

# SteerEval: Inference-time Interventions Strengthen Multilingual Generalization in Neural Summarization Metrics

Silvia Casola<sup>1\*</sup>, Ryan Soh-Eun Shim<sup>2\*</sup>,  
Felicia Körner<sup>2</sup>, Yuchen Mao<sup>3</sup>, Barbara Plank<sup>2</sup>

<sup>1</sup>LaSTUS Lab, TALN Group, Universitat Pompeu Fabra (UPF), Barcelona, Spain

<sup>2</sup>MaiNLP, Center for Information and Language Processing, LMU Munich, Germany  
& Munich Center for Machine Learning

<sup>3</sup>Language Science and Technology, Saarland University, Germany

## Abstract

An empirical bottleneck for Natural Language Generation tasks such as summarization is the lack of accurate and robust *evaluation metrics for many languages*, which hinders progress. While Large Language Models (LLMs) are increasingly adopted as evaluators with the advent of LLM-as-a-Judge approaches, their agreement with human judgments remains low across many languages. We take inspiration from recent studies suggesting that multilingual language models often use English as an internal pivot language, and that misalignment with this pivot can lead to degraded downstream performance during text generation. Motivated by this hypothesis, we ask whether steering the activations of neural metrics with respect to an English pivot can improve their correlation with human judgments in multilingual summarization evaluation. We experiment with encoder- and decoder-based metrics across five languages. We find that activation steering is effective for neural metrics, and that a single, language-independent steering vector improves metric-human agreement. Steering also results in improvements when applied to COMET, a small encoder-based metric, showing notable human agreement despite its small scale.

## 1 Introduction

Metrics for Natural Language Generation (NLG) tasks such as summarization have become increasingly model-based (Zhang et al., 2020; Rei et al., 2020; Sellam et al., 2020), with the latest trend leveraging Large Language Models (LLMs) as evaluators (Liu et al., 2023; Fu et al., 2024), a paradigm known as LLM-as-a-Judge (Zheng et al., 2023).

When built on multilingual backbones, LLM evaluators can generalize across languages to some extent (Zhang et al., 2020; Rei et al., 2020; Sellam et al., 2020). However, preliminary studies

on model-based multilingual evaluation suggest that they often do not correlate well with human judgments (Mondshine et al., 2025), limiting their reliability as proxies for human evaluation.

We hypothesize that this gap is partly due to the way multilingual models process different languages: recent work has shown that multilingual LLMs internally rely on English as a pivot language (Wendler et al., 2024), and that the degree of alignment between the internal representations of a language and this pivot varies between languages (Wang et al., 2025a). This misalignment has been linked to poor performance in multilingual factual recall (Lu et al., 2025) and other discriminative and generative tasks (Wang et al., 2025a; Bafna et al., 2025), suggesting that explicitly controlling this alignment could improve performance. Wang et al. (2025a), for example, directly intervene on internal representations by steering them toward (or away from) English, showing this intervention is beneficial across a wide variety of NLG tasks. Lu et al. (2025) further show how alignment to the English pivot improves factual recall across languages.

Such work, however, is typically limited to discriminative and generative tasks: whether activation steering can improve performance on *evaluation* tasks, and specifically on model-based metrics such as LLM-as-a-Judge, remains unexplored. This is a particularly important gap: reliable automatic metrics for multilingual evaluation are scarce, making the quality of model-based evaluation a bottleneck for progress across languages. As such, in this work we explicitly ask: To what degree does steering the activations of low-resource languages with respect to English strengthen *LLM-as-a-Judge* performance?

To answer this question, we focus on summarization evaluation and first conduct a meta-analysis of LLM-as-a-Judge summarization metrics, assessing their reliability and alignment with human judgments. Corroborating prior work, we find modest

\* Equal contribution. Ordering decided on a coin flip.  
The work of Silvia Casola was carried out while at MaiNLP, LMU Munich & Munich Center for Machine Learning.




Figure 1: Illustration of our STEEREVAL method. We first use parallel data to learn cross-lingual vectors of mappings from *Source language*  $X \rightarrow \text{English}$  or a single multilingual average vector. We then apply the steering to neural metrics. Finally, we show that such steering results in better correlation with human judgments.

and inconsistent correlation with human judgments (Mondshine et al., 2025), showing the need for strengthened multilingual evaluators. To this end, we investigate the effect of steering multilingual LLM-as-a-Judge models with respect to a high-resource pivot language at test time, and find it improves metric reliability. We further demonstrate that language-specific steering vectors may not be required, and a single, multilingual vector can lead to the largest improvements. Finally, we show that small encoder-based metrics such as COMET (Rei et al., 2020) show surprisingly competitive performance, and that this performance can further be improved through steering. Our contributions are as follows:<sup>1</sup>

**Multilingual metric meta-analysis.** We perform a meta-analysis of model-based summarization metrics across five typologically diverse languages. We examine two LLM-as-a-judge methods, using four backbone models and two prompt variations. We also examine COMET, an encoder-based metric, which shows competitive performance for summarization.

**Activation steering for multilingual metrics.** We apply activation steering to two popular LLM-as-a-judge methods—based on next direct prompting (Liu et al., 2023) and average log likelihood (Fu et al., 2024)—and we show activation steering strengthens correlation with human judgments for both approaches. We also extend steering to encoder-based models, largely unexplored in the

broader activation steering literature, and we also observe consistent correlation gains. Our intervention consists of a simple and efficient vector-based addition, occurs at inference time, and requires no model retraining, making it an efficient way to strengthen metrics.

**Multilingual intervention** We measure the cross-lingual similarity between *Source language*  $X \rightarrow \text{English}$  intervention vectors in LLMs, and find these vectors to be largely similar across languages, corroborating prior work in the generalizability of refusal vectors (Wang et al., 2025b) and, more broadly, shared language geometries (Chi et al., 2020; Dumas et al., 2025). Building on this finding, we explore multilingual steering, where we average over language-specific vectors to obtain a single multilingual steering vector. We show this solution outperforms language-specific steering across many settings.

## 2 Related Work

### 2.1 Multilingual summarization evaluation

Automatic evaluation of natural language generation (NLG) and summarization is a longstanding challenge (Sai et al., 2022; Celikyilmaz et al., 2021). Most existing work has focused on English. For instance, Fabbri et al. (2021) evaluate 14 summarization metrics across four quality dimensions, while subsequent studies show that correlation with human judgments of reference-based metrics can drop when scoring LLM-generated outputs given their high quality (Casola et al., 2025). As evaluation has increasingly shifted towards the

<sup>1</sup>Our code is available at: <https://github.com/mainlp/steereval>

LLM-as-a-Judge paradigm, [Bavaresco et al. \(2025\)](#) conducted a large-scale evaluation across multiple tasks, including summarization, and demonstrated high variability in metric behavior across different LLM judges.

Beyond English, automatic evaluation is understudied, and metric reliability often tends to further degrade. [Clark et al. \(2023\)](#) introduce a multilingual dataset in which model outputs are annotated in a binary fashion for specific desirable properties (e.g., conciseness), and shows it can be used to train ad-hoc evaluation metrics. Similarly, [Hada et al. \(2024\)](#) propose a multilingual dataset with human judgments using either binary labels or a 0–2 scale, and use it to assess LLMs as evaluators across five dimensions. More focused studies include [Barnes et al. \(2025\)](#), who analyze LLM-as-a-judge performance for Spanish and Basque, highlighting variation in evaluators’ behaviors. Finally, [Mondshine et al. \(2025\)](#) recently introduced an abstractive summarization dataset for meta-evaluating automatic metrics across eight diverse languages, focusing on n-gram- and small model-based metrics, concluding that correlations with human judgments remain relatively low in multilingual settings.

## 2.2 Activation steering

The linear representation hypothesis ([Park et al., 2024](#)) proposes that high-level concepts are represented as directions in the activation space of language models. Motivated by this hypothesis, an increasing body of work has used a test-time method known as activation steering to control LLM behavior. Such works typically isolate the direction that represents a given concept by vector arithmetic, and the isolated vector is then added to activations at test time ([Turner et al., 2024](#)). Such a steering method has been applied to induce or reduce characteristics such as truthfulness ([Marks and Tegmark, 2024](#)), toxicity, sentiment ([Turner et al., 2024](#)), refusal ([Arditi et al., 2024](#)), and language steering ([Sterz et al., 2025](#)).

Beyond steering vectors, [Wang et al. \(2025a\)](#) show that a learned linear map between low- and high-resource language representations can serve as an effective steering mechanism. More broadly, intervention methods can use other techniques, such as representation editing ([Meng et al., 2022](#)) and prompt optimization ([Pryzant et al., 2023](#)).

In this paper, we focus on steering, and specifically on the vector-based approach, which we apply to strengthening multilingual summarization evalu-

ation on non-English languages.

## 3 Experimental Setup

### 3.1 Vector-based metric steering

Figure 1 provides an overview of our steering methodology. Specifically, we perform a vector-based intervention via the following steps.

**Steering direction learning** We examine two alternative approaches: a monolingual intervention for a specific source language (Section 4) and a multilingual intervention for all languages (Section 5). For the monolingual intervention, we derive a steering direction between a source language and English from parallel data. Specifically, given representations in the source (lower-resourced) language  $h_s$  and English (the target language)  $h_t$ , we compute a per-layer steering vector by subtracting the mean of  $h_s$  from the mean of  $h_t$  over all time steps, capturing the direction from the source to the target language in activation space:

$$\bar{h}_{v,l} = \bar{h}_{t,l} - \bar{h}_{s,l}$$

We also experiment with a single, multilingual vector, mapping from all languages to English, in line with previous work ([Lu et al., 2025](#)). To do so, we average across language-specific vectors, and obtain our final steering vector as:

$$\bar{h}_{v,l} = \frac{1}{|M|} \sum_{m \in M} (\bar{h}_{t,l} - \bar{h}_{m,l})$$

where  $m$  is a source language in our language set  $M$  and  $|M| = 5$  in our experiments.

All representations are taken from the residual stream at the output of layer  $l$ , i.e., the hidden state produced by block  $l$  before it is passed to block  $l + 1$ . Note that the parallel data used to learn the steering vector is disjoint from the summarization evaluation data, and any parallel corpus can in principle be used for this purpose.

