Natural Language Understanding, Generation, and Machine Translation

Lecture 21: Translation and LLMs

7 March 2025 (week 7)

Alexandra Birch

Do we still need MT?

- MT Central to NLP:

  big data, probabilistic modelling,

  encoders-decoders, attention, subwords

- Convergence of NLP on a unified deep learning framework - still train MT models
- And now just ask GPT: Translate "X" to Y

Neural MT

What did we mean by NMT?

- Transformer Encoder-Decoder
- Focus on parallel data
- Bilingual or Multilingual
- Large (but not that large?)
  - MBART 600M
  - NLLB MOE 54.5B parameters and FLOPs similar to that of a 3.3B dense model
  - JDExplore won many WMT22 4.7B

Neural MT

What did we mean by NMT?

No Language Left Behind: Scaling Human-Centered Machine Translation
Costa-jussà et al. 2022

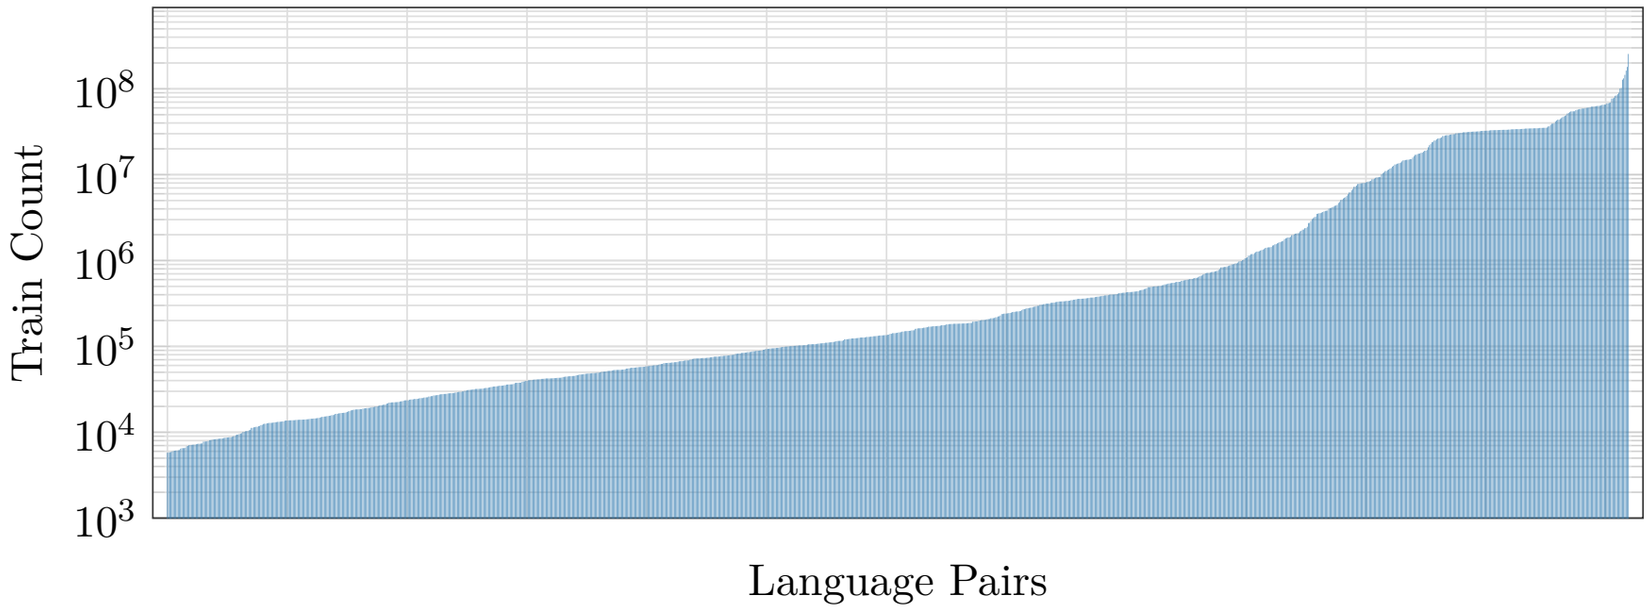

202 language, parallel/mined/BT 1220 language pairs, 18B sentence pairs

Pretrain-FineTune Paradigm

- Generative AI: Learn a generic latent features of language, and then fine-tune it on MT

Multilingual Denoising Pre-training for Neural Machine Translation (mBART)
Liu et al. 2020

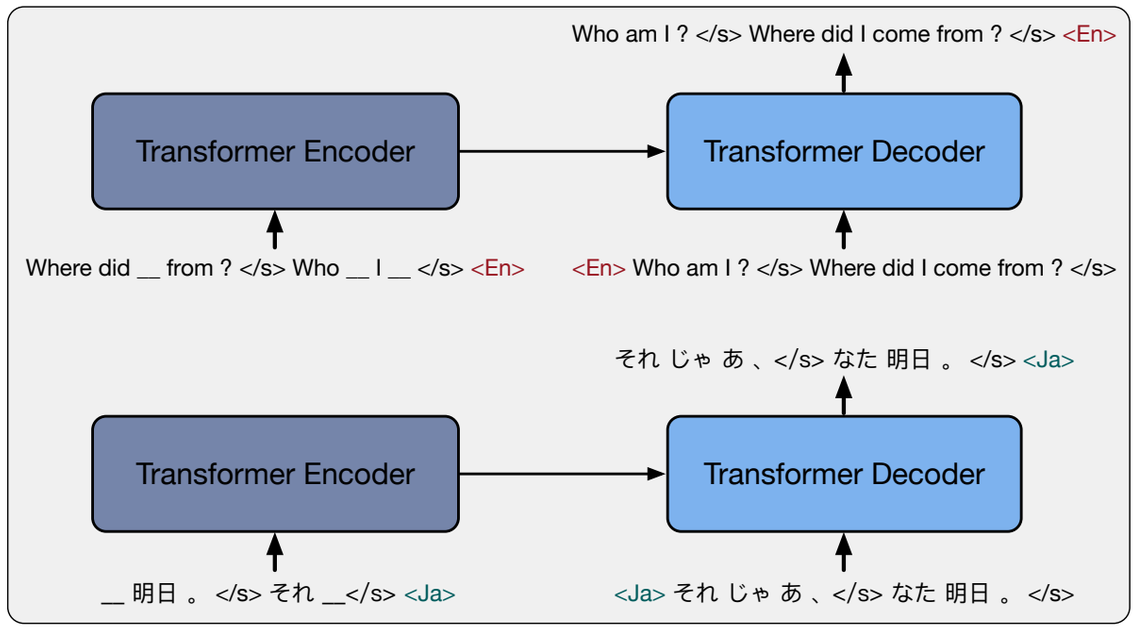

Multilingual Denoising Pre-Training (mBART)

