Reuse Your Rewards:
Reward Model Transfer for Zero-Shot Cross-Lingual Alignment

Zhaofeng Wu [1] **Ananth Balashankar** [2] **Yoon Kim** [1] **Jacob Eisenstein** [2] **Ahmad Beirami** [2]

Abstract

Aligning language models (LMs) based on human-annotated preference data is a crucial step in obtaining practical and performant LM-based systems. However, multilingual human preference data are difficult to obtain at scale, making it challenging to extend this framework to diverse languages. In this work, we evaluate a simple approach for zero-shot cross-lingual alignment, where a reward model is trained on preference data in one source language and directly applied to other target languages. On summarization and open-ended dialog generation, we show that this method is consistently successful under comprehensive evaluation settings, including human evaluation: cross-lingually aligned models are preferred by humans over unaligned models on up to >70% of evaluation instances. We moreover find that a different-language reward model sometimes yields better aligned models than a same-language reward model. We also identify best practices when there is no language-specific data for even supervised finetuning, another component in alignment.

1. Introduction

Alignment has become an indispensable stage for building practical language models (LMs) adjusted to human preferences. This additional step, however, makes it challenging to develop LMs for many languages: unlike for autoregressive language modeling where multilingual unlabeled data may be easy to obtain (Joshi et al., 2020), such as religious texts (Christodouloupoulos & Steedman, 2015), labeled preference data can be expensive to gather. How do we align a LM in a target language without any preference data in that language?
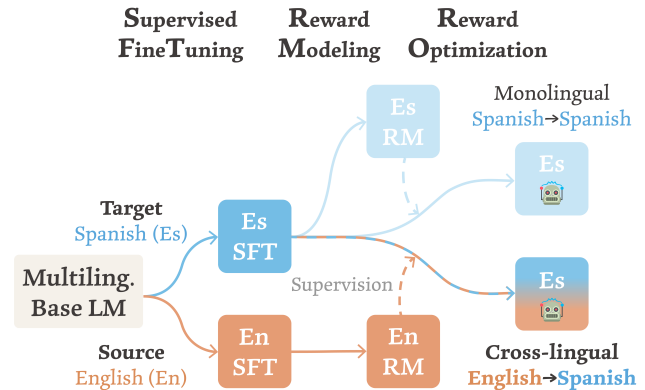
Figure 1. Cross-lingual reward model (RM) transfer. To align in a target language (in this example, Spanish), common monolingual alignment uses a RM for that target language. Instead, we re-purpose a RM for a different source language (in this example, English).

We propose a novel reward model (RM) transfer setup, where we re-purpose a trained RM for some source language to align a LM in a target language (Figure 1), and investigate the effectiveness of this simple recipe. Across two tasks (summarization and open-ended dialog generation), two reward optimization methods (reinforcement learning and best-of-n reranking), and various evaluation settings, we demonstrate substantial and consistent zero-shot cross-lingual utility of RMs. Surprisingly, alignment using a different-language RM sometimes outperforms using a same-language RM, both when judged by humans and LMs. We also show that our RM transfer framework is useful even when target-language data for supervised finetuning (SFT), another component in alignment, is inaccessible.

Our results show that RM signals are generalizable and robust to input distribution changes, which could be leveraged for more future applications. Practically, our findings pave the path towards lowering the costs for training LMs that more equitably serve users around the world.
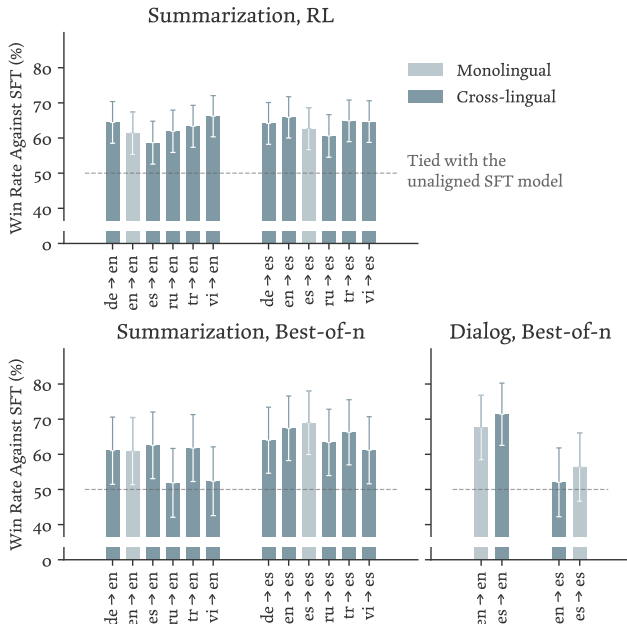
Figure 2. **Performing target-language alignment using a RM for a different source language improves performance, when evaluated exclusively in the target language. This improvement is sometimes even larger than using the target-language RM** (monolingual alignment). Here we measure the win rate against the target-language (unaligned) SFT model judged by humans, and the 95% confidence interval across validation instances. "source→target" denotes using a source-language RM to drive alignment in the target language.

2. Background: Alignment From Human Feedback

In addition to traditional unsupervised LM pretraining, many recent LMs also include an alignment phase to improve helpfulness, harmlessness, etc., supervised by human feedback (Bai et al., 2022a; Ouyang et al., 2022; *i.a.*). A common recipe includes three stages: supervised finetuning (SFT), reward modeling (RM), and reward optimization. We give an overview of each and refer readers to Ouyang et al. (2022) and Bai et al. (2022a) for details. We assume a **base** model already pretrained using a usually next-token prediction objective.

The SFT stage initializes from the base model and takes task inputs $x \in \mathcal{X}$ to train the model to simulate example outputs $y \in \mathcal{Y}$. Specifically, it optimizes the conditional log-likelihood of y given some input x, similar to regular language modeling. We denote the trained SFT model using π_{SFT}.

The RM stage trains a model $r : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ as a proxy for human-judged quality of y under x. It initializes from π_{SFT} and is trained using a dataset of human judgments of

generations. We consider two types of feedback to train the RM:

1. **Pointwise** feedback judges the quality of a single generation; in particular we only consider binary (good or bad) pointwise judgments. Denoting it as $z \in \{0, 1\}$ and letting $\mathcal{D}^{\text{RM}}$ be a dataset of judgments, the RM can be a standard classifier trained using the cross-entropy loss,

$$-\mathbb{E}_{(x,y,z) \sim \mathcal{D}^{\text{RM}}}[z \log \sigma \left(r(x, y) \right) + \\ (1 - z) \log \left(1 - \sigma \left(1 - r(x, y) \right) \right)].$$

2. **Pairwise** feedback chooses a better generation out of two. We denote the chosen one as y_w and the other as y_l. To train a pointwise RM on such data, the Bradley-Terry model (Bradley & Terry, 1952) is often used, maximizing

$$\mathbb{E}_{(x,y_w,y_l) \sim \mathcal{D}^{\text{RM}}}[\log \sigma \left(r(x, y_w) - r(x, y_l) \right)].$$

It is also generalizable to more than two outputs.

The reward optimization stage also initializes from π_{SFT} and further adjusts the model outputs using human feedback (as captured by the RM). Two common methods are reinforcement learning (RL) and best-of-n. Best-of-n is an inference-time procedure that does not change the underlying model, where multiple generations are sampled from π_{SFT} and then reranked using the RM; the highest-scoring generation is returned as the output. In RL, the model itself is changed such that its samples are scored highly by the RM, with the objective

$$\mathbb{E}_{x \sim \mathcal{D}^{\text{RO}}, \tilde{y} \sim \pi_\theta(x)}[r(x, \tilde{y}) - \\ \beta \left(\log \pi_\theta(\tilde{y} \mid x) - \log \pi_{\text{SFT}}(\tilde{y} \mid x) \right)].$$

$\mathcal{D}^{\text{RO}}$ is a dataset of inputs and β is a regularization hyperparameter. The above is typically optimized with PPO (Schulman et al., 2017). While we generally experiment with both methods, in some of our analyses we focus on best-of-n for a clean testbed without confounders from RL training.

3. Reward Model Transfer for Cross-Lingual Alignment

The pipeline in §2 is usually performed monolingually, commonly in English. Aligning for a new language requires both SFT data and RM data in that language. While the former may be relatively easier to obtain due to automatic construction methods, such as by re-purposing existing multilingual datasets (Muennighoff et al., 2023) or by eliciting from LMs (Wang et al., 2023c), RM data for a new language can be more expensive to gather, as it in principle requires human judgments. Additionally, RM data should ideally be periodically re-collected to avoid over-optimization (Bai et al., 2022a), further increasing data demand. Thus, we are mainly interested in alignment without target-language

RM data, though, in §5.3, we investigate dispensing with target-language SFT data too.

We propose to perform reward optimization using a RM trained for a different language (Figure 1). Intuitively, assuming model generation quality transfers cross-lingually (e.g., good English generations are still good when translated into Spanish[1]), a model that can judge the output quality in one language should generalize to others, as long as the RM understands the languages, which is enabled by multilingual *base* model training. This generalizability is often observed for other tasks in the zero-shot cross-lingual transfer literature (Wu & Dredze, 2019; Pires et al., 2019; Conneau et al., 2020b; Hu et al., 2020; *i.a.*), and we expect it to work for RMs too. A simple baseline would be to use automatically translated RM data, to which we compare in §5.1. In this paper, we use **source language** to denote the RM language, and **target language** for the language of the aligned model.

4. Experimental Setup

We consider two tasks: summarization, common in alignment research (Stiennon et al., 2020; Ziegler et al., 2020; Lee et al., 2023; *i.a.*), and open-ended dialog generation, with substantial real-world relevance. §A describes dataset details and §B training details. §G.1 contains our task instructions.

Summarization. The Seahorse dataset (Clark et al., 2023) contains documents and summaries in six languages (German, English, Spanish, Russian, Turkish, and Vietnamese) with pointwise human ratings which we use. For SFT, we gather the data sources of Seahorse: XSum (Narayan et al., 2018), XL-Sum (Hasan et al., 2021), MLSum (Scialom et al., 2020), and WikiLingua (Ladhak et al., 2020). We use mT5-XL (Xue et al., 2021) as our multilingual base model, with 3.7B parameters.

Open-Ended Dialog Generation. We use the OpenAssistant dataset (Köpf et al., 2023) with multilingual, pairwise human-rated chat transcripts. For the SFT data, we use the human-preferred response in each pair to finetune the model. Many languages in OpenAssistant have only limited data, so we only consider three languages with the most amounts of data: English, Spanish, and Russian. We use PaLM-2-XXS as the base model (Anil et al., 2023). The authors of OpenAssistant found RL to be ineffective for this dataset (Köpf et al., 2023), which we confirmed in our experiments (Figure 4). We therefore focus on best-of-n for

this task.

