

Human Judgement as a Compass to Navigate Automatic Metrics for Formality Transfer

Huiyuan Lai, Jiali Mao, Antonio Toral, Malvina Nissim

CLCG, University of Groningen / The Netherlands

{h.lai, jiali.mao, a.toral.ruiz, m.nissim}@rug.nl

Abstract

Although text style transfer has witnessed rapid development in recent years, there is as yet no established standard for evaluation, which is performed using several automatic metrics, lacking the possibility of always resorting to human judgement. We focus on the task of formality transfer, and on the three aspects that are usually evaluated: style strength, content preservation, and fluency. To cast light on how such aspects are assessed by common and new metrics, we run a human-based evaluation and perform a rich correlation analysis. We are then able to offer some recommendations on the use of such metrics in formality transfer, also with an eye to their generalisability (or not) to related tasks.¹

1 Introduction

Text style transfer (TST) is the task of automatically changing the style of a given text while preserving its style-independent content, or theme. Quite different tasks, and thus quite different types of transformations, traditionally fall under the TST label. For example, given the sentence “*i like this screen, it’s just the right size...*”, we may produce its negative counterpart “*i hate this screen, it is not the right size*” for the task defined as *polarity swap* (Shen et al., 2017; Li et al., 2018a), or turn it into the formal “*I like this screen, it is just the right size.*” for the task called *formality transfer* (Rao and Tetreault, 2018).

For the transfer to be considered successful, the output must be written (i) in the appropriate target style; (ii) in a way such that the original content, or theme, is preserved; and (iii) in proper language, hence fluent and grammatical (relative to the desired style). These aspects to be evaluated are usually defined as (i) *style strength*, (ii) *content preservation*, and (iii) *fluency*, and automatic

¹Our analysis code, literature list for Figure 1, and all data are available at <https://github.com/laihuiyuan/eval-formality-transfer>.

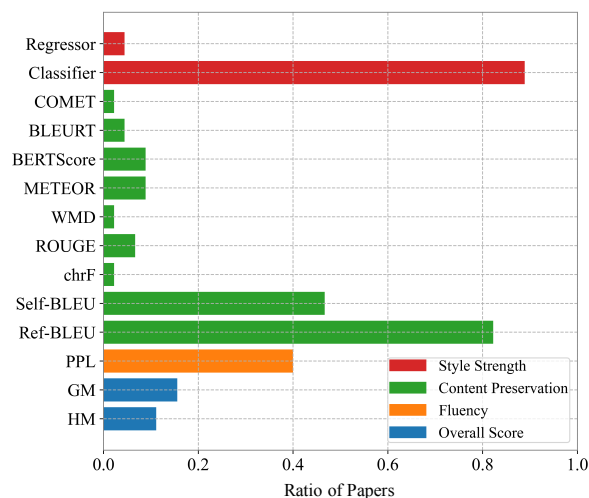


Figure 1: Automatic evaluation metrics in 45 ACL Anthology papers focusing on style transfer and its evaluation in terms of (i) style strength: regressor and classifier; (ii) content preservation: COMET, BLEURT, BERTScore, METEOR, WMD, ROUGE, chrF, Self-BLEU (source-based BLEU) and Ref-BLEU (reference-based BLEU); (iii) fluency: PPL (perplexity); and (iv) overall score: HM (harmonic mean) and GM (geometric mean).

evaluation metrics are used accordingly, lacking the possibility of using human judgement for any given experiment. Figure 1 shows a survey of such metrics (organised by aspect) as used in 45 papers published over the last three years in the ACL Anthology, which focus on TST in general. A classifier or a regressor is used to assess style strength, a variety of content-based metrics target content preservation, perplexity is used to measure fluency, and some overall metrics combining content and style are often reported.

In spite of the attempts to perform careful automatic evaluation, and of some works studying specific aspects of it, such as traditional metrics for polarity swap (Tikhonov et al., 2019; Mir et al., 2019), content preservation for formality transfer (Yamshchikov et al., 2021), and a recent attempt

at correlating automatic metrics and human judgment for some aspects of multilingual formality transfer (Briakou et al., 2021a), the community has not yet reached fully shared standards in evaluation practices. We believe this is due to a concurrence of factors.

First, different tasks are conflated under the TST label while they are not exactly the same, and evaluation is a serious issue. Lai et al. (2021a) have shown that polarity swap and formality transfer cannot be considered alike especially in terms of content preservation, as in the former the meaning of the output is expected to be the opposite of the input rather than approximately the same. Hence, it is difficult to imagine that the same metric would capture well the content aspect in both tasks.

Second, the evaluation setting is not necessarily straightforward: if the content of the input has to be preserved in the output, the quality of the generated text can be assessed either against the input itself or against a human-produced reference, specifically crafted for evaluation. However, not all metrics are equally suitable for both assessments. For instance, BLEU (Papineni et al., 2002) is the metric most commonly used for evaluating content preservation (Fig. 1). Intuitively, this n -gram based metric should be appropriate for comparing the output and the human reference, but is much less suitable for comparing the model output and the source sentence, since the whole task is indeed concerned with changing the surface realisation towards a more appropriate target style. On the contrary, neural network-based metrics should also work between the model output and the source sentence. This leads to asking what the best way is to use and possibly combine these metrics under which settings. Closely related to this point, it is not fully clear what the used metrics actually measure and what desirable scores are. For example, comparing source and reference for metrics that measure content similarity should yield high scores, but we will see in our experiments that this is not the case. Recent research has only compared using the reference and the source sentence for one metric: BLEU (Briakou et al., 2021a), and introduced some embeddings-based metrics only to compare the output to the source. A comprehensive picture of a large set of metrics in the two different evaluation conditions (output to source and output to reference) is still missing and provided in this contribution.

Lastly, and related to the previous point, it is yet unclear whether and how the used metrics correlate to human judgements under different conditions (e.g. not only the given source/reference used for evaluation but also different transfer directions, as previous work has assessed human judgement over the informal to formal direction (Briakou et al., 2021a) only), and how they differ from one another. This does not only affect content preservation, as discussed above, but also style strength and fluency.

Focusing on formality transfer, where the aspect of content preservation is clear, we specifically pose the following research questions:

- **RQ1** What is the difference in using a classifier or a regressor to assess style strength and how do they correlate with human judgement?
- **RQ2** How do different content preservation metrics fare in comparison to human judgement, and how do they behave when used to compare TST outputs to source or reference sentences?
- **RQ3** Is fluency well captured by perplexity, and what if the target style is informal?

To address these questions we conduct a human evaluation for a set of system outputs, collecting judgments over the three evaluation aspects, and unpack each of them by means a thorough correlation analysis with automatic metrics.

Contributions Focusing on formality transfer, we offer a comprehensive analysis of this task and the nature of each aspect of its evaluation. Thanks to the analysis of correlations with human judgements, we uncover which automatic metrics are more reliable for evaluating TST systems and which metrics might not be suitable for this task under specific conditions. Since it is not feasible to always have access to human evaluation, having a clearer picture of which metrics better correlate with human evaluation is an important step towards a better systematisation of the task’s evaluation.

