

MATCHA 🍵: Matching Text via Contrastive Semantic Alignment

Siran Li Ece Sena Etoglu Carsten Eickhoff Seyed Ali Bahrainian

University of Tübingen

siran.li@uni-tuebingen.de, ece.etoglu@guest.uni-tuebingen.de

{carsten.eickhoff, seyed-ali.bahreinian}@uni-tuebingen.de

Abstract

Reliable evaluation is essential for understanding large language model (LLM) performance, yet today’s go-to metrics, namely token-overlap scores (e.g., ROUGE) and embedding-based measures (e.g., BERTScore), often misjudge semantic similarity of documents. Our study shows that both token-overlap metrics and embedding-based metrics routinely assign nearly identical scores to texts that directly contradict each other, thereby potentially masking fundamental errors. We introduce MATCHA, an automatic metric that jointly rewards semantic agreement with a reference and penalizes contradictions. MATCHA employs a dual-view perspective that measures (i) proximity to the gold text and (ii) distance from an adversarially generated counterfactual contradiction. In eight public benchmarks, MATCHA outperforms popular metrics, compared with human annotations on question-answering, image caption generation, natural language inference, summarization, and semantic textual similarity tasks. On the TruthfulQA dataset (i.e., a dataset without a training set, where no embedding-based metrics could locally train on), this improvement in terms of matching texts with a reference reaches 18.38% over ROUGE-L and 20.82% over BERTScore. Both quantitative comparison and qualitative human assessments confirm the efficacy and validity of MATCHA and uncover fundamental weaknesses in pre-existing metrics. Compared with 23 embedding models, including top state-of-the-art ones, used as a metric similar to BERTScore, MATCHA remains the most accurate in distinguishing correct from incorrect statements solely based on a reference. Our code and metric are publicly available ¹.

1 Introduction

In text generation research and applications, evaluation constitutes the main bottleneck in assess-

ing genuine progress, mainly because consistently and objectively measuring the quality, relevance, or correctness of generated outputs remains a highly challenging task, particularly in knowledge-intensive and multimodal settings (Mehri and Eskenazi, 2020; Pang et al., 2020; Rei et al., 2020; Li et al., 2025). Various applications of text generation attract thousands of publications each year at prestigious conferences such as ACL, EMNLP, NeurIPS, and others. Despite rigorous peer review processes, reliably evaluating the validity of the evaluation sections in these research papers remains a challenge for both reviewers and post-acceptance readers (Sai et al., 2022). Although substantial effort is directed toward ensuring the robustness of evaluation methodologies, an essential, yet overlooked, concern is whether popular evaluation metrics themselves are adequate (Reiter, 2018), or whether shifts in the field have rendered historically significant metrics increasingly inadequate (Dai et al., 2024).

Evaluation metrics for text generation applications typically fall into two main categories: (1) Lexical overlap matching methods, which evaluate the quality of generated texts by counting overlapping textual units, such as n-grams, word sequences, or word pairs, between the generated candidate outputs and human-annotated references. Notable examples of this category include widely-used metrics such as ROUGE-N, ROUGE-L (Lin, 2004), and Meteor (Banerjee and Lavie, 2005). (2) Model-based evaluation methods, which measure semantic similarity by mapping texts onto continuous embedding spaces. Prominent metrics in this category include BERTScore (Zhang et al., 2020), BLEURT (Sellam et al., 2020), SimCSE (Gao et al., 2021), and MAUVE (Pillutla et al., 2021). Although metrics such as ROUGE, originating in the early 2000s, initially served evaluations in extractive summarization (i.e., extracting exact sentences from a document that together form a sum-

¹<https://github.com/Siran-Li/MATCHA>

Example	Metric	Correct	Incorrect
Reference: Growth is not affected by caffeine consumption.	RL-F1	28.57	15.38
	METEOR	36.51	7.25
	EmbSim	81.00	77.19
Correct: Drinking coffee does not affect your growth.	BERTScore	82.00	77.20
	BLEURT	-34.88	-44.74
Incorrect: Drinking coffee will stunt your growth.	SimCSE	88.88	75.21
	Human Rating	A (best)	E (worst)

Table 1: Comparison of evaluation metrics and human judgment for correct vs. incorrect pairs. RL-F1 denotes ROUGE-L F1. Human Rating refers to judgments from three annotators on a five-point A–E scale.

mary) when text generation capabilities were limited within the NLP literature, they have persisted and adapted to contemporary systems (Karpinska et al., 2022). Conversely, newer metrics such as BERTScore (Zhang et al., 2020) have emerged alongside the rise of Transformer-based models, rapidly gaining traction and becoming virtual standards for assessing text generation across various NLP tasks. However, in recent years, new datasets with contrastive samples have opened new pathways both for building new robust evaluation metrics, as well as assessing existing text generation evaluation metrics that compare a model-generated response against a human-written one (Guan et al., 2021). For example, we will show that an evaluation metric such as BERTScore assigns very high similarity scores not only to two contradictory documents but also to two documents that are semantically unrelated². Table 1 shows the scores assigned by various popular text generation metrics to two contradicting sentences, one being correct and the other incorrect. We observe that model-based metrics such as BERTScore, Embedding Cosine Similarity (EmbSim) computed using Sentence-Transformer (Reimers and Gurevych, 2019), and SimCSE assign very similar scores to both correct and incorrect sentences, and as a standalone metric, they are not sufficiently discriminative and are also hardly human-interpretable. This trend continues and worsens as we generalize to large test datasets such as MultiNLI (Williams et al., 2018) and TruthfulQA (Lin et al., 2022), as demonstrated in Table 2. It is noteworthy that a distribution-based metric such as MAUVE cannot reliably evaluate

²For the semantically unrelated experiment, please see results on COCO-Caption in Table 3

Metric	MultiNLI		TruthfulQA	
	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$
RL-F1	(38.07, 30.38)	7.69	(40.72, 35.78)	4.94
METEOR	(35.03, 26.03)	9.00	(42.08, 33.60)	8.48
EmbSim	(71.59, 56.79)	7.40	(70.77, 63.11)	3.83
BERTScore	(84.06, 80.62)	3.44	(83.80, 81.29)	2.51
BLEURT	(-39.16, -75.55)	18.20	(-29.73, -63.46)	16.87
SimCSE	(76.29, 49.55)	13.37	(68.98, 55.77)	6.61
MAUVE	(56.01, 49.20)	6.81	(96.81, 86.60)	10.21

Table 2: Average similarity scores for correct and incorrect pairs on two datasets. $N\Delta$ denotes the correct–incorrect gap after rescaling to $[0, 1]$.

single examples, as the authors explained³ and instead requires larger sample sizes such as entire test sets to evaluate, which further adds to its limitation of not being able to score-wise distinguish between correct and incorrect, as demonstrated in Table 1.

Metrics such as BERTScore and MAUVE do not explicitly guard against Type I errors, i.e., false positives where two texts are judged similar even though they diverge significantly. Because these scores rely on dense embedding-space distances (where most well-formed sentences cluster fairly closely together) (Gao et al., 2019), they routinely assign high similarity to almost any pair of coherent documents. This tendency masks important semantic mismatches, so relying on these metrics alone can yield an overly optimistic view of a system’s faithfulness or factual accuracy.

In this paper, we study existing text generation evaluation metrics in terms of reliability and robustness and demonstrate their inherent limitations. Subsequently, we introduce a new evaluation metric, MATCHA (Matching Text via Contrastive Semantic Alignment), which does not suffer from these shortcomings.