Evaluation. We assess model quality across several settings. First, we use the target-language RM, which is by design finetuned to judge target-language generation quality. But because of potential RM biases (Gao et al., 2023; Coste et al., 2023; Eisenstein et al., 2023), we also include two zero-shot-prompted evaluation models with much larger backbones—GPT-4 (OpenAI, 2023) and PaLM-2-L (Anil et al., 2023). This latter evaluation setup is common in prior work and has been demonstrated to correlate well with human judgments (Lee et al., 2023; Rafailov et al., 2023; An et al., 2023; Mu et al., 2023; *i.a.*). We also confirm its validity in §5.1 and §C. Finally, we also perform human evaluations by self-reported native or advanced speakers, though only for a subset of language pairs and 250 (RL) / 100 (best-of-n) instances per pair due to its cost. For both human and LM evaluation, we elicit pairwise judgments to compare responses from the aligned model and the SFT model (Bai et al., 2022b; Lee et al., 2023; *i.a.*). We measure the *win rate*, i.e., how often the judge prefers the former. A 50% win rate indicates no improvement from alignment. §G.2 includes more details such as the evaluation prompts and positional bias control.

5. Results

Here we report the results of cross-lingual alignment. See §H for numerical results that correspond to the plots in this section.

5.1. Cross-Lingual Alignment Is Effective

When evaluated by the finetuned target-language RM, Figure 3 shows that monolingual best-of-n or RL always improves model quality, as expected. Encouragingly, cross-lingual reward optimization improves over the SFT model in all cases too. Similarly, when judged by a general-purpose LM, PaLM-2-L in Figure 4 and GPT-4 in §D, in-language and cross-lingual reward optimization both generally improve model quality. Importantly, we observe high agreement between the two LMs: on an instance level, they agree >70% across setups (see §D); if we consider how often they agree in the relative ranking of two source languages, they agree 78% for summarization (both best-of-n and RL) and 100% for dialog generation (best-of-n). This indicates the reliability of a LM judge.

Human evaluation (Figure 2) reveals the same trend, though with larger confidence intervals due to the cost. It also justifies LM-based evaluation. For summarization, PaLM-2-L (GPT-4) agrees with humans 65% (69%) of the time in English and 66% (62%) in Spanish, matching the 63% human-human agreement for English reference summaries and 67% for Spanish in Seahorse (Clark, personal commu-

---

[1]We believe this is a weak assumption, though for tasks and instances more subject to culture-specific factors, generations may be judged more differently across languages (Costa et al., 2014; Hershcovich et al., 2022; Shwartz, 2022).

nication, April 15, 2024). For dialog, PaLM-2-L (GPT-4) agrees with humans 69% (59%) of the time in English and 62% (60%) in Spanish, again similar to the 63% human-human agreement in Bai et al. (2022a) and 66% in Dubois et al. (2024). With further evidence in §C, we believe our LM judges reasonably reflect output quality.

We also compare our cross-lingual transfer setup to an alternative strategy, sometimes dubbed "translate-train" (Conneau et al., 2018; *i.a.*), that first trains a silver target-language RM by automatically translating the source-language data and then using the silver RM for target-language alignment. Averaged across all 30 ($= 6^2 - 6$) cross-lingual language pairs, under best-of-n and judged by PaLM-2-L, our RM transfer strategy outperforms translate-train[2] (average win rate 58.8 vs. 57.5; see Table 6 and 17 for raw numbers). RM transfer also has an efficiency advantage: to align in multiple target languages, it suffices to train one source-language RM, rather than different ones for each target language. In §F, we also explore alignment using bilingual RMs with two source languages (Mulcaire et al., 2019), though without noticeable improvements.

5.2. Cross-Lingual Alignment Sometimes Outperforms Monolingual Alignment

Remarkably, cross-lingual reward optimization often yields an even better model than using the target-language RM. This is validated by (1) the consistent trend when evaluated by PaLM-2-L, GPT-4, and humans, (2) their instance-level and ranking-level agreement (§5.1), and (3) the small confidence intervals. This may be due to a regularization effect: the target-language RM may possess language-specific spurious artifacts, to which the target-language policy model overfits (Gao et al., 2023) more than artifacts in a different language in the source-language RM. Suppose, for example, that the target-language RM assigns higher rewards when the generation contains certain target-language words (due to bias in the RM training data). A different-language policy model is unlikely to exploit this, as it rarely generates these words, but a same-language policy model may.

This hypothesis is consistent with our observed patterns. First, there are many fewer cases of cross-lingual reward optimization outperforming the monolingual setting when measured by the finetuned target-language RM evaluator than the prompted LM evaluators (Figure 3): under this hypothesis, the finetuned evaluator RMs would be more susceptible to such artifacts and (incorrectly) assign higher scores in the monolingual settings. The underperformance of the translate-train baseline (§5.1) also provides weak evidence: in principle, a source-language RM and a source-translated-into-target-language RM should capture the same reward signal, as they are derived from the same data source, and
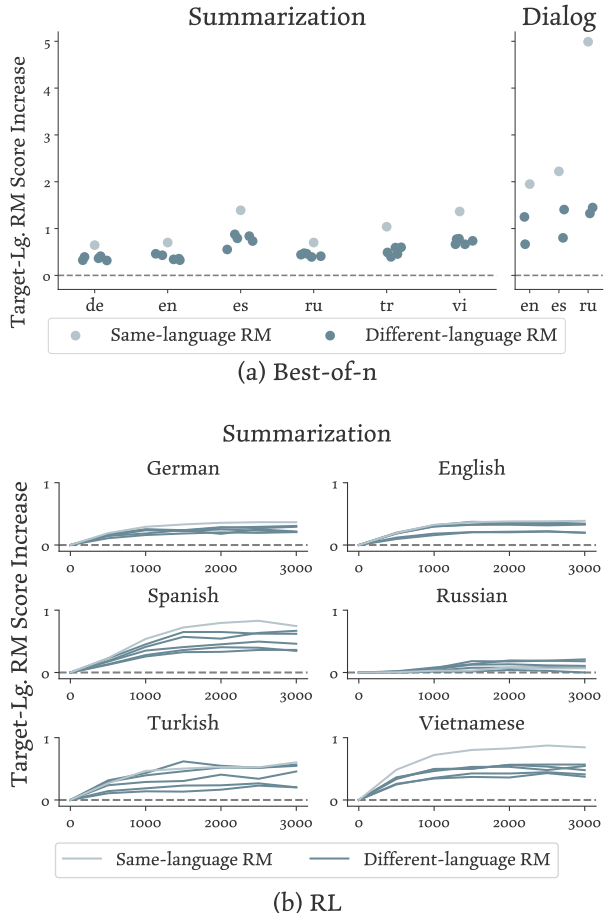
(a) Best-of-n

(b) RL

Figure 3. Cross-lingual alignment effectiveness judged by a finetuned target-language RM evaluator, measured in its score increase between the aligned model and the target-language SFT model. Each group in (a) and subplot in (b) represents one target language, and different dots/lines within each represent different source languages. RL is difficult to train for OpenAssistant (§4), so we omit it here. **In most cases, the RM evaluator score improves for cross-lingually aligned models.**

would lead to similar downstream performance. However, the former is less susceptible to reward over-optimization due to the language mismatch, leading to better performance, though this is admittedly confounded by translation quality.

Corroborating this hypothesis, we also find that when used monolingually, the RMs behave more like a bag-of-word (BoW) model. We take each of the 6 summarization RMs and infer on the validation set of each dataset in each language (Table 1). In every setting, we fit a BoW linear regressor to predict the RM-assigned score for each instance and compute the R^2 across instances as a proxy for the RM's similarity to a BoW model in that setting. For each dataset, and for every source language that differs from the dataset's language, we check whether inferring using the source-language RM or the dataset-language RM results in
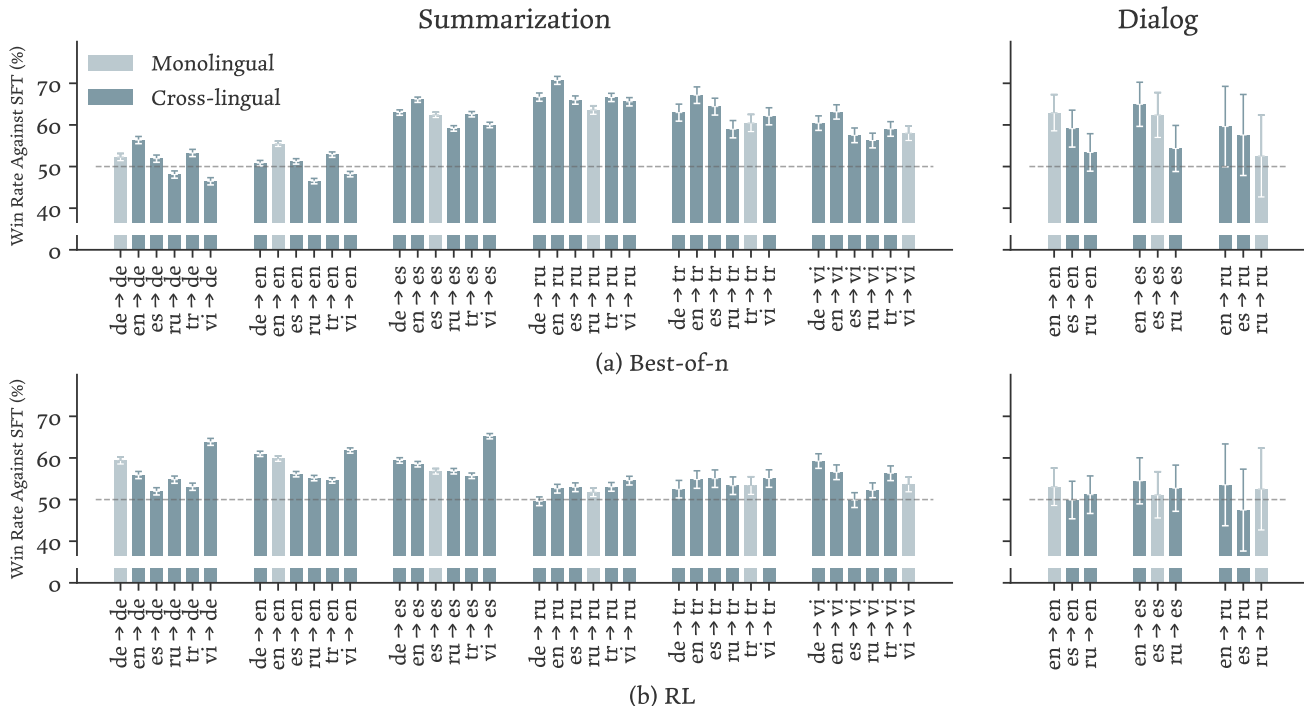
Figure 4. Alignment effectiveness, compared to the target-language SFT model judged by PaLM-2-L, and the 95% confidence interval across validation instances. "source→target" denotes a source-language RM driving alignment in the target language. **Cross-lingual alignment is generally effective, sometimes outperforming monolingual alignment**. RL is hard to train for OpenAssistant, in line with what its authors found ([Köpf et al., 2023](#)).

a larger R^2. The latter monolingual usage has a higher R^2 (0.65 vs. 0.63), so it is more likely that the RMs overfit to lexical patterns when used in-language.

5.3. Cross-Lingual Alignment Without Target-Language SFT Data

So far we assumed access to target-language SFT data since, as §3 argues, SFT data could be more easily obtained than RM data. We now relax this assumption and instead translate the source-language SFT data into the target language using Google Translate. We investigate if it, combined with RM transfer, still enables cross-lingual alignment. As a case study, we only consider summarization and when English is the source or target language.

