

Reasoning Transfer for an Extremely Low-Resource and Endangered Language: Bridging Languages Through Sample-Efficient Language Understanding

Khanh-Tung Tran*, Barry O’Sullivan, Hoang D. Nguyen

Research Ireland Centre for Research Training in Artificial Intelligence and Insight Research Ireland Centre for Data Analytics
School of Computer Science and Information Technology, University College Cork, Ireland
123128577@umail.ucc.ie, b.osullivan@cs.ucc.ie, hn@cs.ucc.ie

Abstract

Recent advances have enabled Large Language Models (LLMs) to tackle reasoning tasks by generating chain-of-thought (CoT) rationales, yet these gains have largely applied to high-resource languages, leaving low-resource languages behind. In this work, we first investigate CoT techniques in extremely low-resource scenarios through previous prompting, model-editing, and fine-tuning approaches. We introduce *English-Pivoted CoT Training*, leveraging the insight that LLMs internally operate in a latent space aligned toward the dominant language. Given input in a low-resource language, we perform supervised fine-tuning to generate CoT in English and output the final response in the target language. Across mathematical reasoning benchmarks, our approach outperforms other baselines with up to 28.33% improvement in low-resource scenarios. Our analysis and additional experiments, including Mixed-Language CoT and Two-Stage Training, show that explicitly separating language understanding from reasoning enhances cross-lingual reasoning abilities. To facilitate future work, we also release *LC2024*, the first benchmark for mathematical tasks in Irish, an extremely low-resource and endangered language. Our results and resources highlight a practical pathway to multilingual reasoning without extensive retraining in every extremely low-resource language, despite data scarcity.

Code, Dataset, Appendix —

<https://github.com/ReML-AI/english-pivoted-cot>

Introduction

Large Language Models (LLMs) have showcased exceptional performance in a wide-range of task (Zhao et al. 2025; Zhang, Jin, and Zhuo 2024), particularly in English, due to the vast amount of available data during pre-training (Qin et al. 2024; Guo et al. 2023; Kim, Koo, and Lim 2024). Recent advances in post-training techniques, such as reasoning scaling (OpenAI 2024b), which explicitly trains models to generate chain-of-thought (CoT) reasoning, have significantly enhanced model accuracy in tasks such as mathematics and programming. However, such advances have unintentionally widened the performance gap between popular

languages like English and low, extremely low-resource languages (Ji et al. 2025). This disparity arises primarily because post-training methods typically require high-quality, manually curated datasets, predominantly available only in English, and misalignment and inherent biases within multilingual training corpora (Ghosh et al. 2025).

Figure 1 illustrates this phenomenon by comparing outputs from a small-size state-of-the-art reasoning model (r1-distill-Llama-8B (DeepSeek-AI 2025)) given the same mathematical problem presented in English and in Irish, classified endangered by UNESCO (UNESCO 2010), less than 0.01% appearance in Common Crawl (Common Crawl 2025). When the question is presented in English, the model reasoning process can lead to correct solution, however, with the same problem in Irish, the model misunderstands, unable to converse in target language, causing the CoT process to fail and resulting in an incorrect answer. Our proposed solution addresses this issue by aligning the model’s reasoning process across languages, allowing it to “think” internally in its dominant language (English).

Enabling CoT reasoning in limited-resource languages is critical for developing AI systems that function effectively across diverse linguistic contexts. Despite its importance, this remains a relatively unexplored domain, leaving low-resource and extremely low-resource languages underrepresented. Current low-resource CoT techniques can be split into three broad categories: (1) multilingual fine-tuning (Lai and Nissim 2024; Anand et al. 2024), (2) prompting strategies (translation-based (Shi et al. 2023) or self alignment-based prompting (Qin et al. 2023)), and (3) model editing (model merging (Yadav et al. 2023; Tao et al. 2024), layer swapping (Bandarkar et al. 2025)). In general, while these methods show promise, they often require extra training data, additional modules, or incur translation errors, limiting their practical usability in truly low-resource settings.

Inspired by recent advances in reasoning scaling, where models explicitly generate CoT traces without the need for carefully crafted reasoning prompts, and recent findings on the dynamics of language models across layers (Shi et al. 2023; Wendler et al. 2024; Tran, O’Sullivan, and Nguyen 2024) (where distinct layers/neurons specialize in language understanding and others in reasoning, typically biased towards English), we explore the connections between reasoning and multilinguality, particularly how reasoning can

*Corresponding author.