2 Related Work

Text Style Transfer In the recent tradition of TST, many related tasks have been proposed by researchers. Xu et al. (2012) employ machine translation techniques to transform modern English into Shakespearean English. Sennrich et al. (2016) propose a task that aims to control the level of politeness via side constraints at test time. Polarity swap (Shen et al., 2017; Li et al., 2018b) is a task

of transforming sentences, swapping their polarity while preserving their theme. Political slant is the task that preserves the intent of the commenter but modifies their observable political affiliation (Prabhumoye et al., 2018). Formality transfer is the task of reformulating an informal sentence into formal (or viceversa) (Rao and Tetreault, 2018; Briakou et al., 2021b). Cao et al. (2020) propose an expertise style transfer that aims to simplify the professional language in medicine to the level of laypeople descriptions using simple words. Jin et al. (2021) provide an overview for different TST tasks.

Automatic Evaluation In Figure 1 we see that more than 80% of papers employ a style classifier to assess the attributes of transferred text for the aspect of style strength. For content preservation, BLEU is by far the most popular automatic metric, but recent work has also employed other metrics, including string-based (e.g. METEOR (Mir et al., 2019; Lyu et al., 2021; Briakou et al., 2021a)) and neural-based (e.g. BERTScore (Reid and Zhong, 2021; Lee et al., 2021; Briakou et al., 2021a)). In order to further increase the capturing of semantic information beyond the lexical level, Lai et al. (2021b,a) recently also employed BLEURT (Selam et al., 2020) and COMET (Rei et al., 2020) to evaluate their systems. These *learnable metrics* attempt to directly optimize the correlation with human judgments, and have shown promising results in machine translation evaluation. For fluency, a language model (LM) trained on the training data is used to calculate the perplexity of the transferred text (John et al., 2019; Sudhakar et al., 2019; Huang et al., 2020). Geometric mean and harmonic mean of style accuracy and BLEU are often used for overall performance (Xu et al., 2018; Luo et al., 2019; Krishna et al., 2020; Lai et al., 2021a,b).

Evaluation Practices Although some previous work has run correlations of human judgements and automatic metrics (Rao and Tetreault, 2018; Luo et al., 2019), this was not the focus of the contribution and no deeper analysis or comparison was run. On the other hand, Yamshchikov et al. (2021) examined 13 content-related metrics in the context of formality transfer and paraphrasing, and show that none of the metrics is close enough to the human judgment. Briakou et al. (2021a) have recently evaluated automatic metrics on the task of multilingual formality transfer against human judgement. We also examine automatic metrics in

terms of correlation with human judgement, but there are some core differences between our contribution and their work. First, for style strength, they focus on comparing two different architectures in a cross-lingual setting using the correlation on human judgement for regression, and they do not provide this analysis for style classification, rather an evaluation against the gold label. In contrast, we adopt an architecture that provides regression and classification comparisons in fitting human judgments. Second, regarding content, Briakou et al. (2021a) focus on similarity (and therefore metrics) to the source sentence, while we stress the importance of triangulation also with the reference². Also, we introduce two learnable metrics in the evaluation setup, which correlation with human judgement shows to be the most informative. Third, they compare perplexity, likelihood, and pseudo-likelihood scores for fluency evaluation, while we provide a deeper evaluation of just perplexity considering though the two directions (Briakou et al. (2021a) evaluate only the informal-to-formal direction) and highlight differences that point to a potential benefit in using different approaches or evaluation strategies for the two directions.

In addition, we (i) use a continuous scale setting for human judgement which, unlike a discrete Likert scale, allows to normalize judgments (Graham et al., 2013), hence increasing homogeneity of the assessments; (ii) evaluate eight existing, published systems of different sorts (including state-of-the-art models) for both transfer directions, thereby potentially enabling a reconsideration of results as reported in previous work; (iii) study the nature of each evaluation aspect and the corresponding automatic metrics, analyzing the differences in the correlation between metric and human judgements that might arise under different conditions (e.g. looking at high-quality systems).

3 Data

We use GYAFC (Rao and Tetreault, 2018), a formality transfer dataset for English that contains aligned formal and informal sentences from two domains: Entertainment & Music and Family & Relationships. Figure 2 shows an example for alignment, transformation, and evaluation relations be-

²Although the reference is not always available, using it in studying evaluation metrics in comparison with how they behave when the source is used provides insights into the overall behaviour of such metrics and how they should best be employed even in the absence of a reference.

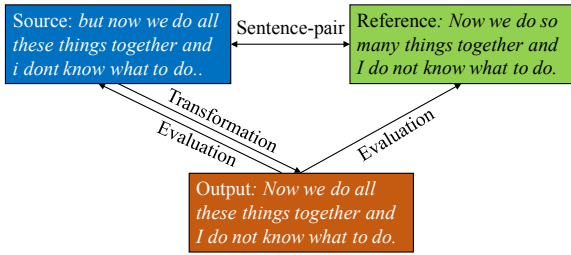


Figure 2: Alignment, transformation, evaluation pairs.

tween input, output, and reference. We run a human evaluation and a battery of automatic metrics on a selection of human- and machine-produced texts.

Source and Reference Texts The source and reference texts we use are from the Family & Relationships domain. The test set contains 1,332 and 1,019 sentences in “informal to formal” and “formal to informal” directions, respectively. There are four human references for each test sentence. We randomly select 80 source sentences (40 for each transfer direction) from the test set, as well as their corresponding human references. For each source sentence, we obtain the corresponding transformations as produced by eight different systems.

System Outputs The evaluation results are often affected by the system’s outputs, since if the evaluated systems are of different types, they may exhibit different error patterns so that various automatic evaluation metrics can be differently sensitive to these patterns (Ma et al., 2019; Mathur et al., 2020). To fully examine the evaluation methods, the systems we use are all from previous work, both supervised and unsupervised approaches.³ Overall, the eight systems yield a total of 640 output sentences (80 per system, 40 in each direction).

4 Methodologies

4.1 Human Evaluation

To facilitate the annotation and obtain a manageable size for each annotator, we split the 80 source sentences (Section 3) into four different surveys with 20 sentences each (10 for each transfer direction), and their corresponding system outputs plus one reference.

We recruited eight highly proficient English speakers for this task, i.e. two per survey, so that two annotations for each target sentence can be collected; from these we can use the average score

³Details of the systems are in Appendix A.1.

assigned, and also calculate inter-annotator agreement. The task is to rate the transferred sentence on a continuous scale (0-100), inspired by Direct Assessment (Graham et al., 2013, 2015), in terms of three evaluation aspects: (i) style strength (does the transformed sentence fit the target style?); (ii) content preservation (is the content of the transformed sentence the same as the original sentence?), and (iii) fluency (considering the target style, could the transformed sentence have been written by a native speaker?).