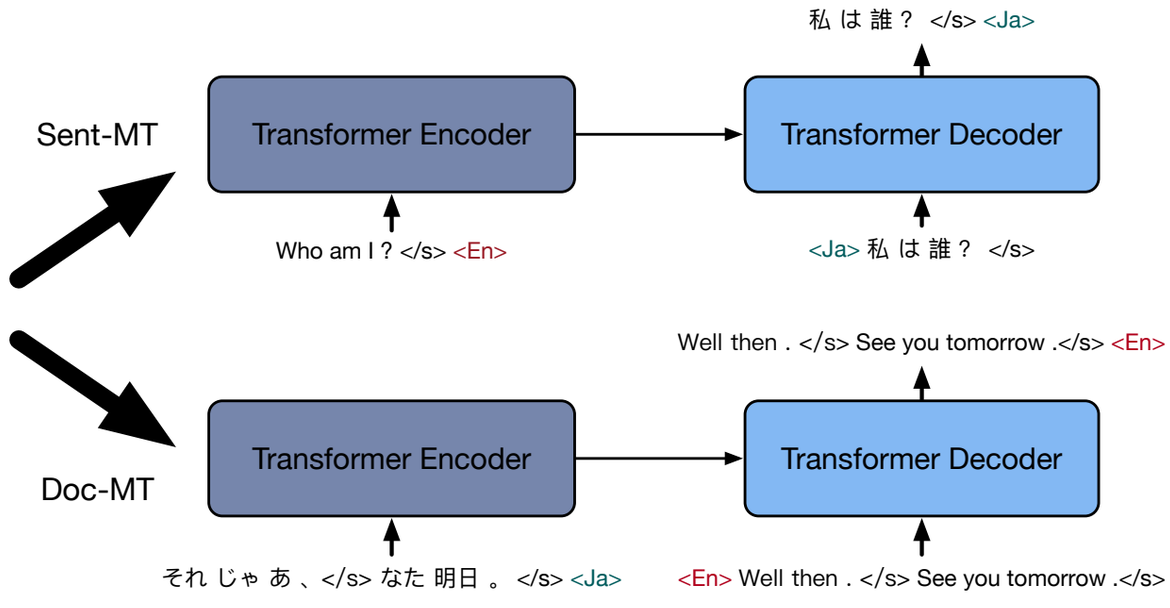

Fine-tuning on Machine Translation

600M param

Pretrain-Prompt Paradigm

- When models are large enough - don't need to fine-tune!
- Just Pretrain and then Prompt!
- Don't need an Encoder - Decoder only architecture
- Mostly trained to predict next word
- Models are very large: > 7B parameters, up to 200B
- Data and compute very large - no longer in reach apart from a handful of groups

How multilingual are LLMs?

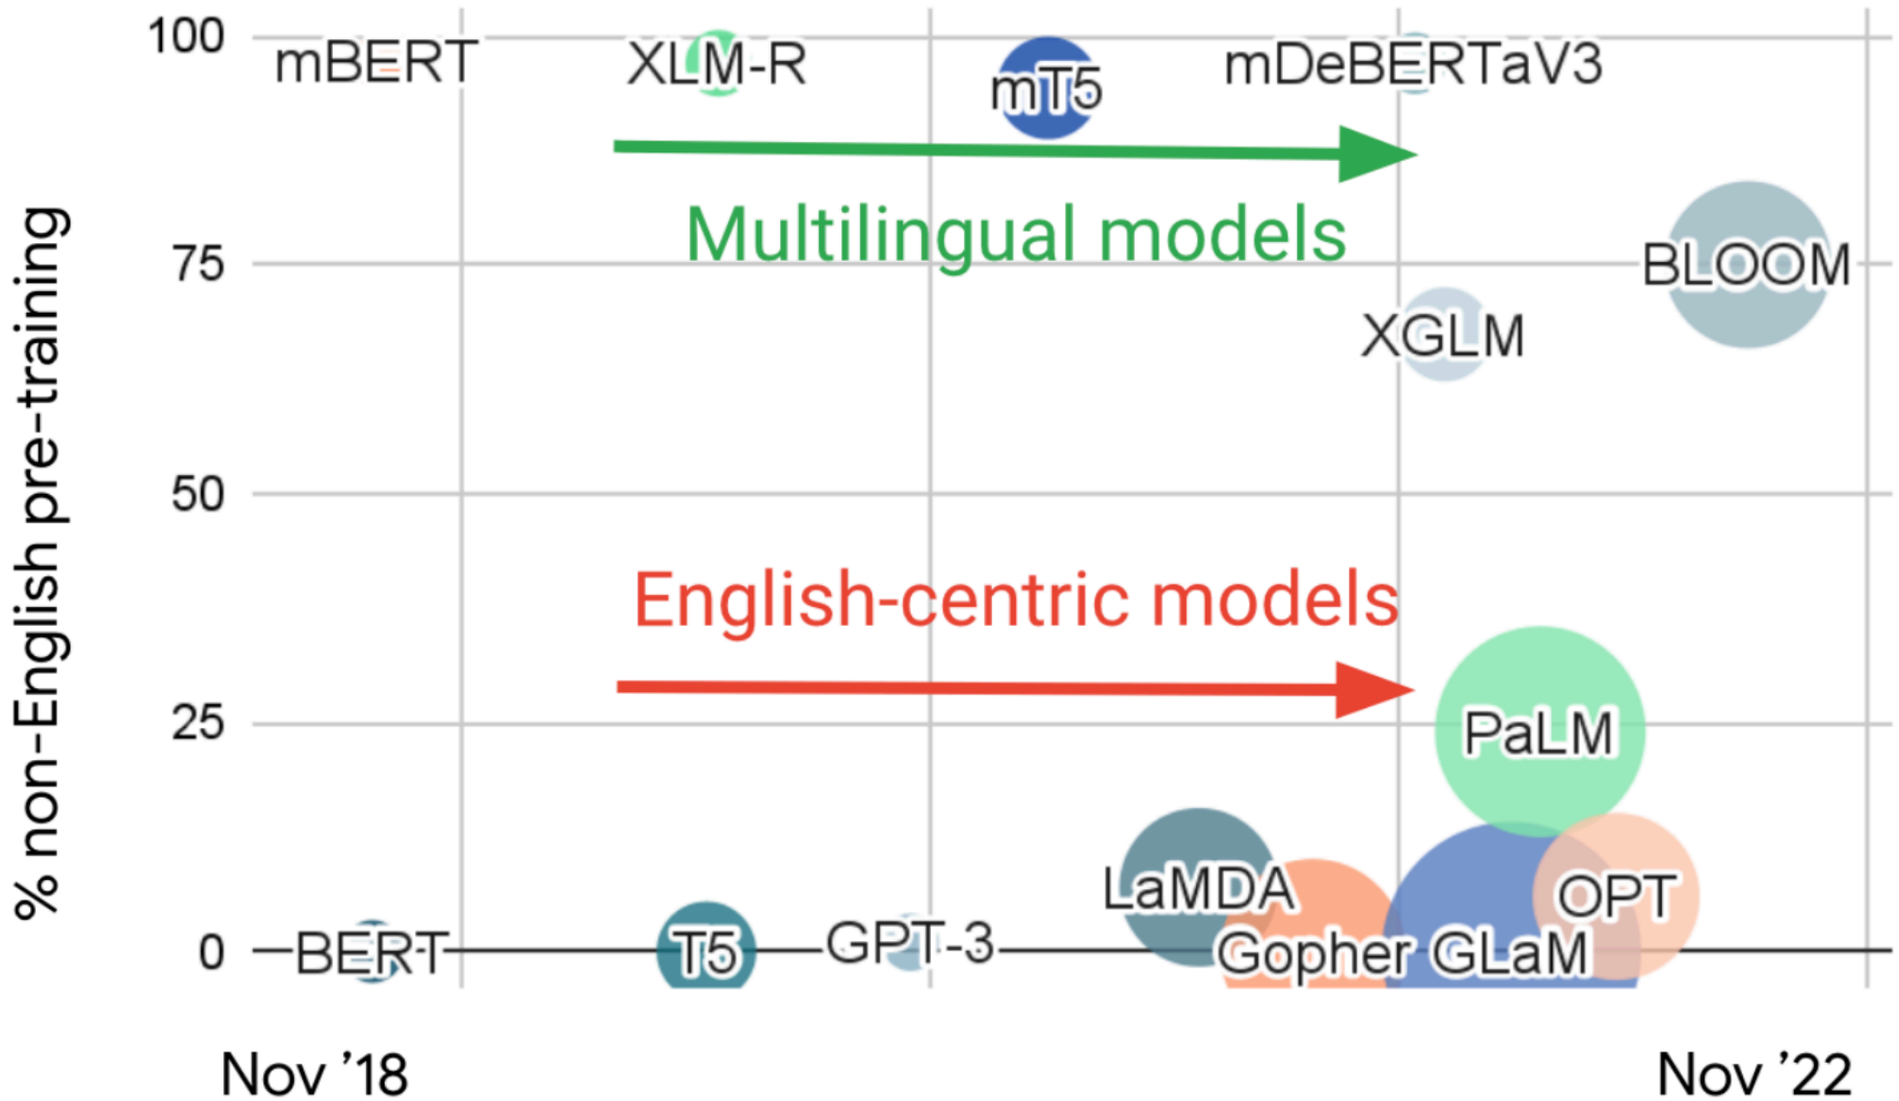

From: https://www.ruder.io/state-of-multilingual-ai/
adapted from Noah Constant

Pretrain-Prompt Paradigm

Language models are few shot learners (GPT3)
Brown et al. 2020

Setting	En→Fr	Fr→En	En→De	De→En	En→Ro	Ro→En
SOTA (Supervised)	**45.6**[a]	35.0 [b]	**41.2**[c]	40.2[d]	**38.5**[e]	**39.9**[e]
XLM [LC19]	33.4	33.3	26.4	34.3	33.3	31.8
MASS [STQ+19]	37.5	34.9	28.3	35.2	35.2	33.1
mBART [LGG+20]	-	-	29.8	34.0	35.0	30.5
GPT-3 Zero-Shot	25.2	21.2	24.6	27.2	14.1	19.9
GPT-3 One-Shot	28.3	33.7	26.2	30.4	20.6	38.6
GPT-3 Few-Shot	32.6	39.2	29.7	40.6	21.0	39.5

175B, 7% non English

Prompt Engineering

What is the best way to prompt for translation?