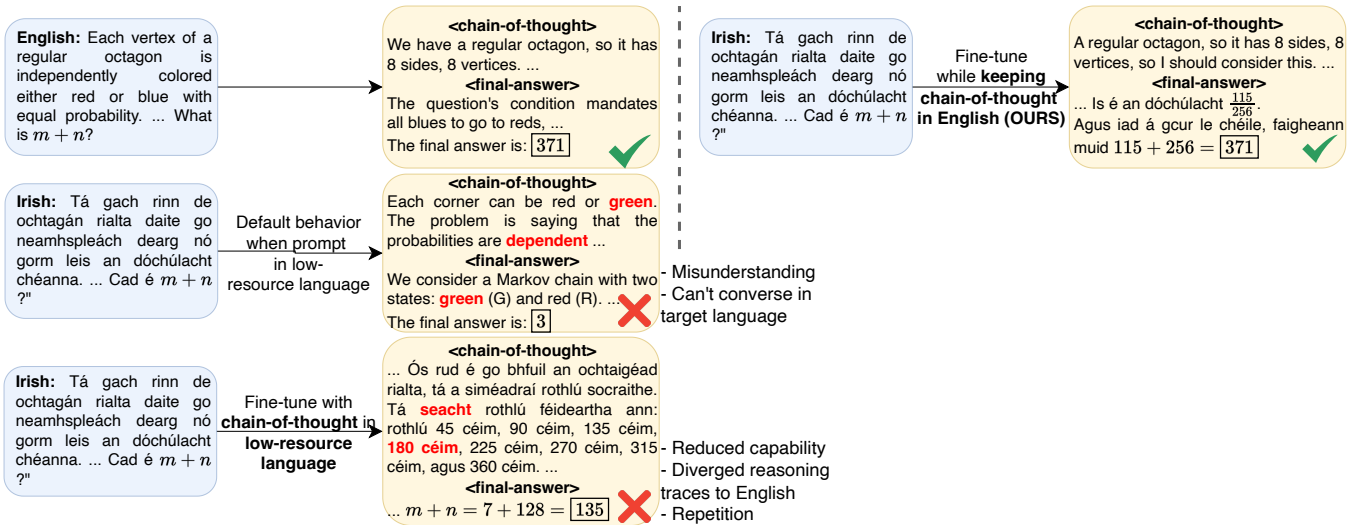


Figure 1: Illustrative example of model behavior (r1-distill-Llama-8B) when prompted with the same problem in English (robust reasoning) versus a low-resource language - Irish (reduced understanding and conversational ability). Training with an Irish chain-of-thought diverges from baseline, while training with English chain-of-thoughts achieves the best of both worlds.

be transferred across languages by keeping the explicit CoT in English, while boosting the low-resource language understanding and generation of the model while keeping inputs and final responses in the target language. Our method allows the LLM to interpret a problem in the target language, perform CoT in English, where it is most robust, and then subsequently generate accurate final responses in the target language. Our contributions are summarized as follows:

- *English-Pivoted CoT Training*, a novel and effective method for adapting LLM’s reasoning capabilities to extremely low-resource scenarios. Our approach explicitly aligns reasoning process across languages, addressing linguistic misalignment and boosting performances.
- Our results demonstrate successful transfer of reasoning capability to Irish, an extremely low-resource language, by leveraging model’s internal reasoning representations across languages, upto 28.33% improvement.
- Our analysis and ablations, including generalization to medium and high-resource languages, mixed CoT language and two-stage training further highlight the benefit of separating language understanding and reasoning, demonstrating a practical pathway for multilingual reasoning without extensive retraining in each language.
- *LC2024*, the first-ever dataset for evaluating mathematical reasoning in Irish, an endangered language.

Related Works

CoT Reasoning for Low-Resource Languages

Reasoning is formally described as the cognitive process of logically analyzing available information to reach conclusions, enabling both humans and AI systems to address complex problems (Ghosh et al. 2025). Chain-of-thought (Wei et al. 2022; Zelikman et al. 2022) has emerged as a powerful technique for improving reasoning in LLMs by gen-

erating explicit step-by-step rationales. Recent works explore CoT in low-resource languages (Shi et al. 2023; Qin et al. 2023). (Shi et al. 2023) finds that few-shot prompting with CoT in English consistently achieves competitive performance to prompting in the native language. To bridge the gap between languages, alignment techniques have been carried out to align representation between low-resource languages and English through learning with parallel sentences (Zhu et al. 2024; Li et al. 2024b) or additional multilingual encoders (Huang et al. 2024). On the other hand, there are also attempts to directly fine-tune LLMs in reasoning tasks across multiple languages, by obtaining reasoning data in low-resource languages, mostly through neural machine translation (Anand et al. 2024; Lai and Nissim 2024).

While prior works highlight stronger performance when aligning language models to reason in English for tasks in low-resource languages, they typically leverage prompting techniques or require additional modules or parallel corpus. In this work, we explicitly train LLMs to separate their reasoning and language understanding capabilities by separating the languages of internal CoT and final response.

Alignment of Internal Representations in Multilingual LLMs

There has been a large interest in investigating and aligning how multilingual LLMs organize and share knowledge between languages internally. A key question is whether LLMs “think” in English, or rely on English-centric representations or rationales even when operating in other languages. There is evidence that multilingual models often map other languages into an implicit English latent space (Li et al. 2024a; Zhu et al. 2024). Recent works analyze intermediate embeddings across transformer layers and found distinct phases of operation, where middle layers perform reasoning in a latent space steer toward dominant language (English) (Wendler




Figure 2: Loss curves over the training process for Native CoT Training (Left) and our approach, English-Pivoted CoT Training (right), normalized with exponential moving average with a smoothing weight of 0.95. Our method shows lower initial reasoning loss and slower decline, indicating effective separation of English reasoning from target-language responses.

et al. 2024; Tran, O’Sullivan, and Nguyen 2024).

These alignment efforts (Tao et al. 2024; Bandarkar et al. 2025) seek to improve cross-lingual generalization, allowing knowledge learned in one language to transfer to others. Our work is informed by these studies but takes a distinct direction. While prior works often relies on external models or separate bilingual modules, our approach aligns the reasoning process within the model’s own representations. This allows us to transfer reasoning skills to low-resource languages without retraining from scratch, advancing the understanding of how multilingual LLMs can maintain consistent reasoning across diverse languages.

