

Searching for a higher power in the human evaluation of MT

Johnny Tian-Zheng Wei*
University of Southern California
jtwei@usc.edu

Tom Kocmi and Christian Federmann
Microsoft
{tom.kocmi, chrife}@microsoft.com

Abstract

In MT evaluation, pairwise comparisons are conducted to identify the better system. In conducting the comparison, the experimenter must allocate a budget to collect Direct Assessment (DA) judgments. We provide a cost effective way to spend the budget, but show that typical budget sizes often do not allow for solid comparison. Taking the perspective that the basis of solid comparison is in achieving statistical significance, we study the power (rate of achieving significance) on a large collection of pairwise DA comparisons. Due to the nature of statistical estimation, power is low for differentiating less than 1-2 DA points, and to achieve a notable increase in power requires at least 2-3x more samples. Applying variance reduction alone will not yield these gains, so we must face the reality of undetectable differences and spending increases. In this context, we propose interim testing, an “early stopping” collection procedure that yields more power per judgment collected, which adaptively focuses the budget on pairs that are borderline significant. Interim testing can achieve up to a 27% efficiency gain when spending 3x the current budget, or 18% savings at the current evaluation power.

1 Introduction

In machine translation (MT), pairwise evaluations are conducted to identify the better system over a test domain. MT has long taken intrinsic quality as an object of interest, and assumes it can be determined directly from the output (Gatt and Krahnert, 2018). Most practitioners accept that human judgments reflect such quality, and take human evaluation as the gold standard (Bojar et al., 2016). In conducting an evaluation, the experimenter must allocate a budget to collect human judgments, and so evaluation can be an expensive endeavor. No one in the history of MT research has ever been satisfied with the cost or reliability of human evaluation (Graham et al., 2017; Chaganty et al., 2018;

* Work done at Microsoft.




Figure 1: A graphical representation of evaluation with different testing procedures. Currently, our evaluation uses fixed testing, and our current budgets (depicted) often result in underpowered comparison (§5). To get a notable increase in power, we will need to spend more (§6), and interim testing is a way to spend efficiently. Interim testing allows for early stopping by trading off power for additional peeks. In MT, such a tradeoff is a favorable and can yield more power per judgment (§7).

Saldías Fuentes et al., 2022, inter alia). Likewise, we were keen to find savings, upon the foundation of statistically rigorous inference.

Evaluation is a noisy process, and we may not expect a repeat experiment to declare the same winners. For one, we may want a holistic answer of the best system over the entire test domain, but we can only evaluate on a small and finite set of input source sentences (Koehn, 2004; Dror et al., 2018). This introduces a sample bias that our conclusion must be wary of. For another, human judgments on the same output may diverge, so we assume that humans are only a noisy reflection of the true intrinsic quality (Graham et al., 2015). This introduces additional noise when drawing a conclusion from our observations. Intuitively, using a larger test set or averaging over more human judgments should yield more consistency in pairwise comparison.

Inferential statistics is necessary in MT evaluation to declare “winning” MT systems under un-

certainty. Basic usage of statistical testing covers the use case of pairwise MT system comparison (Mathur et al., 2020). After data collection is complete, we can declare significance by computing a p-value (statistical primer in §3). When the p-value is low, a real effect is likely to exist. When the p-value is high, repeat experiments will be inconsistent (effectively tossing a coin), and no good decisions can be made even if you used the gold standard Direct Assessment (DA; Graham et al., 2015) annotation. Significance is the meta-analysis that guards against falsely declaring winners due to noise, with some level of guarantee.

Our work takes the perspective that the basis of solid comparison is in achieving significance. The rate/likelihood an experiment will observe significance is the *power*, and we would like it to be high. At the same time, we would like to minimize human effort and keep costs low. This paper investigates several aspects of the relationship between power and cost in human evaluation:

1. *How can we reason about the power of an evaluation?* We recommend a sensitivity perspective to evaluation, where we characterize an evaluation by its minimum detectable effect (MDE), or the smallest pairwise difference the evaluation will reliably yield significance on. By retrospectively analyzing significance in pairwise comparisons, we can derive an empirical MDE. **Our evaluations can reliably detect up to 2-3 point of DA difference, but comparisons often exhibit even smaller differences.**
2. *How can we notably increase the sensitivity of an evaluation?* With the appropriate power analysis, we can get a rough estimate of the number of samples required to achieve an acceptable sensitivity. To increase the sensitivity to the desired level, we might hope variance reduction techniques can give us the necessary sample efficiency. **If we wanted half of the past comparisons to reliably achieve significance, we needed at least 2x more samples, far beyond the ~1.2x sample efficiency variance reduction offers.**
3. *How can we spend more money efficiently?* If we are not satisfied with the power of our current evaluation, increasing the budget and collecting more judgments is necessary. Crucially, if we accept that small differences can't

be known, our evaluation can be more efficient by focusing the budget elsewhere. **We verify that an “early stopping” procedure (interim testing) can achieve up to a 27% efficiency gain when spending 3x our current budget, or 18% savings at our current evaluation power.**

