

PEAR: Pairwise Evaluation for Automatic Relative Scoring in Machine Translation

Lorenzo Proietti^{1*} Roman Grundkiewicz² Matt Post²

¹Sapienza University of Rome ²Microsoft

lproietti@diag.uniroma1.it {rogrundk,mattpost}@microsoft.com

Abstract

We present PEAR (Pairwise Evaluation for Automatic Relative Scoring), a supervised quality estimation (QE) metric family that reframes reference-free machine translation (MT) evaluation as a graded pairwise comparison. Given a source segment and two candidate translations, PEAR predicts the direction and magnitude of their quality difference. The metrics are trained using pairwise supervision derived from differences in human judgments, with an additional regularization term that encourages sign inversion under candidate order reversal. On the WMT24 meta-evaluation benchmark, PEAR outperforms strictly matched single-candidate QE baselines trained with the same data and backbones, isolating the benefit of the proposed pairwise formulation. Despite using substantially fewer parameters than recent large metrics, PEAR surpasses far larger QE models and reference-based metrics. Our analysis further indicates that PEAR yields a less redundant evaluation signal relative to other top metrics. Finally, we show that PEAR is an effective utility function for minimum Bayes risk (MBR) decoding, reducing pairwise scoring cost at negligible impact.

1 Introduction

Automatic metrics are a primary tool for comparing modern machine translation (MT) systems. Shared task evaluations and much of the research literature rely heavily on metric scores (Marie et al., 2021; Kocmi et al., 2021, 2024c,a, 2025b,a). Human evaluation remains the highest-quality signal, but it is expensive and difficult to scale across the growing number of systems, domains, and language pairs of interest (Zouhar et al., 2025a). As MT quality improves, differences between strong systems become subtle, making both human and automatic discrimination harder (Proietti et al., 2025a,b; Zouhar et al., 2025b).

*Work conducted during an internship at Microsoft.

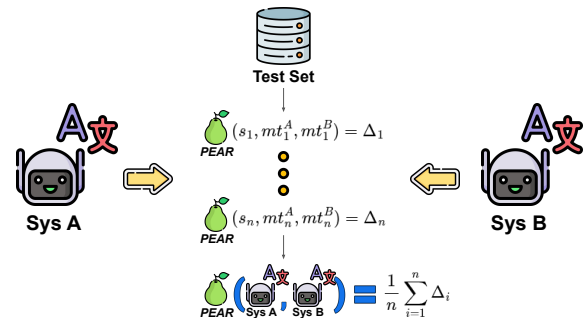


Figure 1: Segment- and system-level PEAR evaluation. For each source segment s_i , PEAR compares the system outputs (mt_i^A, mt_i^B) and predicts a relative score Δ_i . The sign indicates which translation is preferred ($\Delta_i > 0$: mt_i^A ; $\Delta_i < 0$: mt_i^B), while the magnitude $|\Delta_i|$ reflects the strength of the preference, with smaller magnitudes corresponding to weaker preferences approaching a tie. The system-level PEAR score is then the arithmetic average of the segment-level PEAR scores.

Despite their diversity, most MT evaluation metrics, including quality estimation (QE) metrics that do not require references, share a structural property: they evaluate one candidate translation at a time and output an absolute scalar score. Typically, a metric conditions on the source segment, an optional reference translation, and a single candidate translation, and outputs a scalar score. We posit that this single-candidate perspective is mismatched to several aspects of modern MT evaluation, especially in the high-quality settings where differences are subtle. In addition, one of the dominant downstream use cases is comparative, including model selection and candidate translations ranking (Kocmi et al., 2021, 2024c; Perrella et al., 2024a). Motivated by evidence that comparative human judgments can be more consistent than absolute ratings in MT and related settings (Karpinska et al., 2021; Song et al., 2025), we ask whether supervised QE for comparing translations can be better framed as a graded pairwise comparison task.

We present **PEAR** (**P**airwise **E**valuation for **A**utomatic **R**elative **S**coring), a novel supervised QE metric family that reframes reference-free MT evaluation as graded pairwise comparison. Given a source segment and two candidate translations, PEAR predicts the direction and magnitude of their quality difference. The metrics are trained on pairwise supervision constructed from differences in human segment-level judgments, with an additional regularization term that encourages sign inversion under candidate order reversal, yielding a relative score as output.

Our main contributions are as follows:

- We introduce PEAR, a novel supervised QE formulation for MT evaluation based on graded pairwise relative scoring.
- We provide controlled comparisons showing that the proposed pairwise formulation outperforms matched single-candidate QE under the same training setup, and that PEAR also surpasses much larger QE models and reference-based metrics.
- We show that PEAR avoids the need for quadratic system-to-system comparisons via a reference-anchored inference mode that does not rely on human references, remaining effective even when anchored to an MT output.
- We examine how PEAR correlates with other metrics, providing evidence of reduced redundancy.
- We demonstrate PEAR’s effectiveness and efficiency as a utility function for minimum Bayes risk (MBR) decoding.

Figure 1 illustrates the PEAR execution flow at both the segment and system levels. We release code and trained checkpoints of PEAR metrics.¹

2 Background and Related Work

Automatic MT evaluation metrics. Early MT metrics estimate quality via surface overlap against a reference translation, with BLEU (Papineni et al., 2002) and chrF (Popović, 2015) as widely used examples. Learned metrics based on pretrained representations improve agreement with human judgments by fine-tuning on human annotation signals, including reference-based models such as COMET and BLEURT (Rei et al., 2020; Selam et al., 2020a,b) and more recent large-scale

approaches such as XCOMET and MetricX (Guerreiro et al., 2024; Juraska et al., 2024). In parallel, QE metrics remove the dependency on references and score translations conditioned on the source and candidate only, including CometKiwi and its larger variants, as well as QE models of more recent learned metrics (Rei et al., 2022b, 2023; Guerreiro et al., 2024; Juraska et al., 2024). Judging approaches based on large language models, including GEMBA prompting for MQM or ESA, have shown strong performance but are often expensive and may raise reproducibility concerns (Kocmi and Federmann, 2023b,a).

Pairwise and preference-based formulations.

Pairwise evaluation has a long history in MT, including ranking-based training with engineered features and structured models (Ye et al., 2007; Duh, 2008; Guzmán et al., 2014, 2015). Several learned metrics leverage comparative supervision during training while still producing single-candidate scores at inference time, such as COMET-RANK (Rei et al., 2020) and reward-modeling approaches based on pairwise preferences (Tan and Monz, 2025). By comparison, COMET-poly incorporates additional context beyond the single translation at inference time, grounding the evaluation of a candidate in other candidates for the same source (Züfle et al., 2025). Closer to our inference setup, MT-RANKER formulates reference-free MT evaluation as binary classification: given a source and two candidates, it predicts which translation is better; consequently, its output space does not represent ties or the strength of the preference (Moosa et al., 2024). PEAR differs by predicting graded relative differences without ruling out ties by design, and encouraging sign inversion under candidate order reversal.

Comparative human judgments. Ranking-based protocols have long been used in MT (Callison-Burch et al., 2008; Bojar et al., 2016; Sakaguchi and Van Durme, 2018), and recent human evaluations increasingly adopt comparative setups, such as side-by-side MQM, motivated in part by improved consistency and agreement when annotators compare two candidates directly (Song et al., 2025). The same trend is also reflected in recent platforms for human evaluation in MT that support multi-candidate assessment (Zouhar and Kocmi, 2026). More generally, this mirrors broader observations that comparative judgments reduce subjectivity in open-ended generation evaluation

¹<https://github.com/prosho-97/pear>

(Karpinska et al., 2021). PEAR draws on this motivation but targets the automatic setting, treating comparative scoring as the prediction target for supervised QE in MT.

3 PEAR: Pairwise Evaluation for Automatic Relative Scoring

In this section, we introduce PEAR. We begin by briefly recalling the standard regression-based formulation, where a metric assigns an absolute quality score to a single candidate translation. We then describe PEAR, which instead predicts a graded quality difference between two candidate translations of the same source text, together with the model architecture, training objective, and inference procedure.