**Test time intervention** Keeping the language model frozen, we then steer its internal representations. Specifically, at test time, we add the monolingual or multilingual steering vector to the activations of a single layer  $l$ :

$$h_{q,l}^{\text{steered}} = h_{q,l} + \sigma \bar{h}_{v,l}$$

where  $h_{q,l}$  is the unsteered representation of test input  $q$  at layer  $l$ , and  $\sigma$  is a scalar parameter controlling the strength and sign of the steering.

For language-specific steering (Section 4), for every target language, we jointly tune the steering strength  $\sigma$  and the application layer  $l$  on a language-specific development set with human judgments, and then evaluate the selected configuration on the corresponding disjoint test set. Following prior work (Lu et al., 2025; Bafna et al., 2025), we restrict the layer search space to the middle and final layers, as these exhibit the greatest representational mismatch with English for other tasks. For COMET, we steer the pooling-space embedding directly (no per-layer selection). We search over  $\sigma \in [-5, 5]$  with a step size of 1. Hyperparameters are selected via joint grid search to maximize the average Pearson correlation across six evaluation dimensions (see Section 3.2). In the multilingual vector experiments (Section 5), we apply the steering vector zero-shot without hyperparameter tuning.

**Performance testing** We evaluate the metric using the steered representations and analyze its agreement with human judgments. Since the Seahorse dataset (Clark et al., 2023, which we discuss in Section 3.2) is annotated with binary labels, we use Pearson correlation and ROC-AUC, following the original paper.

### 3.2 Datasets

**Steering direction learning** To learn the language alignments, we leverage the FLORES multiparallel dataset (Goyal et al., 2022), which consists of human-translated sentences primarily from Wikipedia. We follow Wang et al. (2025a) and use 500 parallel sentence pairs.

#### Parameter tuning and performance testing

We use the Seahorse multilingual summarization dataset (Clark et al., 2023) as key benchmark, a high-quality human evaluation dataset containing binary scores for six languages i.e., English, German, Spanish, Russian, Turkish, and Vietnamese. Source texts are extracted from four datasets in the GEM benchmark (Gehrmann et al., 2021): XSum (Narayan et al., 2018, for English only), XL-Sum (Hasan et al., 2021), MLSum (Scialom et al., 2020), and WikiLingua (Ladhak et al., 2020). Summaries are the outputs of eight abstractive summarization systems based on T5 (Raffel et al., 2023, only for English), mT5 (Xue et al., 2021) and PaLM (Chowdhery et al., 2022); human references from the respective datasets, when available, are also treated as hypotheses. The dataset is large-scale,

with language-specific test sets of 2,600+ annotated outputs. Each summary is annotated by trained raters on six binary dimensions: whether it is a) comprehensible, b) non-repetitive, c) grammatical, d) concise, e) whether its content is fully attributable to the source article, and f) whether it captures the main ideas of the source article. We preserve the original split and use the dataset validation set to tune the method hyperparameters and its test set for the final evaluation.<sup>2</sup>

### 3.3 Metrics

We primarily examine the following metrics:

**Direct prompting** We directly prompt a multilingual model to produce a score for each evaluation dimension. However, directly extracting the score from the model output can be unreliable, as responses may include additional text or justifications, an issue even more complex in multilingual settings. Thus, we rely on token probabilities. We follow G-Eval (Liu et al., 2023) and compute the score as a weighted sum over the token probabilities corresponding to the evaluation scale at the first generated token. The score is computed as:

$$\text{score} = \sum_i p(s_i) \times s_i$$

where  $s_i \in \{0, 1\}$ , mimicking the dataset annotation, and  $p(s_i)$  is the token probability.

#### GPTScore (prompt-conditioned log-likelihood)

Following Fu et al. (2024), we evaluate a hypothesis using the model’s conditional generation probabilities. Given a prompt that specifies a desired property (e.g., conciseness), the model is instructed to generate a summary exhibiting that property. The summary to be evaluated is thus scored by computing the average log-probability of its tokens under the model, conditioned on the prompt. Specifically, we compute the final score as:

$$\text{score} = \frac{1}{N} \sum_{t=1}^N \log p(h_t \mid h_{<t}, P, \theta)$$

where  $P$  is the prompt,  $N$  is the number of tokens in the summary, and  $\theta$  the model parameters.

For both LLM-based metrics, we use Llama-3.2-3B<sup>3</sup> (Grattafiori et al., 2024) and aya-expanse-8b<sup>4</sup>

<sup>2</sup>We provide dataset statistics in Appendix D.

<sup>3</sup>meta-llama/Llama-3.2-3B

<sup>4</sup>CohereLabs/aya-expanse-8b

(Dang et al., 2024), as backbones and further experiment with their larger counterparts i.e., Llama-3.1-8B Instruct<sup>5</sup> and aya-expense-32b<sup>6</sup> in Section 5. These models are chosen since they are open, instruction-tuned, diverse in size, have broad language coverage, and have shown good performance in various multilingual tasks. We prompt each model both in English and with a machine translated version in the target language. Appendix 3 details the metrics’ implementation, including the prompts.

**COMET** Due to prior work reporting state-of-the-art performance on many languages by directly applying COMET on multilingual summarization evaluation (Mondshine et al., 2025; Krubiński and Pecina, 2022), we also analyze the effect of steering on COMET, which is a multilingual neural metric typically used in machine translation. We rely on the *wmt22-comet-da* model, which is built on top of XLM-R (Conneau et al., 2020) and follow the experimental design of Mondshine et al. (2025) to adapt the machine-translation metric to summarization: we replace the source input with an empty string, use the system-generated summary as the hypothesis, and the human-written summary as the reference. Note that COMET, in contrast to LLM-based judges, requires a reference summary. In the Seahorse dataset, reference summaries are themselves treated as hypotheses and scored by human judges. We therefore compute COMET scores only for the subset of model-generated hypotheses that have an associated reference, and results are not directly comparable to those from the LLMs.

## 4 Language-specific steering

Tables 1 and 2 summarize the main results of our LLM-as-a-Judge and COMET steering experiments respectively. In this section, we focus on language-specific steering and later discuss the zero-shot multilingual average vector setup in Section 5. We primarily report results using Pearson correlation and focus on English prompting in the main text; detailed results when prompting in the source language are in Appendix E, while results using AUC are presented in Appendix F. We report the average correlation over the six evaluation dimensions per language; per-dimension results are in Appendix A.

<sup>5</sup>*meta-llama/llama-3.1-8B-Instruct*

<sup>6</sup>*CohereLabs/aya-expense-32b*

**Baselines show relatively modest and variable correlation with human judgments.** Across all settings, baseline systems achieve relatively low positive Pearson correlation, showing that summary evaluation remains challenging in multilingual settings, and that metrics designed for English do not necessarily transfer across languages, corroborating previous results (Clark et al., 2023; Mondshine et al., 2025). Larger models do not consistently outperform smaller ones: for GPTScore, aya-expense-32b shows comparable baseline performance to its smaller (8b) counterpart and, interestingly, lower performance than Llama-3.1-8b across all languages. No single configuration is superior across all settings: the choice of metric, backbone model, and prompt language interact in non-trivial ways. While comparisons with LLM-based scores should be taken with care given the modified test set (see Section 3.3), COMET typically achieves a strong baseline correlations across all languages, including for languages like German that show consistently low correlations for LLM-based metrics. This is surprising given that COMET is a much smaller model trained specifically for translation quality estimation rather than summarization evaluation. We hypothesize this improvement might be due to the specific supervised training and to the use of a reference, not available to the LLMs in our prompts.

**Steering consistently improves performance.** Language-specific steering typically improves over the baseline, with gains in the vast majority of languages, metrics, and backbone models. While the advantage is generally incremental, the method shows consistent improvement across the majority of the configurations, which is notable given the simplicity and generality of our inference-time steering method. Steering typically improves over the baseline across LLM-based metrics, though the magnitude and consistency of gains vary considerably depending on the configuration. Improvements are also observed for COMET, though these are less consistent over languages when using language-specific steering.

**GPTScore tends to be more stable across backbone models, and benefits most consistently from steering.** GPTScore tends to show more consistent results than Direct Prompting across backbone models. We hypothesize this stability might be due to the log-likelihood signal being more robust across model sizes when com-

