

Favi-Score: A Measure for Favoritism in Automated Preference Ratings for Generative AI Evaluation

Pius von Däniken and Jan Deriu and Don Tuggener and Mark Cieliebak

Centre for Artificial Intelligence
ZHAW School of Engineering
{vode,deri,tuge,ciel}@zhaw.ch

Abstract

Generative AI systems have become ubiquitous for all kinds of modalities, which makes the issue of the evaluation of such models more pressing. One popular approach is preference ratings, where the generated outputs of different systems are shown to evaluators who choose their preferences. In recent years the field shifted towards the development of automated (trained) metrics to assess generated outputs, which can be used to create preference ratings automatically. In this work, we investigate the evaluation of the metrics themselves, which currently rely on measuring the correlation to human judgments or computing sign accuracy scores.

These measures only assess how well the metric agrees with the human ratings. However, our research shows that this does not tell the whole story. Most metrics exhibit a disagreement with human system assessments which is often skewed in favor of particular text generation systems, exposing a degree of favoritism in automated metrics. This paper introduces a formal definition of favoritism in preference metrics, and derives the Favi-Score, which measures this phenomenon. In particular we show that favoritism is strongly related to errors in final system rankings. Thus, we propose that preference-based metrics ought to be evaluated on both sign accuracy scores and favoritism.

1 Introduction

With the rise of Generative AI for text, images, code, music, and other modalities, the question of how to evaluate the quality of the outputs of these systems is becoming more and more important. *Preference Ratings* have become a ubiquitous strategy for this task (Belz and Kow, 2010; Deriu et al., 2020; Stiennon et al., 2020; Freitag et al., 2021, 2022). The general idea is that the outputs of two generative systems are shown to a (typically human) rater who decides which output they prefer, or whether both outputs are of equal quality.

The ratings are then aggregated over a test set to compute in how many cases the outputs of one system were preferred over the other and whether the difference is significant.

Since human-based evaluations are expensive, there has been a growing movement towards training automated metrics to assess generated content. (Celikyilmaz et al., 2020; Deriu et al., 2021; Yeh et al., 2021). So instead of showing the outputs of two generative systems to human raters, one uses an automatic metric and hopes that the evaluation will resemble the ratings of the humans.¹ This raises the question of how these automated metrics themselves ought to be evaluated. Currently, the most popular method is measuring the correlation of the metrics' ratings to human judgments, or in the case of preference ratings, *sign-accuracy* on both the sample and system level is computed (i.e., on how many preference ratings do human and metrics agree, or on how many system-pairs do human and metrics agree) (Kocmi et al., 2021). For some tasks, such as machine translation, we have seen that according to these, metrics are steadily improving (Freitag et al., 2022).

While both measures are important in the assessment of automated metrics, they only count *how many* mistakes the metric typically makes, ignoring how these mistakes are distributed. That is, whether the mistakes are systematically skewed in favor of one generative system under investigation or another. To underscore the significance of this information, let us consider the assessment of two generative systems, denoted as π_a and π_b , using a metric that exhibits only a 10% error rate, which are all skewed towards π_a . This means that in 10% of cases, the metric's ratings will deviate from human judgments and in all cases the deviation is in favor π_a . Furthermore, assume that according

¹For this paper, we treat human ratings as gold-standard, but we are aware of the issues with human ratings (Amidei et al., 2018; Belz et al., 2021)

to humans π_a is slightly worse than π_b , then the favoritism of the metric will cause π_a to be rated as being better than π_b according to the metric.

Thus, the research question that we tackle in this work is how we can measure favoritism, and investigate its impact on evaluation.

Contributions. This work has three main contributions:

- We motivate and formally define the problem of favoritism in automated preference metrics for generative systems.
- We introduce the Favi-Score, an easy-to-compute score to quantify the favoritism exhibited by an automated preference metric in comparison to human judgments.
- We apply the Favi-Score to a variety of text generation tasks,² showcasing the interplay between the Favi-Score and existing measurements of preference-metric quality.


Our main finding is that favoritism causes mistakes in the rankings of systems according to the metric: even a metric with a high sign accuracy can lead to wrong evaluation if it has a strong favoritism. Or stated differently, a metric with low sign accuracy and no favoritism can still yield the correct ranking of systems.³

2 Impact of Favoritism on System Rankings

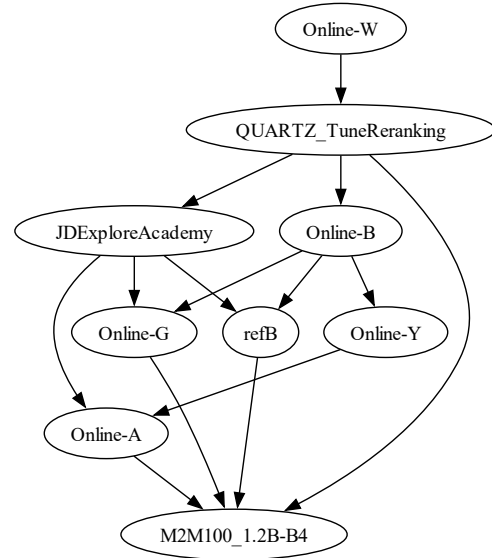
To demonstrate the need for the Favi-Score, we use the data of the WMT-22 metrics shared task (Freitag et al., 2022). This contains ratings for 9 machine translation systems. Based on this data, we produce one ranking using human ratings and one ranking using COMET-22 (Rei et al., 2022), a well-established automated metric for machine translation outputs. We visualize their assessments using directed acyclic graphs (DAGs), as depicted in Figure 1. For this, the preference ratings are computed for each pair of systems using the human ratings and the corresponding COMET-22 ratings. Then, we put an edge between two systems if one system is statistically significantly preferred to the other system according to a sign test (Coakley and Heise, 1996). For better visibility, we omit transitive edges. Note that the system-level agreement lies at 61.1%,

²Note that the formal definition of the Favi-Score could also be applied to other modalities. Thus, being agnostic to the task/domain at hand (e.g., music, images, or code generation).

³We provide an implementation of the Favi-Score: <https://github.com/vodezhaw/faviscore>



(a) Human



(b) COMET-22-refA

Figure 1: Visualization of the consequences of unfair metrics by comparing the human ranking to the COMET ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.

i.e., in 61.1% the edges of the COMET-22 DAG correspond to the ones of the human DAG.

We see in Figure 1a that according to humans, *ref-B* and *Online-W* are both at the top. However, in Figure 1b according to COMET-22, *ref-B* is ranked much worse, while *QUARTZ_TuneReranking* is disproportionately favored by COMET-22 compared to the human ranking. This is a direct consequence of favoritism.

Note that *only* considering the DAG does not reveal the whole story. For example, if a metric unduly favors a system that humans also rank highly, the overall ranking might not change. Yet, the favored system could appear to win by a much larger

margin than it deserves, exaggerating its performance. For a concrete example, if we consider the outcome percentages between *Online-W* and *JDExploreAcademy*, then under human evaluation, *Online-W* is rated better in 24.1% of cases and worse in 18.6% of cases (giving a margin of 5.5%). Under COMET-22, it is rated better in 57.5% of cases and worse in 42.2% (giving a margin of 15.3%). Even though *Online-W* is rated better than *JDExploreAcademy* in both cases, the perceived difference is much more pronounced under COMET-22, which is unfaithful to the human evaluation even though there is sign agreement. A concrete consequence is that, for example, a sign test will report much higher significance (lower p-value) for evaluation under COMET-22 than for human evaluation. This example shows the need for a mechanism that measures the inherent favoritism of a metric.