3.1 From Absolute to Relative Scoring

Supervised QE metrics are trained to predict an absolute score from a source segment s and a single candidate translation mt . In reference-based evaluation, the metric is additionally provided with a reference translation r . Accordingly, these metrics take one of the following two forms:

$$\begin{aligned}\hat{s}(s, mt) &= g_\theta(s, mt), \\ \hat{s}(s, mt, r) &= g_\theta(s, mt, r).\end{aligned}$$

They are typically optimized to fit human segment-level judgments $s_h(s, mt)^2$ via a regression loss (Sellam et al., 2020a; Rei et al., 2020, 2022b; Guerreiro et al., 2024; Juraska et al., 2024);³ other work casts automatic MT evaluation as reward modeling, learning from human preferences while still producing absolute scores for individual candidate translations at inference time (Tan and Monz, 2025). When comparing two candidates mt_a and mt_b , an implicit relative score can then be obtained by subtraction, i.e., by computing $\hat{s}(s, mt_a) - \hat{s}(s, mt_b)$ in reference-less evaluation, or $\hat{s}(s, mt_a, r) - \hat{s}(s, mt_b, r)$ in reference-based evaluation.

Rather than deriving a comparison by subtracting two independently predicted absolute scores, an alternative approach is to predict the preference directly from a joint encoding of the source

²Several recent human evaluation protocols for MT are reference-free: annotators assess system outputs only with respect to the source text, which reduces dependence on any particular reference and thereby mitigates reference bias (Freitag et al., 2021; Kocmi et al., 2022, 2024b).

³In practice, this regression loss is commonly implemented as Mean Squared Error (MSE) between the predicted scores and the human judgments.

and the two candidate translations. MT-RANKER follows this approach in the QE setting, taking (s, mt_a, mt_b) as input and predicting which candidate translation is better through a binary classification framing (Moosa et al., 2024). However, this formulation collapses comparison to a two-way decision: it cannot quantify the strength of the preference, and it rules out ties by design, since they are not part of the output label space. PEAR instead treats graded comparison as the target of prediction. Given a source text, it scores a pair of candidate translations in a shared input context, producing a signed scalar preference.

3.2 Task Formulation

Let s be a source segment and let mt_a, mt_b be two candidate translations of s . A PEAR model f_θ predicts a real-valued relative score

$$\hat{\Delta}_{ab} = f_\theta(s, mt_a, mt_b) \in \mathbb{R},$$

interpreted as the predicted quality difference between mt_a and mt_b as translations of s . Positive values favor mt_a , negative values favor mt_b , and the closer the score is to zero, the more similar the predicted quality of the two candidates becomes, with the preference approaching a tie.

PEAR is trained to approximate human pairwise differences. Specifically, given human segment-level scores $s_h(s, mt)$, we construct supervision for a candidate pair as:

$$\Delta_{ab}^* = s_h(s, mt_a) - s_h(s, mt_b).$$

At the system level, let $\{s_i\}_{i=1}^n$ be a test set, and let $\{(mt_i^A, mt_i^B)\}_{i=1}^n$ be the corresponding output pairs produced by two MT systems S_A and S_B . PEAR produces segment-level predictions $\hat{\Delta}_i = f_\theta(s_i, mt_i^A, mt_i^B)$, which are aggregated by arithmetic mean to yield a system-level graded preference:

$$\hat{\Delta}(S_A, S_B) = \frac{1}{n} \sum_{i=1}^n \hat{\Delta}_i,$$

as also illustrated in Figure 1.

3.3 Model Architecture and Input Formatting

PEAR adopts a cross-encoder setup that jointly encodes the source text and the two candidate translations. We instantiate the encoder with InfoXLM

Large⁴ for PEAR (Chi et al., 2021) and with XLM-RoBERTa-XL⁵ for PEAR-XL (Goyal et al., 2021). Throughout, we describe the scoring computation for a single input sample, from input construction to the final relative score.

Input serialization. Given a source segment s and two candidate translations mt_a and mt_b , we serialize them into a single sequence:

$$x = [\text{BOS}] s [\text{SEP}] mt_a [\text{SEP}] mt_b [\text{EOS}],$$

where [BOS], [SEP], and [EOS] denote the encoder’s special tokens.⁶ We additionally build binary span masks $\mathbf{m}_{\text{src}}, \mathbf{m}_a, \mathbf{m}_b \in \{0, 1\}^T$ over token positions in x that select the content tokens of s , mt_a , and mt_b .⁷

Encoder representations. We denote the contextual token representations extracted from the encoder’s final layer by:

$$\mathbf{H} = \text{Enc}_\phi(x) \in \mathbb{R}^{T \times d},$$

where T denotes the number of tokens in x and d is the encoder hidden dimensionality. We apply masked mean pooling to obtain span representations:

$$\mathbf{h}_\ell = \frac{\sum_{t=1}^T m_{\ell,t} \mathbf{H}_t}{\max\left(\sum_{t=1}^T m_{\ell,t}, \varepsilon\right)} \quad \text{for } \ell \in \{\text{src}, a, b\},$$

where $\mathbf{H}_t \in \mathbb{R}^d$ is the token vector at position t and $\varepsilon > 0$ avoids division by zero.

Pairwise head. We construct a source-aware representation for each candidate translation $k \in \{a, b\}$:

$$\varphi_k = [\mathbf{h}_k; \mathbf{h}_k \odot \mathbf{h}_{\text{src}}; |\mathbf{h}_k - \mathbf{h}_{\text{src}}|] \in \mathbb{R}^{3d},$$

where \odot and $|\cdot|$ are applied elementwise. Using shared parameters, we map φ_k to a scalar utility term:

$$\begin{aligned} \mathbf{a}_k &= \text{Proj}(\varphi_k) \in \mathbb{R}^d, \\ u_k &= \text{FFN}(\mathbf{a}_k) \in \mathbb{R}. \end{aligned}$$

⁴<https://huggingface.co/microsoft/infoclm-large>

⁵<https://huggingface.co/facebook/xlm-roberta-xl>

⁶For InfoXLM and XLM-R encoders, these correspond to `<s>`, `</s>`, and `</s>`, respectively.

⁷Special tokens are masked out in this process.

Proj is a linear projection followed by GELU (Hendrycks and Gimpel, 2016), layer normalization (Ba et al., 2016), and dropout, while FFN is a multilayer feed-forward head with three hidden linear layers,⁸ each followed by GELU and dropout. We then form a comparison logit by subtraction:

$$z = u_a - u_b.$$

Note that only the difference is used for supervision and inference, so the individual u_k values are not intended as absolute quality scores. Then, the final output is a scaled relative score:

$$\hat{\Delta}_{ab} = \alpha z, \quad \alpha = \text{softplus}(\alpha_{\text{raw}}) + \varepsilon,$$

where $\alpha_{\text{raw}} \in \mathbb{R}$ is a learned scalar parameter. The softplus reparameterization, together with the addition of a small constant $\varepsilon > 0$, ensures $\alpha > 0$.

3.4 Training Objective

A key property of pairwise relative scoring is sign inversion under candidate order reversal: reversing the candidate order should negate the predicted difference while preserving its magnitude, i.e., $f_\theta(s, mt_a, mt_b) = -f_\theta(s, mt_b, mt_a)$. To encourage this behavior, for each training instance (s, mt_a, mt_b) we also consider the reversed order (s, mt_b, mt_a) and denote the corresponding prediction by $\hat{\Delta}_{ba} = f_\theta(s, mt_b, mt_a)$. Specifically, we train with a Huber loss (Huber, 1964) on the human difference together with an antisymmetry regularization term:

$$\begin{aligned} \mathcal{L}_{\text{diff}} &= \ell_\delta(\hat{\Delta}_{ab} - \hat{\Delta}_{ba}^*), \\ \mathcal{L}_{\text{flip}} &= (\hat{\Delta}_{ab} + \hat{\Delta}_{ba}^*)^2, \\ \mathcal{L} &= \mathcal{L}_{\text{diff}} + \lambda_{\text{flip}} \mathcal{L}_{\text{flip}}, \end{aligned}$$

where λ_{flip} is a hyperparameter controlling the strength of the antisymmetry objective.⁹ Here ℓ_δ denotes the Huber loss:

$$\ell_\delta(r) = \begin{cases} \frac{1}{2}r^2 & \text{if } |r| \leq \delta, \\ \delta(|r| - \frac{1}{2}\delta) & \text{otherwise,} \end{cases}$$

where $\delta > 0$ is a hyperparameter that sets the residual magnitude at which the penalty transitions from quadratic to linear.¹⁰

⁸For PEAR, the hidden sizes are $512 \rightarrow 256 \rightarrow 128$; for PEAR-XL, they are $2048 \rightarrow 1024 \rightarrow 512$.