Our key contributions are: (1) We present an extensive analysis of existing text generation evaluation metrics and compare them robustly. (2) We present MATCHA, a new text semantic matching metric which is the closest to human judgments as compared with pre-existing metrics. MATCHA consistently achieves the best separation of contradictory statements across diverse NLP evaluation datasets; MATCHA learns a dual-perspective of matching similar documents while detecting dissimilar/contradictory documents, providing sharper semantic boundaries; MATCHA offers fine-grained

³<https://krishnap25.github.io/mauve>

interpretability by highlighting meaningful token-level semantic differences; and MATCHA shows stronger correlation with human similarity ratings, as well as, GPT-4 compared with previous metrics, confirming its reliability for detailed semantic evaluation. (3) We compare MATCHA with 23 top embedding models and show that it outperforms all of these models in distinguishing correct and incorrect statements using a reference.

2 Related Work: Evaluation Metrics for Natural Language Generation

Automatic evaluation metrics are essential for assessing natural language generation (NLG) quality. Traditional approaches range from n-gram-based methods like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) to neural models such as BERTScore (Zhang et al., 2020) and BLEURT (Selam et al., 2020). While these metrics offer valuable approximations, they often fall short in capturing semantic nuances, contextual relevance, and factual correctness. This section reviews key evaluation paradigms, highlights their shortcomings, and outlines the motivation for new alternatives.

Lexical overlap metrics provide a fast baseline but lack semantic awareness. BLEU and ROUGE compare n-gram overlap between generated and reference texts. Although efficient and widely adopted, they poorly reflect human preferences in tasks involving paraphrasing (Bahrainian et al., 2024) or abstraction (Liu et al., 2023; Phy et al., 2020). METEOR (Banerjee and Lavie, 2005) improves this by incorporating synonymy and stemming, but its reliance on static lexical resources limits its scalability across domains (Dziri et al., 2019).

Semantic similarity metrics improve alignment with human judgments but struggle with nuance and reliability. Embedding-based methods leverage semantic similarity through pretrained models. EmbSim (Reimers and Gurevych, 2019), BERTScore (Zhang et al., 2020), BLEURT (Selam et al., 2020), and SimCSE (Gao et al., 2021) achieve stronger alignment with human judgments. MAUVE (Pillutla et al., 2021) further assesses divergence between machine and human text distributions. However, these methods are often insensitive to factual errors, truncations, or adversarial inputs (He et al., 2023; Deutsch et al., 2022; Goyal and Durrett, 2020). MAUVE also struggles with fine-grained evaluation, requiring large text samples (Pimentel et al., 2023). Thus, we need task-

aware, robustness-oriented evaluators that penalize factual errors and adversarial perturbations.

Current evaluation metrics lack semantic precision, interpretability, and robustness. Despite steady progress, current evaluation metrics for NLG often fall short in capturing critical aspects such as semantic contradiction, factual inaccuracy, and over-reliance on fluency, where outputs appear fluent and well-structured despite containing incorrect or unsupported information (Howcroft et al., 2020; Schmidtova et al., 2024; Golovanevsky et al., 2025). Studies have shown that many metrics reward surface-level similarity while failing to penalize outputs that are semantically incorrect or inconsistent with a reference (He et al., 2023; Gehrmann et al., 2021; Fabbri et al., 2021). This is particularly problematic in settings where detecting subtle false positives or contradictory statements is essential. Moreover, model-based and embedding-based methods frequently offer limited interpretability, making it difficult to diagnose errors or understand model decisions (Jiang et al., 2024). As shown in Table 2, these issues persist across domains and tasks, calling for metrics that are both semantically precise and diagnostically useful. To address these challenges, we introduce MATCHA, a semantic matching metric that combines contrastive learning with token-level alignment to deliver accurate, interpretable, and robust NLG evaluation.

3 MATCHA

MATCHA is a contrastive matching metric that compares two documents by encoding them into contextualized representations, projecting them onto a contrastive-learned matching space, and computing their similarity. Inspired by the contrastive learning paradigm (Chen et al., 2020; Liu et al., 2022; Mi et al., 2022), MATCHA encourages semantically similar inputs to have closer representations, while pushing dissimilar inputs apart. The overall architecture of MATCHA is illustrated in Figure 1. Our novel method consists of three main processing stages: Token Embedding, Representation Alignment, and Similarity Computation, offering both document-level matching and fine-grained token attribution.

3.1 Token Embedding

Given an input document of tokens $\mathbf{x} = (x_1, x_2, \dots, x_L)$, we initialize the embedding layer using the pre-trained GPT-2 word embedding




Figure 1: Overview of the MATCHA architecture. Each input Correct Candidate (C), Reference (R), and Incorrect Candidate (I) is first encoded into embeddings \mathbf{E} via an embedding layer initialized from the GPT-2 decoder. A lightweight feed-forward projector with a conversion matrix \mathbf{W} maps \mathbf{E} into a shared semantic space, yielding context representations \mathbf{h} , which are compared via cosine similarity.

layer (Radford et al., 2019). Each token x_i is mapped to a D -dimensional vector:

$$\mathbf{E} = \text{EmbeddingLayer}(\mathbf{x}) \in \mathbb{R}^{L \times D} \quad (1)$$

where L is the sequence length and D is the hidden dimensionality. Only the embedding layer is transferred from the GPT-2 decoder; all subsequent layers are trained from scratch.

3.2 Representation Projection

To adapt the embeddings to our task-specific representation space, we project each token embedding into N_c context vectors via a lightweight feed-forward projection:

$$\mathbf{S} = \text{FeedForward}(\mathbf{E}) \in \mathbb{R}^{N_c \times L \times D} \quad (2)$$

Next, the context vectors are projected into a common representation space using a learned linear conversion matrix, $\mathbf{W} \in \mathbb{R}^{D \times D}$, defined as:

$$\mathbf{S}' = \mathbf{S} \times \mathbf{W} \quad (3)$$

where the multiplication is applied across the last dimension.

Finally, a document-level representation is obtained by mean pooling across both token and context dimensions:

$$\mathbf{h} = \frac{1}{N_c L} \sum_{i=1}^{N_c} \sum_{j=1}^L \mathbf{S}'_{i,j} \quad (4)$$

This yields a fixed-size representation $\mathbf{h} \in \mathbb{R}^{1 \times D}$ for each input document.

3.3 Similarity Computation, Training, and Inference

We describe how MATCHA computes similarity, learns contrastive data, and scores candidates.

Similarity Score. Given two documents with embeddings \mathbf{h}_1 and \mathbf{h}_2 , their semantic similarity is measured using cosine similarity:

$$\text{sim}(\mathbf{h}_1, \mathbf{h}_2) = \frac{\mathbf{h}_1 \cdot \mathbf{h}_2}{\|\mathbf{h}_1\| \|\mathbf{h}_2\|} \quad (5)$$

where \cdot denotes the dot product and $\|\cdot\|$ is the Euclidean norm.

Training Objective. MATCHA is trained to maximize the distance between correct and incorrect candidates in the embedding space. Given a reference embedding \mathbf{h}_R , a positive candidate \mathbf{h}_C , and a negative candidate \mathbf{h}_I , their cosine similarities are:

$$\text{sim}_C = \text{sim}(\mathbf{h}_R, \mathbf{h}_C), \quad \text{sim}_I = \text{sim}(\mathbf{h}_R, \mathbf{h}_I) \quad (6)$$

A margin-based contrastive loss enforces sim_C to be higher than sim_I :

$$\mathcal{L} = \mathbb{E} [\max(0, m + \text{sim}_I - \text{sim}_C)] \quad (7)$$

where m is a fixed margin hyperparameter (default $m = 1$). This encourages the model to pull reference-positive pairs closer together while pushing reference-negative pairs farther apart.

Inference. At inference time, all model parameters are fixed. For a reference R and a candidate T , MATCHA encodes both into document embeddings and computes their similarity score:

$$\text{Score}(R, T) = \text{sim}(\mathbf{h}_R, \mathbf{h}_T) \quad (8)$$

4 Evaluation

In this section, we evaluate MATCHA across a broad range of tasks and demonstrate that it consistently outperforms existing metrics. Additional details on training, preprocessing, implementation details, and additional ablation analysis on data composition and optimization strategy are provided in Appendices A.1, A.2, A.3, and D.

