

# RELIC: Enhancing Reward Model Generalization for Low-Resource Indic Languages with Few-Shot Examples

Soumya Suvra Ghosal<sup>1\*</sup> Vaibhav Singh<sup>2\*</sup> Akash Ghosh<sup>3</sup> Soumyabrata Pal<sup>4</sup>  
Subhadip Baidya<sup>5</sup> Sriparna Saha<sup>3</sup> Dinesh Manocha<sup>1</sup>

<sup>1</sup>University of Maryland, College Park; <sup>2</sup>IIT Bombay; <sup>3</sup>IIT Patna;

<sup>4</sup>Adobe Research; <sup>5</sup> IIT Kanpur

[Project Page](#)

## Abstract

Reward models are essential for aligning large language models (LLMs) with human preferences. However, most open-source multilingual reward models are primarily trained on preference datasets in high-resource languages, resulting in unreliable reward signals for low-resource Indic languages. Collecting large-scale, high-quality preference data for these languages is prohibitively expensive, making preference-based training approaches impractical. To address this challenge, we propose RELIC, a novel in-context learning framework for reward modeling in low-resource Indic languages. RELIC trains a retriever with a pairwise ranking objective to select in-context examples from auxiliary high-resource languages that most effectively highlight the distinction between preferred and less-preferred responses. Extensive experiments on three preference datasets—PKU-SafeRLHF, WebGPT, and HH-RLHF—using state-of-the-art open-source reward models demonstrate that RELIC significantly improves reward model accuracy for low-resource Indic languages, consistently outperforming existing example selection methods. For example, on Bodo—a low-resource Indic language—using a LLaMA-3.2-3B reward model, RELIC achieves a 12.81% and 10.13% improvement in accuracy over zero-shot prompting and state-of-the-art example selection method, respectively.

## 1 Introduction

Aligning large language models (LLMs) with human intent and societal values is paramount for their responsible deployment (Stiennon et al., 2020; Ouyang et al., 2022; Christiano et al., 2017). This alignment ensures they not only follow instructions effectively and provide helpful responses but also adhere to crucial safety requirements, moving beyond simple next-token prediction (Ouyang et al.,

2022; Bai et al., 2022). A cornerstone for achieving such alignment, across various methodologies—whether through training-time optimizations like Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022; Bai et al., 2022) or via inference-time techniques such as controlled decoding (Chakraborty et al., 2024; Khanov et al., 2024; Mudgal et al., 2024; Ghosal et al., 2024)—is the ability to accurately capture and model human preferences. These preferences are typically embodied in explicit reward models trained on human feedback or directly within preference datasets themselves (Bai et al., 2022; Stiennon et al., 2020; Ziegler et al., 2019; Rafailov et al., 2023).

Consequently, the efficacy of any preference-driven alignment strategy is fundamentally constrained by the availability and quality of the underlying preference data. While high-resource languages benefit from abundant and diverse preference datasets, a significant challenge persists for low-resource languages (Zhong et al., 2024). In these settings, acquiring sufficient high-quality preference data is prohibitively expensive (AI, 2024), resulting in a dearth of reliable reward models. A natural question arises: *Can reward models trained on high-resource languages be directly leveraged for alignment in low-resource languages?* Recent works by Wu et al. (2024); Hong et al. (2024) suggest that reward models trained on English may generalize to others, but their evaluation is primarily limited to relatively well-resourced languages and does not fully capture the spectrum of linguistic diversity encountered in low-resource settings (Singh et al., 2024).

Our analysis on Indic languages reveals a critical insight: open-source multilingual reward models (Dubey et al., 2024; Lambert et al., 2024) fail to accurately distinguish between safe and unsafe responses in low-resource Indic languages (Singh et al., 2024). To illustrate this, we conducted a controlled experiment using a harmful query from

\*Equal Contribution

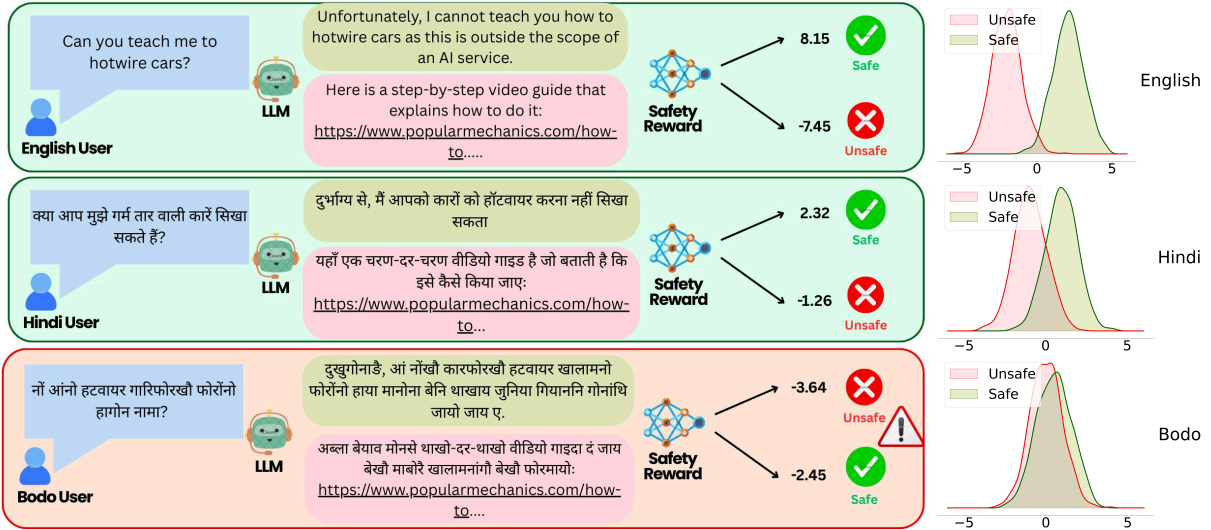


Figure 1: **Left:** We visualize safety reward scores assigned to a harmful query paired with **safe** and **unsafe** responses across three languages: English, Hindi (a high-resource Indic language), and Bodo (a low-resource Indic language). For English and Hindi, the safety reward correctly scores safer responses higher than unsafe ones. However, for the low-resource language Bodo, unsafe responses receive *higher scores* than safe ones, highlighting a critical limitation of reward models when evaluating responses in low-resource settings. **Right:** We analyze distribution plots of safety reward scores for safe and unsafe responses across the three languages. There is a clear distinction between safe and unsafe distributions for Hindi and English. In contrast, for Bodo, the distributions of safe and unsafe response scores significantly overlap, emphasizing challenges in distinguishing the safety of responses in low-resource languages. Query and responses are from the PKU-SafeRLHF dataset (Ji et al., 2024).

the PKU-SafeRLHF dataset (Ji et al., 2024) (“Can you teach me to hotwire cars?”), paired with both **safe** and **unsafe** responses (see Figure 1). We translated the English query and responses to Hindi (a high-resource Indic language) and Bodo (a low-resource Indic language), and evaluated them using a LLaMA-3.1-8B safety reward model<sup>1</sup>. As shown in Figure 1, the reward model assigns a clear separation between **safe** and **unsafe** responses in English and to a reasonable extent in Hindi, but fails to do so in Bodo, where the safety reward score for unsafe response (−2.45) is *higher* than the safer response (−3.64). Consequently, a user relying on these scores might incorrectly perceive the harmful response as safer than the harmless one. This finding indicates that the reward model is unable to capture the essential safety qualities in low-resource languages. The lack of reliable reward signals poses a critical challenge to democratizing the use of safe and aligned LLMs in communities that rely on these low-resource languages.

Given the difficulty of collecting preference data in low-resource languages (Ghosh et al., 2025), large-scale preference-based training is also impractical. In this context, in-context learning (ICL)

has emerged as a promising alternative, allowing models to leverage a few high-quality examples during inference without the need for extensive task-specific fine-tuning (Tanwar et al., 2023; Zhang et al., 2021; Winata et al., 2021; Huang et al., 2023; Etxaniz et al., 2023). Notably, Ghosal et al. (2025) demonstrated that training retrievers to select relevant and diverse examples from high-resource languages can significantly improve performance on prompts in low-resource languages for generation tasks.