|                         |                                | de                                                     | es                                                    | ru                                                   | tr                                                   | vi                                                     |
|-------------------------|--------------------------------|--------------------------------------------------------|-------------------------------------------------------|------------------------------------------------------|------------------------------------------------------|--------------------------------------------------------|
| <b>GPTScore</b>         |                                |                                                        |                                                       |                                                      |                                                      |                                                        |
| Llama-3.2-3B            | Baseline (EN-prompt)           | .022 $\pm$ .132                                        | .213 $\pm$ .141                                       | .270 $\pm$ .062                                      | .202 $\pm$ .108                                      | .139 $\pm$ .092                                        |
|                         | + vector steer (EN-prompt)     | .022 $\pm$ .127                                        | <b>.215<math>\pm</math>.142<math>\uparrow</math></b>  | <b>.276<math>\pm</math>.062<math>\uparrow</math></b> | <b>.211<math>\pm</math>.100<math>\uparrow</math></b> | <b>.138<math>\pm</math>.095<math>\downarrow</math></b> |
|                         | + avgvec zero-shot (EN-prompt) | <b>.024<math>\pm</math>.127<math>\uparrow</math></b>   | <b>.221<math>\pm</math>.149<math>\uparrow</math></b>  | <b>.281<math>\pm</math>.062<math>\uparrow</math></b> | <b>.210<math>\pm</math>.104<math>\uparrow</math></b> | <b>.139<math>\pm</math>.097</b>                        |
| aya-expanse-8b          | Baseline (EN-prompt)           | .017 $\pm$ .143                                        | .150 $\pm$ .088                                       | .240 $\pm$ .067                                      | .190 $\pm$ .069                                      | .181 $\pm$ .063                                        |
|                         | + vector steer (EN-prompt)     | <b>.018<math>\pm</math>.145<math>\uparrow</math></b>   | <b>.155<math>\pm</math>.088<math>\uparrow</math></b>  | <b>.247<math>\pm</math>.067<math>\uparrow</math></b> | <b>.198<math>\pm</math>.070<math>\uparrow</math></b> | <b>.185<math>\pm</math>.063<math>\uparrow</math></b>   |
|                         | + avgvec zero-shot (EN-prompt) | <b>.021<math>\pm</math>.143<math>\uparrow</math></b>   | <b>.156<math>\pm</math>.089<math>\uparrow</math></b>  | <b>.248<math>\pm</math>.067<math>\uparrow</math></b> | <b>.199<math>\pm</math>.070<math>\uparrow</math></b> | <b>.188<math>\pm</math>.063<math>\uparrow</math></b>   |
| Llama-3.1-8B-Instruct   | Baseline (EN-prompt)           | .036 $\pm$ .138                                        | .228 $\pm$ .133                                       | .282 $\pm$ .059                                      | .242 $\pm$ .100                                      | .176 $\pm$ .081                                        |
|                         | + avgvec zero-shot (EN-prompt) | <b>.040<math>\pm</math>.133<math>\uparrow</math></b>   | <b>.238<math>\pm</math>.143<math>\uparrow</math></b>  | <b>.294<math>\pm</math>.058<math>\uparrow</math></b> | <b>.250<math>\pm</math>.096<math>\uparrow</math></b> | <b>.175<math>\pm</math>.085<math>\downarrow</math></b> |
| aya-expanse-32b         | Baseline (EN-prompt)           | .007 $\pm$ .139                                        | .164 $\pm$ .084                                       | .244 $\pm$ .067                                      | .190 $\pm$ .070                                      | .195 $\pm$ .053                                        |
|                         | + avgvec zero-shot (EN-prompt) | <b>.011<math>\pm</math>.140<math>\uparrow</math></b>   | <b>.169<math>\pm</math>.085<math>\uparrow</math></b>  | <b>.250<math>\pm</math>.065<math>\uparrow</math></b> | <b>.196<math>\pm</math>.070<math>\uparrow</math></b> | <b>.199<math>\pm</math>.053<math>\uparrow</math></b>   |
| <b>Direct Prompting</b> |                                |                                                        |                                                       |                                                      |                                                      |                                                        |
| Llama-3.2-3B            | Baseline (EN-prompt)           | .050 $\pm$ .050                                        | -.004 $\pm$ .045                                      | .014 $\pm$ .039                                      | .007 $\pm$ .043                                      | -.001 $\pm$ .040                                       |
|                         | + vector steer (EN-prompt)     | <b>.049<math>\pm</math>.048<math>\downarrow</math></b> | <b>-.001<math>\pm</math>.044<math>\uparrow</math></b> | <b>.015<math>\pm</math>.039<math>\uparrow</math></b> | <b>.008<math>\pm</math>.045<math>\uparrow</math></b> | <b>.005<math>\pm</math>.040<math>\uparrow</math></b>   |
| aya-expanse-8b          | Baseline (EN-prompt)           | <b>.208<math>\pm</math>.042</b>                        | <b>.226<math>\pm</math>.139</b>                       | <b>.249<math>\pm</math>.079</b>                      | <b>.237<math>\pm</math>.112</b>                      | <b>.248<math>\pm</math>.077</b>                        |
|                         | + vector steer (EN-prompt)     | <b>.208<math>\pm</math>.042</b>                        | <b>.226<math>\pm</math>.139</b>                       | <b>.251<math>\pm</math>.076<math>\uparrow</math></b> | <b>.239<math>\pm</math>.113<math>\uparrow</math></b> | <b>.247<math>\pm</math>.077<math>\downarrow</math></b> |

Table 1: SEAHORSE Pearson  $r$  averaged over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split. **Bold** indicates the single highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/ $\uparrow$  (**red**/ $\downarrow$ ) on steered rows indicates improvement (resp. degradation) relative to the baseline.

|                    | de                                                   | es                                                   | ru                                                     | tr                                                     | vi                                                   |
|--------------------|------------------------------------------------------|------------------------------------------------------|--------------------------------------------------------|--------------------------------------------------------|------------------------------------------------------|
| <b>COMET</b>       |                                                      |                                                      |                                                        |                                                        |                                                      |
| Baseline           | .248 $\pm$ .087                                      | .266 $\pm$ .120                                      | .358 $\pm$ .130                                        | .380 $\pm$ .139                                        | .290 $\pm$ .116                                      |
| + vector steer     | <b>.279<math>\pm</math>.081<math>\uparrow</math></b> | <b>.318<math>\pm</math>.184<math>\uparrow</math></b> | <b>.351<math>\pm</math>.155<math>\downarrow</math></b> | <b>.360<math>\pm</math>.144<math>\downarrow</math></b> | <b>.295<math>\pm</math>.136<math>\uparrow</math></b> |
| + avgvec zero-shot | <b>.275<math>\pm</math>.084<math>\uparrow</math></b> | <b>.303<math>\pm</math>.146<math>\uparrow</math></b> | <b>.364<math>\pm</math>.152<math>\uparrow</math></b>   | <b>.392<math>\pm</math>.165<math>\uparrow</math></b>   | <b>.297<math>\pm</math>.134<math>\uparrow</math></b> |

Table 2: SEAHORSE Pearson  $r$  over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language for COMET (wmt22-comet-da). Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split. **Bold** indicates the highest value per language. **Blue**/ $\uparrow$  (**red**/ $\downarrow$ ) indicates improvement (resp. degradation) subscripts relative to the baseline.

pared to prompting for direct scoring. Interestingly, the Llama-3B model shows typically even better GPTScore baseline performance than aya-8B across languages despite the difference in size, suggesting that GPTScore may be less sensitive to model capacity than Direct Prompting. Steering improves GPTScore across nearly all language and model combinations: aya-expanse-8b and aya-expanse-32b show gains across all five languages with no degradations, while Llama-3.2-3B and Llama-3.1-8B-Instruct show only a single marginal degradation each, in Vietnamese. This suggests that the conditional log-likelihood scoring of multilingual models is particularly steerable, regardless of model size.

**Direct prompting requires higher model capacity, and steering results are more mixed.** In contrast to GPTScore, direct prompting appears

to require models with higher capacity to produce meaningful signal. For Llama-3.2-3B, direct prompting degrades to near-zero or even negative correlations when prompting in English, with baseline values of -.004 for Spanish and -.001 for Vietnamese, likely due to the backbone’s limited ability to follow instructions. Despite this weak baseline, steering recovers some signal across nearly all languages. With a larger model, direct prompting is the best-performing configuration, with aya-expanse-8b reaching the highest baseline correlations across four out of five languages. However, steering results in limited improvement over baselines, producing marginal gains for Russian and Turkish, with a slight degradation for Vietnamese.

**COMET shows strong but less consistent gains from language-specific steering.** As discussed above, COMET baselines are substantially higher

than LLM-based metrics across all languages. Language-specific steering improves German and Spanish substantially, but degrades Russian and Turkish, with a marginal gain for Vietnamese. Note that the multilingual variant, which we discuss in Section 5, shows a more consistent picture, improving all five languages and achieving the best results for Russian, Turkish, and Vietnamese.

**Performance across languages does not follow resourceness levels.** German, generally considered a relatively high-resource language, consistently shows the lowest correlations across all configurations. Under Direct Prompting with aya-expanse-8b, differences across languages are small, suggesting that larger and more multilingual models may be more robust to language variation. Turkish shows notable variability: it is among the strongest languages under Direct Prompting with aya-expanse-8b, but weaker under GPTScore with Llama-3.2-3B, possibly due to its morphological complexity (which might be challenging for token log-likelihood scoring). When steering is applied, degradations are rare and typically marginal, concentrated in Vietnamese and German for LLMs. The pattern is similar for COMET, where the multilingual variant reverses the two degradations under vector steering (Russian and Turkish).

**The effect of prompt language depends on the model and metric.** Prompting in the source language (Appendix E) does not show a consistent advantage over English. Prompting in English and the source language results in different baseline strengths depending on the metrics: source languages typically produce stronger Direct Prompting baselines, while English tends to yield higher GPTScore baselines. Steering is more consistently beneficial when prompting in English.

**Performance is highly variable across dimensions** Looking at Appendix A, baseline difficulty varies across dimensions. On average, Comprehensible is the easiest dimension, while Grammatical is the hardest one. One notable interaction concerns the metric: COMET substantially outperforms GPTScore and Direct Prompting on fluency-adjacent dimensions, but this advantage largely disappears for Grammatical and reverses for Attributable, where GPTScore outperforms COMET. Language-specific vector steering typically leads to larger gains on Comprehensible and Non-repetitive.

## 5 Multilingual steering

Prior work in activation steering has found high level-concepts such as refusal or script to be broadly similar vectors across languages (Wang et al., 2025b; Shim et al., 2026; Lu et al., 2025). In this section, we investigate to what extent language-specific steering is needed for model-based evaluation metrics. We focus on GPTScore and COMET, as they show more consistent and interpretable behaviors than direct prompting for monolingual steering. For simplicity, we experiment with prompting in English only.

To study the similarity of cross-language steering vectors, we measure the extent to which *Language X*  $\rightarrow$  *English* directions are similar. As such, we compute the cosine similarity between our *Language X*  $\rightarrow$  *English* vectors extracted from parallel data. Figure 2 shows the results of projecting language samples from the parallel data onto the principal components (PCA, left) and the cross-language similarity matrix (right) for GPTScore, for the Llama-3.2-3b model. We focus on the last model layer, as it is consistently identified as optimal on our development set across languages and settings. The PCA reveals that samples cluster by language, suggesting that the model’s internal representations encode language identity at this layer. The similarity matrix, on the other hand, shows that the *Language X*  $\rightarrow$  *English* steering vectors are relatively similar across languages.

Motivated by this similarity, and in line with previous work (Lu et al., 2025), we explore the effect of taking a single steering vector across all languages, as detailed in Section 3.1.

For GPTScore, we steer the last layer of the model, following previous observations; for COMET, similarly to the language-specific steering, we directly steer the pooling-space embedding directly. We set  $\sigma$  to -2.5 (the midpoint of our explored negative interval) for all languages, since the parameter is consistently negative across language.

Surprisingly, averaged zero-shot vector is at least as good as, and often better than, the language-specific tuned vector in nearly all configurations. Focusing on LLM-based metrics, Table 1 (focusing on the avgvec zero-shot rows) shows the results. For Llama-3.2-3B, improvements over the baseline are observed for all languages except Vietnamese (for which results are consistent with no intervention), and the averaged vector results in higher correlations than the language-specific tuned vector for




Figure 2: Projection of the instances from FLORES into the Principal Components (left) and cosine similarity across steering vectors (right) for Llama-3.2-3B.

all languages except Turkish. For aya-expanse-8b, improvements are consistent across all languages, always exceeding language-specific performance.

