

Evaluating the Evaluation Metrics for Style Transfer: A Case Study in Multilingual Formality Transfer

Eleftheria Briakou[✉], Sweta Agrawal[✉], Joel Tetreault[✉], Marine Carpuat[✉]

[✉]University of Maryland, [✉]Dataminr, Inc.

ebriakou@cs.umd.edu, sweagraw@cs.umd.edu,
jtetreault@dataminr.com, marine@cs.umd.edu

Abstract

While the field of style transfer (ST) has been growing rapidly, it has been hampered by a lack of standardized practices for automatic evaluation. In this paper, we evaluate leading ST automatic metrics on the oft-researched task of formality style transfer. Unlike previous evaluations, which focus solely on English, we expand our focus to Brazilian-Portuguese, French, and Italian, making this work the first multilingual evaluation of metrics in ST. We outline best practices for automatic evaluation in (formality) style transfer and identify several models that correlate well with human judgments and are robust across languages. We hope that this work will help accelerate development in ST, where human evaluation is often challenging to collect.

1 Introduction

Textual style transfer (ST) is defined as a generation task where a text sequence is paraphrased while controlling one aspect of its style (Jin et al., 2020). For instance, the informal sentence in Italian “*in bocca al lupo!*” (i.e., “good luck”) is rewritten to the formal version “*Ti rivolgo un sincero augurio!*” (i.e., “I send you a sincere wish!”). Despite the growing attention on ST in the NLP literature (Jin et al., 2020), progress is hampered by a lack of standardized and reliable automatic evaluation metrics. Standardizing the latter would allow for quicker development of new methods and comparison to prior art without relying on time and cost-intensive human evaluation that is currently employed by more than 70% of ST papers (Briakou et al., 2021a).

ST is usually evaluated across three dimensions: style transfer (i.e., has the style of the generated output changed as intended?), meaning preservation (i.e., are the semantics of the input preserved?), and fluency (i.e., is the output well-formed?). As we will see, a wide range of automatic evaluation metrics and models has been used to quantify each of

these dimensions. For example, prior work has employed as many as nine different automatic systems to rate formality alone (see Table 1). However, it is not clear how different automatic metrics compare to each other and how well they agree with human judgments. Furthermore, previous studies of automatic evaluation have exclusively focused on the English language (Yamshchikov et al., 2021; Pang, 2019; Pang and Gimpel, 2019; Tikhonov et al., 2019; Mir et al., 2019); yet, ST requires evaluation methods that generalize reliably beyond English.

We address these limitations by conducting a controlled empirical comparison of commonly used automatic evaluation metrics. Concretely, for all three evaluation dimensions, we compile a list of different automatic evaluation approaches used in prior ST work and study how well they correlate with human judgments. We choose to build on available resources as collecting human judgments across the evaluation dimensions is a costly process that requires recruiting fluent speakers in each language addressed in evaluation. While there are many stylistic transformations in ST, we conduct our study through the lens of formality style transfer (FOST), which is one of the most popular style dimensions considered by past ST work (Jin et al., 2020; Briakou et al., 2021a) and for which reference outputs and human judgments are available for four languages: English, Brazilian-Portuguese, French, and Italian.

- We contribute a meta-evaluation study that is not only the first *large-scale* comparison of automatic metrics for ST but is also the first work to investigate the robustness of these metrics in *multilingual* settings.
- We show that automatic evaluation approaches based on a formality regression model fine-tuned on XLM-R and the chrF metric correlate well with human judgments for style transfer and meaning preservation, re-

spectively, and propose that the field adopts their usage. These metrics are shown to work well *across* languages, and not just in English.

- We show that framing style transfer evaluation as a binary classification task is problematic and propose that the field treats it as a regression task to better mirror human evaluation.
- Our analysis code and meta-evaluation files with system outputs are made public to facilitate further work in developing better automatic metrics for ST: <https://github.com/Elbria/xformal-FoST-meta>.