English-Pivoted CoT Training for Reasoning Transfer

We introduce *English-Pivoted CoT Training*, leveraging the model’s robust reasoning capabilities developed predominantly in English, by explicitly constraining the intermediate CoT reasoning steps to English, while maintaining inputs and final responses in the target low-resource language. Formally, let x denote an input problem expressed in the target (low-resource) language, and let y represent the final answer also in the target language. We denote the chain-of-thought reasoning trace as C , which is constrained to be in English. The rationale behind this design is that models typically have been extensively trained in English for reasoning tasks, and we hypothesize that their latent spaces are more robustly aligned with English reasoning processes. Our method is designed to learn the conditional probability distribution:

$$P_{\theta}(C, y|x) = P_{\theta}(C|x)P_{\theta}(y|x, C) \quad (1)$$

where θ represents the parameters of the language model, $C = (s_1, \dots, s_n)$ is the chain-of-thought, comprising natural language reasoning steps s_i . Following recent works (Zelikman et al. 2022; OpenAI 2024a), we explicitly separate C and y by a special token (e.g. “<think>”) to mark the transition between the reasoning trace and the final answer.

Given a training dataset $D = \{(x, C, y)\}$ consisting of tuples where:

- $x \in S_T$: input in the target low-resource language.

- $C \in S_{en}$: reasoning trace in English.
- $y \in S_T$: final answer in the target low-resource language.

we optimize the following objective during training:

$$\mathcal{L}(\theta) = \sum_{x, C, y \in D} \log P_{\theta}(C, y|x) \quad (2)$$

$$= \sum_{x, C, y \in D} [\log P_{\theta}(C|x) + \log P_{\theta}(y|x, \hat{C})] \quad (3)$$

$$= \sum_{x, C \in D} \log P_{\theta}(C|x) + \sum_{x, y \in D} \log P_{\theta}(y|x, \hat{C}) \quad (4)$$

$$= \alpha \sum_{x, C \in D} \log P_{\theta}(C|x) + \beta \sum_{x, y \in D} \log P_{\theta}(y|x, \hat{C}) \quad (5)$$

where $\log P_{\theta}(C|x)$ optimizes the model to generate English reasoning traces given input in different languages, and $\log P_{\theta}(y|x, \hat{C})$ with $\hat{C} = \theta(x)$ pushes the model to generate final response based on both the input (in target language) and the generated reasoning trace. Additionally, each objective can be balanced with hyperparameter weights α and β . By default, we set both the hyperparameters α and β to 1.0.

Effectively, this training approach allows the model to “reason” in English (through forcing the ground-truth CoT via the loss function $\log P_{\theta}(C|x)$), a language where it has been extensively trained for reasoning. In other words, this enhances the representation alignment between input prompts in different languages, where they will lead to the same traces of reasoning. The training method also allows the model to understand and provide responses in the low-resource target language. This not only boosts performance on low-resource languages but also simplifies the training process by reducing the need for large amount of data in those languages. In Figure 2, we visualize the changes of the two loss components in Eq. (5) over the first training epoch. The figure indicates that our English-Pivoted CoT training makes it easier for the model to maintain reasoning in English, as evidenced by the lower initial reasoning loss and slower subsequent decrease. The increasing gap between reasoning and response losses shows our method effectively separates English reasoning from target-language generation, enabling better understanding of the low-resource language. In contrast, Native CoT starts with higher initial rea-