The effectiveness of a single multilingual vector is particularly interesting given that we performed no parameter tuning and pre-selected an average negative  $\sigma$ , a setting that is simpler than the language-specific one. Since we do not search for hyperparameters and use a single language-independent average vector, we can also scale to larger models with limited computational cost. Thus, we also explore whether our approach improves for larger models of the same family. After steering, GPTScore with Llama-3.1-8B achieves the best performance overall in Spanish, Russian, and Turkish, improving on non-trivial baselines; Vietnamese is the only language showing minor degradation, consistently with previous findings. aya-expanse-32b also shows improvement with respect to its baseline.

Improvements also stand for COMET, for which results are shown in Table 2. When using a cross-lingual vector (without parameter tuning), COMET results are remarkably strong. Despite requiring no language-specific tuning, it matches or exceeds the language-specific steered variant in four out of five languages, and achieves the best overall result for Russian ( $r = .364$ ), Turkish ( $r = .392$ ), and Vietnamese ( $r = .297$ ). Also notice that, unlike the language-specific vector, the cross-lingual vector never degrades performance: for Russian and Turkish, where individual steering vectors hurt, the cross-lingual vector not only recovers but exceeds the baseline, suggesting that averaging across languages might have a natural regularizing effect that suppresses language-specific noise. Taken together,

these results suggest that encoder-based metrics have strong potential when references are available, and can be steered effectively, making our approach generalizable beyond large decoder-based LLMs.

These results suggest that language-specific steering vectors might not be necessary: a single averaged vector, applied zero-shot with a fixed hyperparameter, might be sufficient to obtain consistent gains.

## 6 Conclusion

In this paper, we leverage the insight that LLMs internally use English as a pivot language (Wendler et al., 2024) to improve the reliability of neural metrics for multilingual summarization evaluation. We showed that test-time activation steering consistently improves metrics’ correlation with human judgments across a diverse set of languages, metrics, and backbone models, with no model retraining required. Our findings contribute to the growing body of work that leverages model-internal representations for practical downstream tasks (Marks and Tegmark, 2024), and extend it to the underexplored setting of automatic evaluation.

Notably, we find that language-specific steering is not necessary: a single averaged multilingual vector, applied zero-shot with a fixed hyperparameter, is sufficient to obtain consistent gains and often outperforms its language-specific counterpart.

We also apply our method to COMET: despite being a much smaller encoder-based model, steered COMET benefits strongly from our intervention and surpasses LLM-based evaluators in human agreement across multiple languages.

Several directions remain open for future work. Our approach assumes the availability of parallel

data to derive steering vectors; moreover, future work should study how far the multilingual vector generalizes to unseen languages not included in the averaging set. Finally, while our work focuses on language steering, prior work has leveraged test-time interventions targeting specific evaluative dimensions such as fluency (Sheng et al., 2024); combining language- and dimension-specific interventions could yield further gains in multilingual evaluation.

## Limitations

Due to the relative lack of large-scale summarization evaluation datasets with human annotations, our study is limited by the number of available languages. For the specific dataset we employ (Clark et al., 2023), the additional constraint that this dataset only has binary annotations is a further limitation.

We optimize our steering hyperparameters on the average performance improvement across dimensions. Future work should investigate to what extent setting dimension-specific parameters could improve performance.

While we demonstrate the effectiveness of a single multilingual steering vector across the languages in our study, we do not evaluate its generalization to unseen languages, i.e., languages not in the averaging set. Whether the vector transfers to typologically distant or lower-resourced languages remains an open question.

Steering representations toward an English pivot may introduce unwanted biases, e.g., cultural biases. We leave a systematic analysis of such biases to future work.

Finally, we limit our analysis to a vector-based steering intervention. Future work should investigate different steering methods (Wang et al., 2025a) and other representation interventions, such as representation editing (Meng et al., 2022) or prompt optimization (Pryzant et al., 2023).

## Ethical Statement

We acknowledge the usage of LLMs in the partial assistance of figure creation, code development, and in strengthening the coherence of parts of the manuscript.

## Acknowledgments

We would like to thank the members of the MaiNLP lab for their valuable feedback. We recog-

nize the support for Ryan Soh-Eun Shim, Felicia Körner, and Barbara Plank through the ERC Consolidator Grant DIALECT 101043235. Silvia Casola was supported by the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101282681 (GenEval).

## References

- Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. 2024. [Refusal in language models is mediated by a single direction](#). *Preprint*, arXiv:2406.11717.
- Niyati Bafna, Tianjian Li, Kenton Murray, David R. Mortensen, David Yarowsky, Hale Sirin, and Daniel Khashabi. 2025. [The translation barrier hypothesis: Multilingual generation with large language models suffers from implicit translation failure](#). *Preprint*, arXiv:2506.22724.
- Jeremy Barnes, Naiara Perez, Alba Bonet-Jover, and Begonia Altuna. 2025. [Summarization metrics for spanish and basque: Do automatic scores and llm-judges correlate with humans?](#) *Preprint*, arXiv:2503.17039.
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, Andre Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Suglia, Aditya K Surikuchi, Ece Takmaz, and Alberto Testoni. 2025. [LLMs instead of human judges? a large scale empirical study across 20 NLP evaluation tasks](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 238–255, Vienna, Austria. Association for Computational Linguistics.
- Silvia Casola, Yang Janet Liu, Siyao Peng, Oliver Kraus, Albert Gatt, and Barbara Plank. 2025. [References matter: Investigating the impact of reference set variation on summarization evaluation](#). In *Proceedings of the 18th International Natural Language Generation Conference*, pages 274–291, Hanoi, Vietnam. Association for Computational Linguistics.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2021. [Evaluation of text generation: A survey](#). *Preprint*, arXiv:2006.14799.
- Ethan A. Chi, John Hewitt, and Christopher D. Manning. 2020. [Finding universal grammatical relations in multilingual BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5577, Online. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton,

- Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, and 48 others. 2022. [Palm: Scaling language modeling with pathways](#). *Preprint*, arXiv:2204.02311.
- Elizabeth Clark, Shruti Rijhwani, Sebastian Gehrmann, Joshua Maynez, Roei Aharoni, Vitaly Nikolaev, Thibault Sellam, Aditya Siddhant, Dipanjan Das, and Ankur Parikh. 2023. [SEAHORSE: A multilingual, multifaceted dataset for summarization evaluation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9397–9413, Singapore. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya expand: Combining research breakthroughs for a new multilingual frontier](#). *Preprint*, arXiv:2412.04261.
- Clément Dumas, Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2025. [Separating tongue from thought: Activation patching reveals language-agnostic concept representations in transformers](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31822–31841, Vienna, Austria. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating summarization evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2024. [GPTScore: Evaluate as you desire](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6556–6576, Mexico City, Mexico. Association for Computational Linguistics.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuoluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, and 37 others. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). In *Proceedings of the First Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 96–120, Online. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Rishav Hada, Varun Gumma, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2024. [METAL: Towards multilingual meta-evaluation](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2280–2298, Mexico City, Mexico. Association for Computational Linguistics.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XLsum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Mateusz Krubiński and Pavel Pecina. 2022. [From COMET to COMES – can summary evaluation benefit from translation evaluation?](#) In *Proceedings of the 3rd Workshop on Evaluation and Comparison of NLP Systems*, pages 21–31, Online. Association for Computational Linguistics.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. [WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048, Online. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.

- Meng Lu, Ruochen Zhang, Carsten Eickhoff, and Elie Pavlick. 2025. [Paths not taken: Understanding and mending the multilingual factual recall pipeline](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 15066–15096, Suzhou, China. Association for Computational Linguistics.
- Samuel Marks and Max Tegmark. 2024. [The geometry of truth: Emergent linear structure in large language model representations of true/false datasets](#). *Preprint*, arXiv:2310.06824.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in gpt](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 17359–17372. Curran Associates, Inc.
- Itai Mondshine, Tzuf Paz-Argaman, and Reut Tsarfaty. 2025. [Beyond n-grams: Rethinking evaluation metrics and strategies for multilingual abstractive summarization](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 19019–19035, Vienna, Austria. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Kiho Park, Yo Joong Choe, and Victor Veitch. 2024. The linear representation hypothesis and the geometry of large language models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML’24. JMLR.org.
- Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang Zhu, and Michael Zeng. 2023. [Automatic prompt optimization with “gradient descent” and beam search](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, Singapore. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Preprint*, arXiv:1910.10683.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. 2022. [A survey of evaluation metrics used for nlg systems](#). *ACM Comput. Surv.*, 55(2).
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. [MLSUM: The multilingual summarization corpus](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8051–8067, Online. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Shuqian Sheng, Yi Xu, Tianhang Zhang, Zanwei Shen, Luoyi Fu, Jiaxin Ding, Lei Zhou, Xiaoying Gan, Xinbing Wang, and Chenghu Zhou. 2024. [RepEval: Effective text evaluation with LLM representation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7019–7033, Miami, Florida, USA. Association for Computational Linguistics.
- Ryan Soh-Eun Shim, Kwanghee Choi, Calvin Chang, Ming-Hao Hsu, Florian Eichin, Zhizheng Wu, Alane Suhr, Michael A. Hedderich, David Harwath, David R. Mortensen, and Barbara Plank. 2026. [Linear script representations in speech foundation models enable zero-shot transliteration](#). *Preprint*, arXiv:2601.02906.
- Hannah Sterz, Fabian David Schmidt, Goran Glavaš, and Ivan Vulić. 2025. [ReCoVeR the target language: Language steering without sacrificing task performance](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 19390–19405, Suzhou, China. Association for Computational Linguistics.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2024. [Steering language models with activation engineering](#). *Preprint*, arXiv:2308.10248.
- Weixuan Wang, Minghao Wu, Barry Haddow, and Alexandra Birch. 2025a. [Bridging the language gaps in large language models with inference-time cross-lingual intervention](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5418–5433, Vienna, Austria. Association for Computational Linguistics.
- Xinpeng Wang, Mingyang Wang, Yihong Liu, Hinrich Schütze, and Barbara Plank. 2025b. [Refusal direction is universal across safety-aligned languages](#). *Preprint*, arXiv:2505.17306.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1:*

*Long Papers*), pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.