2 Background

2.1 Limitations of Automatic Evaluation

Recent work highlights the need for research to improve evaluation practices for ST along multiple directions. Not only does ST lack standardized evaluation practices (Yamshchikov et al., 2021), but commonly used methods have major drawbacks which hamper progress in this field. Pang (2019) and Pang and Gimpel (2019) show that the most widely adopted automatic metric, BLEU, can be gamed. They observe that untransferred text achieves the highest BLEU score for the task of sentiment transfer, questioning complex models’ ability to surpass this trivial baseline. Mir et al. (2019) discuss the inherent trade-off between ST evaluation aspects and propose that models are evaluated at specific points of their trade-off plots. Tikhonov et al. (2019) argue that, despite their cost, human-written references are needed for future experiments with style transfer. They also show that comparing models without reporting error margins can lead to incorrect conclusions as state-of-the-art models sometimes end up within error margins from one another.

2.2 Structured Review of ST Evaluation

We systematically review automatic evaluation practices in ST with formality as a case study. We select FoST for this work since it is one of the most frequently studied styles (Jin et al., 2020) and there is human annotated data including human references available for these evaluations (Rao and Tetreault, 2018; Briakou et al., 2021b). Tables 1 and 2 summarize evaluation details for all FoST methods in papers from the ST survey by Jin et al. (2020).¹ Most works employ automatic evaluation

¹The complete list is hosted at: <https://github.com/fuzhenxin/Style-Transfer-in-Text>

for *style* (87%) and *meaning* preservation (83%). *Fluency* is the least frequently evaluated dimension (43%), while 74% of papers employ automatic metrics to assess the *overall* quality of system outputs that captures all desirable aspects.

Across dimensions, papers also frequently rely on human evaluation (55%, 58%, 60%, and 40% for style, meaning, fluency, and overall). However, human judgments and automatic metrics do not always agree on the best-performing system. In 60% of evaluations, the top-ranked system is the same according to human and automatic evaluation (marked as ✓ in Table 1), and their ranking disagrees in 40% of evaluations (marked as ✗ in Table 1). When there is a disagreement, human evaluation is trusted more and viewed as the standard. This highlights the need for a systematic evaluation of automatic evaluation metrics.

Finally, almost all papers (91%) consider FoST for English (EN), as summarized in Table 2. There are only two exceptions: Korotkova et al. (2019) study FoST for Latvian (LV) and Estonian (ET) in addition to EN, while Briakou et al. (2021b) study FoST for 3 Romance languages: Brazilian Portuguese (BR-PT), French (FR), and Italian (IT). The former provides system output samples as a means of evaluation, and the latter employs human evaluations, highlighting the challenges of automatic evaluation in multilingual settings.

Next, we review the automatic metrics used for each dimension of evaluation in FoST papers. As we will see, a wide range of approaches is used. Yet, it remains unclear how they compare to each other, what their respective strengths and weaknesses are, and how they might generalize to languages other than English.

2.3 Automatic Metrics for FoST

Formality Style transfer is often evaluated using model-based approaches. The most frequent method consists of training a binary classifier on human written formal vs. informal pairs. The classifier is later used to predict the percentage of generated outputs that match the desired attribute per evaluated system—the system with the highest percentage is considered the best performing with respect to style. Across methods, the corpus used to train the classifier is the GYAFC parallel-corpus (Rao and Tetreault, 2018) consisting of 105K parallel informal-formal human-generated excerpts. This corpus is curated for FoST in EN,