Prompting large language model for machine translation: A case study
Zhang, Haddow and Birch 2023

ID	Template (in English)	English		German		Chinese	
		w/o	w/	w/o	w/	w/o	w/
A	`[src]: [input] ◇ [tgt]:`	**38.78**	**31.17**	-26.15	-16.48	**14.82**	**-1.08**
B	`[input] ◇ [tgt]:`	-88.62	-85.35	-135.97	-99.65	-66.55	-85.84
C	`[input]` ◇ Translate to `[tgt]:`	-87.63	-68.75	-106.30	-73.23	-63.38	-70.91
D	`[input]` ◇ Translate from `[src]` to `[tgt]:`	-113.80	-89.16	-153.80	-130.65	-76.79	-67.71
E	`[src]: [input]` ◇ Translate to `[tgt]:`	20.81	16.69	**-24.33**	**-5.68**	-8.61	-30.38
F	`[src]: [input]` ◇ Translate from `[src]` to `[tgt]:`	-27.14	-6.88	-34.36	-9.22	-32.22	-44.95

GLM-130B En,Zh, COMET

Preference for simple English prompt

In Context Learning

How many examples do we need?

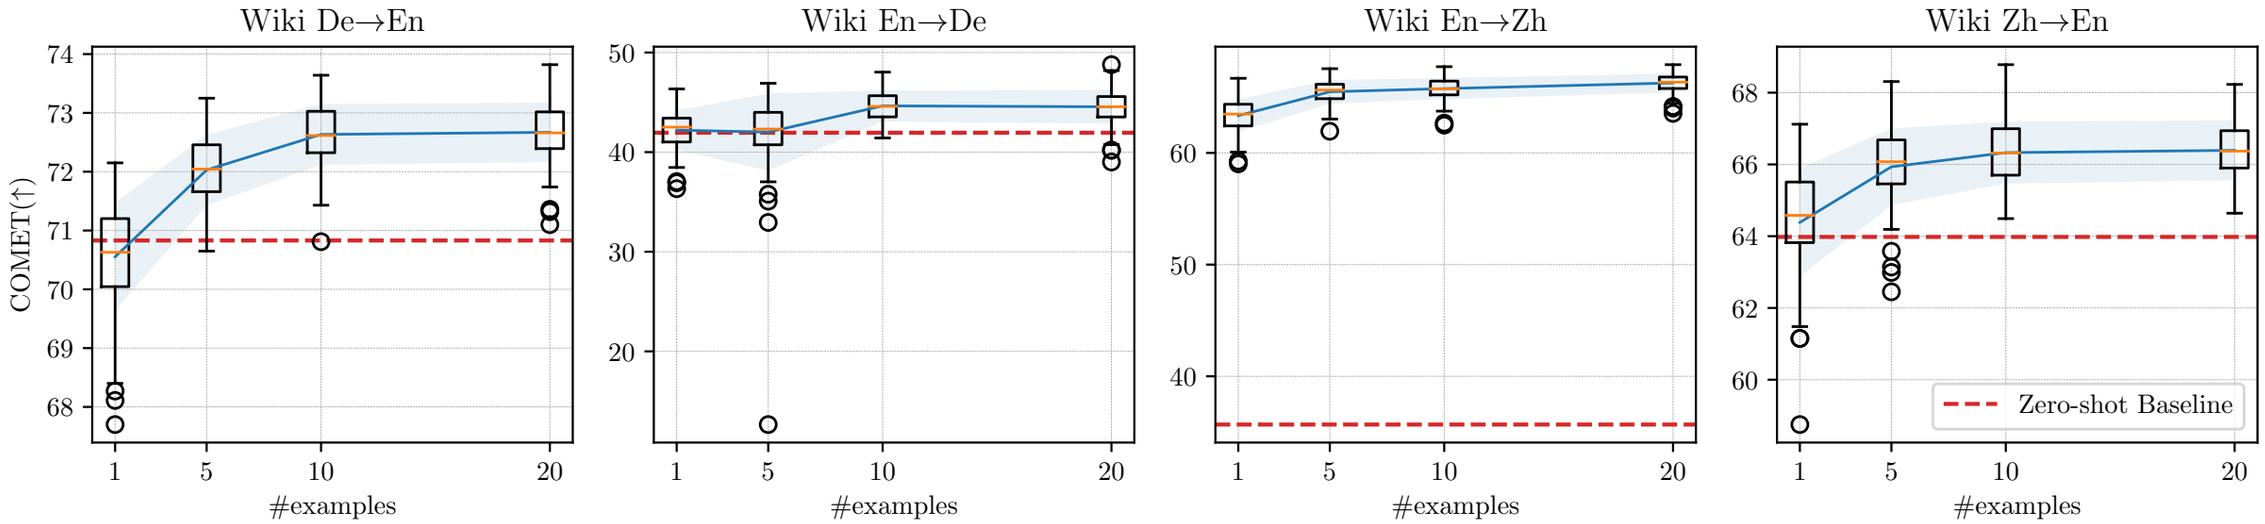

In Context Learning

Does the quality of example matter?

Method	Wiki		WMT	
	BLEU	COMET	BLEU	COMET
Zero-Shot	24.08	33.92	20.38	17.97
1-Shot Translation (high-quality pool)				
Random	26.31	48.29	21.27	30.70
SemScore	26.73	49.34	21.82	31.28
LMScore	26.48	47.92	21.59	30.81
TLength	26.54	48.73	21.29	30.68
5-Shot Translation (high-quality pool)				
Random	**27.46**	51.11	21.82	33.87
SemScore	27.36	**51.66**	**22.37**	34.30
LMScore	27.17	50.65	22.04	**35.19**
TLength	27.08	50.50	21.75	34.29
1-shot Translation (Low-quality Pool)				
Random	24.75	38.86	22.06	30.70
Ours	24.94	39.88	22.23	30.87

Problems remain

Source	根据三江源国家公园管理局长江源园区可可西里管理处统计，藏羚羊回迁数量总体呈逐年上升态势，2019年藏羚羊回迁数量为4860只，比2018年增加338只。
Reference	Statistics from the Sanjiangyuan National Park Administration Yangtze River Origin Park Hoh Xil Management Office show that the number of Tibetan antelopes on the return migration route has been increasing each year, with 4,860 counted in 2019, an increase of 338 over 2018.
GLM-130B (1-shot)	According to the三江源国家公园管理局长江源园区可可西里管理处, the total number of re-migration of the Tibetan antelope has been on the rise since 2018, with 4,860 re-migrating in 2109, an increase of 338 compared to 2808.
Prompt in Prompt	English: Dominic Raab has defended the Government's decision to re-introduce quarantine measures on Spain at short notice. **Translate from English to Chinese:** Chinese:
Reference	针对政府突然做出重新对西班牙实施隔离措施的决定，Dominic Raab 做出了辩解。从英文翻译成中文：
GLM-130B (zero-shot)	多米尼克·拉布(Dominic Raab)对政府决定重新引入西班牙的检疫措施表示支持。**Translate from English to Chinese:**

Errors: copying, dates, misunderstanding, **prompt trap**

Are LLMs competitive?