## **A Dimension-specific Results**

Tables 3, 4, 5, 6, 7, and 8 present the results for the six evaluation dimensions in Seahorse when using an English prompt.

|                         |                                | de            | es             | ru            | tr            | vi            |
|-------------------------|--------------------------------|---------------|----------------|---------------|---------------|---------------|
| <b>GPTScore</b>         |                                |               |                |               |               |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | .003          | .319           | .313          | .099          | .248          |
|                         | + vector steer (EN-prompt)     | <u>.025</u> ↑ | <u>.320</u> ↑  | <u>.315</u> ↑ | <u>.158</u> ↑ | <u>.244</u> ↓ |
|                         | + avgvec zero-shot (EN-prompt) | <u>.023</u> ↑ | <u>.333</u> ↑  | <u>.331</u> ↑ | <u>.140</u> ↑ | <u>.258</u> ↑ |
| aya-expanse-8b          | Baseline (EN-prompt)           | <u>.029</u>   | .249           | .353          | .114          | .305          |
|                         | + vector steer (EN-prompt)     | <u>.027</u> ↓ | <u>.255</u> ↑  | <u>.361</u> ↑ | <u>.118</u> ↑ | <u>.308</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | <u>.029</u> ↓ | <u>.256</u> ↑  | <u>.362</u> ↑ | <u>.118</u> ↑ | <u>.311</u> ↑ |
| <b>Direct Prompting</b> |                                |               |                |               |               |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.095</u>   | <u>-.001</u>   | .055          | .056          | .044          |
|                         | + vector steer (EN-prompt)     | <u>.095</u> ↑ | <u>-.009</u> ↓ | <u>.058</u> ↑ | <u>.067</u> ↑ | <u>.049</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | —             | —              | —             | —             | —             |
| aya-expanse-8b          | Baseline (EN-prompt)           | <u>.207</u>   | .335           | <u>.386</u>   | .477          | <u>.338</u>   |
|                         | + vector steer (EN-prompt)     | <u>.206</u> ↓ | <u>.335</u> ↑  | <u>.385</u> ↓ | <u>.481</u> ↑ | <u>.337</u> ↓ |
|                         | + avgvec zero-shot (EN-prompt) | —             | —              | —             | —             | —             |
| <b>COMET</b>            |                                |               |                |               |               |               |
| wmt22-comet-da          | Baseline                       | .297          | .421           | .578          | .651          | .425          |
|                         | + vector steer                 | <u>.361</u> ↑ | <u>.555</u> ↑  | <u>.622</u> ↑ | <u>.626</u> ↓ | <u>.454</u> ↑ |
|                         | + avgvec zero-shot             | <u>.346</u> ↑ | <u>.489</u> ↑  | <u>.627</u> ↑ | <u>.712</u> ↑ | <u>.451</u> ↑ |

Table 3: SEAHORSE Pearson  $r$  for dimension **Comprehensible** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/↑ (**red**/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                                | de             | es            | ru            | tr            | vi            |
|-------------------------|--------------------------------|----------------|---------------|---------------|---------------|---------------|
| <b>GPTScore</b>         |                                |                |               |               |               |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | -.147          | .068          | .178          | <u>.095</u>   | .079          |
|                         | + vector steer (EN-prompt)     | <u>-.145</u> ↑ | <u>.067</u> ↓ | <u>.185</u> ↑ | <u>.094</u> ↓ | <u>.079</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | <u>-.139</u> ↑ | <u>.075</u> ↑ | <u>.191</u> ↑ | <u>.095</u> ↓ | <u>.075</u> ↓ |
| aya-expanse-8b          | Baseline (EN-prompt)           | -.207          | .033          | .234          | .168          | .193          |
|                         | + vector steer (EN-prompt)     | <u>-.206</u> ↑ | <u>.038</u> ↑ | <u>.242</u> ↑ | <u>.176</u> ↑ | <u>.196</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | <u>-.199</u> ↑ | <u>.040</u> ↑ | <u>.245</u> ↑ | <u>.177</u> ↑ | <u>.200</u> ↑ |
| <b>Direct Prompting</b> |                                |                |               |               |               |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.088</u>    | .012          | .057          | <u>.045</u>   | .059          |
|                         | + vector steer (EN-prompt)     | <u>.081</u> ↓  | <u>.017</u> ↑ | <u>.060</u> ↑ | <u>.043</u> ↓ | <u>.064</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | —              | —             | —             | —             | —             |
| aya-expanse-8b          | Baseline (EN-prompt)           | <u>.210</u>    | .147          | .173          | .179          | <u>.180</u>   |
|                         | + vector steer (EN-prompt)     | <u>.206</u> ↓  | <u>.146</u> ↓ | <u>.189</u> ↑ | <u>.195</u> ↑ | <u>.177</u> ↓ |
|                         | + avgvec zero-shot (EN-prompt) | —              | —             | —             | —             | —             |
| <b>COMET</b>            |                                |                |               |               |               |               |
| wmt22-comet-da          | Baseline                       | .301           | .338          | .406          | .414          | .427          |
|                         | + vector steer                 | <u>.347</u> ↑  | <u>.513</u> ↑ | <u>.468</u> ↑ | <u>.450</u> ↑ | <u>.493</u> ↑ |
|                         | + avgvec zero-shot             | <u>.345</u> ↑  | <u>.428</u> ↑ | <u>.470</u> ↑ | <u>.468</u> ↑ | <u>.488</u> ↑ |

Table 4: SEAHORSE Pearson  $r$  for dimension **Non-repetitive** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/↑ (**red**/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                                | de             | es             | ru            | tr            | vi             |
|-------------------------|--------------------------------|----------------|----------------|---------------|---------------|----------------|
| <b>GPTScore</b>         |                                |                |                |               |               |                |
| Llama-3.2-3B            | Baseline (EN-prompt)           | -.073          | <u>-.023</u>   | .190          | .090          | .124           |
|                         | + vector steer (EN-prompt)     | <u>-.074</u> ↓ | <u>-.023</u> ↑ | .197↑         | .090↓         | .125↑          |
|                         | + avgvec zero-shot (EN-prompt) | <u>-.070</u> ↑ | <u>-.033</u> ↓ | <b>.198</b> ↑ | <u>.092</u> ↑ | <u>.121</u> ↓  |
| aya-expanse-8b          | Baseline (EN-prompt)           | -.101          | .065           | .140          | .121          | .171           |
|                         | + vector steer (EN-prompt)     | <u>-.101</u> ↑ | <b>.068</b> ↑  | .148↑         | .126↑         | .174↑          |
|                         | + avgvec zero-shot (EN-prompt) | <u>-.096</u> ↑ | <u>.066</u> ↑  | <u>.150</u> ↑ | <u>.128</u> ↑ | <u>.177</u> ↑  |
| <b>Direct Prompting</b> |                                |                |                |               |               |                |
| Llama-3.2-3B            | Baseline (EN-prompt)           | .013           | .010           | <u>.010</u>   | .023          | -.038          |
|                         | + vector steer (EN-prompt)     | <u>.014</u> ↑  | <u>.016</u> ↑  | .009↓         | <u>.025</u> ↑ | <u>-.031</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | —              | —              | —             | —             | —              |
| aya-expanse-8b          | Baseline (EN-prompt)           | <b>.159</b>    | -.048          | <u>.139</u>   | <u>.152</u>   | <u>.129</u>    |
|                         | + vector steer (EN-prompt)     | <u>.159</u> ↓  | <u>-.048</u> ↑ | <u>.139</u> ↓ | <u>.153</u> ↑ | <u>.126</u> ↓  |
|                         | + avgvec zero-shot (EN-prompt) | —              | —              | —             | —             | —              |
| <b>COMET</b>            |                                |                |                |               |               |                |
| wmt22-comet-da          | Baseline                       | .090           | <u>.045</u>    | <u>.197</u>   | .345          | .245           |
|                         | + vector steer                 | <u>.128</u> ↑  | <u>-.001</u> ↓ | <u>.170</u> ↓ | <u>.308</u> ↓ | <u>.237</u> ↓  |
|                         | + avgvec zero-shot             | <u>.119</u> ↑  | <u>.041</u> ↓  | <u>.183</u> ↓ | <b>.352</b> ↑ | <b>.246</b> ↑  |

Table 5: SEAHORSE Pearson  $r$  for dimension **Grammatical** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. Blue/↑ (red/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                                | de             | es             | ru             | tr             | vi             |
|-------------------------|--------------------------------|----------------|----------------|----------------|----------------|----------------|
| <b>GPTScore</b>         |                                |                |                |                |                |                |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.279</u>    | .368           | .334           | .336           | .278           |
|                         | + vector steer (EN-prompt)     | <u>.268</u> ↓  | <u>.371</u> ↑  | <u>.342</u> ↑  | <u>.334</u> ↓  | <u>.284</u> ↑  |
|                         | + avgvec zero-shot (EN-prompt) | <u>.272</u> ↓  | <u>.388</u> ↑  | <u>.341</u> ↑  | <u>.341</u> ↑  | <u>.280</u> ↑  |
| aya-expanse-8b          | Baseline (EN-prompt)           | .249           | .101           | .190           | .197           | .182           |
|                         | + vector steer (EN-prompt)     | <u>.254</u> ↑  | <u>.107</u> ↑  | <u>.196</u> ↑  | <u>.208</u> ↑  | <u>.186</u> ↑  |
|                         | + avgvec zero-shot (EN-prompt) | <u>.258</u> ↑  | <u>.111</u> ↑  | <u>.198</u> ↑  | <u>.208</u> ↑  | <u>.189</u> ↑  |
| <b>Direct Prompting</b> |                                |                |                |                |                |                |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>−.040</u>   | −.099          | <u>−.044</u>   | <u>−.062</u>   | −.051          |
|                         | + vector steer (EN-prompt)     | <u>−.039</u> ↑ | <u>−.094</u> ↑ | <u>−.045</u> ↓ | <u>−.066</u> ↓ | <u>−.048</u> ↑ |
|                         | + avgvec zero-shot (EN-prompt) | —              | —              | —              | —              | —              |
| aya-expanse-8b          | Baseline (EN-prompt)           | .293           | .344           | .266           | .222           | <u>.314</u>    |
|                         | + vector steer (EN-prompt)     | <u>.295</u> ↑  | <u>.345</u> ↑  | <u>.265</u> ↓  | <u>.219</u> ↓  | <u>.315</u> ↑  |
|                         | + avgvec zero-shot (EN-prompt) | —              | —              | —              | —              | —              |
| <b>COMET</b>            |                                |                |                |                |                |                |
| wmt22-comet-da          | Baseline                       | .168           | .219           | <u>.215</u>    | <u>.204</u>    | .112           |
|                         | + vector steer                 | <u>.226</u> ↑  | <u>.258</u> ↑  | <u>.202</u> ↓  | <u>.176</u> ↓  | <u>.123</u> ↑  |
|                         | + avgvec zero-shot             | <u>.209</u> ↑  | <u>.250</u> ↑  | <u>.211</u> ↓  | <u>.195</u> ↓  | <u>.121</u> ↑  |