Metric	SNLI		MultiNLI		TruthfulQA		Climate-Fever		COCO-Caption		NEWTS	
	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$	(Corr, Incorr)	$N\Delta$
R1-F1	(41.34, 28.58)	12.76	(45.32, 34.67)	10.65	(45.12, 38.87)	6.25	(26.78, 23.05)	3.73	(27.10, 12.49)	14.61	(19.18, 8.70)	10.48
R2-F1	(19.07, 10.26)	8.81	(22.74, 14.96)	7.78	(24.40, 21.53)	2.87	(7.55, 5.67)	1.88	(12.13, 1.60)	10.53	(2.91, 0.39)	2.52
RL-F1	(38.24, 26.44)	11.80	(38.07, 30.38)	7.69	(40.72, 35.78)	4.94	(19.62, 18.12)	1.50	(23.27, 11.32)	11.95	(12.74, 6.80)	5.94
METEOR	(30.08, 20.78)	9.30	(35.03, 26.03)	9.00	(42.08, 33.60)	8.48	(19.01, 15.33)	3.68	(13.19, 5.51)	7.68	(10.41, 3.74)	6.67
EmbSim	(65.70, 33.36)	16.17	(71.59, 56.79)	7.40	(70.77, 63.11)	3.83	(60.56, 55.74)	2.41	(64.33, 9.25)	27.54	(30.89, 5.69)	12.61
BERTScore	(83.05, 79.34)	3.71	(84.06, 80.62)	3.44	(83.80, 81.29)	2.51	(78.91, 77.74)	1.17	(83.01, 73.30)	9.71	(73.44, 66.14)	7.30
BLEURT	(-64.06, -99.63)	17.79	(-39.16, -75.55)	18.20	(-29.73, -63.46)	16.87	(-41.65, -53.76)	6.06	(-76.28, -124.05)	23.89	(-108.56, -135.69)	13.57
SimCSE	(69.47, 33.43)	18.03	(76.29, 49.55)	13.37	(68.98, 55.77)	6.61	(65.63, 51.82)	6.91	(70.95, 10.20)	30.38	(39.45, 11.53)	13.96
MAUVE	(41.50, 47.67)	-6.17	(56.01, 49.20)	6.81	(96.81, 86.60)	10.21	(70.95, 51.99)	18.96	(6.87, 6.16)	0.71	(1.64, 0.75)	0.89
MATCHA	(71.14, 1.24)	34.95	(64.77, -9.76)	37.27	(62.50, 15.85)	23.33	(52.92, 12.57)	20.18	(81.11, 9.59)	35.75	(44.64, 16.53)	14.06

Table 3: Average similarity scores for correct and incorrect pairs across six benchmarks. $N\Delta$ denotes the normalized correct-incorrect gap, with EmbSim, BLEURT, SimCSE, and MATCHA rescaled to $[0, 1]$ via $(\text{score} + 1)/2$. Bold indicates statistical significance over the next highest score, based on a paired t-test with a confidence of 95%.




Figure 2: Comparison of different metrics' performance on six datasets. The y-axis shows the percentage of samples satisfying $(\text{Sim}_C - \text{Sim}_I) > \text{Threshold}$, with the middle point as the decision boundary. Metric scores are linearly transformed to $[0, 1]$.

Metric	SNLI	MultiNLI	TruthfulQA	Climate	COCO	NEWTS
R1-F1	57.69	58.15	55.04	37.11	33.67	33.71
R2-F1	42.35	43.15	44.29	34.07	33.33	33.33
RL-F1	54.67	50.90	51.53	34.48	33.33	33.52
METEOR	46.68	50.51	55.20	36.93	33.33	33.43
EmbSim	37.96	33.33	33.95	33.33	52.79	60.89
BERTScore	33.33	33.33	33.33	33.33	33.33	33.33
BLEURT	44.72	54.71	56.03	40.23	33.67	33.33
SimCSE	39.20	33.71	34.65	33.33	55.04	49.76
MATCHA	72.59	74.51	61.08	64.08	68.41	57.99

Table 4: Macro-F1 classification performance under a fixed midpoint threshold across six benchmarks.

4.1 Experimental Setup

Baselines and Datasets. We compare MATCHA against nine widely used metrics: ROUGE-1 (R1-F1), ROUGE-2 (R2-F1), ROUGE-L (RL-F1), METEOR, EmbSim, BERTScore, BLEURT, SimCSE, and MAUVE. To demonstrate the broad applicability of MATCHA, we evaluate it on eight datasets, including SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018), TruthfulQA (Lin et al., 2022), Climate-Fever (Diggelmann et al., 2020), COCO-Caption (Lin et al., 2015), NEWTS (Bahrainian et al., 2022), MedNLI (Romanov and Shivade, 2018) and STS-B

(English) dataset (Cer et al., 2017), spanning natural language inference (NLI), question-answering, factuality verification, image captioning, summarization, and semantic textual similarity (STS).

Evaluation Methods. To assess the discriminative power of different metrics, we measure the average correct scores Sim_C and incorrect scores Sim_I and their gaps $N\Delta$, defined as the average of $(\text{Sim}_C - \text{Sim}_I)$ rescaled to $[0, 1]$, the percentage of cases where $(\text{Sim}_C - \text{Sim}_I) > \text{Threshold}$, and the Wasserstein distance (Villani, 2009). Further, we evaluate alignment with human judgment using Rank@1 ($R@1$), defined as the percentage of cases where a metric's score is closest to the human score, average Discounted Cumulative Gain (DCG) (Järvelin and Kekäläinen, 2002), and concordance correlation coefficient (CCC) (Lawrence and Lin, 1989). Unlike Pearson or Spearman, CCC captures not only rank correlation but also calibration, ensuring that metric scores align with human ratings in both ordering and scale. We also analyze token-level interpretability using Integrated Gradients (Sundararajan et al., 2017) to compute token-level attributions in Appendix C.

4.2 Experimental Results

To assess MATCHA’s performance, we conduct four sets of experiments: (i) we compare it with existing metrics on contrastive datasets to distinguish texts with contrastive/irrelevant meanings (Section 4.2.1); (ii) we evaluate MATCHA against 23 state-of-the-art embedding models, examining their semantic separation ability (Section 4.2.2); (iii) we evaluate out-of-domain generalization (Section 4.2.3); and (iv) via human evaluation (Section 4.2.4), we show that MATCHA aligns with human judgments more than the baselines.

4.2.1 Metric Semantic Differentiation Power

We test the semantic differentiation power of MATCHA and baseline metrics on the following three aspects: (1) average similarity gap and macro-F1 classification performance (Table 3 and Table 4), (2) proportion of samples with clear separation (Figure 2), and (3) Wasserstein distribution distance performance (Table 5).

Average similarity gap. We measure the normalized similarity gap, which reflects how well a metric separates correct from incorrect samples generally (Table 3). MATCHA achieves the highest gap on six datasets, substantially outperforming baseline metrics, with a gap of 34.95 on SNLI and 37.27 on MultiNLI, significantly ahead of the next best, SimCSE and BLEURT (18.03 and 18.20, respectively). On TruthfulQA, MATCHA outperforms BLEURT by +6.46%. It further outperforms MAUVE by +1.22% on Climate-Fever, +5.37% on COCO-Caption, and +0.10% on NEWTS. These results demonstrate MATCHA’s strong and consistent semantic differentiation ability, even on challenging out-of-domain tasks.