2 Related work

There is a tradition of using test sets to estimate system performance over the general domain in machine learning (Hastie et al., 2001). There have been calls for statistically rigorous evaluation in natural language processing using significance testing (Dror et al., 2018), however its adoption in reporting has been mixed. For a classic task such as part-of-speech tagging, evaluation is generally significant/consistent even for small gains (Gorman and Bedrick, 2019). In MT, even moderate differences in metric gains (e.g. DA, MQM) may not be consistent, so there is a stronger need for significance testing. Historically, MT evaluation has been heavily based on statistical significance (Koehn, 2004).

MDEs have been used to describe the power of experiments in contexts such as education (which program results in increased test scores?) and sociology (Bloom, 1995). Berg-Kirkpatrick et al. (2012) empirically investigate the conventional wisdom that a certain metric gain corresponds to significance (e.g. 0.5 for BLEU). This threshold is exactly an evaluation’s MDE. They find that a threshold has strong empirical backing, but a few experimental parameters affect this threshold. In our work, we propose taking a sensitivity perspective to evaluation, and reporting the expected MDE of an experiment instead of the other experimental parameters.

Any statistical technique that reduces the cost of human evaluation is, in another view, improving the power offered by some fixed budget. Chaganty et al. (2018) first proposed applying control variates to human evaluation. Control variates increase the sensitivity of an evaluation by leveraging information from a metric. This formulation conveniently allows us to analytically understand its performance based on the experimental conditions. In realistic experimental conditions, they found that the sample efficiency gain is at most 20%, which is in line with results reported in MT (Saldías Fuentes et al., 2022). Mendonça et al. (2021) propose using

online learning to adaptively spend the evaluation budget on determining the best MT systems. However, their technique lacks in statistical rigor for decision making.

Knowing when to “early stop” an evaluation allows us to adaptively spend the budget on difficult pairs and save on easily distinguished pairs. It is known that peeking at the p-value while data collection is ongoing is problematic. Peeking inflates the chance of observing significance and the chance that such significant observation is incorrect (Albers, 2019). While always valid p-values can be calculated that adjust for this error and can be reported at any time, they are mathematically difficult to apply (Johari et al., 2015). Interim testing has been used in medical trials, where experimenters have an ethical consideration in stopping the experiment early (O’Brien and Fleming, 1979). By planning the number of peeks in advance, interim testing can offer rigorous statistical inference while potentially saving time and effort, packaged in an easy to understand technique (Lakens et al., 2021). Our work investigates whether the tradeoff between power and savings is favorable for MT evaluation.

3 A primer on inferential statistics

We consider pairwise comparisons as the basic unit of evaluation echoing calls from Mathur et al. (2020) and Kocmi et al. (2021). Pairwise comparisons are more interpretable than correlations, and more practical for production deployment scenarios. In a pairwise comparison we test the difference between two systems A and B. If you were just to collect a number of DA judgments for each system and declare a winner, a repeat experiment could yield different results due to experimental noise.

A statistical test guards against making an incorrect conclusion due to experimental noise. To do this, we assume a null hypothesis (that A is better than B) and examine how likely we could have made observed our data under this assumption. There are two outcomes of conducting a test:

- (i) there is evidence of a significant difference which rejects the null hypothesis, or
- (ii) the evidence is insufficient and we are unable to reject the null hypothesis.

In the case of (i), a significance test usually guarantees a false detection rate of at most α , where usually $\alpha = 0.05$. Therefore, the best outcome

of statistical testing is the presence of significance, where our inferences enjoy a low false detection rate. The rate at which we can declare significance is called an experiment’s *power* (typically denoted as $1 - \beta$, where β is the false negative rate). In pairwise comparison, our evaluation should have an accuracy $(1 - \alpha)(1 - \beta)$ against the true, pairwise judgment.¹

Intuitively, statistical testing can be loosely thought of as reducing the width of two confidence intervals, spaced by the true system difference of A and B (Krzywinski and Altman, 2013). The power of an experiment is then a function of these three aspects:

- (a) First, the true system difference plays a role in the power. When the distance between the true scores is large relative to the noise, noise is unlikely to obfuscate the true pairwise ranking of the systems.
- (b) Second, the variance of human judgment. The larger the variance in a single judgment, the more judgments that will be needed in an average to get a consistent estimate.
- (c) Finally, the sample size or the budget. The number of judgments you collect shrinks the confidence intervals by a factor of \sqrt{N} from the single judgment variance.

The more judgements you can collect the smaller these confidence intervals will be. When the confidence intervals don’t overlap, the comparison is likely to achieve significance. These three factors all play a role in whether the intervals will be narrow enough.