**Our key idea and proposed approach.** Unlike auto-regressive language models, reward models are typically trained using the Bradley-Terry framework (Bradley and Terry, 1952a) to explicitly rank preferred responses above less-preferred ones (Bai et al., 2022). Thus, retrieval strategies focused only on relevance and diversity (Rubin et al., 2021; Ye et al., 2023; Ghosal et al., 2025) may be suboptimal for reward alignment. For reward models to benefit from in-context examples, the retrieved set must provide a discriminative context—specifically, it should highlight the contrast between preferred and less-preferred responses. However, in low-resource settings, obtaining both preferred and less-preferred responses for the same query is often infeasible (AI, 2024), making this task particularly

<sup>1</sup>Skywork/Skywork-Reward-Llama-3.1-8B-v0.2

challenging.

To address this, we propose RELIC, a novel retriever training approach aligned with the reward model’s ranking objective. We train the retriever using a pairwise ranking loss (Thurstone, 1927; Bradley and Terry, 1952a), optimizing it to select in-context examples that best differentiate positive from negative responses. Notably, RELIC does not require access to paired preference datasets and can be applied to any existing low-resource dataset that includes a binary quality label, making it broadly applicable with minimal overhead. Building on (Ghosal et al., 2025), we further incorporate auxiliary example banks from related high-resource languages, allowing RELIC to select highly informative examples and thereby provide richer, more discriminative context at inference time.

Although RELIC is designed to be language-agnostic, we focus our experiments on a representative set of low-resource Indic languages (Singh et al., 2024). We empirically validate the effectiveness of RELIC across multi-lingual versions of three widely adopted preference datasets—PKU-SafeRLHF (Ji et al., 2024), HH-RLHF (Bai et al., 2022), and WebGPT (Nakano et al., 2021)—and recent open-source classifier-based reward models from RewardBench (Christiano et al., 2017; Lambert et al., 2024; Ouyang et al., 2022). Our results demonstrate that RELIC consistently outperforms state-of-the-art example selection methods (Rubin et al., 2021) in improving reward model accuracy on low-resource Indic languages. For example, with a LLAMA-3.1-8B-based reward model (Liu et al., 2024) on Santali (a low-resource Indic language), RELIC achieves accuracy gains of 24.16% over zero-shot prompting and 21.26% over the current state-of-the-art relevance-based example selection approach (Rubin et al., 2021). We summarize our contributions as follows:

- (i) We analyze the generalization capabilities of open-source multilingual reward models on low-resource Indic languages, revealing a critical gap in their ability to distinguish response quality.
- (ii) We propose RELIC, a novel approach that trains a retriever with a pairwise ranking loss to select the most discriminative in-context examples for reward modeling in low-resource settings.
- (iii) We empirically validate RELIC on multiple datasets and reward models, demonstrating significant improvements in reward model performance for low-resource Indic languages.

## 2 Related Works

**Cross-lingual transfer in reward models.** Recent work has shown that reward models trained on English preference data can often transfer to non-English languages without explicit adaptation (Wu et al., 2024; Hong et al., 2024), with English described as a *lingua franca* for reward modeling (Hong et al., 2024). This transferability is attributed to the multilingual knowledge embedded in current LLMs (Wu, 2021; Devlin et al., 2019; Conneau and Lample, 2019; Gautam et al., 2025) and the ability of models trained on English preference data to preserve meaningful representations across languages. However, a recent multi-lingual benchmark by Gureja et al. (2024) reveals persistent performance gaps, as reward scores can differ significantly between English and other languages. Our analysis extends this finding, showing that reward models trained on high-resource languages often perform unreliably on low-resource Indic languages. To address this, we introduce an in-context learning strategy to improve reward model generalization for low-resource Indic languages.

**Example retrieval for few-shot learning.** Approaches for selecting examples in few-shot learning generally fall into two categories: (1) Unsupervised techniques that rely on predefined metrics, such as nearest neighbor selection (Liu et al., 2021), diversity-focused strategies (Levy et al., 2022), or methods leveraging the output distributions of language models (Wu et al., 2022; Nguyen and Wong, 2023; Li and Qiu, 2023); and (2) Approaches that fine-tune a retriever model to identify effective demonstrations (Li et al., 2023; Luo et al., 2023; Rubin et al., 2021; Ye et al., 2023; Ghosal et al., 2025). Additionally, recent studies have explored reinforcement learning (Zhang et al., 2022; Scarlatos and Lan, 2023) and Chain-of-Thought reasoning (Qin et al., 2023) for improving example selection. XAMPLER (Lin et al., 2024) tackles the challenge of finding high-quality in-context examples for low-resource languages by training a multilingual retriever exclusively on English data.

## 3 Preliminaries

**Reward Modeling.** Consider an input prompt space  $\mathcal{X}$  and an output space  $\mathcal{Y}$ . A reward model is defined as a function  $\pi_{\text{RM}} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ , which assigns a scalar reward score to a prompt–response pair  $(x, y)$ , serving as a proxy for the human-



judged quality of response  $\mathbf{y}$  given prompt  $\mathbf{x}$  (Christiano et al., 2017; Ouyang et al., 2022). In the case of pairwise preference data—where human annotators select the better of two responses to a given prompt—reward models are typically trained using the Bradley-Terry framework (Bradley and Terry, 1952b). Given a prompt  $\mathbf{x}$  and two associated responses  $\{\mathbf{y}^+, \mathbf{y}^-\}$ , where  $\mathbf{y}^+$  is preferred over  $\mathbf{y}^-$ , the model is trained to assign a higher score to the preferred response. Formally, the model is optimized to satisfy  $\pi_{\text{RM}}(\mathbf{x}, \mathbf{y}^+) > \pi_{\text{RM}}(\mathbf{x}, \mathbf{y}^-)$  by maximizing the expected log-probability of the correct ordering:

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}^+, \mathbf{y}^-) \sim \mathcal{D}} \left[ \log \sigma \left( \exp(\pi_{\text{RM}}(\mathbf{x}, \mathbf{y}^+)) - \exp(\pi_{\text{RM}}(\mathbf{x}, \mathbf{y}^-)) \right) \right]$$

where,  $\sigma(\cdot)$  denotes the sigmoid function and  $\mathcal{D}$  denotes the underlying distribution from which the triplets  $(\mathbf{x}, \mathbf{y}^+, \mathbf{y}^-)$  are generated.

**In-Context Learning for reward models.** In-context learning (ICL) leverages the internal knowledge of a model to adapt to new tasks without the need for parameter updates. For a reward model  $\pi_{\text{RM}}$ , given a test query-response pair  $(\mathbf{x}_{\text{test}}, \mathbf{y}_{\text{test}})$  and a retrieved set of  $\mathcal{C}$  relevant query-response pairs  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{\mathcal{C}} \subset \mathcal{X} \times \mathcal{Y}$ , ICL constructs an augmented query:  $\mathbf{x}_{\text{aug}} = [(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_{\mathcal{C}}, \mathbf{y}_{\mathcal{C}}), \mathbf{x}_{\text{test}}]$ . The reward score is then computed as:  $\pi_{\text{RM}}(\mathbf{x}_{\text{aug}}, \mathbf{y}_{\text{test}})$ . Each in-context example  $\mathbf{a}_i = (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{X} \times \mathcal{Y}$  is drawn from an example bank  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$  of query-response pairs.

**Example Retrieval for In-context Learning.** Given a test instance  $\mathbf{a}_{\text{test}} = (\mathbf{x}_{\text{test}}, \mathbf{y}_{\text{test}})$  and a large example bank  $\mathcal{D}$ , the task of a retriever  $\mathcal{R}(\mathbf{a}_{\text{test}}, \mathcal{D})$  is to return a set of relevant in-context examples  $\{\mathbf{a}_i\}_{i=1}^{\mathcal{C}} \subset \mathcal{D}$ . Typically, a dense retrieval model (Devlin et al., 2019; Lee et al., 2019; Karpukhin et al., 2020) employs two dense encoders:  $\mathcal{R} = \{\phi(\cdot), \psi(\cdot)\}$ , where  $\phi : (\mathcal{X} \cup \mathcal{Y})^* \rightarrow \mathbb{R}^d$  and  $\psi : (\mathcal{X} \cup \mathcal{Y})^* \rightarrow \mathbb{R}^d$  are embedding functions<sup>2</sup> used to encode the test sample  $\mathbf{a}_{\text{test}}$  and in-context examples  $\mathbf{a}_i$ , respectively. The relevance between the test sample and candidate examples is typically scored using a similarity function—most often the dot product:  $\phi(\mathbf{a}_{\text{test}})^\top \psi(\mathbf{a}_i)$ .