⁹Appendix A analyzes the impact of this additional training objective on antisymmetry and additive transitivity in PEAR.

¹⁰Compared to MSE, the Huber loss is less sensitive to large residuals (Huber, 1964); in Appendix B, we report an ablation where we show that it improves performance over MSE in our setting.

3.5 Inference Modes

We use PEAR in two inference configurations. In its default pairwise QE mode, it compares two candidate translations of the same source segment without access to reference translations. When a human reference is available, we additionally consider a reference-anchored mode by fixing one side of the pair to the reference.

Pairwise QE mode (PEAR). Given a source segment s and two candidate translations mt_a and mt_b , PEAR returns a relative score:

$$\hat{\Delta}_{ab} = f_{\theta}(s, mt_a, mt_b).$$

To further reduce sensitivity to the candidates’ order, we optionally also score the swapped order and combine the two predictions:

$$\begin{aligned}\hat{\Delta}_{ba} &= f_{\theta}(s, mt_b, mt_a), \\ \tilde{\Delta}_{ab} &= \frac{1}{2}(\hat{\Delta}_{ab} - \hat{\Delta}_{ba}).\end{aligned}$$

This bidirectional variant (denoted PEAR_{both}) averages the two relative scores obtained from both input orders.

Reference-anchored mode (PEAR_{ref}). When a human reference translation r is available, we anchor the comparison by fixing one side of the input to the reference:

$$\hat{\Delta}(mt, r) = f_{\theta}(s, mt, r).$$

This does not make PEAR a reference-based metric in the usual sense; rather, it instantiates the same relative-scoring function in a reference-anchored configuration. As above, we can also score the inverted order $f_{\theta}(s, r, mt)$ and combine the two predictions with the same rule.

A practical advantage of this reference-anchored configuration is computational. When evaluating N systems on an evaluation set $\{(s_i, r_i)\}_{i=1}^n$ with system outputs $\{mt_i^{(j)}\}_{i=1}^n$ for each system S_j , it yields one reference-anchored score per system:

$$\hat{\Delta}(S_j, r) = \frac{1}{n} \sum_{i=1}^n f_{\theta}(s_i, mt_i^{(j)}, r_i).$$

This requires N system scores, avoiding the $N(N - 1)/2$ system-to-system comparisons needed by fully pairwise evaluation.

4 Experimental Setup

We now summarize the empirical setting for our main experiments. We describe the training and evaluation procedures and report the trained PEAR models. Details about data, training hyperparameters, and baselines are reported in Appendix C.

4.1 Training and Evaluation Data

WMT supervision. We train PEAR on WMT human evaluation data, converting absolute segment-level human assessments into pairwise supervision via score differencing (Section 3.2). Our training data combines DA, DA+SQM, and MQM annotations, with MQM offering the most fine-grained supervision. Following common practices in MT metrics training, we adopt a two-stage schedule (Rei et al., 2022a; Juraska et al., 2023, 2024; Perrella et al., 2024b; Proietti et al., 2025b). In the first training stage, we pre-train on DA and DA+SQM judgments released in the WMT evaluation campaigns from WMT16 to WMT23, providing broad coverage across language directions and translation quality. In the second training stage, we fine-tune exclusively on MQM supervision from WMT20 to WMT23,¹¹ additionally including the IndicMT Eval MQM dataset for English→Indic directions (Sai B et al., 2023), which has been used to train XCOMET (Guerreiro et al., 2024), a metric included in our comparisons.

Scaling via distilled MQM supervision. MQM gold data are available for relatively few language pairs and are expensive to collect. To stress-test whether the proposed pairwise QE framing scales favorably with additional supervision, we also test PEAR models whose second training stage is augmented with MQM annotations distilled from GPT-4.1-mini (OpenAI, 2025) on language pairs not covered by MQM, prompting it with a GEMBA-MQM V2 approach (Junczys-Dowmunt, 2025). Appendix C provides additional details on these data.

Evaluation benchmark. Our primary benchmark is the MQM test set released with the WMT24 Metrics Shared Task (Freitag et al., 2024). We adopt their official evaluation toolkit.¹² Following

¹¹This is done not only because MQM offers the most fine-grained MT quality assessment, but also to align the training signal with our main target evaluation setting, since WMT24 meta-evaluation is centered on MQM.

¹²<https://github.com/google-research/mt-metrics-eval>

the WMT24 setup, we use Soft Pairwise Accuracy (SPA) at the system level (Thompson et al., 2024) and pairwise accuracy with tie calibration (acc_{eq}^*) at the segment level (Deutsch et al., 2023). Since PEAR produces pairwise scores by design, we integrate it with the toolkit by producing segment-level scores for each system pair. We average SPA and acc_{eq}^* over the three language pairs featuring human MQM annotations,¹³ and we also report average correlation (Avg Corr), defined as the mean of these two averages.

4.2 Trained Models

PEAR models. We train and evaluate two PEAR variants: **PEAR**, instantiated with In-foXLM Large (Chi et al., 2021), and **PEAR-XL**, instantiated with XLM-RoBERTa-XL (Goyal et al., 2021).

Matched absolute-scoring baselines. To disentangle the effect of the proposed pairwise QE formulation from backbone capacity and training data exposure, we also train matched absolute-scoring QE baselines (**Single-QE** and **Single-QE-XL**) that take only the source segment and a single candidate translation as input, share the same encoder backbones and training data, and predict an absolute scalar quality score for that candidate. These baselines let us validate the proposed pairwise formulation at the methodological level, testing it as an alternative to conventional absolute-score QE under matched training conditions.

5 Results and Analyses

This section reports empirical evidence for PEAR across several settings. We start with reporting results on the WMT24 Metrics Shared Task benchmark, and continue by showing additional analyses.

5.1 Pairwise vs. Single-Candidate QE at Scale

We separate the effect of the proposed pairwise QE formulation from confounding factors such as backbone capacity, training data exposure, and model selection comparing PEAR against strictly matched single-candidate QE baselines trained with the same data, hyperparameters, and backbone model. Checkpoint selection is based on held-out WMT23 MQM data,¹⁴ and results are reported on

¹³English–German, English–Spanish, and Japanese–Chinese.

¹⁴The WMT23 MQM dataset includes English–German, Chinese–English, and Hebrew–English.

Metric	θ	SPA	acc_{eq}^*	Avg Corr	Stat Sign Rank
Single-QE-XL _{KD}	3.5B	80.9	57.9	69.4	2
Single-QE-XL	3.5B	80.4	57.6	69.0	3
Single-QE _{KD}	560M	80.6	57.4	69.0	3
Single-QE	560M	80.0	57.2	68.6	4
PEAR-XL _{both,KD}	3.5B	82.1	58.1	70.1	1
PEAR-XL _{KD}	3.5B	82.0	58.2	70.1	1
PEAR-XL _{both}	3.5B	81.4	58.2	69.8	1
PEAR-XL	3.5B	81.5	58.1	69.8	1
PEAR _{both,KD}	560M	81.9	58.1	70.0	1
PEAR _{KD}	560M	81.8	58.2	70.0	1
PEAR _{both}	560M	81.2	58.0	69.6	2
PEAR	560M	80.9	57.9	69.4	2

Table 1: Controlled comparison of PEAR against matched single-candidate QE baselines on the WMT24 MQM test set. SPA and acc_{eq}^* are averaged over English–German, English–Spanish, and Japanese–Chinese; Avg Corr is the mean of these two averages. Stat Sign Rank reports rankings computed with the PERM-BOTH hypothesis test (Deutsch et al., 2021), following the WMT24 Metrics Shared Task setup (Freitag et al., 2024). Bold indicates the best score in each column. For PEAR, *both* denotes the bidirectional pairwise QE configuration. *KD* denotes models fine-tuned with additional MQM supervision distilled from GPT-4.1-mini.

the WMT24 MQM benchmark.

