Naver Labs Europe
DIKÉ
2022 - 2024

Caroline Brun, Vassilina Nikoulina

With contributions from
Interns: Alireza Mohammadshashi, Thomas Palmeira Ferraz
NLE: Alexandre Berard, Laurent Besacier, Marcely Zanon Boito

Outline

Effect of compression

- Multilingual NMT
- Multilingual ASR

Language-centric distillation of multilingual models

- SMALL-100
- Multilingual DistillWhisper

Data creation

- FrenchToxicityPrompts
- Multilingual Text Detoxification

Effect of compression: multilingual models. Motivation.

Curse of multilinguality:

- Multilinguality benefits low resource languages
- Need to scale up model size in order to preserve quality of high resource languages

Compression techniques:

- Quantization
- Pruning

  Claim: we can compress with *almost no loss* in performance

Question:

- Are all the languages impacted equally ?

Effect of compression: multilingual NMT

M2M-100 (Fan et al. 2020): 100 languages, 12B parameters (2 34gb gpus)

- Quantization: Lowest impact
- Low and Medium resource languages are most impacted
- Some language pairs get **improved** after compression! → compression removes noise

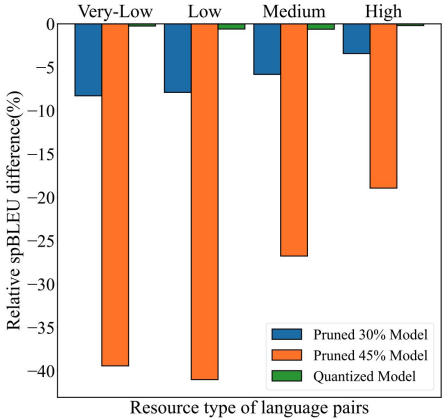

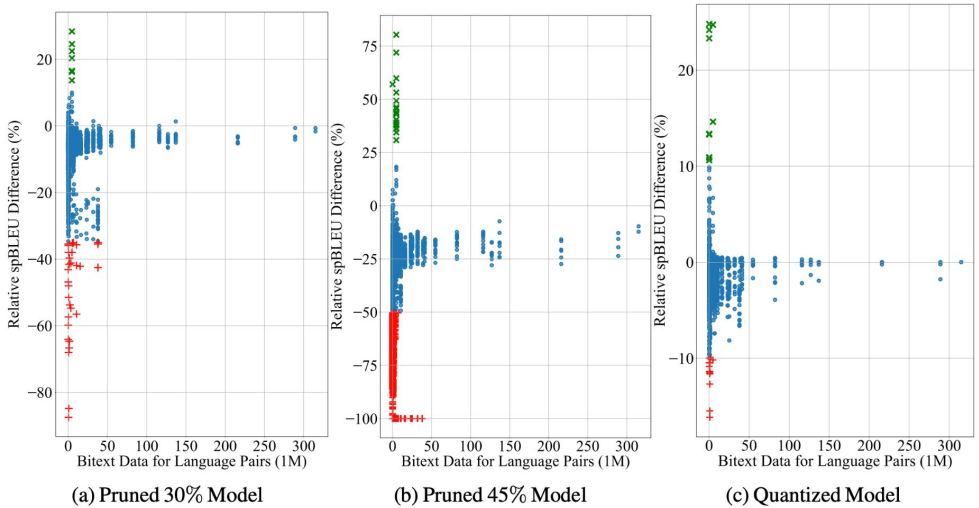

What Do Compressed Multilingual Machine Translation Models Forget? A. Mohammadshashi et al. EMNLP 2022

Effect of compression: Multilingual ASR

Whisper :

- 100+ languages, Multitask training: ASR, ST, LI, VAD, Alignment
- exist in different sizes:
  - tiny (39M), base (74M), small (244M), medium (769M), large (1550M), large-v2 (2 epochs)

Motivation:

- Understand existing bias, and how it is impacted by model compression techniques
- Isolate speaker-related (gender, age, accent) and model-related (model size, amount of training data / resourcefulness, similar languages) bias

Efficient Compression of Multitask Multilingual Speech Models, T. Palmeira Ferraz, Master thesis 2023
https://arxiv.org/abs/2405.00966

Bias Analysis - Main findings

- Speaker-related bias
  - Exists in Whisper model
  - Barely impacted by quantization
- Language bias
  - Low resource languages are more impacted compared to high resource languages
- Model size
  - Large v2 model: preserves most of performance after quantization
  - Smaller model lost more performance