Before starting the rating task, we provided annotators with detailed guidelines and examples of transformed sentences along with plausible assessments for each aspect.⁴ We also reminded the annotators that such examples are only indicative of what we believe to be plausible judgements but there are many possible correct answers, of course.

4.2 Automatic Evaluation

We test a wide range of commonly used as well as new automatic metrics on the three aspects.

Style Strength The most commonly used method for assessing style strength is a style classifier, with the problem cast as a binary classification task (formal vs informal in formality transfer). Briakou et al. (2021a) have recently shown that a style regressor fine-tuned with English rating data correlates better with human judgments in other languages (Italian, French, and Portuguese). To run a proper comparison, we use BERT (Devlin et al., 2019) as our base model, and fine-tune it with style labelled data (GYAFC) and the rating data of PT16 (Pavlick and Tetreault, 2016) to obtain a style classifier (C-GYAFC) and a regressor (R-P16), respectively. Following Rao and Tetreault (2018), we collect sentences from PT16 with human rating from -3 to +1 as informal and the rest as formal, and train a style classifier on them (C-PT16). C-GYAFC and C-PT16 achieve an accuracy of 94.4% and 58.6% on the test sets, respectively.

Content Preservation We consider the following metrics, including both surface-based and embedding-based approaches:⁵

- BLEU (Papineni et al., 2002) It compares a given text to others (reference) by using a precision-oriented approach based on n -gram overlap;

⁴Screenshots of our annotation guidelines and interface are in Appendix A.3.

⁵The implementation details for automatic metrics are in Appendix A.2.

- chrF (Popović, 2015) It measures the similarity of sentences using the character n -gram F-score;
- ROUGE (Lin, 2004) It compares a given text to others (human reference) by using n -gram/the longest co-occurring in sequence overlap and a recall-oriented approach;
- WMD (Kusner et al., 2015) It measures the dissimilarity between two texts as an optimal transport problem which is based on word embedding.
- METEOR (Banerjee and Lavie, 2005) It computes the similarity score of two texts by using a combination of unigram-precision, unigram-recall, and some additional measures like stemming and synonymy matching.
- BERTScore (Zhang* et al., 2020) It computes a similarity score for each token in the candidate sentence with each token in the reference sentence. Instead of exact matches, it computes token similarity using contextual embeddings.
- BLEURT (Sellam et al., 2020) It is a learned evaluation metric based on BERT (Devlin et al., 2019), trained on human judgements. It is trained with a pre-training scheme that uses millions of synthetic examples to help the model generalize.
- COMET (Rei et al., 2020) It is a learnable metric which leverages cross-lingual pretrained language modeling resulting in multilingual machine translation evaluation models that exploit both source and reference sentences.

For assessing content preservation in the output, we can exploit both the source and the reference (see Fig. 2). When comparing our output to the source, we want to answer the following question: (a) how close in content is the generated text to the original text?, which addresses naturally the content preservation aspect of the task. When comparing our output to the human-produced reference, we want to answer a different question: (b) how similar is the automatically generated text to the human written one? Both are valid strategies, but by answering different questions they are likely to react differently to, and require, different metrics.

The advantages of the (a) approach are that evaluation is possible even without a human reference, it is the most natural way of assessing the task, and it does not incur reference bias (Fomicheva and Specia, 2016). The core problem lies in the use and interpretation of metrics: surface-based metrics (like BLEU) would score highest if nothing has

changed from input to output (if the model doesn't perform the task, basically), so aiming for a high score is pointless. A very low score is undesirable, too, however. For more sophisticated metrics, the problem is similar in the sense the highest score would be achieved if the two texts are identical, but since it is not fully clear what they measure exactly in terms of similarity, what to aim for isn't straightforward (an indication is provided by using metrics to compare source and reference).

The main advantage of the (b) approach is that metrics can be used in a more standard way: tending to the highest possible score is good for any of them, since getting close to the human solution is desirable. However, the gold reference is only one of many possible realisations, and while high scores are good, low scores can be somewhat meaningless, as proper meaning-preserving outputs may be very different from the human-produced ones, especially at surface level.

While we have as yet no specific solution to this, this study contributes substantially to a better understanding of automatic metrics, especially for content preservation, possibly leading to a combined metric which considers mainly the source, and possibly the reference(s) in a learning phase.

Fluency In formality transfer, both informal and formal outputs must be evaluated. Intuitively, the latter should be more fluent and grammatical than the former so that evaluating the fluency of informal sentences might be more challenging, both for humans and automatic metrics. We use the perplexity of the language model GPT-2 (Radford et al., 2019) fine-tuned with style labelled texts. Specifically, we fine-tune two GPT-2 models on informal sentences and formal sentences respectively, and then we use the target-style model to calculate the perplexity of the generated sentence. Finally, we provide a separate correlation analysis between automatic metrics and human judgements for the two transfer directions.

4.3 Correlation Methods

Pearson Correlation We employ Pearson correlation (r) as our main evaluation measure for system-/segment-level metrics:

$$r = \frac{\sum_{i=1}^n (H_i - \bar{H})(M_i - \bar{M})}{\sqrt{\sum_{i=1}^n (H_i - \bar{H})^2 \sum_{i=1}^n (M_i - \bar{M})^2}} \quad (1)$$

where H_i is the human assessment score, M_i is the corresponding score as predicted by a given metric.

Survey	N	Content	Style	Fluency	Overall
Survey 1	160	0.90	0.45	0.71	0.70
Survey 2	160	0.84	0.48	0.63	0.66
Survey 3	160	0.83	0.68	0.70	0.72
Survey 4	160	0.81	0.62	0.63	0.68
Overall	640	0.86	0.52	0.66	0.70

Table 1: Inter-Annotator Agreement.

\bar{H} and \bar{M} are the their means, respectively.

Kendall’s Tau-like formulation We follow the WMT17 Metrics Shared Task (Bojar et al., 2017) and take the official Kendall’s Tau-like formulation, τ , as the our main evaluation measure for segment-level metrics. A true pairwise comparison is likely to lead to more stable results for segment-level evaluation (Vazquez-Alvarez and Huckvale, 2002). The Kendall’s Tau-like formulation τ is as follows:

$$\tau = \frac{\text{Concordant} - \text{Discordant}}{\text{Concordant} + \text{Discordant}} \quad (2)$$

Where *Concordant* is the number of times for which a given metric suggests a higher score to the “better” hypothesis judged by human and *Discordant* is the number of times for which a given metric suggests a higher score to the “worse” hypothesis judged by human.

Most automatic metrics, like BLEU, aim to achieve a strong positive correlation with human assessment, with the exception of WMD and perplexity, where the smaller is better. We thereby employ absolute correlation value for WMD and perplexity in the following analysis.

5 Results and Analysis

In this section, we first measure the inter-annotator agreement of the human evaluation, then discuss both system-level and sentence-level evaluation results on the three aforementioned evaluation aspects, so as to provide a different perspective on the correlation between automatic metrics and human judgements under different conditions.