We report results for PEAR in its default configuration (one input order) and for the bidirectional variant that combines predictions from both input orders, matching the inference mode described in Section 3.5. We further consider settings with and without knowledge distillation augmentation in the second training stage, to probe whether the advantage of our pairwise QE formulation persists—and potentially widens—as the amount of MQM supervision is scaled up via distilled annotations.

Table 1 reports the comparison results. PEAR yields better performance than the matched single-candidate baselines, indicating that the proposed pairwise QE framing improves over single-candidate QE for comparing candidate translations. The statistical significance analysis further supports this conclusion: each PEAR model outperforms its corresponding Single-QE counterpart with a better statistical significance rank. In addition, the near-identical performance of the single-order and *both* variants suggests that PEAR largely satisfies sign inversion under candidate order reversal.

5.2 Comparison with MT-RANKER

MT-RANKER is the closest prior approach to PEAR in terms of formulation, but it predicts a strict pairwise preference and does not model ties (Moosa et al., 2024). To enable a fair comparison, we repurpose the WMT24 MQM test set into a

Group	Metric	θ	Ref?	SPA	acc_{eq}^*	Avg Corr	Stat Sign Rank
WMT24 Metrics	MetricX-24-Hybrid-QE-XXL	13B	×	84.9	58.0	71.4	1
	MetricX-24-Hybrid-QE-XL	3.7B	×	83.4	56.5	69.9	3
	MetricX-24-Hybrid-QE-Large	1.2B	×	80.6	56.1	68.3	6
	XCOMET-QE	24B [‡]	×	83.3	55.7	69.5	4
	CometKiwi-XXL	10.5B	×	85.4	55.2	70.3	3
	CometKiwi	560M	×	73.3	54.7	64.0	7
	GEMBA-ESA	GPT-4	×	84.6	57.6	71.1	1
	MetricX-24-Hybrid-Large	1.2B	✓	84.0	57.0	70.5	2
	COMET-22	560M	✓	82.4	55.4	68.9	5
	BLEURT-20	579M	✓	82.1	55.0	68.6	5
Ours	PEAR-XL _{ref,KD}	3.5B	✓ [†]	82.7	58.9	70.8	1
	PEAR-XL _{both,KD}	3.5B	×	82.7	58.3	70.5	2
	PEAR-XL _{ref}	3.5B	✓ [†]	81.6	59.2	70.4	2
	PEAR-XL _{both}	3.5B	×	81.6	58.8	70.2	3
	PEAR _{ref,KD}	560M	✓ [†]	82.6	58.7	70.7	2
	PEAR _{both,KD}	560M	×	82.0	58.8	70.4	2
	PEAR _{ref}	560M	✓ [†]	81.5	59.0	70.3	3
	PEAR _{both}	560M	×	81.2	59.0	70.1	3

Table 2: WMT24 MQM meta-evaluation. SPA and acc_{eq}^* are averaged over English–German, English–Spanish, and Japanese–Chinese; Avg Corr is the mean of these two averages. Stat Sign Rank reports rankings computed with the PERM-BOTH hypothesis test (Deutsch et al., 2021), following the WMT24 Metrics Shared Task setup (Freitag et al., 2024). Bold indicates the best score in each column. † denotes an ensemble that combines two 10.7B checkpoints and one 3.5B checkpoint. For PEAR, *ref* (†) denotes the reference-anchored configuration, *both* denotes the bidirectional pairwise QE configuration, and *KD* denotes models fine-tuned with additional MQM supervision distilled from GPT-4.1-mini.

Metric	θ	En-De	En-Es	Ja-Zh	Avg Pair Acc
PEAR _{KD}	560M	69.7	68.9	68.2	68.9
PEAR	560M	69.2	67.5	67.2	68.0
MT-RANKER-XXL	5.7B	64.7	65.0	67.8	65.8

Table 3: Comparison with MT-RANKER-XXL on the subset of the WMT24 MQM test set obtained after removing translation pairs with tied human MQM assessments.

binary classification test set by removing all candidate translation pairs whose human MQM assessments are tied. This yields 49,666 pairs for English–German, 24,777 for English–Spanish, and 34,045 for Japanese–Chinese. We then compare PEAR and PEAR_{KD} in their default single-order configuration against the best-performing MT-RANKER model, i.e., MT-RANKER-XXL. Table 3 reports segment-level pairwise accuracy on non-tied pairs across the three WMT24 language pairs featuring human MQM annotations, together with macro-average. Both PEAR models outperform MT-RANKER-XXL while using substantially fewer parameters.

5.3 WMT24 Results

Table 2 reports results on the WMT24 meta-evaluation benchmark. For PEAR, we report only the bidirectional configuration (*both*), since Table 1

shows that single-pass and *both* yield very similar performance. All PEAR variants in Table 2 are trained with the same hyperparameters used in the controlled comparison with the only difference that we also include WMT23 data in training.

Better performance with fewer parameters

Among reference-free metrics, PEAR attains high average correlation with substantially smaller models. PEAR-XL_{both} (3.5B) exceeds MetricX-24-Hybrid-QE-XL (3.7B) and XCOMET-QE (24B) (70.2 vs. 69.9 and 69.5), while using roughly $7\times$ fewer parameters than XCOMET-QE. Distilled supervision yields a further improvement, raising average correlation to 70.5.

The same trend holds with the smaller PEAR models. PEAR_{both} (560M) remains higher than MetricX-24-Hybrid-QE-XL (3.7B) and XCOMET-QE (24B) on average correlation (70.1 vs. 69.9 and 69.5), despite using about $7\times$ and $40\times$ fewer parameters, respectively. The only QE metric comparable in size to PEAR is CometKiwi (560M), yet it yields a markedly lower result (64.0 vs. 70.1). With distilled supervision, PEAR_{both,KD} also edges out CometKiwi-XXL (70.4 vs. 70.3), while being about $20\times$ smaller (560M vs. 10.5B).

Comparison against reference-based metrics

Even without references, PEAR compares favor-

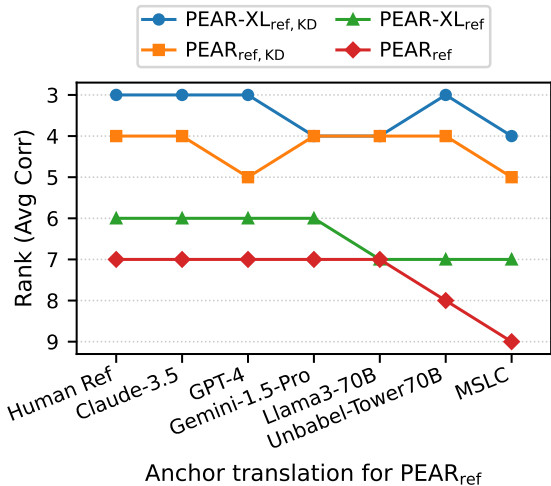


Figure 2: Rank stability of PEAR_{ref} across anchors. The leftmost anchor (*Human Ref*) uses the human reference as the fixed comparison translation, matching Table 2. The remaining anchors use an MT output; for each MT anchor, that system is removed from the benchmark before recomputing meta-evaluation. Ranks are computed by Avg Corr (lower is better).

ably to reference-based metrics. For example, $\text{PEAR-XL}_{\text{both,KD}}$ matches average correlation of MetricX-24-Hybrid-Large (70.5 vs. 70.5), and $\text{PEAR}_{\text{both}}$ exceeds COMET-22 and BLEURT-20 (70.1 vs. 68.9 and 68.6).

5.4 Does PEAR_{ref} Require References?

Table 2 shows that the reference-anchored configuration, PEAR_{ref} , slightly outperforms its corresponding QE (*both*) variant on WMT24. We now test whether these performances depend on anchoring to a human reference, or whether the same behavior holds when the anchor is an MT output.