If we know two of (a), (b), or (c), we can use the appropriate power analysis to deduce the third. Typically, we will observe the (b) human judgment variance, and make a guess at what the true difference (a) would be, to compute what (c) the budget we would have to spend is. When providing estimates for the budget, we would provide estimates under a range of guesses at what the true difference is (Card et al., 2020). Alternatively, we may also ask what the minimum detectable effect is for some fixed budget. Wei and Jia (2021) conducted power analysis in MT and found that small differences

¹This pairwise accuracy holds if you assume that different MT systems always have different quality. By randomizing the systems, the null hypothesis will be true exactly half of the time.




Figure 2: The minimum detectable effect (MDE) is illustrated in the ENU → FRA language pair. Each point represents a pairwise comparison conducted for this language pair. When evaluating pairs exhibiting differences larger than the MDE, 95% of pairs will achieve significance at the $\alpha = .05$ level, which totals to a pairwise accuracy of 90%. Unfortunately, most pairs are on the left hand side of this line. This is also the case for many other language pairs in the ShipData.

require an infeasible amount of budget. This gives a hint that most of our MT evaluation is underpowered. Consistently conducting underpowered experiments run the risk of inflating the error rate in significant observations (Ioannidis, 2005).

4 Dataset

MT evaluation has an established tradition of conducting human evaluation and releasing public datasets. At the time of writing, the current annotation method of choice in MT is *Direct Assessment* (DA; Graham et al., 2015; Akhbardeh et al., 2021). Direct Assessment asks annotators to rate a translation’s quality on a sliding point scale from 0-100. We study the **ShipData** presented in Kocmi et al. (2021), which is the largest human evaluation dataset of pairwise comparisons, accumulated over two years from internal evaluation campaigns at Microsoft Translator. No text is contained i.e. source, references, or outputs, but the raw DA scores are sufficient for our purposes. We focus on this dataset because it is large and often contain comparisons between state-of-the-art systems. It contains 4004 pairwise comparisons between two systems, where each system pair contains about 600 human judgments per system (1200 for both systems).

5 The sensitivity approach to evaluation

The basis of solid comparison is significance. Therefore, we need a way to reason about the power of an experiment. In this section, we recommend

	Significant / insignificant	Obs. MDE	Median difference
ENU → FRA	30 / 153	3.8	1.2
ENU → DEU	19 / 151	3.5	0.7
FRA → ENU	3 / 140	2.4	0.6
DEU → ENU	27 / 130	1.9	0.6
JPN → ENU	78 / 127	2.9	3.2
ENU → JPN	40 / 94	3.8	1.8
ITA → ENU	2 / 81	2.8	0.5
CHS → ENU	30 / 78	2.6	1.5
ENU → PTB	28 / 74	1.0	0.6
ENU → SVE	31 / 73	4.4	1.4

Table 1: Significance and MDE results in the top-10 language pairs (by number of comparisons). Significance is calculated at the $\alpha = 0.05$ level. Observed MDEs are calculated for 90% pairwise accuracy. The median system difference is observed from the data. For most language pairs, less than half of the pairs had a significant observation. MDEs are small but most of the system differences appear to be even smaller.

a sensitivity approach to evaluation, and retrospectively deduce the power of previous evaluations. By looking at the observed effect sizes we can also set a meaningful target power.

5.1 Minimum detectable effects (MDEs)

The pairwise evaluation of two MT systems is not a one-size fits all procedure, even though the MT literature uses a consistent annotation method (Feldermann, 2018). Rather, an evaluation is our best attempt to answer which MT system is better with the evaluation annotation budget at hand. How much budget to allocate should depend on the circumstantial factors. Statistical inference can give us a probabilistic answer to this question with whatever evidence we are able to collect.

In the best case scenario, a significant result is observed and a winner is declared after the data is collected. However, significance depends on the conditions of the experiment (see §3), where the size of the pairwise difference, annotation variance, and number of samples all play a role. The pairwise difference and annotation variance are determined by the annotation method. Since most prefer to use a widely accepted annotation such as Direct Assessment (DA; Graham et al., 2015), these are factors we may not be able to change. However, we can increase the budget, and the larger the budget, the more likely we will be able to achieve significance for some fixed difference.




Figure 3: Power analysis for the total number of judgments required to achieve an MDE with 90% inference accuracy. These figures are calculated through simulation with distributional assumptions on the human scoring function (see §6.1). Compared to the observed MDEs, figures here serve as a lower bound. As the differences decrease linearly, the number of samples required increases exponentially.

We recommend to think about an MT evaluation in terms of its sensitivity. With a fixed budget and annotation method, there is some deducible minimum detectable effect (MDE; Bloom, 1995), where evaluating differences larger than the MDE will enjoy a comfortable level of power (rate of significance). Alternatively, if we did not observe significance for some experiment, we may suspect that the true difference is likely to be lower than the experiment’s MDE. With a sensitivity perspective, our consideration is now to conduct DA evaluations with a budget large enough to exhibit an appropriate MDE. Ideally, our evaluation exhibits an MDE small enough where we believe any smaller differences are not practically meaningful (more in §6.1). Realistically, we would set up an evaluation with MDEs as small as our budgets allow.