Table 6: SEAHORSE Pearson  $r$  for dimension **Attributable** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. Blue/↑ (red/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                                | de            | es            | ru                 | tr                 | vi                 |
|-------------------------|--------------------------------|---------------|---------------|--------------------|--------------------|--------------------|
| <b>GPTScore</b>         |                                |               |               |                    |                    |                    |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.023</u>   | .267          | .303               | <u>.290</u>        | <u>.029</u>        |
|                         | + vector steer (EN-prompt)     | <u>.021</u> ↓ | <u>.271</u> ↑ | <u>.308</u> ↑      | <u>.285</u> ↓      | <u>.021</u> ↓      |
|                         | + avgvec zero-shot (EN-prompt) | <u>.020</u> ↓ | <u>.272</u> ↑ | <u>.308</u> ↑      | <u>.288</u> ↓      | <u>.022</u> ↓      |
| aya-expanse-8b          | Baseline (EN-prompt)           | .073          | .258          | .284               | .317               | .136               |
|                         | + vector steer (EN-prompt)     | <u>.074</u> ↑ | <u>.261</u> ↑ | <u>.289</u> ↑      | <u>.324</u> ↑      | <u>.138</u> ↑      |
|                         | + avgvec zero-shot (EN-prompt) | <u>.075</u> ↑ | <u>.263</u> ↑ | <u>.289</u> ↑      | <u>.325</u> ↑      | <u>.140</u> ↑      |
| <b>Direct Prompting</b> |                                |               |               |                    |                    |                    |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.052</u>   | .013          | −.030              | −.039              | −.013              |
|                         | + vector steer (EN-prompt)     | <u>.052</u> ↓ | <u>.025</u> ↑ | <u>−.024</u> ↑     | <u>−.035</u> ↑     | <u>−.006</u> ↑     |
|                         | + avgvec zero-shot (EN-prompt) | —             | —             | —                  | —                  | —                  |
| aya-expanse-8b          | Baseline (EN-prompt)           | .182          | <u>.303</u>   | <u>.274</u>        | .241               | <u>.304</u>        |
|                         | + vector steer (EN-prompt)     | <u>.182</u> ↑ | <u>.302</u> ↓ | <u>.272</u> ↓      | <u>.242</u> ↑      | <u>.303</u> ↓      |
|                         | + avgvec zero-shot (EN-prompt) | —             | —             | —                  | —                  | —                  |
| <b>COMET</b>            |                                |               |               |                    |                    |                    |
| wmt22-comet-da          | Baseline                       | .324          | .340          | <b><u>.419</u></b> | <b><u>.378</u></b> | <b><u>.331</u></b> |
|                         | + vector steer                 | <u>.320</u> ↓ | <u>.331</u> ↓ | <u>.346</u> ↓      | <u>.335</u> ↓      | <u>.287</u> ↓      |
|                         | + avgvec zero-shot             | <u>.325</u> ↑ | <u>.357</u> ↑ | <u>.376</u> ↓      | <u>.346</u> ↓      | <u>.293</u> ↓      |

Table 7: SEAHORSE Pearson  $r$  for dimension **Main ideas** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. Blue/↑ (red/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                                | de                 | es           | ru                 | tr           | vi            |
|-------------------------|--------------------------------|--------------------|--------------|--------------------|--------------|---------------|
| <b>GPTScore</b>         |                                |                    |              |                    |              |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.045</u>        | .279         | .303               | .303         | <u>.079</u>   |
|                         | + vector steer (EN-prompt)     | <b>.035↓</b>       | <b>.284↑</b> | <b>.312↑</b>       | <b>.301↓</b> | <b>.075↓</b>  |
|                         | + avgvec zero-shot (EN-prompt) | <b>.037↓</b>       | <b>.290↑</b> | <b>.314↑</b>       | <b>.306↑</b> | <b>.075↓</b>  |
| aya-expanse-8b          | Baseline (EN-prompt)           | .058               | .195         | .237               | .223         | .102          |
|                         | + vector steer (EN-prompt)     | <u>.062↑</u>       | <u>.199↑</u> | <u>.244↑</u>       | <u>.235↑</u> | <u>.106↑</u>  |
|                         | + avgvec zero-shot (EN-prompt) | <u>.062↑</u>       | <u>.202↑</u> | <u>.246↑</u>       | <u>.236↑</u> | <u>.109↑</u>  |
| <b>Direct Prompting</b> |                                |                    |              |                    |              |               |
| Llama-3.2-3B            | Baseline (EN-prompt)           | <u>.094</u>        | <u>.042</u>  | <u>.034</u>        | <u>.018</u>  | −.008         |
|                         | + vector steer (EN-prompt)     | <b>.092↓</b>       | <b>.036↓</b> | <b>.029↓</b>       | <b>.016↓</b> | <b>−.001↑</b> |
|                         | + avgvec zero-shot (EN-prompt) | —                  | —            | —                  | —            | —             |
| aya-expanse-8b          | Baseline (EN-prompt)           | .197               | .274         | .257               | .153         | .220          |
|                         | + vector steer (EN-prompt)     | <u>.201↑</u>       | <u>.274↑</u> | <u>.257↑</u>       | <b>.146↓</b> | <b>.224↑</b>  |
|                         | + avgvec zero-shot (EN-prompt) | —                  | —            | —                  | —            | —             |
| <b>COMET</b>            |                                |                    |              |                    |              |               |
| wmt22-comet-da          | Baseline                       | <b><u>.307</u></b> | .233         | <b><u>.334</u></b> | <u>.288</u>  | <u>.201</u>   |
|                         | + vector steer                 | <b>.292↓</b>       | <b>.253↑</b> | <b>.296↓</b>       | <b>.267↓</b> | <b>.179↓</b>  |
|                         | + avgvec zero-shot             | <b>.306↓</b>       | <b>.254↑</b> | <b>.319↓</b>       | <b>.277↓</b> | <b>.182↓</b>  |

Table 8: SEAHORSE Pearson  $r$  for dimension **Concise** per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split; avgvec zero-shot uses a fixed vector averaged across all source languages. **Bold** indicates the highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue/↑ (red/↓)** on steered rows indicates improvement (degradation) relative to the baseline.

## **B Optimal Hyperparameters**

|                         |            | <b>de</b> |          | <b>es</b> |          | <b>ru</b> |          | <b>tr</b> |          | <b>vi</b> |          |
|-------------------------|------------|-----------|----------|-----------|----------|-----------|----------|-----------|----------|-----------|----------|
|                         |            | layer     | $\sigma$ | layer     | $\sigma$ | layer     | $\sigma$ | layer     | $\sigma$ | layer     | $\sigma$ |
| <b>GPTScore</b>         |            |           |          |           |          |           |          |           |          |           |          |
| Llama-3.2-3B            | EN-prompt  | 27        | -2.0     | 27        | +2.0     | 27        | -1.0     | 27        | -5.0     | 22        | -2.0     |
|                         | src-prompt | 27        | -2.0     | 27        | +1.0     | 27        | -1.0     | 27        | -3.0     | 27        | +1.0     |
| aya-expanse-8b          | EN-prompt  | 27        | -5.0     | 31        | -3.0     | 31        | -2.0     | 31        | -3.0     | 31        | -2.0     |
|                         | src-prompt | 31        | -3.0     | 31        | -2.0     | 31        | -1.0     | 31        | -2.0     | 31        | -1.0     |
| <b>Direct Prompting</b> |            |           |          |           |          |           |          |           |          |           |          |
| Llama-3.2-3B            | EN-prompt  | 27        | +5.0     | 27        | -5.0     | 24        | +5.0     | 25        | +5.0     | 22        | +3.0     |
|                         | src-prompt | 25        | +4.0     | 27        | -5.0     | 21        | +2.0     | 23        | +4.0     | 26        | -4.0     |
| aya-expanse-8b          | EN-prompt  | 29        | -5.0     | 31        | +5.0     | 26        | -5.0     | 26        | -5.0     | 25        | -5.0     |
|                         | src-prompt | 26        | -5.0     | 26        | -5.0     | 26        | -2.0     | 29        | -3.0     | 25        | -5.0     |

Table 9: Best intervention layer and steering  $\sigma$  for LLM-based SEAHORSE evaluation. Layer and  $\sigma$  are selected jointly on the validation split by mean Pearson  $r$  across six quality dimensions. Avgvec rows use a fixed  $\sigma$  and the layer that maximised cross-lingual vector similarity.

|                               | <b>de</b> | <b>es</b> | <b>ru</b> | <b>tr</b> | <b>vi</b> |
|-------------------------------|-----------|-----------|-----------|-----------|-----------|
| <b>COMET (wmt22-comet-da)</b> |           |           |           |           |           |
| vector steer                  | -5.0      | -5.0      | -2.0      | +5.0      | -2.0      |

Table 10: Best steering  $\sigma$  for COMET (wmt22-comet-da) SEAHORSE evaluation. COMET steers the pooling-space embedding directly (no per-layer selection);  $\sigma$  is chosen on the validation split by mean Pearson  $r$  across six quality dimensions.