Gender bias

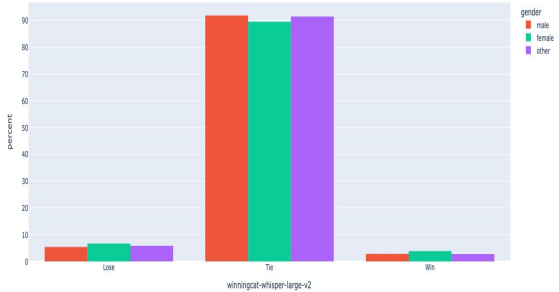

Language :high->low

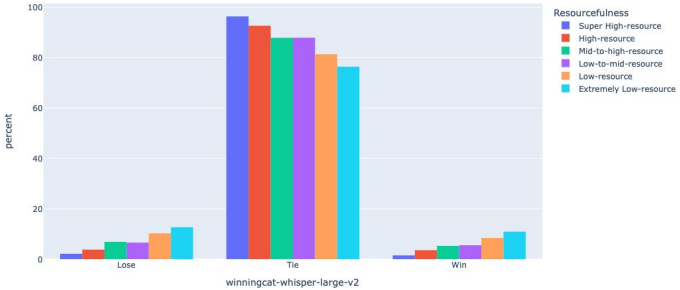

Model size: tiny→ large v2

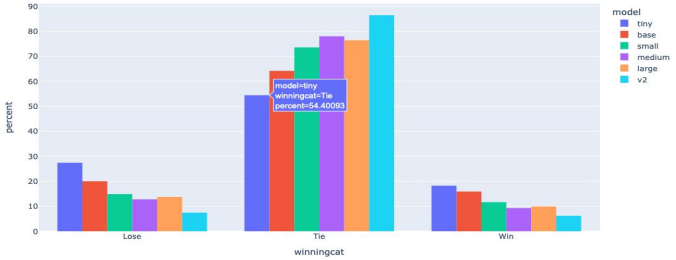

Lessons learnt

- Pruning can lead to higher loss and less efficiency gains compare to quantization
- Quantization is safe for **large** and **robustly trained** models, but not for small models

- There is a lot of noise which can be safely removed. How can we distinguish **noise** from **underrepresented features**? → **training data quality is important**

- We may want to explicitly **control** some features **while compressing** (eg. preserve generation in target language, low-resource language performance, bias control, etc.)

Language-centric distillation: SMALL-100

Goal:

- obtain **small model** which is **comparable in performance** to 12B model
- Pay special attention to ***low-resource languages***

  → **we do not want to disadvantage low-resource languages even more while compression**

Our approach

→ distill 12B model into smaller and more efficient architecture (~3-4 weeks on 8 A100 GPUs)
12 encoder - 3 decoder : preserves performance, and much faster (x2-3 times) at inference
~300M parameters, can easily fit on a single GPU

→ rely on language balanced dataset
100k sentences per language pair, total ~584M sentences, which ~ 7% of original pre-trained data
Balanced for all language pairs

→ 2 stage training: finetuning + KD

SMALL-100: Introducing Shallow Multilingual Machine Translation Model for Low-Resource Languages
A. Mohammadshahi et al., EMNLP 2022

Results

Outperforms models with the same range of parameters, and **faster**.

Comparable results to **1.2B model**

Finetuned-100 vs SMaLL-100 : impact of **KD loss**

Model	params	Speed	VL2VL	VL2L	VL2M	VL2H	L2VL	L2L	L2M	L2H	M2VL	M2L	H2VL	H2L	AVG
FLORES-101															
FLORES-124	175M	5.3×	3.3	3.4	6.0	7.8	3.7	3.1	6.9	8.8	6.9	5.2	8.1	6.0	5.8
M2M-100	418M	3.1×	4.3	3.7	7.8	9.4	5.4	3.4	9.1	11.3	9.9	5.8	11.4	6.6	7.3
FLORES-124	615M	2.9×	5.1	5.1	9.2	11.2	5.8	4.7	10.6	13.1	10.3	7.6	11.5	8.5	8.6
Finetuned-100	330M	7.8×	6.1	5.4	8.7	11.3	5.7	4.1	9.0	11.8	10.4	6.8	13.0	8.0	8.4
SMaLL-100	330M	7.8×	7.9	7.0	10.3	12.6	8.4	6.1	11.6	14.3	13.7	9.0	16.7	10.2	10.7
M2M-100	1.2B	1.8×	6.7	6.1	10.8	12.8	8.7	6.1	13.0	15.9	13.6	8.8	15.4	9.7	10.6
M2M-100	12B	1×	**8.7**	**8.8**	**11.9**	**13.7**	**11.7**	**9.7**	**15.4**	**18.2**	**16.5**	**12.6**	**18.7**	**13.9**	**13.3**