We additionally evaluate binary discrimination performance using macro-F1 under a fixed midpoint threshold (Table 4). MATCHA delivers the strongest overall classification performance, achieving the highest macro-F1 on five of the six benchmarks. Compared with the best baseline on each dataset, MATCHA improves performance by +14.90 on SNLI, +16.36 on MultiNLI, +5.05 on TruthfulQA, +23.85 on Climate-Fever, and +13.37 on COCO-Caption. On NEWTS, MATCHA remains highly competitive at 57.99. These findings show that MATCHA’s larger semantic score margins translate into substantially more reliable decision boundaries for distinguishing correct from incorrect text pairs, while preserving robustness under challenging unseen-domain transfer settings.

Metric	SNLI	MultiNLI	TruthfulQA	Climate	COCO	NEWTS
R1-F1	12.77	10.65	6.26	3.87	14.61	10.48
R2-F1	8.81	7.77	2.98	1.88	10.53	2.52
RL-F1	11.80	7.69	4.95	1.66	11.95	5.94
METEOR	9.30	8.99	8.48	3.68	7.68	6.67
EmbSim	16.17	7.40	3.85	2.42	27.54	12.60
BERTScore	3.72	3.43	2.51	1.17	9.71	7.30
BLEURT	17.79	18.20	16.91	6.06	23.89	13.56
SimCSE	18.02	13.37	6.61	6.91	30.38	13.96
MATCHA	34.95	37.27	23.33	20.17	35.76	14.05

Table 5: Wasserstein distribution distance between correct and incorrect scores. Climate and COCO denote Climate-Fever and COCO-Caption, respectively.

Margin between correct and incorrect scores.

As shown in Figure 2, MATCHA consistently achieves larger margins across thresholds. In contrast, ROUGE, METEOR, and BERTScore drop off quickly, showing limited separation. EmbSim and SimCSE perform better than n-gram metrics but still lag behind MATCHA, especially on MultiNLI and TruthfulQA, where samples often share the same topic with high wording overlap, making contrastive distinctions more difficult than in other datasets.

Wasserstein distribution distance performance. We evaluate separation between correct and incorrect scores using Wasserstein distribution distance, a non-parametric metric that captures how far apart two distributions are without assuming any specific shape. As shown in Table 5, MATCHA consistently records the highest distances, with especially large margins on SNLI (34.95) and MultiNLI (37.27), nearly doubling the strongest baselines. It also surpasses the strongest baselines by +6.42% on TruthfulQA, +13.26% on Climate-Fever, and +5.38% on COCO-Caption, while maintaining competitive performance on NEWTS, confirming its ability to avoid score clustering and provide more reliable separation between correct and incorrect samples.

Traditional metrics such as ROUGE, METEOR, and BERTScore consistently show low similarity gaps, weak separation margins, and poor macro-F1 discrimination performance. BLEURT assigns negative scores on average to both correct and incorrect cases, although it remains competitive on some datasets. EmbSim and SimCSE perform better than the previous metrics but still fall behind MATCHA, particularly in classification robustness and cross-domain generalization. MAUVE, which measures distributional similarity rather than direct semantic

alignment, also fails to provide meaningful separation in SNLI and COCO-Caption. In contrast, MATCHA consistently delivers larger distinction gaps, stronger separation margins, higher Wasserstein distribution distances, and the best macro-F1 scores across datasets, establishing itself as a more reliable measure of semantic distinction.

4.2.2 Embedding-Based Semantic Separation

We evaluate MATCHA in the embedding space by comparing it with 23 embedding models, testing their ability to distinguish between correct and incorrect pairs using cosine similarity. The evaluated models form one of **the most comprehensive comparisons to date**, including top STS performers from the MTEB benchmark (Muennighoff et al., 2023)⁴, as well as baseline transformers and traditional word embeddings.

We include the following models: Bilingual-Embedding-Large (Lajavaness, 2023), Jina-Embeddings-V3 (Sturua et al., 2024), SFR-Embedding-Mistral (Salesforce, 2023), Speed-Embedding-7B-Instruct (Haon-Chen, 2023), Multilingual-E5-Large-Instruct (Wang et al., 2024a), Linq-Embed-Mistral (Linq AI Research, 2023), Jasper and Stella (Zhang et al., 2025), GTE-Large (The NLP Group, 2023), and the instruction-tuned LLM E5-Mistral-7B-Instruct (Intfloat, 2023). Additionally, we evaluate widely used pre-trained sentence encoders such as MiniLM (Wang et al., 2021), MPNet (Song et al., 2020), Sentence-T5 (Ni et al., 2022), E5-Large-v2 (Wang et al., 2024b), BGE-Large (BAAI, 2023), and BERT (base) (Devlin et al., 2019), along with DistilBERT-NLI (Sanh et al., 2020) included as a relevant semantic baseline, due to its fine-tuning on NLI datasets. We also examine embeddings derived from LLMs such as LLaMA-2 (Touvron et al., 2023), LLaMA-3.1-8B (Meta AI, 2024), and Mistral-7B (Jiang et al., 2023). Lastly, our study includes SNPMI (Bahrainian et al., 2021), a co-occurrence-based embedding constructed from normalized pointwise mutual information (NPMI), alongside traditional word embedding baselines such as GloVe (Pennington et al., 2014) and Word2Vec (Mikolov et al., 2013).

Average similarity gap among 23 LM embeddings. Across SNLI, MultiNLI, and TruthfulQA, MATCHA produces the largest contrast between correct and incorrect scores (Figure 3).

⁴<https://huggingface.co/spaces/mteb/leaderboard>, see the English leaderboard

Metric	(Corr, Incorr)	$N\Delta$	macro-F1	Wasserstein
R1-F1	(14.22, 11.65)	2.57	37.45	2.58
R2-F1	(4.66, 3.12)	1.54	34.14	1.54
RL-F1	(13.31, 11.08)	2.23	36.56	2.25
METEOR	(8.36, 6.62)	1.74	35.14	1.75
EmbSim	(46.14, 40.47)	2.84	33.43	2.85
BERTScore	(74.96, 74.17)	0.79	33.33	1.00
BLEURT	(-122.43, -132.13)	4.85	34.94	4.85
SimCSE	(50.43, 37.00)	6.72	35.35	6.72
MAUVE	(1.29, 1.40)	-0.11	–	–
MATCHA	(50.79, 14.24)	18.28	59.49	18.27

Table 6: Performance on MedNLI, reporting average similarity scores for correct and incorrect pairs, $N\Delta$, macro-F1, and Wasserstein distance.

While most embeddings assign uniformly high similarities, yielding small gaps, MATCHA clearly separates the two and is the only model to give negative average similarity to incorrect samples (−9.76 on MultiNLI) while preserving high similarity for correct ones. This strong calibration enables MATCHA to overcome the common “over-similarity” problem in large pretrained embeddings (Ethayarajh, 2019), making it a more human-aligned and trustworthy metric than baselines.

Based on Figures 3 and the semantic separation curves in Appendix B, MATCHA achieves uniquely large similarity gaps and robust margins, addressing the shortcomings of both general-purpose embeddings and fine-tuned STS/NLI models. Due to space constraints, we present results on NLI datasets and TruthfulQA, with other datasets showing similar patterns (Figure 5). These findings demonstrate that MATCHA establishes sharper semantic boundaries and delivers more reliable discriminative performance across diverse tasks.

4.2.3 Out-of-Domain Generalization

To evaluate MATCHA’s robustness under domain shift, we further test it on MedNLI (Romanov and Shivade, 2018), a clinical-domain natural language inference benchmark derived from real medical notes. MedNLI contains specialized medical terminology and domain-specific reasoning patterns that differ substantially from the general-domain datasets used during training, making it a strong out-of-distribution evaluation setting.