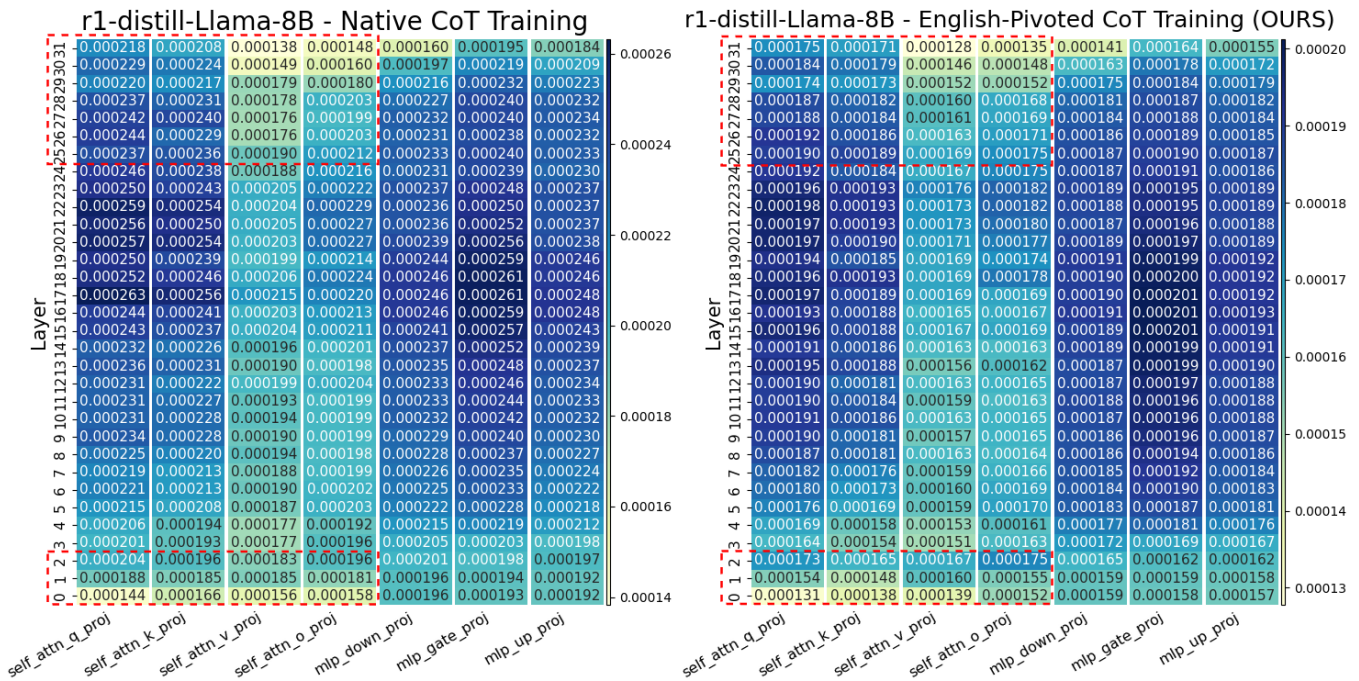


Figure 3: Parameter updates (mean absolute differences) of Native CoT Training (Left) and English-Pivoted CoT Training (right) for r1-distill-Llama-8B. English-Pivoted CoT Training focuses more on language comprehension layers (red boxes).

soning loss that decreases rapidly, resulting in a smaller gap between the two losses, indicating less separation between language understanding and reasoning processes.

This effect can also be seen in Figure 3, comparing the parameter update patterns of Native CoT training (multilingual reasoning fine-tune with both CoT traces and answers in target language) and our proposed English-Pivoted CoT training for r1-distill-Llama-8B. The figure illustrates key differences in adaptation dynamics. Firstly, considering the absolute magnitude of parameter changes, Native CoT introduces larger updates throughout the model (approximately 1.3 times larger), suggesting a more significant departure from the baseline model’s parameters, which can lead to substantially different reasoning behavior. Secondly, examining the relative changes across model layers, English-Pivoted CoT also concentrates updates in the first and last few layers, particularly within attention-related matrices (highlighted in red boxes). These layers have been shown in prior research to be in charge of language understanding and generation tasks (Tran, O’Sullivan, and Nguyen 2024; Wendler et al. 2024). Consequently, our method strategically targets these layers, enabling the model to effectively understand and generate final responses in the low-resource language without deviating from the original model’s reasoning capability.

Experiments

Experiment settings

Baselines. We select r1-distill-Llama-8B (DeepSeek-AI 2025) and DeepHermes-3-Llama-3-8B (Teknium et al. 2025) due to their state-of-the-art performance and open-source availability. We perform fine-tuning on these models

following our proposed approach (denote *English-Pivoted CoT Training*). We compare the performance of our method against several existing multilingual adaptation techniques, including few-shot prompting (Shi et al. 2023), multilingual reasoning fine-tuning (denote *Native CoT Training*) (Anand et al. 2024; Lai and Nissim 2024), model merging (Yadav et al. 2023; Tao et al. 2024) and layer swap (Bandarkar et al. 2025) with Llama-3.1-8B-Instruct (Grattafiori et al. 2024) following the default hyperparameter setup.

Given the limited availability of reasoning data in low-resource languages, following recent works (Tao et al. 2024; Liu et al. 2025a), we use NLLB-3.3B (Costa-Jussà et al. 2022) to translate existing reasoning datasets: Bespoke (Labs 2025) into Irish, pensez (Ha 2025) into French. For Chinese, a high-resource language, we leverage conglu (Liu et al. 2025b), and translate to English for our approach. We sample the amount of data in target language to be $N = 8,000$ samples, and add in 8,000 English samples from the same source to prevent forgetting. In our approach, reasoning traces originally in target languages are translated into English for training purposes. While the amount of data we leverage for adaptation to low-resource languages is relatively small (e.g., compared to 800,000 samples used in (DeepSeek-AI 2025)), our evaluation results illustrate a clear improvement, highlighting its effectiveness. We note that all evaluation datasets are created or manually verified by humans, ensuring accuracy and quality.

Training setup. Our training implementation employs HuggingFace Transformers and DeepSpeed. To manage memory constraints efficiently, we set the maximum input sequence length to 16,384 tokens. Models are trained using

Model	Setting	AIME2024 (en)	AIME2024 (ga)	LC2024 (ga)
r1-distill-Llama-8B	Zero-shot Prompting	43.33	6.67	63.64
r1-distill-Llama-8B	Few-shot Prompting - Irish CoT	46.67	10.00	45.45
r1-distill-Llama-8B	Few-shot Prompting - English CoT	46.67	6.67	32.73
r1-distill-Llama-8B	TIES-merging (Yadav et al. 2023; Tao et al. 2024)	26.67	3.33	34.55
r1-distill-Llama-8B	Layer Swap (Bandarkar et al. 2025)	36.67	16.67	30.91
r1-distill-Llama-8B	Native CoT Training (Anand et al. 2024; Lai and Nissim 2024)	48.33	21.67	37.14
r1-distill-Llama-8B	English-Pivoted CoT Training (OURS)	53.33	35.00	73.33
DeepHermes-3-Llama-3-8B	Zero-shot Prompting	1.67	6.67	52.73
DeepHermes-3-Llama-3-8B	Few-shot Prompting - Irish CoT	6.67	0.00	23.64
DeepHermes-3-Llama-3-8B	Few-shot Prompting - English CoT	6.67	3.33	32.73
DeepHermes-3-Llama-3-8B	TIES-merging (Yadav et al. 2023; Tao et al. 2024)	3.33	0.00	14.55
DeepHermes-3-Llama-3-8B	Layer Swap (Bandarkar et al. 2025)	3.33	3.33	25.45
DeepHermes-3-Llama-3-8B	Native CoT Training (Anand et al. 2024; Lai and Nissim 2024)	1.67	8.33	40.01
DeepHermes-3-Llama-3-8B	English-Pivoted CoT Training (OURS)	5.00	10.00	54.55

Table 1: Performance comparison of our English-Pivoted CoT Training approach against baseline methods across English (en) and Irish (ga) reasoning benchmarks. Bold scores indicate the best performance per benchmark.