Using translated SFT data substantially degrades the quality of the SFT model (Figure 5(a)) and the best-of-n-aligned LM (Figure 5(b)). There are however two factors: (1) quality loss due to translation, and (2) domain/style mismatch. For (2), we note that different languages have SFT data composed of different datasets, following Seahorse (Table 1).[3] And these datasets differ stylistically: for example, while XSum includes news articles, WikiLingua consists of how-

---

[3]SFT data quantity may also be a confounder, but we consider directions both from and to English, and the degradation is substantial in both. So quantity is not the biggest factor.

to articles and with formulaic summaries. There would thus be a domain difference between organic target-language SFT data vs. data translated from a different language.

To account for this, we employ round-trip back-translation, first translating the target-language SFT data into the source language and then back to the target language. This setup is not practically useful but it upper-bounds the effect of translation errors alone. In Figure 5(a), we see that this bridges most of the gap, in some cases leading to models that win over the SFT model >50% of the time. Alternatively, we artificially control for domain by repeating our experiments solely using WikiLingua for both SFT and RM as it is present for all languages. From Figure 5(c), the gap indeed reduces, with the translated SFT models sometimes even outperforming the original, and back-translation is no longer consistently beneficial.

Other than genre control, we also hypothesize that the gap would be smaller for RL than best-of-n because the RM, whose transferability we verified (§5), intuitively plays a bigger role in the RL pipeline. Best-of-n, on the other hand, is more reliant on the SFT model quality, as reflected by the high resemblance between the transfer performance patterns in Figure 5(b) and the SFT model quality in Figure 5(a). Figure 5(d) indeed shows that the translated models have little performance drop, except for cases where the former
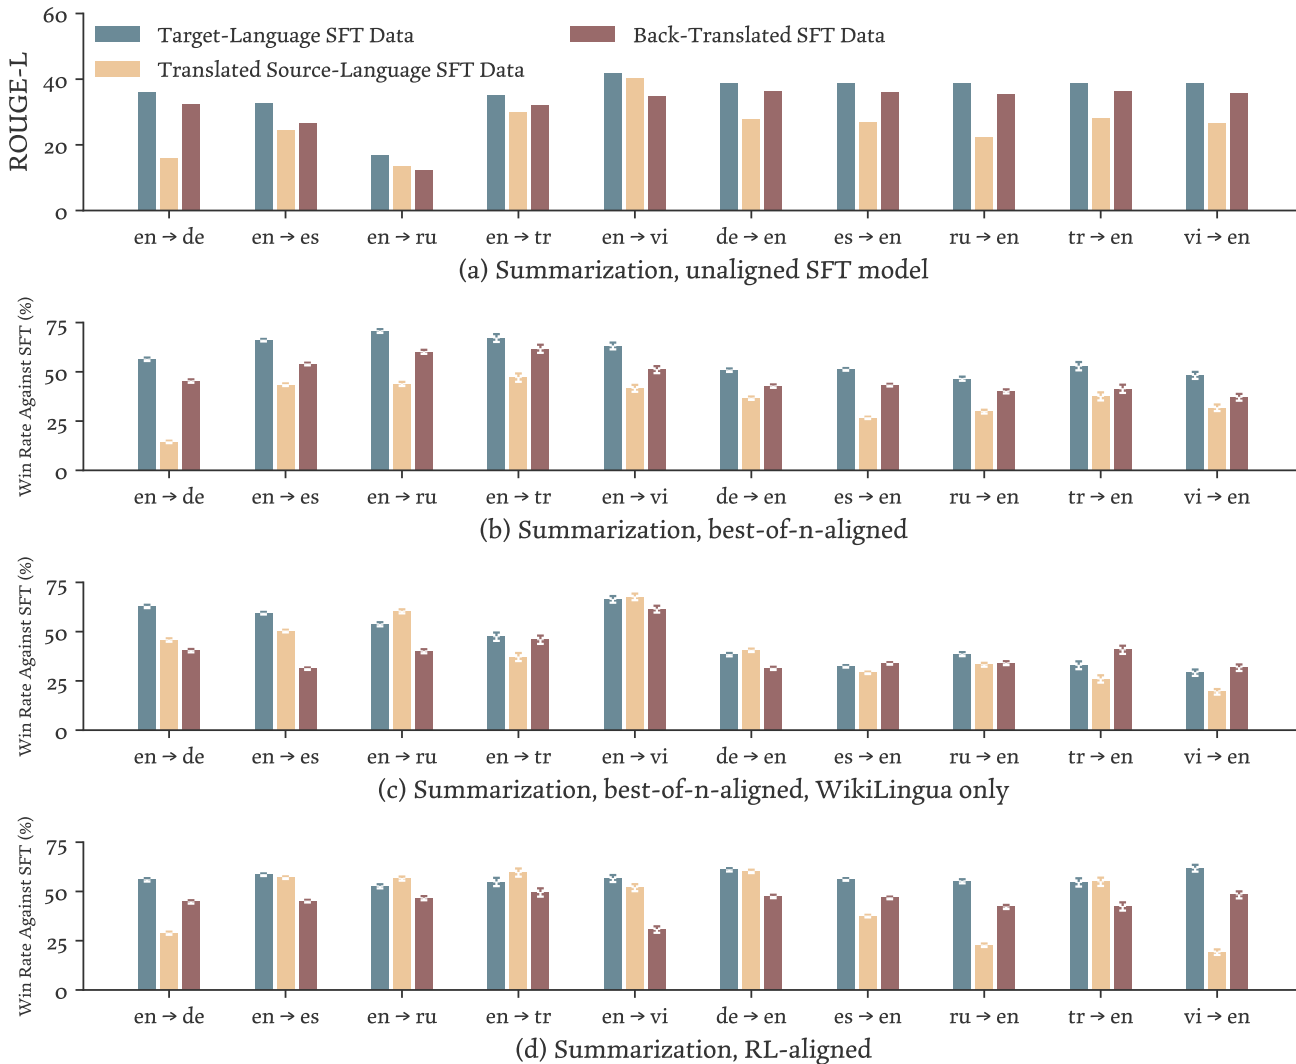
Figure 5. Cross-lingual alignment results without target-language SFT data using various strategies and on different data. **Training the SFT model using data translated from another language can be helpful when aligning using RL ((d)), but domain match is important for best-of-n ((c)) and the back-translation results).**

degenerates.[4] Again, apart from the degenerate cases, back-translation is not helpful.

To summarize,[5] cross-lingual alignment could still be helpful even without target-language SFT data, though care needs to be taken when training the surrogate SFT model. While we only experimented on summarization, we believe there will be larger text diversity for dialog generation in the wild, for which this issue warrants greater attention.

---

[4]Which we believe is due to a lack of careful case-by-case hyperparameter tuning, which we did not perform as it would be very expensive to tune for each transfer pair.

[5]No pun intended.

5.4. Practical Recommendations

Our findings suggest that, for SFT, it is always beneficial to use organic target-language data, but when inaccessible, automatic translation may be a remedy, though one should be mindful of the data distribution match between the data source and the application, or relying more on RL.

For RM, cross-lingual transfer is often successful, but how does one select the source RM language to align in a new target language? In Figure 6, we show the source languages ranked by transfer effectiveness for each target language. The rankings across target languages are generally stable, especially for best-of-n: source languages effective for one target language are usually effective for others too. Thus, one may select the source language by extrapolating from its performance on other target languages. In practice, English
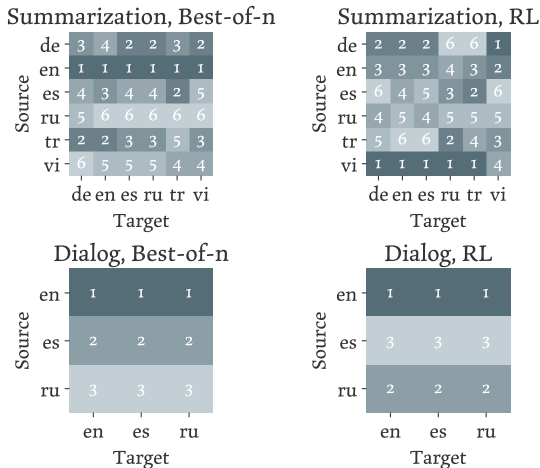
Figure 6. PaLM-2-L-judged rankings of source language effectiveness when driving alignment in different target languages. English is generally a good source.

RMs are usually the most accessible. Our results show that it is a decent strategy to use them as the source: English is often a highly-ranked source language, most frequently the best, perhaps due to the relatively higher annotator quantity and quality (Yu et al., 2022) or implicit modeling assumptions (Dyer et al., 2019). Beyond this empirical observation, we try to causally predict the pairwise transferability from various features in §6, but without success.

6. Analysis

The effectiveness of cross-lingual alignment motivates us to better understand how it relates to various factors. We show that while RM generalizability within the original reward modeling task is a prerequisite, it does not uniquely explain the downstream success. Similarly, we also show that the pairwise win rates (judged by PaLM-2-L unless otherwise mentioned) cannot be fully explained by, and thereby not predictable from, language features or the KL-divergence from the SFT model.

6.1. Impact of RM Generalizability Within Reward Modeling

The RMs' cross-lingual utility in downstream alignment is predicated on their generalizability within the original reward modeling task, but the latter is not sufficient for the former. So how much does this generalizability explain the alignment success? We analyze this generalizability following the cross-lingual transfer tradition, zero-shot applying a source-language RM to the target-language validation data and computing accuracy (Wu & Dredze, 2019; 2020; Pires et al., 2019; *i.a.*). We also consider a majority baseline and a length baseline to check if the RMs are only superficially capturing generation length (Wang et al., 2023b;

Singhal et al., 2023). To compute this length baseline: for dialog generation, a pairwise task, all longer, or shorter, responses in each pair are chosen, depending on which (long or short) yields higher training set accuracy. For summarization, a pointwise task, all responses longer (or shorter) than a threshold are chosen. The direction (long or short) and the threshold are also selected using the training set.

Figure 7 confirms cross-lingual RM generalizability: cross-lingual RMs often perform above the majority baseline for summarization and random performance (50%) for dialog. §E verifies this generalizability with another setup.

Nevertheless, the improvements over the majority/random baselines are modest. The dialog models even sometimes underperform the length baseline (though this does not mean the RMs only rely on length[6]). Part of this is due to the high subjectivity of the reward modeling task: the RM accuracies here are near the human agreement level for Seahorse (Clark et al., 2023), plotted in Figure 7, and generally match the human agreement numbers in dialog generation work (Bai et al., 2022a; Dubois et al., 2024). But it is still interesting that seemingly weak RMs, like the Vietnamese RM which performs similarly to the majority baseline when used monolingually or the dialog RMs which are often surpassed by the length baseline, can achieve high cross-lingual alignment effectiveness (Figure 4).

Furthermore, the results here do not match their downstream utility, regardless of whether we consider the quality of the RMs as measured by their in-language validation accuracy (Turkish, for example, is the best in Figure 7, but not so in Figure 6), the generalizability of the RMs which we operationalize as the difference between in-language training and validation loss (or accuracy—they yield the same ranking: Russian, German, English, Turkish, Vietnamese, and Spanish, from the least amount of overfitting to the most, again different from Figure 6), or the specific pairwise transfer effectiveness (for each target language, we compare the effectiveness of source languages ranked by the reward modeling task generalizability here vs. by downstream alignment win rate; on summarization, averaged across target languages, Kendall's τ=0.1 (same when using best-of-n or RL), indicating low ranking agreement). Overall, while cross-lingual alignment depends on RM generalizability on the original task, there are other factors at play too.