5.2 Observed MDEs

In this section, we attempt to retrospectively understand the MDEs/sensitivity of our past evaluations. Refer to Figure 2 for graphical intuition. We can empirically estimate (as opposed to making assumptions and simulating, see Card et al., 2020) an (observed) minimum detectable effect by sorting all the pairs by their observed absolute system difference, and choosing the difference where comparisons with a larger system difference (effect size) will have at least 95% of experiments showing significance (corresponding to experimental power $1 - \beta = 0.95$) at a level of $\alpha = 0.05$ by the Mann

		Variance (std. dev.)	Reducible variance
WMT21	*-en	866.2 (29.4)	23.1%
pSQM	zh-en	683.2 (26.1)	9.8%
pSQM	en-de	705.4 (26.5)	53.4%

Table 2: Total annotation variance and the reducible proportion of that variance. pSQM scores are provided by Freitag et al. (2021) and are collected from professional annotators. WMT21 scores are provided by Akhbardeh et al. (2021) and are collected from crowdworkers. pSQM scores are normalized from 0-100 for ease of interpretation. At least half of the variance is irreducible.

ρ	WMT21	pSQM(zh-en)	pSQM(en-de)
1.0	1.30	1.20	4.33
0.5	1.06	1.12	3.09
0.2	1.01	1.11	2.94

Table 3: Data efficiencies for the control variates estimator under different conditions. Each column represents a different condition of reducible variance, instantiated from observed statistics from Table 2. ρ is the correlation of the metric that would be used in the control variates estimator. With the exception in pSQM en-de, variance reduction is far from giving us the 2x-10x multiplier we need.

Whitney U (MWU) test. This ensures that at least $(1 - \alpha)(1 - \beta) \approx 0.9$ of the pairs should be accurate (Wei and Jia, 2021). We can interpret this as the threshold at which our experiments will stop being accurate at the 90% level.

The minimum detectable effects (MDE) are small, but differences between systems are even smaller. Refer to Table 1. Our evaluations have been able to detect up 1 or 2 points of system-level DA difference, but often a third of the comparisons are still not significant. Looking at the density of the differences (see the x-axis in Figure 2) we see that most of the pairs exhibit small differences. An immediate consequence is that most of the budget is being spent to declare ties. Most of our comparisons are underpowered, and where the p-value is high the experiments are not much better than a coin toss. The median difference provides a target MDE if we want half of our evaluations to show significance (alternatively, declaring ties in half of the evaluations is acceptable).

6 Known unknowns

Now that we have established a way to reason about experimental power, we conduct power analysis to understand how much more gain we need to improve our power to a desired sensitivity. We investigate whether variance reduction techniques are sufficient, and conclude that the only way forward is to increase the annotation budget.

6.1 Power analysis for the desired sensitivity

As suggested in Card et al. (2020), we can roughly determine the number of samples for a fixed power using simulation. As with any power analysis, we must make some assumptions to estimate the number of samples needed. Here we assume that the judgments for a given system’s translation is distributed as $s \sim 100 - \text{Gamma}(k, \theta)$ where $k = \frac{\mu^2}{\sigma^2}$ and $\theta = \frac{\sigma^2}{\mu}$ are fit to match the average mean and variance of a system for that language pair. We choose the use of the Gamma distribution because the resulting scoring distribution is such that most of the scores are high, and the more severe the translation error the more rare it is, which matches what we observe in Kocmi et al. (2021). We then use the bisection method to determine the integer whose power has the closest match to our desired β value. We find that the simulation reasonably matches empirically observed MDEs.

Power analysis shows that most pairs needs not a little, but a lot more judgments. Refer to Figure 3. Comparing to the observed MDEs, the power analysis is *optimistic*, where the figures we provide can be seen as a lower bound. Even a reduction of our MDE to 1 point can require up to 2x times more judgments (than originally used in the ShipData). We highlight the fact that as differences get linearly smaller, the number of samples is an exponential growth. The nature of statistical estimation is that smaller differences are increasingly elusive.

In the search for higher power, we must also keep in mind that arbitrarily small differences require arbitrarily large budgets. Therefore, for modern state-of-the-art comparisons, some differences will be left unknown. We can not fantasize about detecting every single small difference out there just by spending more budget or applying some strong statistical technique (see §6.2). Perhaps this may be taken in stride, as mathematicians learned to accept the existence of unprovable theorems nearly a century ago (Gödel, 1934). Many other important

fields such as domain adaption also grapple with their unknowns (Ben-David et al., 2010).