Model	Setting	concepts & skills	contexts & applications
r1-distill-Llama-8B	Zero-shot Prompting	84.85	31.82
r1-distill-Llama-8B	Native CoT Training	47.50	31.82
r1-distill-Llama-8B	English-Pivoted CoT Training (OURS)	82.82	59.09
DeepHermes-3-Llama-3-8B	Zero-shot Prompting	57.58	45.45
DeepHermes-3-Llama-3-8B	Native CoT Training	57.58	13.64
DeepHermes-3-Llama-3-8B	English-Pivoted CoT Training (OURS)	57.58	50.00

Table 2: Comparisons of low-resource language understanding on LC2024, where the exam is split into 2 parts: concepts & skills and contexts & applications. The latter part has additional contextual and real-world descriptions, requiring a higher level of Irish language understanding to comprehend the input. Bold scores indicate the best performance per benchmark.

the AdamW optimizer (Loshchilov and Hutter 2019) for 3 epochs, with a maximum learning rate of 1×10^{-5} . Training is distributed across two Nvidia A100 GPUs (80GB each), with a total batch size of 24.

Benchmarks. We evaluate the reasoning capabilities of all models on standard English benchmarks, including AIME2024 (Veeraboina 2023) (challenging Olympiad-level math) and MGSM (Shi et al. 2023) (high school-level math). Additionally, we use MGSM for evaluating performance in French and Chinese (as it is a multilingual benchmark dataset), supplemented by MATH-hard (Mohamad Alhajar 2024) for French, follow the French LLM leaderboard (Mohamad Alhajar 2024). For the Irish language, due to lack of existing datasets, we contribute two new ones:

- Irish version of AIME2024: translated and verified by two Irish speakers.
- Leaving Certificate 2024 Math exam - Higher Level (LC2024): derived from official Irish Leaving Certificate (Wikipedia 2025) examination with 2 splits: con-

cepts&skills and contexts&applications/ We extract individual questions that have concrete answers (excluding proof-based questions) and do not require modalities beyond text (e.g., geometric diagrams). The questions are formatted in LaTeX with a total of 55 samples.

We adopt each benchmark’s default hyperparameters, leveraging a decoding temperature of 0.6, top-k sampling with $k = 0.95$, limiting generation outputs to 16,384 tokens.

Results and analysis

Reasoning capability can be effectively transferred to low-resource languages. Table 1 presents the performance of models trained with our approach (English-Pivoted CoT Training) compared to baselines on English and Irish benchmarks. Our method demonstrates a strong improvement: on the Irish version of AIME2024, while all baselines perform poorly, our approach obtains a clear gap of upto 28.33% on r1-distill-Llama-8B baseline. More specifically, we achieve accuracies of 35.00% and 10.00% when fine-




Figure 4: Representation retrieval accuracy between *Left*: questions, and *Right*: questions and generated CoT traces of the same questions in different languages.

Model	Setting	AIME2024 (en)	MSGM (en)	MATH- hard (fr)	MSGM (fr)
r1-distill-Llama-8B	Zero-shot Prompting	43.33	79.6	49.74	54.8
r1-distill-Llama-8B	Native CoT Training	45.00	82.0	31.90	61.2
r1-distill-Llama-8B	English-Pivoted CoT Training (OURS)	50.00	89.6	70.01	83.2
DeepHermes-3-Llama-3-8B	Zero-shot Prompting	1.67	85.2	3.69	21.6
DeepHermes-3-Llama-3-8B	Native CoT Training	3.33	70.8	6.26	44.8
DeepHermes-3-Llama-3-8B	English-Pivoted CoT Training (OURS)	8.33	76.8	32.27	77.2

Table 3: Ablation study of fine-tuning on French - a medium resource language, and benchmark across English and French reasoning datasets. Bold scores indicate the best performance per benchmark.

tuned on r1-distill-Llama-8B and DeepHermes-3-Llama-3-8B, respectively, matching performance on the English version. On LC2024, which is relatively less challenging than AIME2024, although r1-distill-Llama-8B achieves an acceptable accuracy of 63.64%, our approach surpasses it by a 10% margin (73.33%). Further inspection shows that both baseline models generate all CoTs and final answers in English, regardless of prompting in which languages. This default behavior aligns with the English prompting approaches in (Qin et al. 2023; Shi et al. 2023). Interestingly, performances on LC2024 (authored by Irish speakers) drops compared to the translated Irish AIME2024 for all baseline approaches. We hypothesize the model relies on familiar math notation (same across all languages) to infer the task and then applies English reasoning, whereas interventions such as prompting, CoT fine-tuning or model merging disrupt this implicit strategy, causing parameter interference in reasoning in an extremely low-resource language.