Multilingual Machine Translation with Large Language Models: Empirical Results and Analysis
Wenhao Zhu et al. 2023

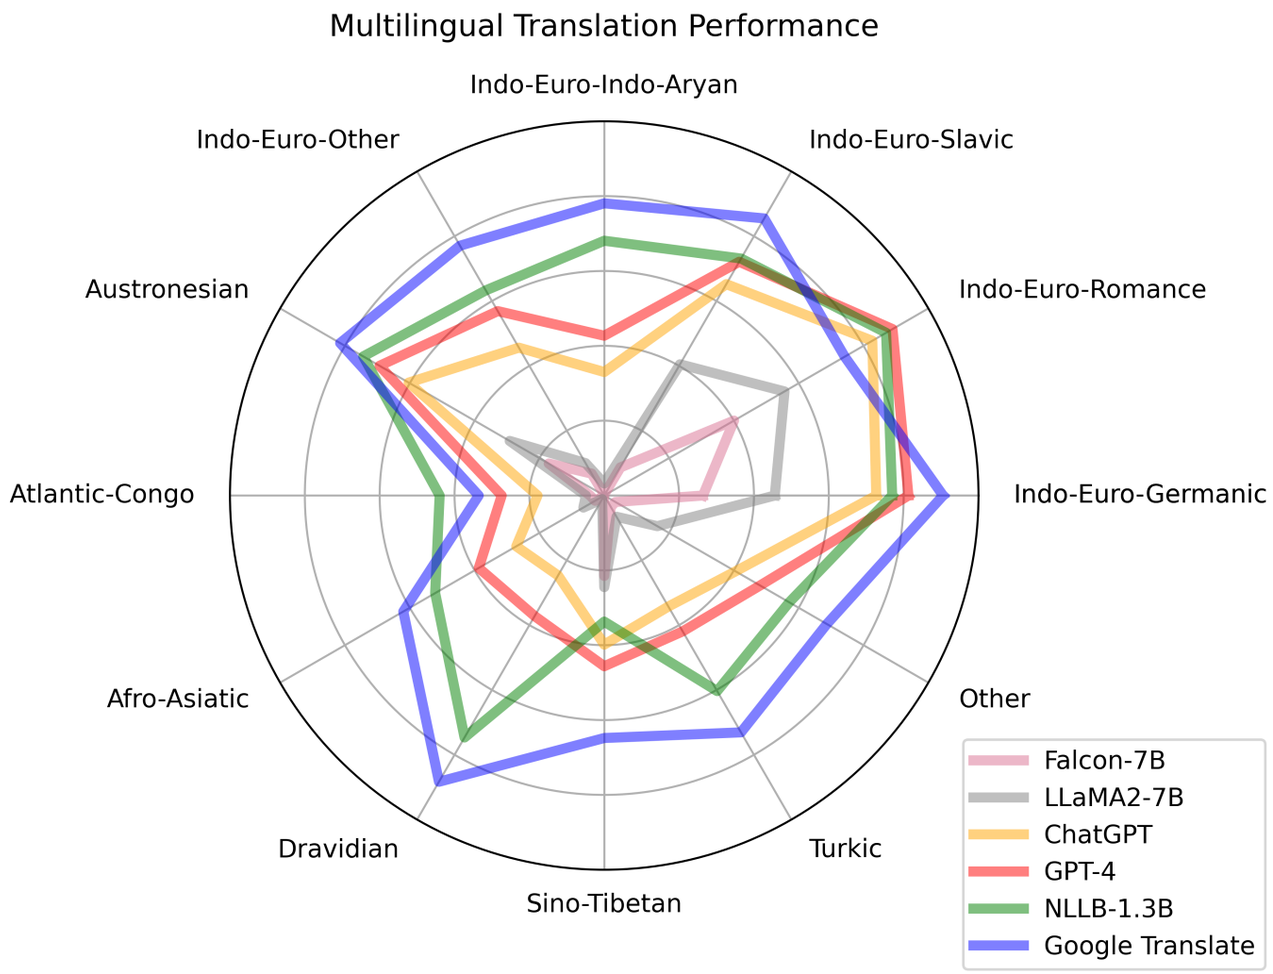

En-X, 8 in context examples, 101 Flores languages

Behaviour LLMs vs MT?

How Good Are GPT Models at Machine Translation? A Comprehensive Evaluation Hendy et al. 2023

Parallel data Bias:

- Noise from parallel data
- Data from strange domains with different distributions

Monolingual Bias

- Instructions might fail to override LLM training
- Lack of teacher forcing supervision means might not be faithful to source sentence
- Favour fluency over accuracy eg, introducing undesirable punctuation or removing tokens which have been unseen

Behaviour LLMs vs MT?

Sequence Type	Translation Instance	Phenomenon
Source	Bis auf die E 95 02 wurden **alle Lokomotiven zerlegt**.	
MS Translator	With the exception of E 95 02, **all locomotives were dismantled**.	**Non-Monotonicity (NM)**
GPT	**All locomotives were dismantled** except for the E 95 02.	
Source	**Oder ist sie** ganz aus dem Sortiment genommen?	
MS Translator	**Or is it** completely removed from the range?	**Fluency (F)**
GPT	**Or has it been** completely removed from the range?	
Source	Sehen Sie bitte im Screenshot was der Kollege geschrieben hat	
MS Translator	Please see in the screenshot what the colleague wrote	**Punctuation Insertion (PI)**
GPT	Please see the screenshot for what the colleague wrote**.**	
Source	Die Email zur Stornierung wurde am 26.12.**#NUMBER#** versendet.	
MS Translator	The cancellation email was sent on 26.12.**#NUMBER#**.	**Dropped Content (USW)**
GPT	The cancellation email was sent on December 26th.	
Source	"We won’t accept the **CAA** and that is for sure.	
MS Translator	“我们不会接受**CAA**，这是肯定的。	**Inserted Content (UTW)**
GPT	“我们不会接受《公民法》，这是肯定的。	

Behaviour LLMs vs MT?