We instantiate the anchor slot with (i) the human reference (*Human Ref*) and (ii) the output of six MT systems. For each MT anchor, we re-run WMT24 meta-evaluation after removing the anchor system from the benchmark, since its outputs are used as the fixed comparison target. Among the six MT anchors, we also include MSLC as a more extreme stress test: in the English–German MQM human evaluation reported by Freitag et al. (2024), for example, Llama3-70B has a system-level MQM score of -3.62, whereas MSLC obtains -13.46, making it the lowest-quality system by a large margin.

Figure 2 shows that the ranks¹⁵ of PEAR_{ref} vari-

¹⁵Since the set of evaluated systems changes with the chosen MT anchor, the resulting Avg Corr values are not directly comparable across anchors.

ants are stable across anchors. This picture remains largely unchanged even when PEAR_{ref} is anchored to MSLC, a substantially lower-quality MT system according to the MQM human evaluation results of Freitag et al. (2024). This suggests that PEAR_{ref} does not rely on human references to retain its performance.

5.5 Segment-level Correlation Between PEAR and Other Metrics

PEAR is trained with a pairwise relative-scoring objective, unlike most supervised metrics that use single-candidate regression. To assess how distinct PEAR’s evaluation signal is, we analyze its segment-level correlation with other metrics.

For this analysis, we compute the Pearson correlation over pairwise segment-level difference scores. For PEAR, we use its predicted pairwise scores directly. For single-candidate metrics m that assign a scalar score to each candidate translation (i.e., $m(s, mt) \in \mathbb{R}$), we obtain pairwise differences by subtracting their segment-level scores for each MT system pair: $\Delta_m(s, mt_a, mt_b) = m(s, mt_a) - m(s, mt_b)$.¹⁶ We compute correlations separately for each WMT24 MQM language pair.

Figures 3, 4, and 5 in Appendix D show the correlation matrices for English–German, English–Spanish, and Japanese–Chinese. Across all language pairs, $\text{PEAR}_{\text{both}}$ and $\text{PEAR}_{\text{both,KD}}$ exhibit consistently lower correlation with other top-performing WMT24 metrics compared to supervised and LLM-as-a-Judge approaches, suggesting that PEAR may provide a distinct evaluation signal rather than mirroring existing metrics. For example, the correlation of these two PEAR variants with MetricX-24-Hybrid-QE¹⁷ is moderate on English–German ($r \approx 0.71$), drops on English–Spanish ($r \approx 0.51$), and is much lower on Japanese–Chinese ($r \approx 0.26$). Excluding lexical and unsupervised baselines, i.e., BLEU (Papineni et al., 2002), chrF (Popović, 2015), and BERTScore (Zhang et al., 2020), PEAR metrics are the least correlated with the rest of the metric suite. Importantly, this does not reflect weaker metric performance, but rather suggests a different evaluation behavior. Such lower correlation can be useful in MT fine-tuning, where diverse metrics help avoid over-optimizing

¹⁶For simplicity, we use reference-free notation here. In the correlation analysis, we also include reference-based metrics that additionally condition on a human reference.

¹⁷It denotes the XXL checkpoint.

Utility	LP	XCOMET-XL	CometKiwi-XL	MetricX-XL
PEAR (full)	En-De	0.855	0.731	-5.2
PEAR (sym.)	En-De	0.854	0.731	-5.4
COMET-22	En-De	0.844	0.730	-6.6
BLEURT-20	En-De	0.842	0.728	-6.3
PEAR (full)	En-Ja	0.810	0.685	-6.4
PEAR (sym.)	En-Ja	0.809	0.685	-6.7
COMET-22	En-Ja	0.798	0.684	-7.8
BLEURT-20	En-Ja	0.796	0.683	-6.9

Table 4: MBR decoding on 100-best lists for WMT24 En-De and En-Ja. PEAR (sym.) computes only one triangle of the $N \times N$ utility matrix (with N candidates; here $N = 100$) and fills the remainder by antisymmetry.

for a single signal. We leave a deeper investigation of the sources and implications of these correlations to future work.

5.6 PEAR for MBR Decoding

Reference-based metrics such as COMET-22 (Rei et al., 2020; Freitag et al., 2024) and BLEURT-20 (Sellam et al., 2020b) are trained to condition on human references, whereas in MBR decoding they are applied to candidate-to-candidate comparisons (Freitag et al., 2022), creating a slight mismatch between training and use. PEAR, by contrast, is trained explicitly for reference-free pairwise comparison, making it naturally suited for this setting. Here, we compare PEAR, COMET-22, and BLEURT-20 as MBR utility functions.

Candidate lists. We translate the English source side of the WMT24 test set into German and Japanese with a transformer-based multilingual MT system trained on public and internal data with teacher-student knowledge distillation. We decode with beam search and retain a 100-best list for each source segment.

MBR utilities. We run MBR over each 100-best list using three utilities: PEAR,¹⁸ and the reference-based metrics COMET-22 and BLEURT-20.¹⁹ For PEAR, we compare two implementations. The first computes the full utility matrix, while the second exploits antisymmetry and evaluates only one triangle of the matrix, using $u(h_j, h_i) = -u(h_i, h_j)$ to fill the other half, thereby halving the number of forward passes. This tests whether PEAR behaves as an approximately antisymmetric utility during decoding, while en-

¹⁸InfoXLM Large checkpoint fine-tuned without distilled supervision. We use the default inference configuration, not the *both* variant.

¹⁹Their parameter counts are comparable to PEAR’s.

abling a cheaper MBR procedure by reducing the number of required pairwise comparisons.

Results. We evaluate the resulting MBR outputs with three metrics: XCOMET-XL, CometKiwi-XL, and MetricX-24-Hybrid-XL. Table 4 shows that PEAR yields nearly identical performance under the full and antisymmetry-reduced MBR implementations, indicating that the antisymmetry shortcut is effective in practice. Moreover, PEAR-based MBR improves over COMET-22 and BLEURT-20 under XCOMET-XL and MetricX-24-Hybrid-XL, while gains under CometKiwi-XL are relatively smaller. These results suggest that PEAR, trained explicitly for pairwise comparison, can be an effective utility for MBR decoding.

6 Conclusion

We presented PEAR, a supervised QE metric family that models MT evaluation as graded pairwise relative scoring. Unlike standard metrics that score one candidate at a time, PEAR compares two translations jointly and predicts both the direction and strength of preference, which mirrors the way metrics are often used in practice, namely to decide which MT system or candidate is better. In controlled experiments, PEAR consistently outperforms matched single-candidate QE baselines, indicating that the pairwise formulation is a more effective approach for comparing candidate translations.

On WMT24 MQM meta-evaluation, PEAR achieves higher correlation than the largest QE submissions, despite its smaller model size. It further outperforms reference-based metrics, despite not relying on human references. We also showed that PEAR is an effective utility function for minimum Bayes risk (MBR) decoding, enabling a substantially cheaper pairwise scoring procedure with negligible impact on decoding quality. We further found that PEAR_{ref} remains effective when the anchor translation is an MT output rather than a human reference, and that PEAR’s segment-level pairwise scores are less correlated with other top metrics, suggesting that PEAR captures a different evaluation signal.

Future work will focus on understanding which phenomena drive this divergence, and on leveraging pairwise relative scoring in other evaluation settings.

Limitations

Model scale. Our experiments do not fully characterize how PEAR performance scales with substantially larger models. The largest PEAR checkpoint we fine-tune in this work is PEAR-XL (3.5B parameters), and we therefore do not test whether the performance gains we attribute to the pairwise formulation persist, widen, or saturate with larger models. Recent works suggest that scaling up the underlying model and fine-tuning larger backbones can be a strong driver of performance for supervised metrics (Rei et al., 2023; Guerreiro et al., 2024; Juraska et al., 2024, 2025; Tan and Monz, 2025). Extending PEAR to larger model sizes is a natural next step. In addition, the pairwise setup could be extended from scalar relative scoring to MQM sequence tagging that explicitly predicts error spans in each candidate translation (Perrera et al., 2022; Kocmi and Federmann, 2023a; Junczys-Dowmunt, 2025), which is well aligned with side-by-side MQM protocols based on comparative judgment (Song et al., 2025).

