

Gradual Code-Switching as Inference-Time Cross-Lingual Representational Alignment for LLMs

Haneul Yoo[♡] Jiho Jin[♣] Kyunghyun Cho[♡] Alice Oh[♣]
[♡]New York University [♣]KAIST
 haneul.yoo@nyu.edu

Abstract

While large language models (LLMs) have achieved notable progress in multilingual settings, their performance remains uneven across languages as LLMs often rely on English-centric latent representations. In this work, we introduce code-switching in-context learning (**CSICL**), an inference-time mechanism for cross-lingual representational alignment. Rather than relying on an abrupt translation pivot, **CSICL** explicitly scaffolds the reasoning trajectory by gradually transitioning from a target language to English, aligning non-English inputs with an English-centric reasoning space. Across 4 LLMs, 6 datasets, and 10 languages, **CSICL** consistently outperforms cross-lingual in-context learning baselines, yielding average gains of 6.0pp and 4.8pp in target and unseen languages, respectively. The improvements generalize across language families and are even more pronounced in low-resource settings, with gains of 14.7pp in target and 5.3pp in unseen languages. These findings establish code-switching as a robust and effective approach for reducing cross-lingual misalignment during inference, moving LLMs towards more equitable and effective multilingual systems.

<https://github.com/haneul-yoo/csicl/>




Figure 1: Overview of **CSICL**, which gradually transitions from a target language to English to better align cross-lingual representations.

1 Introduction

Large language models (LLMs) have demonstrated remarkable multilingual capabilities, powering diverse tasks such as question answering (Tjauatja et al., 2024), translation (Xu et al., 2024), and reasoning (Shi et al., 2023). However, recent studies reveal that this competence is far from language-agnostic, as LLMs typically rely on English-centric latent representations (Wendler et al., 2024; Zhong et al., 2025). Consequently, cross-lingual misalignment,

such as failures in internal translations into English, can lead to a sharp drop in non-English performance (Bafna et al., 2025). This highlights that multilingual competence in LLMs remains uneven, thereby limiting the inclusive deployment of these models across diverse linguistic communities around the world.

A common approach to improving multilingual performance is cross-lingual in-context learning (X-ICL), which uses demonstrations in high-resource languages (e.g., English) (Lin et al., 2022) to elicit task behavior in low-resource (Cahyawijaya et al., 2024) or unseen (Winata et al., 2022) languages. However, existing X-ICL often transfers superficial task patterns (e.g., output format, label mapping, etc.), employing monolingual demonstrations as an abrupt single-step pivot without necessarily reducing the underlying cross-lingual representational gap.

Meanwhile, code-switching, a natural alternation between languages, has been shown to benefit cross-lingual transfer in training phases (Wang et al., 2025c; Yoo et al., 2025a; Chai et al., 2025), pointing to its potential to improve cross-lingual alignment at inference time. In this paper, we propose code-switching in-context learning (**CSICL**), a novel inference-time representational alignment mechanism. Instead of relying on an abrupt single-step transition pivot, **CSICL** explicitly scaffolds the reasoning trajectory of LLMs through a gradual transition from a target language to English during inference (Figure 1). Specifically, **CSICL** begins in the target language, progressively introduces English tokens via code-switching, and converges to a full English equivalent. This gradual shift acts as a *linguistic bridge*, nudging LLMs to dynamically align cross-lingual representations directly, rather than relying solely on an abrupt latent translation.

We evaluate **CSICL** across 4 multilingual LLMs in 10 languages on 6 datasets spanning knowledge-intensive and reasoning-oriented tasks. We observe that **CSICL** consistently outperforms X-ICL baselines, yielding improvements of 6.0pp and 4.8pp in target and unseen languages, respectively. We generalize this to diverse tasks and highlight its effectiveness, particularly in translation and reasoning tasks. These findings position **CSICL** as a robust, effective mechanism for mitigating cross-lingual misalignment that controls gradual cross-lingual transitions at inference time.

Our contributions are as follows:

- We propose **CSICL**, an inference-time cross-lingual representation alignment mechanism that gradually bridges target languages with English representations during inference.
- We provide systematic empirical evidence across 4 LLMs, 6 tasks, and 10 languages that **CSICL** consistently improves cross-lingual transfer over standard X-ICL baselines, especially in unseen and low-resource languages.
- We visualize hidden states where **CSICL** systematically aligns non-English inputs with the English-centric latent space. We further analyze key design factors of **CSICL**, including granularity, directionality, pivot language, and generation strategy.

2 Code-Switching In-Context Learning

We propose **CSICL**, a training-free inference-time method for improving cross-lingual transfer of LLMs. Figure 1 illustrates an abstract example of **CSICL**. Given (1) a gradual translation instruction and (2) gradual code-switching few-shot demonstrations, **CSICL** constructs a stepwise transition from a target language to English.

2.1 Inference-Time Transition via Gradual Code-Switching

Instruction. We instruct LLMs to process a non-English input through a gradual transition into English. Specifically, we ask them to

1. Repeat a short, single-sentence instruction that guides a model to gradual translation (i.e., “Let’s gradually translate this non-English query into English, then think in English, and finally answer the question.”);
2. Explicitly show their step-by-step, progressive translation process into English;

3. Provide the final answer in the target language, following the specified format (e.g., “The answer is ...”).

Demonstrations. In addition to the instruction, **CSICL** provides few-shot demonstrations that instantiate the same transition pattern, where each demonstration:

1. Begins with a query in a target language (*i.e.*, En 0%);
2. Progressively transitions to English by leveraging code-switching whose matrix language is the target language and the embedded language is English (*i.e.*, En 25→50→75%);
3. Finally concludes with the full English equivalent (*i.e.*, En 100%).

2.2 Demonstration Construction

For each target language, we randomly sample 5 instances as demonstrations and convert them into gradual code-switching sequences. Figure 2 illustrates the two-step pipeline to construct these demonstrations in **CSICL**. Following Kim et al. (2025), we first generate the code-switching sentences given parallel inputs, and then construct a gradual code-switching sequence increasing the English proportion. While our default implementation uses GPT-5 for demonstration construction, **CSICL** is not fundamentally restricted to an external oracle. We evaluate alternative strategies, including self-generated and rule-based demonstrations.




Figure 2: Two-step pipeline to generate gradual code-switching few-shot demonstrations.

3 Experimental Setup

3.1 Evaluation Protocol

Following Lin et al. (2022), we use English prompts with cross-lingual 5-shot demonstrations. For each experiment, we set 3 target languages in total—*i.e.*, French (*high*), Korean (*mid*), Yoruba (*low*)—and report performance in: (1) the target language and (2) random, unseen languages that do not appear in the demonstrations. Evaluating unseen languages is particularly important, as real-world multilingual systems are often deployed in settings where the test language rarely appears in the training or demonstration phase.

We compare **CSICL** against five conventional X-ICL baselines following Asai et al. (2024):

- **Monolingual** prompting in either English or a target language (Lin et al., 2022; Tanwar et al., 2023);
- **Parallel** demonstrations in both the target language and English (Kim et al., 2024);
- **Translation**-based prompting from the target language into either English (Huang et al., 2023) or a random pivot language (Chai et al., 2025).

3.2 Evaluation Models

We use four state-of-the-art multilingual LLMs (*i.e.*, two open and two proprietary models): Qwen3-32B (Yang et al., 2025), deepseek-chat-v3.1 (DeepSeek-AI et al., 2025), grok-4-fast (xAI, 2025), and Gemini 2.5 Flash (Comanici et al., 2025).

3.3 Datasets and Metrics

We use Global MMLU (Singh et al., 2025) as a primary benchmark, covering 10 languages and diverse subject domains. In particular, the languages include English, three target languages from three distinct language families, and six unseen languages across six different language families spanning resource levels—Chinese, Spanish (*high*); Indonesian, Turkish (*mid*); Swahili, Telugu (*low*).

We further evaluate cross-lingual transfer across tasks and knowledge domains using five additional datasets, spanning translation, reasoning, and knowledge tasks. We fix Spanish as the target and select one mid- to high-resource unseen language per dataset (shown in parentheses):

- FLORES+ (Costa-jussà et al., 2024) for machine translation to English (Japanese);
- MedExpQA (Alonso et al., 2024) for domain-specific (*i.e.*, medical) reasoning (French);
- PolyMath (Wang et al., 2025b) for mathematical reasoning (Chinese);
- BLENd (Myung et al., 2024) for cultural knowledge (Korean);
- MBBQ (Neplenbroek et al., 2024) for social bias (Dutch).

We use accuracy, exact match (EM), and COMET (Rei et al., 2022) as evaluation metrics for multiple-choice questions (Global MMLU and MBBQ), short-answer questions (MedExpQA, PolyMath, and BLENd), and translation (FLORES+), respectively.

3.4 Statistical Significance Testing

Due to limited computational resources, we conduct the following experiments as a single run with a fixed seed. Instead, we apply bootstrap resampling with 2,000 iterations over the evaluation set to rigorously evaluate whether **CSICL** significantly outperforms the baseline systems. In each iteration, we compute the performance difference between **CSICL** and a given baseline and derive a 95% percentile bootstrap confidence interval (CI) from the resulting distribution.

4 Results and Analyses

4.1 Main Results

Table 1 reports the macro-averaged results of **CSICL** on Global MMLU, where the full results are stated in Appendix D.4. Table 9 shows additional experimental results controlling the combination of demonstration and instruction settings. **CSICL** outperforms existing X-ICL baselines in both target and unseen languages. Specifically, **CSICL** produces 6.0pp higher performance than monolingual demonstrations in the target language, while the performance gap in English remains within 0.2pp. In addition, **CSICL** shows substantial cross-lingual transfer in unseen, low-resource languages (+4.8pp), highlighting its practical value in multilingual scenarios. We observe that monolingual demonstrations in the target language yield higher performance, whereas the benefit diminishes in unseen languages. This contrast highlights a key limitation of monolingual prompting: it overfits the demonstrated language rather than fostering cross-lingual alignment. In contrast, **CSICL** explicitly facilitates this alignment through gradual code-switching, resulting in stronger transfer to unseen settings. We do not observe significant differences in the *out-of-format* ratio, which remains below 0.05% for all X-ICL settings. We further demonstrate the effectiveness of **CSICL** across model sizes in Appendix D.1 (Figure 6). We also control for prompt budget and structural differences among X-ICL baselines in Appendix D.2, all of which still underperform **CSICL**.

Method	X-ICL setting		En	Tgt.*	Unseen Lang.		
	Demonstration	Instruction			High*	Mid*	Low*
Zero-shot learning	✗	✗	88.6	68.6	86.2	62.1	39.4
Few-shot learning	✓ Monolingual (En)	✗	88.8	70.8	86.5	62.8	41.2
	✓ Monolingual (Tgt.)	✗	88.8	72.0	86.9	62.1	38.7
	✓ Parallel	✗	88.7	72.7	87.1	63.0	41.4
Zero-shot CoT	✗	✓ Translation (Tgt.→En)	88.8	74.5	87.4	63.7	42.0
	✗	✓ Translation (Tgt.→Rnd.)	88.6	73.8	<u>87.5</u>	<u>63.8</u>	<u>42.3</u>
CSICL	✓ Gradual CS (Tgt.→En)	✓ Gradual Translation (Tgt.→En)	88.6	76.8	87.8	64.9	46.0

Table 1: Experimental results comparing **CSICL** to existing X-ICL baselines using Qwen 3. X-ICL setting describes each method in terms of a combination of its few-shot demonstrations and instruction. **CSICL** outperforms X-ICL baselines in both target and unseen languages, yielding stable results in English.

Throughout this paper, **bold** and underline denote the best and the second-best results, respectively. Tgt. and Rnd. denote a target language and a random language, respectively. Asterisk indicates statistical significance against each and every baseline.

Models. Figure 3a displays experimental results of each model with X-ICL approaches in target languages on Global MMLU. All four models achieve similar trends across X-ICL settings, where **CSICL** significantly outperforms existing baselines.

Subject categories. Figure 3b shows experimental results of X-ICL approaches using Qwen 3 for six subject categories on Global MMLU. **CSICL** significantly surpasses baselines across all categories.




Figure 3: Experimental results of **CSICL** compared to X-ICL approaches.

Languages. Figure 4 presents performance differences (pp) of each X-ICL setting compared to zero-shot learning, whose target languages are French, Korean, and Yoruba, respectively. **CSICL** outperforms X-ICL baselines, particularly in target languages and mid- to low-resource unseen languages, while English performance remains stable. The benefits of **CSICL** extend beyond script similarity or language family, implying that it fosters language-agnostic cross-lingual alignment rather than relying on superficial linguistic overlap.

4.2 Task and Domain Knowledge

Table 2 presents experimental results of X-ICL approaches using Qwen 3 across five additional tasks spanning translation, reasoning, and knowledge evaluation. Full experimental results with the other three models are presented in Appendix D.4. In line with earlier results, **CSICL** consistently outperforms existing baselines across settings. Following the descriptions provided in the original papers, we categorize tasks into three types: machine translation




Figure 4: Performance differences (pp) of **CSICL** and X-ICL baselines compared to zero-shot learning setting per a target language using Qwen 3.

(FLORES+); reasoning-oriented tasks (MedExpQA, PolyMath); and knowledge-intensive tasks (BLENd, MBBQ).

Task type	Translation		Reasoning-oriented						Knowledge-intensive					
			Domain-specific			Math			Cultural			Social Bias		
X-ICL setting	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.
Monolingual (En)	76.6	72.3	46.4	38.9	32.1	51.2	49.3	47.4	83.7	77.8	73.2	92.0	88.9	86.3
Monolingual (Tgt.)	75.3	71.6	45.8	38.5	30.3	50.5	49.7	46.1	83.1	78.9	71.5	91.6	89.3	85.8
Parallel	78.1	73.0	<u>46.3</u>	<u>40.2</u>	31.9	51.4	49.7	47.7	83.8	<u>79.2</u>	72.0	91.5	89.5	86.0
Translation (Tgt.→En)	74.8	72.1	46.1	39.6	32.3	50.8	50.8	48.2	83.3	78.5	71.7	92.1	90.1	86.4
Translation (Tgt.→Rnd.)	73.2	<u>73.7</u>	46.3	40.1	<u>32.7</u>	50.3	<u>51.0</u>	<u>47.9</u>	83.6	78.1	72.4	91.9	<u>89.8</u>	87.3
CSICL (Tgt.→En)	83.4	75.3	46.2	44.4	35.2	50.9	54.5	51.8	83.5	81.3	74.5	91.8	90.5	<u>87.2</u>

Table 2: Experimental results on 5 diverse tasks and domain knowledge using Qwen 3. **CSICL** achieves larger improvements in translation and reasoning tasks.

Translation. **CSICL** achieves the largest gains in the translation task—gains of +6.8pp and +3.0pp COMET in the target and the unseen languages, respectively, over monolingual demonstrations. This strong effect is expected, as progressive transition directly scaffolds the latent translation process, aligning the input with English-centric representations. It is noteworthy that **CSICL** is effective even in a purely generative setting.

Reasoning. Beyond translation, **CSICL** also improves reasoning-oriented tasks, with average gains of +5.4pp and +3.8pp of EM over monolingual baselines. We suppose that gradual transitions mitigate representational interference: instead of reasoning entirely in a non-English space, the demonstrations nudge the model to “think in English,” where reasoning capabilities are strongest. This highlights the potential of **CSICL** as a lightweight test-time alignment mechanism, complementary to scaling methods such as self-consistency.

Knowledge. For cultural knowledge and social bias, the improvements are modest but still consistent (+2.6pp and +1.1pp in the target and the unseen languages, respectively, on average). This indicates that **CSICL** not only facilitates translation-like tasks but also improves general factual consistency across languages.

4.3 What Makes CSICL Effective?

Layer-wise visualization of CSICL. We analyze whether **CSICL** helps an LLM resolve cross-lingual misalignment by explicitly facilitating its latent transition to English representations. We evaluate three input conditions: Korean (Ko), English equivalent (En), and **CSICL** applied to the same Korean set (**CSICL**). For each condition, we run forward passes and extract hidden representations from every layer, focusing on the test query span without few-shot contexts via span-based pooling to obtain a single vector per layer per example. We then visualize how these representations evolve across depth by projecting layer-wise vectors with PCA, plotting matched triples (*i.e.*, En, Ko, **CSICL**) in a shared space. Figure 5 visualizes

the representations of Qwen3-32B, compared to English and Korean results as anchors. We reveal that **CSICL** systematically shifts the non-English inputs toward the English anchor.




Figure 5: PCA visualizations of Qwen 3 with **CSICL** on Global MMLU, compared to English results and Korean results as anchors.

We further examine which design factors contribute to the gains of **CSICL** and report the experimental results on Global MMLU using Qwen 3 in Table 3.

Demonstrations vs. instruction. We observe that both gradual code-switching demonstrations and translation instruction contribute to cross-lingual transfer, by yielding 3.4pp and 3.8% higher accuracy compared to the monolingual baseline. In particular, gradual code-switching consistently achieves higher performances than code-switching (+0.5pp), parallel (+1.5pp), and monolingual demonstrations (+3.4pp). Zero-shot and gradual translation instructions yield minimal differences within 0.1pp in the target language, as the model fails to understand the concept of gradual translation and to follow it without explicit demonstrations. Specifically, zero-shot gradual translation instruction often collapses into abrupt translation, where the model outputs two sentences entirely in Korean, followed by three equivalent sentences entirely in English.

Directionality. Furthermore, transitioning from a target language to English outperforms the reverse direction, supporting our hypothesis that **CSICL** alleviates the translation barrier by scaffolding the latent translation process of LLMs. This asymmetry indicates that LLMs benefit more when demonstrations converge to English representations, aligning with their latent space, rather than diverging away from it.

Method	X-ICL setting		En	Tgt.*	Unseen Lang.		
	Demonstration	Instruction			High*	Mid*	Low*
Few-shot learning	✓ CS (En+Tgt.)	✗	88.6	73.6	87.1	62.9	43.3
	✓ CS (Tgt.+En)	✗	88.7	73.7	87.0	62.8	43.1
	✓ Gradual CS (En→Tgt.)	✗	88.7	74.0	87.6	63.8	44.5
	✓ Gradual CS (Tgt.→En)	✗	88.6	74.2	87.7	64.3	45.7
Zero-shot CoT	✗	✓ Gradual Translation (Tgt.→En)	88.7	74.6	87.2	63.8	42.0
CSICL	✓ Gradual CS (En→Tgt.)	✓ Gradual Translation (En→Tgt.)	88.7	75.0	87.6	64.0	45.2
	✓ Gradual CS (Tgt.→En)	✓ Gradual Translation (Tgt.→En)	88.6	76.8	87.8	64.9	46.0

Table 3: Ablation results to examine the effectiveness of each design factor in **CSICL**. Gradual transition demonstrations, instruction, and direction contribute to cross-lingual transfer.

Transition granularity. Table 4 shows experimental results of **CSICL** on Global MMLU using Qwen 3, controlling the number of gradual transition steps and the number of few-shot demonstrations. We also compare linear vs. non-linear schedules of gradual code-switching. We observe that the more steps in gradual transition and the larger number of demonstrations result in better target and unseen language performance. However, the performance gap between 5-step and 7-step transitions is minimal, and even mixed. This suggests that excessively fine-grained switching may exceed the model’s controllable generation capability, while gradual alignment is beneficial. Similarly, performance increases

up to 5 shots and then saturates, indicating that gradual transition behavior stabilizes once sufficiently demonstrated. Additionally, linear transition consistently outperforms non-linear variants, indicating that smooth interpolation is more effective than abrupt jumps.

CSICL (5 shots)	En	Tgt.	Unseen Lang.			CSICL (5 steps)	En	Tgt.	Unseen Lang.		
			High	Mid	Low				High	Mid	Low
0-25-100 (3 steps)	88.6	74.5	87.0	63.9	42.4	0-25-50-75-100 (0 shot)	88.7	74.6	87.2	63.8	42.0
0-75-100 (3 steps)	88.6	74.3	87.2	63.8	42.3	0-25-50-75-100 (1 shot)	88.7	75.3	87.3	64.2	44.6
0-50-100 (3 steps)	88.7	75.1	87.7	64.0	44.4	0-25-50-75-100 (3 shots)	88.8	76.2	87.6	64.7	45.3
0-25-75-100 (4 steps)	88.6	76.0	87.7	64.5	45.3	0-25-50-75-100 (5 shots)	<u>88.6</u>	76.8	87.8	64.9	46.0
0-25-50-75-100 (5 steps)	<u>88.6</u>	76.8	87.8	64.9	46.0	0-25-50-75-100 (7 shots)	<u>88.6</u>	<u>77.0</u>	<u>87.7</u>	<u>65.0</u>	<u>46.1</u>
0-16-33-50-66-83-100 (7 steps)	88.5	76.9	87.5	65.1	45.8	(b) # of demonstrations					
(a) # of steps											

Table 4: Experimental results of **CSICL** using Qwen 3, controlling the number of transition steps and the number of few-shot demonstrations.