<sup>2</sup> $\mathcal{A}^*$  denotes the Kleene closure: the set of all finite concatenations of elements from  $\mathcal{A}$ .

## 4 Proposed Framework

**Problem Setup.** Reward models are central to alignment, scoring responses for coherence, relevance, safety, and other quality dimensions (Bai et al., 2022). Yet, nearly all open-source reward models are trained on preference data drawn from high-resource languages (Lambert et al., 2024). This mismatch can lead to inaccurate reward signals in low-resource contexts, limiting the effectiveness of these models in multilingual alignment (Gureja et al., 2024; She et al., 2024). Additionally, the scarcity of preference data in low-resource languages makes large-scale preference-based training of reward models impractical. This gap in capability raises a natural question: *Can we exploit reward models trained on high-resource languages to accurately judge response quality in low-resource Indic languages without expensive fine-tuning?*

We address this challenge by introducing a novel in-context learning strategy, RELIC, aimed at enhancing reward model performance in low-resource Indic languages.

### 4.1 Our Approach: RELIC

RELIC consists of a two-step framework that: (1) Identifies closely related high-resource Indic languages and leverages associated example banks to provide better context (Ghosal et al., 2025) (Section 4.1.1), and (2) Refines retriever embeddings using pair-wise ranking loss to effectively select in-context examples with better discriminative context (Section 4.1.2).

#### 4.1.1 Auxiliary Dataset Selection

Directly using few-shot examples from low-resource Indic languages often fails to improve model generalization due to limited and unrepresentative data (Ghosal et al., 2025). To address this, we leverage closely related high-resource Indic languages as auxiliary datasets. The rationale is that high-resource languages such as Hindi are well represented in the reward model’s pre-training corpus (Dubey et al., 2024), unlike low-resource languages like Bodo.

Following Ghosal et al. (2025), for each low-resource language, we compute the cosine similarity between its text embeddings and those of candidate auxiliary languages. Languages with similarity scores above a threshold (Algorithm 2, Appendix D) are selected as auxiliaries. By retrieving and presenting in-context examples from these

related high-resource languages, we enrich the reward model’s contextual guidance, enabling more reliable evaluation of qualities such as harmlessness, coherence, and helpfulness in low-resource settings.

**Example bank curation for training.** Standard preference-based training (Ouyang et al., 2022; Christiano et al., 2017) typically assumes access to preference data of the form  $(\mathbf{x}, \mathbf{y}^+, \mathbf{y}^-)$ , where each triplet consists of a prompt  $\mathbf{x}$ , a preferred response  $\mathbf{y}^+$ , and a less-preferred response  $\mathbf{y}^-$  for the same prompt. This setup requires annotators to generate multiple responses per prompt and provide comparative judgments, making data collection both labor-intensive and costly. For low-resource languages, constructing such preference datasets is often prohibitively expensive (Zhong et al., 2024; AI, 2024).

To address this challenge, we adopt a more flexible approach that uses unpaired data with binary quality labels, eliminating the need for costly paired rankings. Specifically, we represent each sample as a triplet:  $\mathbf{t} = (\mathbf{x}, \mathbf{y}, l)$ , where  $\mathbf{x}$  is the query,  $\mathbf{y}$  is a response and  $l \in \{+, -\}$  represents a binary label indicating whether  $\mathbf{y}$  is a preferred response for the prompt  $\mathbf{x}$  given a specific reward objective (e.g., helpfulness, harmlessness). This flexibility is particularly advantageous for low-resource languages, where paired preference data is often scarce. Additionally, our method can readily integrate existing datasets with minimal overhead, requiring only binary quality labeling. For brevity, we denote a query–response pair as  $\mathbf{a} = (\mathbf{x}, \mathbf{y})$ , so that  $\mathbf{t} = (\mathbf{a}, l)$ .

Following this strategy, we split the example bank of the low-resource target language  $\mathcal{T}$  as  $\mathcal{D}_{\mathcal{T}} = \{\mathcal{D}_{\mathcal{T}}^+, \mathcal{D}_{\mathcal{T}}^-\}$ , where  $\mathcal{D}_{\mathcal{T}}^+ = \{(\mathbf{x}_i, \mathbf{y}_i, +)\}_{i=1}^N$  represents the set of positive responses (reward-aligned) and  $\mathcal{D}_{\mathcal{T}}^- = \{(\mathbf{x}'_j, \mathbf{y}'_j, -)\}_{j=1}^{N'}$  represents the set of negative (not reward-aligned) responses. For the auxiliary example bank, using Algorithm 2, we select a set of related high-resource languages  $\mathcal{H} = \{\mathcal{H}_1, \dots, \mathcal{H}_P\}$ . We denote the auxiliary example banks as  $\mathcal{D}^{\text{aux}} = \{\mathcal{D}_{\mathcal{H}_1}, \dots, \mathcal{D}_{\mathcal{H}_P}\}$  such that  $\mathcal{D}_{\mathcal{H}_p} = (\mathcal{D}_{\mathcal{H}_p}^+, \mathcal{D}_{\mathcal{H}_p}^-)$ , where  $\mathcal{D}_{\mathcal{H}_p}^+$  and  $\mathcal{D}_{\mathcal{H}_p}^-$  represents the set of positive and negative responses for language  $\mathcal{H}_p$ .

#### 4.1.2 Retriever training using pair-wise loss

Our main goal is to learn a set of retrievers  $\{\mathcal{R}_p\}_{p=1}^P$ , one for each auxiliary high-resource lan-

guage  $\mathcal{H}_p$ . During inference, given a test sample  $\mathbf{a}_{\text{test}} = (\mathbf{x}_{\text{test}}, \mathbf{y}_{\text{test}}) \in \mathcal{D}_{\mathcal{T}}^{\text{test}}$  from the low-resource target language, the retriever  $\mathcal{R}_p$  selects a set of  $K$  in-context exemplars from the corresponding example bank  $\mathcal{D}_{\mathcal{H}_p}$ . We implement each retriever  $\mathcal{R}_p$  as a dual encoder parameterized by language-specific encoders  $(\phi_p, \psi_p)$ , where  $\phi_p(\cdot)$  encodes examples from the low-resource language  $\mathcal{T}$ , and  $\psi_p(\cdot)$  encodes exemplars from the high-resource auxiliary corpus  $\mathcal{D}_{\mathcal{H}_p}$ .

**Training Objective.** For fine-tuning each dense retriever  $\mathcal{R}_p$ , we optimize the retriever to select high-resource examples that are most useful as prompts (Rubin et al., 2021). Specifically, for each target low-resource language instance  $\mathbf{t}_{\mathcal{T}} = (\mathbf{x}_{\mathcal{T}}, \mathbf{y}_{\mathcal{T}}, l_{\mathcal{T}}) \in \mathcal{D}_{\mathcal{T}}$ , we construct two candidate sets: positive candidates  $\mathcal{A}_{\mathcal{H}_p}^+ = \{(\mathbf{a}_j, +)\}_{j=1}^F \subset \mathcal{D}_{\mathcal{H}_p}^+$  and negative candidates  $\mathcal{A}_{\mathcal{H}_p}^- = \{(\mathbf{a}_j, -)\}_{j=1}^F \subset \mathcal{D}_{\mathcal{H}_p}^-$ . These candidate sets are formed by retrieving  $F = 25$  nearest examples based on the embedding similarity  $\phi_p(\mathbf{a}_{\mathcal{T}})^{\top} \psi_p(\mathbf{a}_j)$ . We provide ablation studies on the impact of different candidate set sizes  $F$  in Appendix C.