6.2 Variance reduction is inadequate

Generally, we assume that a human evaluator scores a segment with the true segment level quality score, plus some noise. If $H(x)$ is the human scoring function on system translations x , there are 2 parts to the scoring variance. We can decompose the variance of H to

$$\begin{aligned} \text{Var}(H(x)) &= \mathbb{E}[\text{Var}(H(x)|x)] \\ &+ \text{Var}(\mathbb{E}[H(x)|x]) \end{aligned} \quad (1)$$

by the law of total variance. The first part is the variance of the true translation quality scores, capturing the real difference in quality across output translations, and the second part is the rest of the variance. The second term, which we broadly term annotator noise, can include annotator biases, preferences, and even mood.

Using repeat judgments we can estimate the second term (annotator noise), which is similar to an inter-annotator agreement (Wei and Jia, 2021). Since the ShipData doesn’t contain any repeat judgments, we provide estimate of the second term from a few similar datasets (Akhbardeh et al., 2021; Freitag et al., 2021). Refer to Table 2. In designing variance reduction techniques, we usually leverage metric scores to reduce the first term, but not annotator information to reduce the annotator noise (second term), as it is too difficult (Saldías Fuentes et al., 2022).

With variance reduction (VR) techniques, we can achieve a higher power with the current budget by leveraging side information (Owen, 2013). However, VR is not arbitrarily powerful, and its effectiveness is constrained by the amount of reducible variance present, and how much of the reducible variance you can actually reduce. Here, we look at the control variates technique² which leverages the linear information in a metric for the estimation of system quality. The data efficiency in Chaganty et al. (2018) describes how many times a control variates estimator improves over the regular sample mean estimate, and is characterized by

$$\text{DE} := \frac{\text{Var}(\hat{\mu}_{\text{mean}})}{\text{Var}(\hat{\mu}_{\text{cv}})} = \frac{1 + \gamma}{1 - \rho^2 + \gamma} \quad (2)$$

²Equal proportion stratified sampling is a special case of control variates, so these results also apply (Owen, 2013). Any technique which uses a metric to bin outputs, where the same number of outputs are sampled for scoring within each bin, are constrained by these results as well.




Figure 4: The average power of each pairwise comparison for fixed testing at 1200 against interim-futility testing at 2300. Each point represents a pairwise comparison. When planning for 2300 judgments with interim-futility, the actual amount of judgments collected in our simulation is about 1200. For the same budget, we see that interim-futility testing boosts the power of moderate to high-powered pairs, but drops that of the lower powered pairs.

where ρ is the sentence-level Pearson correlation of the metric and

$$\gamma = \frac{\sigma_a^2}{\sigma_f^2} = \frac{\mathbb{E}[\text{Var}(H(x)|x)]}{\text{Var}(\mathbb{E}[H(x)|x])} \quad (3)$$

Refer to Table 3. **With the optimistic assumption of a perfect metric, we often only get a ~1.2x efficiency gain from VR, far from the 2-10x multiplier we need to obtain significant comparisons.** The gains we predict for VR is consistent with the practical results presented in Saldias et al. (2022). These reduction techniques work, but is far from achieving what we need, echoing the narrative of Chaganty et al. (2018).

7 Spending effectively

To have a notable gain in sensitivity, variance reduction alone is inadequate. Therefore, spending is necessary in the search for higher power. This section describes a simple yet statistically rigorous way of “early stopping” in a human evaluation campaign. Interim testing adaptively allocates the budget to borderline significant pairs, and can be seen as an efficient way to spend.

7.1 Peek-a-boo! Planning interim peeks

Savings can be achieved if we can stop data collection as soon as a result can be concluded. If the experimenter runs the preferred statistical test (at false detection rate $\alpha = 0.05$) periodically while




Figure 5: The average number of judgments collected by each sampling method. For interim and interim-futility, 1200 judgments were planned, and the actual judgments collected are strictly less. As the system differences grow larger, both methods have the potential to stop early. For interim-futility, pairs with small differences also incurred less judgments.

data collection is on-going, the final process will have a false detection far higher than the α intended (Albers, 2019). There are a class of sequential sampling techniques, which allow you to test after every single sample while maintaining the false detection rate constant, but are mathematically difficult to apply (Johari et al., 2015).

A simpler solution is to use interim sampling and apply a correction for multiple testing (Lakens et al., 2021). For instance, the Pocock correction (Pocock et al., 1987) is appropriate when multiple comparisons are made, but we want a false detection to be maintained at a desired α .³ Refer to Figure 1. For interim testing, we can plan in advance to collect batches of data, and test between each batch. To maintain a final false detection rate to the fixed procedure, your interim tests must have an α_0 appropriately adjusted with the Pocock correction. The downside is that this correction is conservative, and each test has less power.