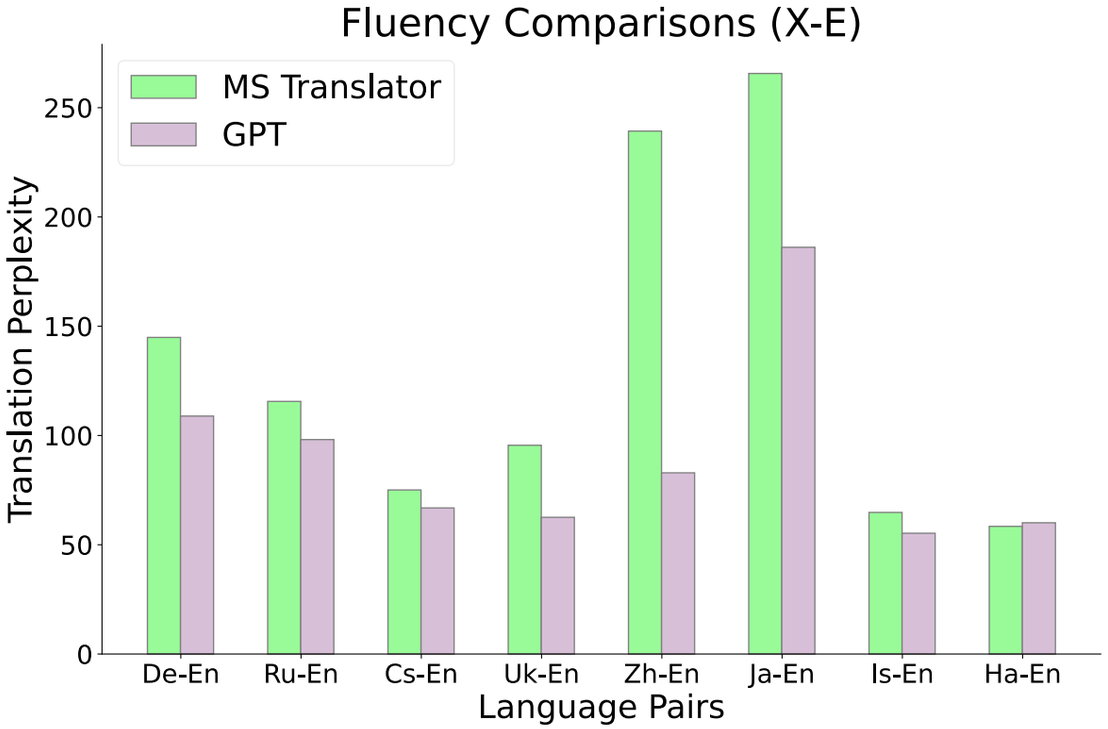

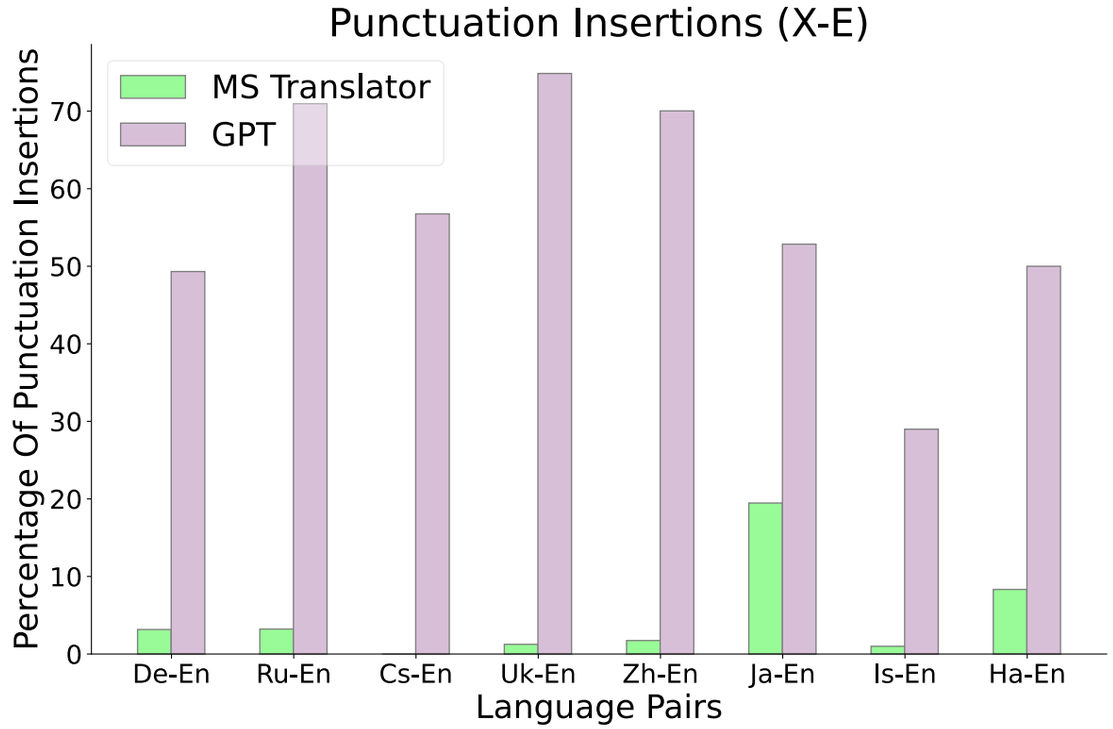

Behaviour LLMs vs MT?

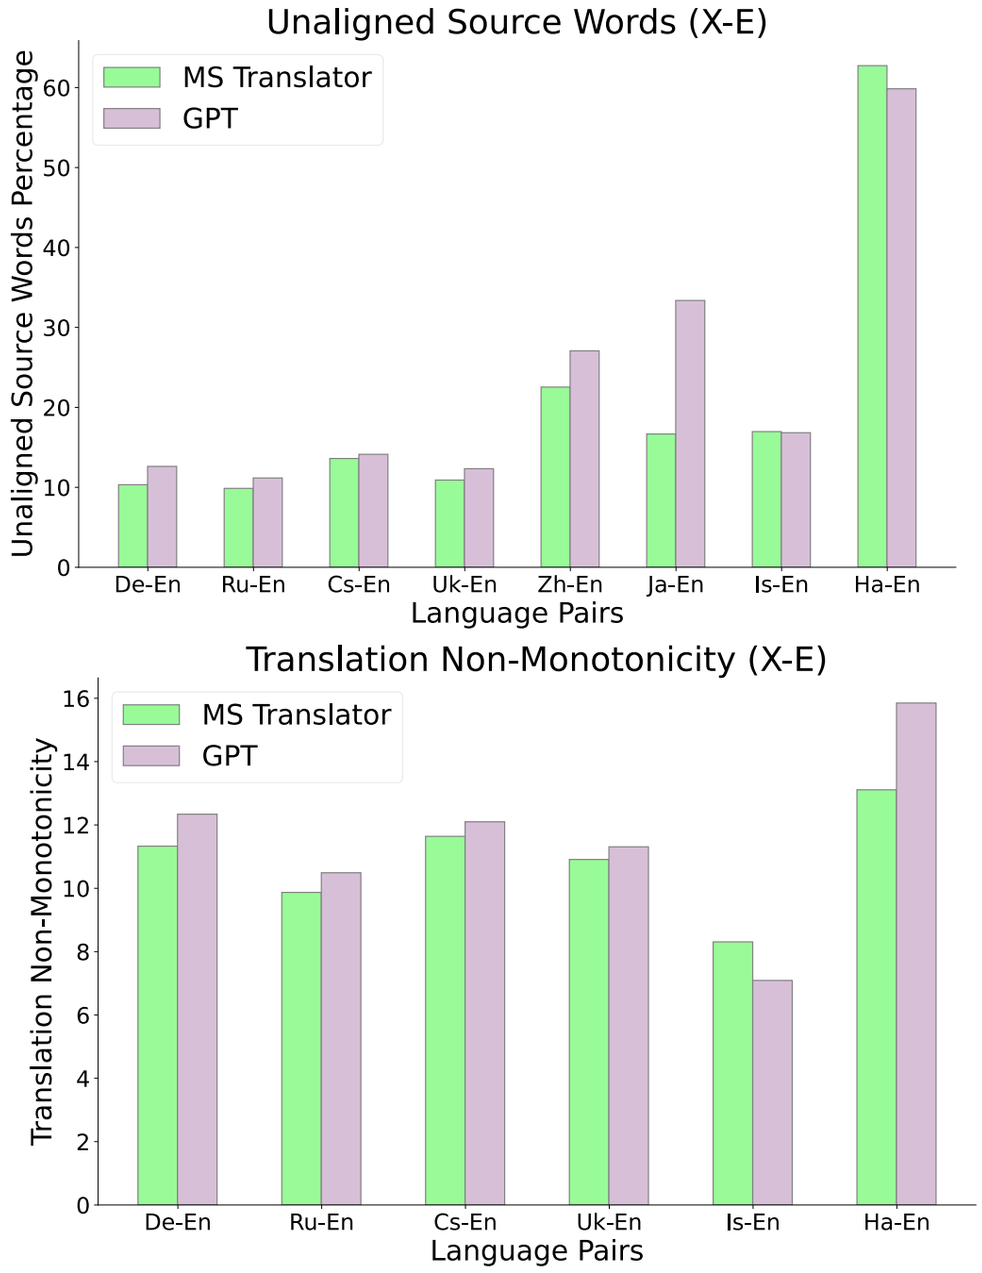

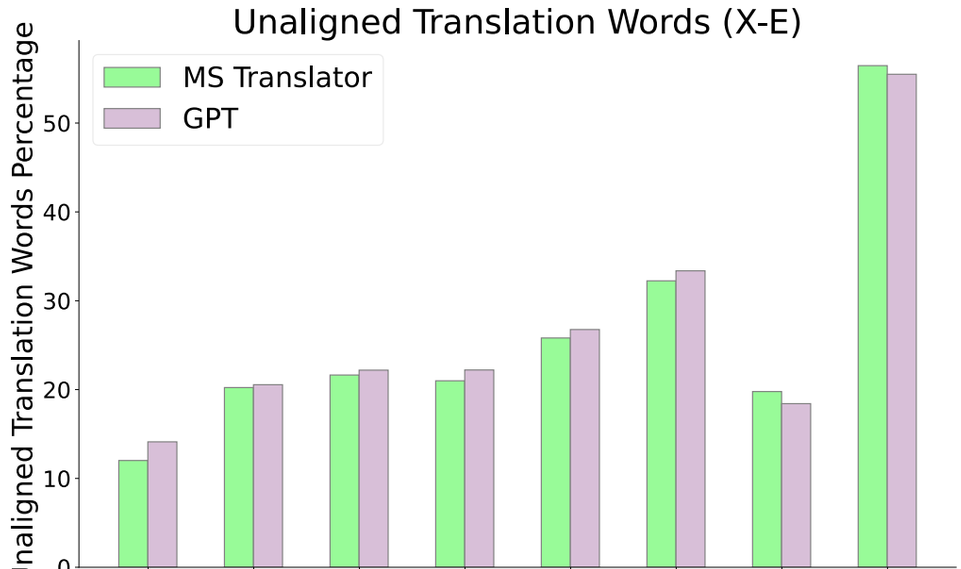

- Ignore more source - no more inserted target
- Less literal
- Better for figurative text