Informally, the Favi score is a relative measure between the human ratings and the metric ratings. *Instead of measuring how much the humans and metrics agree on their preference, it is a function of those samples where the humans and metrics disagree. The main idea is to count and weigh the cases in which the metric disproportionately assigns more samples in favor of one system than the human judges.* Thus, the Favi score yields a more fine-grained view compared to the DAG, as it also can measure the cases where there are no changes to the DAG, but one system is disproportionately favored (or disfavored) by the metric (see Examples in Appendix D). The Favi score can be used as a diagnostic tool to discover those cases. In the subsequent sections, we will formally introduce the Favi score.

3 Related Work

We are not aware of any prior work that defines or analyses favoritism in preference ratings. Most work on the analysis of metrics for text generation tasks is concerned with the overall performance, which is measured in terms of correlation to human judgments or sign accuracy (Deriu et al., 2021; Fabbri et al., 2021; Yeh et al., 2021; Freitag et al., 2022). The main conclusion of these analyses is that automated metrics do not yet match human ratings, and that the performance of the metrics vary depending on the domain and task. Moreover, there is a growing body of work that analyses individual metrics for generative tasks and their problems

(Schluter, 2017; Freitag et al., 2020; Amrhein and Sennrich, 2022; Hanna and Bojar, 2021, inter alia), but these studies do not compare or assess different metrics.

Novikova et al. (2017) go a step further and analyse specific weaknesses of various metrics for data-to-text. They showed that the metrics had difficulties distinguishing medium from good-quality outputs. Similarly, Kryscinski et al. (2019) showed that in the case of automated text summarization, automated metrics behave differently depending on whether the systems are extractive or abstractive. Thus, there is evidence in the literature that automated metrics are biased toward certain types of outputs and have difficulties with certain features. However, we are not aware of a fundamental analysis and discussion on the topic of structurally preferring the outputs of one system over others.

A different line of work is concerned with adversarial attacks on metrics. Sai et al. (2019) showed that ADEM (Lowe et al., 2017), a trained metric, is susceptible to specific changes in the inputs. Deriu et al. (2022) devised a method based on reinforcement learning to find nonsensical adversarial samples that elicit perfect scores from metrics for conversational dialogue systems. These findings indicate a lack of robustness and raise the question why certain outputs are preferred.

Finally, the last type of work related to ours is concerned with a more fundamental analysis of metrics. Wei and Jia (2021) apply the bias-variance-noise decomposition from Domingos (2000) to assess statistical bias of automated metrics used in the evaluation of machine translation and summarization. Their analysis produces a ranking of automated metrics based on their statistical bias, but it does not indicate the direction of the bias. Also, their approach does not investigate the relation between the bias strength and the systems under scrutiny. While they find that automated metrics for Natural Language Generation (NLG) are biased, we extend this notion and show that the strengths of the biases vary on the system level and affects system rankings. Chaganty et al. (2018) analyse the bias and variance of automated metrics. They find that automated metrics correlate poorly with human judgments, and that there is a correlation bias, i.e. the metrics often agree with humans in judging bad outputs but correlate poorly with humans on average or good quality outputs.

To summarize, there is strong evidence that automated metrics exhibit behavior which can benefit

or harm certain systems. To our knowledge, our work provides a definition and a measure for quantifying the amount of this benefit or harm for the first time.

4 Preference Ratings

In this section, we formalize the notion of a preference rating and define the confusion matrix which shows the amount of disagreement of two preference rating schemes (e.g. human judgments vs. an automated metric).

We first start by formally defining a *Generative System (GS)* as a function from an input space to an output space.

Definition 1 (Generative System (GS)). *Let \mathcal{I} denote the set of all possible inputs to a GS, and \mathcal{O} the set of all possible outputs. Then, a GS is a function π that takes an input and generates an output*

$$\pi : \mathcal{I} \rightarrow \mathcal{O} \quad (1)$$

Note that the input and output spaces are defined quite abstractly and are task-specific. For instance, in machine translation, the input space might be the set of utterances in the source language and the output space the set of utterances in the target language, whereas in news summarization, the input space would be the set of all news articles and the outputs is the set of all texts.

Next, we define the *Preference Rating* as a function of a triple of one input and two outputs to a sign.

Definition 2 (Preference Rating). *Let \mathcal{I} denote the set of all possible inputs to a GS, and \mathcal{O} the set of all possible outputs. Then, we define a preference rating as:*

$$\mathcal{R} : \mathcal{I} \times \mathcal{O} \times \mathcal{O} \rightarrow \{+, =, -\} \quad (2)$$

where $+$ denotes that the first output is preferred over the second, $=$ denotes that both outputs are of equal quality, and $-$ denotes that the second output is preferred.

Thus, for two generative systems π_1 and π_2 , $r[i] = \mathcal{R}(i, \pi_1(i), \pi_2(i))$ is the preference rating that compares the outputs of system π_1 and system π_2 for the input $i \in \mathcal{I}$. In the following, we denote a human-based preference rating as $r^H[i] = \mathcal{R}^H(i, \pi_1(i), \pi_2(i))$ and an automated metric-based preference rating as $r^A[i] = \mathcal{R}^A(i, \pi_1(i), \pi_2(i))$, respectively.

In order to run an evaluation and apply a preference rating \mathcal{R} to decide whether π_1 performs better than π_2 , we first define the dataset that defines the evaluation setting.

Definition 3 (Evaluation Setting). *Given a test-set $\mathcal{T} \subseteq \mathcal{I}$, a pair of text generation systems π_1 and π_2 , and two preference ratings \mathcal{R}^H and \mathcal{R}^A , we define the evaluation setting as:*

$$\mathcal{E} = \{i, \pi_1(i), \pi_2(i), r^H[i], r^A[i] | i \in \mathcal{T}\} \quad (3)$$

Thus, the evaluation setting \mathcal{E} is composed of the inputs of the test set, the outputs of both TG systems, and the two preference ratings (e.g. according to humans and according to an automated metric).

For a given evaluation setting \mathcal{E} , we can compute the confusion matrix between \mathcal{R}^A and \mathcal{R}^H .

Definition 4 (Confusion Matrix). *Given an evaluation setting \mathcal{E} , we define the confusion matrix C as*

$$C = \begin{pmatrix} C_{++} & C_{+=} & C_{+-} \\ C_{=+} & C_{==} & C_{=-} \\ C_{-+} & C_{-=} & C_{--} \end{pmatrix} \quad (4)$$

where $C_{mn} = \sum_{i \in \mathcal{T}} \mathbb{I}[r^H[i] = m \wedge r^A[i] = n]$ and \mathbb{I} is the indicator function.

This corresponds to the usual confusion matrix known from classification settings, and the entries C_{mn} count how many times a rating of type m is classified as type n .

We call the cases where the automated rating disagrees with the human rating *errors*.

Definition 5 (Error). *Given an evaluation setting \mathcal{E} , $(i, \pi_1(i), \pi_2(i), r^H[i], r^A[i]) \in \mathcal{E}$ is an error (of the automated rating with respect to the human rating) if $r^A[i] \neq r^H[i]$.*

Definition 6 (Total Error). *Given an evaluation setting \mathcal{E} , the total error is given by $E = \sum_{i \in \mathcal{T}} \mathbb{I}[r^A[i] \neq r^H[i]]$.*

To decide which of the two systems is better, we define the *Outcome* of the evaluation, by counting the preferences.