As shown in Table 6, MATCHA consistently outperforms all baselines across average score separation ($N\Delta$), macro-F1, and Wasserstein distance. In particular, MATCHA achieves the highest normalized separation gap (18.28), improving over the strongest baseline SimCSE by +11.56%. The




Figure 3: Average similarity scores and gaps between correct and incorrect pairs across embedding models on SNLI, MultiNLI, and TruthfulQA.

R1-F1	R2-F1	RL-F1	MTR	EmbSim	BSc	BLT	SimCSE	MATCHA
50.73	35.78	50.56	50.95	34.37	9.16	59.76	31.81	61.08

Table 7: Concordance correlation coefficient for metrics on STS-B. MTR, BSc, and BLT denote METEOR, BERTScore, and BLEURT, respectively.

same trend holds for macro-F1 under a fixed midpoint threshold, where MATCHA reaches 59.49, outperforming the best baseline R1-F1 by +22.04%. MATCHA also records the strongest Wasserstein distance (18.27), exceeding the best baseline SimCSE by +11.55%. These findings demonstrate that MATCHA maintains robust semantic separation under substantial domain shift and generalizes effectively to domains with unseen vocabulary and specialized reasoning requirements.

While threshold-based discrimination can suffice for binary decision making, large and well-calibrated score margins are especially desirable for evaluation metrics because they improve robustness and interpretability. In practice, evaluation metrics are often used across datasets, models, and tasks without a fixed or carefully tuned threshold; small score gaps make them highly sensitive to noise, calibration shifts, and domain changes. As shown in our experiments, metrics with narrow margins frequently assign similarly high scores to both correct and contradictory candidates, masking fundamental semantic errors. In contrast, MATCHA’s larger margins provide more stable separation, stronger alignment with human judgments, and more reliable comparisons across models and settings.

4.2.4 Human Evaluation

As Liu et al. (2016) explains, NLG evaluation metrics must agree with human judgment as their ultimate test. We present a comprehensive human evaluation of MATCHA against the baselines through

Metric	SNLI		MultiNLI		TruthfulQA		Overall	
	R@1	DCG	R@1	DCG	R@1	DCG	R@1	DCG
R1-F1	5.50	50.29	8.00	47.88	2.50	46.19	5.33	48.12
R2-F1	6.00	40.94	10.50	43.53	14.00	49.00	10.17	44.49
RL-F1	1.00	48.08	3.00	45.86	3.50	47.46	2.50	47.13
METEOR	11.50	47.65	11.50	46.64	9.00	48.14	10.67	47.48
EmbSim	12.50	49.05	8.00	47.00	8.00	44.06	9.50	46.71
BERTScore	15.50	47.97	15.00	48.60	8.50	43.44	13.00	46.67
BLEURT	6.50	39.62	4.00	40.38	13.50	48.40	8.00	42.80
SimCSE	7.00	46.74	10.50	49.13	11.00	46.35	9.50	47.40
MATCHA	18.50	52.45	23.00	56.08	19.00	50.49	20.17	53.01

Table 8: Comparison of metric agreement with human judgments. Higher R@1 and DCG are better.

two complementary studies: (1) a large-scale study on the STS-B dataset, as in previous literature; and (2) a focused study designed to examine how evaluation metrics score contrastive sentence pairs relative to human judgments. Although MATCHA is designed to capture semantic alignment rather than surface similarity, we include STS-B for direct comparability with established baselines. Only for this analysis, we fine-tune MATCHA on the STS-B training split in a separate run that is fully isolated from all other experiments (Table 7).

Large-scale human evaluation on STS task. Prior evaluation metrics like SimCSE (Gao et al., 2021), commonly report results on STS, using STS-B as a standard benchmark. Table 7 shows that MATCHA attains the highest concordance correlation coefficient with STS-B human annotations, outperforming prior NLG metrics. These results indicate that MATCHA better captures semantic similarity, with document similarity scores that are most aligned with human judgments.

Human evaluation focused on contrastive correct/incorrect pairs. In order to best demonstrate MATCHA’s performance in distinguishing correct from incorrect while also emphasizing false-negative detection, we conduct a focused human




Figure 4: Concordance correlation coefficient among metrics and between metrics and human judgment.

study on contrastive pairs. This study is designed to be more focused than the large-scale one, as STS-B does not provide contrastive correct/incorrect pairs.

Study Design. We randomly sample 200 examples each from test sets of SNLI, MultiNLI, and TruthfulQA, and combine them for Overall evaluation. Human annotators rated the semantic similarity of randomly presented sentence pairs on a 1–5 scale, focusing on semantic meaning rather than surface wording. We recruited eight annotators, primarily scientists from various fields, who were not directly involved in this research. The evaluation interface is publicly available online⁵. Both human ratings and metric scores are linearly scaled to the [0,1] range. To assess alignment with human judgments, we report Rank@1 percentage (R@1), DCG, and concordance correlation coefficient.

Ranking Methods by Alignment with Human Judgments. As shown in Table 8, MATCHA consistently outperforms all baseline metrics on the overall dataset. It achieves the highest overall Rank@1 score (20.17, +7.17% over BERTScore), and highest DCG (53.01, +4.89% over R1-F1). These results suggest that MATCHA produces similarity scores more closely aligned with human perception in ranking quality.

Correlation with human ratings. Figure 4 shows that n-gram metrics (R1-F1, R2-F1, RL-F1, METEOR) correlate strongly with each other (60–90) but poorly with human judgments (13–27),

consistent with Liu et al. (2016). Embedding-based methods such as BLEURT (23.80) and SimCSE (25.23) perform better but still struggle with calibration and generalization. Notably, BERTScore yields the weakest correlation with human ratings (7.83), with scores clustered around 0.70–0.80 regardless of quality. In contrast, MATCHA achieves the highest correlation with human ratings (45.35), confirming its ability to capture nuanced, human-perceived meaning.

As a further reference, we compare MATCHA to FActScore (Min et al., 2023), a recently proposed metric that evaluates factual accuracy by assessing whether each claim is supported by an external knowledge source. While FActScore achieves a higher correlation (66.19), it is based on an industry-grade LLM (GPT-4 (Achiam et al., 2023)), with live Internet access, and explicitly optimized for factual consistency rather than semantic similarity, making it an indirect comparison. Despite being much smaller, MATCHA achieves the strongest correlation with FActScore among all metrics, validating its reliability.

5 Conclusions and Future Work

Motivated by the limited capability of existing text similarity metrics in distinguishing contradictory or sometimes even semantically unrelated texts from one another, we introduced MATCHA, a contrastive semantic similarity metric which is highly aligned with human judgment. We validate MATCHA via extensive experiments on eight datasets, comparing it with top text similarity metrics across tasks such as question answering, natural language inference, caption generation, summarization, and semantic textual similarity. In conclusion, MATCHA achieves a new state-of-the-art performance with a huge improvement. Additionally, we show that MATCHA can outperform 23 state-of-the-art embedding models in its ability to distinguish correct from incorrect using a reference. The metric is applicable to a wide range of text generation systems, including LLM steering (Braun et al., 2026), across dimensions such as topic (Ravenda et al., 2025) and sentiment (Bahrainian et al., 2014; Bahrainian and Dengel, 2015). Future work may explore extending MATCHA to multilingual and cross-modal contexts, as well as integrating it into training-time objectives to guide generation models toward more semantically faithful outputs.

⁵<https://evalmetric.streamlit.app>

Limitations

Despite MATCHA’s strong performance and alignment with human judgment across a variety of tasks, it still has some limitations that require further exploration.

(1) MATCHA has been developed and evaluated primarily on English-language datasets, and its ability to generalize to multilingual contexts remains untested. This may limit its effectiveness in diverse linguistic settings, where variations in grammar, semantics, and cultural nuance can impact similarity judgments.