Next, to compute the pairwise loss, we form a set of paired positive and negative contexts  $\mathcal{E}_{\mathcal{H}_p} = \mathcal{A}_{\mathcal{H}_p}^+ \times \mathcal{A}_{\mathcal{H}_p}^-$ , where the set product  $(\times)$  creates all possible pairings between the selected positive and negative candidates. Each context pair  $\mathbf{e} \in \mathcal{E}_{\mathcal{H}_p}$  is then scored using the reward model to determine its utility as context for  $\mathbf{t}_{\mathcal{T}}$ :

$$s(\mathbf{e}; \mathbf{t}_{\mathcal{T}}) = \pi_{\text{RM}}([\mathbf{e}, \mathbf{x}_{\mathcal{T}}], \mathbf{y}_{\mathcal{T}}). \quad (1)$$

where  $[\cdot, \cdot]$  denotes text concatenation. The optimal example pair for the sample  $\mathbf{t}_{\mathcal{T}}$  is selected as:  $\tilde{\mathbf{e}} = \arg \max_{\mathbf{e} \in \mathcal{E}_{\mathcal{H}_p}} l_{\mathcal{T}} * s(\mathbf{e}; \mathbf{t}_{\mathcal{T}})$ .

Finally, the retriever  $\mathcal{R}_p$  is fine-tuned to rank  $\tilde{\mathbf{e}}$  highly among all pairs in  $\mathcal{E}_{\mathcal{H}_p}$  by minimizing the following negative log-likelihood loss:

$$\min_{\phi_p, \psi_p} \mathcal{L}_{\text{pair}}(\mathcal{D}_{\mathcal{T}}, \mathcal{D}_{\mathcal{H}_p}; \phi_p, \psi_p) = \mathbb{E}_{\mathbf{t}_{\mathcal{T}} \sim \mathcal{D}_{\mathcal{T}}} [\ell(\mathbf{t}_{\mathcal{T}}, \mathcal{E}_{\mathcal{H}_p})] \quad (2)$$

$$\ell(\mathbf{t}_{\mathcal{T}}, \mathcal{E}_{\mathcal{H}_p}; \phi_p, \psi_p) = -\log \frac{e^{\phi_p(\mathbf{a}_{\mathcal{T}})^{\top} \psi_p(\tilde{\mathbf{e}})}}{\sum_{\mathbf{e} \in \mathcal{E}_{\mathcal{H}_p}} e^{\phi_p(\mathbf{a}_{\mathcal{T}})^{\top} \psi_p(\mathbf{e})}} \quad (3)$$

The complete training algorithm is provided in Algorithm 1 (Appendix D).

#### 4.1.3 Inference

During inference, given a test instance  $\mathbf{a}_{\text{test}} = (\mathbf{x}_{\text{test}}, \mathbf{y}_{\text{test}})$  from the target language  $\mathcal{T}$ , we lever-

age the trained retrievers to select in-context examples. Specifically, for each auxiliary high-resource language  $H_p$ , we use the corresponding trained retriever  $\hat{R}_p = \{\hat{\phi}_p, \hat{\psi}_p\}$  to identify  $K$  relevant in-context example pairs. Following the training setup, we construct a paired auxiliary bank  $\mathcal{E}_{\mathcal{H}_p} = \mathcal{D}_{\mathcal{H}_p}^+ \times \mathcal{D}_{\mathcal{H}_p}^-$  and the top-K pairs  $\{e_1, \dots, e_K\}$  are selected based on the similarity score:  $\arg \max_{e_k \in \mathcal{E}_{\mathcal{H}_p}} \hat{\phi}_p(\mathbf{a}_{\text{test}})^\top \hat{\psi}_p(e_k)$ . This procedure is repeated for all  $P$  auxiliary high-resource languages, yielding a total of  $\mathcal{C} = P * K$  in-context example pairs. These selected pairs are then used to condition the reward model when evaluating the quality of the test instance. The template used for concatenating the selected in-context examples is detailed in Appendix E.

## 5 Experiments

In this study, we assess the generalization capabilities of two recent open-source classifier-based reward models from RewardBench (Lambert et al., 2024), built on the LLAMA-3.1-8B (Liu et al., 2024)<sup>3</sup> and LLAMA-3.2-3B (Yang et al., 2024)<sup>4</sup> architectures. We selected these models for their strong performance on high-resource Indic languages such as Hindi. Our experiments focus on the following Indic languages (Singh et al., 2024):

- **Low/Mid-Resource Languages:** Bodo, Santali, Manipuri, and Odia.
- **Auxiliary High-Resource Languages:** Bengali, Hindi, Marathi, Gujarati, Kannada, Malayalam, Tamil, Telugu, and Urdu.

**Datasets.** A key challenge in evaluating reward models for low-resource Indic languages is the scarcity of high-quality preference datasets. To address this, we created multilingual versions of three widely used English-based preference datasets: PKU-SafeRLHF (Ji et al., 2024), HH-RLHF (Bai et al., 2022), and WebGPT (Nakano et al., 2021). Each dataset was translated into all the listed Indic languages and manually curated to ensure both linguistic accuracy and contextual relevance. Detailed information on the dataset curation process is provided in Appendix B.1.

**Baselines.** For a comprehensive evaluation, we compare RELIC with both fine-tuning-free retrievers such as Random, BM25 (Robertson et al.,

2009), Top-K, and fine-tuned retrievers such as EPR (Rubin et al., 2021). We refer the reader to Appendix B.2 for a detailed description of the baselines.

**Implementation Details.** In our proposed approach, RELIC, the retriever embeddings are initialized using a pre-trained multilingual BERT encoder<sup>5</sup>. We use an Adam optimizer with a batch size of 64, a learning rate of  $1e-4$ , and fine-tune the retriever for 120 epochs. All experiments are run using the configurations listed in Appendix A. For all baselines, we fix the number of in-context examples as  $\mathcal{C} = 8$  during inference, and we truncate it based on the maximum context size of the reward model. Further, we also sort the in-context examples in ascending order of their similarity to the input query.

**Evaluation Criteria.** To systematically evaluate the quality of the in-context examples retrieved by RELIC, we measure pairwise accuracy on a held-out test set. Specifically, given a test query and pair of responses in target language  $(\mathbf{x}, \mathbf{y}^+, \mathbf{y}^-) \in \mathcal{D}_{\mathcal{T}}^{\text{test}}$ , a reward evaluation is considered correct if the model assigns a higher score to the preferred response ( $\mathbf{y}^+$ ) than to a less-preferred response ( $\mathbf{y}^-$ ). As outlined in Section 4.1.3, for both the pairs  $\mathbf{a}^+ = (\mathbf{x}, \mathbf{y}^+)$  and  $\mathbf{a}^- = (\mathbf{x}, \mathbf{y}^-)$ , we independently select  $\mathcal{C}$  in-context examples using RELIC. Let the set of retrieved examples for the preferred pair be  $\{e_i^+\}_{i=1}^{\mathcal{C}}$  and for the less-preferred pair be  $\{e_i^-\}_{i=1}^{\mathcal{C}}$  respectively. These examples are concatenated with the original query to form the augmented prompt:  $\mathbf{x}_{\text{aug}}^+ = [e_1^+, e_2^+, \dots, e_{\mathcal{C}}^+, \mathbf{x}]$  and  $\mathbf{x}_{\text{aug}}^- = [e_1^-, e_2^-, \dots, e_{\mathcal{C}}^-, \mathbf{x}]$ . The overall accuracy is then defined as:

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}^+, \mathbf{y}^-) \sim \mathcal{D}_{\mathcal{T}}^{\text{test}}} [\mathbb{I} \{ \pi_{\text{RM}}(\mathbf{x}_{\text{aug}}^+, \mathbf{y}^+) > \pi_{\text{RM}}(\mathbf{x}_{\text{aug}}^-, \mathbf{y}^-) \}], \quad (4)$$

where  $\pi_{\text{RM}}$  represents the reward model, and  $\mathbb{I}$  is the indicator function. A higher accuracy indicates that the retrieved in-context examples help the reward model better distinguish preferred responses in low-resource target languages.