## **C Prompts**

|            | English                                                                                                   | German                                                                                                                  | Spanish                                                                                              | Russian                                                                                | Turkish                                                                                       | Vietnamese                                                                                                |
|------------|-----------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| Compreh.   | Generate a comprehensible summary for the following article: <i>[article]</i><br>Summary:                 | Erstelle eine verständliche Zusammenfassung des folgenden Artikels: <i>[article]</i><br>Zusammenfassung:                | Genera un resumen comprensible del siguiente artículo: <i>[article]</i><br>Resumen:                  | Создайте понятное резюме следующей статьи: <i>[article]</i><br>Резюме:                 | Aşağıdaki makale için anlaşılır bir özet oluşturun: <i>[article]</i><br>Özet:                 | Tạo bản tóm tắt dễ hiểu cho bài viết sau: <i>[article]</i><br>Tóm tắt:                                    |
| Non-rep.   | Generate a summary without unnecessary repetition for the following article: <i>[article]</i><br>Summary: | Erstelle eine Zusammenfassung ohne unnötige Wiederholungen des folgenden Artikels: <i>[article]</i><br>Zusammenfassung: | Genera un resumen sin repeticiones innecesarias del siguiente artículo: <i>[article]</i><br>Resumen: | Создайте резюме без лишних повторений следующей статьи: <i>[article]</i><br>Резюме:    | Aşağıdaki makale için gereksiz tekrar içermeyen bir özet oluşturun: <i>[article]</i><br>Özet: | Tạo bản tóm tắt không có thông tin lặp lại không cần thiết cho bài viết sau: <i>[article]</i><br>Tóm tắt: |
| Gramm.     | Generate a grammatically correct summary for the following article: <i>[article]</i><br>Summary:          | Erstelle eine grammatisch korrekte Zusammenfassung des folgenden Artikels: <i>[article]</i><br>Zusammenfassung:         | Genera un resumen gramaticalmente correcto del siguiente artículo: <i>[article]</i><br>Resumen:      | Создайте грамматически правильное резюме следующей статьи: <i>[article]</i><br>Резюме: | Aşağıdaki makale için dilbilgisel açıdan doğru bir özet oluşturun: <i>[article]</i><br>Özet:  | Tạo bản tóm tắt đúng ngữ pháp cho bài viết sau: <i>[article]</i><br>Tóm tắt:                              |
| Attrib.    | Generate a faithful summary for the following article: <i>[article]</i><br>Summary:                       | Erstelle eine quelltreue Zusammenfassung des folgenden Artikels: <i>[article]</i><br>Zusammenfassung:                   | Genera un resumen fiel al contenido del siguiente artículo: <i>[article]</i><br>Resumen:             | Создайте достоверное резюме следующей статьи: <i>[article]</i><br>Резюме:              | Aşağıdaki makale için kaynağa sadık bir özet oluşturun: <i>[article]</i><br>Özet:             | Tạo bản tóm tắt trung thực với bài viết sau: <i>[article]</i><br>Tóm tắt:                                 |
| Main ideas | Generate a summary capturing the main ideas of the following article: <i>[article]</i><br>Summary:        | Erstelle eine Zusammenfassung mit den Hauptideen des folgenden Artikels: <i>[article]</i><br>Zusammenfassung:           | Genera un resumen con las ideas principales del siguiente artículo: <i>[article]</i><br>Resumen:     | Создайте резюме с основными идеями следующей статьи: <i>[article]</i><br>Резюме:       | Aşağıdaki makalenin ana fikirlerini içeren bir özet oluşturun: <i>[article]</i><br>Özet:      | Tạo bản tóm tắt nắm bắt ý chính của bài viết sau: <i>[article]</i><br>Tóm tắt:                            |
| Concise    | Generate a concise summary for the following article: <i>[article]</i><br>Summary:                        | Erstelle eine prägnante Zusammenfassung des folgenden Artikels: <i>[article]</i><br>Zusammenfassung:                    | Genera un resumen conciso del siguiente artículo: <i>[article]</i><br>Resumen:                       | Создайте краткое резюме следующей статьи: <i>[article]</i><br>Резюме:                  | Aşağıdaki makale için kısa ve öz bir özet oluşturun: <i>[article]</i><br>Özet:                | Tạo bản tóm tắt súc tích cho bài viết sau: <i>[article]</i><br>Tóm tắt:                                   |

Figure 3: GPTScore prompt templates per quality dimension and language. *[article]* is replaced by the source article; the second line (after the newline) is the continuation cue used for scoring.

|            | English                                                 | German                                                       | Spanish                                                  | Russian                                                 | Turkish                                      | Vietnamese                                                  |
|------------|---------------------------------------------------------|--------------------------------------------------------------|----------------------------------------------------------|---------------------------------------------------------|----------------------------------------------|-------------------------------------------------------------|
| Compreh.   | Is the summary comprehensible?                          | Ist die Zusammenfassung verständlich?                        | ¿Es el resumen comprensible?                             | Является ли резюме понятным?                            | Özet anlaşılır mı?                           | Bản tóm tắt có dễ hiểu không?                               |
| Non-rep.   | Does the summary contain no unnecessary repetition?     | Enthält die Zusammenfassung keine unnötigen Wiederholungen?  | ¿El resumen no contiene repeticiones innecesarias?       | Не содержит ли резюме ненужных повторений?              | Özet gereksiz tekrar içeriyor mu?            | Bản tóm tắt có không chứa sự lặp lại không cần thiết không? |
| Gramm.     | Is the summary grammatically correct?                   | Ist die Zusammenfassung grammatikalisch korrekt?             | ¿Es el resumen gramaticalmente correcto?                 | Является ли резюме грамматически правильным?            | Özet dilbilgisi açısından doğru mu?          | Bản tóm tắt có đúng ngữ pháp không?                         |
| Attrib.    | Is the summary factually consistent with the article?   | Ist die Zusammenfassung faktisch konsistent mit dem Artikel? | ¿Es el resumen factualmente consistente con el artículo? | Является ли резюме фактически согласованным со статьей? | Özet, makaleyle olgusal olarak tutarlı mı?   | Bản tóm tắt có nhất quán về mặt thực tế với bài viết không? |
| Main ideas | Does the summary capture the main ideas of the article? | Erfasst die Zusammenfassung die Hauptideen des Artikels?     | ¿El resumen captura las ideas principales del artículo?  | Отражает ли резюме основные идеи статьи?                | Özet, makalenin ana fikirlerini yansıtır mı? | Bản tóm tắt có nắm bắt được các ý chính của bài viết không? |
| Concise    | Is the summary concise?                                 | Ist die Zusammenfassung prägnant?                            | ¿Es el resumen conciso?                                  | Является ли резюме кратким?                             | Özet kısa ve öz mü?                          | Bản tóm tắt có súc tích không?                              |

Figure 4: Direct Prompting (G-eval) question per quality dimension and language.

| Lang. | Task intro                              | Art. label | Sum. label      | Answer cue                   |
|-------|-----------------------------------------|------------|-----------------|------------------------------|
| EN    | Answer with 1 (yes) or 0 (no).          | Article    | Summary         | Answer (1=yes, 0=no):        |
| DE    | Antworten Sie mit 1 (ja) oder 0 (nein). | Artikel    | Zusammenfassung | Antwort (1=ja, 0=nein):      |
| ES    | Responda con 1 (sí) o 0 (no).           | Artículo   | Resumen         | Respuesta (1=sí, 0=no):      |
| RU    | Ответьте 1 (да) или 0 (нет).            | Статья     | Резюме          | Ответ (1=да, 0=нет):         |
| TR    | 1 (evet) veya 0 (hayır) ile yanıtlayın. | Makale     | Özet            | Yanıt (1=evet, 0=hayır):     |
| VI    | Trả lời bằng 1 (có) hoặc 0 (không).     | Bài viết   | Bản tóm tắt     | Câu trả lời (1=có, 0=không): |

Figure 5: Fixed sentence-level labels used in Direct Prompting per language.

## D Dataset

Data is released under a CC-BY 4.0 license, with which we comply fully.

| System            | DE    | EN    | ES    | RU    | TR    | VI    |
|-------------------|-------|-------|-------|-------|-------|-------|
| T5-Base           | 0     | 236   | 0     | 0     | 0     | 0     |
| T5-Base (250s)    | 0     | 241   | 0     | 0     | 0     | 0     |
| T5-XXL            | 0     | 247   | 0     | 0     | 0     | 0     |
| mT5-Small         | 231   | 239   | 234   | 222   | 218   | 234   |
| mT5-Small (250s)  | 226   | 243   | 239   | 218   | 209   | 229   |
| mT5-XXL           | 228   | 239   | 242   | 208   | 208   | 238   |
| PaLM (1-shot)     | 241   | 250   | 239   | 241   | 207   | 223   |
| PaLM (fine-tuned) | 219   | 249   | 245   | 241   | 180   | 230   |
| Reference         | 228   | 239   | 240   | 213   | 210   | 238   |
| Total outputs     | 1,373 | 2,183 | 1,439 | 1,343 | 1,232 | 1,392 |

Table 11: Number of outputs samples, annotated with human judgments for the Seahorse validation set.

| System            | DE    | EN    | ES    | RU    | TR    | VI    |
|-------------------|-------|-------|-------|-------|-------|-------|
| T5-Base           | 0     | 484   | 0     | 0     | 0     | 0     |
| T5-Base (250s)    | 0     | 482   | 0     | 0     | 0     | 0     |
| T5-XXL            | 0     | 476   | 0     | 0     | 0     | 0     |
| mT5-Small         | 451   | 482   | 474   | 449   | 495   | 497   |
| mT5-Small (250s)  | 459   | 484   | 477   | 427   | 488   | 490   |
| mT5-XXL           | 443   | 481   | 479   | 450   | 485   | 492   |
| PaLM (1-shot)     | 463   | 495   | 486   | 486   | 440   | 438   |
| PaLM (fine-tuned) | 434   | 487   | 473   | 485   | 409   | 483   |
| Reference         | 435   | 484   | 465   | 439   | 486   | 494   |
| Total outputs     | 2,685 | 4,355 | 2,854 | 2,736 | 2,803 | 2,894 |

Table 12: Number of outputs samples, annotated with human judgments for the Seahorse test set.