Improved language understanding capability. To investigate whether our approach also enhances language understanding or simply overfits to mathematical phrases, we split LC2024 into its two original sections: *concepts & skills* and *contexts & applications* - with more contextual and real-world application descriptions, requires deeper Irish language understanding. Due to space limitation, we

present the results of our method and 2 strongest baselines, with more results in the Appendix. As shown in Table 2, Native CoT training reduces the model’s performance for DeepHermes-3-Llama-3-8B from 45.45 to 13.64 on *concepts & skills*, while maintaining the same performance on r1-distill-Llama-8B. This shows that training to force the model to reason in low-resource languages affects its ability to generate effective CoTs. In contrast, our approach yields substantially gain in *contexts & applications*, indicating that not only does it enable effective reasoning in low-resource languages but also improves language understanding.

Diverged reasoning traces when forced to reason in different languages. We assess cross-language alignment on AIME2024 by retrieving English embeddings for Irish questions (with and without CoTs) using average token representations across layers. As illustrated in Figure 4, target-language-only models show near-perfect question-only alignment (up to 100%), but this falls to around 35% once CoTs are included, revealing divergent internal representations. Figure 1 illustrates how minor embedding shifts (e.g., confusing “independence” and “dependence”) can alter interpretation. In contrast, English-Pivoted CoT Training maintains almost 100% alignment for both questions and reasoning traces, demonstrating stable, language-agnostic embeddings.

Training Setting	Prompting Setting	AIME2024	LC2024 - concepts&skills	LC2024 - contexts&applications
Mixed CoT Lang.	Irish CoT	33.33	3.33	54.55
Mixed CoT Lang.	English CoT	33.33	30.00	61.82
Two-stage	Irish CoT	40.00	6.67	45.45
Two-stage	English CoT	33.33	33.33	58.18

Table 4: Irish task performance by CoT language for r1-distill-Llama-8B trained on English and Irish CoTs.

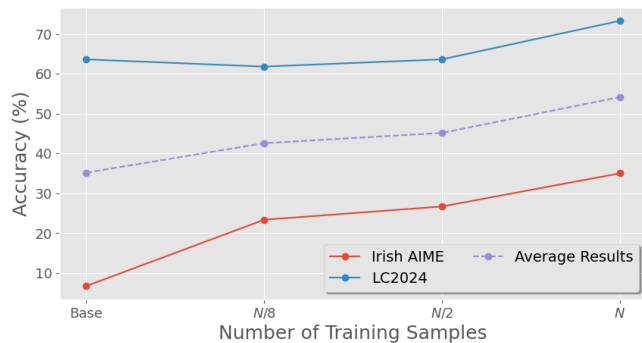


Figure 5: Performance of English-Pivoted CoT training on Irish AIME and LC2024 with varying number of training samples on r1-distill-Llama-8B (default is N).

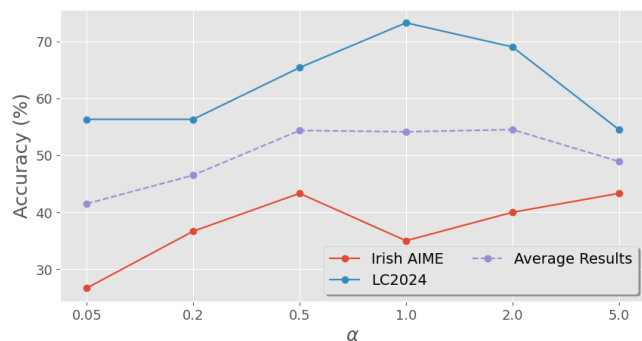


Figure 6: Performance of English-Pivoted CoT training with different α (default: 1.0) during training ($\beta = 1.0$) on r1-distill-Llama-8B. x-axis is evenly spaced for visualization.

Additionally, we experience with **controlling the CoT language at test time**. We introduce 2 additional training settings: 1) Mixed CoT Language: 50% of CoTs in Irish, 50% in English, 2) Two-Stage: first stage uses English-Pivoted CoT to align comprehension, then a second stage fine-tunes on native Irish CoTs to refine reasoning in the target language. We prepend each rationale with “Let’s think in X language” and fix this prompt at inference to control the CoT language. Table 4 reports that on Irish AIME2024 (olympiad-level), all methods perform similarly, but English-CoT prompted models clearly outperform on both LC2024 subsets. This suggests that in data-limited scenarios, directing model to think in English remains the most effective strategy for robust multilingual reasoning.




Figure 7: Changes in correction rate from r1-distill-Llama-8B when finetuned with our proposed approach across 3 data resources, low (Irish), medium (French), and high (Chinese).

Sample efficiency of English-Pivoted CoT Training. By decoupling reasoning from language learning, our method needs only comprehension fine-tuning. As Figure 5 shows, with just 1,000 samples ($N/8$) it beats the base model by over 5%, and accuracy climbs further with more data. This highlights its strong performance in low-resource settings.

Learning dynamic of English-Pivoted CoT Training. Varying the CoT-loss weight α (with $\beta = 1$) in Eq. 5 (Figure 6) reveals that both reducing and over-emphasizing α degrade performance. This highlights the need to balance the English-CoT and target-response objectives. These results underscore the need to balance the reasoning trace and final response objectives in English-Pivoted CoT Training.

Generalizability to medium- and high-resource languages. Experiments on Chinese and French - medium to high-resource languages (Zeman et al. 2024; Chang et al. 2024), present in Table 3 and Table 5, indicate the generalizability: on French MGSM, r1-distill-Llama-8B scores 54.8%, while our approach boosts it to 83.2%; on Chinese, English pivoting yields only marginal gains over baseline 65.2%, likely because the model naturally reasons in Chinese, so forcing English reasoning can cause interference. Furthermore, the percentage of the new model corrects an incorrect answer by the base model and vice versa are shown in Figure 7. We can see that on MGSM (zh), the correct → incorrect rate is the highest, reflecting these conflicts. Nevertheless, our method still achieves an average improvement of 14.30% across the 3 resource regimes (low, medium, and high-resources), demonstrating its broad applicability.