(2) Although we did not present results using large language models, MATCHA was intentionally designed to remain as small as possible while still achieving strong performance. This decision reflects our consideration that future users of MATCHA may not have access to advanced GPU hardware. While this choice enhances MATCHA’s accessibility and practical utility for the community, it could also be viewed as a limitation, since models with several billion parameters or more were not evaluated.

(3) As with all supervised and learned evaluation metrics, the scoring behavior may reflect biases inherited from the training data and task-specific definitions of correctness. This is a general limitation of data-driven metrics rather than one unique to MATCHA. To reduce this risk, we design MATCHA with broad multi-domain data collection, contrastive supervision across diverse tasks, and an interleaved training strategy that discourages over-specialization to any single benchmark. We further validate its robustness across multiple benchmarks and domains, although no trained metric can fully account for all possible linguistic variations and contextual notions of correctness.

Acknowledgments

The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for their support.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

BAAI. 2023. Bge-large-en-v1.5. Available at <https://huggingface.co/BAAI/bge-large-en-v1.5>.

Seyed Ali Bahrainian and Andreas Dengel. 2015. Sentiment analysis of texts by capturing underlying sentiment patterns. In *Web Intelligence*, volume 13, pages 53–68. SAGE Publications Sage UK: London, England.

Seyed Ali Bahrainian, Jonathan Dou, and Carsten Eickhoff. 2024. *Text simplification via adaptive teaching*. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6574–6584, Bangkok, Thailand. Association for Computational Linguistics.

Seyed Ali Bahrainian, Sheridan Feucht, and Carsten Eickhoff. 2022. *NEWS: A corpus for news topic-focused summarization*. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 493–503, Dublin, Ireland. Association for Computational Linguistics.

Seyed Ali Bahrainian, Martin Jaggi, and Carsten Eickhoff. 2021. *Self-supervised neural topic modeling*. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3363–3370. Association for Computational Linguistics.

Seyed Ali Bahrainian, Marcus Liwicki, and Andreas Dengel. 2014. *Fuzzy subjective sentiment phrases: A context sensitive and self-maintaining sentiment lexicon*. In *2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, volume 1, pages 361–368.

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. *A large annotated corpus for learning natural language inference*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. Association for Computational Linguistics.

Joschka Braun, Carsten Eickhoff, and Seyed Ali Bahrainian. 2026. *Beyond multiple choice: Evaluating steering vectors for summarization*. In *Findings of the Association for Computational Linguistics: EACL 2026*, pages 3849–3884, Rabat, Morocco. Association for Computational Linguistics.

Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. *SemEval-2017 Task 1: Semantic textual similarity multilingual and crosslingual focused evaluation*. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.

- William Coster and David Kauchak. 2011. Simple English Wikipedia: a new text simplification task. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 665–669.
- Xiang Dai, Sarvnaz Karimi, and Biaoyan Fang. 2024. [A critical look at meta-evaluating summarisation evaluation metrics](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14795–14808, Miami, Florida, USA. Association for Computational Linguistics.
- DataCanary, hilfialkaff, Lili Jiang, Meg Risdal, Nikhil Dandekar, and tomtung. 2017. Quora question pairs. <https://kaggle.com/competitions/quora-question-pairs>. Kaggle.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2022. [On the limitations of reference-free evaluations of generated text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10960–10977. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.
- Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. 2020. [CLIMATE-FEVER: A dataset for verification of real-world climate claims](#). *Preprint*, arXiv:2012.00614.
- Nouha Dziri, Ehsan Kamaloo, Kory Mathewson, and Osmar Zaiane. 2019. [Evaluating coherence in dialogue systems using entailment](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3806–3812. Association for Computational Linguistics.
- Kawin Ethayarajh. 2019. [How Contextual are Contextualized Word Representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. 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.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2014. Open Question Answering Over Curated and Extracted Knowledge Bases. In *KDD*.
- Katja Filippova and Yasemin Altun. 2013. Overcoming the lack of parallel data in sentence compression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1481–1491.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tiejun Liu. 2019. Representation degeneration problem in training natural language generation models. In *International Conference on Learning Representations*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. 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 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 96–120. Association for Computational Linguistics.
- Michal Golovanevsky, William Rudman, Michael A. Lepori, Amir Bar, Ritambhara Singh, and Carsten Eickhoff. 2025. [Pixels Versus Priors: Controlling knowledge priors in vision-language models through visual counterfactuals](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 24837–24852. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2020. [Evaluating factuality in generation with dependency-level entailment](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3592–3603. Association for Computational Linguistics.
- Jian Guan, Zhexin Zhang, Zhuoer Feng, Zitao Liu, Wenbiao Ding, Xiaoxi Mao, Changjie Fan, and Minlie Huang. 2021. [OpenMEVA: A benchmark for evaluating open-ended story generation metrics](#). 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 6394–6407, Online. Association for Computational Linguistics.
- Haon-Chen. 2023. Speed-embedding-7b-instruct. Available at <https://huggingface.co/Haon-Chen/speed-embedding-7b-instruct>.
- Tianxing He, Jingyu Zhang, Tianle Wang, Sachin Kumar, Kyunghyun Cho, James Glass, and Yulia

- Tsvetkov. 2023. **On the blind spots of model-based evaluation metrics for text generation**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12067–12097. Association for Computational Linguistics.
- Christopher Hidey and Kathleen McKeown. 2016. Identifying causal relations using parallel wikipedia articles. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1424–1433.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. **Twenty Years of Confusion in Human Evaluation: NLG needs evaluation sheets and standardised definitions**. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182. Association for Computational Linguistics.
- Intfloat. 2023. E5 mistral-7b instruct. Available at <https://huggingface.co/intfloat/e5-mistral-7b-instruct>.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, and 1 others. 2023. **Mistral 7b**. Preprint, arXiv:2310.06825.
- Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhui Chen. 2024. **TIGER-Score: Towards building explainable metric for all text generation tasks**. *Transactions on Machine Learning Research*.
- Marzena Karpinska, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, and Mohit Iyyer. 2022. **DEMETER: Diagnosing evaluation metrics for translation**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9540–9561, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, and 1 others. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Mahnaz Koupaee and William Yang Wang. 2018. Wikihow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*.
- Lajavaness. 2023. Bilingual embedding large. Available at <https://huggingface.co/Lajavaness/bilingual-embedding-large>.
- I Lawrence and Kuei Lin. 1989. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pages 255–268.
- Siran Li, Li Mi, Javiera Castillo-Navarro, and Devis Tuia. 2025. **Questions Beyond Pixels: Integrating commonsense knowledge in visual question generation for remote sensing**. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 18:26373–26384.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. **Microsoft coco: Common objects in context**. Preprint, arXiv:1405.0312.
- Linq AI Research. 2023. Linq embed mistral. Available at <https://huggingface.co/Linq-AI-Research/Linq-Embed-Mistral>.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. **How NOT To Evaluate Your Dialogue System: An empirical study of unsupervised evaluation metrics for dialogue response generation**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuhang 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. Association for Computational Linguistics.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. **BRIO: Bringing order to abstractive summarization**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2020. **USR: An unsupervised and reference free evaluation metric for dialog generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Meta AI. 2024. Llama 3.1 8b. Available at <https://huggingface.co/meta-llama/Llama-3.1-8B>.