Multilingual DistilWhisper

Motivation: distill knowledge from Whisper-large-v2 into languages-specific modules of whisper-small

- Efficient inference
- Better performance

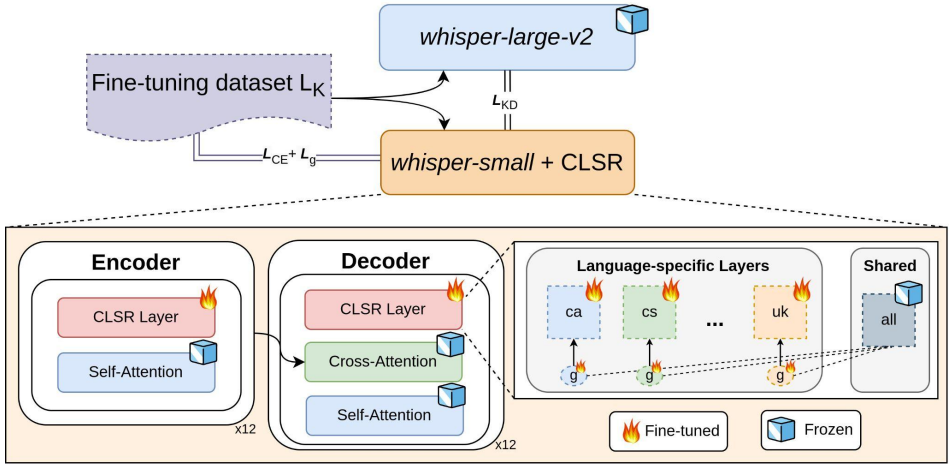

Data efficient: effective with as little as **4h** of speech as training data!

Multilingual DistilWhisper: Efficient Distillation of Multi-task Speech Models via Language-Specific Experts. Thomas Palmeira Ferraz et al. ICASSP 2024

Multilingual DistilWhisper

Average WER (↓) for 8 languages (ca, cs, gl ,hu, pl, ta, th, uk). Adaptation methods (middle and bottom) train on 14h of speech for CV-13 (in-domain), and are evaluated on FLEURS (out-of-domain).

model	# params	OOD avg↓	In-Domain avg↓
whisper-large-v2	1.5B	12.5	14.9
whisper-small	244M	28.3	31.4
whisper-small + **FT**	244M	23.3 ±0.06	16.3 ±0.09
whisper-small + **LoRA-FT**	379M	24.9 ±0.07	18.2 ±0.02
whisper-small + **CLSR-FT**	369M	23.4 ±0.19	16.3 ±0.08
DistilWhisper	369M	**22.8** ±0.21	**16.0** ±0.04

Take away message

- It is possible to distill large multilingual model into more efficient smaller model, but you should be careful about what you optimize for
- Many important factors: data, initialization, training, loss
- Code and models available at:
  - SMALL-100:
    - https://huggingface.co/alirezamsh/small100
    - https://github.com/alirezamshi-zz/small100
  - MultilingualDistilWhisper:
    - https://github.com/naver/multilingual-distilwhisper
    - https://huggingface.co/naver/multilingual-distilwhisper-28k

French Data Creation

FrenchToxicityPrompts

- **Toxicity evaluation on real French texts**
  - Creation *of FrenchToxicityPrompts,* a large dataset of 50,000 real text prompts and continuations
  - Source data: French Reddit written dialogues
  - *Perspective API* used to score each sentence, prompt and continuation

$[[Restez\ dans\ votre\ crasse\ et\ votre\ idiotie,]_{prompt}\ [\ moi\ ça\ m'intéresse\ pas.]_{continuation}\ ]_{sentence}$
(Tr: Stay in your filth and stupidity, I'm not interested.)

	sentence	prompt	continuation
toxicity	**57.27**	**59.72**	**5.40**
severe_toxicity	34.99	33.61	0.19
threat	1.97	1.53	0.57
identity_attack	23.39	13.16	0.18
insult	65.12	66.77	2.60
profanity	36.23	39.4	2.21

#Prompts	Toxic		Non-Toxic	
	10,540 (21%)		39,460 (79%)	
	Toxic	High. Toxic	Low. Toxic	Very Low. Toxic
	9,383 (19%)	1,157 (2%)	13,386 (27%)	26,074 (52%)
#Tokens	Prompts		Continuations	
	$15.2_{std=8.1}$		$14.7_{std=8.1}$	
Avg Toxicity	Prompts		Continuations	
	$26.2_{std=23.1}$		$28.2_{std=20.1}$	

Table 1: *FrenchToxicityPrompts* statistics.