6.2. Impact of Language Features

Can the cross-lingual alignment performance be predicted from simple language features, such as their frequency in the pretraining corpus or typological similarity? The summa-

---

[6]The RMs agree with the length baseline on 72.6% of the validation instances, higher than the baseline agreement level of 56.6% (how often two random models at their accuracy levels agree on average), but far from full agreement.
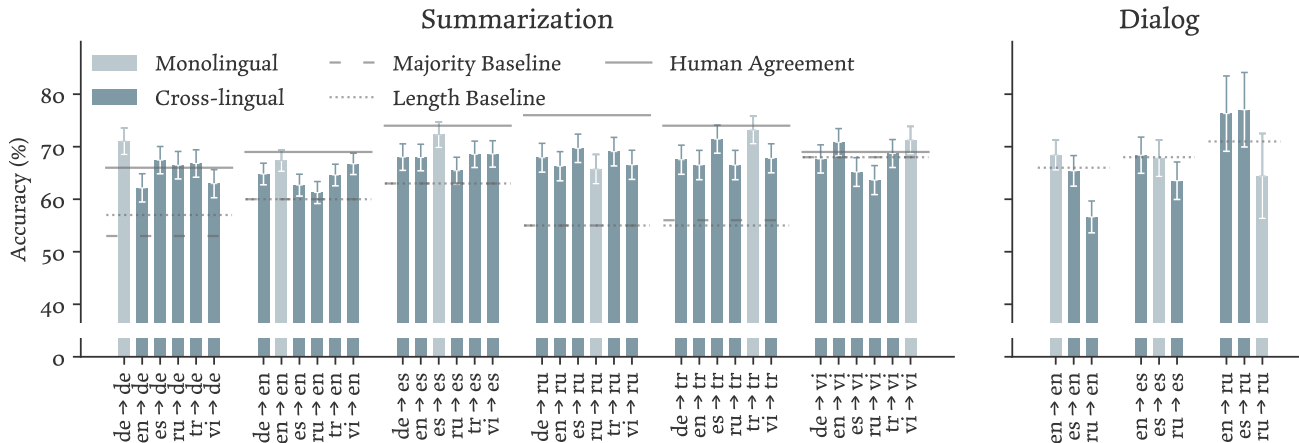
Figure 7. Source-language RM generalizability within the original reward modeling task and the 95% confidence interval across validation instances. "source→target" denotes training a source-language RM and measuring its accuracy on the target language validation data. The baselines are explained in §6.1. Dialog generation, a pairwise task, does not have a majority baseline; the dataset authors also did not report human agreement. **RMs generally exhibit cross-lingual generalizability, exceeding the majority baseline and often the length baseline**.

rization languages ranked by frequency in the mT5 corpus, the base model for this task, are: English, Russian, Spanish, German, Turkish, Vietnamese (Xue et al., 2021). This does not match the transfer utility ranking in Figure 6. Similarly, neither does the ranking match the SFT data quantity or RM data quantity (in §A).

Linguistic typology and orthography are also common predictors of cross-lingual transferability (Gerz et al., 2018; K et al., 2020; Muller et al., 2021; *i.a.*). This, however, is not the case for us either: for summarization RL, for example, English benefits from Vietnamese the most, but they belong to disparate language families. Orthography may be playing a role: Russian overall does not transfer well to other languages, and it is the only language that does not use the Latin script, but this trend is not clear. Systematically, we compute the correlation between alignment utility and WALS features of linguistic typology (Dryer & Haspelmath, 2013). For each WALS feature present for all 6 summarization languages, we divide all win rates into two groups: those between language pairs that have the same, or different, feature values. Under a one-sided unpaired t-test, no feature shows statistical significance at $\alpha = 0.05$ with Bonferroni correction (Dunn, 1961).[7] Therefore, alignment utility does not strongly correlate with such language features.
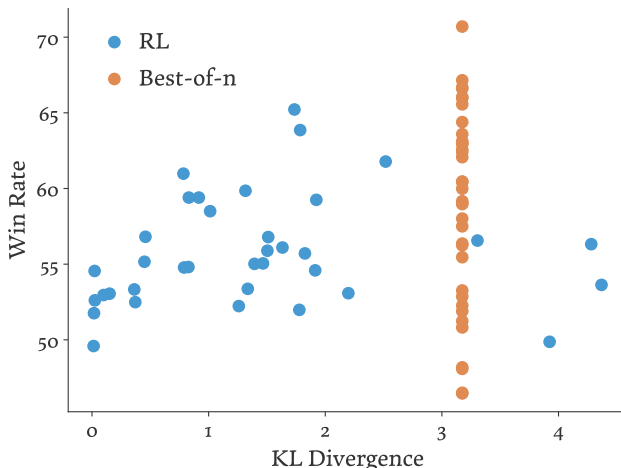
Figure 8. Win rate (PaLM-2-L-judged) vs. KL-divergence for summarization across different (source, target) language pairs. For best-of-n, we use the upper bound formula in Stiennon et al. (2020), Beirami et al. (2024), *i.a.*, which is a function of n and thus appears as a vertical line. **KL-divergence does not fully explain the final alignment performance**.

6.3. Impact of Policy Divergence

From a learning angle, it has been shown that the reward that a learned policy can obtain strongly correlates with its KL-divergence from the base (SFT) policy (Bai et al., 2022a). This could be concerning, if the model deviates from the base policy to "hack" the reward (Gao et al., 2023; Coste et al., 2023; Eisenstein et al., 2023), but not if the eval-

---

[7]Even without correction, only 4 show statistical significance at $\alpha = 0.05$ out of 123: 1A, 3A, 37A, and 54A. The first two are phonological features, and the other two minor syntactic features, thus likely being spurious correlations.

uation metric is robust. As we perform human evaluation and also verified that our LM judges correlate with human judgments, this is less of a problem for us. Nevertheless, in Figure 8, we plot the correlation between the win rates and the KL-divergence of the aligned models. There is not a clear correlation, and hence we do not observe reward over-optimization.

7. Related Work

Zero-shot cross-lingual transfer. There is a long line of research on cross-lingual representation generalizability, such as with sentence embeddings (Conneau et al., 2018) or more recently, LMs (Wu & Dredze, 2019; 2020; Pires et al., 2019; Siddhant et al., 2020). Commonly, a multilingual LM (Devlin et al., 2019; Conneau & Lample, 2019; Conneau et al., 2020a; *i.a.*) is finetuned on a task in a source language and evaluated on the task's test set in a different language. This is generally effective. Our RM transfer setup can be viewed under this framework, but we go further and show that this generalizability is useful for *downstream* tasks, in our case alignment. Shaham et al. (2024) and Chirkova & Nikoulina (2024) are close to us in studying cross-lingual generalizability in alignment, but only focusing on SFT and only using translated data.

Multilingual Alignment. For SFT, it is common to assemble existing multilingual task datasets into instruction datasets (Muennighoff et al., 2023; Asai et al., 2023; Ahuja et al., 2023). Some have directly collected SFT data for non-English languages, either on a per-language basis (Zhang et al., 2023; Xu et al., 2023; Ni et al., 2023; *i.a.*) or multilingually (Zhao et al., 2024; Singh et al., 2024), though this can be expensive. Past work has also used automatic translation for SFT (Li et al., 2023a; Lai et al., 2023; Shaham et al., 2024; *i.a.*) and RM data (Lai et al., 2023; Shen et al., 2024). We also use translation for SFT, but showed that cross-lingual transfer outperforms translation for RM.

8. Conclusion

We showed through two different tasks that we can perform alignment using a different-language RM. Surprisingly, we find this to be sometimes more effective than using a same-language RM. We also identified issues and remedies when we dispense with target-language SFT data. We hope our findings can motivate future work to build better LMs for more languages. Adapting our RM transfer setup to other settings such as domain generalization would also be exciting future directions.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Ahuja, K., Diddee, H., Hada, R., Ochieng, M., Ramesh, K., Jain, P., Nambi, A., Ganu, T., Segal, S., Ahmed, M., Bali, K., and Sitaram, S. MEGA: Multilingual evaluation of generative AI. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 4232–4267, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023. emnlp-main.258. URL https://aclanthology.org/2023.emnlp-main.258.

An, C., Gong, S., Zhong, M., Zhao, X., Li, M., Zhang, J., Kong, L., and Qiu, X. L-Eval: Instituting standardized evaluation for long context language models, 2023.

Anil, R., Dai, A. M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., Chu, E., Clark, J. H., Shafey, L. E., Huang, Y., Meier-Hellstern, K., Mishra, G., Moreira, E., Omernick, M., Robinson, K., Ruder, S., Tay, Y., Xiao, K., Xu, Y., Zhang, Y., Abrego, G. H., Ahn, J., Austin, J., Barham, P., Botha, J., Bradbury, J., Brahma, S., Brooks, K., Catasta, M., Cheng, Y., Cherry, C., Choquette-Choo, C. A., Chowdhery, A., Crepy, C., Dave, S., Dehghani, M., Dev, S., Devlin, J., Díaz, M., Du, N., Dyer, E., Feinberg, V., Feng, F., Fienber, V., Freitag, M., Garcia, X., Gehrmann, S., Gonzalez, L., Gur-Ari, G., Hand, S., Hashemi, H., Hou, L., Howland, J., Hu, A., Hui, J., Hurwitz, J., Isard, M., Ittycheriah, A., Jagielski, M., Jia, W., Kenealy, K., Krikun, M., Kudugunta, S., Lan, C., Lee, K., Lee, B., Li, E., Li, M., Li, W., Li, Y., Li, J., Lim, H., Lin, H., Liu, Z., Liu, F., Maggioni, M., Mahendru, A., Maynez, J., Misra, V., Moussalem, M., Nado, Z., Nham, J., Ni, E., Nystrom, A., Parrish, A., Pellat, M., Polacek, M., Polozov, A., Pope,

R., Qiao, S., Reif, E., Richter, B., Riley, P., Ros, A. C., Roy, A., Saeta, B., Samuel, R., Shelby, R., Slone, A., Smilkov, D., So, D. R., Sohn, D., Tokumine, S., Valter, D., Vasudevan, V., Vodrahalli, K., Wang, X., Wang, P., Wang, Z., Wang, T., Wieting, J., Wu, Y., Xu, K., Xu, Y., Xue, L., Yin, P., Yu, J., Zhang, Q., Zheng, S., Zheng, C., Zhou, W., Zhou, D., Petrov, S., and Wu, Y. PaLM 2 technical report, 2023.

Asai, A., Kudugunta, S., Yu, X. V., Blevins, T., Gonen, H., Reid, M., Tsvetkov, Y., Ruder, S., and Hajishirzi, H. BUFFET: Benchmarking large language models for few-shot cross-lingual transfer, 2023.

Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., and Kaplan, J. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022a.

Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C., Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., Kerr, J., Mueller, J., Ladish, J., Landau, J., Ndousse, K., Lukosuite, K., Lovitt, L., Sellitto, M., Elhage, N., Schiefer, N., Mercado, N., DasSarma, N., Lasenby, R., Larson, R., Ringer, S., Johnston, S., Kravec, S., Showk, S. E., Fort, S., Lanham, T., Telleen-Lawton, T., Conerly, T., Henighan, T., Hume, T., Bowman, S. R., Hatfield-Dodds, Z., Mann, B., Amodei, D., Joseph, N., McCandlish, S., Brown, T., and Kaplan, J. Constitutional AI: Harmlessness from AI feedback, 2022b.

Beirami, A., Agarwal, A., Berant, J., D'Amour, A., Eisenstein, J., Nagpal, C., and Suresh, A. T. Theoretical guarantees on the best-of-n alignment policy, 2024.

Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952. ISSN 00063444. URL http://www.jstor.org/stable/2334029.

Chirkova, N. and Nikoulina, V. Zero-shot cross-lingual transfer in instruction tuning of large language model, 2024.

Christodoulopoulos, C. and Steedman, M. A massively parallel corpus: the Bible in 100 languages. *Language Resources and Evaluation*, 49(2):375–395, Jun 2015. ISSN 1574-0218. doi: 10.1007/s10579-014-9287-y. URL https://doi.org/10.1007/s10579-014-9287-y.

Clark, E., Rijhwani, S., Gehrmann, S., Maynez, J., Aharoni, R., Nikolaev, V., Sellam, T., Siddhant, A., Das, D., and Parikh, A. SEAHORSE: A multilingual, multifaceted dataset for summarization evaluation. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 9397–9413, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.584. URL https://aclanthology.org/2023.emnlp-main.584.

Conneau, A. and Lample, G. Cross-lingual language model pretraining. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/c04c19c2c2474dbf5f7ac4372c5b9af1-Paper.pdf.

Conneau, A., Rinott, R., Lample, G., Williams, A., Bowman, S., Schwenk, H., and Stoyanov, V. XNLI: Evaluating cross-lingual sentence representations. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2475–2485, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1269. URL https://aclanthology.org/D18-1269.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. Unsupervised cross-lingual representation learning at scale. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8440–8451, Online, July 2020a. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.747. URL https://aclanthology.org/2020.acl-main.747.

Conneau, A., Wu, S., Li, H., Zettlemoyer, L., and Stoyanov, V. Emerging cross-lingual structure in pretrained language models. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6022–6034, Online, July 2020b. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.536. URL https://aclanthology.org/2020.acl-main.536.

Costa, A., Foucart, A., Hayakawa, S., Aparici, M., Apesteguia, J., Heafner, J., and Keysar, B. Your morals depend on language. *PLOS ONE*, 9(4):1–7, 04 2014. doi: 10.1371/journal.pone.0094842. URL https://doi.org/10.1371/journal.pone.0094842.

Coste, T., Anwar, U., Kirk, R., and Krueger, D. Reward model ensembles help mitigate overoptimization, 2023.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL https://aclanthology.org/N19-1423.

Dryer, M. S. and Haspelmath, M. (eds.). *WALS Online (v2020.3)*. Zenodo, 2013. doi: 10.5281/zenodo.7385533. URL https://doi.org/10.5281/zenodo.7385533.

Dubois, Y., Li, X., Taori, R., Zhang, T., Gulrajani, I., Ba, J., Guestrin, C., Liang, P., and Hashimoto, T. B. Alpacafarm: A simulation framework for methods that learn from human feedback, 2024.

Dunn, O. J. Multiple comparisons among means. *Journal of the American Statistical Association*, 56(293):52–64, 1961. ISSN 01621459. URL http://www.jstor.org/stable/2282330.

Dyer, C., Melis, G., and Blunsom, P. A critical analysis of biased parsers in unsupervised parsing, 2019.

Eisenstein, J., Nagpal, C., Agarwal, A., Beirami, A., D'Amour, A., Dvijotham, D., Fisch, A., Heller, K., Pfohl, S., Ramachandran, D., Shaw, P., and Berant, J. Helping or herding? Reward model ensembles mitigate but do not eliminate reward hacking, 2023.

Gao, L., Schulman, J., and Hilton, J. Scaling laws for reward model overoptimization. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 10835–10866. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/gao23h.html.

Gehrmann, S., Adewumi, T., Aggarwal, K., Ammanamanchi, P. S., Aremu, A., Bosselut, A., Chandu, K. R., Clinciu, M.-A., Das, D., Dhole, K., Du, W., Durmus, E., Dušek, O., Emezue, C. C., Gangal, V., Garbacea, C., Hashimoto, T., Hou, Y., Jernite, Y., Jhamtani, H., Ji, Y., Jolly, S., Kale, M., Kumar, D., Ladhak, F., Madaan, A., Maddela, M., Mahajan, K., Mahamood, S., Majumder, B. P., Martins, P. H., McMillan-Major, A., Mille, S., van Miltenburg, E., Nadeem, M., Narayan, S., Nikolaev, V., Niyongabo Rubungo, A., Osei, S., Parikh, A., Perez-Beltrachini, L., Rao, N. R., Raunak, V., Rodriguez, J. D.,

Santhanam, S., Sedoc, J., Sellam, T., Shaikh, S., Shimorina, A., Sobrevilla Cabezudo, M. A., Strobelt, H., Subramani, N., Xu, W., Yang, D., Yerukola, A., and Zhou, J. The GEM benchmark: Natural language generation, its evaluation and metrics. In Bosselut, A., Durmus, E., Gangal, V. P., Gehrmann, S., Jernite, Y., Perez-Beltrachini, L., Shaikh, S., and Xu, W. (eds.), *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pp. 96–120, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.gem-1.10. URL https://aclanthology.org/2021.gem-1.10.

Gerz, D., Vulić, I., Ponti, E. M., Reichart, R., and Korhonen, A. On the relation between linguistic typology and (limitations of) multilingual language modeling. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 316–327, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1029. URL https://aclanthology.org/D18-1029.

Hasan, T., Bhattacharjee, A., Islam, M. S., Mubasshir, K., Li, Y.-F., Kang, Y.-B., Rahman, M. S., and Shahriyar, R. XL-sum: Large-scale multilingual abstractive summarization for 44 languages. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4693–4703, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.413. URL https://aclanthology.org/2021.findings-acl.413.

Hershcovich, D., Frank, S., Lent, H., de Lhoneux, M., Abdou, M., Brandl, S., Bugliarello, E., Cabello Piqueras, L., Chalkidis, I., Cui, R., Fierro, C., Margatina, K., Rust, P., and Søgaard, A. Challenges and strategies in cross-cultural NLP. In Muresan, S., Nakov, P., and Villavicencio, A. (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6997–7013, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.482. URL https://aclanthology.org/2022.acl-long.482.

Hu, J., Ruder, S., Siddhant, A., Neubig, G., Firat, O., and Johnson, M. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4411–4421. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/hu20b.html.

Joshi, P., Santy, S., Budhiraja, A., Bali, K., and Choudhury, M. The state and fate of linguistic diversity and inclusion

in the NLP world. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6282–6293, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.560. URL https://aclanthology.org/2020.acl-main.560.

K, K., Wang, Z., Mayhew, S., and Roth, D. Cross-lingual ability of multilingual BERT: An empirical study. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=HJeT3yrtDr.

Köpf, A., Kilcher, Y., von Rütte, D., Anagnostidis, S., Tam, Z.-R., Stevens, K., Barhoum, A., Duc, N. M., Stanley, O., Nagyfi, R., ES, S., Suri, S., Glushkov, D., Dantuluri, A., Maguire, A., Schuhmann, C., Nguyen, H., and Mattick, A. OpenAssistant conversations – democratizing large language model alignment, 2023.

Ladhak, F., Durmus, E., Cardie, C., and McKeown, K. Wiki-Lingua: A new benchmark dataset for cross-lingual abstractive summarization. In Cohn, T., He, Y., and Liu, Y. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4034–4048, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.360. URL https://aclanthology.org/2020.findings-emnlp.360.

Lai, V., Nguyen, C., Ngo, N., Nguyen, T., Dernoncourt, F., Rossi, R., and Nguyen, T. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In Feng, Y. and Lefever, E. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 318–327, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-demo.28. URL https://aclanthology.org/2023.emnlp-demo.28.

Lee, H., Phatale, S., Mansoor, H., Mesnard, T., Ferret, J., Lu, K., Bishop, C., Hall, E., Carbune, V., Rastogi, A., and Prakash, S. RLAIF: Scaling reinforcement learning from human feedback with ai feedback, 2023.

Li, H., Koto, F., Wu, M., Aji, A. F., and Baldwin, T. Bactrian-X: Multilingual replicable instruction-following models with low-rank adaptation, 2023a.

Li, X., Zhang, T., Dubois, Y., Taori, R., Gulrajani, I., Guestrin, C., Liang, P., and Hashimoto, T. B. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval, 2023b.

Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL https://aclanthology.org/W04-1013.

Mu, J., Li, X. L., and Goodman, N. Learning to compress prompts with gist tokens, 2023.

Muennighoff, N., Wang, T., Sutawika, L., Roberts, A., Biderman, S., Le Scao, T., Bari, M. S., Shen, S., Yong, Z. X., Schoelkopf, H., Tang, X., Radev, D., Aji, A. F., Almubarak, K., Albanie, S., Alyafeai, Z., Webson, A., Raff, E., and Raffel, C. Crosslingual generalization through multitask finetuning. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15991–16111, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.891. URL https://aclanthology.org/2023.acl-long.891.

Mulcaire, P., Kasai, J., and Smith, N. A. Polyglot contextual representations improve crosslingual transfer. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3912–3918, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1392. URL https://aclanthology.org/N19-1392.

Muller, B., Anastasopoulos, A., Sagot, B., and Seddah, D. When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models. In Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y. (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 448–462, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.38. URL https://aclanthology.org/2021.naacl-main.38.

Narayan, S., Cohen, S. B., and Lapata, M. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1797–1807, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1206. URL https://aclanthology.org/D18-1206.

Ni, J., Xue, F., Deng, Y., Phang, J., Jain, K., Shah, M. H., Zheng, Z., and You, Y. Instruction in the wild: A user-based instruction dataset. https://github.com/XueFuzhao/InstructionWild, 2023.

OpenAI. GPT-4 technical report, 2023.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., and Lowe, R. Training language models to follow instructions with human feedback. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 27730–27744. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf.

Pezeshkpour, P. and Hruschka, E. Large language models sensitivity to the order of options in multiple-choice questions, 2023.

Pires, T., Schlinger, E., and Garrette, D. How multilingual is multilingual BERT? In Korhonen, A., Traum, D., and Màrquez, L. (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4996–5001, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1493. URL https://aclanthology.org/P19-1493.

Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=HPuSIXJaa9.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms, 2017.

Scialom, T., Dray, P.-A., Lamprier, S., Piwowarski, B., and Staiano, J. MLSUM: The multilingual summarization corpus. In Webber, B., Cohn, T., He, Y., and Liu, Y. (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8051–8067, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.647. URL https://aclanthology.org/2020.emnlp-main.647.

Shaham, U., Herzig, J., Aharoni, R., Szpektor, I., Tsarfaty, R., and Eyal, M. Multilingual instruction tuning with just a pinch of multilinguality, 2024.

Shazeer, N. and Stern, M. Adafactor: Adaptive learning rates with sublinear memory cost. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4596–4604. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/shazeer18a.html.