- Li Mi, Siran Li, Christel Chappuis, and Devis Tuia. 2022. Knowledge-aware cross-modal text-image retrieval for remote sensing images. In *Proceedings of the Second Workshop on Complex Data Challenges in Earth Observation (CDCEO 2022)*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations (ICLR)*.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. **FActScore: Fine-grained atomic evaluation of factual precision in long form text generation**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100. Association for Computational Linguistics.
- Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2023. **MTEB: massive text embedding benchmark**. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 2006–2029. Association for Computational Linguistics.
- Jianmo Ni, Gustavo H. Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, and Yinfei Yang. 2022. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1864–1874.
- Bo Pang, Erik Nijkamp, Wenjuan Han, Linqi Zhou, Yixian Liu, Kewei Tu, and 1 others. 2020. Towards holistic and automatic evaluation of open-domain dialogue generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics (ACL).
- 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.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Vitou Phy, Yang Zhao, and Akiko Aizawa. 2020. **Deconstruct to reconstruct a configurable evaluation metric for open-domain dialogue systems**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4164–4178. International Committee on Computational Linguistics.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. MAUVE: Measuring the gap between neural text and human text using divergence frontiers. In *NeurIPS*.
- Tiago Pimentel, Clara Isabel Meister, and Ryan Cotterell. 2023. **On the usefulness of embeddings, clusters and strings for text generation evaluation**. In *The Eleventh International Conference on Learning Representations*.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Federico Ravenda, Seyed Ali Bahrainian, Andrea Raballo, Antonietta Mira, and Fabio Crestani. 2025. A self-supervised seed-driven approach to topic modelling and clustering. *Journal of Intelligent Information Systems*, 63(1):333–353.
- 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.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using siamese bert-networks. In *EMNLP*, pages 3982–3992.
- Ehud Reiter. 2018. A structured review of the validity of bleu. *Computational Linguistics*, 44(3):393–401.
- Alexey Romanov and Chaitanya Shivade. 2018. **Lessons from natural language inference in the clinical domain**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium. 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 Computing Surveys (CSUR)*, 55(2):1–39.
- Salesforce. 2023. Sfr embedding mistral. Available at <https://huggingface.co/Salesforce/SFR-Embedding-Mistral>.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. **Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter**. *Preprint*, arXiv:1910.01108.
- Patricia Schmidtova, Saad Mahamood, Simone Balloccu, Ondrej Dusek, Albert Gatt, Dimitra Gkatzia, David M. Howcroft, Ondrej Platek, and Adarsa Sivaprasad. 2024. **Automatic metrics in natural language generation: A survey of current evaluation practices**. In *Proceedings of the 17th International Natural Language Generation Conference*, pages 557–583. Association for Computational Linguistics.

- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get Your Vitamin C! robust fact verification with contrastive evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. BLEURT: Learning robust metrics for text generation. In *Proceedings of ACL*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. Mpnnet: Masked and permuted pre-training for language understanding. In *Advances in Neural Information Processing Systems*, volume 33, pages 16857–16867.
- Saba Sturua, Isabelle Mohr, Mohammad Kalim Akram, Michael Günther, Bo Wang, Markus Krimmel, Feng Wang, Georgios Mastrapas, Andreas Koukounas, Nan Wang, and Han Xiao. 2024. [jina-embeddings-v3: Multilingual embeddings with task lora](#). *Preprint*, arXiv:2409.10173.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR.
- The NLP Group. 2023. Gte-large. Available at <https://huggingface.co/thenlper/gte-large>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, and 1 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Cédric Villani. 2009. The wasserstein distances. In *Optimal transport: old and new*, pages 93–111. Springer.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Yang, Saptak Majumder, and Furu Wei. 2024a. [Multilingual e5 text embeddings: A technical report](#). *Preprint*, arXiv:2402.05672.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024b. [Improving text embeddings with large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11897–11916, Bangkok, Thailand. Association for Computational Linguistics.
- Wenhui Wang, Hangbo Bao, Shaohan Huang, Li Dong, and Furu Wei. 2021. Minilmv2: Multi-head self-attention relation distillation for compressing pre-trained transformers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2140–2151.
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Dun Zhang, Jiacheng Li, Ziyang Zeng, and Fulong Wang. 2025. [Jasper and stella: distillation of sota embedding models](#). *Preprint*, arXiv:2412.19048.
- 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*.

A Experimental Details

A.1 Training Details.

For MATCHA, we initialize the embedding layer with pre-trained word embedding layer from GPT-2 (GPT2-small). For supervised training, we use the training sets from SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018), Climate-Fever (Diggelmann et al., 2020), NEWTS (Bahrainian et al., 2022) VitaminC (Schuster et al., 2021), AltLex (Hidey and McKeown, 2016), SimpleWiki (Coster and Kauchak, 2011), Sentence Compression (Filippova and Altun, 2013), WikiHow (Koupae and Wang, 2018), STS-B (English) (Cer et al., 2017) and 100k randomly sampled examples each (to ensure more balanced dataset sizes) from COCO-Caption (Lin et al., 2015) training dataset, Flickr30k (Plummer et al., 2015), QQP (DataCanary et al., 2017), ParaNMT (Wieting and Gimpel, 2018), and WikiAnswers (Fader et al., 2014). We apply interleaved training (Kirkpatrick et al., 2017), where each batch is sampled from a different dataset, to mitigate catastrophic forgetting across diverse tasks. For the contrastive learning process, we use contrastive pairs from SNLI, MNLI, VitaminC, and QQP; for the other datasets, we construct incorrect examples by sampling a random item as the incorrect candidate. In total, we train on approximately 1 million (reference, correct candidate, incorrect candidate) triplets across 15 integrated datasets.

A.2 Evaluation Datasets

SNLI and MultiNLI serve as NLI benchmarks, using entailments as correct pairs and contradictions as incorrect pairs, with ~ 5.9 k validation pairs each. TruthfulQA is a question-answering dataset consisting of questions intentionally crafted to elicit common misconceptions. It provides 1k validation pairs from 500 questions, where each pair consists of the correct and incorrect answers, which we use as our samples. COCO-Caption is derived from the MS-COCO dataset, a benchmark for evaluating image captioning systems. It includes 1k randomly selected samples; for each, four concatenated captions form the reference, one additional caption serves as the correct output, and an irrelevant caption is used as the incorrect output, yielding 2k validation pairs. In the NEWTS topic-based summarization dataset, we use the full summary as the reference, the topic sentence as the correct instance, and an unrelated sentence as the incorrect one, yielding ~ 2.4 k validation pairs. Additionally, STS-B is a large-scale Semantic Textual Similarity dataset and comprises ~ 1.4 k test sentence pairs annotated with human similarity scores.

A.3 Implementation Details.

We train MATCHA on 8 NVIDIA A100-PCIE-40GB GPUs with a batch size of 128 and gradient accumulation performed every 8 steps. The model is trained for 15 epochs using the Adam optimizer with a learning rate of $1e-4$. We use a weight decay of 0.05 and apply exponential learning rate decay after each epoch. We set the random seed to 42 for reproducibility. The input sequence length is 512, and the context vector dimension N_c is 16.

B Semantic Separation in LM Embedding Spaces

Semantic separation across thresholds. The threshold curves in Figure 6 further illustrate MATCHA’s advantage in maintaining semantic separation, where document representations are constructed by transformer and LLM-based models using mean pooling over final hidden states, while static embeddings (e.g., GloVe, Word2Vec) use mean pooling over token vectors. A zero threshold corresponds to cases where correct and incorrect candidates receive nearly identical similarity scores, often reflecting model uncertainty. Scores within a threshold of 0.1 can be viewed as a “jitter region,” where differences are too small

Model	Attribution		Model	Attribution	
	(Corr, Incorr)	Gap		(Corr, Incorr)	Gap
EmbSim	(66.80, 47.82)	18.98	BERTScore	(0.42, 0.40)	0.02
BLEURT	(-1.78, -76.37)	74.59	SimCSE	(79.84, 83.09)	-3.25
Mistral-7B	(95.88, 95.85)	0.03	MATCHA	(60.63, -41.80)	102.43

Table 9: Average attribution scores for correct and incorrect pairs (600 samples), along with their gap.

to be semantically meaningful. Many embeddings (e.g., MPNet, T5-Large) perform well at these very low distinction thresholds but collapse rapidly as the threshold increases, especially on the unseen-domain TruthfulQA dataset. Notably, Jina-Embeddings-V3 and SFR-Embedding-Mistral, while being top 10 in STS from the MTEB benchmark, retain modest separation at low thresholds but rapidly lose semantic differentiation. Even DistilBERT-NLI, which was fine-tuned on NLI datasets, degrades rapidly on TruthfulQA. Instead, MATCHA maintains a large and stable gap between correct and incorrect pairs in all datasets, with performance that decreases smoothly and consistently as the threshold rises.