Opportunities

Controllable:

Hire a Linguist!: Learning Endangered Languages with In-Context Linguistic Descriptions. Zhang et al. 2024

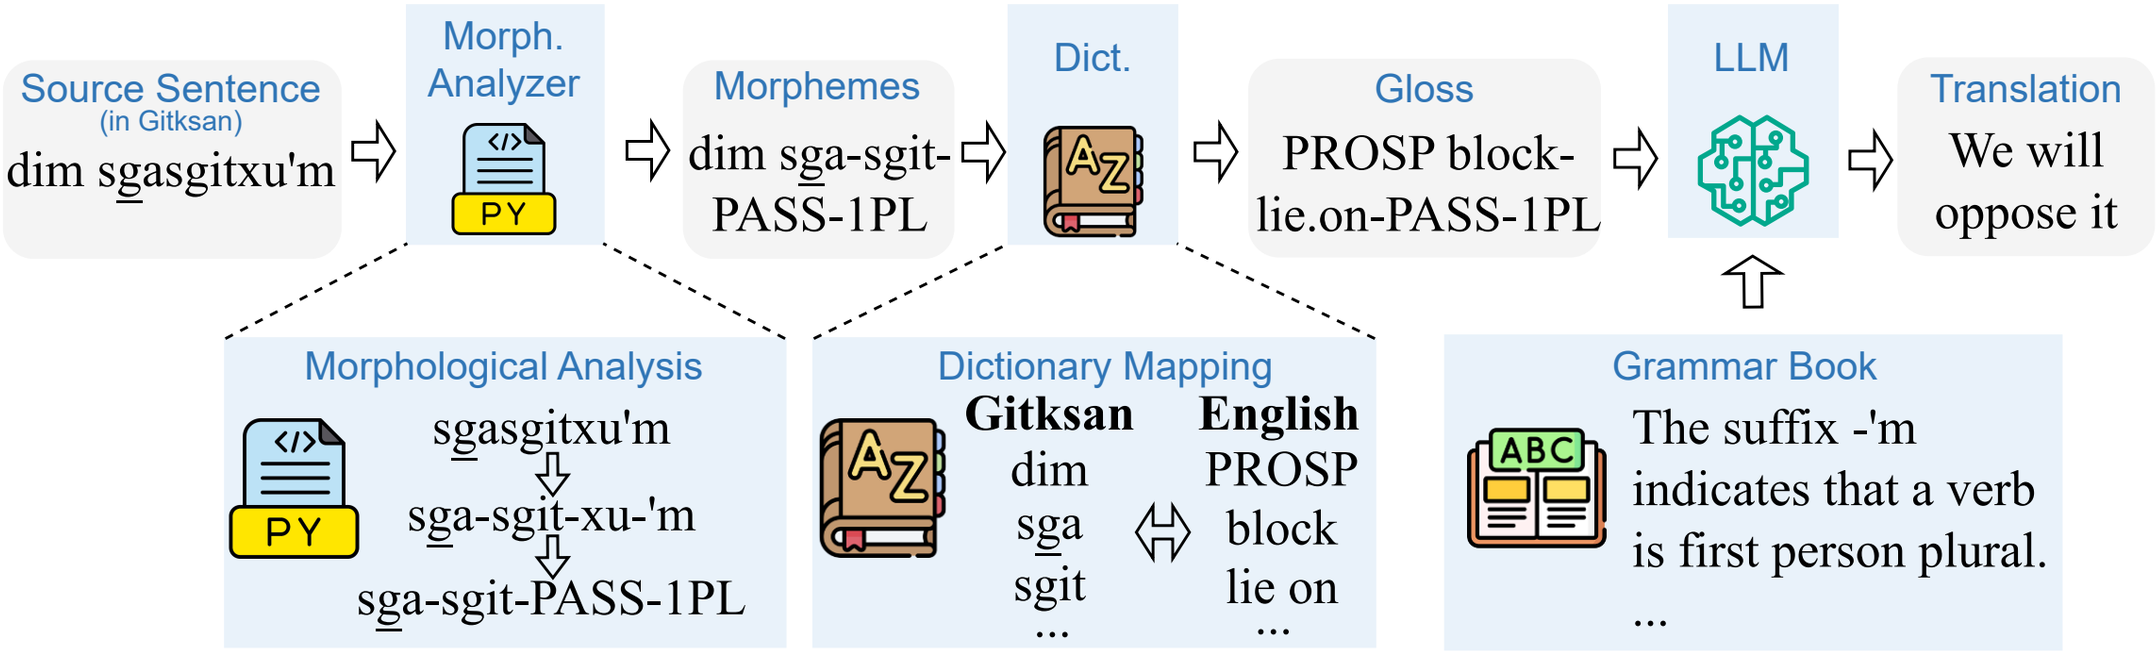

Opportunities

Controllable:

Towards Effective Disambiguation for Machine Translation with Large Language Models. Iyer, Chen and Birch 2023

Source	The horse had a blaze between its eyes.
DeepL	那匹马的两眼之间有一团火焰。 (There is a flame between the horse's eyes.)
BLOOMZ (176B)	这匹马的眼睛之间有一道白线。 (There is a white line between the horse's eyes.)

Opportunities

Controllable:

Towards Effective Disambiguation for Machine Translation with Large Language Models. Iyer, Chen and Birch 2023

System	1-shot		3-shot		5-shot	
	Rand.	Sim.	Rand.	Sim.	Rand.	Sim.
DeepL	—63.91—					
NLLB-200 54B	—***61.33***—					
LLaMA 7B	53.64	**54.01**	**55.53**	52.52	**56.33**	54.45
LLaMA 65B	56.57	**59.38**	59.83	**62.44**	60.78	**63.74**
BLOOM 176B	**63.66**	62.44	64.52	**66.19**	65.53	**68.22**
BLOOMZ 176B	64.35	**69.57**	67.31	**71.15**	68.55	**71.33**

EuroLLM

- Multilingual Support all official EU languages plus selected major world languages. **Pretrain from scratch** with best tokenisation!
- High Performance Competitive with similar sized open-weights models.
- Open Source No usage restrictions, code and data made available.