4.4 Is CSICL Practically Useful?

Cost-performance trade-off. To quantify the cost-performance trade-off, we compute the average number of input/output tokens for X-ICL approaches using the Qwen3-32B tokenizer on Global MMLU. Table 5 summarizes cost-performance trade-off. As expected, the full **CSICL** configuration with five few-shot demonstrations and a five-step transition uses more tokens than simpler X-ICL baselines, while yielding the highest gains. However, its average token budget of 1,446.7 remains well within practical context limits for Qwen 3, whose native context window is 32,768 tokens and can be extended to 131,072 with YaRN. While longer prompts directly affect latency and costs in API-based deployments, the overhead is comparable to that of other few-shot approaches rather than an order-of-magnitude increase. Importantly, our ablation with fewer-shot variants shows that even 0-shot and 1-shot versions of **CSICL** already outperform strong X-ICL baselines while using fewer tokens than existing X-ICL alternatives. These results suggest that **CSICL** offers a favorable and tunable efficiency-accuracy trade-off: 0-1 shot is attractive for budget-constrained settings, whereas 3+ shots are preferable when maximizing accuracy is the priority.

X-ICL setting	Avg. input tokens	Avg. output tokens	Tgt. performance
Zero-shot learning	86.9	4.8	68.6
Parallel	651.7	131.6	72.7
Translation (Tgt.→En)	438.8	133.2	74.5
CSICL (0 shot)	231.5	154.7	74.6
CSICL (1 shot)	387.6	184	75.3
CSICL (3 shots)	868.4	181.8	76.2
CSICL (5 shots)	1,265.1	181.6	76.8

Table 5: Cost-performance trade-off of X-ICL approaches considering both required tokens and target language performance.

Demonstration generation. We examine whether **CSICL** depends on a frontier model for few-shot demonstration generation. Here, we set up three alternative generation strategies: (1) using the same model (Qwen3-32B) of the X-ICL inference setup, (2) a small language model (gemma-3-1b) to lower the cost for few-shot generation, and (3) a rule-based method without relying on any external model. Table 6 reports the experimental results of **CSICL** with alternative strategies for few-shot demonstration generation. Self-generated demonstrations achieve nearly identical performance to GPT-5. This implies that **CSICL** enables a truly end-to-end multilingual system regardless of whether external models are used. While small model and rule-based approaches yield slightly lower performances, they still outperform other X-ICL baselines. This demonstrates and supports the feasibility and robustness of **CSICL**. Importantly, **CSICL** requires only five demonstrations to improve 14.7pp of target language performance. These handful of demonstrations can be produced automatically, as we have shown above, or even manually curated with minimal effort.

5 Related Work

5.1 Cross-lingual Transfer in LLMs

Recent studies suggest that LLMs tend to rely on English as a pivot for cross-lingual processing. Zhao et al. (2024) showed that LLMs exhibit language-specific neurons to process multilingual inputs, and Schut et al. (2025) demonstrated that LLMs internally translate multilingual inputs into English representations and think in English. This reliance on internal translation introduces a critical vulnerability: *translation barrier*, where a failure in the initial translation stage is propagated to the quality of the final outputs (Bafna et al., 2025). While external machine translation can slightly mitigate this issue (Etxaniz et al., 2024; Intrator et al., 2024), it undermines the goal of using LLMs as seamless, end-to-end multilingual systems. This motivates the exploration of mechanisms that can directly activate cross-lingual capabilities of LLMs without relying on latent English translation.

To bridge this gap, cross-lingual transfer has been widely studied (Pallucchini et al., 2025). For example, Tanwar et al. (2023); Huang et al. (2023); Li et al. (2024); Lin et al. (2025) investigated X-ICL approaches to better align cross-lingual representations. However, these approaches are primarily based on an abrupt single-step pivot, requiring a more effective cross-lingual signal that can mitigate this misalignment during inference.

5.2 Code-switching for Cross-lingual Transfer

Code-switching, which has been studied over decades by the natural language processing community (Doğruöz et al., 2021; Winata et al., 2023), is a promising candidate for such a direct signal. Code-switching has shown its efficacy for cross-lingual transfer in various training stages, including pre-training (Wang et al., 2025c), continual pre-training (Yoo et al., 2025a), supervised fine-tuning (SFT) (Lee et al., 2024; Chai et al., 2025; Son et al., 2026), and selective constrained decoding (Li et al., 2025). In particular, Chai et al. (2025) demonstrated that (1) SFT using inter-sentential code-switching data and (2) chain-of-thought prompting, which translates a non-English query into a random language and answers in the target language, enhances cross-lingual transfer. Li et al. (2025) introduced probe-guided decoding with an additional classifier to enhance bilingual LLM reasoning with a case study of English-Chinese code-switching. Son et al. (2026) also demonstrated that SFT using code-switching instruction data enhances LLM reasoning. Kang et al. (2026) suggested to selectively incorporate English translations into the latent reasoning of LLMs using additional classifiers.

However, these approaches require substantial additional training resources, while exploration remains limited at the inference level, often focusing only on a specific task and domain knowledge. For machine reading comprehension, Kim et al. (2024) provided a passage in English and a QA pair in a target language, which is a form of intra-sentential code-switching. Kim et al. (2025) showcased a case study where Korean-English inter-sentential code-switching can unlock Korean cultural knowledge that remains inaccessible in English queries. Wang et al. (2025a) investigated patterns, impact, and internal causes of code-switching in reasoning language models. Our work, in contrast, moves beyond utilizing code-switching merely as a static context. We introduce an inference-time representational alignment mechanism that systematically controls the reasoning trajectory. By gradually transitioning from the target language to English, CSICL explicitly scaffolds the latent space of LLMs, improving cross-lingual transfer during inference.

Demonstrations	En	Tgt.	Unseen Lang.		
			High	Mid	Low
GPT-5	88.6	76.8	87.8	64.9	46.0
Same model	88.7	76.7	87.8	64.6	45.8
SLM	88.6	75.5	87.5	64.2	44.3
Rule-based	88.6	75.2	87.5	64.4	44.0

Table 6: Experimental results of CSICL using Qwen 3 with few-shot demonstrations generated by three alternative strategies: (1) the same model of the X-ICL inference setup, (2) a small language model (SLM), and (3) a rule-based method without any external model.

6 Conclusion

We address cross-lingual misalignment in multilingual LLMs—their over-reliance on English-centric latent representations that often leads to degraded performance in non-English. To mitigate this issue, we introduce **CSICL**, an inference-time method that improves cross-lingual transfer by gradually transitioning from a target language to English through structured code-switching. Across 4 multilingual LLMs, 6 datasets, and 10 languages, **CSICL** consistently outperforms standard X-ICL baselines, yielding average gains of 6.0pp and 4.8pp in both target and unseen languages, respectively. We demonstrate its pronounced improvements in low-resource settings, with gains of 14.7pp in target and 5.3pp in unseen languages. Beyond empirical improvements, this paper positions gradual code-switching as a new lens for cross-lingual alignment during LLM inference.

Acknowledgments

We thank Dongkwan Kim and Guijin Son for their valuable advice. This work was supported by Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. RS-2024-00509258 and No. RS-2024-00469482, Global AI Frontier Lab). This research was conducted as part of the Sovereign AI Foundation Model Project(GPU Track), organized by the Ministry of Science and ICT(MSIT) and supported by the National IT Industry Promotion Agency(NIPA), S.Korea. (PJT-26-010017). This research was supported by the MSIT(Ministry of Science, ICT), Korea, under the Top-Tier AI Global HRD invitation program (RS-2025-25461932) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation).

We use ChatGPT, Gemini, and Copilot for writing and coding assistance. We also use LLMs to generate few-shot demonstrations in the experiments (*e.g.*, gradual code-switching, translation, and paraphrasing).

Ethics Statement

All datasets used in this study are publicly available and employed solely for research purposes, consistent with their intended licenses. We prevent data leakage by excluding the samples used as few-shot demonstrations from the corresponding test set. **CSICL** aims to improve the multilingual performance of LLMs, with the intended use for both research and practical applications, which is not used for malicious purposes.

Reproducibility Statement

For open-source LLM inference, we use 4 H200 GPUs with 564 GB memory and 4 RTX 8000 GPUs with 188 GB memory. For closed-source LLMs, we access them via OpenRouter¹. All experiments and sample selections are conducted with a random seed of 42. All COMET scores in this paper are COMET-22 (unbabel/wmt22-comet-da) (Rei et al., 2022). For reliable reproduction, we use greedy decoding with a temperature of 0.0, if possible. All system prompts used in the experiments are provided in Section B.

All evaluation datasets used in this paper are publicly available. Global MMLU (Singh et al., 2025) is released under the Apache-2.0 license. BLENd (Myung et al., 2024), FLORES+ (Costa-jussà et al., 2024), and MedExpQA (Alonso et al., 2024) are distributed under the CC-BY-SA-4.0 and CC-BY-4.0 licenses, respectively. The licenses of MBBQ (Neplenbroek et al., 2024) and PolyMath (Wang et al., 2025b) are not explicitly specified. The samples used as few-shot demonstrations are excluded from the corresponding test set to avoid data leakage.

We use a balanced set of 36,000 instances from Global MMLU by randomly sampling 600 questions per subject category and language. For MBBQ, we randomly select ten samples

¹<https://openrouter.ai/>

(five ambiguous contexts and five unambiguous contexts) per template in the MBBQ dataset, totaling 980 samples, due to limited computational resources for processing the full 10k+ instances. We employ the entire set for the other datasets.

In Appendix 4.4, we follow the experimental setup of Yoo et al. (2025a) for the rule-based method. Specifically, we employ a part-of-speech (POS) tagger and randomly replace a specified proportion of Korean nouns in Korean sentences with English words using a bilingual dictionary. In Appendix D.2, we follow the implementations of Huang et al. (2023) for the *translate-then-reason* baseline.

References

- Samir Abdaljalil, Erchin Serpedin, Khalid Qaraqe, and Hasan Kurban. Evaluating multilingual and code-switched alignment in LLMs via synthetic natural language inference. *arXiv preprint arXiv:2508.14735*, 2025. URL <https://arxiv.org/abs/2508.14735>.
- Iñigo Alonso, Maite Oronoz, and Rodrigo Agerri. MedExpQA: Multilingual benchmarking of large language models for medical question answering. *Artificial Intelligence in Medicine*, 155:102938, 2024. ISSN 0933-3657. doi: <https://doi.org/10.1016/j.artmed.2024.102938>. URL <https://www.sciencedirect.com/science/article/pii/S0933365724001805>.
- Akari Asai, Sneha Kudugunta, Xinyan Yu, Terra Blevins, Hila Gonen, Machel Reid, Yulia Tsvetkov, Sebastian Ruder, and Hannaneh Hajishirzi. BUFFET: Benchmarking large language models for few-shot cross-lingual transfer. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 1771–1800, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.100. URL <https://aclanthology.org/2024.naacl-long.100/>.
- Peter Auer. *Code-switching in conversation*. Routledge, London, England, 1998.
- Niyati Bafna, Tianjian Li, Kenton Murray, David R. Mortensen, David Yarowsky, Hale Sirin, and Daniel Khashabi. The translation barrier hypothesis: Multilingual generation with large language models suffers from implicit translation failure. In Kentaro Inui, Sakriani Sakti, Haofen Wang, Derek F. Wong, Pushpak Bhattacharyya, Biplab Banerjee, Asif Ekbal, Tanmoy Chakraborty, and Dharendra Pratap Singh (eds.), *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pp. 1541–1568, Mumbai, India, December 2025. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics. ISBN 979-8-89176-298-5. doi: 10.18653/v1/2025.ijcnlp-long.83. URL <https://aclanthology.org/2025.ijcnlp-long.83/>.
- Samuel Cahyawijaya, Holy Lovenia, and Pascale Fung. LLMs are few-shot in-context low-resource language learners. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 405–433, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.24. URL <https://aclanthology.org/2024.naacl-long.24/>.
- Linzhen Chai, Jian Yang, Tao Sun, Hongcheng Guo, Jiaheng Liu, Bing Wang, Xinnian Liang, Jiaqi Bai, Tongliang Li, Qiyao Peng, and Zhoujun Li. XCOT: Cross-lingual instruction tuning for cross-lingual chain-of-thought reasoning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(22):23550–23558, Apr. 2025. doi: 10.1609/aaai.v39i22.34524. URL <https://ojs.aaai.org/index.php/AAAI/article/view/34524>.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, Krishna Haridasan, Ahmed Omran, Nikunj Saunshi, Dara Bahri, Gaurav Mishra, Eric Chu, Toby Boyd, Brad Hekman, Aaron Parisi, Chaoyi Zhang,

Kornrathop Kawintiranon, Tania Bedrax-Weiss, Oliver Wang, Ya Xu, Ollie Purkiss, Uri Mendlovic, Ilai Deutel, Nam Nguyen, Adam Langley, Flip Korn, Lucia Rossazza, Alexandre Ramé, Sagar Waghmare, Helen Miller, Nathan Byrd, Ashrith Sheshan, Raia Hadsell, Sangnie Bhardwaj, Pawel Janus, Tero Rissa, Dan Horgan, Alvin Abdagic, Lior Belenki, James Allingham, Anima Singh, Theo Guidroz, Srivatsan Srinivasan, Herman Schmit, Kristen Chiafullo, Andre Elisseeff, Nilpa Jha, Prateek Kolhar, Leonard Berrada, Frank Ding, Xiance Si, Shrestha Basu Mallick, Franz Och, Sofia Erell, Eric Ni, Tejasi Latkar, Sherry Yang, Petar Sirkovic, Ziqiang Feng, Robert Leland, Rachel Hornung, Gang Wu, Charles Blundell, Hamidreza Alvari, Po-Sen Huang, Cathy Yip, Sanja Deur, Li Liu, Gabriela Surita, Pablo Duque, Dima Damen, Johnson Jia, Arthur Guez, Markus Mircea, Animesh Sinha, Alberto Magni, Paweł Stradowski, Tal Marian, Vlado Galić, Wenhua Chen, Hisham Husain, Achintya Singhal, Dominik Grewe, François-Xavier Aubet, Shuang Song, Lorenzo Blanco, Leland Rechis, Lewis Ho, Rich Munoz, Kelvin Zheng, Jessica Hamrick, Kevin Mather, Hagai Taitelbaum, Eliza Rutherford, Yun Lei, Kuangyuan Chen, Anand Shukla, Erica Moreira, Eric Doi, Berivan Isik, Nir Shabat, Dominika Rogozińska, Kashyap Kolipaka, Jason Chang, Eugen Vušak, Srinivasan Venkatachary, Shadi Noghbi, Tarun Bharti, Younghoon Jun, Aleksandr Zaks, Simon Green, Jeshwanth Challagundla, William Wong, Muqthar Mohammad, Dean Hirsch, Yong Cheng, Iftekhar Naim, Lev Proleev, Damien Vincent, Aayush Singh, Maxim Krikun, Dilip Krishnan, Zoubin Ghahramani, Aviel Atlas, Rajeev Aggarwal, Christo Kirov, Dimitrios Vytiniotis, Christy Koh, Alexandra Chronopoulou, Pawan Dogra, Vlad-Doru Ion, Gladys Tyen, Jason Lee, Felix Weissenberger, Trevor Strohman, Ashwin Balakrishna, Jack Rae, Marko Velic, Raoul de Liedekerke, Oded Elyada, Wentao Yuan, Canoe Liu, Lior Shani, Sergey Kishchenko, Bea Alessio, Yandong Li, Richard Song, Sam Kwei, Orion Jankowski, Aneesh Pappu, Youhei Namiki, Yenai Ma, Nilesch Tripuraneni, Colin Cherry, Marissa Ikonmidis, Yu-Cheng Ling, Colin Ji, Beka Westberg, Auriel Wright, Da Yu, David Parkinson, Swaroop Ramaswamy, Jerome Connor, Soheil Hassas Yeganeh, Snchit Grover, George Kenwright, Lubo Litchev, Chris Apps, Alex Tomala, Felix Halim, Alex Castro-Ros, Zefei Li, Anudhyan Boral, Pauline Sho, Michal Yarom, Eric Malmi, David Klinghoffer, Rebecca Lin, Alan Ansell, Pradeep Kumar S, Shubin Zhao, Siqi Zuo, Adam Santoro, Heng-Tze Cheng, Solomon Demmessie, Yuchi Liu, Nicole Brichtova, Allie Culp, Nathaniel Braun, Dan Graur, Will Ng, Nikhil Mehta, Aaron Phillips, Patrik Sundberg, Varun Godbole, Fangyu Liu, Yash Katariya, David Rim, Mojtaba Seyedhosseini, Sean Ammirati, Jonas Valfridsson, Mahan Malihi, Timothy Knight, Andeep Toor, Thomas Lampe, Abe Ittycheriah, Lewis Chiang, Chak Yeung, Alexandre Fréchette, Jinmeng Rao, Huisheng Wang, Himanshu Srivastava, Richard Zhang, Rocky Rhodes, Ariel Brand, Dean Weesner, Ilya Figotin, Felix Gimeno, Rachana Fellingner, Pierre Marcenac, José Leal, Eyal Marcus, Victor Cotruta, Rodrigo Cabrera, Sheryl Luo, Dan Garrette, Vera Axelrod, Sorin Baltateanu, David Barker, Dongkai Chen, Horia Toma, Ben Ingram, Jason Riesa, Chinmay Kulkarni, Yujing Zhang, Hongbin Liu, Chao Wang, Martin Polacek, Will Wu, Kai Hui, Adrian N Reyes, Yi Su, Megan Barnes, Ishaan Malhi, Anfal Siddiqui, Qixuan Feng, Mihai Damaschin, Daniele Pighin, Andreas Steiner, Samuel Yang, Ramya Sree Boppana, Simeon Ivanov, Arun Kandoor, Aditya Shah, Asier Mujika, Da Huang, Christopher A. Choquette-Choo, Mohak Patel, Tianhe Yu, Toni Creswell, Jerry, Liu, Catarina Barros, Yasaman Razeghi, Aurko Roy, Phil Culliton, Binbin Xiong, Jiaqi Pan, Thomas Strohmann, Tolly Powell, Babi Seal, Doug DeCarlo, Pranav Shyam, Kaan Katircioglu, Xuezhi Wang, Cassidy Hardin, Immanuel Odisho, Josef Broder, Oscar Chang, Arun Nair, Artem Shtefan, Maura O'Brien, Manu Agarwal, Sahitya Potluri, Siddharth Goyal, Amit Jhinal, Saksham Thakur, Yury Stuken, James Lyon, Kristina Toutanova, Fangxiaoyu Feng, Austin Wu, Ben Horn, Alek Wang, Alex Cullum, Gabe Taubman, Disha Shrivastava, Chongyang Shi, Hamish Tomlinson, Roma Patel, Tao Tu, Ada Maksudaj Oflazer, Francesco Pongetti, Mingyao Yang, Adrien Ali Taiga, Vincent Perot, Nuo Wang Pierse, Feng Han, Yoel Drori, Iñaki Iturrate, Ayan Chakrabarti, Legg Yeung, Dave Dopson, Yi ting Chen, Apoorv Kulshreshtha, Tongfei Guo, Philip Pham, Tal Schuster, Junquan Chen, Alex Polozov, Jinwei Xing, Huanjie Zhou, Praneeth Kacham, Doron Kukliansky, Antoine Miech, Sergey Yaroshenko, Ed Chi, Sholto Douglas, Hongliang Fei, Mathieu Blondel, Preethi Myla, Lior Madmoni, Xing Wu, Daniel Keysers, Kristian Kjems, Isabela Albuquerque, Lijun Yu, Joel D'sa, Michelle Plantan, Vlad Ionescu, Jaume Sanchez Elias, Abhirut Gupta, Manish Reddy Vuyyuru, Fred Alcober, Tong Zhou, Kaiyang Ji, Florian Hartmann, Subha Puttagunta, Hugo Song, Ehsan Amid, Anca Stefanoiu, Andrew Lee, Paul Pucciarelli,