Definition 7 (Outcome). *Given an evaluation setting \mathcal{E} , we define the outcome d as*

$$d = (d_+, d_-, d_0) \quad (5)$$

where $d_n = \sum_{i \in \mathcal{T}} \mathbb{I}[r^H[i] = n]$. We call $mar(d) = d_+ - d_-$ the *outcome margin*.

The outcome aggregates the preference ratings and counts how often one system is preferred over the other. For simplicity, we denote \mathbf{d} as the outcome computed using the human ratings, and $\hat{\mathbf{d}}$ as the outcome according to the metric.

5 Measuring Favoritism

In this section, we define favoritism and derive a score to quantify it. Our view of favoritism is fundamentally relative, since it is based on comparing the behaviour of a metric with respect to the ground-truth provided by humans. In theory, one could compute favoritism compared to some other set of ground-truth labels (e.g., comparing two sets of human ratings). However, to keep our discussion simple, we treat human ratings as the unbiased ground truth, and leave deviations of this assumption to future work.

Our measure of favoritism, called Favi-Score, is based on two observations about errors of an automated rating with respect to the human rating. First, every error is necessarily in favor of a specific system. For example, changing a + to a =, is favoring π_2 . Second, not every error has the same severity. Switching a preference from one system to another (i.e., + to a -) is worse than classifying a draw as a preference (i.e., = to a -), since the former has a larger impact on the outcome margin than the latter. This means that the favoritism measure should depend both on the proportion of errors in favor of one system and the severity of those errors.

$$\begin{array}{cc} \begin{pmatrix} 100 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 100 & \mathbf{0} \\ \mathbf{10} & \mathbf{0} & 90 \end{pmatrix} & \begin{pmatrix} 100 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & 100 & \mathbf{0} \\ \mathbf{0} & \mathbf{10} & 90 \end{pmatrix} \\ \text{(a) } \mathbf{C}_1, \hat{\mathbf{d}}_1 & \text{(b) } \mathbf{C}_2, \hat{\mathbf{d}}_2 \\ \begin{pmatrix} 90 & \mathbf{0} & \mathbf{10} \\ \mathbf{0} & 100 & \mathbf{0} \\ \mathbf{10} & \mathbf{0} & 90 \end{pmatrix} & \begin{pmatrix} 90 & \mathbf{10} & \mathbf{0} \\ \mathbf{0} & 100 & \mathbf{0} \\ \mathbf{10} & \mathbf{0} & 90 \end{pmatrix} \\ \text{(c) } \mathbf{C}_3, \hat{\mathbf{d}}_3 & \text{(d) } \mathbf{C}_4, \hat{\mathbf{d}}_4 \end{array}$$

Figure 2: Example Confusion Matrices \mathbf{C} and outcomes $\hat{\mathbf{d}}$ according to automated metrics for the human outcome $\mathbf{d} = (100, 100, 100)$.

To better illustrate the interplay of these criteria, let us consider the confusion matrices depicted in Figure 2 for an outcome of $\mathbf{d} = (100, 100, 100)$

alongside the corresponding outcome according to the metric $\hat{\mathbf{d}}$. We can see the error counts in the off-diagonal entries of \mathbf{C}_i , where the lower triangular sub-matrix, indicated in magenta and red, corresponds to errors in favor of π_1 and the upper triangular sub-matrix, indicated in cyan and blue, corresponds to errors in favor of π_2 . Let us first compare \mathbf{C}_1 and \mathbf{C}_2 . In both cases there are 10 errors that are all in favor of π_1 . Nevertheless, \mathbf{C}_1 shows more favoritism since its errors are more severe, favoring π_1 more strongly. The consequence is visible in the outcome margin, where $\text{mar}(\hat{\mathbf{d}}_1) = 110 - 90 = 20$ and $\text{mar}(\hat{\mathbf{d}}_2) = 100 - 90 = 10$. \mathbf{C}_3 has more total errors, namely 20, but they are equally distributed in favor of each system and of equal severity. Therefore, \mathbf{C}_3 does not favor either system, which is also visible in the outcome margin $\text{mar}(\hat{\mathbf{d}}_3) = 0$. Finally, \mathbf{C}_4 also has 20 errors distributed equally in favor of either system. But since the errors in favor of π_1 are more severe (i.e., mistaking + and -) than the ones in favor of π_2 (i.e., mistaking + and =), \mathbf{C}_4 favors π_1 overall. Thus, favoritism depends on how the errors are distributed, and how the severity of the mistakes are counted. The severity can be quantified in terms of how much they change the outcome according to the metric. Mistaking + for - leads to a change of 2 in the outcome margin, whereas mistaking + with = changes the outcome margin by 1. Thus, we weigh these mistakes accordingly and define the error cost:

Definition 8 (Directed Error Cost). *Let the following matrix be the error cost matrix:*

$$\mathbf{W} = \begin{pmatrix} 0 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \end{pmatrix}$$

, where the magnitude of each entry denotes the cost for each mistake, and the sign denotes the direction of the favoritism.

Thus the cost of an error reflects its impact on the outcome margin, and the errors in favor of π_1 have a positive weight and errors in favor of π_2 have a negative weight. In cases where $m = n$ the cost $W_{mn} = 0$. Based on these observations, we will define a measure of favoritism, called Favi-Score Φ , that corresponds to the expected weight of the errors of an automated rating with respect to the human rating.

Definition 9 (Favi-Score). *Given an evaluation setting \mathcal{E} and its confusion matrix \mathbf{C} the Favi-Score*

$$C_5 = \begin{pmatrix} 360 & 180 & 60 \\ 20 & 40 & 40 \\ 90 & 90 & 120 \\ 470 & 310 & 220 \end{pmatrix}$$

Figure 3: Another example confusion matrix for illustration for $\mathbf{d} = (600, 100, 300)$.

is

$$\Phi(C) = \sum_{m,n} W_{mn} \frac{C_{mn}}{E}$$

The final score ranges from -2 to 2, where positive values indicate favoritism of π_1 and negative values indicate favoritism of π_2 .

The term $\frac{C_{mn}}{E}$ corresponds to the probability of a given type of error in cases where $m \neq n$, meaning $\frac{C_{mn}}{E} = P(r^H[i] = m \wedge r^A[i] = n | r^A[i] \neq r^H[i])$.

Thus, this brings us to the following interpretation of the Favi-Score:

The Favi-Score is the expected Error Cost.

There is an equivalent definition of the Favi-Score based on the outcome margins.

Lemma 1. *The following equality holds:*

$$\Phi(C) = \frac{\text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d})}{E} \quad (6)$$

Thus, the Favi-Score can also be interpreted as the average change in outcome margin per error.⁴

Coming back to our examples, the Favi-Scores for the confusion matrices in Figure 2 are therefore: $\Phi(C_1) = 2$, $\Phi(C_2) = 1$, $\Phi(C_3) = 0$, and $\Phi(C_4) = \frac{1}{2}$. In Figure 3 we show a more complex example. In this case more errors are in favor of π_2 (280 vs. 200), but more of the severe errors are in favor of π_1 (90 vs. 60). The overall Favi-Score $\Phi(C_5) = -0.104$ shows slight favor for π_2 . The score can be interpreted as each error contributing around 0.104 increase in margin in favor of π_2 .

6 Favi-Score vs. Sign Accuracy

In this section we explore the relation of the Favi-Score to sign accuracy. We show that the Favi-Score and sign accuracies measure different properties of a metric, and that the scores are complementary to each other.