https://huggingface.co/utter-project/EuroLLM-9B-Instruct

Data Mix

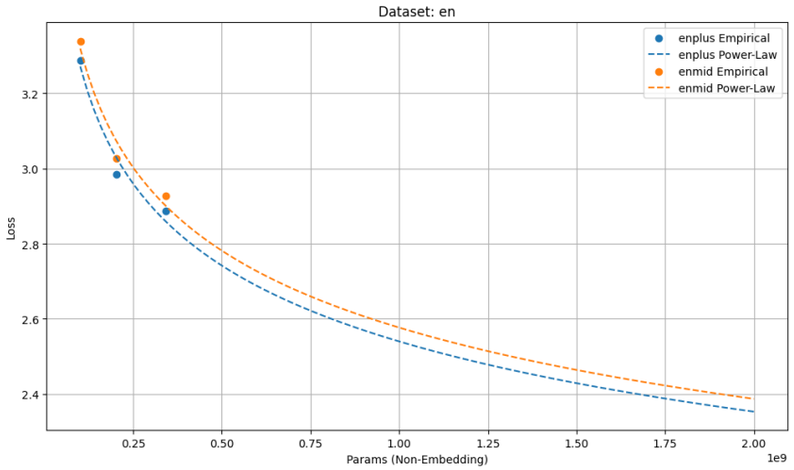

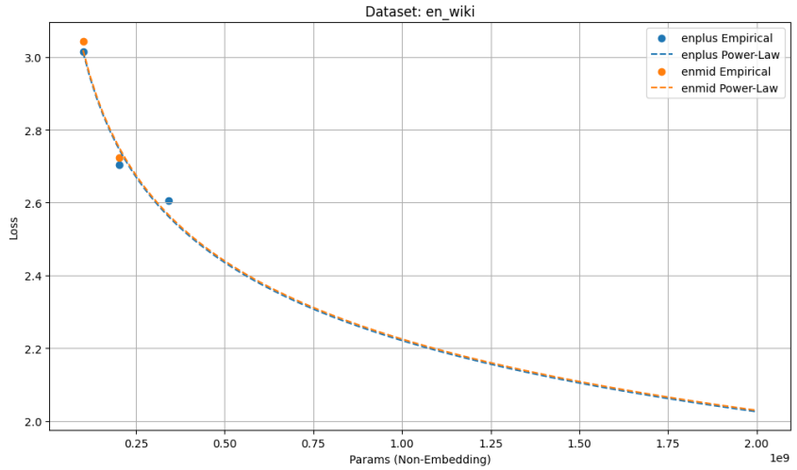

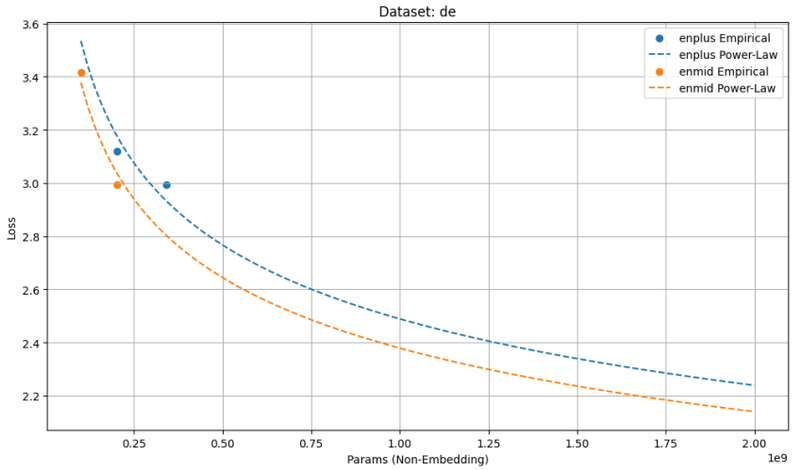

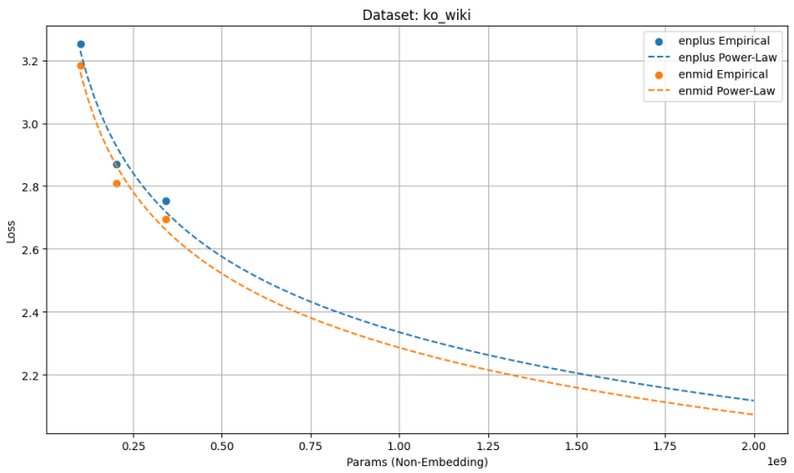

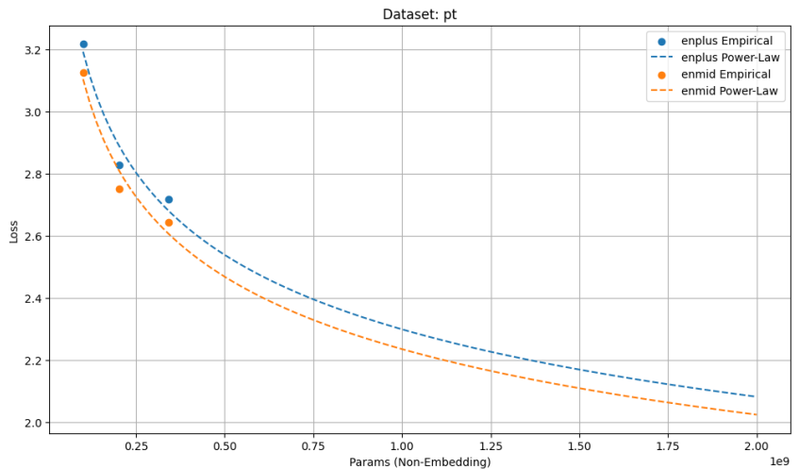

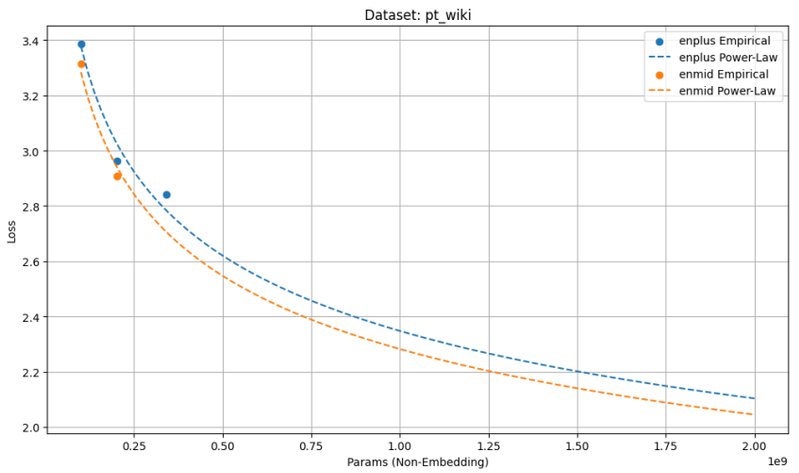

Scaling law: How much English?