Emma Wang, Amit Raul, Slav Petrov, Isaac Tian, Valentin Anklin, Nana Nti, Victor Gomes, Max Schumacher, Grace Vesom, Alex Panagopoulos, Konstantinos Bousmalis, Daniel Andor, Josh Jacob, Yuan Zhang, Bill Rosgen, Matija Kecman, Matthew Tung, Alexandra Belias, Noah Goodman, Paul Covington, Brian Wieder, Nikita Saxena, Elnaz Davoodi, Muhuan Huang, Sharath Maddineni, Vincent Roulet, Folawiyo Campbell-Ajala, Pier Giuseppe Sessa, Xintian Wu, Guangda Lai, Paul Collins, Alex Haig, Vytenis Sakenas, Xiaowei Xu, Marissa Giustina, Laurent El Shafey, Pichi Charoenpanit, Shefali Garg, Joshua Ainslie, Boone Severson, Montse Gonzalez Arenas, Shreya Pathak, Sujee Rajayogam, Jie Feng, Michiel Bakker, Sheng Li, Nevan Wichers, Jamie Rogers, Xinyang Geng, Yeqing Li, Rolf Jagerman, Chao Jia, Nadav Olmert, David Sharon, Matthew Mauger, Sandeep Mariserla, Hongxu Ma, Megha Mohabey, Kyuyeun Kim, Alek Andreev, Scott Pollom, Juliette Love, Vihan Jain, Priyanka Agrawal, Yannick Schroecker, Alisa Fortin, Manfred Warmuth, Ji Liu, Andrew Leach, Irina Blok, Ganesh Poomal Girirajan, Roei Aharoni, Benigno Uria, Andrei Sozanschi, Dan Goldberg, Lucian Ionita, Marco Tulio Ribeiro, Martin Zlocha, Vighnesh Birodkar, Sami Lachgar, Liangzhe Yuan, Himadri Choudhury, Matt Ginsberg, Fei Zheng, Gregory Dibb, Emily Graves, Swachhand Lokhande, Gabriel Rasskin, George-Cristian Muraru, Corbin Quick, Sandeep Tata, Pierre Sermanet, Aditya Chawla, Itay Karo, Yan Wang, Susan Zhang, Orgad Keller, Anca Dragan, Guolong Su, Ian Chou, Xi Liu, Yiqing Tao, Shruthi Prabhakara, Marc Wilson, Ruibo Liu, Shibo Wang, Georgie Evans, David Du, Alfonso Castaño, Gautam Prasad, Mona El Mahdy, Sebastian Gerlach, Machel Reid, Jarrod Kahn, Amir Zait, Thanumalayan Sankaranarayanan Pillai, Thatcher Ulrich, Guanyu Wang, Jan Wassenberg, Efrat Farkash, Kiran Yalasangi, Congchao Wang, Maria Bauza, Simon Bucher, Ting Liu, Jun Yan, Gary Leung, Vikas Sindhwani, Parker Barnes, Avi Singh, Ivan Jurin, Jichuan Chang, Niket Kumar Bhumihar, Sivan Eiger, Gui Citovsky, Ben Withbroe, Zhang Li, Siyang Xue, Niccolò Dal Santo, Georgi Stoyanov, Yves Raimond, Steven Zheng, Yilin Gao, Vít Listík, Sławek Kwasiborski, Rachel Saputro, Adnan Ozturel, Ganesh Mallya, Kushal Majmudar, Ross West, Paul Caron, Jinliang Wei, Lluís Castrejon, Sharad Vikram, Deepak Ramachandran, Nikhil Dhawan, Jiho Park, Sara Smoot, George van den Driessche, Yochai Blau, Chase Malik, Wei Liang, Roy Hirsch, Cicero Nogueira dos Santos, Eugene Weinstein, Aaron van den Oord, Sid Lall, Nicholas FitzGerald, Zixuan Jiang, Xuan Yang, Dale Webster, Ali Elqursh, Aedan Pope, Georges Rotival, David Raposo, Wanzheng Zhu, Jeff Dean, Sami Alabed, Dustin Tran, Arushi Gupta, Zach Gleicher, Jessica Austin, Edouard Rosseel, Megh Umekar, Dipanjan Das, Yinghao Sun, Kai Chen, Karolis Misiunas, Xiang Zhou, Yixian Di, Alyssa Loo, Josh Newlan, Bo Li, Vinay Ramasesh, Ying Xu, Alex Chen, Sudeep Gandhe, Radu Soricut, Nikita Gupta, Shuguang Hu, Seliem El-Sayed, Xavier Garcia, Idan Brusilovsky, Pu-Chin Chen, Andrew Bolt, Lu Huang, Alex Gurney, Zhiying Zhang, Alexander Pritzel, Jarek Wilkiewicz, Bryan Seybold, Bhargav Kanagal Shamanna, Felix Fischer, Josef Dean, Karan Gill, Ross McIlroy, Abhishek Bhowmick, Jeremy Selier, Antoine Yang, Derek Cheng, Vladimir Magay, Jie Tan, Dhriti Varma, Christian Walder, Tomas Kocisky, Ryo Nakashima, Paul Natsev, Mike Kwong, Ionel Gog, Chiyuan Zhang, Sander Dieleman, Thomas Jimma, Andrey Ryabtsev, Siddhartha Brahma, David Steiner, Dayou Du, Ante Žužul, Mislav Žanić, Mukund Raghavachari, Willi Gierke, Zeyu Zheng, Dessie Petrova, Yann Dauphin, Yuchuan Liu, Ido Kessler, Steven Hand, Chris Duvarney, Seokhwan Kim, Hyo Lee, Léonard Hussenot, Jeffrey Hui, Josh Smith, Deepali Jain, Jiawei Xia, Gaurav Singh Tomar, Keyvan Amiri, Du Phan, Fabian Fuchs, Tobias Weyand, Nenad Tomasev, Alexandra Cordell, Xin Liu, Jonathan Mallinson, Pankaj Joshi, Andy Crawford, Arun Suggala, Steve Chien, Nick Fernando, Mariella Sanchez-Vargas, Duncan Williams, Phil Crone, Xiyang Luo, Igor Karpov, Jyn Shan, Terry Thurk, Robin Strudel, Paul Voigtlaender, Piyush Patil, Tim Dozat, Ali Khodaei, Sahil Singla, Piotr Ambroszczyk, Qiying Wu, Yifan Chang, Brian Roark, Chaitra Hegde, Tianli Ding, Angelos Filos, Zhongru Wu, André Susano Pinto, Shuang Liu, Saarthak Khanna, Aditya Pandey, Siobhan Mcloughlin, Qiujia Li, Sam Haves, Allan Zhou, Elena Buchatskaya, Isabel Leal, Peter de Boursac, Nami Akazawa, Nina Anderson, Terry Chen, Krishna Somandepalli, Chen Liang, Sheela Goenka, Stephanie Winkler, Alexander Grushetsky, Yifan Ding, Jamie Smith, Fan Ye, Jordi Pont-Tuset, Eric Li, Ruichao Li, Tomer Golany, Dawid Wegner, Tao Jiang, Omer Barak, Yuan Shangguan, Eszter Vértés, Renee Wong, Jörg Bornschein, Alex Tudor, Michele Bevilacqua, Tom Schaul, Ankit Singh Rawat, Yang Zhao, Kyriakos Axiotis, Lei Meng, Cory McLean, Jonathan Lai, Jennifer Beattie, Nate Kushman, Yaxin Liu, Blair Kutzman, Fiona Lang, Jingchen Ye, Praneeth Netrapalli, Pushkar Mishra, Myriam Khan,

Megha Goel, Rob Willoughby, David Tian, Honglei Zhuang, JD Chen, Zak Tsai, Tasos Kementsietsidis, Arjun Khare, James Keeling, Keyang Xu, Nathan Waters, Florent Altché, Ashok Popat, Bhavishya Mittal, David Saxton, Dalia El Badawy, Michael Mathieu, Zheng Zheng, Hao Zhou, Nishant Ranka, Richard Shin, Qingnan Duan, Tim Salimans, Ioana Mihailescu, Uri Shaham, Ming-Wei Chang, Yannis Assael, Nishanth Dikkala, Martin Izard, Vincent Cohen-Addad, Cat Graves, Vlad Feinberg, Grace Chung, DJ Strouse, Danny Karmon, Sahand Sharifzadeh, Zoe Ashwood, Khiem Pham, Jon Blanton, Alex Vasiloff, Jarred Barber, Mark Geller, Aurick Zhou, Fedir Zubach, Tzu-Kuo Huang, Lei Zhang, Himanshu Gupta, Matt Young, Julia Proskurnia, Ronny Votel, Valentin Gabeur, Gabriel Barcik, Aditya Tripathi, Hongkun Yu, Geng Yan, Beer Changpinyo, Filip Pavetić, Amy Coyle, Yasuhisa Fujii, Jorge Gonzalez Mendez, Tianhao Zhou, Harish Rajamani, Blake Hechtman, Eddie Cao, Da-Cheng Juan, Yi-Xuan Tan, Valentin Dalibard, Yilun Du, Natalie Clay, Kaisheng Yao, Wenhao Jia, Dimple Vijaykumar, Yuxiang Zhou, Xinyi Bai, Wei-Chih Hung, Steven Pecht, Georgi Todorov, Nikhil Khadke, Pramod Gupta, Preethi Lahoti, Arnaud Autef, Karthik Duddu, James Lee-Thorp, Alexander Bykovsky, Tautvydas Misiunas, Sebastian Flennerhag, Santhosh Thangaraj, Jed McGiffin, Zack Nado, Markus Kunesch, Andreas Noever, Amir Hertz, Marco Liang, Victor Stone, Evan Palmer, Samira Daruki, Arijit Pramanik, Siim Pöder, Austin Kyker, Mina Khan, Evgeny Sluzhaev, Marvin Ritter, Avraham Ruderman, Wenlei Zhou, Chirag Nagpal, Kiran Vodrahalli, George Necula, Paul Barham, Ellie Pavlick, Jay Hartford, Izhak Shafran, Long Zhao, Maciej Mikula, Tom Eccles, Hidetoshi Shimokawa, Kanav Garg, Luke Vilnis, Hanwen Chen, Ilia Shumailov, Kuang-Huei Lee, Abdelrahman Abdelhamed, Meiyan Xie, Vered Cohen, Ester Hlavnova, Dan Malkin, Chawin Sitawarin, James Lottes, Pauline Coquinot, Tianli Yu, Sandeep Kumar, Jingwei Zhang, Aroma Mahendru, Zafarali Ahmed, James Martens, Tao Chen, Aviel Boag, Daiyi Peng, Coline Devin, Arseniy Klimovskiy, Mary Phuong, Danny Vainstein, Jin Xie, Bhuvana Ramabhadran, Nathan Howard, Xinxin Yu, Gitartha Goswami, Jingyu Cui, Sam Shleifer, Mario Pinto, Chih-Kuan Yeh, Ming-Hsuan Yang, Sara Javanmardi, Dan Ethier, Chace Lee, Jordi Orbay, Suyog Kotecha, Carla Bromberg, Pete Shaw, James Thornton, Adi Gerzi Rosenthal, Shane Gu, Matt Thomas, Ian Gemp, Aditya Ayyar, Asahi Ushio, Aarush Selvan, Joel Wee, Chenxi Liu, Maryam Majzoubi, Weiren Yu, Jake Abernethy, Tyler Liechty, Renke Pan, Hoang Nguyen, Qiong, Hu, Sarah Perrin, Abhinav Arora, Emily Pitler, Weiye Wang, Kaushik Shivakumar, Flavien Prost, Ben Limonchik, Jing Wang, Yi Gao, Timothee Cour, Shyamal Buch, Huan Gui, Maria Ivanova, Philipp Neubeck, Kelvin Chan, Lucy Kim, Huizhong Chen, Naman Goyal, Da-Woon Chung, Lu Liu, Yao Su, Anastasia Petrushkina, Jiajun Shen, Armand Joulin, Yuanzhong Xu, Stein Xudong Lin, Yana Kulizhskaya, Ciprian Chelba, Shobha Vasudevan, Eli Collins, Vasilisa Bashlovkina, Tony Lu, Doug Fritz, Jongbin Park, Yanqi Zhou, Chen Su, Richard Tanburn, Mikhail Sushkov, Michelle Rasquinha, Jinning Li, Jennifer Prendki, Yiming Li, Pallavi LV, Shriya Sharma, Hen Fitoussi, Hui Huang, Andrew Dai, Phuong Dao, Mike Burrows, Henry Prior, Danfeng Qin, Golan Pundak, Lars Lowe Sjoesund, Art Khurshudov, Zhenkai Zhu, Albert Webson, Elizabeth Kemp, Tat Tan, Saurabh Agrawal, Susie Sargsyan, Liqun Cheng, Jim Stephan, Tom Kwiatkowski, David Reid, Arunkumar Byravan, Assaf Hurwitz, Michael, Nicolas Heess, Luowei Zhou, Sonam Goenka, Viral Carpenter, Anselm Levskaya, Bo Wang, Reed Roberts, Rémi Leblond, Sharat Chikkerur, Stav Ginzburg, Max Chang, Robert Riachi, Chuqiao, Xu, Zalán Borsos, Michael Pliskin, Julia Pawar, Morgane Lustman, Hannah Kirkwood, Ankit Anand, Aditi Chaudhary, Norbert Kalb, Kieran Milan, Sean Augenstein, Anna Goldie, Laurel Prince, Karthik Raman, Yanhua Sun, Vivian Xia, Aaron Cohen, Zhouyuan Huo, Josh Camp, Seher Ellis, Lukas Zilka, David Vilar Torres, Lisa Patel, Sho Arora, Betty Chan, Jonas Adler, Kareem Ayoub, Jacky Liang, Fayaz Jamil, Jiepu Jiang, Simon Baumgartner, Haitian Sun, Yael Karov, Yaroslav Akulov, Hui Zheng, Irene Cai, Claudio Fantacci, James Rubin, Alex Rav Acha, Mengchao Wang, Nina D'Souza, Rohit Sathyanarayana, Shengyang Dai, Simon Rowe, Andrey Simanovsky, Omer Goldman, Yuheng Kuang, Xiaoyue Pan, Andrew Rosenberg, Tania Rojas-Esponda, Praneet Dutta, Amy Zeng, Irina Jurenka, Greg Farquhar, Yamini Bansal, Shariq Iqbal, Becca Roelofs, Ga-Young Joung, Parker Beak, Changwan Ryu, Ryan Poplin, Yan Wu, Jean-Baptiste Alayrac, Senaka Buthpitiya, Olaf Ronneberger, Caleb Habtegebriel, Wei Li, Paul Cavallaro, Aurora Wei, Guy Bensusky, Timo Denk, Harish Ganapathy, Jeff Stanway, Pratik Joshi, Francesco Bertolini, Jessica Lo, Olivia Ma, Zachary Charles, Geta Sampemane, Himanshu Sahni, Xu Chen, Harry Askham, David Gaddy, Peter Young, Jiewen Tan, Matan Eyal, Arthur

Bražinskas, Li Zhong, Zhichun Wu, Mark Epstein, Kai Bailey, Andrew Hard, Kamyu Lee, Sasha Goldshtein, Alex Ruiz, Mohammed Badawi, Matthias Lochbrunner, JK Kearns, Ashley Brown, Fabio Pardo, Theophane Weber, Haichuan Yang, Pan-Pan Jiang, Berkin Akin, Zhao Fu, Marcus Wainwright, Chi Zou, Meenu Gaba, Pierre-Antoine Manzagol, Wendy Kan, Yang Song, Karina Zainullina, Rui Lin, Jeongwoo Ko, Salil Deshmukh, Apoorv Jindal, James Svensson, Divya Tyam, Heri Zhao, Christine Kaeser-Chen, Scott Baird, Pooya Moradi, Jamie Hall, Qiuchen Guo, Vincent Tsang, Bowen Liang, Fernando Pereira, Suhas Ganesh, Ivan Korotkov, Jakub Adamek, Sridhar Thiagarajan, Vinh Tran, Charles Chen, Chris Tar, Sanil Jain, Ishita Dasgupta, Taylan Bilal, David Reitter, Kai Zhao, Giulia Vezzani, Yasmin Gehman, Pulkit Mehta, Lauren Beltrone, Xerxes Dotiwalla, Sergio Guadarrama, Zaheer Abbas, Stefani Karp, Petko Georgiev, Chun-Sung Ferng, Marc Brockschmidt, Liqian Peng, Christoph Hirsenschall, Vikas Verma, Yingying Bi, Ying Xiao, Avigail Dabush, Kelvin Xu, Phil Wallis, Randall Parker, Qifei Wang, Yang Xu, Ilkin Safarli, Dinesh Tewari, Yin Zhang, Seungyeon Kim, Andrea Gesmundo, Mackenzie Thomas, Sergey Levi, Ahmed Chowdhury, Kanishka Rao, Peter Garst, Sam Conway-Rahman, Helen Ran, Kay McKinney, Zhisheng Xiao, Wenhao Yu, Rohan Agrawal, Axel Stjerngren, Catalin Ionescu, Jingjing Chen, Vivek Sharma, Justin Chiu, Fei Liu, Ken Franko, Clayton Sanford, Xingyu Cai, Paul Michel, Sanjay Ganapathy, Jane Labanowski, Zachary Garrett, Ben Vargas, Sean Sun, Bryan Gale, Thomas Buschmann, Guillaume Desjardins, Nimesh Ghelani, Palak Jain, Mudit Verma, Chulayuth Asawaroengchai, Julian Eisenschlos, Jitendra Harlalka, Hideto Kazawa, Don Metzler, Joshua Howland, Ying Jian, Jake Ades, Viral Shah, Tynan Gangwani, Seungji Lee, Roman Ring, Steven M. Hernandez, Dean Reich, Amer Sinha, Ashutosh Sathe, Joe Kovac, Ashleah Gill, Ajay Kannan, Andrea D'olimpio, Martin Sevenich, Jay Whang, Been Kim, Khe Chai Sim, Jilin Chen, Jiageng Zhang, Shuba Lall, Yossi Matias, Bill Jia, Abe Friesen, Sara Nasso, Ashish Thapliyal, Bryan Perozzi, Ting Yu, Anna Shekhawat, Safeen Huda, Peter Grabowski, Eric Wang, Ashwin Sreevatsa, Hilal Dib, Mehadi Hassen, Parker Schuh, Vedrana Milutinovic, Chris Welty, Michael Quinn, Ali Shah, Bangju Wang, Gabe Barth-Maron, Justin Frye, Natalie Axelsson, Tao Zhu, Yukun Ma, Irene Giannoumis, Hanie Sedghi, Chang Ye, Yi Luan, Kevin Aydin, Bilva Chandra, Vivek Sampathkumar, Ronny Huang, Victor Lavrenko, Ahmed Eleryan, Zhi Hong, Steven Hansen, Sara Mc Carthy, Bidisha Samanta, Domagoj Čevd, Xin Wang, Fangtao Li, Michael Voznesensky, Matt Hoffman, Andreas Terzis, Vikash Sehwal, Gil Fidel, Luheng He, Mu Cai, Yanzhang He, Alex Feng, Martin Nikoltchev, Samrat Phatale, Jason Chase, Rory Lawton, Ming Zhang, Tom Ouyang, Manuel Tragut, Mehdi Hafezi Manshadi, Arjun Narayanan, Jiaming Shen, Xu Gao, Tolga Bolukbasi, Nick Roy, Xin Li, Daniel Golovin, Liviu Panait, Zhen Qin, Guangxing Han, Thomas Anthony, Sneha Kudugunta, Viorica Patraucean, Aniket Ray, Xinyun Chen, Xiaochen Yang, Tanuj Bhatia, Pranav Talluri, Alex Morris, Andrija Ražnatović, Bethanie Brownfield, James An, Sheng Peng, Patrick Kane, Ce Zheng, Nico Duduta, Joshua Kessinger, James Noraky, Siqi Liu, Keran Rong, Petar Veličković, Keith Rush, Alex Goldin, Fanny Wei, Shiva Mohan Reddy Garlapati, Caroline Pantofaru, Okwan Kwon, Jianmo Ni, Eric Noland, Julia Di Trapani, Françoise Beaufays, Abhijit Guha Roy, Yinlam Chow, Aybuke Turker, Geoffrey Cideron, Lantao Mei, Jon Clark, Qingyun Dou, Matko Bošnjak, Ralph Leith, Yuqing Du, Amir Yazdanbakhsh, Milad Nasr, Chester Kwak, Suraj Satishkumar Sheth, Alex Kaskasoli, Ankesh Anand, Balaji Lakshminarayanan, Sammy Jerome, David Bieber, Chun-Te Chu, Alexandre Senges, Tianxiao Shen, Mukund Sridhar, Ndaba Ndebele, Benjamin Beyret, Shakir Mohamed, Mia Chen, Markus Freitag, Jiaxian Guo, Luyang Liu, Paul Roit, Heng Chen, Shen Yan, Tom Stone, JD Co-Reyes, Jeremy Cole, Salvatore Scellato, Shekoofeh Azizi, Hadi Hashemi, Alicia Jin, Anand Iyer, Marcella Valentine, András György, Arun Ahuja, Daniel Hernandez Diaz, Chen-Yu Lee, Nathan Clement, Weize Kong, Drew Garmon, Ishaan Watts, Kush Bhatia, Khyatti Gupta, Matt Miecznikowski, Hugo Vallet, Ankur Taly, Edward Loper, Saket Joshi, James Atwood, Jo Chick, Mark Collier, Fotis Iliopoulos, Ryan Trostle, Beliz Gunel, Ramiro Leal-Cavazos, Arnar Mar Hrafnkelsson, Michael Guzman, Xiaoen Ju, Andy Forbes, Jesse Emond, Kushal Chauhan, Ben Caine, Li Xiao, Wenjun Zeng, Alexandre Moufarek, Daniel Murphy, Maya Meng, Nitish Gupta, Felix Riedel, Anil Das, Elijah Lawal, Shashi Narayan, Tiberiu Sosea, James Swirhun, Linda Friso, Behnam Neyshabur, Jing Lu, Sertan Girgin, Michael Wunder, Edouard Yvinec, Aroonlalok Pyne, Victor Carbune, Shruti Rijhwani, Yang Guo, Tulsee Doshi, Anton Briukhov, Max Bain, Ayal Hitron, Xuanhui Wang, Ashish Gupta, Ke Chen, Cosmo Du, Weiyang Zhang, Dhruv