First, we show the relationship between the Favi-Score and the system level sign accuracy. For this,

⁴The proof is in Appendix A.

we first define the system-level sign-accuracy using our notation.

Definition 10 (System-level sign accuracy). *Given $N = \binom{|\Pi|}{2}$ different pairs of systems. The formula for the system-level sign accuracy is given by:*

$$\frac{1}{N} \sum_{\pi^a, \pi^b} \mathbb{I}[\text{sgn}(d_+ - d_-) = \text{sgn}(\hat{d}_+ - \hat{d}_-)] \quad (7)$$

That is, the system level sign accuracy measures the number of times where the outcomes of the metric and the human evaluation agree on which system had more outputs of superior quality. The relationship between the Favi-Score and the System-level sign accuracy is that if there is no favoritism, the sign accuracy will be 1.

Lemma 2 (No System Level Sign Change). *If $\Phi(C) = 0$, that is, there is no favoritism, then the system-level sign will remain the same under the metric. That is: $\mathbb{I}[\text{sgn}(d_+ - d_-) = \text{sgn}(\hat{d}_+ - \hat{d}_-)] = 1$.*

Thus, a metric with no favoritism ensures that the outcome will lead to the same ranking. However, note that as discussed in Section 1, the converse might not hold, i.e., a system level sign accuracy of 1 does not imply that there is no favoritism. This brings us to the following insight:

Metrics with a low Favi-Score will tend to preserve the correct ranking regardless of their sign-accuracy.

The difference between the Favi-Score and the sample-level sign accuracy is apparent by considering its definition:

Definition 11 (Sample-level sign accuracy). *The sample-level sign accuracy is the fraction of samples, which are correctly predicted, i.e., the sum of the diagonal of C divided by the number of samples.*

$$\frac{1}{|\mathcal{T}|} \sum_{m \in \{+, =, -\}} C_{m,m} = \frac{|\mathcal{T}| - E}{|\mathcal{T}|} \quad (8)$$

Thus, sign accuracy is a formula based on the diagonal entries of C , i.e., the fraction of samples where the metric and humans agree. Favoritism, on the other hand, is formulated via the difference of error types in favor of one system or the other, that is the upper and lower triangular sub-matrices. To illustrate this, we come back to our example in Figure 3, where $\Phi(C) = -0.104$, and the sample-level sign accuracy is 0.52. A different way to

showcase the difference is by assuming a fixed number of errors E , then the sign accuracy is 0 for $|\mathcal{T}| = E$, and approaches 1 for $|\mathcal{T}| \rightarrow \infty$, however, $|\mathcal{T}|$ does not influence $\Phi(C)$. Thus, sign accuracy tells us how often we can expect the human and metric to agree on a single sample, whereas favoritism tells us in which direction the mistakes are skewed and by how much.

7 Application of the Favi Score

For the empirical application, we require data in the form highlighted by the evaluation setting $\mathcal{E} = \{i, \pi_1(i), \pi_2(i), r^H[i], r^A[i] | i \in \mathcal{T}\}$. That is where we have a set of text generation systems with paired ratings from humans and automated metrics. For this, we use the data collected by [Deriu et al. \(2023\)](#), which consists of three different tasks (machine translation, summarization, and dialogue systems). We extended their dataset with data from the WMT-22 metrics shared task (see Appendix C for details). Here we provide an overview of the data (see Table 1). Note that for most metrics and human ratings there are either scalar or Likert outputs provided, which we transform to preference ratings by comparing the values (Appendix C).

Chatbot. The chatbot domain consists of the outputs of 5 different chatbots systems and one human reference automatically rated by 5 different metrics for the BlendedSkillTask (BST) dataset ([Smith et al., 2020](#)). The metrics consist of: DEB ([Sai et al., 2020](#)), GRADE ([Huang et al., 2020](#)), HolisticEval ([Pang et al., 2020](#)), MAUDE ([Sinha et al., 2020](#)), and USL-H ([Phy et al., 2020](#)). There are 50 human ratings for each pair of systems, yielding $|\mathcal{E}| = 50$ paired ratings.

Summarization. The summarization domain is based on the data of the SummEval framework ([Fabbri et al., 2021](#)), which provides the outputs of 16 different summarization tools on the CNN/DailyMail corpus ([Nallapati et al., 2016](#)) rated by 7 different automated metrics: BertScore ([Zhang et al., 2019](#)), BLANC ([Vasilyev et al., 2020](#)), CIDEr ([Vedantam et al., 2015](#)), RougeL ([Lin, 2004](#)), S3 ([Peyrard et al., 2017](#)), SummaQA ([Scialom et al., 2019](#)), and SUPERT ([Gao et al., 2020](#)). For 100 outputs, there are expert Likert ratings for each of the features: relevance, coherence, consistency, and fluency, thus $|\mathcal{E}| = 100$.

Machine Translation. For machine translation, we use the WMT-21 metrics task data ([Freitag et al., 2021](#)), and the WMT-22 metrics task data (Fre-

	Chatbot	SummEval	WMT21	WMT22
Data	BST	CNN/DM	News EN->DE	EN->DE
Metrics	5	7	4	5
TG Systems	6	16	11	9
$ \mathcal{E} $	50	100	500	1315

Table 1: Overview of Data Statistics. The ratings refer to the number of ratings available for each pair of TG systems.

[itag et al., 2022](#)). For WMT-21 we use the English to German language pair and the news domain where eight machine translation systems were evaluated, plus three human references. We used the four most prominent metrics: BleuRT ([Sellam et al., 2020](#)), COMET ([Rei et al., 2020](#)), C-SPEC ([Takahashi et al., 2021](#)), and sentence-level BLEU ([Papineni et al., 2002](#)). There are 500 ratings by expert translators for each of the translation systems, thus, $|\mathcal{E}| = 500$. For WMT-22, we use the English to German language pair and selected eight machine translation systems plus one human reference, and five metrics. Namely, BLEU, BleuRT, COMET22 ([Rei et al., 2022](#)), UniTE ([Wan et al., 2022](#)), and METRICX_XL.⁵ We additionally added *ChatGPT*⁶ as a preference metric (see Appendix C.2). For 1315 samples, expert ratings were provided, thus $|\mathcal{E}| = 1315$.

7.1 Results & Discussion

Overall Results. We compare the Favi-Score to the sample- and system-level sign accuracy. Since the Favi-Score is computed for each pair of systems, we report the absolute average Favi-Score over each pair of systems alongside the standard-deviation, while for the two sign accuracy scores, we report the overall score. Figure 4 shows the three scores for each domain and metric. Note that for the Favi-Score (green) we also report the standard deviation (green area). First, we note that each metric exhibits a certain degree of favoritism, with some degree of variance with regards to the system pairs. The magnitude of the Favi-Score varies depending on the domain, for instance, machine-translation metrics exhibit a lower amount of favoritism, while the summarization metrics exhibit a larger amount as well as a larger variance.

System Pair Analysis. To get a better understanding of the interplay between the three measures, we compute the sign-accuracy on the system-pair level (for the system-level sign accuracy this means that we return whether the signs agree or

⁵No reference available.