Metric divergence. In Section 5.5, we report that PEAR exhibits lower correlation with other strong WMT24 metrics than is typical among top-performing metrics. In this work, we treat this primarily as evidence that PEAR may induce a complementary evaluation signal, but we do not deeply investigate how this could be exploited. For example, lower correlation could be beneficial when combining metrics in an ensemble, or when other metrics are used as optimization targets, such as rewards in reinforcement learning, since it may reduce over-optimization to a single signal and support fairer automatic assessment, but it could also indicate sensitivity to specific phenomena that warrant targeted diagnosis. A more detailed analysis of this divergence is left for future work.

Targeted synthetic data. Unlike some recent metrics that incorporate targeted synthetic examples designed to address known failure modes (e.g., fluent but unrelated translations, undertranslation, and related specific errors), PEAR is not fine-tuned with comparable hand-designed perturbation suites (Juraska et al., 2023, 2024; Guerreiro et al., 2024). Although we experiment with scaling supervision via distilled MQM annotations, we do not study whether PEAR particularly benefits from synthetic pair construction that is explicitly contrastive by design. PEAR’s pairwise formulation makes tar-

geted contrastive pair construction more natural and potentially more cost-effective than for other traditional supervised metrics. We leave a systematic investigation of such targeted synthetic data strategies, and how they interact with training on graded pairwise comparisons, for future work.

Ethics Statement

We foresee no ethical issues with our work.

Acknowledgements

This work was carried out while Lorenzo Proietti was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome. The authors thank Marcin Junczys-Dowmunt for his valuable help with the knowledge distillation process from GPT-4.1-mini.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer Normalization](#). *Preprint*, arXiv:1607.06450.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névéal, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 Conference on Machine Translation](#). in *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Chris Callison-Burch, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2008. [Further Meta-Evaluation of Machine Translation](#). in *Proceedings of the Third Workshop on Statistical Machine Translation*, pages 70–106, Columbus, Ohio. Association for Computational Linguistics.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training](#). in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2021. [A Statistical Analysis of Summarization Evaluation Metrics Using Resampling Methods](#). *Transactions of the Association for Computational Linguistics*, 9:1132–1146.

- Daniel Deutsch, George Foster, and Markus Freitag. 2023. [Ties Matter: Meta-Evaluating Modern Metrics with Pairwise Accuracy and Tie Calibration](#). in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12914–12929, Singapore. Association for Computational Linguistics.
- Kevin Duh. 2008. [Ranking vs. Regression in Machine Translation Evaluation](#). in *Proceedings of the Third Workshop on Statistical Machine Translation*, pages 191–194, Columbus, Ohio. Association for Computational Linguistics.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, Errors, and Context: A Large-Scale Study of Human Evaluation for Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. [High Quality Rather than High Model Probability: Minimum Bayes Risk Decoding with Neural Metrics](#). *Transactions of the Association for Computational Linguistics*, 10:811–825.
- Markus Freitag, Nitika Mathur, Daniel Deutsch, Chikiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Frederic Blain, Tom Kocmi, Jiayi Wang, David Ifeoluwa Adelani, Marianna Buchicchio, Chrysoula Zerva, and Alon Lavie. 2024. [Are LLMs Breaking MT Metrics? Results of the WMT24 Metrics Shared Task](#). in *Proceedings of the Ninth Conference on Machine Translation*, pages 47–81, Miami, Florida, USA. Association for Computational Linguistics.
- Naman Goyal, Jingfei Du, Myle Ott, Giri Anantharaman, and Alexis Conneau. 2021. [Larger-Scale Transformers for Multilingual Masked Language Modeling](#). in *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 29–33, Online. Association for Computational Linguistics.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. [xCOMET: Transparent Machine Translation Evaluation through Fine-grained Error Detection](#). *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Francisco Guzmán, Shafiq Joty, Lluís Màrquez, Alessandro Moschitti, Preslav Nakov, and Massimo Nicosia. 2014. [Learning to Differentiate Better from Worse Translations](#). in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 214–220, Doha, Qatar. Association for Computational Linguistics.
- Francisco Guzmán, Shafiq Joty, Lluís Màrquez, and Preslav Nakov. 2015. [Pairwise Neural Machine Translation Evaluation](#). in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 805–814, Beijing, China. Association for Computational Linguistics.
- Dan Hendrycks and Kevin Gimpel. 2016. [Gaussian Error Linear Units \(GELUs\)](#). *arXiv preprint arXiv:1606.08415*.
- Peter J. Huber. 1964. [Robust Estimation of a Location Parameter](#). *The Annals of Mathematical Statistics*, 35(1):73–101.
- Marcin Junczys-Dowmunt. 2025. [GEMBA V2: Ten Judgments Are Better Than One](#). in *Proceedings of the Tenth Conference on Machine Translation*, pages 926–933, Suzhou, China. Association for Computational Linguistics.
- Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. 2024. [MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task](#). in *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.
- Juraj Juraska, Tobias Domhan, Mara Finkelstein, Tetsuji Nakagawa, Geza Kovacs, Daniel Deutsch, Pidong Wang, and Markus Freitag. 2025. [MetricX-25 and GemSpanEval: Google Translate Submissions to the WMT25 Evaluation Shared Task](#). in *Proceedings of the Tenth Conference on Machine Translation*, pages 957–968, Suzhou, China. Association for Computational Linguistics.
- Juraj Juraska, Mara Finkelstein, Daniel Deutsch, Aditya Siddhant, Mehdi Mirzazadeh, and Markus Freitag. 2023. [MetricX-23: The Google Submission to the WMT 2023 Metrics Shared Task](#). in *Proceedings of the Eighth Conference on Machine Translation*, pages 756–767, Singapore. Association for Computational Linguistics.
- Marzena Karpinska, Nader Akoury, and Mohit Iyyer. 2021. [The Perils of Using Mechanical Turk to Evaluate Open-Ended Text Generation](#). in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1265–1285, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tom Kocmi, Sweta Agrawal, Ekaterina Artemova, Eleftherios Avramidis, Eleftheria Briakou, Pinzhen Chen, Marzieh Fadaee, Markus Freitag, Roman Grundkiewicz, Yupeng Hou, Philipp Koehn, Julia Kreutzer, Saab Mansour, Stefano Perrella, Lorenzo Proietti, Parker Riley, Eduardo Sánchez, Patricia Schmidtova, Mariya Shmatova, and Vilém Zouhar. 2025a. [Findings of the WMT25 Multilingual Instruction Shared Task: Persistent Hurdles in Reasoning, Generation, and Evaluation](#). in *Proceedings of the Tenth Conference on Machine Translation*, pages 414–435, Suzhou, China. Association for Computational Linguistics.

