar
TwitUser	59	20	13	social
MassiveSumm	77	24	13	news
UD2.12	54	11	17	medical, news, academic, wiki, legal, nonfiction, learner-essays, fiction, social, grammar-examples, reviews, religious, spoken
Total	2176/ 7850	51/ 163	145/ 298	

Models

- Heuristics: textcat
- Naive Bayes: langid.py
- Embeddings: FastText
- Neural: BiLSTM
- CLM: Glot500

Size

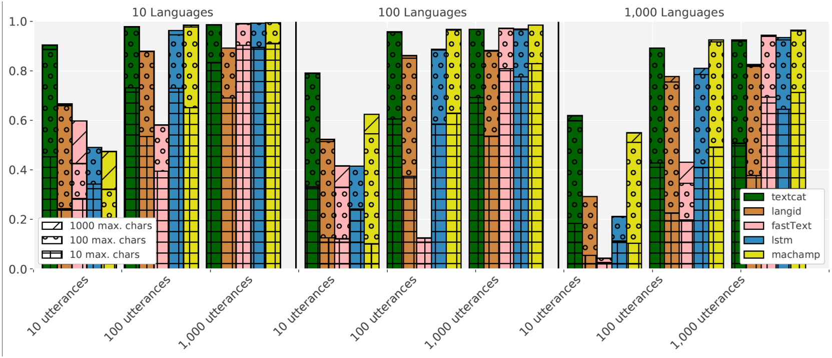
10 Languages
100 Languages
1,000 Languages
1.0
0.8
0.6
0.4
0.2
0.0
1000 max. chars
100 max. chars
10 max. chars
textcat
langid
fastText
lstm
machamp
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances

Size

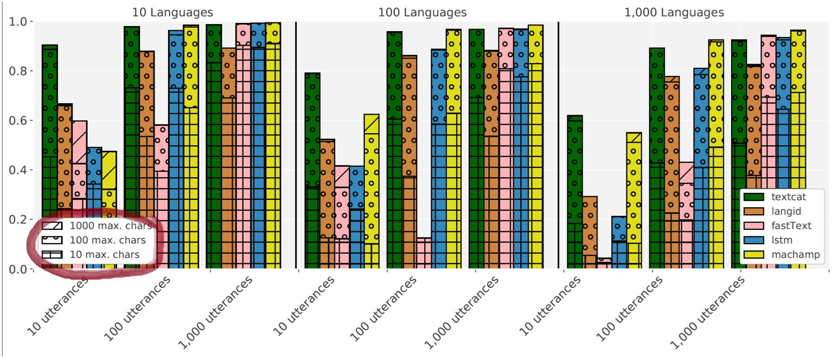
10 Languages
100 Languages
1,000 Languages
1.0
0.8
0.6
0.4
0.2
0.0
1000 max. chars
100 max. chars
10 max. chars
textcat
langid
fastText
lstm
machamp
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances

Size

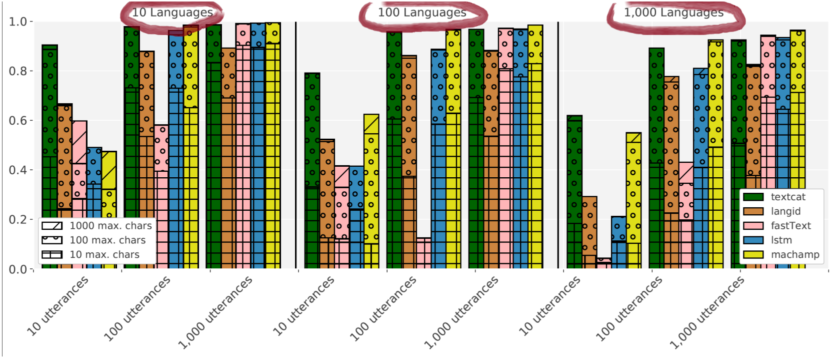
10 Languages
100 Languages
1,000 Languages
1.0
0.8
0.6
0.4
0.2
0.0
1000 max. chars
100 max. chars
10 max. chars
textcat
langid
fastText
lstm
machamp
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances

Size

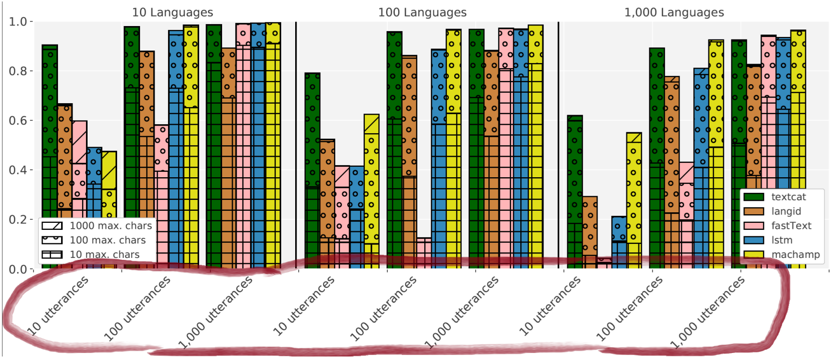
10 Languages
100 Languages
1,000 Languages
1.0
0.8
0.6
0.4
0.2
0.0
1000 max. chars
100 max. chars
10 max. chars
textcat
langid
fastText
lstm
machamp
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances

Size

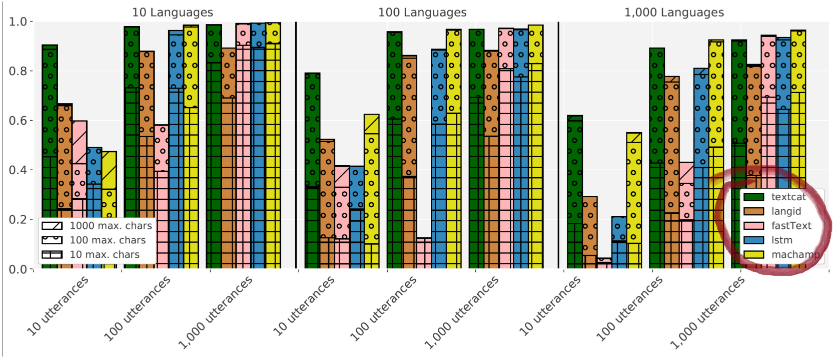

10 Languages
100 Languages
1,000 Languages
1.0
0.8
0.6
0.4
0.2
0.0
1000 max. chars
100 max. chars
10 max. chars
textcat
langid
fastText
lstm
machamp
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances
10 utterances
100 utterances
1,000 utterances

Size: takeaways

- 100 characters is enough
- # of languages is not very influential when there are enough (100) utterances
- Glot500 most robust
- Character n-gram overlap still impressively good