Shah, Arjun Akula, Max Dylla, Ashyana Kachra, Weicheng Kuo, Tingting Zou, Lily Wang, Luyao Xu, Jifan Zhu, Justin Snyder, Sachit Menon, Orhan Firat, Igor Mordatch, Yuan Yuan, Natalia Ponomareva, Rory Blevins, Lawrence Moore, Weijun Wang, Phil Chen, Martin Scholz, Artur Dwornik, Jason Lin, Sicheng Li, Diego Antognini, Te I, Xiaodan Song, Matt Miller, Uday Kalra, Adam Raveret, Oscar Akerlund, Felix Wu, Andrew Nystrom, Namrata Godbole, Tianqi Liu, Hannah DeBalsi, Jewel Zhao, Buhuang Liu, Avi Caciularu, Lauren Lax, Urvashi Khandelwal, Victoria Langston, Eric Bailey, Silvio Lattanzi, Yufei Wang, Neel Kovelamudi, Sneha Mondal, Guru Guruganesh, Nan Hua, Ofir Roval, Paweł Wesolowski, Rishikesh Ingale, Jonathan Halcrow, Tim Sohn, Christof Angermueller, Bahram Raad, Eli Stickgold, Eva Lu, Alec Kosik, Jing Xie, Timothy Lillicrap, Austin Huang, Lydia Lihui Zhang, Dominik Paulus, Clement Farabet, Alex Wertheim, Bing Wang, Rishabh Joshi, Chu ling Ko, Yonghui Wu, Shubham Agrawal, Lily Lin, XiangHai Sheng, Peter Sung, Tyler Breland-King, Christina Butterfield, Swapnil Gawde, Sumeet Singh, Qiao Zhang, Raj Apte, Shilpa Shetty, Adrian Hutter, Tao Li, Elizabeth Salesky, Federico Lebron, Jonni Kanerva, Michela Paganini, Arthur Nguyen, Rohith Vallu, Jan-Thorsten Peter, Sarmishta Velury, David Kao, Jay Hoover, Anna Bortsova, Colton Bishop, Shoshana Jakobovits, Alessandro Agostini, Alekh Agarwal, Chang Liu, Charles Kwong, Sasan Tavakkol, Ioana Bica, Alex Greve, Anirudh GP, Jake Marcus, Le Hou, Tom Duerig, Rivka Moroshko, Dave Lacey, Andy Davis, Julien Amelot, Guohui Wang, Frank Kim, Theofilos Strinopoulos, Hui Wan, Charline Le Lan, Shankar Krishnan, Haotian Tang, Peter Humphreys, Junwen Bai, Idan Heimlich Shtacher, Diego Machado, Chenxi Pang, Ken Burke, Dangyi Liu, Renga Aravamudhan, Yue Song, Ed Hirst, Abhimanyu Singh, Brendan Jou, Liang Bai, Francesco Piccinno, Chuyuan Kelly Fu, Robin Alazard, Barak Meiri, Daniel Winter, Charlie Chen, Mingda Zhang, Jens Heitkaemper, John Lambert, Jinhyuk Lee, Alexander Frömmgen, Sergey Rogulenko, Pranav Nair, Paul Niemczyk, Anton Bulyenov, Bibo Xu, Hadar Shemtov, Morteza Zadimoghaddam, Serge Toropov, Mateo Wirth, Hanjun Dai, Sreenivas Gollapudi, Daniel Zheng, Alex Kurakin, Chansoo Lee, Kalesha Bullard, Nicolas Serrano, Ivana Balazevic, Yang Li, Johan Schalkwyk, Mark Murphy, Mingyang Zhang, Kevin Sequeira, Romina Datta, Nishant Agrawal, Charles Sutton, Nithya Attaluri, Mencher Chiang, Wael Farhan, Gregory Thornton, Kate Lin, Travis Choma, Hung Nguyen, Kingshuk Dasgupta, Dirk Robinson, Iulia Comşa, Michael Riley, Arjun Pillai, Basil Mustafa, Ben Golan, Amir Zandieh, Jean-Baptiste Lespiau, Billy Porter, David Ross, Sujeewan Rajayogam, Mohit Agarwal, Subhashini Venugopalan, Bobak Shahriari, Qiqi Yan, Hao Xu, Taylor Tobin, Pavel Dubov, Hongzhi Shi, Adrià Recasens, Anton Kovsharov, Sebastian Borgeaud, Lucio Dery, Shanthal Vasanth, Elena Gribovskaya, Linhai Qiu, Mahdis Mahdiah, Wojtek Skut, Elizabeth Nielsen, CJ Zheng, Adams Yu, Carrie Grimes Bostock, Shaleen Gupta, Aaron Archer, Chris Rawles, Elinor Davies, Alexey Svyatkovskiy, Tomy Tsai, Yoni Halpern, Christian Reisswig, Bartek Wydrowski, Bo Chang, Joan Puigcerver, Mor Hazan Taege, Jian Li, Eva Schnider, Xinjian Li, Dragos Dena, Yunhan Xu, Umesh Telang, Tianze Shi, Heiga Zen, Kyle Kastner, Yeongil Ko, Neesha Subramaniam, Aviral Kumar, Pete Blois, Zhuyun Dai, John Wieting, Yifeng Lu, Yoel Zeldes, Tian Xie, Anja Hauth, Alexandru Țifrea, Yuqi Li, Sam El-Husseini, Dan Abolafia, Howard Zhou, Wen Ding, Sahra Ghalebikesabi, Carlos Guía, Andrii Maksai, Ágoston Weisz, Sercan Arik, Nick Sukhanov, Aga Świetlik, Xuhui Jia, Luo Yu, Weiyue Wang, Mark Brand, Dawn Bloxwich, Sean Kirmani, Zhe Chen, Alec Go, Pablo Sprechmann, Nithish Kannan, Alen Carin, Paramjit Sandhu, Isabel Edkins, Leslie Nooteboom, Jai Gupta, Loren Maggiore, Javad Azizi, Yael Pritch, Pengcheng Yin, Mansi Gupta, Danny Tarlow, Duncan Smith, Desi Ivanov, Mohammad Babaeizadeh, Ankita Goel, Satish Kambala, Grace Chu, Matej Kastelic, Michelle Liu, Hagen Soltau, Austin Stone, Shivani Agrawal, Min Kim, Kedar Soparkar, Srinivas Tadepalli, Oskar Bunyan, Rachel Soh, Arvind Kannan, DY Kim, Blake JianHang Chen, Afief Halumi, Sudeshna Roy, Yulong Wang, Olcan Sercinoglu, Gena Gibson, Sijal Bhatnagar, Motoki Sano, Daniel von Dincklage, Qingchun Ren, Blagoj Mitrevski, Mirek Olšák, Jennifer She, Carl Doersch, Jilei, Wang, Bingyuan Liu, Qijun Tan, Tamar Yakar, Tris Warkentin, Alex Ramirez, Carl Lebsack, Josh Dillon, Rajiv Mathews, Tom Cobley, Zelin Wu, Zhuoyuan Chen, Jon Simon, Swaroop Nath, Tara Sainath, Alexei Bendebury, Ryan Julian, Bharath Mankalale, Daria Ćurko, Paulo Zacchello, Adam R. Brown, Kiranbir Sodhia, Heidi Howard, Sergi Caelles, Abhinav Gupta, Gareth Evans, Anna Bulanova, Lesley Katzen, Roman Goldenberg, Anton Tsitsulin, Joe Stanton, Benoit Schillings, Vitaly Kovalev, Corey Fry, Rushin

Shah, Kuo Lin, Shyam Upadhyay, Cheng Li, Soroush Radpour, Marcello Maggioni, Jing Xiong, Lukas Haas, Jenny Brennan, Aishwarya Kamath, Nikolay Savinov, Arsha Nagrani, Trevor Yacovone, Ryan Kappedal, Kostas Andriopoulos, Li Lao, YaGuang Li, Grigory Rozhdestvenskiy, Kazuma Hashimoto, Andrew Audibert, Sophia Austin, Daniel Rodriguez, Anian Ruoss, Garrett Honke, Deep Karkhanis, Xi Xiong, Qing Wei, James Huang, Zhaoqi Leng, Vittal Premachandran, Stan Bileschi, Georgios Evangelopoulos, Thomas Mensink, Jay Pavagadhi, Denis Teplyashin, Paul Chang, Linting Xue, Garrett Tanzer, Sally Goldman, Kaushal Patel, Shixin Li, Jeremy Wiesner, Ivy Zheng, Ian Stewart-Binks, Jie Han, Zhi Li, Liangchen Luo, Karel Lenc, Mario Lučić, Fuzhao Xue, Ryan Mullins, Alexey Guseynov, Chung-Ching Chang, Isaac Galatzer-Levy, Adam Zhang, Garrett Bingham, Grace Hu, Ale Hartman, Yue Ma, Jordan Griffith, Alex Irpan, Carey Radebaugh, Summer Yue, Lijie Fan, Victor Ungureanu, Christina Sorokin, Hannah Teufel, Peiran Li, Rohan Anil, Dimitris Paparas, Todd Wang, Chu-Cheng Lin, Hui Peng, Megan Shum, Goran Petrovic, Demetra Brady, Richard Nguyen, Klaus Macherey, Zhihao Li, Harman Singh, Madhavi Yenugula, Mariko Iinuma, Xinyi Chen, Kavya Kopparapu, Alexey Stern, Shachi Dave, Chandu Thekkath, Florence Perot, Anurag Kumar, Fangda Li, Yang Xiao, Matthew Bilotti, Mohammad Hossein Bateni, Isaac Noble, Lisa Lee, Amelio Vázquez-Reina, Julian Salazar, Xiaomeng Yang, Boyu Wang, Ela Gruzewska, Anand Rao, Sindhu Raghuram, Zheng Xu, Eyal Ben-David, Jieru Mei, Sid Dalmia, Zhaoyi Zhang, Yuchen Liu, Gagan Bansal, Helena Pankov, Steven Schwarcz, Andrea Burns, Christine Chan, Sumit Sanghai, Ricky Liang, Ethan Liang, Antoine He, Amy Stuart, Arun Narayanan, Yukun Zhu, Christian Frank, Bahar Fatemi, Amit Sabne, Oran Lang, Indro Bhattacharya, Shane Settle, Maria Wang, Brendan McMahan, Andrea Tacchetti, Livio Baldini Soares, Majid Hadian, Serkan Cabi, Timothy Chung, Nikita Putikhin, Gang Li, Jeremy Chen, Austin Tarango, Henryk Michalewski, Mehran Kazemi, Hussain Masoom, Hila Sheftel, Rakesh Shivanna, Archita Vadali, Ramona Comanescu, Doug Reid, Joss Moore, Arvind Neelakantan, Michaël Sander, Jonathan Herzig, Aviv Rosenberg, Mostafa Dehghani, JD Choi, Michael Fink, Reid Hayes, Eric Ge, Shitao Weng, Chia-Hua Ho, John Karro, Kalpesh Krishna, Lam Nguyen Thiet, Amy Skerry-Ryan, Daniel Eppens, Marco Andreetto, Navin Sarma, Silvano Bonacina, Barcu Karagol Ayan, Megha Nawhal, Zhihao Shan, Mike Dusenberry, Shantanu Thakoor, Sagar Gubbi, Duc Dung Nguyen, Reut Tsarfaty, Samuel Albanie, Jovana Mitrović, Meet Gandhi, Bo-Juen Chen, Alessandro Epasto, Georgi Stephanov, Ye Jin, Samuel Gehman, Aida Amini, Jack Weber, Feryal Behbahani, Shawn Xu, Miltos Allamanis, Xi Chen, Myle Ott, Claire Sha, Michal Jastrzebski, Hang Qi, David Greene, Xinyi Wu, Abodunrinwa Toki, Daniel Vlasic, Jane Shapiro, Ragha Kotikalapudi, Zhe Shen, Takaaki Saeki, Sirui Xie, Albin Cassirer, Shikhar Bharadwaj, Tatsuya Kiyono, Srinadh Bhojanapalli, Elan Rosenfeld, Sam Ritter, Jieming Mao, João Gabriel Oliveira, Zoltan Egyed, Bernd Bandemer, Emilio Parisotto, Keisuke Kinoshita, Juliette Pluto, Petros Maniatis, Steve Li, Yaohui Guo, Golnaz Ghiasi, Jean Tarbouriech, Srimon Chatterjee, Julie Jin, Katrina Xu, Jennimaria Palomaki, Séb Arnold, Madhavi Sewak, Federico Piccinini, Mohit Sharma, Ben Albrecht, Sean Purser-haskell, Ashwin Vaswani, Chongyan Chen, Matheus Wisniewski, Qin Cao, John Aslanides, Nguyet Minh Phu, Maximilian Sieb, Lauren Agubuzu, Anne Zheng, Daniel Sohn, Marco Selvi, Anders Andreassen, Krishan Subudhi, Prem Erubetne, Oliver Woodman, Tomas Mery, Sebastian Krause, Xiaoqi Ren, Xiao Ma, Jincheng Luo, Dawn Chen, Wei Fan, Henry Griffiths, Christian Schuler, Alice Li, Shujian Zhang, Jean-Michel Sarr, Shixin Luo, Riccardo Patana, Matthew Watson, Dani Naboulsi, Michael Collins, Sailesh Sidhwani, Emiel Hoogeboom, Sharon Silver, Emily Caveness, Xiaokai Zhao, Mikel Rodriguez, Maxine Deines, Libin Bai, Patrick Griffin, Marco Tagliasacchi, Emily Xue, Spandana Raj Babbula, Bo Pang, Nan Ding, Gloria Shen, Elijah Peake, Remi Crocker, Shubha Srinivas Raghvendra, Danny Swisher, Woohyun Han, Richa Singh, Ling Wu, Vladimir Pchelin, Tsendsuren Munkhdalai, Dana Alon, Geoff Bacon, Efren Robles, Jannis Bulian, Melvin Johnson, George Powell, Felipe Tiengo Ferreira, Yaoyiran Li, Frederik Benzing, Mihajlo Velimirović, Hubert Soyer, William Kong, Tony, Nguyễn, Zhen Yang, Jeremiah Liu, Joost van Amersfoort, Daniel Gillick, Baochen Sun, Nathalie Rauschmayr, Katie Zhang, Serena Zhan, Tao Zhou, Alexey Frolov, Chengrun Yang, Denis Vnukov, Louis Rouillard, Hongji Li, Amol Mandhane, Nova Fallen, Rajesh Venkataraman, Clara Huiyi Hu, Jennifer Brennan, Jenny Lee, Jerry Chang, Martin Sundermeyer, Zhufeng Pan, Rosemary Ke, Simon Tong, Alex Fabrikant, William Bono, Jindong Gu, Ryan Foley, Yiran Mao, Manolis Delakis, Dhruva Bhaswar, Roy Frostig, Nick Li, Avital Zipori, Cath Hope,

Olga Kozlova, Swaroop Mishra, Josip Djolonga, Craig Schiff, Majd Al Merey, Eleftheria Briakou, Peter Morgan, Andy Wan, Avinatan Hassidim, RJ Skerry-Ryan, Kuntal Sengupta, Mary Jasarevic, Praveen Kallakuri, Paige Kunkle, Hannah Brennan, Tom Lieber, Hassan Mansoor, Julian Walker, Bing Zhang, Annie Xie, Goran Žužić, Adaeze Chukwuka, Alex Druinsky, Donghyun Cho, Rui Yao, Ferjad Naeem, Shiraz Butt, Eunyoung Kim, Zhipeng Jia, Mandy Jordan, Adam Lelkes, Mark Kurzeja, Sophie Wang, James Zhao, Andrew Over, Abhishek Chakladar, Marcel Prasetya, Neha Jha, Sriram Ganapathy, Yale Cong, Prakash Shroff, Carl Saroufim, Sobhan Miryoosefi, Mohamed Hammad, Tajwar Nasir, Weijuan Xi, Yang Gao, Young Maeng, Ben Hora, Chin-Yi Cheng, Parisa Haghighi, Yoad Lewenberg, Caden Lu, Martin Matysiak, Naina Raisinghani, Huiyu Wang, Lexi Baugher, Rahul Sukthankar, Minh Giang, John Schultz, Noah Fiedel, Minmin Chen, Cheng-Chun Lee, Tapomay Dey, Hao Zheng, Shachi Paul, Celine Smith, Andy Ly, Yicheng Wang, Rishabh Bansal, Bartek Perz, Susanna Ricco, Stasha Blank, Vaishakh Keshava, Deepak Sharma, Marvin Chow, Kunal Lad, Komal Jalan, Simon Osindero, Craig Swanson, Jacob Scott, Anastasija Ilić, Xiaowei Li, Siddhartha Reddy Jonnalagadda, Afzal Shama Soudagar, Yan Xiong, Bat-Orgil Batsaikhan, Daniel Jarrett, Naveen Kumar, Maulik Shah, Matt Lawlor, Austin Waters, Mark Graham, Rhys May, Sabela Ramos, Sandra Lefdal, Zeynep Cankara, Nacho Cano, Brendan O'Donoghue, Jed Borovik, Frederick Liu, Jordan Grimstad, Mahmoud Alnahlawi, Katerina Tsihlias, Tom Hudson, Nikolai Grigorev, Yiling Jia, Terry Huang, Tobenna Peter Igwe, Sergei Lebedev, Xiaodan Tang, Igor Kriukon, Frankie Garcia, Melissa Tan, Eric Jia, Peter Stys, Shikhar Vashishth, Yu Liang, Balaji Venkatraman, Chenjie Gu, Anastasios Kementsietsidis, Chen Zhu, Junehyuk Jung, Yunfei Bai, Mohammad Javad Hosseini, Faruk Ahmed, Aditya Gupta, Xin Yuan, Shereen Ashraf, Shitij Nigam, Gautam Vasudevan, Pranjal Awasthi, Adi Mayrav Gilady, Zelda Mariet, Ramy Eskander, Haiguang Li, Hexiang Hu, Guillermo Garrido, Philippe Schlattner, George Zhang, Rohun Saxena, Petar Dević, Kritika Muralidharan, Ashwin Murthy, Yiqian Zhou, Min Choi, Arissa Wongpanich, Zhengdong Wang, Premal Shah, Yuntao Xu, Yiling Huang, Stephen Spencer, Alice Chen, James Cohan, Junjie Wang, Jonathan Tompson, Junru Wu, Ruba Haroun, Haiqiong Li, Blanca Huergo, Fan Yang, Tongxin Yin, James Wendt, Michael Bendersky, Rahma Chaabouni, Javier Snaider, Johan Ferret, Abhishek Jindal, Tara Thompson, Andrew Xue, Will Bishop, Shubham Milind Phal, Archit Sharma, Yunhsuan Sung, Prabakar Radhakrishnan, Mo Shomrat, Reeve Ingle, Roopal Vij, Justin Gilmer, Mihai Dorin Istina, Sam Sobell, Yang Lu, Emily Nottage, Dorsa Sadigh, Jeremiah Willcock, Tingnan Zhang, Steve Xu, Sasha Brown, Katherine Lee, Gary Wang, Yun Zhu, Yi Tay, Cheolmin Kim, Audrey Gutierrez, Abhanshu Sharma, Yongqin Xian, Sungyong Seo, Claire Cui, Elena Pochernina, Cip Baetu, Krzysztof Jastrzebski, Mimi Ly, Mohamed Elhawaty, Dan Suh, Eren Sezener, Pidong Wang, Nancy Yuen, George Tucker, Jiahao Cai, Zuguang Yang, Cindy Wang, Alex Muzio, Hai Qian, Jae Yoo, Derek Lockhart, Kevin R. McKee, Mandy Guo, Malika Mehrotra, Artur Mendonça, Sanket Vaibhav Mehta, Sherry Ben, Chetan Tekur, Jiaqi Mu, Muye Zhu, Victoria Krakovna, Hongrae Lee, AJ Maschinot, Sébastien Cevey, HyunJeong Choe, Aijun Bai, Hansa Srinivasan, Derek Gasaway, Nick Young, Patrick Siegler, Dan Holtmann-Rice, Vihari Piratla, Kate Baumli, Roey Yogeve, Alex Hofer, Hado van Hasselt, Svetlana Grant, Yuri Chervonyi, David Silver, Andrew Hogue, Ayushi Agarwal, Kathie Wang, Preeti Singh, Four Flynn, Josh Lipschultz, Robert David, Lizzetth Bellot, Yao-Yuan Yang, Long Le, Filippo Graziano, Kate Olszewska, Kevin Hui, Akanksha Maurya, Nikos Parotsidis, Weijie Chen, Tayo Oguntebi, Joe Kelley, Anirudh Baddepudi, Johannes Mauere, Gregory Shaw, Alex Siegmán, Lin Yang, Shravya Shetty, Subhrajit Roy, Yunting Song, Wojciech Stokowiec, Ryan Burnell, Omkar Savant, Robert Busa-Fekete, Jin Miao, Samrat Ghosh, Liam MacDermed, Phillip Lippe, Mikhail Dektiarev, Zach Behrman, Fabian Mentzer, Kelvin Nguyen, Meng Wei, Siddharth Verma, Chris Knutsen, Sudeep Dasari, Zhipeng Yan, Petr Mitrichev, Xingyu Wang, Virat Shejwalkar, Jacob Austin, Srinivas Sunkara, Navneet Potti, Yan Virin, Christian Wright, Gaël Liu, Oriana Riva, Etienne Pot, Greg Kochanski, Quoc Le, Gargi Balasubramaniam, Arka Dhar, Yuguo Liao, Adam Bloniarz, Divyansh Shukla, Elizabeth Cole, Jong Lee, Sheng Zhang, Sushant Kafle, Siddharth Vashishtha, Parsa Mahmoudieh, Grace Chen, Raphael Hoffmann, Pranesh Srinivasan, Agustin Dal Lago, Yoav Ben Shalom, Zi Wang, Michael Elabd, Anuj Sharma, Junhyuk Oh, Suraj Kothawade, Maigo Le, Marianne Monteiro, Shentao Yang, Kaiz Alarakya, Robert Geirhos, Diana Mincu, Håvard Garnes, Hayato Kobayashi, Soroosh Mariooryad, Kacper Krasowiak, Zhixin Lai, Shibl Mourad, Mingqiu Wang, Fan