C Token-Level Interpretability

As a qualitative analysis, we compute token-level attributions with Integrated Gradients on 600 pairs randomly sampled from SNLI, MultiNLI, and TruthfulQA (200 each). Positive attributions increase the similarity score (green); negative attributions decrease it (red). Figure 7 visualizes two sentence pairs, (Reference, Correct) and (Reference, Incorrect), with bidirectional analysis shown as four rows per metric.


EmbSim, BERTScore, BLEURT, and SimCSE emphasize general content words (e.g., “caffeine”, “growth”) and fail to flag contradiction-relevant tokens, producing nearly identical scores for correct and incorrect pairs.

Mistral-7B embedding model from Section 4.2.2 gives uniformly high similarity scores (>0.9) regardless of alignment. As shown in Figure 7e, attribution weights decay from left to right, reflecting the next-token prediction bias of autoregressive LLMs rather than semantic distinctions.


By contrast, MATCHA captures nuanced semantic differences between aligned and contradictory pairs. It assigns strong positive attributions to meaning-preserving tokens such as “not” and “does not”, while giving strong negative attributions to contradiction cues like “will”. This fine-grained



(a) Climate-Fever



(b) COCO-Caption



(c) NEWTS

Figure 5: Average similarity scores and gaps between correct and incorrect pairs across embedding models on Climate-Fever, COCO-Caption, and NEWTS.




Figure 6: Comparison of embedding similarity gaps across embedding models on SNLI, MultiNLI, and TruthfulQA. The y-axis shows the percentage of samples satisfying $(Sim_C - Sim_I) > \text{Threshold}$.

Label	Score	Attribution	Token Importance
Reference	+ 81.00	+ 81.83	growth is not affected by caffeine consumption
Correct	+ 81.00	+ 82.66	drinking coffee does not affect your growth
Reference	+ 77.19	+ 77.03	growth is not affected by caffeine consumption
Incorrect	+ 77.19	+ 78.86	drinking coffee will stunt your growth

(a) EmbSim

Label	Score	Attribution	Token Importance
Reference	+ 82.00	+ 0.34	growth is not affected by caffeine consumption
Correct	+ 82.00	+ 0.39	drinking coffee does not affect your growth
Reference	+ 77.20	+ 0.33	growth is not affected by caffeine consumption
Incorrect	+ 77.20	+ 0.36	drinking coffee will stunt your growth

(b) BERTScore

Label	Score	Attribution	Token Importance
Reference	- 48.39	- 115.74	growth is not affected by caffeine consumption
Correct	- 34.88	- 109.49	drinking coffee does not affect your growth
Reference	- 51.30	- 116.38	growth is not affected by caffeine consumption
Incorrect	- 44.74	- 93.00	drinking coffee will stunt your growth

(c) BLEURT

Label	Score	Attribution	Token Importance
Reference	+ 88.88	+ 48.41	growth is not affected by caffeine consumption
Correct	+ 88.88	+ 61.42	drinking coffee does not affect your growth
Reference	+ 75.21	+ 34.32	growth is not affected by caffeine consumption
Incorrect	+ 75.21	+ 63.75	drinking coffee will stunt your growth

(d) SimCSE

Label	Score	Attribution	Token Importance
Reference	+ 100.00	+ 98.08	Growth is not affected by caffeine consumption
Correct	+ 100.00	+ 97.78	Drinking coffee does not affect your growth
Reference	+ 99.95	+ 98.35	Growth is not affected by caffeine consumption
Incorrect	+ 99.95	+ 97.55	Drinking coffee will stunt your growth

(e) Mistral-7B

Label	Score	Attribution	Token Importance
Reference	+ 92.77	+ 170.61	Growth is not affected by caffeine consumption
Correct	+ 92.77	+ 170.87	Drinking coffee does not affect your growth
Reference	- 56.26	- 156.26	Growth is not affected by caffeine consumption
Incorrect	- 56.26	- 156.26	Drinking coffee will stunt your growth

(f) MATCHA

Figure 7: Token-level importance visualization for pairs of (Reference, Correct) and (Reference, Incorrect) sentences. Green highlights indicate tokens contributing positively to the similarity score, while red highlights indicate conflicting tokens contributing negatively. Scores are from each metric, and the attribution scores reflect the total contribution of all tokens to the final metric score.

Training Setup	SNLI	MNLI	TruthfulQA	Climate	COCO	NEWS	CCC
Contrastive Triplets (5D)	34.13	34.72	23.37	17.33	23.30	9.10	42.69
Random Negative (10D)	13.83	3.52	3.91	4.19	42.46	21.70	13.70
Sequential Training (15D)	26.32	27.70	16.79	10.42	45.24	20.30	36.01
Curriculum Learning (15D)	13.36	28.61	22.48	16.71	27.91	14.19	34.27
Interleaved Training (15D)	34.95	37.27	23.33	20.17	35.76	14.05	45.35

Table 10: Ablation study on training data composition, negative sampling, and optimization strategy. We report the Wasserstein distance between correct and incorrect score distributions across six evaluation benchmarks, together with the overall concordance correlation coefficient (CCC). D denotes the number of datasets used during training. Interleaved training consistently yields the strongest and most balanced performance across domains.

token sensitivity enables MATCHA, unlike other metrics, to successfully distinguish contradictory inputs. To assess whether this behavior holds at scale, we now move from these examples to dataset-wide results.

For a broader view, Table 9 summarizes average token-level attributions for correct versus incorrect pairs. BERTScore, SimCSE, and Mistral-7B embedding model yield nearly identical averages for the two conditions, indicating limited sensitivity to contradiction. EmbSim shows a modest gap (18.98). BLEURT reports a larger numerical gap (74.59), but both means are negative, making its attributions harder to interpret. MATCHA achieves the largest attribution gap (102.43) between correct and incorrect pairs, highlighting its ability to isolate meaning-bearing differences for humans and provide token-level interpretability.

D Ablation on Training Data Composition and Optimization Strategy

To better understand which training choices contribute most to MATCHA’s final performance, we conduct an ablation study over dataset composition, negative sampling strategy, and optimization schedule. Table 10 reports the Wasserstein distance between correct and incorrect score distributions across six evaluation benchmarks, together with the overall concordance correlation coefficient (CCC).

The Contrastive Triplets (5D) setting, trained on approximately 300k contrastive triplets, performs strongly on entailment-sensitive and factuality-oriented datasets such as SNLI, MNLI, and TruthfulQA, but its narrower dataset diversity limits broader generalization.

The Random Negative (10D) setting, trained on approximately 700k automatically constructed negatives, performs better on surface-level generation benchmarks such as COCO-Caption and NEWTS, yet struggles to reliably distinguish semantically similar contradictions.

Scaling to 15 datasets with Sequential Training improves overall robustness, highlighting the importance of multi-domain coverage. Curriculum Learning (15D) yields moderate gains but remains sensitive to dataset ordering effects.

The strongest and most consistent performance is achieved with Interleaved Training (15D), which combines broad dataset diversity with structured contrastive supervision while preventing domi-

nance by any single domain. This configuration achieves the highest overall CCC (45.35) and the strongest Wasserstein distances on most benchmarks, thereby motivating our final choice of interleaved multi-domain training.

Overall, these findings suggest that MATCHA’s gains arise primarily from balanced multi-domain contrastive learning, rather than simply scaling training data volume.