5.1 Inter-Annotator Agreement

There are two human judgements for each sentence and we measure their inter-annotator agreement (IAA) by computing the Pearson Correlation coefficient, instead of the commonly used Cohen’s K , since judgements are given on a continuous scale.

Table 1 presents the results of IAA for each aspect in each single survey and overall. Across the four surveys annotators have the highest agreement

	N	R-PT16	C-PT16	C-GYAFC
System-level (r)	8	<u>0.93</u>	<u>0.93</u>	<u>0.97</u>
Segment-level (τ)	640	0.33	0.39	0.42

Table 2: Correlation of automatic metrics in style strength with human judgements. The underlined scores indicate $p < 0.01$.

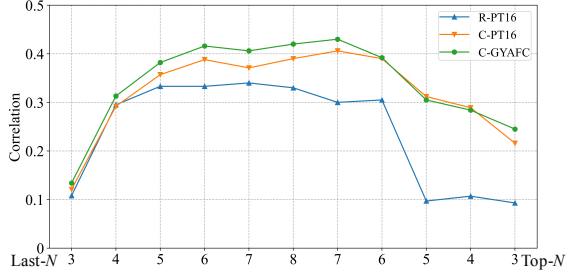


Figure 3: Kendall’s Tau-like correlation in style strength computed over the top-/last- N systems which are sorted by human judgements.

on the content aspect, followed by fluency, with style yielding the lowest scores, suggesting that annotators have more varied perceptions of sentence style than content. Overall, we achieve reasonable agreement for all surveys and evaluation aspects.

5.2 Style Strength

Table 2 shows the correlation of automatic metrics in style strength with human judgements. We see that C-GYAFC achieves the highest correlation at both system- and segment-level, R-PT16 and C-PT16 have the same system-level correlation score while the former has a slightly lower score at segment-level. Given that C-PT16 and C-GYAFC have close correlation scores while their performances on the test set are quite different, we also employ Pearson correlation to compute the segment-level result, and see rather different correlation scores (C-PT16 with 0.33 and C-GYAFC with 0.67). We think that evaluating the system outputs for a given source using C-PT16 and C-GYAFC results in similar scores ranking, so their Kendall’s Tau-like correlations are very close.

In general, it is easier to evaluate systems which have large differences in quality, while it is more difficult when systems have similar quality. To assess the reliability of automatic metrics for close-quality systems, we first sort the systems based on human judgements, and plot the correlation of the top-/last- N systems, with N ranging from all systems to the best/worst three systems (Fig. 3). We see that the correlation between automatic metrics

Systems	AVE. z	Source Sentence								Human Reference							
		BLEU	chrF	ROUGE-L	WMD	METEOR	BERTScore	BLEURT	COMET-w	BLEU	chrF	ROUGE-L	WMD	METEOR	BERTScore	BLEURT	COMET-w
Reference	0.009	0.291	0.492	0.501	1.334	0.487	0.605	0.235	0.314	-	-	-	-	-	-	-	
HIGH	0.542	0.608	0.775	0.758	0.672	0.808	0.880	0.851	0.895	0.366	0.547	0.582	1.086	0.554	0.643	0.347	0.400
NIU	0.491	0.637	0.772	0.769	0.652	0.808	0.873	0.818	0.899	0.376	0.560	0.605	1.036	0.567	0.649	0.373	0.418
BART	0.370	0.514	0.688	0.692	0.840	0.724	0.798	0.687	0.752	0.382	0.555	0.596	1.053	0.573	0.646	0.388	0.425
IBT	0.337	0.543	0.711	0.717	0.782	0.749	0.838	0.744	0.813	0.373	0.550	0.582	1.094	0.574	0.635	0.350	0.391
RAO	0.328	0.649	0.778	0.791	0.608	0.815	0.833	0.751	0.822	0.336	0.525	0.561	1.145	0.533	0.601	0.234	0.305
ZHOU	-0.659	0.610	0.717	0.765	0.758	0.770	0.739	0.189	0.318	0.253	0.461	0.494	1.351	0.469	0.508	-0.200	-0.125
YI	-0.669	0.547	0.684	0.731	0.823	0.728	0.716	0.148	0.320	0.288	0.483	0.517	1.307	0.491	0.524	-0.154	-0.059
LUO	-0.749	0.472	0.638	0.660	1.034	0.681	0.646	0.020	0.034	0.222	0.416	0.445	1.514	0.434	0.453	-0.289	-0.278

Table 3: Human evaluation (z-score) and automatic metrics in content preservation. Notes: (i) ↓ indicates the lower the score the better; (ii) COMET-w indicates that the input setting is not used.

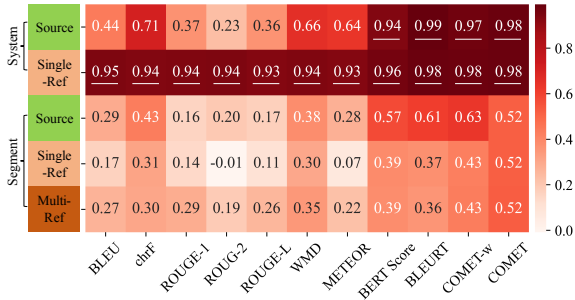


Figure 4: Correlations of automatic metrics computed against source/reference in content preservation with human judgments. Underlining indicates $p < 0.01$.

	BLEU	chrF	ROUGE-1	ROUGE-2	ROUGE-L	WMD	METEOR	BERTScore	BLEURT	COMET-w
Reference 2	0.28	0.37	0.33	0.10	0.36	0.46	0.21	0.59	0.61	0.61
Reference 3	0.25	0.41	0.37	0.12	0.35	0.47	0.34	0.60	0.60	0.55
Reference 4	0.37	0.41	0.46	0.24	0.46	0.49	0.31	0.60	0.56	0.62

Table 4: Kendall’s Tau-like correlation between using the first human reference and other references for evaluation content preservation at segment-level.

and human judgements decreases as we decrease N for both top- N and last- N systems, especially R-PT16 in the top- N systems. Again we observe that C-GYAFC and C-PT16 have similar scores over the top-/last- N systems. Overall, C-GYAFC appears to be the most stable model.