Bu, Ophir Aharoni, Guanjie Chen, Abhimanyu Goyal, Vadim Zubov, Ankur Bapna, Elahe Dabir, Nisarg Kothari, Kay Lamerigts, Nicola De Cao, Jeremy Shar, Christopher Yew, Nitish Kulkarni, Dre Mahaarachchi, Mandar Joshi, Zhenhai Zhu, Jared Lichtarge, Yichao Zhou, Hannah Muckenhirn, Vittorio Selo, Oriol Vinyals, Peter Chen, Anthony Brohan, Vaibhav Mehta, Sarah Cogan, Ruth Wang, Ty Geri, Wei-Jen Ko, Wei Chen, Fabio Viola, Keshav Shivam, Lisa Wang, Madeleine Clare Elish, Raluca Ada Popa, Sébastien Pereira, Jianqiao Liu, Raphael Koster, Donnie Kim, Gufeng Zhang, Sayna Ebrahimi, Partha Talukdar, Yanyan Zheng, Petra Poklukar, Ales Mikhalap, Dale Johnson, Anitha Vijayakumar, Mark Omernick, Matt Dibb, Ayush Dubey, Qiong Hu, Apurv Suman, Vaibhav Aggarwal, Ilya Kornakov, Fei Xia, Wing Lowe, Alexey Kolganov, Ted Xiao, Vitaly Nikolaev, Steven Hemingray, Bonnie Li, Joana Iljazi, Mikołaj Rybiński, Ballie Sandhu, Peggy Lu, Thang Luong, Rodolphe Jenatton, Vineetha Govindaraj, Hui, Li, Gabriel Dulac-Arnold, Wonpyo Park, Henry Wang, Abhinit Modi, Jean Pouget-Abadie, Kristina Greller, Rahul Gupta, Robert Berry, Prajit Ramachandran, Jinyu Xie, Liam McCafferty, Jianling Wang, Kilol Gupta, Hyeontaek Lim, Blaž Bratanič, Andy Brock, Ilia Akolzin, Jim Sproch, Dan Karliner, Duhyeon Kim, Adrian Goedeckemeyer, Noam Shazeer, Cordelia Schmid, Daniele Calandriello, Parul Bhatia, Krzysztof Choromanski, Ceslee Montgomery, Dheeru Dua, Ana Ramalho, Helen King, Yue Gao, Lynn Nguyen, David Lindner, Divya Pitta, Oleaser Johnson, Khalid Salama, Diego Ardila, Michael Han, Erin Farnese, Seth Odoom, Ziyue Wang, Xiangzhuo Ding, Norman Rink, Ray Smith, Harshal Tushar Lehri, Eden Cohen, Neera Vats, Tong He, Parthasarathy Gopavarapu, Adam Paszke, Miteyan Patel, Wouter Van Gansbeke, Lucia Loher, Luis Castro, Maria Voitovich, Tamara von Glehn, Nelson George, Simon Niklaus, Zach Eaton-Rosen, Nemanja Rakićević, Erik Jue, Sagi Perel, Carrie Zhang, Yuval Bahat, Angéline Pouget, Zhi Xing, Fantine Huot, Ashish Shenoy, Taylor Bos, Vincent Coriou, Bryan Richter, Natasha Noy, Yaqing Wang, Santiago Ontanon, Siyang Qin, Gleb Makarchuk, Demis Hassabis, Zhuowen Li, Mandar Sharma, Kumaran Venkatesan, Iurii Kemaev, Roxanne Daniel, Shiyu Huang, Saloni Shah, Octavio Ponce, Warren, Chen, Manaal Faruqui, Jialin Wu, Slavica Andačić, Szabolcs Payrits, Daniel McDuff, Tom Hume, Yuan Cao, MH Tessler, Qingze Wang, Yinan Wang, Ivor Rendulic, Eirikur Agustsson, Matthew Johnson, Tanya Lando, Andrew Howard, Sri Gayatri Sundara Padmanabhan, Mayank Daswani, Andrea Banino, Michael Kilgore, Jonathan Heek, Ziwei Ji, Alvaro Caceres, Conglong Li, Nora Kassner, Alexey Vlasov, Zeyu Liu, Alex Grills, Yanhan Hou, Roykrong Sukkerd, Gowoon Cheon, Nishita Shetty, Larisa Markeeva, Piotr Stanczyk, Tejas Iyer, Yuan Gong, Shawn Gao, Keerthana Gopalakrishnan, Tim Blyth, Malcolm Reynolds, Avishkar Bhoopchand, Misha Bilenko, Dero Gharibian, Vicky Zayats, Aleksandra Faust, Abhinav Singh, Min Ma, Hongyang Jiao, Sudheendra Vijayanarasimhan, Lora Aroyo, Vikas Yadav, Sarah Chakera, Ashwin Kakarla, Vilobh Meshram, Karol Gregor, Gabriela Botea, Evan Senter, Dawei Jia, Geza Kovacs, Neha Sharma, Sebastien Baur, Kai Kang, Yifan He, Lin Zhuo, Marija Kostelac, Itay Laish, Songyou Peng, Louis O'Bryan, Daniel Kasenberg, Girish Ramchandra Rao, Edouard Leurent, Biao Zhang, Sage Stevens, Ana Salazar, Ye Zhang, Ivan Lobov, Jake Walker, Allen Porter, Morgan Redshaw, Han Ke, Abhishek Rao, Alex Lee, Hoi Lam, Michael Moffitt, Jaeyoun Kim, Siyuan Qiao, Terry Koo, Robert Dadashi, Xinying Song, Mukund Sundarajan, Peng Xu, Chizu Kawamoto, Yan Zhong, Clara Barbu, Apoorv Reddy, Mauro Verzetti, Leon Li, George Papamakarios, Hanna Klimczak-Plucińska, Mary Cassin, Koray Kavukcuoglu, Rigel Swavelly, Alain Vaucher, Jeffrey Zhao, Ross Hemsley, Michael Tschannen, Heming Ge, Gaurav Menghani, Yang Yu, Natalie Ha, Wei He, Xiao Wu, Maggie Song, Rachel Sterneck, Stefan Zinke, Dan A. Calian, Annie Marsden, Alejandro Cruzado Ruiz, Matteo Hessel, Almog Gueta, Benjamin Lee, Brian Farris, Manish Gupta, Yunjie Li, Mohammad Saleh, Vedant Misra, Kefan Xiao, Piermaria Mendolicchio, Gavin Buttmore, Varvara Krayvanova, Nigamaa Nayakanti, Matthew Wiethoff, Yash Pande, Azalia Mirhoseini, Ni Lao, Jasmine Liu, Yiqing Hua, Angie Chen, Yury Malkov, Dmitry Kalashnikov, Shubham Gupta, Kartik Audhkhasi, Yuexiang Zhai, Sudhindra Kopalle, Prateek Jain, Eran Ofek, Clemens Meyer, Khushen Baatarsukh, Hana Strejček, Jun Qian, James Freedman, Ricardo Figueira, Michal Sokolik, Olivier Bachem, Raymond Lin, Dia Kharrat, Chris Hidey, Pingmei Xu, Dennis Duan, Yin Li, Muge Ersoy, Richard Everett, Kevin Cen, Rebeca Santamaria-Fernandez, Amir Taubenfeld, Ian Mackinnon, Linda Deng, Polina Zablotskaia, Shashank Viswanadha, Shivanker Goel, Damion Yates, Yunxiao Deng, Peter Choy, Mingqing Chen, Abhishek Sinha, Alex Mossin, Yiming Wang, Arthur Szlam, Susan Hao, Paul Kishan Rubenstein, Metin

Toksoz-Exley, Miranda Aperghis, Yin Zhong, Junwhan Ahn, Michael Isard, Olivier Lacombe, Florian Luisier, Chrysovalantis Anastasiou, Yogesh Kalley, Utsav Prabhu, Emma Dunleavy, Shaan Bijwadia, Justin Mao-Jones, Kelly Chen, Rama Pasumarthi, Emily Wood, Adil Dostmohamed, Nate Hurley, Jiri Simsa, Alicia Parrish, Mantas Pajarskas, Matt Harvey, Ondrej Skopek, Yony Kochinski, Javier Rey, Verena Rieser, Denny Zhou, Sun Jae Lee, Trilok Acharya, Guowang Li, Joe Jiang, Xiaofan Zhang, Bryant Gipson, Ethan Mahintorabi, Marco Gelmi, Nima Khajehnouri, Angel Yeh, Kayi Lee, Loic Matthey, Leslie Baker, Trang Pham, Han Fu, Alex Pak, Prakhar Gupta, Cristina Vasconcelos, Adam Sadovsky, Brian Walker, Sissie Hsiao, Patrik Zochbauer, Andreea Marzoca, Noam Velan, Junhao Zeng, Gilles Baechler, Danny Driess, Divya Jain, Yanping Huang, Lizzie Tao, John Maggs, Nir Levine, Jon Schneider, Erika Gemzer, Samuel Petit, Shan Han, Zach Fisher, Dustin Zelle, Courtney Biles, Eugene Ie, Asya Fadeeva, Casper Liu, Juliana Vicente Franco, Adrian Collister, Hao Zhang, Renshen Wang, Ruizhe Zhao, Leandro Kieliger, Kurt Shuster, Rui Zhu, Boqing Gong, Lawrence Chan, Ruoxi Sun, Sujoy Basu, Roland Zimmermann, Jamie Hayes, Abhishek Bapna, Jasper Snoek, Weel Yang, Puranjay Datta, Jad Al Abdallah, Kevin Kilgour, Lu Li, SQ Mah, Yennie Jun, Morgane Rivière, Abhijit Karmarkar, Tammo Spalink, Tao Huang, Lucas Gonzalez, Duc-Hieu Tran, Averil Nowak, John Palowitch, Martin Chadwick, Ellie Talius, Harsh Mehta, Thibault Sellam, Philipp Fränken, Massimo Nicosia, Kyle He, Aditya Kini, David Amos, Sugato Basu, Harrison Jobe, Eleni Shaw, Qiantong Xu, Colin Evans, Daisuke Ikeda, Chaochao Yan, Larry Jin, Lun Wang, Sachin Yadav, Ilia Labzovsky, Ramesh Sampath, Ada Ma, Candice Schumann, Aditya Siddhant, Rohin Shah, John Youssef, Rishabh Agarwal, Natalie Dabney, Alessio Tonioni, Moran Ambar, Jing Li, Isabelle Guyon, Benny Li, David Soergel, Boya Fang, Georgi Karadzhov, Cristian Udrescu, Trieu Trinh, Vikas Raunak, Seb Noury, Dee Guo, Sonal Gupta, Mara Finkelstein, Denis Petek, Lihao Liang, Greg Billock, Pei Sun, David Wood, Yiwen Song, Xiaobin Yu, Tatiana Matejovicova, Regev Cohen, Kalyan Andra, David D'Ambrosio, Zhiwei Deng, Vincent Nallatamby, Ebrahim Songhori, Rumen Dangovski, Andrew Lampinen, Pankil Botadra, Adam Hillier, Jiawei Cao, Nagabhushan Baddi, Adhi Kuncoro, Toshihiro Yoshino, Ankit Bhagatwala, Marcáurelio Ranzato, Rylan Schaeffer, Tianlin Liu, Shuai Ye, Obaid Sarvana, John Nham, Chenkai Kuang, Isabel Gao, Jinoo Baek, Shubham Mittal, Ayzaan Wahid, Anita Gergely, Bin Ni, Josh Feldman, Carrie Muir, Pascal Lamblin, Wolfgang Macherey, Ethan Dyer, Logan Kilpatrick, Víctor Campos, Mukul Bhutani, Stanislav Fort, Yanif Ahmad, Aliaksei Severyn, Kleopatra Chatziprimou, Oleksandr Ferludin, Mason Dimarco, Aditya Kusupati, Joe Heyward, Dan Bahir, Kevin Vilella, Katie Millican, Dror Marcus, Sanaz Bahargam, Caglar Unlu, Nicholas Roth, Zichuan Wei, Siddharth Gopal, Deepanway Ghoshal, Edward Lee, Sharon Lin, Jennie Lees, Dayeong Lee, Anahita Hosseini, Connie Fan, Seth Neel, Marcus Wu, Yasemin Altun, Honglong Cai, Enrique Piqueras, Josh Woodward, Alessandro Bissacco, Salem Haykal, Mahyar Bordbar, Prasha Sundaram, Sarah Hodgkinson, Daniel Toyama, George Polovets, Austin Myers, Anu Sinha, Tomer Levinboim, Kashyap Krishnakumar, Rachita Chhaparia, Tatiana Sholokhova, Nitesh Bharadwaj Gundavarapu, Ganesh Jawahar, Haroon Qureshi, Jieru Hu, Nikola Momchev, Matthew Rahtz, Renjie Wu, Aishwarya P S, Kedar Dhamdhare, Meiqi Guo, Umang Gupta, Ali Eslami, Mariano Schain, Michiel Blokzijl, David Welling, Dave Orr, Levent Bolelli, Nicolas Perez-Nieves, Mikhail Sirotenko, Aman Prasad, Arjun Kar, Borja De Balle Pigem, Tayfun Terzi, Gellért Weisz, Dipankar Ghosh, Aditi Mavalankar, Dhruv Madeka, Kaspar Dagaard, Hartwig Adam, Viraj Shah, Dana Berman, Maggie Tran, Steven Baker, Ewa Andrejczuk, Grishma Chole, Ganna Raboshchuk, Mahdi Mirzazadeh, Thais Kagohara, Shimu Wu, Christian Schallhart, Bernett Orlando, Chen Wang, Alban Rrustemi, Hao Xiong, Hao Liu, Arpi Vezer, Nolan Ramsden, Shuo yiin Chang, Sidharth Mudgal, Yan Li, Nino Vieillard, Yedid Hoshen, Farooq Ahmad, Ambrose Slone, Amy Hua, Natan Potikha, Mirko Rossini, Jon Stritar, Sushant Prakash, Zifeng Wang, Xuanyi Dong, Alireza Nazari, Efrat Nehoran, Kaan Tekelioglu, Yinxiao Li, Kartikeya Badola, Tom Funkhouser, Yuanzhen Li, Varun Yerram, Ramya Ganeshan, Daniel Formoso, Karol Langner, Tian Shi, Huijian Li, Yumeya Yamamori, Amayika Panda, Alaa Saade, Angelo Scorza Scarpatti, Chris Breaux, CJ Carey, Zongwei Zhou, Cho-Jui Hsieh, Sophie Bridgers, Alena Butryna, Nishesh Gupta, Vaibhav Tulsyan, Sanghyun Woo, Evgenii Eltyshov, Will Grathwohl, Chanel Parks, Seth Benjamin, Rina Panigrahy, Shenil Dodhia, Daniel De Freitas, Chris Sauer, Will Song, Ferran Alet, Jackson Tolins, Cosmin Paduraru, Xingyi Zhou, Brian Albert, Zizhao Zhang, Lei Shu, Mudit Bansal, Sarah Nguyen, Amir Globerson, Owen Xiao, James