Model	Setting	AIME2024 (en)	MSGM (en)	MSGM (zh)
r1-distill-Llama-8B	Zero-shot Prompting	43.33	79.6	65.2
r1-distill-Llama-8B	Native CoT Training	46.33	92.4	81.6
r1-distill-Llama-8B	English-Pivoted CoT Training (OURS)	56.67	92.4	70.0
DeepHermes-3-Llama-3-8B	Zero-shot Prompting	1.67	85.2	50.8
DeepHermes-3-Llama-3-8B	Native CoT Training	11.67	83.2	61.6
DeepHermes-3-Llama-3-8B	English-Pivoted CoT Training (OURS)	13.33	89.6	62.4

Table 5: Ablation study of fine-tuning on Chinese - a high resource language, where baselines have high level of understanding and tend to output CoT traces in Chinese. Bold scores indicate the best performance per benchmark.

Conclusion

In this work, we propose a novel approach named English-Pivoted CoT Training to effectively transfer the reasoning capabilities of LLMs to low-resource languages. Through training to explicitly aligning reasoning processes across languages by forcing CoT traces to be in the dominant language, our approach achieves significant performance gains (up to 28.33%) over existing techniques, and is generalizable to other resource regimes (medium and high-resource languages). Furthermore, our analysis provides insights into multilingual LLM behavior, particularly the benefit of separating language understanding from reasoning. By contributing the first-ever Irish mathematical reasoning benchmark (LC2024), we also aim to support future research in multilingual reasoning. Future work can be interested in extending to other low-resource languages and tasks, and developing strategies to further enhance cross-lingual reasoning.

Ethics Statement

Our work advances language technologies for low-resource and endangered languages, with a particular focus on Irish. By leveraging English-Pivoted CoT Training, we aim to enhance the accessibility and usability of these languages, thereby promoting linguistic diversity and preserving cultural heritage.

All training and evaluation data are drawn from publicly available sources. Our newly released LC2024 dataset is based on Regulations on Re-use of Public Sector Information 2005, SI 279 of 2005, and is licensed for non-commercial research and educational use only.

Acknowledgments

We would like to acknowledge CloudCIX Limited for the generous collaborative support of computing resources on their NVIDIA HGX/H100 GPU cluster. This publication has emanated from research supported in part by grants from Research Ireland under Grant [12-RC-2289-P2] and [18/CRT/6223] which is co-funded under the European Regional Development Fund. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

- Anand, A.; Prasad, K.; Kirtani, C.; Nair, A. R.; Nema, M. K.; Jaiswal, R.; and Shah, R. R. 2024. Multilingual Mathematical Reasoning: Advancing Open-Source LLMs in Hindi and English. *arXiv:2412.18415*.
- Bandarkar, L.; Muller, B.; Yuvraj, P.; Hou, R.; Singhal, N.; Lv, H.; and Liu, B. 2025. Layer Swapping for Zero-Shot Cross-Lingual Transfer in Large Language Models. In *The Thirteenth International Conference on Learning Representations*.
- Chang, T. A.; Arnett, C.; Tu, Z.; and Bergen, B. 2024. When Is Multilinguality a Curse? Language Modeling for 250 High- and Low-Resource Languages. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 4074–4096. Miami, Florida, USA: Association for Computational Linguistics.
- Common Crawl. 2025. Common Crawl—Open Repository of Web Crawl Data. Accessed: 2025-07-29.
- Costa-Jussà, M. R.; Cross, J.; Çelebi, O.; Elbayad, M.; Heafield, K.; Heffernan, K.; Kalbassi, E.; Lam, J.; Licht, D.; Maillard, J.; et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv:2501.12948*.
- Ghosh, A.; Datta, D.; Saha, S.; and Agarwal, C. 2025. The Multilingual Mind : A Survey of Multilingual Reasoning in Language Models. *arXiv:2502.09457*.
- Grattafiori, A.; et al. 2024. The Llama 3 Herd of Models. *arXiv:2407.21783*.
- Guo, P.; Hu, Y.; Ren, Y.; Li, Y.; Zhang, J.; and Zhang, X. 2023. Mitigating Long-Tail Language Representation Collapsing via Cross-Lingual Bootstrapped Unsupervised Fine-Tuning. In *ECAI 2023 - 26th European Conference on Artificial Intelligence*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, 940–947. IOS Press.
- Ha, H. H. 2025. Pensez: Less Data, Better Reasoning – Rethinking French LLM. *arXiv:2503.13661*.
- Huang, Z.; Zhu, W.; Cheng, G.; Li, L.; and Yuan, F. 2024. MindMerger: Efficiently Boosting LLM Reasoning in non-English Languages. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