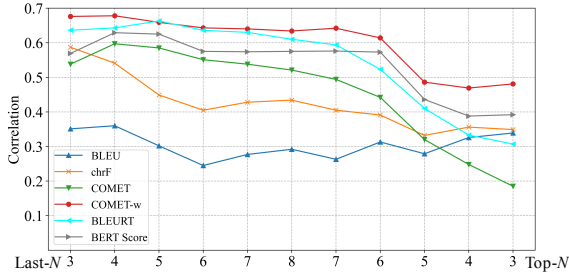
5.3 Content Preservation

As mentioned in the Introduction, since a style-transformed output should not alter the meaning of the input, content preservation can be measured against the input itself, or against a human reference in the expected target style. However, metrics cannot be used interchangeably (Section 4.2), as, for instance, the output is expected to have a higher

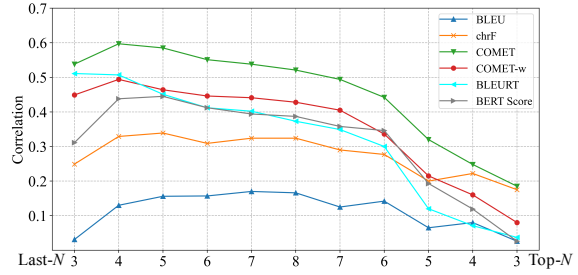
n -gram overlap with the reference, while this is not desirable with respect to the input.

Table 3 presents the results of human and automatic evaluation: all systems have a higher n -gram overlap (BLEU, chrF) with the source sentence than the human reference, indicating that existing models tend to copy from the input and lack diverse rewriting abilities. We also report the results for the reference against the source. Bearing in mind that the reference can be conceived as an optimal output, it is interesting to see that it does not score high in any metric, not even the learnable ones. This leaves some crucial open questions: how can these metrics be best used to assess content preservation in generated outputs? What are desirable scores? We also observe that RAO’s system has the highest scores of surface-based metrics (e.g. BLEU) with the source sentence while its scores with learnable metrics (e.g. BLEURT) are lower than some other systems (e.g. HIGH). In the evaluation against the human reference, the system BART and NIU achieve better results on most metrics.

Figure 4 shows the correlations of content preservation metrics with human judgments. For the system-level results, there is a big gap in correlation between source sentence and human reference for surface-based metrics (e.g. BLEU), but not for neural network based ones (e.g. COMET). Using the latter therefore seems to open up the possibility of automatically evaluating content without a human reference. It is interesting to see that the correlations of using source sentences at segment-level are all higher than using the human reference, and surface-based metrics of the latter correlate particularly poorly with human scores. We suggest two main reasons: (i) existing systems lack diverse rewriting ability given the source sentences, and



(a) Automatic metrics results against source sentence.



(b) Automatic metrics results against human reference.

Figure 5: Kendall’s Tau-like correlation in content preservation computed over the top-/last- N systems which are sorted by human judgements.

the annotators rate the generated sentences comparing them to the source sentence, not to a reference; (ii) human references are linguistically more diverse (e.g. word choice and order). The first one is not within the scope of this work. For the second aspect, we exploit the fact that we have multiple references available, and run the evaluation in a multi-reference setting; we observe that correlations for surface-based metrics improve as more variety is included, but not for neural ones. In Table 4, we see that learnable metrics using the first reference have higher correlation with other references than surface-based metrics. Overall, learnable metrics always have the highest correlation scores in evaluating content preservation using source sentences or human references, while surface-based metrics generally require a multi-reference setting.

Similar to style strength, we plot the correlation of the top-/last- N systems sorted by human judgements for the content aspect (Fig. 5). The correlation score between automatic metrics and human scores decreases as we decrease N for the top- N systems while this shows stability for the last- N systems. This suggests that evaluating high-quality TST systems is more challenging than evaluating low-quality systems. Again, we see that the correlation when using the source sentence has better stability than when using human references. Although BLEU and charF show stable performances, their correlations are lower than those by other metrics in most cases. Regardless of whether we use human references or source sentences, COMET(-w) generally has the highest correlation scores with human judgements under different conditions.

5.4 Fluency

Table 5 shows the absolute correlation of fluency metrics with human judgements. Unsurprisingly,

	N	Informal-to-Formal	Formal-to-Informal
System-level (r)	8	<u>0.96</u>	0.65
Segment-level (τ)	320	<u>0.52</u>	0.35

Table 5: Absolute correlation of automatic metrics in fluency with human judgements. The underlined scores indicate $p < 0.01$.

	Informal-to-Formal			Formal-to-Informal		
	GPT2-Inf	GPT2-For	r	GPT2-Inf	GPT2-For	r
Source	76	143	-	87	68	-
Reference	60	37	0.21	115	270	0.13
BART	34	26	0.33	24	28	0.02
IBT	32	26	0.32	33	40	0.17
NIU	43	37	0.30	71	75	0.03
HIGH	41	35	0.62	80	75	0.00
RAO	54	57	0.33	54	55	0.02
ZHOU	189	218	0.36	103	111	0.42
YI	160	182	0.31	205	436	0.27
LUO	128	152	0.43	6962	8191	0.17

Table 6: Results of GPT-2 based perplexity scores and their absolute Pearson correlation with human judgements at segment-level. Notes: (i) GPT2-Inf and GPT2-For are fine-tuned with informal sentences and formal sentences, respectively; (ii) the correlation is calculated using the perplexity of GPT-2 in the target style with human judgment.

we see that GPT-2 based perplexity correlates better with human scores in the direction informal-to-formal than in the opposite one, at both system- and segment-level. In general, a “good” formal sentence should be fluent, while an informal sentence might as well not be, and there can be varied perceptions by people. Indeed, we see higher IAA scores in the informal-to-formal direction (informal-to-formal: 0.70 vs informal-to-formal: 0.63). Table 6 presents the results of correlations and perplexity scores of GPT-2 in the two transfer directions for each system. The perplexity scores for most sentences are in the *correct* place, i.e. the scores from GPT2-Inf are higher than those from GPT2-For for the informal sentences, and viceversa. However,

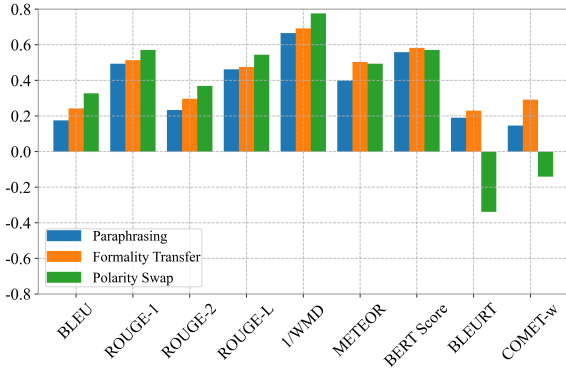


Figure 6: The distance between the source and target sentences as measured by content-related metrics.

we also observe that the correlations of informal-to-formal for each system (except ZHOU) are higher than those for the formal-to-informal direction. This confirms our hypothesis that assessing the fluency of informal sentences is not that obvious even for humans.

5.5 Broader Implications for Style Transfer

We have focused here on formality transfer, but polarity swap is also commonly defined as a style transfer task. In previous work, we have suggested that these tasks are intrinsically different, especially in terms of content preservation, since while formality transfer is somewhat akin to paraphrasing, in polarity swap the meaning is substantially altered (Lai et al., 2021a). This would imply that content-measuring metrics could not be used in the same way in the two tasks.