Manyika, Tom Hennigan, Rong Rong, Josip Matak, Anton Bakalov, Ankur Sharma, Danila Sinopalnikov, Andrew Pierson, Stephen Roller, Geoff Brown, Mingcen Gao, Toshiyuki Fukuzawa, Amin Ghafouri, Kenny Vassigh, Iain Barr, Zhicheng Wang, Anna Korsun, Rajesh Jayaram, Lijie Ren, Tim Zaman, Samira Khan, Yana Lunts, Dan Deutsch, Dave Uthus, Nitzan Katz, Masha Samsikova, Amr Khalifa, Nikhil Sethi, Jiao Sun, Luming Tang, Uri Alon, Xianghong Luo, Dian Yu, Abhishek Nayyar, Bryce Petrini, Will Truong, Vincent Hellendoorn, Nikolai Chinaev, Chris Alberti, Wei Wang, Jingcao Hu, Vahab Mirrokni, Ananth Balashankar, Avia Aharon, Aahil Mehta, Ahmet Iscen, Joseph Kready, Lucas Manning, Anhad Mohananey, Yuankai Chen, Anshuman Tripathi, Allen Wu, Igor Petrovski, Dawsen Hwang, Martin Baeuml, Shreyas Chandrakaladharan, Yuan Liu, Rey Coaguila, Maxwell Chen, Sally Ma, Pouya Tafti, Susheel Tatineni, Terry Spitz, Jiayu Ye, Paul Vicol, Mihaela Rosca, Adrià Puigdomènech, Zohar Yahav, Sanjay Ghemawat, Hanzhao Lin, Phoebe Kirk, Zaid Nabulsi, Sergey Brin, Bernd Bohnet, Ken Caluwaerts, Aditya Srikanth Veerubhotla, Dan Zheng, Zihang Dai, Petre Petrov, Yichong Xu, Ramin Mehran, Zhuo Xu, Luisa Zintgraf, Jiho Choi, Spurthi Amba Hombaiah, Romal Thoppilan, Sashank Reddi, Lukasz Lew, Li Li, Kellie Webster, KP Sawhney, Lampros Lamprou, Siamak Shakeri, Mayank Lunayach, Jianmin Chen, Sumit Bagri, Alex Salcianu, Ying Chen, Yani Donchev, Charlotte Magister, Signe Nørly, Vitor Rodrigues, Tomas Izo, Hila Noga, Joe Zou, Thomas Köppe, Wenxuan Zhou, Kenton Lee, Xiangzhu Long, Danielle Eisenbud, Anthony Chen, Connor Schenck, Chi Ming To, Peilin Zhong, Emanuel Taropa, Minh Truong, Omer Levy, Danilo Martins, Zhiyuan Zhang, Christopher Sementurs, Kelvin Zhang, Alex Yakubovich, Pol Moreno, Lara McConnaughey, Di Lu, Sam Redmond, Lotte Weerts, Yonatan Bitton, Tiziana Refice, Nicolas Lacasse, Arthur Conmy, Corentin Talleg, Julian Odell, Hannah Forbes-Pollard, Arkadiusz Socala, Jonathan Hoech, Pushmeet Kohli, Alanna Walton, Rui Wang, Mikita Sazanovich, Kexin Zhu, Andrei Kapishnikov, Rich Galt, Matthew Denton, Ben Murdoch, Caitlin Sikora, Kareem Mohamed, Wei Wei, Uri First, Tim McConnell, Luis C. Cobo, James Qin, Thi Avrahami, Daniel Balle, Yu Watanabe, Annie Louis, Adam Kraft, Setareh Ariaifar, Yiming Gu, Eugénie Rives, Charles Yoon, Andrei Rusu, James Cobon-Kerr, Chris Hahn, Jiaming Luo, Yuvein, Zhu, Niharika Ahuja, Rodrigo Benenson, Raphaël Lopez Kaufman, Honglin Yu, Lloyd Hightower, Junlin Zhang, Darren Ni, Lisa Anne Hendricks, Gabby Wang, Gal Yona, Lalit Jain, Pablo Barrio, Surya Bhupatiraju, Siva Velusamy, Allan Dafoe, Sebastian Riedel, Tara Thomas, Zhe Yuan, Mathias Bellaïche, Sheena Panthaplackel, Klemen Kloboves, Sarthak Jauhari, Canfer Akbulut, Todor Davchev, Evgeny Gladchenko, David Madras, Aleksandr Chuklin, Tyrone Hill, Quan Yuan, Mukundan Madhavan, Luke Leonhard, Dylan Scandinaro, Qihang Chen, Ning Niu, Arthur Douillard, Bogdan Damoc, Yasumasa Onoe, Fabian Pedregosa, Fred Bertsch, Chas Leichter, Joseph Pagadora, Jonathan Malmaud, Sameera Ponda, Andy Twigg, Oleksii Duzhyi, Jingwei Shen, Miaosen Wang, Roopal Garg, Jing Chen, Utku Evci, Jonathan Lee, Leon Liu, Koji Kojima, Masa Yamaguchi, Arunkumar Rajendran, AJ Piergiovanni, Vinodh Kumar Rajendran, Marco Fornoni, Gabriel Ibagon, Harry Ragan, Sadh MNM Khan, John Blitzer, Andrew Bunner, Guan Sun, Takahiro Kosakai, Scott Lundberg, Ndidi Elue, Kelvin Guu, SK Park, Jane Park, Arunachalam Narayanaswamy, Chengda Wu, Jayaram Mudigonda, Trevor Cohn, Hairong Mu, Ravi Kumar, Laura Graesser, Yichi Zhang, Richard Killam, Vincent Zhuang, Mai Giménez, Wael Al Jishi, Ruy Ley-Wild, Alex Zhai, Kazuki Osawa, Diego Cedillo, Jialu Liu, Mayank Upadhyay, Marcin Sieniek, Roshan Sharma, Tom Paine, Anelia Angelova, Sravanti Addepalli, Carolina Parada, Kingshuk Majumder, Avery Lamp, Sanjiv Kumar, Xiang Deng, Artiom Myaskovsky, Tea Sabolić, Jeffrey Dudek, Sarah York, Félix de Chaumont Quitry, Jiazhong Nie, Dee Cattle, Alok Gunjan, Bilal Piot, Waleed Khawaja, Seojin Bang, Simon Wang, Siavash Khodadadeh, Raghavender R, Praynaa Rawlani, Richard Powell, Kevin Lee, Johannes Griesser, GS Oh, Cesar Magalhaes, Yujia Li, Simon Tokumine, Hadas Natalie Vogel, Dennis Hsu, Arturo BC, Disha Jindal, Matan Cohen, Zi Yang, Junwei Yuan, Dario de Cesare, Tony Bruguier, Jun Xu, Monica Roy, Alon Jacovi, Dan Belov, Rahul Arya, Phoenix Meadowlark, Shlomi Cohen-Ganor, Wenting Ye, Patrick Morris-Suzuki, Praseem Banzal, Gan Song, Pranavaraj Ponnuramu, Fred Zhang, George Scrivener, Salah Zaiem, Alif Raditya Rochman, Kehang Han, Badih Ghazi, Kate Lee, Shahar Drath, Daniel Suo, Antonious Girgis, Pradeep Shenoy, Duy Nguyen, Douglas Eck, Somit Gupta, Le Yan, Joao Carreira, Anmol Gulati, Ruoxin Sang, Daniil Mirylenka, Emma Cooney, Edward Chou, Mingyang Ling, Cindy Fan, Ben Coleman, Guilherme Tubone, Ravin Kumar, Jason Baldridge, Felix Hernandez-Campos, Angeliki

- Lazaridou, James Besley, Itay Yona, Neslihan Bulut, Quentin Wellens, AJ Pierigiovanni, Jasmine George, Richard Green, Pu Han, Connie Tao, Geoff Clark, Chong You, Abbas Abdolmaleki, Justin Fu, Tongzhou Chen, Ashwin Chaugule, Angad Chandorkar, Altaf Rahman, Will Thompson, Penporn Koanantakool, Mike Bernico, Jie Ren, Andrey Vlasov, Sergei Vassilvitskii, Maciej Kula, Yizhong Liang, Dahun Kim, Yangsibo Huang, Chengxi Ye, Dmitry Lepikhin, and Wesley Helmholtz. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. URL <https://arxiv.org/abs/2507.06261>.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahé Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Jeff Wang, and N. L. L. B. Team. Scaling neural machine translation to 200 languages. *Nature*, 630(8018):841–846, Jun 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07335-x. URL <https://doi.org/10.1038/s41586-024-07335-x>.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojuan Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanbiao Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. DeepSeek-V3 technical report. *arXiv preprint arXiv:2412.19437*, 2025. URL <https://arxiv.org/abs/2412.19437>.
- A. Seza Doğruöz, Sunayana Sitaram, Barbara E. Bullock, and Almeida Jacqueline Toribio. A survey of code-switching: Linguistic and social perspectives for language technologies. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1654–1666, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.131. URL <https://aclanthology.org/2021.acl-long.131/>.

- Julen Etxaniz, Gorka Azkune, Aitor Soroa, Oier Lopez de Lacalle, and Mikel Artetxe. Do multilingual language models think better in English? In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pp. 550–564, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-short.46. URL <https://aclanthology.org/2024.naacl-short.46/>.
- Daniel L. Finer. Movement triggers and reflexivization in Korean-English codeswitching. In *Grammatical Theory and Bilingual Codeswitching*. The MIT Press, 12 2014. ISBN 9780262320351. doi: 10.7551/mitpress/8338.003.0004. URL <https://doi.org/10.7551/mitpress/8338.003.0004>.
- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Xin Zhao, Ting Song, Yan Xia, and Furu Wei. Not all languages are created equal in LLMs: Improving multilingual capability by cross-lingual-thought prompting. In Houada Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 12365–12394, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.826. URL <https://aclanthology.org/2023.findings-emnlp.826/>.
- Muhammad Huzaifah, Weihua Zheng, Nattapol Chanpaisit, and Kui Wu. Evaluating code-switching translation with large language models. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 6381–6394, Torino, Italia, May 2024. ELRA and ICCL. URL <https://aclanthology.org/2024.lrec-main.565/>.
- Yotam Intrator, Matan Halfon, Roman Goldenberg, Reut Tsarfaty, Matan Eyal, Ehud Rivlin, Yossi Matias, and Natalia Aizenberg. Breaking the language barrier: Can direct inference outperform pre-translation in multilingual LLM applications? In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pp. 829–844, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-short.75. URL <https://aclanthology.org/2024.naacl-short.75/>.
- Sai Muralidhar Jayanthi, Kavya Nerella, Khyathi Raghavi Chandu, and Alan W Black. CodemixedNLP: An extensible and open NLP toolkit for code-mixing. In Tamar Solorio, Shuguang Chen, Alan W. Black, Mona Diab, Sunayana Sitaram, Victor Soto, Emre Yilmaz, and Anirudh Srinivasan (eds.), *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pp. 113–118, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.calcs-1.14. URL <https://aclanthology.org/2021.calcs-1.14/>.
- Deokhyung Kang, Seonjeong Hwang, Daehui Kim, Hyounghun Kim, and Gary Lee. Why do multilingual reasoning gaps emerge in reasoning language models? In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 31684–31716, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.1586. URL <https://aclanthology.org/2026.findings-acl.1586/>.
- Seoyeon Kim, Huiseo Kim, Chanjun Park, Jinyoung Yeo, and Dongha Lee. Can code-switched texts activate a knowledge switch in LLMs? a case study on English-Korean code-switching. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 22303–22327, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.1215. URL <https://aclanthology.org/2025.findings-emnlp.1215/>.

- Sunkyoung Kim, Dayeon Ki, Yireun Kim, and Jinsik Lee. Cross-lingual QA: A key to unlocking in-context cross-lingual performance. In *ICML 2024 Workshop on In-Context Learning*, 2024. URL <https://openreview.net/forum?id=CzcCUzJQER>.
- Jaeseong Lee, YeonJoon Jung, and Seung-won Hwang. COMMIT: Code-mixing English-centric large language model for multilingual instruction tuning. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 3130–3137, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.198. URL <https://aclanthology.org/2024.findings-naacl.198/>.
- Chong Li, Shaonan Wang, Jiajun Zhang, and Chengqing Zong. Improving in-context learning of multilingual generative language models with cross-lingual alignment. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 8058–8076, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.445. URL <https://aclanthology.org/2024.naacl-long.445/>.
- Yihao Li, Jiayi Xin, Miranda Muqing Miao, Qi Long, and Lyle Ungar. The impact of language mixing on bilingual LLM reasoning. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 32519–32536, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.1654. URL <https://aclanthology.org/2025.emnlp-main.1654/>.
- Peiqin Lin, Andre Martins, and Hinrich Schuetze. XAMPLER: Learning to retrieve cross-lingual in-context examples. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 3968–3977, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl.221. URL <https://aclanthology.org/2025.findings-naacl.221/>.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. Few-shot learning with multilingual generative language models. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 9019–9052, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.616. URL <https://aclanthology.org/2022.emnlp-main.616/>.
- Amr Mohamed, Yang Zhang, Michalis Vazirgiannis, and Guokan Shang. Lost in the mix: Evaluating LLM understanding of code-switched text. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 44922–44938, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-390-6. doi: 10.18653/v1/2026.acl-long.2080. URL <https://aclanthology.org/2026.acl-long.2080/>.
- Carol Myers-Scotton. *Duelling languages: Grammatical structure in codeswitching*. Oxford University Press, 1997.
- Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-Basulto, Yazmín Ibáñez García, Hwaran Lee, Shamsuddeen Hassan Muhammad, Kiwoong Park, Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, Nedjma Ousidhoum, Jose Camacho-Collados, and Alice Oh. BLEnD: A benchmark for LLMs on everyday knowledge in diverse cultures and languages. In A. Globerson,

- L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 78104–78146. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/8eb88844dafefa92a26aaec9f3acad93-Paper-Datasets_and_Benchmarks_Track.pdf.
- Vera Neplenbroek, Arianna Bisazza, and Raquel Fernández. MBBQ: A dataset for cross-lingual comparison of stereotypes in generative LLMs. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=X9yV41FHt4>.
- Filippo Pallucchini, Lorenzo Malandri, Fabio Mercorio, and Mario Mezzanzanica. Lost in alignment: A survey on cross-lingual alignment methods for contextualized representation. *ACM Comput. Surv.*, August 2025. ISSN 0360-0300. doi: 10.1145/3764112. URL <https://doi.org/10.1145/3764112>.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. COMET-22: Unbabel-IST 2022 submission for the metrics shared task. In Philipp Koehn, Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Tom Kocmi, André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, Matteo Negri, Aurélie Névoul, Mariana Neves, Martin Popel, Marco Turchi, and Marcos Zampieri (eds.), *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pp. 578–585, Abu Dhabi, United Arab Emirates (Hybrid), December 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.wmt-1.52/>.
- Mohd Sanad Zaki Rizvi, Anirudh Srinivasan, Tanuja Ganu, Monojit Choudhury, and Sunayana Sitaram. GCM: A toolkit for generating synthetic code-mixed text. In Dimitra Gkatzia and Djamé Seddah (eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pp. 205–211, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-demos.24. URL <https://aclanthology.org/2021.eacl-demos.24/>.
- Lisa Schut, Yarin Gal, and Sebastian Farquhar. Do multilingual LLMs think in English? *arXiv preprint arXiv:2502.15603*, 2025. URL <https://arxiv.org/abs/2502.15603>.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. Language models are multilingual chain-of-thought reasoners. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=fR3wGCK-IXp>.
- Shivalika Singh, Angelika Romanou, Clémentine Fourier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, Enzo Ferrante, Marzieh Fadaee, Beyza Ermis, and Sara Hooker. Global MMLU: Understanding and addressing cultural and linguistic biases in multilingual evaluation. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 18761–18799, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.919. URL <https://aclanthology.org/2025.acl-long.919/>.
- Guijin Son, Donghun Yang, Hitesh Laxmichand Patel, Amit Agarwal, Hyunwoo Ko, Chanuk Lim, Srikant Panda, Minhyuk Kim, Nikunj drolia, Dasol Choi, Kyong-Ha Lee, and Youngjae Yu. Pushing on multilingual reasoning models with language-mixed chain-of-thought. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=ABc5y3741T>.
- Jiayang Song, Yuheng Huang, Zhehua Zhou, and Lei Ma. Multilingual blending: Large language model safety alignment evaluation with language mixture. In Luis Chiruzzo,

- Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 3433–3449, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl.191. URL <https://aclanthology.org/2025.findings-naacl.191/>.
- Sathya Krishnan Suresh, Tanmay Surana, Lim Zhi Hao, and Eng Siong Chng. CS-sum: A benchmark for code-switching dialogue summarization and the limits of large language models. In Yue Dong, Wen Xiao, Haopeng Zhang, Rui Zhang, Ori Ernst, Lu Wang, and Fei Liu (eds.), *Proceedings of The 5th New Frontiers in Summarization Workshop*, pp. 31–47, Hybrid, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-337-1. doi: 10.18653/v1/2025.newsum-main.3. URL <https://aclanthology.org/2025.newsum-main.3/>.
- Eshaan Tanwar, Subhabrata Dutta, Manish Borthakur, and Tanmoy Chakraborty. Multilingual LLMs are better cross-lingual in-context learners with alignment. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6292–6307, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.346. URL <https://aclanthology.org/2023.acl-long.346/>.
- Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. Do LLMs exhibit human-like response biases? a case study in survey design. *Transactions of the Association for Computational Linguistics*, 12:1011–1026, 09 2024. ISSN 2307-387X. doi: 10.1162/tacl_a_00685. URL https://doi.org/10.1162/tacl_a_00685.
- Mingyang Wang, Lukas Lange, Heike Adel, Yunpu Ma, Jannik Strötgen, and Hinrich Schuetze. Language mixing in reasoning language models: Patterns, impact, and internal causes. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2637–2665, Suzhou, China, November 2025a. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.132. URL <https://aclanthology.org/2025.emnlp-main.132/>.
- Yiming Wang, Pei Zhang, Jialong Tang, Haoran Wei, Baosong Yang, Rui Wang, Chenshu Sun, Feitong Sun, Jiran Zhang, Junxuan Wu, Qiqian Cang, Yichang Zhang, Fei Huang, Junyang Lin, Fei Huang, and Jingren Zhou. PolyMath: Evaluating mathematical reasoning in multilingual contexts. *arXiv preprint arXiv:2504.18428*, 2025b. URL <https://arxiv.org/abs/2504.18428>.
- Zhijun Wang, Jiahuan Li, Hao Zhou, Rongxiang Weng, Jingang Wang, Xin Huang, Xue Han, Junlan Feng, Chao Deng, and Shujian Huang. Investigating and scaling up code-switching for multilingual language model pre-training. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 11032–11046, Vienna, Austria, July 2025c. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.575. URL <https://aclanthology.org/2025.findings-acl.575/>.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. Do llamas work in English? on the latent language of multilingual transformers. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15366–15394, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.820. URL <https://aclanthology.org/2024.acl-long.820/>.
- Genta Winata, Shijie Wu, Mayank Kulkarni, Tamar Solorio, and Daniel Preotiuc-Pietro. Cross-lingual few-shot learning on unseen languages. In Yulan He, Heng Ji, Sujian Li, Yang Liu, and Chua-Hui Chang (eds.), *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 777–791, Online only, November 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.aacl-main.59. URL <https://aclanthology.org/2022.aacl-main.59/>.

- Genta Winata, Alham Fikri Aji, Zheng Xin Yong, and Tamar Solorio. The decades progress on code-switching research in NLP: A systematic survey on trends and challenges. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 2936–2978, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.185. URL <https://aclanthology.org/2023.findings-acl.185/>.
- xAI. Grok 4 fast model card, 2025. URL <https://data.x.ai/2025-09-19-grok-4-fast-model-card.pdf>.
- Peng Xie, Xingyuan Liu, Tsz Wai Chan, Yequan Bie, Yangqiu Song, Yang Wang, Hao Chen, and Kani Chen. SwitchLingua: The first large-scale multilingual and multi-ethnic code-switching dataset. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2025. URL <https://openreview.net/forum?id=N5BB70r30g>.
- Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. A paradigm shift in machine translation: Boosting translation performance of large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=farT6XXntP>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. URL <https://arxiv.org/abs/2505.09388>.
- Yilun Yang and Yekun Chai. CodeMixBench: Evaluating code-mixing capabilities of LLMs across 18 languages. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2139–2169, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.109. URL <https://aclanthology.org/2025.emnlp-main.109/>.
- Haneul Yoo, Cheonbok Park, Sangdoo Yun, Alice Oh, and Hwaran Lee. Code-switching curriculum learning for multilingual transfer in LLMs. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 7816–7836, Vienna, Austria, July 2025a. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.407. URL <https://aclanthology.org/2025.findings-acl.407/>.
- Haneul Yoo, Yongjin Yang, and Hwaran Lee. Code-switching red-teaming: LLM evaluation for safety and multilingual understanding. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13392–13413, Vienna, Austria, July 2025b. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.657. URL <https://aclanthology.org/2025.acl-long.657/>.
- Ruochen Zhang, Samuel Cahyawijaya, Jan Christian Blaise Cruz, Genta Winata, and Alham Fikri Aji. Multilingual large language models are not (yet) code-switchers. In Houada Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 12567–12582, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.774. URL <https://aclanthology.org/2023.emnlp-main.774/>.

- Wenbo Zhang, Aditya Majumdar, and Amulya Yadav. CHAI for LLMs: Improving code-mixed translation in large language models through reinforcement learning with AI feedback. *arXiv preprint arXiv:2411.09073*, 2025. URL <https://arxiv.org/abs/2411.09073>.
- Yiran Zhao, Wenxuan Zhang, Guizhen Chen, Kenji Kawaguchi, and Lidong Bing. How do large language models handle multilingualism? In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 15296–15319. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/1bd359b32ab8b2a6bbafa1ed2856cf40-Paper-Conference.pdf.
- Chengzhi Zhong, Qianying Liu, Fei Cheng, Junfeng Jiang, Zhen Wan, Chenhui Chu, Yugo Murawaki, and Sadao Kurohashi. What language do non-English-centric large language models think in? In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 26333–26346, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.1350. URL <https://aclanthology.org/2025.findings-acl.1350/>.

Appendix

A Limitations

Due to computational constraints, most experiments are reported with a single seed; the bootstrap intervals therefore capture uncertainty over evaluation instances rather than full run-to-run variability.