Domains

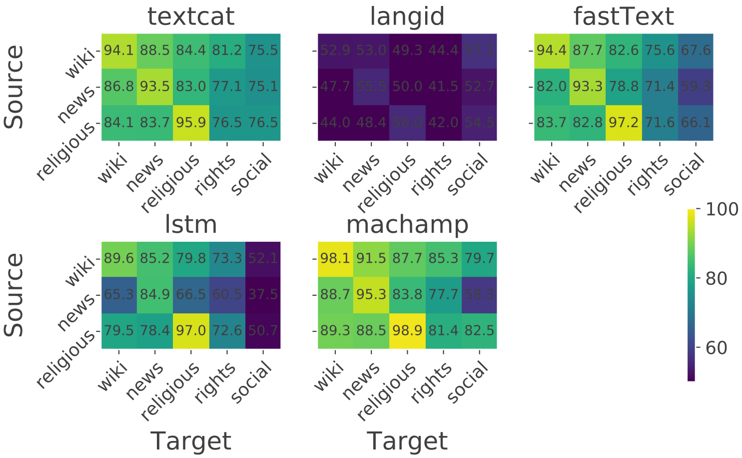
textcat
langid
fastText
lstm
machamp
Source
Target
wiki
news
religious
rights
social
94.1 88.5 84.4 81.2 75.5
86.8 93.5 83.0 77.1 75.1
84.1 83.7 95.9 76.5 76.5
52.9 53.0 49.3 44.4 57.1
47.7 55.5 50.0 41.5 52.7
44.0 48.4 56.0 42.0 54.5
94.4 87.7 82.6 75.6 67.6
82.0 93.3 78.8 71.4 59.3
83.7 82.8 97.2 71.6 66.1
89.6 85.2 79.8 73.3 52.1
65.3 84.9 66.5 60.5 37.5
79.5 78.4 97.0 72.6 50.7
98.1 91.5 87.7 85.3 79.7
88.7 95.3 83.8 77.7 58.3
89.3 88.5 98.9 81.4 82.5
100
80
60

How about LLM's

- Tokenization
- Language classification

Train LLM for Danish

Cultural evaluation of LLMs

Retrain Llama:

- 13B words LM
- 3M instructions (translated)

Cultural evaluation of LLMs

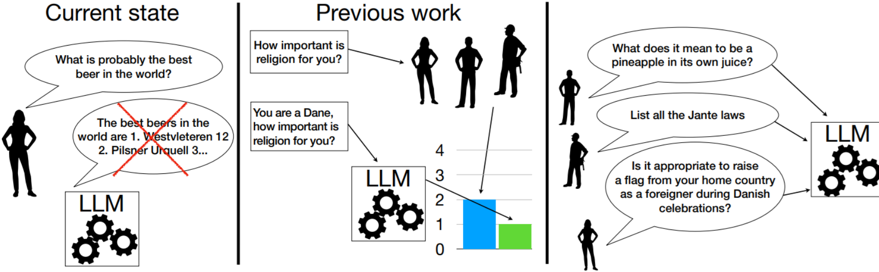

Cultural evaluation of LLMs

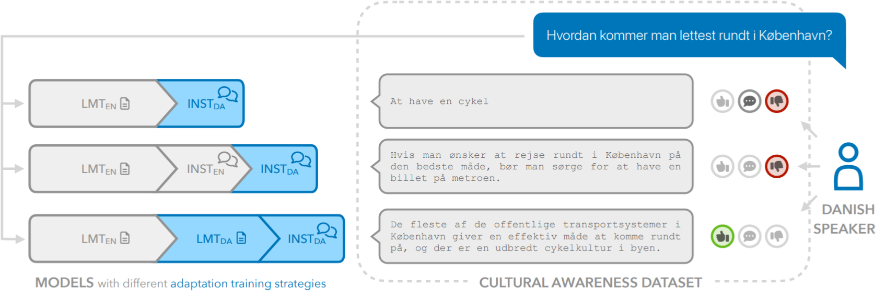

Cultural evaluation of LLMs

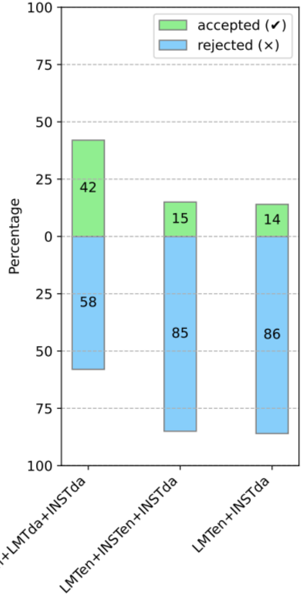

Cultural evaluation of LLMs

为什么中国的龙长得这么长？是不是它们会飞?
为什么中国的月亮节会吃月饼？月饼里面有什么?

Thanks!

Lexical normalization

Multi-task learning

Is X solved?