- Tom Kocmi, Ekaterina Artemova, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Konstantin Dranch, Anton Dvorkovich, Sergey Dukanov, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Howard Lakouagna, Jessica Lundin, Christof Monz, Kenton Murray, Masaaki Nagata, Stefano Perrella, Lorenzo Proietti, Martin Popel, Maja Popović, Parker Riley, Mariya Shmatova, Steinhórfur Steingrímsson, Lisa Yankovskaya, and Vilém Zouhar. 2025b. [Findings of the WMT25 General Machine Translation Shared Task: Time to Stop Evaluating on Easy Test Sets](#). in *Proceedings of the Tenth Conference on Machine Translation*, pages 355–413, Suzhou, China. Association for Computational Linguistics.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popović, Mariya Shmatova, Steinhórfur Steingrímsson, and Vilém Zouhar. 2024a. [Findings of the WMT24 General Machine Translation Shared Task: The LLM Era Is Here but MT Is Not Solved Yet](#). in *Proceedings of the Ninth Conference on Machine Translation*, pages 1–46, Miami, Florida, USA. Association for Computational Linguistics.
- Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 Conference on Machine Translation \(WMT22\)](#). in *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023a. [GEMBA-MQM: Detecting Translation Quality Error Spans with GPT-4](#). in *Proceedings of the Eighth Conference on Machine Translation*, pages 768–775, Singapore. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023b. [Large Language Models Are State-of-the-Art Evaluators of Translation Quality](#). in *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To Ship or Not to Ship: An Extensive Evaluation of Automatic Metrics for Machine Translation](#). in *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.
- Tom Kocmi, Vilém Zouhar, Eleftherios Avramidis, Roman Grundkiewicz, Marzena Karpinska, Maja Popović, Mrinmaya Sachan, and Mariya Shmatova. 2024b. [Error Span Annotation: A Balanced Approach for Human Evaluation of Machine Translation](#). in *Proceedings of the Ninth Conference on Machine Translation*, pages 1440–1453, Miami, Florida, USA. Association for Computational Linguistics.
- Tom Kocmi, Vilém Zouhar, Christian Federmann, and Matt Post. 2024c. [Navigating the Metrics Maze: Reconciling Score Magnitudes and Accuracies](#). in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1999–2014, Bangkok, Thailand. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#). in *International Conference on Learning Representations*.
- Benjamin Marie, Atsushi Fujita, and Raphael Rubino. 2021. [Scientific Credibility of Machine Translation Research: A Meta-Evaluation of 769 Papers](#). in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7297–7306, Online. Association for Computational Linguistics.
- Ibraheem Muhammad Moosa, Rui Zhang, and Wenpeng Yin. 2024. [MT-Ranker: Reference-free machine translation evaluation by inter-system ranking](#). in *The Twelfth International Conference on Learning Representations*.
- OpenAI. 2025. [Introducing GPT-4.1 in the API](https://openai.com/index/gpt-4-1/). <https://openai.com/index/gpt-4-1/>. Accessed: 2026-04-19.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a Method for Automatic Evaluation of Machine Translation](#). in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Stefano Perrella, Lorenzo Proietti, Pere-Lluís Huguet Cabot, Edoardo Barba, and Roberto Navigli. 2024a. [Beyond Correlation: Interpretable Evaluation of Machine Translation Metrics](#). in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 20689–20714, Miami, Florida, USA. Association for Computational Linguistics.
- Stefano Perrella, Lorenzo Proietti, Alessandro Scirè, Edoardo Barba, and Roberto Navigli. 2024b. [Guardians of the Machine Translation Meta-Evaluation: Sentinel Metrics Fall In!](#) in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16216–16244, Bangkok, Thailand. Association for Computational Linguistics.

- Stefano Perrella, Lorenzo Proietti, Alessandro Scirè, Niccolò Campolungo, and Roberto Navigli. 2022. **MaTESe: Machine Translation Evaluation as a Sequence Tagging Problem**. in *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 569–577, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Maja Popović. 2015. **chrF: character n-gram F-score for automatic MT evaluation**. in *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Lorenzo Proietti, Stefano Perrella, and Roberto Navigli. 2025a. **Has Machine Translation Evaluation Achieved Human Parity? The Human Reference and the Limits of Progress**. in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 790–813, Vienna, Austria. Association for Computational Linguistics.
- Lorenzo Proietti, Stefano Perrella, Vilém Zouhar, Roberto Navigli, and Tom Kocmi. 2025b. **Estimating Machine Translation Difficulty**. in *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 24261–24285, Suzhou, China. Association for Computational Linguistics.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. **COMET-22: Unbabel-IST 2022 Submission for the Metrics Shared Task**. in *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ricardo Rei, Nuno M. Guerreiro, José Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André F. T. Martins. 2023. **Scaling up CometKiwi: Unbabel-IST 2023 Submission for the Quality Estimation Shared Task**. in *Proceedings of the Eighth Conference on Machine Translation*, pages 841–848, Singapore. Association for Computational Linguistics.
- 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.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022b. **CometKiwi: IST-Unbabel 2022 Submission for the Quality Estimation Shared Task**. in *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ananya Sai B, Tanay Dixit, Vignesh Nagarajan, Anoop Kunchukuttan, Pratyush Kumar, Mitesh M. Khapra, and Raj Dabre. 2023. **IndicMT Eval: A Dataset to Meta-Evaluate Machine Translation Metrics for Indian Languages**. in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14210–14228, Toronto, Canada. Association for Computational Linguistics.
- Keisuke Sakaguchi and Benjamin Van Durme. 2018. **Efficient Online Scalar Annotation with Bounded Support**. in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 208–218, Melbourne, Australia. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020a. **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.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur Parikh. 2020b. **Learning to Evaluate Translation Beyond English: BLEURT Submissions to the WMT Metrics 2020 Shared Task**. in *Proceedings of the Fifth Conference on Machine Translation*, pages 921–927, Online. Association for Computational Linguistics.
- Yixiao Song, Parker Riley, Daniel Deutsch, and Markus Freitag. 2025. **Enhancing Human Evaluation in Machine Translation with Comparative Judgement**. in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20536–20551, Vienna, Austria. Association for Computational Linguistics.
- Shaomu Tan and Christof Monz. 2025. **ReMedy: Learning Machine Translation Evaluation from Human Preferences with Reward Modeling**. in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4370–4387, Suzhou, China. Association for Computational Linguistics.
- Brian Thompson, Nitika Mathur, Daniel Deutsch, and Huda Khayrallah. 2024. **Improving Statistical Significance in Human Evaluation of Automatic Metrics via Soft Pairwise Accuracy**. in *Proceedings of the Ninth Conference on Machine Translation*, pages 1222–1234, Miami, Florida, USA. Association for Computational Linguistics.
- Yang Ye, Ming Zhou, and Chin-Yew Lin. 2007. **Sentence Level Machine Translation Evaluation as a Ranking**. in *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 240–247, Prague, Czech Republic. 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*.

Vilém Zouhar, Tom Kocmi, and Mrinmaya Sachan. 2025a. **AI-Assisted Human Evaluation of Machine Translation**. in *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4936–4950, Albuquerque, New Mexico. Association for Computational Linguistics.

Vilém Zouhar, Peng Cui, and Mrinmaya Sachan. 2025b. **How to Select Datapoints for Efficient Human Evaluation of NLG Models?** *Preprint*, arXiv:2501.18251.

Vilém Zouhar and Tom Kocmi. 2026. **Pearmut: Human Evaluation of Translation Made Trivial**. *Preprint*, arXiv:2601.02933.

Maike Züfle, Vilém Zouhar, Tu Anh Dinh, Felipe Maia Polo, Jan Niehues, and Mrinmaya Sachan. 2025. **COMET-poly: Machine Translation Metric Grounded in Other Candidates**. in *Proceedings of the Tenth Conference on Machine Translation*, pages 887–904, Suzhou, China. Association for Computational Linguistics.

A Pairwise Consistency of PEAR: Antisymmetry and Transitivity

PEAR is trained to assess graded quality differences between pairs of translations (Section 3). We focus on two consistency properties of relative scoring: *antisymmetry* and *transitivity*. Antisymmetry captures sign inversion under candidate reversal, whereas transitivity constrains score composition over triples. In this section, we quantify PEAR’s adherence to these two consistency properties and how the antisymmetry auxiliary training objective (Section 3.4) affects both.

Setup. We run this analysis on the WMT24 MQM benchmark across English–German, English–Spanish, and Japanese–Chinese. We focus on the most basic PEAR variant,²⁰ trained without distilled MQM data and evaluated in the single-pass inference configuration. We compare two checkpoints that differ only in the antisymmetry regularization strength: $\lambda_{\text{flip}} = 0$ and $\lambda_{\text{flip}} = 0.1$.

Antisymmetry deviation. Let $\Delta(S_A, S_B)$ denote the system-level PEAR score for the ordered system pair (S_A, S_B) , computed as in Section 3.2 by averaging segment-level scores. Antisymmetry requires $\Delta(S_A, S_B) = -\Delta(S_B, S_A)$. We measure

deviation from this property using the absolute *antisymmetry residual*:

$$\epsilon_{\text{AS}}(S_A, S_B) = |\Delta(S_A, S_B) + \Delta(S_B, S_A)|.$$

We compute this over all MT system pairs.

Transitivity deviation. A stronger additive consistency property over triples is:

$$\Delta(S_A, S_C) = \Delta(S_A, S_B) + \Delta(S_B, S_C).$$

We quantify deviation using the absolute *transitivity residual*:

$$\epsilon_{\text{TR}}(S_A, S_B, S_C) = \left| \Delta(S_A, S_C) - (\Delta(S_A, S_B) + \Delta(S_B, S_C)) \right|.$$

We compute this over all triples of distinct MT systems (S_A, S_B, S_C) .