We further peek here into this issue, in view of future work that should evaluate metrics for the assessment of polarity swap, too, and show in Figure 6 the use of different metrics to measure the distance between the source and target sentences for paraphrasing, formality transfer, and polarity swap. Using n -gram based metrics, we see that the distance between source and target sentences in polarity swap is closer than in the other two tasks. With learnable metrics, instead, we see that source and target sentences for polarity swap are quite distant. Formality transfer shows overall the same trend as paraphrasing in all metrics, suggesting that it’s much more of a content-preserving paraphrase-like task than polarity swap, and metrics should be selected accordingly. Future work will explore how to best use them in polarity swap under different settings (using source vs reference, for example).

6 Conclusion

We have considered a wide range of automatic metrics on the three evaluation aspects of formality transfer, and assessed them against human judgements that we have elicited.

For **style strength**, we have compared the style classifiers and regressor in the setting of using the same raw data for training (with a binary label for classification and continuous scores for regression), as well as classifiers with different performances. We have observed that there is little difference among them when evaluating multiple TST systems. However, the style regressor performs worse when evaluating high-quality TST systems. For classifiers with different performances, we recommend the one with the highest performance since it results in the highest overall Pearson correlation with human judgements.

To assess **content preservation**, we have explored different kinds of automatic metrics using the source or reference(s), and have observed the following: (i) if using the source sentence, we strongly recommend employing learnable metrics since their correlation in that condition is much higher than those of traditional surface-based metrics (which are not indicative, since high scores correspond to not changing the input, hence not performing the task); still, the question of how scores should be interpreted and what score ranges are desirable remains open; (ii) most metrics are reliable to be used to measure and compare the performances at system-level when a human reference is available; (iii) however, we do not recommend to use surface-base metrics to measure sentence-level comparisons, especially with only one reference. Overall, learnable metrics seem to provide a more reliable measurement.

For **fluency**, perplexity can be used for evaluating the informal-to-formal direction, either at system- or segment-level, while it is clearly less reliable for the opposite direction, and it remains to be investigated how to best perform evaluation in this transfer direction, considering the wide variability of acceptable outputs.

This study focuses on formality transfer, and offers a better understanding of automatic evaluation thanks to the comprehensive correlations with human judgments herein conducted. However, the findings may not generalise to other tasks usually considered similar, such as polarity swap. To this end, future dedicated work will be required.

Acknowledgments

This work was partly funded by the China Scholarship Council (CSC). We are very grateful to the anonymous reviewers for their useful comments, especially in connection to closely related work, which contributed to strengthening this paper. We would also like to thank the Center for Information Technology of the University of Groningen for their support and for providing access to the Peregrine high performance computing cluster. We thank the annotators as well as Ana Guerberof Arenas and Amy Isard for testing, and helping us to improve, a preliminary version of the survey.

References

- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Ondřej Bojar, Yvette Graham, and Amir Kamran. 2017. [Results of the WMT17 metrics shared task](#). In *Proceedings of the Second Conference on Machine Translation*, pages 489–513, Copenhagen, Denmark. Association for Computational Linguistics.
- Eleftheria Briakou, Sweta Agrawal, Joel Tetreault, and Marine Carpuat. 2021a. [Evaluating the evaluation metrics for style transfer: A case study in multilingual formality transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1321–1336, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Eleftheria Briakou, Di Lu, Ke Zhang, and Joel Tetreault. 2021b. [Olá, bonjour, salve! XFORMAL: A benchmark for multilingual formality style transfer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3199–3216, Online. Association for Computational Linguistics.
- Yixin Cao, Ruihao Shui, Liangming Pan, Min-Yen Kan, Zhiyuan Liu, and Tat-Seng Chua. 2020. [Expertise style transfer: A new task towards better communication between experts and laymen](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1061–1071, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marina Fomicheva and Lucia Specia. 2016. [Reference bias in monolingual machine translation evaluation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 77–82, Berlin, Germany. Association for Computational Linguistics.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. [Continuous measurement scales in human evaluation of machine translation](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41, Sofia, Bulgaria. Association for Computational Linguistics.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2015. [Can machine translation systems be evaluated by the crowd alone](#). *Natural Language Engineering*, 23:3 – 30.
- Yufang Huang, Wentao Zhu, Deyi Xiong, Yiye Zhang, Changjian Hu, and Feiyu Xu. 2020. [Cycle-consistent adversarial autoencoders for unsupervised text style transfer](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2213–2223, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2021. [Deep learning for text style transfer: A survey](#). *arXiv preprint, arXiv:2011.00416*.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. [Disentangled representation learning for non-parallel text style transfer](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434, Florence, Italy. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of the International Conference on Learning Representations*.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. [From word embeddings to document distances](#). In *Proceedings of the 32nd In-*

- ternational Conference on Machine Learning, volume 37 of *Proceedings of Machine Learning Research*, pages 957–966, Lille, France.
- Huiyuan Lai, Antonio Toral, and Malvina Nissim. 2021a. [Generic resources are what you need: Style transfer tasks without task-specific parallel training data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4241–4254, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Huiyuan Lai, Antonio Toral, and Malvina Nissim. 2021b. [Thank you BART! rewarding pre-trained models improves formality style transfer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 484–494, Online. Association for Computational Linguistics.
- Dongkyu Lee, Zhiliang Tian, Lanqing Xue, and Nevin L. Zhang. 2021. [Enhancing content preservation in text style transfer using reverse attention and conditional layer normalization](#). 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 93–102, Online. Association for Computational Linguistics.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018a. [Delete, retrieve, generate: a simple approach to sentiment and style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1865–1874, New Orleans, Louisiana. Association for Computational Linguistics.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018b. [Delete, retrieve, generate: a simple approach to sentiment and style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1865–1874, New Orleans, Louisiana. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019. [A dual reinforcement learning framework for unsupervised text style transfer](#). In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5116–5122.
- Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. [StylePTB: A compositional benchmark for fine-grained controllable text style transfer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2116–2138, Online. Association for Computational Linguistics.
- Qingsong Ma, Johnny Wei, Ondřej Bojar, and Yvette Graham. 2019. [Results of the WMT19 metrics shared task: Segment-level and strong MT systems pose big challenges](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 62–90, Florence, Italy. Association for Computational Linguistics.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020. [Results of the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 688–725, Online. Association for Computational Linguistics.
- Remi Mir, Bjarke Felbo, Nick Obradovich, and Iyad Rahwan. 2019. [Evaluating style transfer for text](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 495–504, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xing Niu, Sudha Rao, and Marine Carpuat. 2018. [Multi-task neural models for translating between styles within and across languages](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1008–1021, Santa Fe, New Mexico, USA. 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.
- Ellie Pavlick and Joel Tetreault. 2016. [An empirical analysis of formality in online communication](#). *Transactions of the Association for Computational Linguistics*, 4:61–74.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. [Style transfer through back-translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876, Melbourne, Australia. Association for Computational Linguistics.

- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3936–3945, Hong Kong, China. Association for Computational Linguistics.
- Sudha Rao and Joel Tetreault. 2018. [Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer.](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 129–140, New Orleans, Louisiana. 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.
- Machel Reid and Victor Zhong. 2021. [LEWIS: Levenshtein editing for unsupervised text style transfer.](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3932–3944, Online. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Controlling politeness in neural machine translation via side constraints.](#) In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40, San Diego, California. Association for Computational Linguistics.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment.](#) In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6833–6844, Red Hook, NY, USA. Curran Associates Inc.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. [“transforming” delete, retrieve, generate approach for controlled text style transfer.](#) 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 3269–3279, Hong Kong, China. Association for Computational Linguistics.
- Alexey Tikhonov, Viacheslav Shibaev, Aleksander Nagaev, Aigul Nugmanova, and Ivan P. Yamshchikov. 2019. [Style transfer for texts: Retrain, report errors, compare with rewrites.](#) 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 3936–3945, Hong Kong, China. Association for Computational Linguistics.
- Yolanda Vazquez-Alvarez and Mark Huckvale. 2002. The reliability of the itu-t p.85 standard for the evaluation of text-to-speech systems.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Jingjing Xu, Xu Sun, Qi Zeng, Xiaodong Zhang, Xuancheng Ren, Houfeng Wang, and Wenjie Li. 2018. [Unpaired sentiment-to-sentiment translation: A cycled reinforcement learning approach.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 979–988, Melbourne, Australia. Association for Computational Linguistics.
- Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, and Colin Cherry. 2012. [Paraphrasing for style.](#) In *Proceedings of COLING 2012*, pages 2899–2914, Mumbai, India. The COLING 2012 Organizing Committee.
- Ivan P. Yamshchikov, Viacheslav Shibaev, Nikolay Khlebnikov, and Alexey Tikhonov. 2021. [Style-transfer and paraphrase: Looking for a sensible semantic similarity metric.](#) In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14213–14220.
- Xiaoyuan Yi, Zhenghao Liu, Wenhao Li, and Maosong Sun. 2020. [Text style transfer via learning style instance supported latent space.](#) In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3801–3807.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert.](#) In *International Conference on Learning Representations*.
- Chulun Zhou, Liangyu Chen, Jiachen Liu, Xinyan Xiao, Jinsong Su, Sheng Guo, and Hua Wu. 2020. [Exploring contextual word-level style relevance for unsupervised style transfer.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7135–7144, Online. Association for Computational Linguistics.

A Appendices:

These Appendices include: (i) evaluated systems (A.1); (ii) implementation details for automatic metrics (A.2); and (iii) annotation guidelines and interface (A.3).

A.1 Evaluated Systems

Table A.1 presents the systems’ ranking based on the human judgements. We use eight published systems of different sorts (including state-of-the-art models). For **supervised approaches**, we include the following systems:

- RAO (Rao and Tetreault, 2018): A copy-enriched NMT model trained on the rule-processed data and the additional forward and backward translations produced by the PBMT model;
- NIU (Niu et al., 2018): A bi-directional model trained on formality-tagged bilingual data using multi-task learning;
- BART (Lai et al., 2021b): Fine-tuning a pre-trained model BART with gold parallel data and reward strategies;
- HIGH (Lai et al., 2021a): Fine-tuning BART with high-quality synthetic parallel data and reward strategies.

For **unsupervised approaches**, we include the following systems:

- LUO (Luo et al., 2019): A dual reinforcement learning framework that directly transforms the style of the text via a one-step mapping model without parallel data;
- YI (Yi et al., 2020): A style instance supported method that learns a more discriminative and expressive latent space to enhance style signals and make a better balance between style and content;
- Zhou (Zhou et al., 2020): An attentional seq2seq model that pre-trains the model to reconstruct the source sentence and re-predict its word-level style relevance;
- IBT (Lai et al., 2021a): An iterative back-translation framework based on the pre-trained seq2seq model BART.

Table A.2 presents automatic evaluation results in content preservation.

A.2 Implementation Details for Automatic Metrics

- BLEU: We adopt `sentence_bleu` of the NLTK library with a smoothing function to compute the segment-level score, and `multi-bleu.perl` with default settings for system-level.⁶
- chrF: Following Briakou et al. (2021a), we use `sentence_chrf` of the open-sourced implementation `sacreBLEU`.⁷
- ROUGE: We use the open-sourced implementations `Rouge`.⁸
- WMD: We employ the `gensim` library and word embedding `googlenews-vectors-negative300.bin`.⁹
- METEOR: We adopt the NLTK library.
- BERTScore: We use the official implementation with a rescaling function.¹⁰
- BLEURT: We use the official checkpoint of `bleurt-large-512`.¹¹
- COMET: We adopt the official checkpoint of `wmt-large-da-estimator-1719`.¹² COMET-QE is a referenceless metric that uses source and output only. But we found that it yielded lower correlations with human judgements than COMET in our evaluations. This may be because the input and output are different languages in COMET-QE training.
- Style and Fluency: All experiments are implemented atop Transformers (Wolf et al., 2020) using BERT base model (cased) for style and GPT-2 base model for fluency. We fine-tune models using the Adam optimiser (Kingma and Ba, 2015) with learning rate of 1e-5 for BERT and 3e-5 for GPT-2, with a batch size of 32 for all experiments.

A.3 Annotation Guidelines and Interface

Figure A.1 show the screenshots of task guidelines and annotation interface.

⁶<https://www.nltk.org/>

⁷<https://github.com/mjpost/sacrebleu>

⁸<https://github.com/pltrdy/rouge>

⁹<https://radimrehurek.com/gensim/index.html>

¹⁰https://github.com/Tiiiger/bert_score

¹¹<https://github.com/google-research/bleurt>

¹²<https://github.com/Unbabel/COMET>

Style				Content				Fluency			
System	Rank	AVE. s	AVE. z	System	Rank	AVE. s	AVE. z	System	Rank	AVE. s	AVE. z
BART	1	82.7	0.494	HIGH	1	92.4	0.542	BART	1	87.8	0.540
REF	2	82.3	0.469	NIU	2	90.7	0.491	IBT	2	86.0	0.491
IBT	3	80.1	0.407	BART	3	86.5	0.370	NIU	3	84.9	0.463
NIU	4	76.9	0.297	IBT	4	85.1	0.337	HIGH	4	83.3	0.420
HIGH	5	76.3	0.293	RAO	5	84.7	0.328	REF	5	82.4	0.385
RAO	6	70.2	0.085	REF	6	73.6	0.009	RAO	6	77.3	0.247
YI	7	51.1	-0.588	ZHOU	7	50.9	-0.659	ZHOU	7	45.1	-0.717
ZHOU	8	47.2	-0.726	YI	8	50.5	-0.669	YI	8	38.6	-0.903
LUO	9	46.7	-0.731	LUO	9	47.6	-0.749	LUO	9	37.9	-0.926

Table A.1: Results based on original human evaluation and z-score.