For evaluation, we largely employ automated metrics (*i.e.*, accuracy, exact match, and COMET), which could be considered a limitation. We adopt automated metrics for consistency and scalability across multilingual settings, which is central to our study. Further, a small-scale human evaluation in Appendix D.3 shows strong agreement with the COMET-based results, indicating that our main findings are robust.

We conduct experiments on 10 languages spanning different resource levels. While this set cannot cover all language families, it offers a broad and representative sample that supports the generality of CSICL.

In our task- and domain-specific evaluation (§4.2), we fix Spanish as the target language and vary the unseen language across datasets due to the limited availability of parallel data in specialized tasks and domains. While this design enables broad coverage of domains, part of the performance differences observed across tasks may stem from the choice of unseen languages rather than from the tasks themselves. It reflects the practical challenges of multilingual evaluation and still provides a diverse and informative basis for analysis.

B System Prompts

To ground the following system prompts, we carefully curate the real-world code-switching examples from [Finer \(2014\)](#).

B.1 Few-shot Demonstrations Generation

Prompt for generating English-Korean code-switching few-shot demonstrations

You are a bilingual rewriting assistant.

[TASK]

- Input : an English sentence (E) and its Korean translation (K)
- Output : the code-switching version of the parallel sentences following Matrix Language Frame (MLF) model
- Replace about 50% of words/phrases in E with their Korean equivalents taken from K
- Keep the original English word order and follow English syntax (S-V-O)
- DO NOT add explanations, examples, tags, prefix or extra sentences
- If there is no suitable Korean equivalent, keep the English word

[EXAMPLE]

<English> I ate dinner quickly.

<Korean> 나는 저녁을 빨리 먹었다.

<Code-Switching> I ate 저녁 빨리.

<English> Hana, put the toys in the basket quickly and go home.

<Korean> 하나야, 바구니에 장난감을 빨리 넣고 집에 가자.

<Code-Switching> Hana, put 장난감 in the basket quickly and 집에 가자.

<English> Dad was about to throw away my tooth.

<Korean> 아빠가 내 이빨을 빼려고 했어.

<Code-Switching> 아빠 was about to 빨래 my 이빨.

<English> I have to wash my hand.

<Korean> 나는 손을 씻어야 해.

<Code-Switching> I have to 닦아 my hand.

<English> Tom thinks Bill likes himself.
 <Korean> 톰은 빌이 자기 자신을 좋아한다고 생각한다.
 <Code-Switching> Tom thinks that Bill이 자기를 좋아한다.

<English> Tom thinks Bill likes himself.
 <Korean> 톰은 빌이 자기 자신을 좋아한다고 생각한다.
 <Code-Switching> Tom이 생각하는 Bill likes himself.

[BEGIN TASK]

Prompt for generating Korean-English code-switching few-shot demonstrations

You are a bilingual rewriting assistant.

[TASK]

- Input : an English sentence (E) and its Korean translation (K)
- Output : the code-switching version of the parallel sentences following Matrix Language Frame (MLF) model
- Replace about 50% of words/phrases in K with their English equivalents taken from E
- Keep the original Korean word order and follow Korean syntax (S-O-V)
- DO NOT add explanations, examples, tags, prefix or extra sentences
- If there is no suitable English equivalent, keep the Korean word

[EXAMPLE]

<English> Meena, put all the toys in the basket quickly, and go home.

<Korean> 미나야, 바구니에 장난감을 다 넣고 빨리 집에 가자.

<Code-Switching> Meena, basket 안에다 all the toys를 빨리 put하고 집에 가자.

<English> Last time, by mistake, they did not renew my driver's license at the DMV.

<Korean> 지난 번에 도로교통공단이 내 운전면허증을 실수로 갱신하지 않았어요.

<Code-Switching> 지난번에 motor vehicle department에서 내 driver's license를 mistake로 갱신하지 않았어요.

<English> They often do a thing like that.

<Korean> 그들이 일을 종종 그렇게 해요.

<Code-Switching> 그들이 일을 often 그렇게 해요.

<English> Tom thinks Bill likes himself.

<Korean> 톰은 빌이 자기 자신을 좋아한다고 생각한다.

<Code-Switching> Tom은 Bill이 himself를 좋아한다고 생각한다.

<English> John wonders what Mary bought yesterday.

<Korean> 존은 메리가 어제 무엇을 샀는지 궁금해한다.

<Code-Switching> John은 Mary가 yesterday 무엇을 샀는지 궁금해한다.

[BEGIN TASK]

Prompt for generating gradual code-switching (En→Ko) few-shot demonstrations

You are a bilingual rewriting assistant.

Your task is to generate five versions of a sentence that gradually transition from English to Korean.

[INPUT]

- One English sentence (E)
- Its Korean translation (K)
- A code-switching version of the sentence (C), where about 50% of English words are replaced by Korean equivalents

[OUTPUT]

Generate a sequence of five sentences showing a smooth progression from English to Korean:

1. English only (100% English, source syntax S-V-O)
2. 75% English + 25% Korean (matrix language: English, embedded language: Korean)
3. 50% English + 50% Korean (matrix language: English, embedded language: Korean)

4. 25% English + 75% Korean (matrix language: English, embedded language: Korean)
5. Korean only (100% Korean, target syntax S-O-V)

[RULES]

Following the Matrix Language Frame (MLF) model,

- Preserve English word order (S-V-O) and syntax until version 5 (full Korean).
- Use Korean equivalents from K when inserting Korean into English sentences.
- Keep the code-switching natural and consistent, not random.
- Do not add explanations, notes, or extra text — output only the five sentences in order.

[EXAMPLES]

EXAMPLE 1

Input:

<English> I ate dinner quickly.

<Korean> 나는 저녁을 빨리 먹었다.

<Code-Switching> I ate 저녁 빨리.

Output:

1. I ate dinner quickly.
2. I ate dinner 빨리.
3. I ate 저녁 빨리.
4. 나는 저녁 빨리 ate.
5. 나는 저녁을 빨리 먹었다.

EXAMPLE 2

Input:

<English> Dad was about to throw away my tooth.

<Korean> 아빠가 내 이빨을 빼려고 했어.

<Code-Switching> 아빠 was about to 빨래 my 이빨.

Output:

1. Dad was about to throw away my tooth.
2. Dad was about to throw away 내 이빨.
3. 아빠 was about to 빨래 my 이빨.
4. 아빠 was about to 내 이빨 빼려고 했어.
5. 아빠가 내 이빨을 빼려고 했어.

Example 3

Input:

<English> Tom thinks Bill likes himself.

<Korean> 톰은 빌이 자기 자신을 좋아한다고 생각한다.

<Code-Switching> Tom thinks that Bill이 자기를 좋아한다.

Output:

1. Tom thinks Bill likes himself.
2. Tom thinks Bill likes 자기.
3. Tom thinks that Bill이 자기를 좋아한다.
4. 톰은 빌이 자기를 좋아한다고 think한다.
5. 톰은 빌이 자기 자신을 좋아한다고 생각한다.

[BEGIN TASK]

Prompt for generating gradual code-switching (Ko→En) few-shot demonstrations

You are a bilingual rewriting assistant.

Your task is to generate five versions of a sentence that gradually transition from Korean to English.

[INPUT]

- One English sentence (E)
- Its Korean translation (K)
- A code-switching version of the sentence (C), where about 50% of Korean words are replaced by English equivalents

[OUTPUT]

Generate a sequence of five sentences showing a smooth progression from Korean to English:

1. Korean only (100% Korean, source syntax S-O-V)
2. 75% Korean + 25% English (matrix language: Korean, embedded language: English)
3. 50% Korean + 50% English (matrix language: Korean, embedded language: English)
4. 25% Korean + 75% English (matrix language: Korean, embedded language: English)
5. English only (100% English, target syntax S-V-O)

[RULES]

Following the Matrix Language Frame (MLF) model,

- Preserve Korean word order (S-O-V) and syntax until version 5 (full English).
- Use English equivalents from E when inserting English into Korean sentences.
- Keep the code-switching natural and consistent, not random.
- Do not add explanations, notes, or extra text — output only the five sentences in order.

[EXAMPLES]

EXAMPLE 1

Input:

<Korean> 미나야, 바구니에 장난감을 다 넣고 빨리 집에 가자.

<English> Meena, put all the toys in the basket quickly, and go home.

<Code-Switching> Meena, basket 안에다 all the toys를 빨리 put하고 집에 가자.

Output:

1. 미나야, 바구니에 장난감을 다 넣고 빨리 집에 가자.
2. Meena, 바구니에 장난감을 다 put하고 빨리 집에 가자.
3. Meena, basket 안에다 all the toys를 빨리 put하고 집에 가자.
4. Meena, put all the toys in the basket quickly, 집에 가자.
5. Meena, put all the toys in the basket quickly, and go home.

EXAMPLE 2

Input:

<Korean> 지난 번에 도로교통공단이 내 운전면허증을 실수로 갱신하지 않았어요.

<English> Last time, by mistake, they did not renew my driver's license at the DMV.

<Code-Switching> 지난번에 motor vehicle department에서 내 driver's license를 mistake로 갱신하지 않았어요.

Output:

1. 지난 번에 도로교통공단이 내 운전면허증을 실수로 갱신하지 않았어요.
2. 지난 번에 motor vehicle department가 내 운전면허증을 실수로 갱신하지 않았어요.
3. 지난번에 motor vehicle department에서 내 driver's license를 mistake로 갱신하지 않았어요.
4. Last time, motor vehicle department에서 my driver's license를 mistake로 renew하지 않았어요.
5. Last time, by mistake, they did not renew my driver's license at the DMV.

Example 3

Input:

<Korean> 존은 메리가 어제 무엇을 샀는지 궁금해한다.

<English> John wonders what Mary bought yesterday.

<Code-Switching> John은 Mary가 yesterday 무엇을 샀는지 궁금해한다.

Output:

1. 존은 메리가 어제 무엇을 샀는지 궁금해한다.
2. John은 메리가 yesterday 무엇을 샀는지 궁금해한다.
3. John은 Mary가 yesterday what을 샀는지 궁금해한다.
4. John wonders what Mary가 yesterday 샀는지 궁금해한다.
5. John wonders what Mary bought yesterday.

[BEGIN TASK]

B.2 Code-switching In-context Learning

Prompt for CSICL

Answer the following questions written in non-English by explicitly showing a step-by-step translation process into English, then provide the final answer.

****Core Behaviors****

1. Mandatory Opening Self-Instruction

- Always begin with the exact sentence: "Let's gradually translate this non-English query into English, then think in English, and finally answer the question."

2. Gradual Code-Switching Output:

- Show the transformation from the original query to English in 5 steps.
- Each step progressively replaces more non-English words with English until the query is fully natural English.
- End with a clean, natural English rendering of the question.

3. Answer Format:

- End with the final answer in the exact format: "The answer is X", where 'X' is a single uppercase letter of the correct choice (e.g., A, B, C, D).
- No explanations or justifications after the answer.

Unlike hidden scratchpad reasoning, note that the translation and reasoning process must be explicitly output. Keep the structure identical across all responses.

{Few-shot demonstrations}

B.3 Evaluation Experiments

Prompt for multiple-choice QA experiments

You will be asked to answer a multiple-choice question. Read the following question and choices then select the single best answer.

- Output only the single uppercase letter of the correct choice (e.g., A, B, C, D).
- Do not output anything else: no reasoning, no explanation, no restating the question, no prefixes like "Answer:", no punctuation, spaces, or newlines.

The final output must be exactly one letter.

Prompt for short-answer QA experiments

You will be asked to answer a short-answer question. Read the following question and provide a single answer without any explanations.

- Output only the single answer.
- Do not output anything else: no reasoning, no explanation, no restating the question, no prefixes like "Answer:".

Prompt for machine translation experiments

You will be asked to translate a {target language} text. Read the following sentence and translate it into English without any explanations.

- Output only the English equivalent.
- Do not output anything else: no reasoning, no explanation, no restating the question, no prefixes like "Answer:".

C Additional Related Work

In this section, we provide further discussions on prior studies related to code-switching.

Code-switching. Code-switching, also known as code-mixing or language alternation, is a common linguistic phenomenon in which a speaker interleaves two or more languages within a single conversational context (Auer, 1998). It occurs in various switching levels: subwords such as at morpheme boundary (*i.e.*, intra-word switching), tag phrases (*i.e.*, tag-switching), words (*i.e.*, intra-sentential switching), and sentences or clauses (*i.e.*, inter-sentential switching). Matrix Language Frame (MLF) model, which distinguishes a grammatically dominant *matrix language* (ML) and an inserted *embedded language* (EL) (Myers-Scotton, 1997), has been widely adopted as a syntactic rule for code-switching. The following examples illustrate code-switching, with variation in the choice of matrix language and embedded language. The sentence originates from Global MMLU (Singh et al., 2025): “*Paper will burn at approximately what temperature in Fahrenheit?*”

- (1) *Paper will brûler at approximately quelle température in Fahrenheit ?*
 ML ML EL ML ML EL EL ML ML

(ML: English, EL: French)

- (2) *En degrés Fahrenheit, à approximately what temperature le papier*
 ML ML ML ML EL EL EL ML ML
burn -t-il ?
 EL ML

(ML: French, EL: English)

Code-switching text generation. Only a limited number of code-switching corpora and labeled datasets exist for specific language pairs, and code-switching among non-English languages is hardly available (Winata et al., 2023). Hence, the synthetic generation of code-switching has become the primary strategy. Early efforts such as Jayanthi et al. (2021); Rizvi et al. (2021) proposed toolkits for Hindi-English, although they lack generalizability across languages. More recently, prompting LLMs (*i.e.*, GPT-4) has been widely adopted for synthetic data generation Yoo et al. (2025a;b); Kim et al. (2025); Yang & Chai (2025). In particular, Yoo et al. (2025a) further provided quantitative and qualitative analyses comparing LLM-generated code-switching text with human bilingual data, showing that such text largely follows natural switching patterns with one distinctive feature—redundant synonyms appearing in both languages—which may actually help LLMs by providing explicit lexical alignment cues. Therefore, a slight “unnaturalness” of synthetic code-switching may not be a weakness but rather a source of beneficial noise that facilitates multilingual transfer. In addition, Xie et al. (2025) introduced LinguaMaster, a LLM multi-agent collaboration framework specifically designed for controlled code-switching generation.

Code-switching for multilingual evaluations. As the multilingual capabilities of LLMs advance, recent benchmarks range from broad NLP tasks (Yang & Chai, 2025; Abdaljalil et al., 2025; Mohamed et al., 2026) to domain-specific evaluations, including translation (Huzaifah et al., 2024; Zhang et al., 2025), dialogue summarization (Suresh et al., 2025), and red-teaming (Song et al., 2025; Yoo et al., 2025b). Despite these efforts, Zhang et al. (2023) reported that LLMs still struggle with code-switching, sometimes underperforming compared to smaller, fine-tuned systems. This suggests that code-switching remains an unsolved challenge for multilingual evaluation, rather than a solved proxy task.

D Additional Results and Analyses

D.1 Ablation Study

Scalability of CSICL. We examine the effectiveness of **CSICL** on small language models (SLMs), which fall behind larger models in instruction-following performance. We test different sizes of Qwen 3 (*i.e.*, 0.6B, 1.7B, 4B, 8B, 14B, and 32B) on the sampled set of Global MMLU in Korean as a target language. Figure 6 reports the target language performance of X-ICL baselines and **CSICL** across model sizes, shown in a log scale. While the gap between **CSICL** and X-ICL baselines is particularly pronounced when the model size is larger than 4B, we observe that **CSICL** consistently outperforms all X-ICL baselines across model sizes.




Figure 6: Experimental results of X-ICL approaches in the target language (Ko) on Global MMLU using different sizes of Qwen 3. X-axis is on a log scale.

Extended Thinking. Several frontier models have strong native chain-of-thought or reasoning capabilities. We investigate whether extended thinking of LLMs diminishes the benefit of **CSICL**, by mitigating the latent translation barrier. Table 7 reports the experimental results of Qwen3-32B on Global MMLU with the thinking mode enabled. As expected, extended thinking raises the baseline, especially for mid- and low-resource languages, thereby reducing **CSICL**’s marginal gains. However, the gains do not disappear. **CSICL**’s gains slightly reduce from +6.0→+5.8pp in target languages, +1.3→+1.0pp in unseen high-resource languages, +2.1→+1.8pp in unseen mid-resource languages, and +4.8→+2.9pp in unseen low-resource languages. **CSICL** still significantly improves performance in the target and unseen mid- to low-resource languages across all X-ICL baselines. We further conduct a qualitative error analysis of the model’s reasoning traces. We observe that the extended thinking often helps the model explicitly translate or paraphrase the non-English query before reasoning. However, in failure cases, the model sometimes mistranslates a key entity, relation, or constraint early in the trace and then reasons consistently but incorrectly from that corrupted English surrogate Bafna et al. (2025). Once the reasoning trajectory enters this wrong path, the model rarely self-corrects later. In contrast, **CSICL** often mitigates these failures by making the cross-lingual transition explicit and gradual: key content is preserved across intermediate code-switched steps before the model fully reasons in English. This suggests that native extended thinking and **CSICL** are partially overlapping but complementary.

Method	X-ICL setting		En	Tgt.*	Unseen Lang.		
	Demonstration	Instruction			High	Mid*	Low*
Zero-shot learning	✗	✗	93.3	73.8	88.1	67.1	47.2
Few-shot learning	✓ Monolingual (En)	✗	93.4	75.3	88.4	67.8	49.1
	✓ Monolingual (Tgt.)	✗	93.1	75.4	86.0	67.1	47.3
	✓ Parallel	✗	93.3	76.1	88.5	67.7	48.8
Zero-shot CoT	✗	✓ Translation (Tgt.→En)	93.2	76.8	88.4	68.5	49.5
	✗	✓ Translation (Tgt.→Rnd.)	93.0	78.8	86.2	65.2	49.4
CSICL	✓ Gradual CS (Tgt.→En)	✓ Gradual Translation (Tgt.→En)	93.2	81.1	89.4	69.6	52.0

Table 7: Ablation results comparing **CSICL** to existing X-ICL baselines using Qwen 3 with thinking mode enabled. X-ICL setting describes each method in terms of a combination of its few-shot demonstrations and instruction. **CSICL** outperforms X-ICL baselines in both target and unseen languages, yielding stable results in English.

Demonstration Sensitivity. We test the sensitivity of the demonstration of **CSICL**. Table 8 reports the experimental results on Global MMLU using Qwen3-32B under five different random seeds and dynamic demonstrations, where the target language is Korean. Specifically,

we use the same balanced Global MMLU subset as in the main experiments, consisting of 3.6k instances. We first convert all Korean instances into candidates for gradual code-switching demonstrations using the **CSICL** construction pipeline. Then, for each evaluation query, we randomly sample 5 gradual code-switching demonstrations from the remaining instances, explicitly excluding the current query itself. We repeat this entire process with 5 different random seeds. Under this randomized demonstration-sampling setting, **CSICL** achieves 81.5 ± 1.2 accuracy, while the strongest X-ICL baseline achieves 78.8. The results show that **CSICL** is not dependent on the particular demonstrations or seeds. **CSICL** outperforms all X-ICL baselines in the target language and mid- to low-resource unseen languages.

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	89.1$\pm$0.4	75.5 $\pm$ 0.4	89.4 $\pm$ 0.2	83.8 $\pm$ 0.3	65.3 $\pm$ 0.5	60.7 $\pm$ 0.2	44.2 $\pm$ 0.5	37.6 $\pm$ 0.3
Monolingual (Tgt.)	88.7 $\pm$ 0.2	77.0 $\pm$ 0.3	90.0 $\pm$ 0.4	83.0 $\pm$ 0.2	64.7 $\pm$ 0.2	<u>61.6$\pm$0.5</u>	42.3 $\pm$ 0.3	35.8 $\pm$ 0.6
Parallel	88.6 $\pm$ 0.4	75.6 $\pm$ 0.2	89.7 $\pm$ 0.4	83.3 $\pm$ 0.2	65.0 $\pm$ 0.4	61.1 $\pm$ 0.4	44.6 $\pm$ 0.3	37.7 $\pm$ 0.3
Translation (Tgt. $\rightarrow$ En)	88.9 $\pm$ 0.6	78.8 $\pm$ 0.2	90.1$\pm$0.2	84.1 $\pm$ 0.3	65.2 $\pm$ 0.2	61.1 $\pm$ 0.4	44.9 $\pm$ 0.2	38.5 $\pm$ 0.2
Translation (Rnd. $\rightarrow$ En)	88.4 $\pm$ 0.4	78.6 $\pm$ 0.4	89.9 $\pm$ 0.1	84.5$\pm$0.3	65.1 $\pm$ 0.5	61.0 $\pm$ 0.2	44.8 $\pm$ 0.3	<u>38.3$\pm$0.2</u>
CSICL (Tgt. $\rightarrow$ En)	88.5 $\pm$ 0.5	81.5$\pm$1.2	90.1$\pm$0.8	<u>84.2$\pm$1.0</u>	67.5$\pm$1.3	63.0$\pm$1.1	48.3$\pm$1.4	41.6$\pm$1.8

Table 8: Ablation results comparing **CSICL** to existing X-ICL baselines using Qwen 3 on Global MMLU with five different random seeds and dynamic demonstrations.