enmid - 33%

enplus - 50%

Tokenisation

EuroLLM	128k
Mistral	32K
Llama3	128k
Teuken/Salamadra	250k

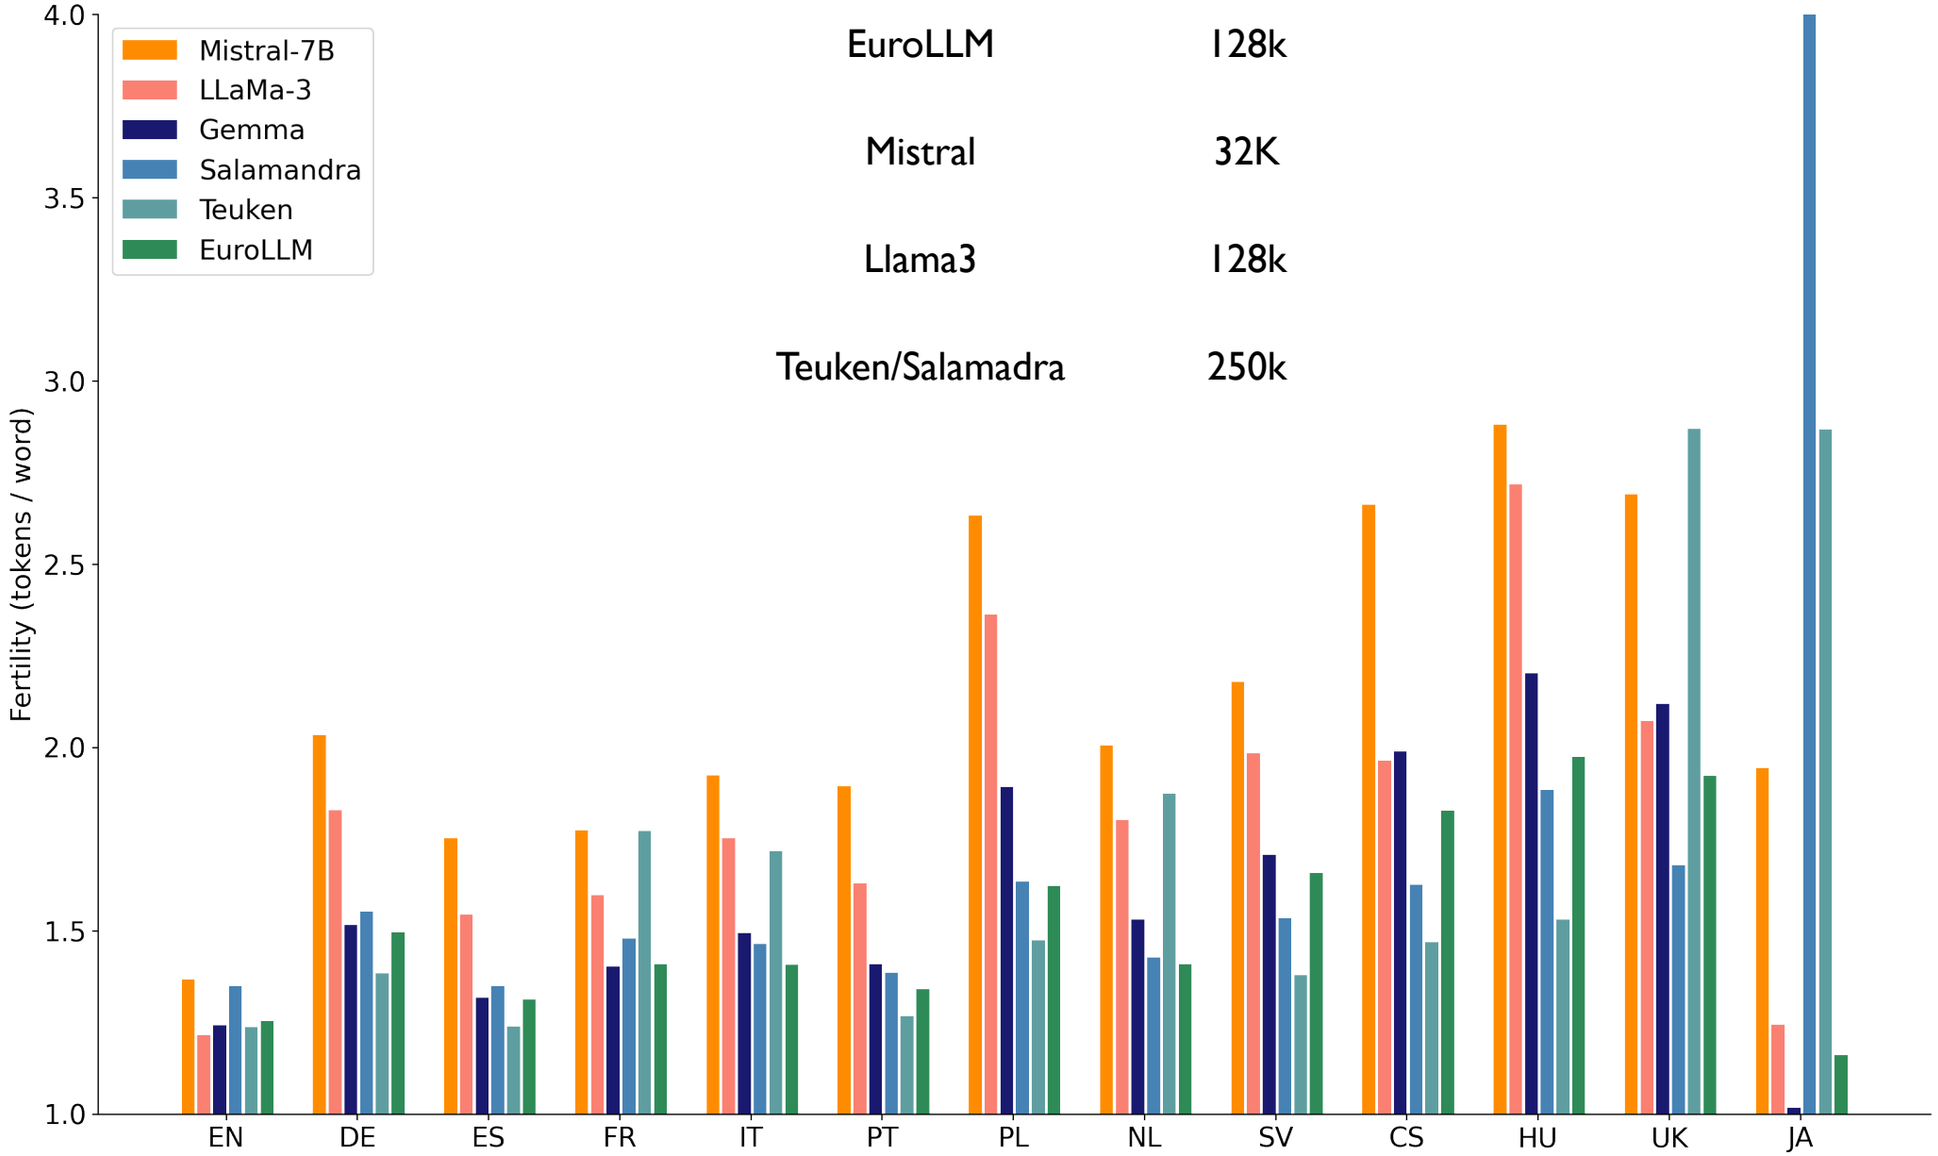

Evaluation

Averaged EU languages, 3 and 46 translation directions

Post-trained	**Arc-C** (25-shot)	**Hellaswag** (10-shot)	**MMLU** (5-shot)	**MMLU-pro** (5-shot)	**MUSR** (0-shot)	**WMT-24** (0-shot)	**FLORES** (0-shot)	**Borda C ↓**
Non-European								
Gemma-2-9B-IT	**57.98**	**66.95**	**63.07**	27.42	**8.38**	79.82	**86.82**	**1.3**
LLaMa-3.1-8B-IT	52.75	62.40	57.53	24.22	4.01	78.94	84.85	3.0
Granite-3-8B-IT	42.44	55.85	50.15	20.10	7.90	72.18	72.25	4.4
Qwen-2.5-7B-IT	47.09	57.73	62.86	**29.68**	7.62	75.96	76.97	3.1
OLMo-2-7B-IT	40.81	52.02	45.65	12.38	4.02	69.24	71.47	5.9
Aya-Expanse-8B	47.40	61.84	53.58	19.77	5.52	**83.01**	77.73	3.3
European								
Mistral-7B-IT	50.39	61.46	50.75	**18.19**	6.94	75.11	77.98	4.0
Ministral-8B-IT	48.67	61.62	51.55	17.41	6.17	77.13	81.34	3.9
Occiglot-7B-eu5-IT	42.13	59.49	42.08	11.77	4.17	75.10	74.40	6.1
Salamandra-7B-IT	44.69	63.60	44.60	7.01	7.17	80.87	87.35	3.9
Pharia-1-LLM-7B-C	40.55	55.22	39.91	10.10	**9.83**	63.80	58.91	6.4
Teuken-7B-IT-R-v0.4	46.84	62.75	39.81	9.29	2.25	77.91	82.63	5.3
Teuken-7B-IT-C-v0.4	46.28	62.73	41.74	9.79	2.94	77.68	84.41	5.0
EuroLLM-9B-IT	**56.55**	**67.53**	**52.97**	17.04	9.02	**83.61**	**88.87**	**1.4**

So do we need MT?

- MT models and LLM models have converged with differences in: size, encode/decoder, amount of monolingual and parallel data in pre-training and finetuning

- LLMs are robust, controllable and produce excellent MT performance when fine tuned, need far less parallel data

- But MT models use far less data overall, and are much smaller, use less compute to train and run and produce highest quality literal translations in the right language

Conclusion

- A huge number of research problems now possible on the border between translations and generation
- MT unique generative NLP task with lessons for the field:
  - Large amounts of labelled data: both translations and evaluations
  - Maturer understanding of evaluation and human interaction
  - Compelling task that can benefit humanity