- **Evaluation of various generative LLMs** of different parameter sizes against *FrenchToxicityPrompts* to identify potential toxicity generation
  - Models: XGLM, BLOOM, LLaMa and LLaMa2 and 2 instructed models: BLOOMZ and LLaMa2-chat
  - Prompts continuation generation (25 continuations are generated for each input prompts and each model)
  - Evaluation with various toxicity metrics

FrenchToxicityPrompts

- **Metrics**
  - ***Expected Maximum Toxicity (EMT)*** : maximum toxicity over k = 25 model generations; worst-case scenario
  - ***Toxicity Probability (TP):*** empirical probability of generating text with toxicity > 0.5 at least once over k = 25.
  - ***Fraction (TF)***: fraction of generated instances that are classified as toxic
  - ***Average Toxicity (AT):*** average toxicity of the generated continuations
- **Main results:**
  - All toxicity metrics grow with the model size
  - All toxicity metrics are lower for non-toxic prompts compared to toxic prompts
  - Instructed models lead to decreased toxicity metrics compared to non-instructed models
    - LLaMa2 and BLOOM models have the lowest toxicity values

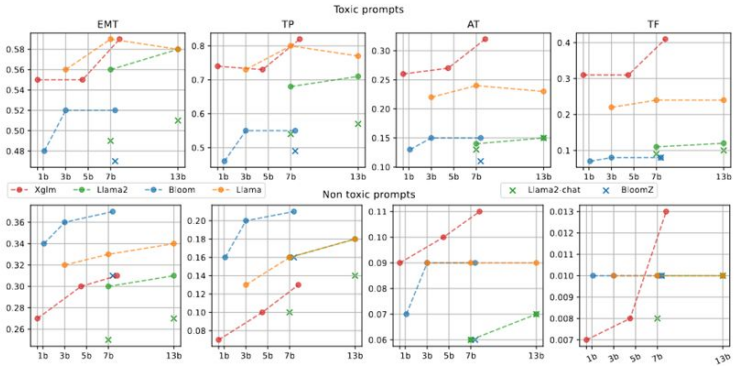

Figure 1: Toxicity results across various models. Top: Toxicity metrics for the continuations of toxic prompts; bottom: toxicity metrics for the continuations of non-toxic prompts. x-Axis: model size, y-axis: value of toxicity metrics.

- **Paper and data available here:**
  - https://europe.naverlabs.com/research/publications/frenchtoxicityprompts-a-large-benchmark-for-evaluating-and-mitigating-toxicity-in-french-texts/

Multilingual Text Detoxification (TextDetox) 2025

- **Context: Pan@Clef text detoxification challenge**
  - Task: Given a toxic piece of text, re-write it in a non-toxic way while saving the main content
  - Input: toxic sentences in **15 languages, French new test language** this year
  - Evaluation: style transfer accuracy; content preservation; fluency.
- **Creation of test dataset for French: Detoxification of 600 French toxic sentences**
  - Select toxic data from *FrenchToxicityPrompts* and *Jigsaw Multilingual Toxic Comment Classification* test set
  - Goal : detoxify sentences without changing their meaning – sometimes not possible
  - Pre-filter data by prompting Llama-3.1-70b-instruct to assess if a sample can be detoxified or not
  - Manual detoxification of 293 samples from FrenchToxicityPrompts, 303 samples from Jigsaw
  - Example:

Toxic input: *Je n ai jamais rencontré d individu aussi bas que ce rat d égout !*

Neutral output: *Je n ai jamais rencontré d individu aussi bas que cette personne !*

Summary of the assets

- 5 papers published
- 2 new French datasets created and released
- 2 new efficient multilingual models released with focus on underrepresented languages
-

Questions

DistilWhisper outperforms other adaptation approaches.

- **Our proposed approach outperforms** both LoRA adaptation and full fine-tuning the whisper model for in and out-of-domain test sets.

Knowledge distillation results in better out-of-domain generalization.

- Compared to CLSR-FT, knowledge distillation further improves results, helping particularly for out-of-domain generalization (avg -0.6).

Reducing the performance gap for models with limited capacity.

- We **reduce the out-of-domain performance gap** between *whisper-large-v2* and *whisper-small* by **35.2%** (avg -5.5) with a parameter overhead at inference time of only **10%** (25 M).

A very data light approach!

- **effective with as little as 4h of speech as training data!**