⁶gpt-3.5-turbo-0613

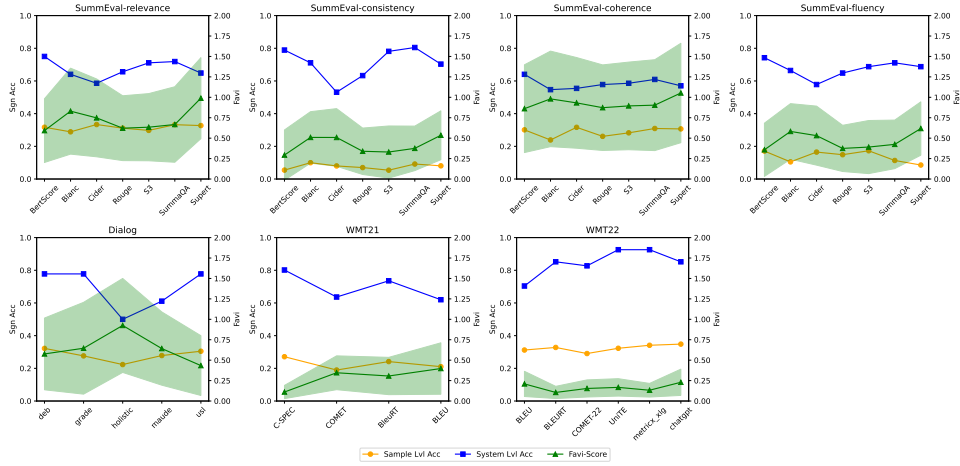




Figure 4: Visualization of the sample-level sign accuracy (orange), system-level sign accuracy (blue), and the average favi-score with a standard deviation (green).



(a) Favi vs System Level Sign Acc



(b) Favi vs Sample Level Sign Acc

Figure 5: The relationship between the Favi-score and both types of sign accuracy. In the upper plot, the boxplot show the distribution of absolute Favi-scores for the cases where the System Level sign agrees for a pair (1) and for the cases where it disagrees (0). The lower plot scatters the Favi-scores vs. the sample level accuracy, where the color showcases the system-level accuracy.

not). First, we investigate whether the system-level sign accuracy correlates to the Favi-Score. In Figure 5a, we show the Favi-Scores depending on




Figure 6: Box Plot of the Favi-Scores for each System for the COMET-22 metric.

whether the sign agrees on the system level (1) or not (0). We observe that system-pairs where the sign does not agree tend to have a larger Favi-Score, which is expected according to Lemma 2. In fact, there is a low-to-moderate inverse correlation between the two scores of Spearman's $\rho = -0.25$. The correlation of the Favi-Score to the sample-level sign accuracy is depicted in Figure 5b, which shows a small correlation (Spearman's $\rho = -0.01$). We note that very high sample-level accuracies are associated with smaller Favi-Scores.

Return to Motivating Example. To return to the motivating example from Section 2, we showcase the Favi-Scores for each system (one-vs-all) according to COMET-22 in Figure 6. That is we computed the Favi-Score for each pair of systems, and display for each system all Favi-Scores in terms of a box-plot. COMET-22 strongly disfavors *ref-B*

(i.e., the human reference), since the Favi-Score is negative for *ref-B*, no matter the system we compare against. This matches the DAG in Figure 1 in which *ref-B* is ranked significantly lower than in the human ranking. The *M2M100* system is also strongly disfavored. However, this is not visible from the DAGs alone, where it is ranked last, both according to COMET-22 and human judgements. This is a case where the Favi-Score uncovers a potentially undesired outcome that could go unnoticed otherwise, highlighting the utility of the Favi-Score as a diagnostic tool. Finally, we observe a strong favoritism with regards to the *QUARTZ* system, which is ranked significantly higher according to COMET-22. Thus, the Favi-Score is also useful as a diagnostic tool and provides an explanation for observed errors in system rankings. See Appendix D for more such examples.

The results of the application of the Favi-Score to evaluate metrics shows that it measures a different aspect of the metrics behaviour than the sign-accuracy scores. Thus, we advocate the usage of the Favi-Score alongside existing scores when a novel metric for generative systems is developed.

8 Conclusion

We motivated and defined the Favi-Score, a simple-to-implement and easy to interpret score, which measures how much an automated preference metric favors the outputs of a generative AI system. Our investigations show that the Favi-Score differs from the sign accuracy scores, and is to be used complementary to them. The empirical application of the Favi-Score to real-world setting reveals that all metrics exhibit some degree of favoritism, which needs to be accounted for when applying these metrics, to better interpret results. We envision that the Favi-Score is used as a diagnostic tool to understand which systems are favored by a metric.

Acknowledgments

This work was supported by the Swiss National Science Foundation within the project "Unified Model for Evaluation of Text Generation Systems (UniVal)" [200020_219819].

Limitations

Relative Measure. Our definition of favoritism is fundamentally relative, meaning we can only say

that some automated preference ratings display favoritism with respect to a specific set of human ratings.

Anchoring. This makes it crucial to prudently select the reference ratings that are used for the comparison. In this work, we treat human ratings as the gold standard reference despite their limitations (Amidei et al., 2018; Belz et al., 2021).

Comparisons. The Favi-Score is defined for a given evaluation setting \mathcal{E} , meaning for a fixed set of automated ratings, human ratings, and system pair. If we replace the the automated ratings, the resulting new Favi-Score is directly comparable. On the other hand, comparing the Favi-Score for one system pair to another, even for the same automated metric, means we have to assume a degree of consistency between reference ratings.

Alternative Definitions. Our definition of favoritism is narrowly defined based on (perceived) errors with respect to human ratings. In particular, we do not consider the content of either the inputs or outputs of various systems under test. Even within our narrow framework, one could imagine alternative definitions taking into account how a change in outcome affects significance tests.

Robustness. Throughout this work, we consider confusion matrices based on reasonably sized evaluation settings. For smaller settings the utility of the Favi-Score can be limited. For example, if there is only 1 error then the only possible values of the Favi-Score are ± 1 and ± 2 . This could potentially be mitigated by a full probabilistic treatment of the evaluation setting, which is outside the scope of this work.

Diagnostic. While we are convinced that the Favi-Score adds a useful diagnostic to a practitioner's toolbox, it does not provide any remedies. For example, to lower the Favi-Score, one can try to distribute errors more evenly, but it is not clear how this should be achieved in a concrete case.

References

- Jacopo Amidei, Paul Piwek, and Alistair Willis. 2018. [Rethinking the agreement in human evaluation tasks](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3318–3329, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Chantal Amrhein and Rico Sennrich. 2022. [Identifying weaknesses in machine translation metrics through minimum Bayes risk decoding: A case study for COMET](#). In *Proceedings of the 2nd Conference*