At each interim point, we can also stop for futility, or when we see that even in completion of the data collection, we are unable to achieve significance. Practically, there are many ways to set up this stopping rule (Lakens et al., 2021), but in our simulation we find that a simple heuristic (checking if the $p > 0.5$) works well for our purposes. An alternative view of futility stopping is that we are unwilling to conduct the analysis of the original

³Here’s why we need a correction: imagine 20 comparisons made at $\alpha = 0.05$ where the null hypothesis is true, then the probability of getting at least 1 significant result is actually $1 - 0.95^{20} \approx 0.63$.

experiment with the corresponding MDE.

7.2 Experimental setup

We compare three different kinds of testing methods. Refer to Figure 1.

- **Fixed** testing is most commonly used in evaluation. In fixed testing, the annotation budget is spent all at once, and the statistical test is performed at the end. The advantage of fixed testing is that only statistical test is performed with the highest (least conservative) alpha threshold (e.g. is $p < 0.05$?).
- **Interim** testing plans to spend the budget in equal sized steps, with an interim analysis between each step. If significance is observed at any point, the data collection is terminated. We always plan for 3 peeks, and use the Pocock correction (e.g. is $p < 0.0221$? at each peek). While the testing threshold is lower (more conservative), the savings obtained from some pairs can be used on others, by planning more judgments for all pairs.
- **Interim-futility** is the same as interim testing but also applies a futility stopping rule at each analysis. If $p > 0.5$ then the experiment is terminated early. Futility stopping does not affect the false detection rate so it does not need to be adjusted. Futility stopping results in strictly less power, but the savings can be used elsewhere, by planning more judgments.

To benchmark these testing procedures against each other, we simulate data collection from the pairs in the ShipData by sampling with replacement. For each pair we simulate each testing procedure 1000 times and record the number of times the procedure is able to achieve significance. For all tests we use the Mann Whitney U test (standard to machine translation; Akhbardeh et al., 2021) with a testing threshold of $\alpha = 0.05$. Within the ShipData, each pair only has about 1200 judgments, from which we often oversample. We note that this is our best faith attempt to study these testing methods in the large budget regime, and actual benchmarking would require infeasible cost, so the simulation can serve as our best synthetic testbed.

7.3 Results

Refer to Figure 5. For a fixed sampling procedure, the number of samples collected is constant for

every effect size. This can be inefficient as pairs with large differences do not need as many judgments to declare significance. Interim testing is adaptive; as the differences get larger, interim testing can declare significance at an early step. For interim-futility, less judgments are also collected for the pairs with the smallest differences, where early steps may declare futility. We will later see that the interim-futility behavior is most favorable.

Refer to Figure 4. When comparing fixed and interim-futility, we compare two procedures that spend the same budget. Since interim sampling spends more on borderline pairs, the power for pairs with moderate to high differences increases. Savings are made on pairs with both large and small differences, with small difference pairs having a decrease in power. We highlight that interim-futility is a different kind of testing. **While the use of fixed testing seeks to best detect every difference no matter how small, the use of interim-futility prioritizes the pairs that have borderline significant differences.**

Refer to Figure 6(a). The main metric we benchmark these methods is by the average power, or the number of significant comparisons over all the ShipData. When comparing over all pairs, interim testing has slightly better performance, but interim-futility gives considerable gains even at current budget sizes. Our results show that to attain the fixed testing power at 1200, interim testing only needed to spend 990 judgments per comparison, which is an 18% saving⁴. The advantage of interim sampling over fixed sampling is even more pronounced when we are spending large budget sizes, where we can gain 28% savings at 3600 judgments (3x). Refer to Figure 6(b). When testing small differences interim sampling is worse than fixed sampling, as it has a stricter significance threshold. However, interim-futility is able to stop on pairs with little hope and prioritize the borderline significant pairs. Refer to Figure 6(c). On pairs with large differences interim sampling is best, with interim-futility achieving similar performance. For large differences futility stopping should rarely trigger, so the two methods should be similar.

We want to highlight that the distribution of the differences is key to the success of the interim-futility testing procedure. Since most of the pairs are concentrated either in the dense

⁴All the results in this paragraph are derived using linear interpolation, akin to using a ruler on Figure 6.