- Ji, Y.; Li, J.; Ye, H.; Wu, K.; Xu, J.; Mo, L.; and Zhang, M. 2025. Test-time Computing: from System-1 Thinking to System-2 Thinking. *arXiv preprint arXiv:2501.02497*.
- Kim, J.; Koo, S.; and Lim, H. 2024. *Revisiting Under-Represented Knowledge of Latin American Literature in Large Language Models*. IOS Press. ISBN 9781643685489.
- Labs, B. 2025. Bespoke-Stratos: The unreasonable effectiveness of reasoning distillation. <https://www.bespokelabs.ai/blog/bespoke-stratos-the-unreasonable-effectiveness-of-reasoning-distillation>. Accessed: 2025-01-22.
- Lai, H.; and Nissim, M. 2024. mCoT: Multilingual Instruction Tuning for Reasoning Consistency in Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12012–12026. Bangkok, Thailand: Association for Computational Linguistics.
- Li, C.; Wang, S.; Zhang, J.; and Zong, C. 2024a. Align after Pre-train: Improving Multilingual Generative Models with Cross-lingual Alignment.
- Li, C.; Wang, S.; Zhang, J.; and Zong, C. 2024b. Improving In-context Learning of Multilingual Generative Language Models with Cross-lingual Alignment. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8058–8076. Mexico City, Mexico: Association for Computational Linguistics.
- Liu, C.; Zhang, W.; Zhao, Y.; Luu, A. T.; and Bing, L. 2025a. Is Translation All You Need? A Study on Solving Multilingual Tasks with Large Language Models. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 9594–9614. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-189-6.
- Liu, C.; et al. 2025b. The Chinese dataset distilled from DeepSeek-R1-671b. <https://huggingface.co/datasets/Congliu/Chinese-DeepSeek-R1-Distill-data-110k>.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Mohamad Alhajar, A. L. 2024. Open LLM French Leaderboard v0.2. <https://huggingface.co/spaces/le-leadboard/OpenLLMFrenchLeaderboard>.
- OpenAI. 2024a. Learning to reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>. [Accessed 19-03-2025].
- OpenAI. 2024b. OpenAI o1 System Card. [arXiv:2412.16720](https://arxiv.org/abs/2412.16720).
- Qin, L.; Chen, Q.; Wei, F.; Huang, S.; and Che, W. 2023. Cross-lingual Prompting: Improving Zero-shot Chain-of-Thought Reasoning across Languages. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2695–2709. Singapore: Association for Computational Linguistics.
- Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; and Yu, P. S. 2024. Multilingual large language model: A survey of resources, taxonomy and frontiers. *arXiv preprint arXiv:2404.04925*.
- Shi, F.; Suzgun, M.; Freitag, M.; Wang, X.; Srivats, S.; Vosoughi, S.; Chung, H. W.; Tay, Y.; Ruder, S.; Zhou, D.; Das, D.; and Wei, J. 2023. Language models are multilingual chain-of-thought reasoners. In *The Eleventh International Conference on Learning Representations*.
- Tao, M.; Zhang, C.; Huang, Q.; Ma, T.; Huang, S.; Zhao, D.; and Feng, Y. 2024. Unlocking the Potential of Model Merging for Low-Resource Languages. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 8705–8720. Miami, Florida, USA: Association for Computational Linguistics.
- Teknum; Jin, R.; Guang, C.; Suphavadeepravit, J.; and Quesnelle, J. 2025. DeepHermes 3 Preview.
- Tran, K.-T.; O’Sullivan, B.; and Nguyen, H. 2024. Irish-based Large Language Model with Extreme Low-Resource Settings in Machine Translation. In Ojha, A. K.; Liu, C.-h.; Vylomova, E.; Pirinen, F.; Abbott, J.; Washington, J.; Oco, N.; Malykh, V.; Logacheva, V.; and Zhao, X., eds., *Proceedings of the Seventh Workshop on Technologies for Machine Translation of Low-Resource Languages (LoResMT 2024)*, 193–202. Bangkok, Thailand: Association for Computational Linguistics.
- UNESCO, ed. 2010. *Atlas of the world’s languages in danger*. Memory of Peoples Series. UNESCO, 3 edition.
- Veeraboina, H. 2023. AIME Problem Set 1983-2024.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wendler, C.; Veselovsky, V.; Monea, G.; and West, R. 2024. Do Llamas Work in English? On the Latent Language of Multilingual Transformers. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15366–15394. Bangkok, Thailand: Association for Computational Linguistics.
- Wikipedia. 2025. Leaving Certificate (Ireland). [https://en.wikipedia.org/wiki/Leaving_Certificate_\(Ireland\)](https://en.wikipedia.org/wiki/Leaving_Certificate_(Ireland)). [Online; accessed 25 July 2025].
- Yadav, P.; Tam, D.; Choshen, L.; Raffel, C.; and Bansal, M. 2023. TIES-Merging: Resolving Interference When Merging Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Zelikman, E.; Wu, Y.; Mu, J.; and Goodman, N. 2022. STaR: Bootstrapping Reasoning With Reasoning. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.

Zeman, D.; et al. 2024. Universal Dependencies 2.15. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Zhang, F.; Jin, K.; and Zhuo, H. H. 2024. *Planning with Logical Graph-Based Language Model for Instruction Generation*. IOS Press. ISBN 9781643685489.

Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; Du, Y.; Yang, C.; Chen, Y.; Chen, Z.; Jiang, J.; Ren, R.; Li, Y.; Tang, X.; Liu, Z.; Liu, P.; Nie, J.-Y.; and Wen, J.-R. 2025. A Survey of Large Language Models. arXiv:2303.18223.

Zhu, W.; Huang, S.; Yuan, F.; She, S.; Chen, J.; and Birch, A. 2024. Question Translation Training for Better Multilingual Reasoning. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 8411–8423. Bangkok, Thailand: Association for Computational Linguistics.