- of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1125–1141, Online only. Association for Computational Linguistics.
- Anja Belz and Eric Kow. 2010. [Comparing rating scales and preference judgements in language evaluation](#). In *Proceedings of the 6th International Natural Language Generation Conference*. Association for Computational Linguistics.
- Anya Belz, Anastasia Shimorina, Shubham Agarwal, and Ehud Reiter. 2021. [The ReProGen shared task on reproducibility of human evaluations in NLG: Overview and results](#). In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 249–258, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. Evaluation of text generation: A survey. *arXiv preprint arXiv:2006.14799*.
- Arun Chaganty, Stephen Mussmann, and Percy Liang. 2018. [The price of debiasing automatic metrics in natural language evaluation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 643–653, Melbourne, Australia. Association for Computational Linguistics.
- Clint W. Coakley and Mark A. Heise. 1996. [Versions of the sign test in the presence of ties](#). *Biometrics*, 52(4):1242–1251.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, 54(1):755–810.
- Jan Deriu, Don Tuggener, Pius von Däniken, Jon Ander Campos, Alvaro Rodrigo, Thiziri Belkacem, Aitor Soroa, Eneko Agirre, and Mark Cieliebak. 2020. [Spot the bot: A robust and efficient framework for the evaluation of conversational dialogue systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3971–3984, Online. Association for Computational Linguistics.
- Jan Deriu, Don Tuggener, Pius Von Däniken, and Mark Cieliebak. 2022. [Probing the robustness of trained metrics for conversational dialogue systems](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 750–761, Dublin, Ireland. Association for Computational Linguistics.
- Jan Deriu, Pius von Däniken, Don Tuggener, and Mark Cieliebak. 2023. [Correction of errors in preference ratings from automated metrics for text generation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6456–6474, Toronto, Canada. Association for Computational Linguistics.
- Pedro M. Domingos. 2000. A unified bias-variance decomposition and its applications.
- 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.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be guilty but references are not innocent](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 61–71, Online. Association for Computational Linguistics.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. [Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021. [Results of the WMT21 metrics shared task: Evaluating metrics with expert-based human evaluations on TED and news domain](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.
- Yang Gao, Wei Zhao, and Steffen Eger. 2020. [SUPERT: Towards new frontiers in unsupervised evaluation metrics for multi-document summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1347–1354, Online. Association for Computational Linguistics.
- Michael Hanna and Ondřej Bojar. 2021. [A fine-grained analysis of BERTScore](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 507–517, Online. Association for Computational Linguistics.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. [GRADE: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9230–9240, Online. Association for Computational Linguistics.
- 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.
- Klaus Krippendorff. 1970. [Estimating the reliability, systematic error and random error of interval](#)

- data. *Educational and Psychological Measurement*, 30(1):61–70.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). 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 540–551, Hong Kong, China. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A Package for Automatic Evaluation of Summaries](#). In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. [Towards an automatic Turing test: Learning to evaluate dialogue responses](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1116–1126, Vancouver, Canada. Association for Computational Linguistics.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. Association for Computational Linguistics.
- OpenAI. 2023. GPT 3.5 Turbo (0613). <https://platform.openai.com/docs/models/gpt-3-5-turbo>. Accessed: 2024-01-31.
- OpenAI. 2023. OpenAI API. <https://platform.openai.com/docs/guides/function-calling>. Accessed: 2024-01-31.
- Bo Pang, Erik Nijkamp, Wenjuan Han, Linqi Zhou, Yixian Liu, and Kewei Tu. 2020. [Towards holistic and automatic evaluation of open-domain dialogue generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3619–3629, Online. Association for Computational Linguistics.
- 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.
- Maxime Peyrard, Teresa Botschen, and Iryna Gurevych. 2017. [Learning to score system summaries for better content selection evaluation](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 74–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Vitou Phyl, 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, Barcelona, Spain (Online). International Committee on 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. 2022. [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, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Ananya B. Sai, Mithun Das Gupta, Mitesh M. Khapra, and Mukundhan Srinivasan. 2019. [Re-evaluating adam: A deeper look at scoring dialogue responses](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6220–6227.
- Ananya B. Sai, Akash Kumar Mohankumar, Siddhartha Arora, and Mitesh M. Khapra. 2020. [Improving Dialog Evaluation with a Multi-reference Adversarial Dataset and Large Scale Pretraining](#). *Transactions of the Association for Computational Linguistics*, 8:810–827.
- Natalie Schluter. 2017. [The limits of automatic summarisation according to ROUGE](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 41–45, Valencia, Spain. Association for Computational Linguistics.
- Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. [Answers unite! unsupervised metrics for reinforced summarization models](#). 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 3246–3256, Hong Kong, China. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages

7881–7892, Online. Association for Computational Linguistics.

Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L. Hamilton, and Joelle Pineau. 2020. [Learning an unreferenced metric for online dialogue evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441, Online. Association for Computational Linguistics.

Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. [Can you put it all together: Evaluating conversational agents’ ability to blend skills](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030, Online. Association for Computational Linguistics.

Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.

Kosuke Takahashi, Yoichi Ishibashi, Katsuhito Sudoh, and Satoshi Nakamura. 2021. [Multilingual machine translation evaluation metrics fine-tuned on pseudo-negative examples for WMT 2021 metrics task](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1049–1052, Online. Association for Computational Linguistics.

Oleg Vasilyev, Vedant Dharnidharka, and John Bohannon. 2020. [Fill in the BLANC: Human-free quality estimation of document summaries](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 11–20, Online. Association for Computational Linguistics.

Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek Wong, and Lidia Chao. 2022. [UniTE: Unified translation evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8117–8127, Dublin, Ireland. Association for Computational Linguistics.

Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. 2024. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompt elicits reasoning in large language models](#). In

Advances in Neural Information Processing Systems, volume 35, pages 24824–24837. Curran Associates, Inc.

Johnny Tian-Zheng Wei and Robin Jia. 2021. The statistical advantage of automatic nlg metrics at the system level. In *Annual Meeting of the Association for Computational Linguistics*.

Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. [A comprehensive assessment of dialog evaluation metrics](#). In *The First Workshop on Evaluations and Assessments of Neural Conversation Systems*, pages 15–33, Online. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

A Proof of Equal Definition

Here, we prove Lemma 1, which states:

$$\sum_{m,n} W_{mn} \frac{C_{mn}}{E} = \frac{\text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d})}{E} \quad (9)$$

We start the proof by expanding both sides of the equation. We start with the left side. Let for simplicity $r_n^H = \mathbb{I}[r^H[i] = n]$

$$\begin{aligned} \Phi(C) &= \sum_{m,n} W_{mn} \frac{C_{mn}}{E} \\ &= \frac{1}{E} \sum_{m,n} W_{mn} * C_{mn} \\ &= \frac{1}{E} \sum_{i \in \mathcal{T}} \sum_{m,n} W_{mn} \mathbb{I}[r^H[i] = m \wedge r^A[i] = n] \\ &= \frac{1}{E} \sum_{i \in \mathcal{T}} \sum_{m,n} W_{mn} \mathbb{I}[r^H[i] = m] \mathbb{I}[r^A[i] = n] \\ &= \frac{1}{E} \sum_{i \in \mathcal{T}} [-r_+^H r_-^A + -2r_+^H r_-^A - r_-^H r_-^A \\ &\quad + r_-^H r_+^A + 2r_-^H r_+^A + r_-^H r_+^A] \end{aligned} \quad (10)$$

Let us now expand the right side:

$$\begin{aligned}
\Phi(C) &= \frac{\text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d})}{E} \\
&= \frac{1}{E}((\hat{d}_+ - \hat{d}_-) - (d_+ - d_-)) \\
&= \frac{1}{E}(\sum_{i \in \mathcal{T}} r_+^A - \sum_{i \in \mathcal{T}} r_-^A - \sum_{i \in \mathcal{T}} r_+^H + \sum_{i \in \mathcal{T}} r_-^H) \\
&= \frac{1}{E} \sum_{i \in \mathcal{T}} (r_+^A - r_-^A - r_+^H + r_-^H)
\end{aligned} \tag{11}$$

It is now apparent that we need to find an equivalence of the terms inside the sum. That is:

$$\begin{aligned}
&-r_+^H r_-^A + -2r_+^H r_-^A - r_-^H r_-^A \\
&+ r_-^H r_+^A + 2r_-^H r_+^A + r_+^H r_+^A \\
&= r_+^A - r_-^A - r_+^H + r_-^H
\end{aligned} \tag{12}$$

We note that on the left hand side only at most one term of the sum can be non-zero. On the right hand side at most one term for the automated metric can be 1 at the same time, and one for the human rating. We complete the proof by simple enumeration of all possible values of r_n^A and r_m^H .

r_+^H	r_-^H	r_+^H	r_+^A	r_-^A	r_-^A	LHS	RHS
0	0	1	0	0	1	0	0
0	0	1	0	1	0	1	1
0	0	1	1	0	0	2	2
0	1	0	0	0	1	-1	-1
0	1	0	0	1	0	0	0
0	1	0	1	0	0	1	1
1	0	0	0	0	1	-2	-2
1	0	0	0	1	0	-1	-1
1	0	0	1	0	0	0	0

Thus, for all possible combinations of r_n^A and r_m^H , both sides yield the same value, thus, showing the equality of the two definitions.