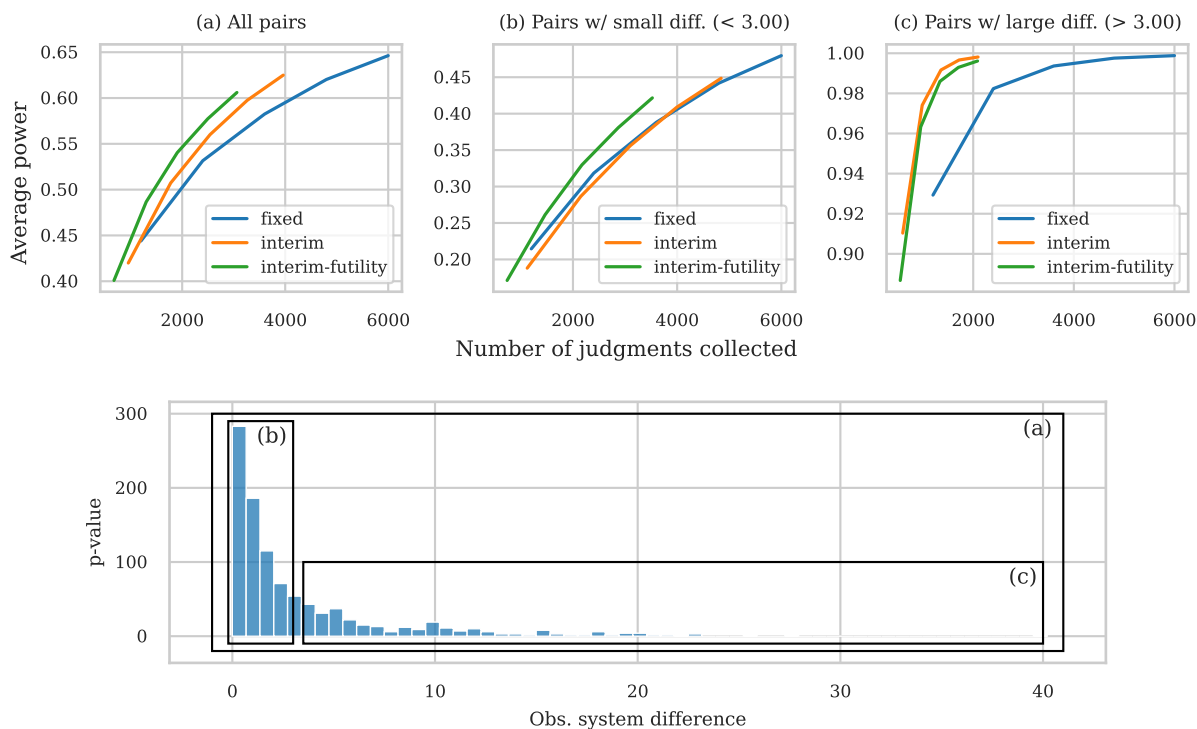


Figure 6: (Top) The average power of each testing procedure across the ShipData for different sized budgets. (a) Shows the average power across all data, and (b) shows it over pairs with large differences and (c) shows it for small differences. (Bottom) The histogram of the true differences in each pairwise comparison. These are true differences due to the simulation we used to test these procedures. Interim-futility is most favorable by average power in (a), (b) and (c). Interim testing is weaker in (c) due to its stricter significance threshold.

region of small differences or in the long tail of large differences, these are areas where interim-futility can early stop. Compared to fixed testing, interim-futility will be able to make savings here to spend elsewhere. Crucially, the application of futility stopping also requires a change in our evaluation mindset, as we must be willing to accept that some small differences are not worth detecting. If we can make this change, then interim-futility is most favorable in terms of average power.

8 Limitations

The most important assumption of our work is in the use of Direct Assessment (DA). While our methods can generalize to any real valued judgment, we analyzed DA because of its widely recognized, gold standard status in MT evaluation. DA is a particularly noisy judgment, and so the power and variance reduction results are pessimistic. However, we believe that the study of annotation will be the most important direction in MT evaluation.

Let’s take Hassan et al. (2018), where one of the first claims of MT-human parity was made. By their evaluation, which was conducted according to the

community standard, no significant difference was found between human and machine translations with a reasonable budget, and so a tie was declared. Toral et al. (2018) reassesses this claim, and essentially presents a series of alternative evaluations and observe significant differences that contradict with Hassan et al. (2018). This is just one of many studies which compels an alternative evaluation with qualitative insight (Läubli et al., 2018, 2020; Freitag et al., 2021).

Our perspective is that significance is only one pillar of MT evaluation. It is our hope that the analyses in this work will further our understanding of significance and evaluation power. However, the second pillar of MT evaluation is in the annotation method. While power is quantitative, the study of annotation methods will be qualitative. Going forward, understanding how we can change the annotation method to increase the power will be crucial. We will need good qualitative understanding to be able to move away from DA and establish new gold standards.

In addition, we showed that interim testing is only effective for pairwise comparison. Future

work should look to make savings in the leaderboard styled evaluation of WMT. This may come in the form of generalizing interim sampling for multiple comparisons or formalizing the bandit results from [Mendonça et al. \(2021\)](#) in terms of statistical inference.

9 Conclusion

Our work is motivated by the cost of human evaluation in machine translation. Before searching for a higher power from our current budget, we determined how much more power was necessary. In doing so, we recommend taking a sensitivity approach to evaluation. From here we came to the conclusion that to achieve the power/sensitivity necessary, variance reduction alone would be insufficient, and spending is our only option. If we decide to allocate larger budgets, interim testing is a more effective way to spend, which can yield 18% savings at the current evaluation power, or 27% savings at 3x the original budget.