Shen, L., Tan, W., Chen, S., Chen, Y., Zhang, J., Xu, H., Zheng, B., Koehn, P., and Khashabi, D. The language barrier: Dissecting safety challenges of llms in multilingual contexts, 2024.

Shwartz, V. Good night at 4 pm?! Time expressions in different cultures. In Muresan, S., Nakov, P., and Villavicencio, A. (eds.), *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 2842–2853, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.224. URL https://aclanthology.org/2022.findings-acl.224.

Siddhant, A., Johnson, M., Tsai, H., Ari, N., Riesa, J., Bapna, A., Firat, O., and Raman, K. Evaluating the cross-lingual effectiveness of massively multilingual neural machine translation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8854–8861, Apr. 2020. doi: 10.1609/aaai.v34i05.6414. URL https://ojs.aaai.org/index.php/AAAI/article/view/6414.

Singh, S., Vargus, F., Dsouza, D., Karlsson, B. F., Mahendiran, A., Ko, W.-Y., Shandilya, H., Patel, J., Mataciunas, D., OMahony, L., Zhang, M., Hettiarachchi, R., Wilson, J., Machado, M., Moura, L. S., Krzemiński, D., Fadaei, H., Ergün, I., Okoh, I., Alaagib, A., Mudannayake, O., Alyafeai, Z., Chien, V. M., Ruder, S., Guthikonda, S., Alghamdi, E. A., Gehrmann, S., Muennighoff, N., Bartolo, M., Kreutzer, J., Üstün, A., Fadaee, M., and Hooker, S. Aya dataset: An open-access collection for multilingual instruction tuning, 2024.

Singhal, P., Goyal, T., Xu, J., and Durrett, G. A long way to go: Investigating length correlations in RLHF, 2023.

Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., and Christiano, P. F. Learning to summarize with human feedback. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 3008–3021. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf.

Wang, P., Li, L., Chen, L., Cai, Z., Zhu, D., Lin, B., Cao, Y., Liu, Q., Liu, T., and Sui, Z. Large language models are not fair evaluators, 2023a.

Wang, Y., Ivison, H., Dasigi, P., Hessel, J., Khot, T., Chandu, K., Wadden, D., MacMillan, K., Smith, N. A., Beltagy, I., and Hajishirzi, H. How far can camels go? Exploring the state of instruction tuning on open resources. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023b. URL https://openreview.net/forum?id=w4zZNC4ZaV.

Wang, Y., Kordi, Y., Mishra, S., Liu, A., Smith, N. A., Khashabi, D., and Hajishirzi, H. Self-instruct: Aligning language models with self-generated instructions. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13484–13508, Toronto, Canada, July 2023c. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.754. URL https://aclanthology.org/2023.acl-long.754.

Wu, S. and Dredze, M. Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 833–844, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1077. URL https://aclanthology.org/D19-1077.

Wu, S. and Dredze, M. Are all languages created equal in multilingual BERT? In Gella, S., Welbl, J., Rei, M., Petroni, F., Lewis, P., Strubell, E., Seo, M., and Hajishirzi, H. (eds.), *Proceedings of the 5th Workshop on Representation Learning for NLP*, pp. 120–130, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.repl4nlp-1.16. URL https://aclanthology.org/2020.repl4nlp-1.16.

Xu, G., Liu, J., Yan, M., Xu, H., Si, J., Zhou, Z., Yi, P., Gao, X., Sang, J., Zhang, R., Zhang, J., Peng, C., Huang, F., and Zhou, J. CValues: Measuring the values of chinese large language models from safety to responsibility, 2023.

Xue, L., Constant, N., Roberts, A., Kale, M., Al-Rfou, R., Siddhant, A., Barua, A., and Raffel, C. mT5: A massively multilingual pre-trained text-to-text transformer. In Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y. (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 483–498, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.41. URL https://aclanthology.org/2021.naacl-main.41.

Yu, X., Chatterjee, T., Asai, A., Hu, J., and Choi, E. Beyond counting datasets: A survey of multilingual dataset construction and necessary resources. In Goldberg, Y., Kozareva, Z., and Zhang, Y. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 3725–3743, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-emnlp.273. URL https://aclanthology.org/2022.findings-emnlp.273.

Zhang, G., Shi, Y., Liu, R., Yuan, R., Li, Y., Dong, S., Shu, Y., Li, Z., Wang, Z., Lin, C., Huang, W., and Fu, J. Chinese open instruction generalist: A preliminary release, 2023.

Zhao, W., Ren, X., Hessel, J., Cardie, C., Choi, Y., and Deng, Y. (InThe)WildChat: 570k ChatGPT interaction logs in the wild. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Bl8u7ZRlbM.

Zheng, C., Zhou, H., Meng, F., Zhou, J., and Huang, M. Large language models are not robust multiple choice selectors, 2023.

Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., and Irving, G. Fine-tuning language models from human preferences, 2020.

Table 1. Number of summarization instances for the SFT and reward optimization stages. The datasets are taken from the GEM release (Gehrmann et al., 2021) and with certain validation instances removed (§A).

		Train	Validation
German	MLSum	220748	8932
	WikiLingua	40839	3699
English	XSum	23206	642
	XL-Sum	306522	9690
	WikiLingua	99020	12021
Spanish	XL-Sum	38110	3170
	MLSum	259888	8374
	WikiLingua	79212	9730
Russian	XL-Sum	62243	5492
	WikiLingua	37028	3209
Turkish	XL-Sum	27176	1953
	WikiLingua	3148	194
Vietnamese	XL-Sum	32111	2341
	WikiLingua	13707	679

Table 2. Number of summarization instances for reward modeling.

	Train	Validation
German	8389	1250
English	14031	2071
Spanish	8741	1310
Russian	7679	1112
Turkish	7855	1096
Vietnamese	7844	1166

Table 3. Number of dialog generation instances for the SFT and reward optimization stages.

	Train	Validation
English	8898	472
Spanish	5681	311
Russian	1884	99

Table 4. Number of dialog generation instances for reward modeling.

	Train	Validation
English	22076	1026
Spanish	13714	699
Russian	2627	135

A. Dataset Details and Statistics

We report dataset statistics in Table 1, 2, 3, and 4. We reuse the SFT data for reward optimization (for both training and evaluation for RL, and for only evaluation for best-of-n since it does not have a training stage), but only the input x, without reference generations y.

The summarization SFT datasets, reported in Table 1, are the original data sources of Seahorse, which we take from the GEM release (Gehrmann et al., 2021). They are evenly mixed at the instance level for both SFT training and RL training. For evaluation of the aligned model, we macro-average the per-dataset metrics (e.g., win rate) for a language-level score. Because the Seahorse dataset was created using the validation and test instances of the original summarization datasets, to be clean, we exclude the Seahorse training instances from these splits when performing SFT and reward optimization. OpenAssistant does not have this issue and has clean split separations. The Seahorse summaries are human-rated along six axes, and we only use the sixth axis for our pointwise reward as it encapsulates previous axes (Clark et al., 2023). We limit the maximum length of model inputs to 1,024 tokens and outputs to 512 tokens. See also §G.1 for instructions we attach to the dataset instances during training and inference.

B. Training Details

SFT. The model is trained using Adafactor (Shazeer & Stern, 2018) with a constant learning rate at 10^{-3} for summarization and 10^{-5} for dialog generation, batch size 32, and dropout 0.1. We perform checkpoint selection using validation ROUGE-L score (Lin, 2004).

RM. The model is trained using Adafactor with a constant learning rate at 10^{-4} after 1,000 linear warm-up steps, batch size 32, and dropout 0.1. We perform checkpoint selection using validation loss.

RL. We use PPO for RL training with a constant learning rate at 10^{-4}, batch size 32, for 3,000 steps for summarization and 2,500 steps for dialog generation. The value model has 1,000 linear warm-up steps and we only start training the policy model after 2,000 steps elapse. We set the regularization coefficient at $\beta = 0.1$.

Best-of-n. We use $n = 64$.

C. LM Judge Accuracy on Ground-truth Reward Modeling Data

We verify the validity of using LM as a judge for our tasks by computing its accuracy on the validation splits of the RM datasets we used. We only consider PaLM-2-L as a case study. For OpenAssistant, a pairwise dataset, we simply

Table 5. The accuracy of evaluating the PaLM-2-L judge on the RM validation data. We also report the number of comparisons based on which the accuracy is calculated.

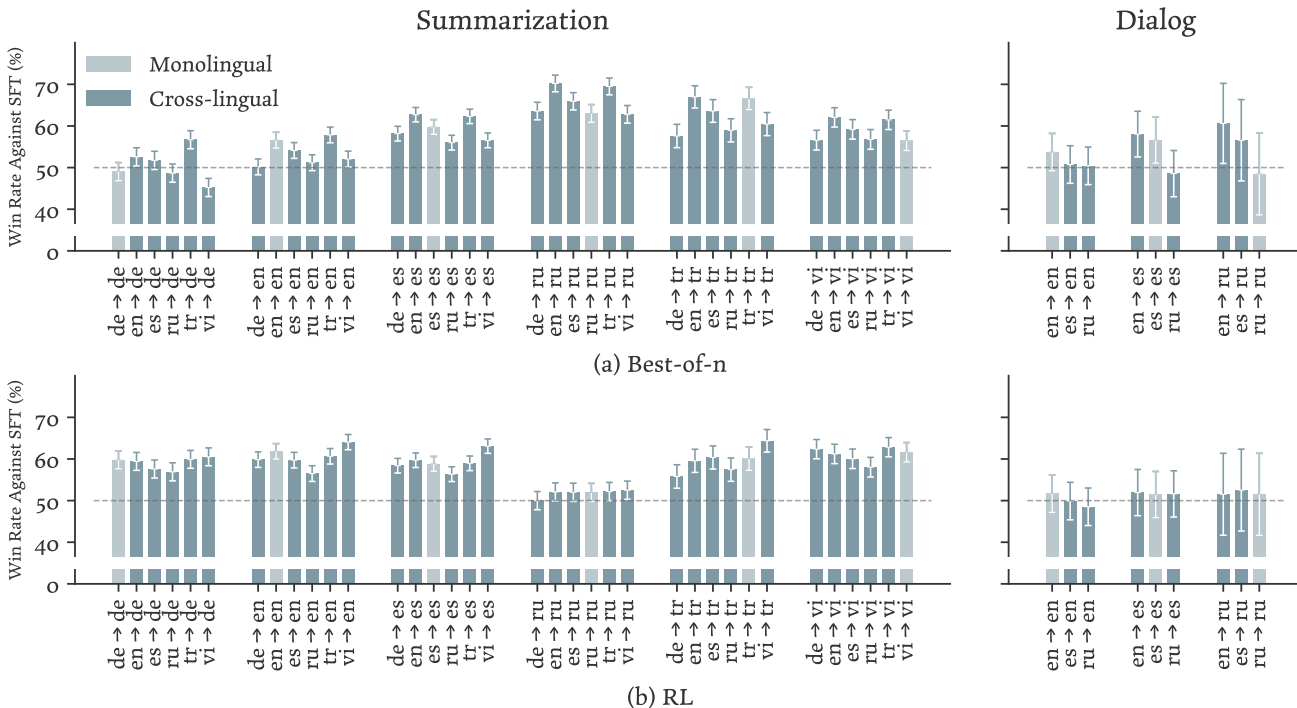|		De	En	Es	Ru	Tr	Vi
Summarization	Acc.	73.5%	73.0%	73.2%	73.7%	73.6%	78.2%
	N	306	1672	295	255	720	349
Dialog	Acc.	–	72.0%	70.8%	73.3%	–	–
	N	–	472	311	99	–	–

Figure 9. Alignment effectiveness, compared to the target-language SFT model judged by GPT-4, and the 95% confidence interval across validation instances. "source→target" denotes a source-language RM driving alignment in the target language. **Cross-lingual alignment is generally effective, sometimes outperforming monolingual alignment**. RL is hard to train for OpenAssistant, in line with what its authors found (Köpf et al., 2023).

check if the RM ranks the candidate generations correctly according to human preference. For Seahorse, a pointwise dataset, we group summaries for the same source document, and for each summary pair in such groups, we compute the ranking correctness.