B Proof of No System Level Sing Change

Here, we proof Lemma 2, i.e., that

$$\begin{aligned}
\Phi(C) &= 0 \\
\implies \mathbb{I}[\text{sgn}(d_+ - d_-), \text{sgn}(\hat{d}_+ - \hat{d}_-)] &= 1
\end{aligned} \tag{13}$$

The proof is a simple application of the alternative definition of the Favi-Score.

$$\begin{aligned}
\Phi(C) &= 0 \\
\iff \frac{\text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d})}{E} &= 0 \\
\iff \text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d}) &= 0 \\
\iff \text{mar}(\hat{\mathbf{d}}) - \text{mar}(\mathbf{d}) &= 0 \\
\iff \text{mar}(\hat{\mathbf{d}}) &= \text{mar}(\mathbf{d}) \\
\iff \hat{d}_+ - \hat{d}_- &= d_+ - d_- \\
\implies \text{sgn}(\hat{d}_+ - \hat{d}_-) &= \text{sgn}(d_+ - d_-)
\end{aligned} \tag{14}$$

Note that the last line of the proof is an implication, which is not reversible. Thus if the sign of both margins are equal, we cannot conclude that there is no favoritism.

C Experimental Setup Details

C.1 Convert Scalar to Preference

Since most metrics return a scalar value, we need to transform them into a preference rating. Similarly, in many cases human ratings are available as scalar or ordinal judgements for a single input and output pair. In these cases we assume that we have paired ratings, meaning that we have system output and rating for each test input and all systems. We can then compare the individual scalar (or ordinal) ratings to derive a preference rating, assuming that higher values indicate higher quality.

Definition 12 (Scalar Metric). *We call real valued functions of inputs and outputs **scalar metrics**: $\mathcal{M}_s : \mathcal{I} \times \mathcal{O} \rightarrow \mathbb{R}$.*

A preference metric can be constructed from a scalar metric as follows:

Definition 13. *The **derived preference metric** \mathcal{M} of a given scalar metric \mathcal{M}_s is defined as*

$$\mathcal{M}(i, o_1, o_2) = \begin{cases} + & \mathcal{M}_s(i, o_1) > \mathcal{M}_s(i, o_2) \\ = & \text{otherwise} \\ - & \mathcal{M}_s(i, o_1) < \mathcal{M}_s(i, o_2) \end{cases}$$

C.2 WMT-22 ChatGPT

Additionally to metrics participating in the WMT-22 metrics shared task (Freitag et al., 2022), we include ratings from GPT3.5-Turbo (0613) (OpenAI, 2023) prompted as a preference metric (Wang et al., 2024). We provided both a static *system prompt* describing the rating task and a *user prompt* for a concrete source sentence and pair of target translations. We used the following system prompt:

You act as an expert translator giving detailed feedback about candidate translations provided by users. Consider in particular the spelling, grammar, accuracy, and fluency of different translations. Make sure you give detailed information why one translation might be better than another.

We used the following user prompt, where the placeholders *source*, *hyp_a*, and *hyp_b* were replaced with a concrete source sentence and translations from systems π_1 and π_2 :

Please give feedback for the following translations:

Original Sentence: {source}

Candidate A: {hyp_a}

Candidate B: {hyp_b}

We used the *function calling* API (OpenAI, 2023) to get a structured output rating and asked the model to provide its preference (*Candidate A*, *Candidate B*, or *No Preference*), specific feedback for each candidate translation, and an explanation for its decision. For our evaluation, we discard the feedback and explanation texts which were included to elicit behavior analogous to *chain of thought* (Wei et al., 2022) prompting. We note that we received preference ratings that were uniquely interpretable as one of the 3 options in almost all cases. For 3 items, the model responded with an empty string as preference rating, which we assigned as *No Preference*.

For our experiments, we assume that the ratings are symmetric with respect to system order, meaning that if we have already collected the ratings for the pair (π_1, π_2) , we invert the ratings (+ to - and vice-versa) for the pair (π_2, π_1) . We enumerated systems in lexicographical order.

C.2.1 Symmetry of Preference Ratings

In the main text we treat the ratings from GPT3.5-Turbo as symmetric. This is mainly to keep the analysis simple. We note that according to Definition 13 the other preference metrics are already symmetric.

We nevertheless also collected ratings for flipped system pairs and noticed discrepancies. We show the confusion matrix in Figure 7. The overall accuracy is 69% and the intra-rater Krippendorff- α (Krippendorff, 1970) is 0.249.




Figure 7: Confusion Matrix of reversed ratings with respect to the original GPT3.5 Turbo ratings.

D Examples of Favoritism Consequences

In this section, we show more examples of DAGs with the system-wise Favi-Score box plots.

D.1 Dialogue

The DAG according to human ratings is depicted in Figure 8a. The human reference is on the top row, while *DialoGPTLarge* is rated as the worst system according to the human ratings. According to the DEB Metric (Sai et al., 2020), human responses are rated much lower than according to the human raters, while the *PolyEncoder* is rated better than the *SeekerDial3B*. This matches the Favi-Scores depicted in Figure 8c, where the human reference is clearly disfavored, and the *PolyEncoder* has an overall positive Favi-Score. The *DialoGPTLarge* is also favored by the DEB metric, however, it does not have an impact on the new ranking. We also note that *PolyEncoder* is rated as being better than *SeekerDial3B* according to the metric, which is also a consequence of favoritism, as *SeekerDial3B* is heavily disfavored. In fact, the Favi-Score between the two systems lies at 0.72 in favor of *PolyEncoder*.

In Figure 9, the analysis for the MAUDE (Sinha et al., 2020) metric is depicted. MAUDE clearly prefers the outputs of *PolyEncoder*, which is also clear from the Favi-Score. Similar to DEB, MAUDE also prefers the outputs of *DialoGPTLarge*, however without changing the ranking. MAUDE has no clear disfavor against the human




Figure 8: Visualization of the consequences of unfair metrics by comparing the human ranking to the DEB ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.