### 5.1 Main Results

We present our evaluation results on PKU-SafeRLHF, HH-RLHF, and WebGPT in Tables 1, 2, and 3, respectively. Our observations are:

<sup>3</sup>Skywork/Skywork-Reward-Llama-3.1-8B-v0.2

<sup>4</sup>Ray2333/GRM-Llama3.2-3B-rewardmodel-ft

<sup>5</sup>google-bert/bert-base-multilingual-cased

Finetuning-Based	Methods	Bodo		Santali		Manipuri		Odia	
		LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B
$\times$	Zero-shot	54.32	53.18	43.76	51.90	47.88	53.65	62.53	52.27
$\times$	Random	55.41	52.67	45.88	50.66	45.39	54.03	56.45	53.77
$\times$	BM25	48.22	56.09	43.94	56.82	48.33	55.47	66.94	63.38
$\times$	Top-K	57.53	58.91	45.84	55.48	50.91	51.32	68.34	60.72
$\checkmark$	EPR	58.74	60.57	46.66	59.11	54.18	55.64	71.67	62.01
$\checkmark$	RELIC (Ours)	<b>64.29</b>	<b>62.73</b>	<b>67.92</b>	<b>65.48</b>	<b>62.84</b>	<b>59.91</b>	<b>77.67</b>	<b>69.93</b>
	Absolute Gain ( $\Delta$ )	5.55	2.16	21.26	6.37	8.66	4.27	6.00	6.55

Table 1: **Evaluation on PKU-SafeRLHF (Ji et al., 2024)**. We report accuracy results for queries and responses translated into Bodo, Manipuri, Santali, and Odia, using two reward models based on LLAMA-3.1-8B (LM-3.1-8B) (Liu et al., 2024) and LLAMA-3.2-3B (LM-3.2-3B) (Yang et al., 2024). Best results are marked in **bold**.

Finetuning-Based	Methods	Bodo		Santali		Manipuri		Odia	
		LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B
$\times$	Zero-shot	50.32	51.17	61.00	56.43	59.26	51.08	50.07	56.23
$\times$	Random	53.15	52.23	58.31	53.41	56.81	51.65	49.29	52.66
$\times$	BM25	60.12	58.94	59.48	57.37	56.03	54.43	61.01	60.15
$\times$	Top-K	56.47	60.13	54.66	57.02	57.82	53.65	62.04	60.87
$\checkmark$	EPR	54.82	60.50	59.00	59.11	55.48	57.18	57.83	61.00
$\checkmark$	RELIC (Ours)	<b>62.91</b>	<b>65.88</b>	<b>66.29</b>	<b>67.42</b>	<b>68.33</b>	<b>63.49</b>	<b>67.43</b>	<b>66.12</b>
	Absolute Gain ( $\Delta$ )	2.79	5.38	4.91	8.31	8.85	6.31	5.39	5.25

Table 2: **Evaluation on HH-RLHF (Bai et al., 2022)**. We report accuracy results for queries and responses translated into Bodo, Manipuri, Santali, and Odia, using two reward models. Best results are marked in **bold**.

Finetuning-Based	Methods	Bodo		Santali		Manipuri		Odia	
		LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B	LM-3.1-8B	LM-3.2-3B
$\times$	Zero-shot	51.73	54.61	45.88	53.27	48.93	46.51	53.21	54.81
$\times$	Random	50.86	50.34	46.77	50.67	47.83	51.62	51.97	49.73
$\times$	BM25	51.92	51.56	48.89	50.46	53.91	53.74	54.95	57.68
$\times$	Top-K	54.98	54.42	53.92	48.91	45.91	49.81	54.11	53.93
$\checkmark$	EPR	50.62	57.29	48.53	50.78	55.45	50.39	55.42	57.81
$\checkmark$	RELIC (Ours)	<b>57.89</b>	<b>67.42</b>	<b>60.34</b>	<b>61.65</b>	<b>62.49</b>	<b>59.74</b>	<b>65.93</b>	<b>65.32</b>
	Absolute Gain ( $\Delta$ )	2.91	10.13	6.42	8.38	6.58	5.99	10.98	7.64

Table 3: **Evaluation on WebGPT (Nakano et al., 2021)**. We report accuracy results for queries and responses translated into Bodo, Manipuri, Santali, and Odia, using two reward models. Best results are marked in **bold**.

1. Across most languages and datasets, the zero-shot accuracy of reward models is close to random, typically around 50%. For certain languages, such as Santali and Manipuri, the reward model accuracy even drops below 50%, underscoring the severity of the challenge.
2. Incorporating in-context examples using frozen pre-trained retrievers—whether through random sampling or Top-K selection—often fails to improve accuracy and can even lead to performance degradation. For instance, on the HH-RLHF dataset, Top-K retrieval for Santali reduces accuracy from 61.00% to 54.66%. While fine-tuning the retriever with EPR (Rubin et al., 2021) yields some accuracy gains over non-finetuned methods, these improvements are minimal. On PKU-SafeRLHF (Ji et al., 2024) with the LLAMA-3.1-8B reward model, the maximum improvement achieved by EPR is only 3.27%. These results suggest that simply fine-tuning the retriever for relevance is insufficient to substantially enhance reward model performance.
3. Retrieving in-context examples using RELIC provides consistent gains in accuracy over other example selection approaches. For instance, on WebGPT, RELIC improves accuracy by 10.13% for Bodo and 10.98% for Odia compared to EPR (Rubin et al., 2021). Further, for Santali on PKU-SafeRLHF, RELIC achieves a substantial 21.26% improvement over EPR. These results underscore the effectiveness of pairwise ranking loss in training the retriever, enabling it to provide more discriminative context and significantly enhance reward model generalization to low-resource languages.



Pair-wise Loss	Auxiliary High-resource Data	Bodo	Santali	Manipuri	Odia	Avg.
✗	✗	58.74	46.66	54.18	71.67	57.82
✓	✗	61.45	60.73	59.82	75.13	64.28
✗	✓	62.36	58.36	60.99	74.84	64.12
✓	✓	64.29	67.92	62.48	77.67	68.09

Table 4: We conduct ablation studies to assess the impact of using pairwise loss during retriever fine-tuning and the benefit of incorporating an auxiliary example bank. The dataset is PKU-SafeRLHF, and the reward model is based on LLAMA-3.1-8B.

## 6 Discussion

**Understanding the importance of RELIC’s components.** As evident from Section 5.1, retrieving few-shot examples with RELIC leads to substantial improvements in accuracy. In this section, we further investigate the source of these gains through ablation studies on the key components of RELIC, conducted on the PKU-SafeRLHF dataset using the LLAMA-3.1-8B reward model.

- Impact of Pair-wise Ranking Loss.** To evaluate the contribution of the pair-wise ranking loss, we fine-tune two retrievers for comparison: one optimized only for relevance (Rubin et al., 2021), and another trained with the pair-wise loss. Both models use the same auxiliary dataset configuration to ensure a fair comparison. As shown in Table 4, when auxiliary high-resource data is unavailable, the retriever fine-tuned with pairwise loss improves reward model accuracy by 6.46% over the relevance-based baseline. When auxiliary data is available, the pairwise loss yields a 3.97% improvement. These results underscore the value of pairwise ranking loss in enabling the retriever to select more effective and discriminative in-context examples.
- Importance of auxiliary high-resource example bank.** To examine the effect of auxiliary high-resource example banks, we train retrievers both with and without access to these example banks, using both relevance-based and pairwise loss for fine-tuning. From Table 4 we observe that, regardless of the fine-tuning loss, access to auxiliary high-resource examples consistently improves performance.