References

- Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydrin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(wmt21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, Online. Association for Computational Linguistics.
- Casper Albers. 2019. The problem with unadjusted multiple and sequential statistical testing. *Nature Communications*, 10(1):1–4.
- Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. 2010. [Impossibility theorems for domain adaptation](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 129–136. JMLR.org.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. [An empirical investigation of statistical significance in NLP](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005, Jeju Island, Korea. Association for Computational Linguistics.
- Howard S. Bloom. 1995. [Minimum detectable effects: A simple way to report the statistical power of experimental designs](#). *Evaluation Review*, 19(5):547–556.
- Ondřej Bojar, Christian Federmann, Barry Haddow, Philipp Koehn, Matt Post, and Lucia Specia. 2016. Ten years of wmt evaluation campaigns: Lessons learnt. In *Proceedings of the LREC 2016 Workshop “Translation Evaluation—From Fragmented Tools and Data Sets to an Integrated Ecosystem*, pages 27–34.
- Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. [With little power comes great responsibility](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9263–9274, Online. Association for Computational Linguistics.
- Arun Tejasvi 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 643–653. Association for Computational Linguistics.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. [The hitchhiker’s guide to testing statistical significance in natural language processing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Christian Federmann. 2018. [Appraise evaluation framework for machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 86–88, Santa Fe, New Mexico. Association for Computational Linguistics.
- Markus Freitag, George F. Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *CoRR*, abs/2104.14478.
- Albert Gatt and Emiel Krahmer. 2018. [Survey of the state of the art in natural language generation: Core tasks, applications and evaluation](#). *J. Artif. Intell. Res.*, 61:65–170.
- Kyle Gorman and Steven Bedrick. 2019. [We need to talk about standard splits](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.

- Yvette Graham, Timothy Baldwin, and Nitika Mathur. 2015. [Accurate evaluation of segment-level machine translation metrics](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1183–1191. The Association for Computational Linguistics.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2017. [Can machine translation systems be evaluated by the crowd alone](#). *Natural Language Engineering*, 23(1):3–30.
- Kurt Gödel. 1934. On undecidable propositions of formal mathematical systems.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. 2018. [Achieving human parity on automatic chinese to english news translation](#). *CoRR*, abs/1803.05567.
- Trevor Hastie, Jerome H. Friedman, and Robert Tibshirani. 2001. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer.
- John P. A. Ioannidis. 2005. [Why most published research findings are false](#). *PLOS Medicine*, 2(8):null.
- Ramesh Johari, Leo Pekelis, and David J. Walsh. 2015. [Always valid inference: Bringing sequential analysis to a/b testing](#).
- 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.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Martin Krzywinski and Naomi Altman. 2013. [Error bars](#). *Nature Methods*, 10(10):921–922.
- Daniel Lakens, Friedrich Pahlke, and Gernot Wassmer. 2021. [Group sequential designs: A tutorial](#).
- Samuel Lübbli, Sheila Castilho, Graham Neubig, Rico Sennrich, Qinlan Shen, and Antonio Toral. 2020. [A set of recommendations for assessing human-machine parity in language translation](#). *J. Artif. Intell. Res.*, 67:653–672.
- Samuel Lübbli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? A case for document-level evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4791–4796. Association for Computational Linguistics.
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020. [Tangled up in BLEU: Reevaluating the evaluation of automatic machine translation evaluation metrics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4984–4997, Online. Association for Computational Linguistics.
- Vânia Mendonça, Ricardo Rei, Luisa Coheur, Alberto Sardinha, and Ana Lúcia Santos. 2021. [Online Learning meets Machine Translation evaluation: Finding the best systems with the least human effort](#). 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 3105–3117, Online. Association for Computational Linguistics.
- Peter C. O’Brien and Thomas R. Fleming. 1979. [A multiple testing procedure for clinical trials](#). *Biometrics*, 35(3):549–556.
- Art B. Owen. 2013. *Monte Carlo theory, methods and examples*.
- Stuart J Pocock, Nancy L Geller, and Anastasios A Tsiatis. 1987. The analysis of multiple endpoints in clinical trials. *Biometrics*, pages 487–498.
- Belén Saldías, George F. Foster, Markus Freitag, and Qijun Tan. 2022. [Toward more effective human evaluation for machine translation](#). *CoRR*, abs/2204.05307.
- Belén Saldías Fuentes, George Foster, Markus Freitag, and Qijun Tan. 2022. [Toward more effective human evaluation for machine translation](#). In *Proceedings of the 2nd Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 76–89, Dublin, Ireland. Association for Computational Linguistics.
- Antonio Toral, Sheila Castilho, Ke Hu, and Andy Way. 2018. [Attaining the unattainable? reassessing claims of human parity in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.
- Johnny Wei and Robin Jia. 2021. [The statistical advantage of automatic NLG metrics at the system level](#). 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 6840–6854, Online. Association for Computational Linguistics.