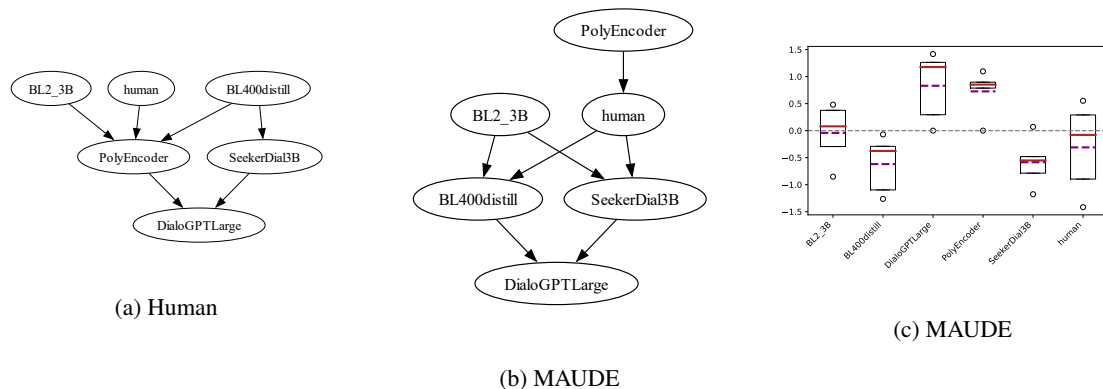


Figure 9: Visualization of the consequences of unfair metrics by comparing the human ranking to the MAUDE ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.

references like the DEB metric. Also *BL2_3B* is favored compared to *BL400distill* with a Favi-Score of 0.375 in favor.

D.2 Summarization

Figure 10 shows the analysis for BertScore (Zhang et al., 2019) on the SummEval data for the Consistency feature. The strongest favoritism is shown towards *M11* (a decoder focused on promoting novelty), which is the worst according to the human ranking, but is in solid mid-field according to the metric. While *M20* (a GPT2 model) is strongly disfavored by the metric. We also note that BertScore has a high system-level sign accuracy (0.789), while having an exceptionally low sample level accuracy (0.05). However, since the

Favi-Score is low as well (mean of 0.29), it shows that the Favi-Score is more influential on the final ranking than the sample level accuracy.

Figure 11 shows the analysis of the ROUGE score. Overall, the ROUGE score has a slightly higher sample-level accuracy (0.07) while having a higher average Favi-Score (0.339), and a much lower system level sign accuracy (0.633). We also note that the favoritism is more pronounced in the ROUGE score, again with *M11* being strongly favored. While *M0* and *M1* are both disfavored, which also is apparent in the final ranking.

D.3 Machine Translation

In Figure 12 the analysis of ChatGPT as metric is shown. ChatGPT has a high system-level accuracy




Figure 10: Visualization of the consequences of unfair metrics by comparing the human ranking to the BertScore ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.

(0.852), and a relatively low Favi-Score (0.229), as well as a low sample level accuracy (0.348). Among the selected WMT-22 metrics, it has one of the lower system-level accuracy scores, the highest sample-level accuracy score, and the highest Favi-Score. Thus, highlighting the influence of favoritism on the final ranking. Similar to other metrics, it disfavors the human reference (*ref-B*). It favors *M2M100*, however without chaining the ranking. These are cases where most mistakes are in favor of *M2M100*, however, the outcome margin is too large for the errors to have an impact on the ranking.

Figure 13 shows the analysis according to the BLEU score. The BLEU score has the lowest system level sign accuracy (0.7), and one of the higher Favi-Scores (0.21), and sample-level sign accuracy comparable to the other metrics (0.31). Again, we observe that the human reference is rated very poorly, with a high Favi-Score against it. Also *Online-W* has a negative Favi-Score, which also reflects its low position in the ranking according to BLEU.

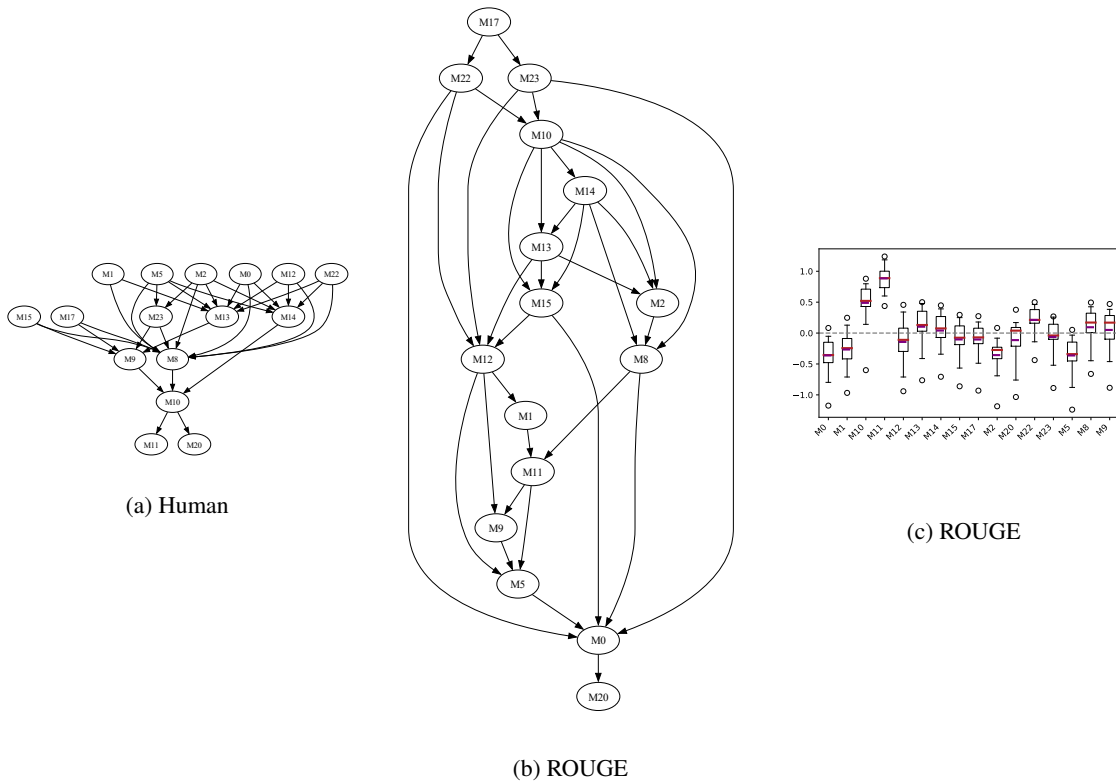


Figure 11: Visualization of the consequences of unfair metrics by comparing the human ranking to the ROUGE ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.

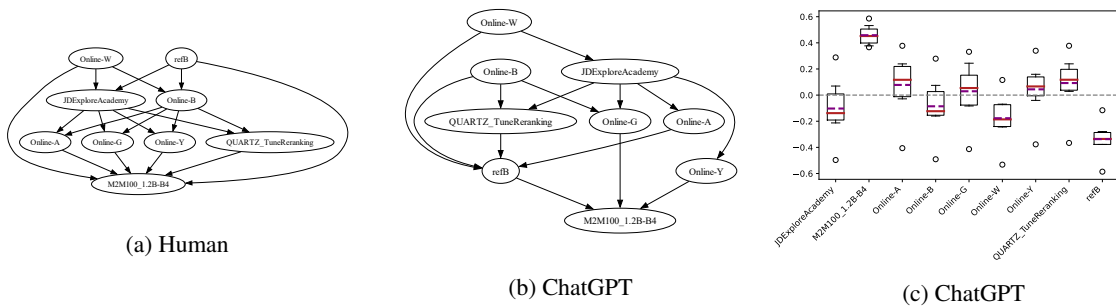





Figure 12: Visualization of the consequences of unfair metrics by comparing the human ranking to the ChatGPT ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.



(a) Human



(b) BLEU



(c) BLEU

Figure 13: Visualization of the consequences of unfair metrics by comparing the human ranking to the BLEU ranking. The ranking of a set of systems is depicted as a Directed Acyclic Graph, where an edge from system A to system B states that system A "wins" against system B. Here, a win is determined by a sign test (Coakley and Heise, 1996) at a 95% confidence threshold.