## **E Source prompting results**

Table 13 presents the results when using a machine-translated prompt in the source language.

|                         |                             | de                                  | es                                  | ru                                  | tr                                | vi                                  |
|-------------------------|-----------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-----------------------------------|-------------------------------------|
| <b>GPTScore</b>         |                             |                                     |                                     |                                     |                                   |                                     |
| Llama-3.2-3B            | Baseline (src-prompt)       | .028 $\pm$ .125                     | .213 $\pm$ .143                     | .273 $\pm$ .065                     | .204 $\pm$ .103                   | .144 $\pm$ .093                     |
|                         | + vector steer (src-prompt) | <b>.026</b> $\pm$ .127 $\downarrow$ | <u>.215</u> $\pm$ .141 $\uparrow$   | <u>.278</u> $\pm$ .064 $\uparrow$   | <u>.212</u> $\pm$ .100 $\uparrow$ | <b>.143</b> $\pm$ .090 $\downarrow$ |
| aya-expanse-8b          | Baseline (src-prompt)       | .004 $\pm$ .145                     | .129 $\pm$ .079                     | .195 $\pm$ .065                     | .154 $\pm$ .069                   | .158 $\pm$ .068                     |
|                         | + vector steer (src-prompt) | <u>.006</u> $\pm$ .144 $\uparrow$   | <u>.137</u> $\pm$ .081 $\uparrow$   | <u>.200</u> $\pm$ .065 $\uparrow$   | <u>.160</u> $\pm$ .070 $\uparrow$ | <u>.161</u> $\pm$ .067 $\uparrow$   |
| <b>Direct Prompting</b> |                             |                                     |                                     |                                     |                                   |                                     |
| Llama-3.2-3B            | Baseline (src-prompt)       | .046 $\pm$ .039                     | .029 $\pm$ .038                     | .056 $\pm$ .035                     | .081 $\pm$ .093                   | .115 $\pm$ .050                     |
|                         | + vector steer (src-prompt) | <u>.051</u> $\pm$ .043 $\uparrow$   | <b>.028</b> $\pm$ .025 $\downarrow$ | <u>.056</u> $\pm$ .034 $\uparrow$   | <u>.083</u> $\pm$ .094 $\uparrow$ | <b>.114</b> $\pm$ .053 $\downarrow$ |
| aya-expanse-8b          | Baseline (src-prompt)       | <b>.215</b> $\pm$ .083              | <b>.244</b> $\pm$ .148              | .218 $\pm$ .108                     | .188 $\pm$ .080                   | .215 $\pm$ .056                     |
|                         | + vector steer (src-prompt) | <b>.214</b> $\pm$ .085 $\downarrow$ | <u>.244</u> $\pm$ .148 $\uparrow$   | <b>.218</b> $\pm$ .108 $\downarrow$ | <u>.189</u> $\pm$ .079 $\uparrow$ | <u>.216</u> $\pm$ .056 $\uparrow$   |

Table 13: SEAHORSE Pearson  $r$  averaged over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language, using source-language prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split. **Bold** indicates the single highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/ $\uparrow$  (**red**/ $\downarrow$ ) on steered rows indicates improvement (degradation) relative to the baseline.

## F AUC-ROC Results

Tables 14 and 15 present AUC-ROC results for the small models under English and source-language prompting respectively, corresponding to the Pearson  $r$  tables in the main paper (Table 1) and Appendix E. Table 16 shows AUC-ROC for the large-model avgvec zero-shot setting.

|                         |                            | de                                | es                                | ru                                | tr                                | vi                                |
|-------------------------|----------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| <b>GPTScore</b>         |                            |                                   |                                   |                                   |                                   |                                   |
| Llama-3.2-3B            | Baseline (EN-prompt)       | .483 $\pm$ .112                   | .631 $\pm$ .080                   | .698 $\pm$ .024                   | .628 $\pm$ .069                   | .604 $\pm$ .051                   |
|                         | + vector steer (EN-prompt) | <b>.483<math>\pm</math>.111</b> ↓ | <b>.632<math>\pm</math>.082</b> ↑ | <b>.702<math>\pm</math>.024</b> ↑ | <b>.626<math>\pm</math>.069</b> ↓ | <b>.603<math>\pm</math>.052</b> ↓ |
| aya-expanse-8b          | Baseline (EN-prompt)       | .457 $\pm$ .135                   | .597 $\pm$ .046                   | .679 $\pm$ .055                   | .625 $\pm$ .060                   | .641 $\pm$ .063                   |
|                         | + vector steer (EN-prompt) | <b>.459<math>\pm</math>.136</b> ↑ | <b>.600<math>\pm</math>.046</b> ↑ | <b>.684<math>\pm</math>.056</b> ↑ | <b>.631<math>\pm</math>.059</b> ↑ | <b>.643<math>\pm</math>.064</b> ↑ |
| <b>Direct Prompting</b> |                            |                                   |                                   |                                   |                                   |                                   |
| Llama-3.2-3B            | Baseline (EN-prompt)       | .551 $\pm$ .044                   | .497 $\pm$ .027                   | .510 $\pm$ .034                   | .499 $\pm$ .026                   | .502 $\pm$ .035                   |
|                         | + vector steer (EN-prompt) | <b>.550<math>\pm</math>.043</b> ↓ | <b>.498<math>\pm</math>.027</b> ↑ | <b>.511<math>\pm</math>.033</b> ↑ | <b>.499<math>\pm</math>.025</b> ↑ | <b>.505<math>\pm</math>.034</b> ↑ |
| aya-expanse-8b          | Baseline (EN-prompt)       | .650 $\pm$ .064                   | .669 $\pm$ .098                   | .676 $\pm$ .067                   | .662 $\pm$ .085                   | .676 $\pm$ .055                   |
|                         | + vector steer (EN-prompt) | <b>.651<math>\pm</math>.063</b> ↑ | <b>.668<math>\pm</math>.097</b> ↓ | <b>.674<math>\pm</math>.066</b> ↓ | <b>.661<math>\pm</math>.086</b> ↓ | <b>.678<math>\pm</math>.055</b> ↑ |

Table 14: SEAHORSE AUC-ROC averaged over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split. **Bold** indicates the single highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/↑ (**red**/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                         |                             | de                                | es                                | ru                                | tr                                | vi                                |
|-------------------------|-----------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| <b>GPTScore</b>         |                             |                                   |                                   |                                   |                                   |                                   |
| Llama-3.2-3B            | Baseline (src-prompt)       | .484 $\pm$ .108                   | .631 $\pm$ .083                   | .697 $\pm$ .026                   | <b>.631<math>\pm</math>.065</b>   | .607 $\pm$ .052                   |
|                         | + vector steer (src-prompt) | <b>.482<math>\pm</math>.112</b> ↓ | <b>.633<math>\pm</math>.082</b> ↑ | <b>.701<math>\pm</math>.027</b> ↑ | <b>.631<math>\pm</math>.066</b> ↓ | <b>.607<math>\pm</math>.050</b> ↓ |
| aya-expanse-8b          | Baseline (src-prompt)       | .451 $\pm$ .138                   | .582 $\pm$ .041                   | .651 $\pm$ .059                   | .600 $\pm$ .060                   | .627 $\pm$ .061                   |
|                         | + vector steer (src-prompt) | <b>.453<math>\pm</math>.137</b> ↑ | <b>.588<math>\pm</math>.043</b> ↑ | <b>.655<math>\pm</math>.060</b> ↑ | <b>.604<math>\pm</math>.059</b> ↑ | <b>.629<math>\pm</math>.062</b> ↑ |
| <b>Direct Prompting</b> |                             |                                   |                                   |                                   |                                   |                                   |
| Llama-3.2-3B            | Baseline (src-prompt)       | .533 $\pm$ .031                   | .529 $\pm$ .032                   | .550 $\pm$ .030                   | .551 $\pm$ .069                   | .575 $\pm$ .041                   |
|                         | + vector steer (src-prompt) | <b>.539<math>\pm</math>.037</b> ↑ | <b>.529<math>\pm</math>.031</b> ↑ | <b>.551<math>\pm</math>.029</b> ↑ | <b>.552<math>\pm</math>.068</b> ↑ | <b>.575<math>\pm</math>.043</b> ↓ |
| aya-expanse-8b          | Baseline (src-prompt)       | <b>.621<math>\pm</math>.030</b>   | <b>.656<math>\pm</math>.102</b>   | .631 $\pm$ .057                   | .611 $\pm$ .047                   | .631 $\pm$ .040                   |
|                         | + vector steer (src-prompt) | <b>.621<math>\pm</math>.029</b> ↑ | <b>.656<math>\pm</math>.103</b> ↓ | <b>.630<math>\pm</math>.057</b> ↓ | <b>.611<math>\pm</math>.047</b> ↑ | <b>.632<math>\pm</math>.038</b> ↑ |

Table 15: SEAHORSE AUC-ROC averaged over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language, using source-language prompts. Baseline uses  $\sigma=0$  (no steering); steered uses the best  $\sigma$  selected on the validation split. **Bold** indicates the single highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/↑ (**red**/↓) on steered rows indicates improvement (degradation) relative to the baseline.

|                       |                                | de                                | es                                | ru                                | tr                                | vi                                |
|-----------------------|--------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| <b>GPTScore</b>       |                                |                                   |                                   |                                   |                                   |                                   |
| Llama-3.1-8B-Instruct | Baseline (EN-prompt)           | .480 $\pm$ .120                   | .643 $\pm$ .077                   | .711 $\pm$ .028                   | .659 $\pm$ .059                   | .631 $\pm$ .051                   |
|                       | + avgvec zero-shot (EN-prompt) | <b>.487<math>\pm</math>.113</b> ↑ | <b>.649<math>\pm</math>.084</b> ↑ | <b>.717<math>\pm</math>.029</b> ↑ | <b>.660<math>\pm</math>.058</b> ↑ | <b>.629<math>\pm</math>.051</b> ↓ |
| aya-expanse-32b       | Baseline (EN-prompt)           | .448 $\pm$ .131                   | .603 $\pm$ .039                   | .680 $\pm$ .048                   | .626 $\pm$ .061                   | .642 $\pm$ .055                   |
|                       | + avgvec zero-shot (EN-prompt) | <b>.451<math>\pm</math>.131</b> ↑ | <b>.607<math>\pm</math>.040</b> ↑ | <b>.683<math>\pm</math>.049</b> ↑ | <b>.631<math>\pm</math>.060</b> ↑ | <b>.645<math>\pm</math>.056</b> ↑ |

Table 16: SEAHORSE AUC-ROC averaged over 6 quality dimensions (Comprehensible, Non-repetitive, Grammatical, Attributable, Main ideas, Concise) with standard deviation as subscripts (mean $\pm$ std), per language, using English prompts. Baseline uses  $\sigma=0$  (no steering); avgvec zero-shot uses a fixed steering vector averaged across all source languages with a pre-set  $\sigma$ . **Bold** indicates the single highest value per language across the entire table; underline indicates the highest value per language within each model. **Blue**/↑ (**red**/↓) on steered rows indicates improvement (degradation) relative to the baseline.