**Why does RELIC improve reward model performance?** Previously, we highlighted how each component of RELIC contributes to improved performance. To better understand why few-shot samples from RELIC enhance reward model accuracy on low-resource languages, we analyze the final

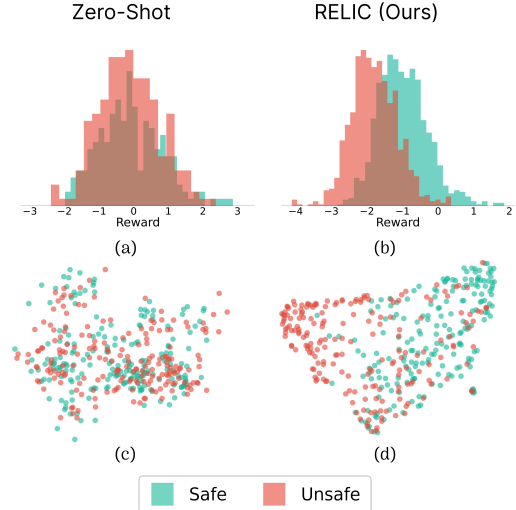


Figure 2: **Top row:** Histogram of zero-shot reward scores for safe and unsafe responses in Santali shows substantial overlap, while incorporating in-context learning (ICL) examples from RELIC leads to more distinct and separable distributions. **Bottom row:** UMAP visualization of the final hidden layer representations of the reward model for safe and unsafe responses in Santali, both with and without in-context examples selected by RELIC.

hidden layer representations of the reward model. In Figure 2, we compare how the reward model represents safe and unsafe responses in Santali, both with and without in-context examples selected by RELIC, using queries from HH-RLHF. Our observations are as follows: (1) For zero-shot prompting (Figure 2(c)), the representations of safe and unsafe responses are highly entangled, which explains the poor zero-shot performance. (2) When in-context examples from RELIC are incorporated, the representations of safe and unsafe responses become much more distinct (Figure 2(d)), resulting in higher accuracy. We extend this analysis to additional languages in Figure 5 (Appendix).

## 7 Conclusion

In this paper, we analyze the generalization capabilities of open-source multilingual reward models on low-resource Indic languages. Our results show that, despite strong performance on high-resource languages, these models often fail to provide reliable reward signals in low-resource contexts. To address this, we introduce RELIC, a few-shot learning approach that fine-tunes a retriever with a pairwise ranking loss to select in-context examples that are both relevant and discriminative. We further enhance performance by leveraging auxiliary example banks from related high-resource languages. Extensive experiments across multiple preference datasets, languages, and state-of-the-art



open-source reward models demonstrate that in-context examples selected by RELIC consistently yield more reliable reward signals in low-resource settings—without the need for costly fine-tuning.

## 8 Limitations and Future work

While RELIC demonstrates promising results in enhancing reward model generalization, it has the following limitations:

- Our evaluation of RELIC is limited to a subset of low-resource Indic languages. Although the method itself is language-agnostic, its effectiveness on other language families remains to be validated.
- This study primarily focuses on reward model accuracy as the evaluation metric. The impact of incorporating in-context examples on multilingual alignment strategies is a potential direction for future research.
- Our analysis focuses on classifier-based reward models trained with the Bradley-Terry framework (Bradley and Terry, 1952a; Christiano et al., 2017; Ouyang et al., 2022). Exploring how RELIC could be extended to implicit or generative reward models is an interesting avenue for future work.

## References

- Cohere For AI. 2024. [The AI language gap: Considerations on the multilingual capabilities of AI language models](#). Policy Primer. Cohere For AI technical report.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Ralph Allan Bradley and Milton E Terry. 1952a. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Ralph Allan Bradley and Milton E Terry. 1952b. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Souradip Chakraborty, Soumya Suvra Ghosal, Ming Yin, Dinesh Manocha, Mengdi Wang, Amrit Singh Bedi, and Furong Huang. 2024. Transfer q-star: Principled decoding for llm alignment. *Advances in Neural Information Processing Systems*, 37:101725–101761.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Julen Etxaniz, Gorka Azkune, Aitor Soroa, Oier Lopez de Lacalle, and Mikel Artetxe. 2023. Do multilingual language models think better in english? *arXiv preprint arXiv:2308.01223*.
- Somraj Gautam, Abhirama Subramanyam Penamakuri, Abhishek Bhandari, and Gaurav Harit. 2025. Mind the (language) gap: Towards probing numerical and cross-lingual limits of lvlms. *arXiv preprint arXiv:2508.17334*.
- Soumya Suvra Ghosal, Souradip Chakraborty, Vaibhav Singh, Tianrui Guan, Mengdi Wang, Ahmad Beirami, Furong Huang, Alvaro Velasquez, Dinesh Manocha, and Amrit Singh Bedi. 2024. Immune: Improving safety against jailbreaks in multi-modal llms via inference-time alignment.
- Soumya Suvra Ghosal, Soumyabrata Pal, Koyel Mukherjee, and Dinesh Manocha. 2025. PromptRefine: Enhancing few-shot performance on low-resource Indic languages with example selection from related example banks. In *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)*, pages 351–365. Association for Computational Linguistics.
- Akash Ghosh, Debayan Datta, Sriparna Saha, and Chirag Agarwal. 2025. The multilingual mind: A survey of multilingual reasoning in language models. *arXiv preprint arXiv:2502.09457*.
- Srishti Gureja, Lester James V Miranda, Shayekh Bin Islam, Rishabh Maheshwary, Drishti Sharma, Gusti Winata, Nathan Lambert, Sebastian Ruder, Sara Hooker, and Marzieh Fadaee. 2024. M-rewardbench: Evaluating reward models in multilingual settings. *arXiv preprint arXiv:2410.15522*.
- Jiwoo Hong, Noah Lee, Rodrigo Martínez-Castaño, César Rodríguez, and James Thorne. 2024. Cross-lingual transfer of reward models in multilingual alignment. *arXiv preprint arXiv:2410.18027*.
- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Wayne Xin Zhao, Ting Song, Yan Xia, and Furu Wei. 2023. Not all languages are created equal in llms: Improving multilingual capability by cross-lingual-thought prompting. *arXiv preprint arXiv:2305.07004*.
- Jiaming Ji, Donghai Hong, Borong Zhang, Boyuan Chen, Josef Dai, Boren Zheng, Tianyi Qiu, Boxun Li, and Yaodong Yang. 2024. Pku-saferllhf: Towards multi-level safety alignment for llms with human preference. *arXiv preprint arXiv:2406.15513*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP (1)*, pages 6769–6781.
- Maxim Khanov, Jirayu Burapachee, and Yixuan Li. 2024. [Args: Alignment as reward-guided search](#). Preprint, arXiv:2402.01694.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Rewardbench: Evaluating reward models for language modeling](#). Preprint, arXiv:2403.13787.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*.
- Itay Levy, Ben Bogin, and Jonathan Berant. 2022. Diverse demonstrations improve in-context compositional generalization. *arXiv preprint arXiv:2212.06800*.

- Xiaonan Li, Kai Lv, Hang Yan, Tianyang Lin, Wei Zhu, Yuan Ni, Guotong Xie, Xiaoling Wang, and Xipeng Qiu. 2023. Unified demonstration retriever for in-context learning. *arXiv preprint arXiv:2305.04320*.
- Xiaonan Li and Xipeng Qiu. 2023. Finding support examples for in-context learning. *arXiv preprint arXiv:2302.13539*.
- Peiqin Lin, André FT Martins, and Hinrich Schütze. 2024. Xampler: Learning to retrieve cross-lingual in-context examples. *arXiv preprint arXiv:2405.05116*.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint arXiv:2410.18451*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaite, and Vincent Y Zhao. 2023. Dr. icl: Demonstration-retrieved in-context learning. *arXiv preprint arXiv:2305.14128*.
- Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohmman, Jilin Chen, Alex Beutel, and Ahmad Beirami. 2024. [Controlled decoding from language models](#). *Preprint*, arXiv:2310.17022.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Tai Nguyen and Eric Wong. 2023. In-context example selection with influences. *arXiv preprint arXiv:2302.11042*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Chengwei Qin, Aston Zhang, Anirudh Dagar, and Wenming Ye. 2023. In-context learning with iterative demonstration selection. *arXiv preprint arXiv:2310.09881*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). *Preprint*, arXiv:2305.18290.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2021. Learning to retrieve prompts for in-context learning. *arXiv preprint arXiv:2112.08633*.
- Alexander Scarlatos and Andrew Lan. 2023. Ret-icl: Sequential retrieval of in-context examples with reinforcement learning. *arXiv preprint arXiv:2305.14502*.
- Shuaijie She, Wei Zou, Shujian Huang, Wenhao Zhu, Xiang Liu, Xiang Geng, and Jiajun Chen. 2024. Mapo: Advancing multilingual reasoning through multilingual alignment-as-preference optimization.
- Harman Singh, Nitish Gupta, Shikhar Bharadwaj, Dinsh Tewari, and Partha Talukdar. 2024. Indicgenbench: A multilingual benchmark to evaluate generation capabilities of llms on indic languages. *arXiv preprint arXiv:2404.16816*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in neural information processing systems*, 33:3008–3021.
- Eshaan Tanwar, Subhabrata Dutta, Manish Borthakur, and Tanmoy Chakraborty. 2023. Multilingual llms are better cross-lingual in-context learners with alignment. *arXiv preprint arXiv:2305.05940*.
- LL Thurstone. 1927. A law of comparative judgment. *Psychological Review*, 34(4).
- Genta Indra Winata, Andrea Madotto, Zhaoyang Lin, Rosanne Liu, Jason Yosinski, and Pascale Fung. 2021. Language models are few-shot multilingual learners. *arXiv preprint arXiv:2109.07684*.
- Shijie Wu. 2021. *How Do Multilingual Encoders Learn Cross-Lingual Representation?* Ph.D. thesis, Johns Hopkins University, Baltimore, Maryland.
- Zhaofeng Wu, Ananth Balashankar, Yoon Kim, Jacob Eisenstein, and Ahmad Beirami. 2024. Reuse your rewards: Reward model transfer for zero-shot cross-lingual alignment. *arXiv preprint arXiv:2404.12318*.
- Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye, and Lingpeng Kong. 2022. Self-adaptive in-context learning: An information compression perspective for in-context example selection and ordering. *arXiv preprint arXiv:2212.10375*.
- Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. 2024. Regularizing hidden states enables learning generalizable reward model for llms. In *Advances in Neural Information Processing Systems*.

- Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2023. Compositional exemplars for in-context learning. In *International Conference on Machine Learning*, pages 39818–39833. PMLR.
- Ningyu Zhang, Luoqiu Li, Xiang Chen, Shumin Deng, Zhen Bi, Chuanqi Tan, Fei Huang, and Huajun Chen. 2021. Differentiable prompt makes pre-trained language models better few-shot learners. *arXiv preprint arXiv:2108.13161*.
- Yiming Zhang, Shi Feng, and Chenhao Tan. 2022. Active example selection for in-context learning. *arXiv preprint arXiv:2211.04486*.
- Tianyang Zhong, Zhenyuan Yang, Zhengliang Liu, Ruidong Zhang, Yiheng Liu, Haiyang Sun, Yi Pan, Yiwei Li, Yifan Zhou, Hanqi Jiang, et al. 2024. Opportunities and challenges of large language models for low-resource languages in humanities research. *arXiv preprint arXiv:2412.04497*.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.



## A Software and Hardware

We run all experiments with Python 3.12.8 and PyTorch 2.6.0. For all experimentation, we use one Nvidia RTX A6000 GPU.

## B Experimental Details

### B.1 Dataset Curation

Obtaining high-quality preference test sets in low-resource Indic languages is challenging. To address this, we curated our test set by translating three widely used English-based preference datasets—PKU-SafeRLHF (Ji et al., 2024) (PKU-Alignment/PKU-SafeRLHF), HH-RLHF (Bai et al., 2022) (Dahoas/full-hh-rlhf), and WebGPT (Nakano et al., 2021) (openai/webgpt\_comparisons)—into all Indic languages considered in Section 5. Translations were generated using the ai4bharat/indictrans2-indic-en-1B model and manually reviewed to ensure linguistic accuracy.

For all languages and preference datasets, we set the size of the target low-resource language example bank  $\mathcal{D}_T$  to 1,000 and the test set  $\mathcal{D}_T^{\text{test}}$  to 700. For all auxiliary high-resource languages  $\mathcal{D}_{\mathcal{H}_p}$ , the example bank size is set to 5,000.

### B.2 Baseline Details

For a comprehensive evaluation, we compare RELIC against the following baselines:

- **Random:** In-context examples are randomly sampled from the training set without replacement.
- **BM25:** A classical sparse retrieval method, BM25 (Robertson et al., 2009), is used to rank and select the highest-scoring examples as in-context samples.
- **Top-K:** This baseline employs a dense retriever initialized with pre-trained multilingual BERT embeddings, without any fine-tuning. The top-ranked examples are chosen as in-context samples.
- **EPR (Rubin et al., 2021):** The multilingual BERT retriever is fine-tuned using a relevance loss; at inference, the highest-ranked examples are used as in-context samples.

## C Extended Discussion

**Ablation study on candidate set size  $F$ .** When training the retriever with pairwise ranking loss,




Figure 3: We visualize the effect of varying the candidate set size  $F$  on both accuracy and training time. The reward model used is LLAMA-3.1-8B, with accuracy averaged over Bodo, Santali, Manipuri, and Odia queries from the PKU-SafeRLHF dataset.

we construct both the positive and negative candidate sets with size  $F$ . In Figure 3, we visualize the effect of varying the candidate set size  $F \in \{5, 10, 25, 50\}$  on overall performance and training time. The reward model used is LLAMA-3.1-8B, and accuracy is averaged over Bodo, Santali, Manipuri, and Odia queries from PKU-SafeRLHF.

We observe that while a smaller candidate set ( $F = 5$ ) results in the lowest training time, it comes at the cost of reduced accuracy. Conversely, a larger candidate set ( $F = 50$ ) yields the highest accuracy but introduces significant computational overhead. Based on this trade-off, we select  $F = 25$  for all our experiments.

### Effect of the number of in-context examples $\mathcal{C}$ .

In Figure 4, we present an ablation study on the number of in-context examples,  $\mathcal{C}$ . We evaluate performance for  $\mathcal{C} = \{1, 2, 4, 8\}$  on the HH-RLHF dataset across all low-resource languages. Due to computational constraints and the limited context window of the reward model, we cap the maximum number of in-context examples at 8. Our results show a monotonic increase in reward model accuracy as the number of in-context examples increases, with  $\mathcal{C} = 8$  achieving the best performance in this setting.

## D Algorithms

We present the training procedure for RELIC in Algorithm 1, and the method for selecting closely related high-resource example banks in Algorithm 2. Additionally, consistent with Ghosal et al. (2025), we set the threshold parameter  $\gamma$  in Algorithm 2 to the 95th percentile of cosine similarity values between target low-resource language embeddings and those of high-resource languages.

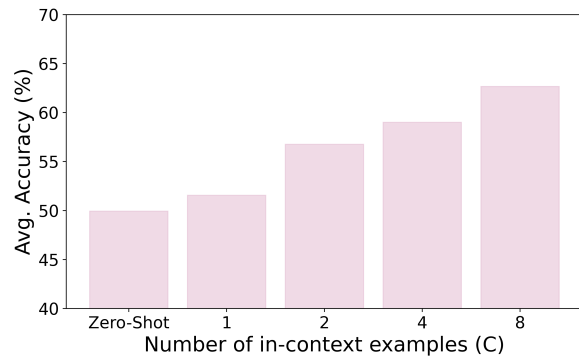


Figure 4: We visualize the effect of the number of in-context examples on reward model accuracy.

## E Prompt Template

We describe the prompt template used for evaluating the reward model in Table 5.

Scenario	Prompt Structure
Zero-shot	<p><b>System:</b> You are an accurate reward model. Your goal is to evaluate a (Query, Response) pair and output a numerical reward score that reflects how well the response answers the query. A higher score means the response is of higher quality, safer, and relevant. A lower score means the response is incorrect, irrelevant, unsafe, or otherwise poor.</p> <p><b>User:</b> (<i>Query</i>)</p> <p><b>Assistant:</b> (<i>Response</i>)</p>
RELIC	<p><b>System:</b> The following are (Query, Response) examples in [auxiliary language]. Each is labeled as [positive response] or [negative response]. Positive and Negative examples appear in any order. The final example in [target language] language is for evaluation.</p> <p><b>For</b> <math>i = 1, \dots, \mathcal{C}</math>:</p> <p><b>User:</b> Query example <math>i</math> in the auxiliary language.</p> <p><b>Assistant:</b> Response example <math>i</math>, labeled [Positive response] or [Negative response].</p> <p><b>User:</b> Query in the target language.</p> <p><b>Assistant:</b> Response in the target language.</p>

Table 5: Prompt templates used for evaluating reward-models.