Systems	BLEU	chrF	ROUGE-1	ROUGE-2	ROUGE-L	WMD ↓	METEOR	BERTScore	BLEURT	COMET-w	BLEU	chrF	ROUGE-1	ROUGE-2	ROUGE-L	WMD ↓	METEOR	BERTScore	BLEURT	COMET-w
	Reference 1										Reference 2									
Reference	0.291	0.492	0.533	0.307	0.501	1.334	0.487	0.605	0.235	0.314	0.231	0.459	0.494	0.259	0.449	1.469	0.444	0.565	0.155	0.202
HIGH	0.366	0.547	0.624	0.401	0.582	1.086	0.554	0.643	0.347	0.400	0.300	0.515	0.564	0.342	0.512	1.260	0.521	0.605	0.317	0.289
NIU	0.376	0.560	0.646	0.434	0.605	1.036	0.567	0.649	0.373	0.418	0.333	0.525	0.578	0.369	0.526	1.202	0.538	0.617	0.329	0.286
BART	0.382	0.555	0.632	0.412	0.596	1.053	0.573	0.646	0.388	0.425	0.305	0.511	0.561	0.349	0.513	1.278	0.526	0.605	0.353	0.279
IBT	0.373	0.550	0.620	0.404	0.582	1.094	0.574	0.635	0.350	0.391	0.291	0.503	0.553	0.335	0.503	1.289	0.512	0.595	0.305	0.271
RAO	0.336	0.525	0.602	0.367	0.561	1.145	0.533	0.601	0.234	0.305	0.297	0.505	0.556	0.344	0.512	1.281	0.512	0.568	0.200	0.196
ZHOU	0.253	0.461	0.536	0.300	0.494	1.351	0.469	0.508	-0.200	-0.125	0.245	0.451	0.495	0.271	0.444	1.488	0.476	0.478	-0.206	-0.212
YI	0.288	0.483	0.551	0.324	0.517	1.307	0.491	0.524	-0.154	-0.059	0.225	0.443	0.497	0.263	0.454	1.475	0.457	0.488	-0.203	-0.167
LUO	0.222	0.416	0.483	0.272	0.445	1.514	0.434	0.453	-0.289	-0.278	0.189	0.381	0.419	0.209	0.378	1.694	0.389	0.425	-0.266	-0.368
Systems	Reference 3										Reference 4									
Reference	0.213	0.442	0.472	0.231	0.434	1.537	0.433	0.567	0.102	0.190	0.231	0.459	0.505	0.261	0.461	1.438	0.466	0.595	0.224	0.293
HIGH	0.316	0.513	0.566	0.340	0.528	1.229	0.506	0.617	0.236	0.326	0.295	0.511	0.585	0.343	0.535	1.227	0.526	0.634	0.327	0.412
NIU	0.325	0.509	0.574	0.351	0.534	1.232	0.505	0.612	0.257	0.309	0.310	0.518	0.607	0.365	0.552	1.173	0.548	0.637	0.349	0.413
BART	0.341	0.517	0.577	0.361	0.539	1.208	0.526	0.617	0.274	0.354	0.327	0.532	0.621	0.384	0.574	1.128	0.565	0.655	0.405	0.447
IBT	0.307	0.514	0.570	0.344	0.531	1.220	0.522	0.614	0.267	0.328	0.316	0.520	0.592	0.363	0.543	1.217	0.534	0.632	0.332	0.388
RAO	0.293	0.499	0.556	0.329	0.511	1.288	0.493	0.574	0.140	0.252	0.293	0.505	0.577	0.336	0.526	1.234	0.541	0.600	0.250	0.315
ZHOU	0.227	0.419	0.478	0.245	0.438	1.496	0.421	0.489	-0.257	-0.186	0.210	0.425	0.507	0.248	0.453	1.451	0.448	0.507	-0.212	-0.162
YI	0.220	0.436	0.487	0.255	0.449	1.477	0.416	0.488	-0.263	-0.149	0.204	0.432	0.501	0.250	0.458	1.466	0.430	0.509	-0.182	-0.086
LUO	0.189	0.380	0.422	0.244	0.390	1.671	0.371	0.431	-0.346	-0.356	0.197	0.393	0.458	0.243	0.410	1.591	0.420	0.451	-0.282	-0.317

Table A.2: Automatic evaluation results in content preservation. Notes: (i) the results of Reference is the distance between source and reference sentence measuring by metrics; (ii) ↓ indicates the lower score is better.



Task Guideline

This task consists of judging sentence changes. To this aim, you will be shown different changes in a given sentence. The changes are related to style: from informal to formal or from formal to informal. We call these changes transformations.

You are asked to judge to what extent the transformations are successful by assessing the **content** (the meaning of the sentence), the **style** (how appropriate is the formal or informal tone), and the **fluency** (how close the sentence is to the language a native speaker uses) of the new sentence. You will see an original sentence and various possible transformations. For each transformation you will have to use sliders (little online buttons to control the level of agreement) to indicate how much you agree with each statement. There are three sliders: one for content, one for style and one for fluency. The more you slide to the right, the higher your level of agreement.

The table below shows you examples of these transformations, so you can have an idea of the type of changes you will see. The list is not exhaustive.

Original informal sentence	Transformed formal sentence
he is wayyyy hottt	He is very attractive.
yes, except for episode iv.	Yes, but not for episode IV.
I've watched it and it is AWESOME!!!!	I viewed it and I believe it is a quality program.
Well... Do you talk to that someone much?	Do you talk to that person often?
Haven't seen the tv series, but R.O.D.	I have not seen the television series, however I have seen the R.O.D
that page did not give me viroses (i think)	I don't think that page gave me viruses.
I didn't know they had an HBO in the 80's	I did not know HBO existed in the 1980s.
my exams r not over yet	My exams are not over yet.
But you will DEFINITELY know when you are in love!	You will definitely know when you are in love.

Before starting the actual task, on the next page you will see examples of transformed sentences with plausible assessments on the sliders for each aspect (content, style, fluency). Please, do remember that such examples are indicative of what we believe to be plausible judgements, but there are many possible correct answers. Your judgement is what we are looking for.

(a) A screenshot of task guidelines.



Transforming informal to formal

Original sentence:

it all depends on when ur ready.

Transformed sentence:

It depends on when you are ready.

Strongly disagree Strongly agree

The content of the transformed sentence is the same as the original sentence.

The transformed sentence fits the target style.

Considering the target style, the transformed sentence could have been written by a native speaker.

Transforming informal to formal

Original sentence:

it all depends on when ur ready.

Transformed sentence:

It all depends on when your ready.

Strongly disagree Strongly agree

The content of the transformed sentence is the same as the original sentence.

The transformed sentence fits the target style.

Considering the target style, the transformed sentence could have been written by a native speaker.

(b) A screenshot of annotation interface.

Figure A.1: Screenshots of our interface.