Relative residual reporting. We report residuals relative to the typical magnitude of system-level PEAR scores. This is preferable to raw residual values because adding or removing the antisymmetry regularization training objective could change the output score distribution, making absolute residuals less comparable across training setups. Concretely, we normalize by the mean absolute system-level score:

$$\mu_{\Delta} = \mathbb{E}_{(S_A, S_B)} [|\Delta(S_A, S_B)|],$$

and report the relative deviations:

$$\rho_{\text{AS}} = \frac{\mathbb{E}_{(S_A, S_B)} [\epsilon_{\text{AS}}(S_A, S_B)]}{\mu_{\Delta}},$$

$$\rho_{\text{TR}} = \frac{\mathbb{E}_{(S_A, S_B, S_C)} [\epsilon_{\text{TR}}(S_A, S_B, S_C)]}{\mu_{\Delta}}.$$

Results. Table 5 shows that enabling the antisymmetry auxiliary training objective substantially reduces antisymmetry deviations, yielding residuals that are only a small fraction of the typical PEAR score magnitude. While transitivity deviations remain larger overall (as expected for an unconstrained real-valued scorer), they are also consistently reduced with $\lambda_{\text{flip}} = 0.1$. This indicates that encouraging sign inversion under candidate reversal improves not only antisymmetry, but also promotes a more consistent behavior across MT system triples.

²⁰InfoXML Large checkpoint.

Model (InfoXLM Large)	$\rho_{AS} \downarrow$	$\rho_{TR} \downarrow$
PEAR ($\lambda_{flip} = 0$)	0.196	0.561
PEAR ($\lambda_{flip} = 0.1$)	0.014	0.189

Table 5: Pairwise consistency deviations on WMT24 MQM (En-De, En-Es, Ja-Zh), computed over all system pairs (antisymmetry) and system triples (transitivity). Residuals are reported as proportions relative to the mean absolute system-level PEAR score, i.e., lower is better.

Loss	SPA \uparrow	$acc_{eq}^* \uparrow$	Avg Corr \uparrow
MSE	80.6	57.6	69.1
Huber	80.9	57.9	69.4

Table 6: WMT24 meta-evaluation for PEAR (InfoXLM Large, 560M) trained with MSE vs. Huber regression on pairwise human-difference supervision. SPA is reported at the system level (Thompson et al., 2024); acc_{eq}^* is segment-level pairwise accuracy with tie calibration (Deutsch et al., 2023).

B Huber Loss vs. MSE for Pairwise Difference Regression

PEAR is trained to regress human quality differences Δ_{ab}^* (Section 3.4). In all the experiments, we use the Huber loss, which is less sensitive to occasional large residuals than MSE (Huber, 1964).

To quantify the impact of this choice, we run an ablation for PEAR²¹ in the same training setup as Section 5.1, changing only the regression loss from Huber to MSE while keeping all other hyperparameters fixed. We evaluate on WMT24 MQM with the official meta-evaluation toolkit, reporting Soft Pairwise Accuracy (SPA) at the system level (Thompson et al., 2024), pairwise accuracy with tie calibration (acc_{eq}^*) at the segment level (Deutsch et al., 2023), and Avg Corr.

Table 6 shows that Huber yields small but consistent gains over MSE across meta-evaluation statistics. We hypothesize that down-weighting large residuals stabilizes learning under heavy-tailed pairwise targets, which can improve calibration for close comparisons among strong MT systems. Further analysis of when and why these gains arise is left to future work.

C Training and Evaluation Setup

C.1 Data

PEAR training set is composed of 7M translation pairs across 51 language pairs for the first train-

²¹We use the InfoXLM Large variant.

ing stage (DA and DA+SQM), and by 6M across 10 language pairs for the second training stage (MQM). We do not apply any transformation to the scores provided by human annotators. Following Proietti et al. (2025b), when multiple MQM scores are available for a translation, we do not aggregate them and instead treat each annotation as a separate training instance. For the KD data, we run the GEMBA-MQM V2 approach on WMT data annotated only with DA and DA+SQM, and on internal MT output data, featuring a total of 2M additional translation pairs, covering six additional language pairs in addition to those present in gold MQM data.

C.2 Hyperparameters

For the experiments outlined in Section 5.1, we train all the models with AdamW (Loshchilov and Hutter, 2019), using a learning rate equal to $2e-5$ until convergence on the development set (WMT23 MQM data) for a maximum of five epochs. For the experiments described in Section 5.3, we keep the same hyperparameters, extend the training set with WMT23, and train for one epoch. The λ_{flip} hyperparameter has been set to 0.1 after preliminary hyperparameter optimization experiments where we did not observe large differences in the $[0.1, 0.5]$ range, with a step size equal to 0.1. The δ hyperparameter has been set to 4.5 after preliminary hyperparameter optimization experiments, where we did not observe large differences in the $[2.0, 8.0]$ range, with a step size equal to 0.5. Additionally, in our initial experiments (Section 5.1), including the antisymmetry auxiliary training objective yielded a higher Avg Corr on the overall WMT24 evaluation. Without it, PEAR’s Avg Corr decreased from 69.4²² to 68.9. The parameter α_{raw} presented in Section 3 is kept frozen for the first training stage on DA and DA+SQM data, and it is learned in the second training stage on MQM, with an initial value of 1.0. This is intended to facilitate adaptation to the training distribution in the second stage, which is derived from MQM annotations rather than DA or DA+SQM.

C.3 Baselines

Since PEAR targets improvements in QE for comparing candidate translations, we primarily benchmark against top-performing QE metric families

²²This refers to the performance obtained by PEAR in Table 1.

submitted as primary entries to the WMT24 Metrics Shared Task (Freitag et al., 2024). For additional context, we also report a set of strong WMT24 reference-based baselines.

C.3.1 QE metrics.

We compare with the following metric families:

- **COMET family.** We include CometKiwi (Rei et al., 2022b), its large-scale variant CometKiwi-XXL (Rei et al., 2023), and XCOMET-QE (Guerreiro et al., 2024).
- **MetricX family.** We include the QE variants of MetricX-24 Hybrid (Juraska et al., 2024), namely MetricX-24-Hybrid-QE-Large, -XL, and -XXL.
- **GEMBA family.** We include GEMBA-ESA, a WMT24 primary submission that prompts GPT-4 to follow the ESA procedure by extracting error spans and then producing a final 0–100 score (Kocmi and Federmann, 2023a,b; Freitag et al., 2024).

We do not include COMET-poly in the WMT24 comparison: the publicly released checkpoints are trained on WMT data up to 2024 (Züfle et al., 2025), which overlaps with the benchmark.

C.3.2 Reference-Based Metrics

To provide an additional performance reference from strong reference-based metrics, we include COMET-22 (Rei et al., 2020; Freitag et al., 2024), BLEURT-20 (Sellam et al., 2020b), and MetricX-24-Hybrid-Large (Juraska et al., 2024). These models score single candidate translations in a reference-based setting and serve as competitive anchors alongside the comparisons focused on QE.

D Correlation Matrices

This appendix section provides the correlation matrices for the WMT24 MQM language pairs En-De (Figure 3), En-Es (Figure 4), and Ja-Zh (Figure 5), which are discussed in Section 5.5.




Figure 3: En-De Pearson correlation matrix between segment-level pairwise difference scores derived from WMT24 metric scores. Single-candidate metrics are converted into pairwise differences by subtraction, while PEAR produces pairwise scores directly. The two PEAR models are instantiated with InfoXLM Large and used in the bidirectional (*both*) inference mode.




Figure 4: En-Es Pearson correlation matrix between segment-level pairwise difference scores derived from WMT24 metric scores. Single-candidate metrics are converted into pairwise differences by subtraction, while PEAR produces pairwise scores directly. The two PEAR models are instantiated with InfoXLM Large and used in the bidirectional (*both*) inference mode.




Figure 5: Ja-Zh Pearson correlation matrix between segment-level pairwise difference scores derived from WMT24 metric scores. Single-candidate metrics are converted into pairwise differences by subtraction, while PEAR produces pairwise scores directly. The two PEAR models are instantiated with InfoXLM Large and used in the bidirectional (*both*) inference mode.