We show the results in Table 5. The accuracies generally match the human agreement in Seahorse (Clark et al., 2023), and while human agreement was not reported in OpenAssistant, they generally match the human agreement numbers in past work on dialog generation (Bai et al., 2022a; Dubois et al., 2024) too (see §5.1 for reference human agreement numbers). Taken together with the LM judges' agreement with human evaluation (§5.1), we believe it is valid to use a LM to assess the generation quality in our setup.

D. GPT-4 as a Judge Results

In this section, we present the alignment evaluation results as judged by GPT-4, specifically the `gpt-4-0125-preview` model. Due to its high cost, we cap the number of evaluation instances for each dataset at 1,000 (i.e., for each row of Table 1 and 3). The results are shown in Figure 9. We observe the same trends as in §5.1, where cross-lingual reward optimization is generally effective, sometimes even more so than when done monolingually. Compared to PaLM-2-L, the two LMs agree on 72% of the instances in English and 71% in Spanish for summarization, and 75% and 73% for these languages for dialog. These are higher than the baseline human-human agreement numbers in §5.1. This shows a sign of homogeneity between LM judges, but also confirms their reliability.
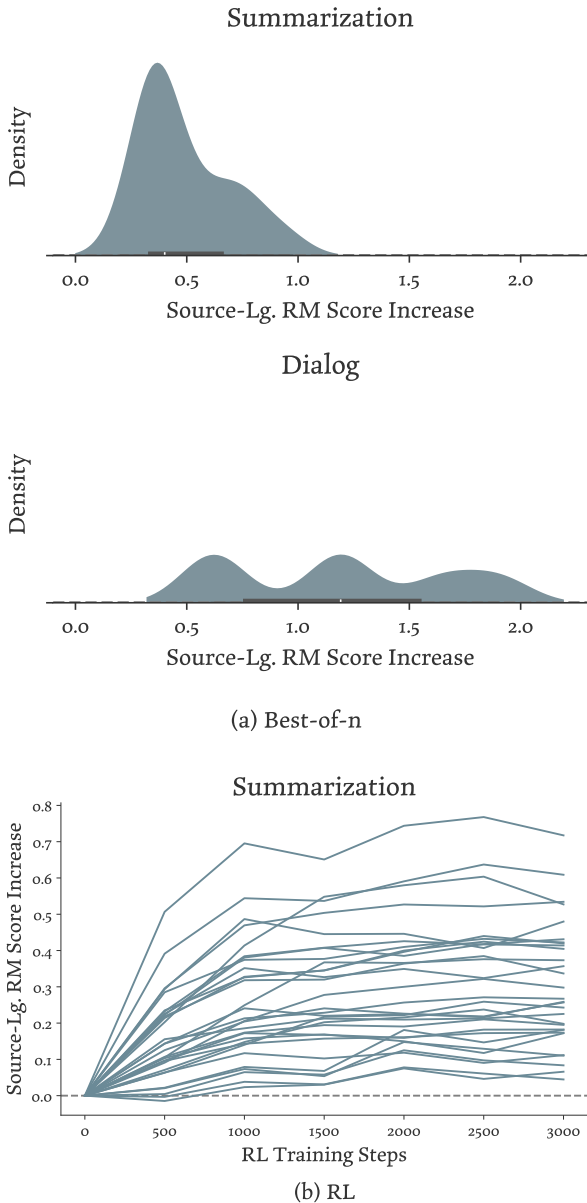
Summarization

Dialog

(a) Best-of-n

Summarization

(b) RL

Figure 10. Source-language RM generalizability evaluated by increases in scores they assign to target-language generations after monolingual target-language alignment (best-of-n or RL). We show all (source, target) language pairs where the two languages differ as density in (a) and lines in (b). RL is difficult to train for OpenAssistant (§4), so we omit it here, since the assumption that the RL'ed model is better would not hold. **In most cases, the source-language RM assigns a higher score (>0 increase) to aligned models, demonstrating cross-lingual RM generalizability.**

E. Verifying RM Transfer for Reward Modeling

In §6.1, we observed RM generalizability on the original reward modeling task, which would be a necessary condition for successful downstream cross-lingual alignment. There, we showed that the source-language RMs assign higher scores to better target-language generations than worse ones. Here, we consider an alternative setup to study the same problem: instead of relying on existing RM datasets for the better and worse generations, we take generations from monolingually-aligned models as better ones than those from unaligned (i.e., SFT) models. The assumption here is that monolingual alignment improves model quality, which is indeed the case as illustrated in Figure 4 and 9. Like in §6.1, we indeed see from Figure 10 that source-language RMs assign higher scores to monolingually-aligned models than unaligned SFT models. Under RL, this score difference also increases throughout training. These results confirm the RMs' cross-lingual generalizability within the reward modeling task.

F. Alignment Using Bilingual RMs

Seeing the benefit of cross-lingual RM transferability in §5, we hypothesize that bilingual RMs could bring further improvements since the resulting reward could be encouraged to be more language agnostic (Mulcaire et al., 2019). It would be computationally expensive to experiment with all possible language configurations (there would be a cubic number of them with pairwise sources), so, for simplicity, we take the best-performing source languages under the summarization best-of-n setup as judged by PaLM-2-L, English and German (Figure 6), and see if a bilingual RM based on them would lead to further performance improvement. Specifically, we first train a bilingual SFT model by pooling the SFT data for both languages, and similarly for the RM, which initializes from this bilingual SFT model.

Figure 11 does not show an improvement from the bilingual RM, which always achieves similar performance to the English RM, the better of the two monolingual RMs. Nevertheless, if this trend holds consistently, that the bilingual RM matches the performance of the better monolingual RM, this could be useful as an alternative to having to perform source language selection. We leave a more systematic validation of this phenomenon to future work.

G. Prompts

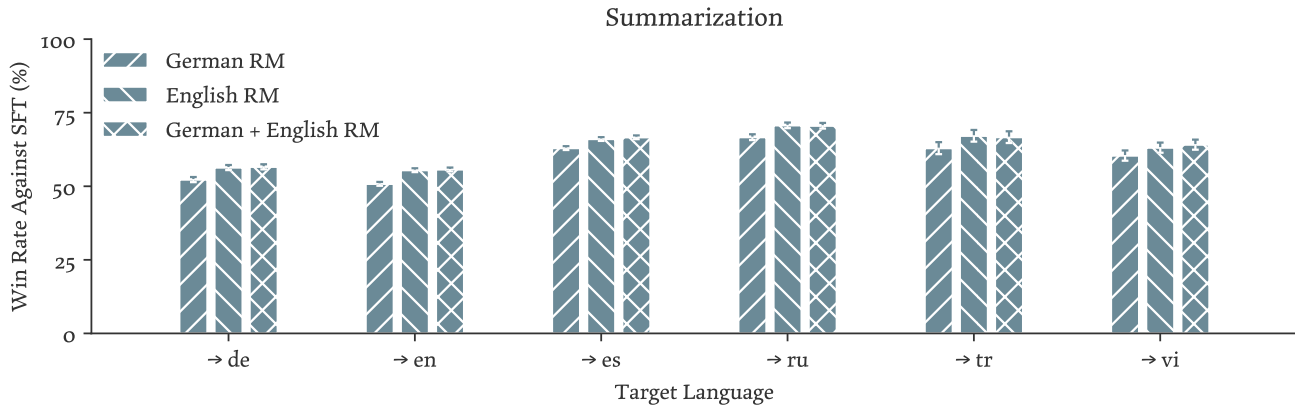In this section, we list all the prompts we used.

Figure 11. Alignment performance, measured in the win rate against the monolingual target-language SFT model, when alignment is driven by a German RM, an English RM, or a bilingual German + English RM. **The bilingual RM does not yield a noticeable improvement**.

G.1. Task Instructions

We prepend the following task-specific instructions to inputs for SFT and reward optimization. All occurrences of [LANGUAGE] are substituted with the target language. The RM stage does not include such prompts, where we simply concatenate the texts with delimiters.

Summarization: Summarize the following text in [LANGUAGE]:

Dialog generation: You are given a dialog between a human and an assistant in [LANGUAGE]. Please write one turn of the assistant side in [LANGUAGE].\n\n"

G.2. Evaluation Prompts

We use the following prompts to elicit pairwise generation judgments for both human and LM judge evaluation. All occurrences of [LANGUAGE], [INPUT], [GENERATION1], and [GENERATION2] are substituted with the respective content. For both tasks, we compare the probability of the tokens "1" and "2". To control for the positional bias of LMs (Wang et al., 2023a; Pezeshkpour & Hruschka, 2023; Zheng et al., 2023) and potentially of our human annotators, we randomly shuffle the two generations for human evaluation and the GPT-4 judge. For the PaLM-2 judge for which we have probability access, we prompt the LM judge twice with both orderings of the generations and compute the accuracy by averaging the probabilities of the "1" and "2" tokens.

Summarization. This prompt is adapted from the one in Lee et al. (2023).

A good summary is a shorter piece of text that has the essence of the original. It tries to accomplish the same purpose and

conveys the key information from the original post. Below we define four evaluation axes for summary quality: coherence, accuracy, coverage, and overall quality.

Coherence: This axis answers the question "how coherent is the summary on its own?" A summary is coherent if it's easy to understand when read on its own and free of English errors. A summary is not coherent if it's difficult to understand what the summary is trying to say. Generally, it's more important that the summary is understandable than it being free of grammar errors.

Accuracy: This axis answers the question "does the factual information in the summary accurately match the post?" A summary is accurate if it doesn't say things that aren't in the article, it doesn't mix up people, and generally is not misleading.

Coverage: This axis answers the question "how well does the summary cover the important information in the post?" A summary has good coverage if it mentions the main information from the post that's important to understand the situation described in the post. A summary has poor coverage if someone reading only the summary would be missing several important pieces of information about the situation in the post. A summary with good coverage

should also match the purpose of the
original post (e.g. to ask for advice).

Overall quality: This axis answers the
question "how good is the summary overall
at representing the post?" This can
encompass all of the above axes of quality,
as well as others you feel are important.
If it's hard to find ways to make the
summary better, the overall quality is
good. If there are lots of different ways
the summary can be made better, the overall
quality is bad.

You are an expert summary rater and are
knowledgeable in [LANGUAGE]. Given a
piece of text in [LANGUAGE] and two of its
possible summaries, also in [LANGUAGE],
output 1 or 2 to indicate which summary
best adheres to coherence, accuracy,
coverage, and overall quality as defined
above.