D.2 Is CSICL More Than Translation or Prompt Expansion?

Table 9 reports experimental results on Qwen3-32B using other X-ICL baselines, controlling the combinations of few-shot demonstrations setting and instruction setting, where **CSICL** outperforms all X-ICL baselines.

Method	X-ICL setting		En	Tgt.*	Unseen Lang.		
	Demonstration	Instruction			High*	Mid*	Low*
Zero-shot CoT	X		88.7	74.4	87.5	63.7	42.3
	✓ Monolingual (Tgt.)	✓ Translation (Tgt. $\rightarrow$ En) + Reason in En	<u>88.7</u>	74.0	<u>87.6</u>	63.7	42.1
Few-shot learning	✓ Monolingual (Tgt.)	✓ Translation (Tgt. $\rightarrow$ En)	88.8	75.3	87.2	63.3	41.8
	✓ Parallel	✓ Translation (Tgt. $\rightarrow$ En)	<u>88.7</u>	<u>75.0</u>	87.4	64.0	<u>43.2</u>
	✓ Parallel	✓ Translation (Tgt. $\rightarrow$ En) + Reason in En	<u>88.7</u>	75.0	87.4	<u>64.2</u>	43.1
	✓ Monolingual (Tgt.)	✓ Translation (Tgt. $\rightarrow$ Rnd.)	88.6	74.8	87.5	63.5	42.0
CSICL	✓ Gradual CS (Tgt. $\rightarrow$ En)	✓ Gradual Translation (Tgt. $\rightarrow$ En)	88.6	76.8	87.8	64.9	46.0

Table 9: Ablation results of other X-ICL baselines using Qwen 3, controlling the few-shot demonstrations setting and the instruction setting.

Translation baselines. We add an ablation that explicitly translates the query into English and instructs the model to reason in English. Translation vs. Translation-then-reason baseline yields comparable results, and both consistently remain below **CSICL**. This suggests that **CSICL** is not merely improving translation quality but facilitating cross-lingual representational alignment.

Pivot language. We also add an ablation using random high-resource, random pivot languages (*i.e.*, German, Russian, Dutch, and Japanese), ensuring they do not overlap with the evaluation languages. Translation to random vs. random high-resource pivot languages yields comparable results, and both are consistently outperformed by **CSICL**.

Paraphrasing control. **CSICL** may benefit from using a larger number of sentences, since **CSICL** includes four additional sentences per demonstration. To control for this, we fix the number of sentences per demonstration to five, the same as in **CSICL**, and compare it against monolingual demonstrations. Specifically, for each shot, we add four paraphrased monolingual demonstrations, generated by GPT-5, and then ask the LLMs to solve the task.

Table 10 describes experimental results of **CSICL** and paraphrased demonstrations on Global MMLU, averaged over Qwen 3 and Gemini 2.5. Although paraphrasing provides a marginal improvement in cross-lingual transfer for both target (+0.2pp) and unseen languages (+0.3pp) on average, its effect is limited; **CSICL** consistently outperforms both baselines (+6.0pp and +4.8pp in the target and the unseen languages, respectively, on average). This suggests that the gains of **CSICL** cannot be attributed to a larger demonstration budget alone. Instead, the gradual transition across languages supplies a distinct cross-lingual signal that paraphrasing fails to capture, reinforcing the role of code-switching in aligning multilingual representations.

We extend this examination by controlling the number of English and Korean few-shot demonstrations. Table 11 shows experimental results on Global MMLU using Qwen 3 under different combinations of paraphrasing baselines. We observe mixed results across the paraphrasing baselines for target language and for unseen languages in high- and mid-resource languages. Interestingly, we find that increasing English demonstrations helps only in unseen low-resource languages. However, **CSICL** consistently outperforms all paraphrasing baselines across target and unseen languages.

X-ICL setting	En	Tgt.*	Unseen Lang.		
			High*	Mid*	Low*
Paraphrasing (En)	88.9	71.0	86.6	63.2	41.6
Paraphrasing (Tgt.)	88.8	<u>72.3</u>	<u>87.0</u>	62.6	39.3
CSICL (Tgt.→En)	<u>88.6</u>	76.8	87.8	64.9	46.0

Table 10: Experimental results comparing **CSICL** to paraphrased, monolingual demonstrations, keeping the number of sentences per shot.

# En	# Tgt.	En	Tgt.*	Unseen Lang.		
				High*	Mid*	Low*
5	0	88.9	71.0	86.6	63.2	41.6
4	1	88.6	71.3	86.8	<u>63.2</u>	<u>41.0</u>
3	2	88.9	71.7	86.9	<u>62.7</u>	40.7
2	3	88.7	72.0	87.1	63.0	40.2
1	4	88.7	72.5	<u>86.7</u>	62.7	39.5
0	5	88.8	<u>72.3</u>	87.0	62.6	39.3
CSICL		88.6	76.8	87.8	64.9	46.0

Table 11: Experimental results comparing **CSICL** to paraphrasing baselines, controlling the number of English and target language few-shot demonstrations. # En and # Tgt. denote the number of English demonstrations and the number of target language demonstrations, respectively.

D.3 Human Validation on COMET

We additionally conduct a pairwise human evaluation to validate the experimental results obtained with COMET. We randomly sample 200 FLORES+ examples and compare **CSICL** against the best-performing baseline for each source language (*i.e.*, Parallel for target language translation and Translation (Tgt.→Rnd.) for unseen language translation, respectively). Following the WMT practice, one author, who is a native speaker of Korean and fluent in English, is shown two anonymized translations (**CSICL** vs. baseline) and asked to choose “A wins / B wins / tie” based on overall adequacy and fluency. **CSICL** is preferred in 63% of cases (and 16% ties) on target-language performance, and in 68% of cases (and 9% ties) on unseen-language performance, respectively. These results provide additional evidence that **CSICL** produces more adequate and fluent translations, as measured by both automatic metrics and human preference.

D.4 Full Results

Tables 12, 13, 14, and 15 present experimental results of X-ICL approaches on Global MMLU using Qwen 3, DeepSeek 3.1, Grok 4, and Gemini 2.5. Table 16 presents experimental results of X-ICL approaches on five different tasks using DeepSeek 3.1, Grok 4, and Gemini 2.5.

X-ICL setting	En	Tgt.*	Zh	Es*	Id*	Tr*	Sw*	Te*
Monolingual (En)	88.6	83.5	89.5	83.2	64.8	60.5	44.6	38.2
Monolingual (Tgt.)	88.7	83.3	90.3	83.8	64.3	59.9	42.8	36.9
Parallel	88.5	83.7	90.4	82.9	65.1	60.6	44.7	38.0
Translation (Tgt.→En)	88.7	83.6	90.5	83.6	65.5	61.1	45.3	38.8
Translation (Rnd.→En)	88.7	<u>83.8</u>	90.2	83.7	<u>66.0</u>	<u>61.2</u>	<u>45.1</u>	<u>39.0</u>
CSICL (Tgt.→En)	88.5	84.9	90.5	83.9	66.6	61.8	48.4	42.8

(a) Target: French (*high*)

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	88.9	75.4	89.3	83.7	65.1	60.8	44.2	37.9
Monolingual (Tgt.)	88.7	77.2	90.0	83.0	64.8	61.3	42.3	35.7
Parallel	88.9	75.6	89.7	83.6	64.9	<u>61.5</u>	44.4	37.6
Translation (Tgt.→En)	88.9	79.4	90.2	84.1	65.5	61.4	44.8	38.6
Translation (Rnd.→En)	88.5	<u>79.1</u>	90.0	<u>84.3</u>	<u>65.2</u>	61.2	<u>45.1</u>	<u>38.3</u>
CSICL (Tgt.→En)	88.6	80.6	90.3	84.4	67.4	62.8	48.5	41.7

(b) Target: Korean (*mid*)

X-ICL setting	En	Tgt.*	Zh*	Es*	Id*	Tr*	Sw*	Te*
Monolingual (En)	<u>88.9</u>	53.5	89.8	83.5	64.8	60.8	44.6	37.7
Monolingual (Tgt.)	89.0	55.5	87.9	<u>86.4</u>	63.9	58.4	43.5	31.0
Parallel	88.7	58.8	90.4	85.6	64.2	61.7	48.8	34.9
Translation (Tgt.→En)	88.8	60.5	90.3	85.7	64.9	63.8	50.0	34.5
Translation (Rnd.→En)	88.6	<u>58.5</u>	<u>90.6</u>	86.2	<u>65.0</u>	<u>64.2</u>	<u>44.9</u>	<u>41.4</u>
CSICL (Tgt.→En)	88.7	64.9	90.8	86.9	66.1	64.7	51.2	43.4

(c) Target: Yoruba (*low*)Table 12: Full experimental results comparing **CSICL** to X-ICL baselines using Qwen3-32B.

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	89.6	84.6	91.2	84.3	66.2	61.5	45.3	38.8
Monolingual (Tgt.)	89.5	84.4	91.8	<u>84.9</u>	65.7	61.0	43.6	37.5
Parallel	89.4	84.7	92.0	84.0	66.5	61.7	45.4	38.6
Translation (Tgt.→En)	89.6	84.8	<u>92.1</u>	84.8	66.8	62.1	46.0	39.2
Translation (Rnd.→En)	89.6	<u>85.0</u>	<u>91.9</u>	<u>84.9</u>	<u>67.2</u>	<u>62.2</u>	<u>45.8</u>	<u>39.3</u>
CSICL (Tgt.→En)	89.4	86.0	92.3	85.1	68.0	62.9	49.1	43.1

(a) Target: French (*high*)

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	89.8	77.3	91.0	84.6	66.4	61.8	45.0	38.6
Monolingual (Tgt.)	89.6	79.1	91.6	83.9	66.0	62.3	43.2	36.5
Parallel	89.8	77.5	91.4	84.5	66.1	<u>62.5</u>	45.2	38.3
Translation (Tgt.→En)	89.8	81.1	91.8	85.0	66.7	62.4	45.6	39.0
Translation (Rnd.→En)	89.4	80.8	<u>91.5</u>	<u>85.1</u>	<u>66.4</u>	62.2	<u>45.9</u>	<u>38.7</u>
CSICL (Tgt.→En)	89.5	82.4	92.0	85.3	68.6	63.6	49.2	41.7

(b) Target: Korean (*mid*)

X-ICL setting	En	Tgt.*	Zh	Es*	Id*	Tr	Sw*	Te*
Monolingual (En)	89.8	55.3	91.5	84.7	66.3	61.9	45.2	38.2
Monolingual (Tgt.)	89.9	57.3	89.6	<u>87.4</u>	65.5	59.5	44.1	32.0
Parallel	89.7	60.6	92.0	86.6	65.8	62.7	49.3	35.9
Translation (Tgt.→En)	89.8	<u>62.3</u>	91.9	86.7	66.6	64.5	50.5	35.6
Translation (Rnd.→En)	<u>89.6</u>	<u>60.2</u>	<u>92.3</u>	87.2	<u>66.7</u>	<u>65.0</u>	<u>45.4</u>	<u>42.2</u>
CSICL (Tgt.→En)	89.7	66.7	92.5	87.8	67.7	65.4	51.6	44.2

(c) Target: Yoruba (*low*)Table 13: Full experimental results comparing **CSICL** to X-ICL baselines using DeepSeek-chat-v3.1.

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	88.3	83.0	89.8	83.6	65.0	60.3	44.0	37.5
Monolingual (Tgt.)	88.1	82.9	90.5	84.1	64.6	60.6	42.2	36.2
Parallel	88.2	83.3	90.6	83.4	65.2	60.9	44.2	37.2
Translation (Tgt.→En)	88.3	83.2	<u>90.7</u>	84.0	65.6	<u>61.3</u>	44.7	37.9
Translation (Rnd.→En)	88.1	<u>83.5</u>	<u>90.4</u>	<u>84.1</u>	<u>65.7</u>	<u>61.3</u>	44.6	<u>38.1</u>
CSICL (Tgt.→En)	88.0	84.5	90.8	84.3	66.5	61.9	47.7	41.7

(a) Target: French (*high*)

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	88.4	74.6	89.6	83.9	65.3	60.6	43.7	37.2
Monolingual (Tgt.)	88.2	76.4	90.2	83.2	65.0	61.1	41.9	35.1
Parallel	88.4	<u>78.6</u>	89.9	83.8	65.1	<u>61.3</u>	44.0	36.9
Translation (Tgt.→En)	88.4	74.8	90.4	84.2	65.7	61.2	44.3	37.8
Translation (Rnd.→En)	88.1	78.3	<u>90.1</u>	<u>84.4</u>	<u>65.4</u>	61.1	<u>44.6</u>	37.5
CSICL (Tgt.→En)	88.2	79.7	90.6	84.6	67.1	62.3	47.8	40.8

(b) Target: Korean (*mid*)

X-ICL setting	En	Tgt.*	Zh	Es*	Id*	Tr*	Sw*	Te*
Monolingual (En)	88.4	52.7	90.1	83.8	65.1	60.7	43.9	36.9
Monolingual (Tgt.)	88.5	54.6	88.2	86.7	64.2	58.2	42.8	30.2
Parallel	88.2	<u>59.5</u>	90.7	85.9	64.7	61.4	47.9	34.3
Translation (Tgt.→En)	88.3	57.8	90.6	86.0	65.4	63.5	49.1	34.0
Translation (Rnd.→En)	88.1	57.6	<u>90.9</u>	<u>86.5</u>	<u>65.5</u>	<u>63.8</u>	<u>44.0</u>	<u>40.6</u>
CSICL (Tgt.→En)	88.2	63.7	91.0	87.0	66.5	64.2	50.3	42.5

(c) Target: Yoruba (*low*)Table 14: Full experimental results comparing **CSICL** to X-ICL baselines using grok-4-fast.

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	90.2	85.6	91.5	85.1	67.2	62.7	46.1	39.6
Monolingual (Tgt.)	90.0	85.4	92.1	85.7	66.7	62.3	44.3	38.3
Parallel	90.1	85.8	92.2	84.9	67.5	63.0	46.3	39.4
Translation (Tgt.→En)	90.2	85.9	92.3	85.6	67.9	63.4	46.9	40.1
Translation (Rnd.→En)	90.1	<u>86.1</u>	<u>92.1</u>	<u>85.7</u>	<u>68.3</u>	<u>63.5</u>	<u>46.8</u>	<u>40.3</u>
CSICL (Tgt.→En)	90.0	87.3	92.5	85.9	69.2	64.1	49.9	44.3

(a) Target: French (*high*)

X-ICL setting	En	Tgt.*	Zh	Es*	Id*	Tr*	Sw*	Te*
Monolingual (En)	90.3	79.3	91.3	85.5	67.6	63.0	45.8	39.4
Monolingual (Tgt.)	90.1	81.1	91.9	84.8	67.2	63.5	44.0	37.3
Parallel	90.3	79.5	91.6	85.4	67.3	<u>63.7</u>	46.0	39.1
Translation (Tgt.→En)	90.3	83.0	92.0	85.9	67.9	63.6	46.4	39.9
Translation (Rnd.→En)	90.0	<u>82.7</u>	<u>91.7</u>	<u>86.1</u>	<u>67.6</u>	63.5	<u>46.7</u>	39.6
CSICL (Tgt.→En)	90.1	84.3	92.2	86.4	69.7	64.8	50.1	42.8

(b) Target: Korean (*mid*)

X-ICL setting	En	Tgt.*	Zh	Es	Id*	Tr*	Sw*	Te*
Monolingual (En)	90.3	56.8	91.8	85.4	67.4	63.1	46.0	39.0
Monolingual (Tgt.)	90.4	58.9	89.8	<u>88.3</u>	66.6	60.8	44.9	31.7
Parallel	90.2	62.1	92.4	87.5	66.9	63.8	50.1	35.6
Translation (Tgt.→En)	90.3	63.9	92.3	87.6	67.7	65.7	51.3	35.3
Translation (Rnd.→En)	<u>90.1</u>	<u>61.8</u>	<u>92.6</u>	88.1	<u>67.8</u>	<u>66.2</u>	<u>46.3</u>	<u>42.0</u>
CSICL (Tgt.→En)	90.2	66.2	92.8	88.6	68.9	66.6	52.5	44.8

(c) Target: Yoruba (*low*)Table 15: Full experimental results comparing **CSICL** to X-ICL baselines using Gemini 2.5 Flash.

Task type	Translation		Reasoning-oriented						Knowledge-intensive					
			Domain-specific			Math			Cultural			Social Bias		
X-ICL setting	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.
Monolingual (En)	78.5	73.1	45.2	38.4	30.6	58.5	52.1	50.3	84.1	78.5	74.1	92.5	89.1	86.8
Monolingual (Tgt.)	76.5	72.4	46.5	40.1	28.5	59.2	49.5	48.8	83.8	72.1	71.9	90.1	92.0	86.2
Parallel	80.1	79.0	48.0	41.5	29.1	57.5	51.0	50.1	84.5	80.1	73.0	90.5	92.1	86.5
Translation (Tgt.→En)	78.2	75.5	47.5	40.5	28.0	58.4	57.1	51.5	84.0	79.5	72.5	91.0	92.5	87.0
Translation (Rnd.→En)	77.4	74.8	46.4	41.0	29.8	58.8	58.2	52.0	84.2	79.0	73.1	92.4	90.1	87.8
CSICL	85.5	84.8	46.8	43.5	36.3	59.0	60.1	56.2	84.0	82.5	75.8	92.4	93.2	88.0

(a) DeepSeek 3.1

Task type	Translation		Reasoning-oriented						Knowledge-intensive					
			Domain-specific			Math			Cultural			Social Bias		
X-ICL setting	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*
Monolingual (En)	75.2	70.8	42.4	44.8	31.0	49.5	47.7	48.1	82.2	77.6	70.3	90.5	87.4	80.1
Monolingual (Tgt.)	73.8	70.2	43.6	45.1	29.3	50.1	48.2	48.7	81.8	78.2	67.8	88.2	90.3	82.1
Parallel	71.6	76.6	44.8	38.6	30.6	49.8	48.1	46.1	82.3	77.7	70.4	88.0	90.0	81.6
Translation (Tgt.→En)	70.6	73.4	44.6	38.0	30.9	49.2	49.2	49.7	81.8	76.9	70.9	90.6	88.6	80.4
Translation (Rnd.→En)	69.8	73.0	44.2	38.4	31.2	48.7	48.8	49.8	81.3	76.5	70.5	90.6	88.3	82.9
CSICL	76.8	81.6	42.8	48.9	35.2	50.1	52.8	50.2	82.0	80.8	72.1	89.0	90.3	83.8

(b) Grok 4

Task type	Translation		Reasoning-oriented						Knowledge-intensive					
			Domain-specific			Math			Cultural			Social Bias		
X-ICL setting	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.*	En	Tgt.*	Uns.
Monolingual (En)	79.5	75.2	48.6	49.2	34.3	54.1	52.0	52.4	86.3	85.7	80.5	94.6	91.5	94.2
Monolingual (Tgt.)	77.8	74.5	48.8	49.6	32.6	54.8	52.6	53.2	85.9	84.5	81.1	92.5	94.5	93.8
Parallel	76.1	81.1	49.1	53.0	35.0	54.3	54.5	50.4	86.5	85.8	79.8	92.2	94.1	92.7
Translation (Tgt.→En)	75.3	77.8	48.8	52.4	35.5	53.6	55.8	54.1	85.9	86.1	80.2	94.7	92.7	94.5
Translation (Rnd.→En)	74.5	77.4	48.4	52.8	36.9	53.1	53.2	54.2	85.4	86.6	79.8	94.8	92.4	92.0
CSICL	86.2	83.3	49.3	56.5	39.8	54.5	57.6	55.8	86.2	88.0	82.2	93.3	94.7	94.3

(c) Gemini 2.5

Table 16: Full experimental results on five diverse tasks and domain knowledge.