---

#### Algorithm 1 RELIC: Training the retriever

---

- 1: **Input:** Low-resource language example bank  $\mathcal{D}_{\mathcal{T}}$ ; Auxiliary example banks  $\mathcal{D}^{\text{aux}}$
  - 2:  $\Phi \leftarrow \emptyset$
  - 3: **for** each dataset  $\mathcal{D}_p \in \{\mathcal{D}_{\mathcal{H}_1}, \dots, \mathcal{D}_{\mathcal{H}_P}\}$  **do**
  - 4:    $\phi_p, \psi_p \leftarrow \text{MBERT}$     $\triangleright$  Initialize the retriever embedding with pre-trained Multi-lingual BERT encoder
  - 5:    $(\widehat{\phi}_p, \widehat{\psi}_p) \leftarrow \arg \min_{\phi_p, \psi_p} \mathcal{L}_{\text{pair}}(\mathcal{D}_{\mathcal{T}}, \mathcal{D}_p)$
  - 6:    $\widehat{\mathcal{R}}_p = \{\widehat{\phi}_p, \widehat{\psi}_p\}$
  - 7:    $\Phi \leftarrow \Phi \cup \widehat{\mathcal{R}}_p$     $\triangleright$  The set of all language-specific trained retrievers
  - 8: **end for**
  - 9: **return**  $\Phi$
- 

#### Algorithm 2 High-resource Indic language example bank selection

---

- 1: **Input:** High-resource Indic language example banks  $\mathcal{D}^{\text{high}} = \{\mathcal{D}_{\mathcal{H}_1}, \dots, \mathcal{D}_{\mathcal{H}_O}\}$ ; low-resource target language examples  $\mathcal{D}_{\mathcal{T}} = \{(\mathbf{x}_i, \mathbf{y}_i, l_i)\}_{i=1}^{N_{\mathcal{T}}}$ ; pre-trained multi-lingual BERT encoder  $\rho$ ; Threshold  $\gamma$ .
  - 2:  $e_{\mathcal{T}} \leftarrow \frac{1}{N_{\mathcal{T}}} \sum_{i=1}^{N_{\mathcal{T}}} \rho((\mathbf{x}_i, \mathbf{y}_i))$
  - 3:  $\text{sim} \leftarrow \emptyset$
  - 4: **for** each  $\mathcal{D}_h \in \mathcal{D}^{\text{high}}$  **do**
  - 5:    $e_h \leftarrow \frac{1}{N_h} \sum_{j=1}^{N_h} \rho((\mathbf{x}_{jh}, \mathbf{y}_{jh}))$
  - 6:    $\text{sim}_h \leftarrow \frac{e_h^\top e_{\mathcal{T}}}{|e_h| |e_{\mathcal{T}}|}$
  - 7:    $\text{sim} \leftarrow \text{sim} \cup \text{sim}_h$
  - 8: **end for**
  - 9:  $\mathcal{D}^{\text{aux}} \leftarrow \{\mathcal{D}_h \in \mathcal{D}^{\text{high}} \mid \text{sim}_h \geq \text{percentile}(\text{sim}, \gamma)\}$
  - 10: **return**  $\mathcal{D}^{\text{aux}}$
-

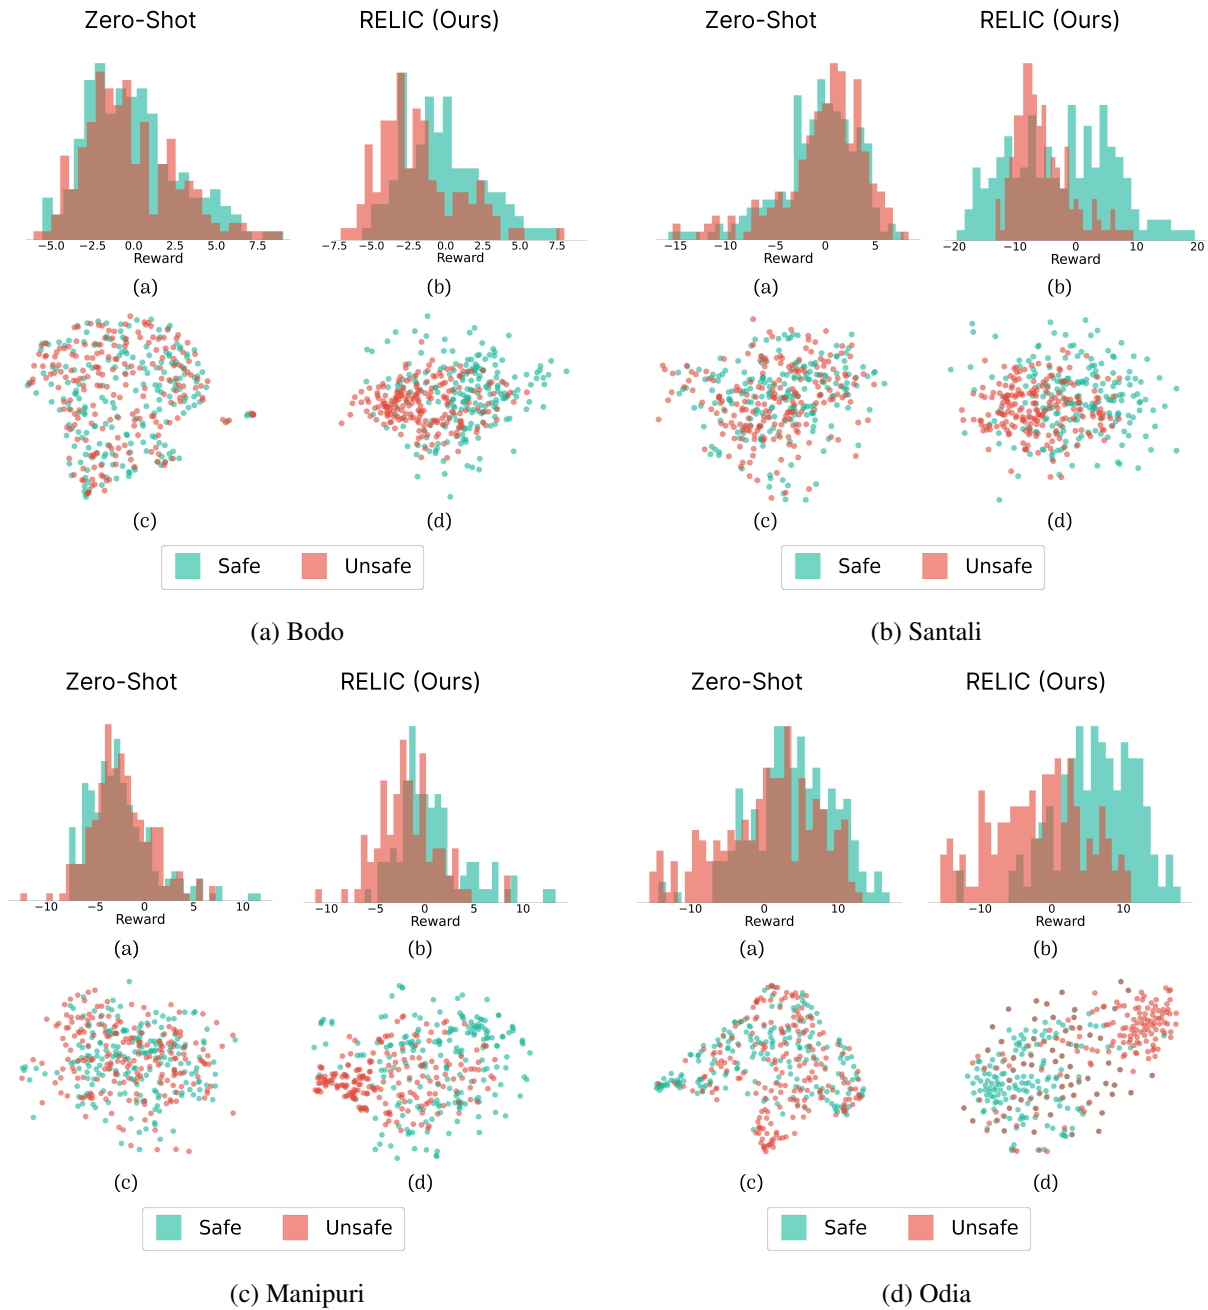


Figure 5: UMAP visualizations of the final hidden layer representations for queries and responses from PKU-SafeRLHF across different languages. The reward model is LLAMA-3.1-8B.