Text - [INPUT]
Summary 1 - [GENERATION1]
Summary 2 - [GENERATION2]

Preferred Summary=

Dialog Generation   This prompt is adapted from the one in Li et al. (2023b).

You are a helpful assistant, that ranks
models by the quality of their answers.
You are also knowledgeable in [LANGUAGE].

I want you to create a leaderboard of
different large-language models. To do
so, I will give you the instructions
(prompts) given to the models, and the
responses of two models. Please rank the
models based on which response would be
preferred by humans. All inputs are
python dictionaries.

Here is the prompt, in [LANGUAGE]:
{{
    "instruction": """[INPUT]""",
}}

Here are the outputs of the models, also
in [LANGUAGE]:
[
    {{

Table 6. Cross-lingual alignment results using **best-of-**n with $n = 64$, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **PaLM-2-L** (Figure 4).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	52.3	50.8	63.0	66.7	63.0	60.4
En	56.4	55.5	66.1	70.7	67.2	63.1
Es	51.9	51.2	62.4	66.0	64.4	57.5
Ru	48.1	46.5	59.2	63.6	59.0	56.3
Tr	53.3	52.9	62.6	66.6	60.4	59.0
Vi	46.5	48.2	60.0	65.6	62.1	58.0

Table 7. Cross-lingual alignment results using **best-of-**n with $n = 64$, for the **dialog generation** task, measured in win rate (%) against the target-language SFT model as judged by **PaLM-2-L** (Figure 4).

Src \ Tgt	En	Es	Ru
En	62.9	65.0	59.6
Es	59.1	62.4	57.6
Ru	53.4	54.3	52.5

```
        "model": "model_1",
        "answer": """[GENERATION1]"""
    }},
    {{
        "model": "model_2",
        "answer": """[GENERATION2]"""
    }}
]
```

Respond 1 or 2 to indicate the better
output. Please provide the ranking that
the majority of humans would give.

Better output=

H. Raw Results

We show the raw numerical results that correspond to our plots in Table 6 to 25.

Table 8. Cross-lingual alignment results using **RL**, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **PaLM-2-L** (Figure 4).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	59.4	61.0	59.4	49.6	52.5	59.3
En	55.9	59.9	58.5	52.6	54.8	56.6
Es	52.0	56.1	56.8	53.0	55.0	49.9
Ru	54.8	55.2	56.8	51.8	53.3	52.2
Tr	53.1	54.6	55.7	53.1	53.4	56.3
Vi	63.9	61.8	65.2	54.6	55.1	53.6

Table 9. Cross-lingual alignment results using **RL**, for the **dialog generation** task, measured in win rate (%) against the target-language SFT model as judged by **PaLM-2-L** (Figure 4).

Src \ Tgt	En	Es	Ru
En	53.1	54.5	53.5
Es	49.9	51.1	47.5
Ru	51.2	52.7	52.5

Table 10. Cross-lingual alignment results using **best-of-**n with $n = 64$, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **GPT-4** (Figure 9).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	49.0	50.2	58.2	63.6	57.6	56.6
En	52.6	56.6	62.7	70.2	67.0	62.1
Es	51.7	54.1	59.8	65.9	63.6	59.2
Ru	48.7	51.2	56.0	63.0	59.0	56.8
Tr	56.7	57.8	62.3	69.5	66.6	61.5
Vi	45.2	52.1	56.6	62.8	60.5	56.5

Table 11. Cross-lingual alignment results using **best-of-**n with $n = 64$, for the **dialog generation** task, measured in win rate (%) against the target-language SFT model as judged by **GPT-4** (Figure 9).

Src \ Tgt	En	Es	Ru
En	53.7	58.0	60.6
Es	50.7	56.6	56.6
Ru	50.4	48.6	48.5

Table 12. Cross-lingual alignment results using **RL**, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **GPT-4** (Figure 9).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	59.8	59.9	58.4	50.0	55.8	62.4
En	59.4	61.8	59.7	52.1	59.6	61.2
Es	57.6	59.7	58.8	52.0	60.4	60.1
Ru	56.9	56.5	56.4	52.0	57.4	58.0
Tr	59.9	60.7	59.0	52.2	60.1	62.8
Vi	60.5	64.1	63.1	52.5	64.4	61.6

Table 13. Cross-lingual alignment results using **RL**, for the **dialog generation** task, measured in win rate (%) against the target-language SFT model as judged by **GPT-4** (Figure 9).

Src \ Tgt	En	Es	Ru
En	51.7	51.9	51.5
Es	49.9	51.5	52.5
Ru	48.5	51.6	51.5

Table 14. Cross-lingual alignment results using **best-of-**n, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **human evaluators** (Figure 2).

Src \ Tgt	En	Es
De	61.0	64.0
En	60.9	67.4
Es	62.6	69.0
Ru	51.9	63.4
Tr	61.8	66.3
Vi	52.3	61.2

Table 15. Cross-lingual alignment results using **RL**, for the **summarization** task, measured in win rate (%) against the target-language SFT model as judged by **human evaluators** (Figure 2).

Src \ Tgt	En	Es
De	64.4	64.2
En	61.4	65.9
Es	58.7	62.7
Ru	61.9	60.6
Tr	63.3	64.9
Vi	66.2	64.7

Table 16. Cross-lingual alignment results using **best-of-**n with $n = 64$, for the **dialog generation** task, measured in win rate (%) against the target-language SFT model as judged by **human evaluators** (Figure 2).

Src \ Tgt	En	Es
En	67.6	52.0
Es	71.4	56.4

Table 17. Alignment quality using RM trained by translating the source language data into the target language using best-of-n with $n = 64$, for the summarization task, measured in win rate (%) against the target-language SFT model as judged by PaLM-2-L (§5.1).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	–	50.0	61.9	66.1	66.1	54.6
En	47.9	–	63.3	64.9	64.5	53.1
Es	50.6	52.9	–	64.1	64.5	59.0
Ru	47.4	51.2	60.3	–	63.3	57.7
Tr	50.6	52.5	61.8	65.6	–	50.8
Vi	42.0	50.8	59.1	64.4	63.6	–

Table 18. RM generalizability within the reward modeling task evaluated by accuracy (%) on in-task validation data for the summarization task, on the six Seahorse languages, as well as the majority baseline and the length baseline (§6.1) (Figure 7).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	71.0	64.8	68.0	67.9	67.5	67.7
En	62.2	67.4	67.9	66.3	66.5	70.8
Es	67.4	62.7	72.3	69.7	71.4	65.2
Ru	66.5	61.3	65.4	65.7	66.5	63.6
Tr	66.8	64.6	68.5	69.1	73.2	68.7
Vi	63.0	66.7	68.6	66.5	67.8	71.3
Majority	52.9	59.5	63.1	55.1	56.2	67.9
Length	56.6	59.5	63.1	55.1	55.2	67.9

Table 19. RM generalizability within the reward modeling task evaluated by accuracy (%) on in-task validation data for the dialog generation task, in three languages, as well as the length baseline (§6.1) (Figure 7).

Src \ Tgt	En	Es	Ru
En	68.4	68.4	76.3
Es	65.4	67.8	77.0
Ru	56.6	63.5	64.4
Length	66.1	68.1	71.1

Table 20. KL-divergence of the RL models from the corresponding target-language SFT model for the summarization task (Figure 8).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	0.92	0.78	0.83	0.01	0.37	1.92
En	1.50	1.32	1.01	0.02	0.83	3.30
Es	1.78	1.63	1.51	0.10	1.39	3.92
Ru	0.79	0.45	0.46	0.02	0.36	1.26
Tr	2.20	1.91	1.83	0.15	1.34	4.28
Vi	1.78	2.52	1.74	0.02	1.47	4.37

Table 21. ROUGE-L score when the SFT model is trained using different strategies, either monolingually, translated from a source language, or back-translated into a source language and then back (Figure 5(a)).

Lg.	De	En	Es	Ru	Tr	Vi
Mono.	36.2	38.9	32.9	16.9	35.2	41.8
Lg→En	27.8	–	27.1	22.4	28.2	26.7
En→Lg	16.1	–	24.6	13.6	29.9	40.3
En→Lg→En	36.5	–	36.1	35.4	36.5	35.8
Lg→En→Lg	32.5	–	26.6	12.2	32.1	34.9

Table 22. Alignment performance using best-of-n, measured in the win rate against the monolingual target language SFT model as judged by PaLM-2-L, when the SFT model is trained using different strategies. The first section uses a SFT model that is trained on target-language datasets (same as Table 6), while the second uses translated or back-translated SFT data (Figure 5(b)).

Lg.	De	Es	Ru	Tr	Vi
Target-language SFT; RM transfer only					
Lg→En	50.8	51.2	46.5	52.9	48.2
En→Lg	56.4	66.1	70.7	67.2	63.1
(Back-)Translated SFT					
Lg→En	36.6	26.6	29.8	37.5	31.8
En→Lg	14.4	43.5	43.9	47.1	41.6
En→Lg→En	42.7	43.2	40.1	41.4	37.1
Lg→En→Lg	45.3	54.0	60.1	61.7	51.1

Table 23. Alignment performance using best-of-n, measured in the win rate against the monolingual target language SFT model as judged by PaLM-2-L, when the SFT model is trained using different strategies. The first section uses a SFT model that is trained on target-language datasets, while the second uses translated or back-translated SFT data. Here, we only consider the WikiLingua dataset for both SFT and RM (Figure 5(c)).

Lg.	De	Es	Ru	Tr	Vi
Target-language SFT; RM transfer only					
Lg→En	38.3	32.4	38.6	32.9	29.2
En→Lg	62.8	59.4	53.7	47.4	66.4
(Back-)Translated SFT					
Lg→En	40.5	29.1	33.2	26.0	19.4
En→Lg	45.7	50.3	60.3	37.1	67.6
En→Lg→En	31.4	33.9	34.0	40.8	31.7
Lg→En→Lg	40.3	31.2	40.1	45.9	61.4

Table 24. Alignment performance using RL, measured in the win rate against the monolingual target language SFT model as judged by PaLM-2-L, when the SFT model is trained using different strategies. The first section uses a SFT model that is trained on target-language datasets, while the second uses translated or back-translated SFT data (Figure 5(d)).

Lg.	De	Es	Ru	Tr	Vi
Target-language SFT; RM transfer only					
Lg→En	61.0	56.1	55.2	54.6	61.8
En→Lg	55.9	58.5	52.6	54.8	56.6
(Back-)Translated SFT					
Lg→En	60.2	37.5	22.7	54.9	19.2
En→Lg	28.8	57.0	56.5	59.6	51.9
En→Lg→En	47.5	46.7	42.1	42.4	48.3
Lg→En→Lg	44.7	45.1	46.6	49.5	30.7

Table 25. Alignment performance using best-of-n, measured in the win rate against the monolingual target language SFT model as judged by PaLM-2-L, when using either a monolingual RM (same as Table 6) or a bilingual RM (Figure 11).

Src \ Tgt	De	En	Es	Ru	Tr	Vi
De	52.3	50.8	63.0	66.7	63.0	60.4
En	56.4	55.5	66.1	70.7	67.2	63.1
De + En	56.6	55.7	66.6	70.6	66.7	64.1