PAPER ID	STYLE			MEANING		FLUENCY			OVERALL		
	metric	arch.		metric	arch.	metric	arch.		metric		
[1]	REG	Linear reg.	✗	CLS	CNN	✓	REG	Linear reg.	✗	r-BLEU	✓
[2]										r-BLEU	–
[3]	CLS	CNN	–	r-BLEU		–				GM(S,M)	–
[4]	CLS	CNN	✗	r-BLEU		✓				GM(S,M)	✓
[5]	CLS	CNN	✗	r-BLEU		✓					
[6]	CLS	LSTM	–	CLS	BERT	–				r-BLEU	–
[7]										r-BLEU	–
[8]	CLS	CNN	–								
[9]	CLS	LSTM	–	EMB-SIM		–	PPL	LM (RNN)	–	F1(S,M)	–
[10]	CLS	ROBERTa	✗	EMB-SIM		✓	PPL	LM (ROBERTa)	–	J(S,M,F)	✓
[11]	CLS	CNN	✗	r-BLEU		✗				F1(S,M)	–
[12]	CLS	GRU	✗				CLS	Linear reg.	✗	r-BLEU	✗
[13]	CLS	BERT	✓	r-BLEU		✓	PPL	LM (KenLM)	✗	GM(S,M,F)	–
[14]										r-BLEU	–
[15]	CLS	FASTTEXT	✓	r-BLEU		✓	PPL	LM (GPT)	✓		
[16]	CLS	CNN	–	r-BLEU		–	PPL	LM (LSTM)	–		
[17]	CLS	CNN	✓	r-BLEU		✓					
[18]	CLS	CNN	✓	r-BLEU		✓				GMHM(S,M)	✓
[19]	CLS	GRU	–	CLS	BERT	–				r-BLEU	✓
[20]	CLS	ROBERTa	–	r-BLEU		–	PPL	LM (GPT)	–	GMHM(S,M)	–
[21]	CLS	CNN	–	r-BLEU		✓	PPL	LM (GPT)	✓		
[22]										r-BLEU	–
[23]	REG	BERT	✗	s-BLEU		✓	PPL	LM (KenLM)	✗	r-BLEU	✗

Table 1: Details on automatic evaluation practices of prior work for FOST. For each dimension, ✓ and ✗ denote whether the best ranking system based on the automatic evaluation agrees or disagrees with the one pointed by the human evaluation; – denotes that no human evaluation is conducted. REG stands for regression, CLS for classification, SIM for similarity, EMB for embedding-based, LM for language model, GM or HM for geometric or harmonic mean, J for corpus-level product, and arch. for architecture. The mappings between each paper ID and its corresponding citation is included in Table 2.

PAPER ID	CITATION	LANGUAGE
[1]	Rao and Tetreault (2018)	EN
[2]	Kajiwara (2019)	EN
[3]	Li et al. (2019)	EN
[4]	Luo et al. (2019)	EN
[5]	Shang et al. (2019)	EN
[6]	Wang et al. (2019)	EN
[7]	Xu et al. (2019)	EN
[8]	Korotkova et al. (2019)	EN, ET, LV
[9]	Gong et al. (2019)	EN
[10]	Krishna et al. (2020)	EN
[11]	Sancheti et al. (2020)	EN
[12]	Wu et al. (2020)	EN
[13]	Yi et al. (2020)	EN
[14]	Zhang et al. (2020b)	EN
[15]	Goyal et al. (2020)	EN
[16]	He et al. (2019)	EN
[17]	Cheng et al. (2020)	EN
[18]	Zhou et al. (2020)	EN
[19]	Wang et al. (2020)	EN
[20]	Li et al. (2020)	EN
[21]	Liu et al. (2021)	EN
[22]	Lyu et al. (2021)	EN
[23]	Briakou et al. (2021b)	BR-PT, FR, IT

Table 2: List of prior works on FOST along with languages (i.e., iso codes) addressed in each of them.

while similar resources are not available for other languages. Different model architectures have been used by prior work (e.g., CNN, LSTM, GRU, fine-tuning on pre-trained language models such as ROBERTa and BERT; Table 1). In most papers, the resulting classifier is evaluated on the test side of the GYAFC corpus, reporting accuracy scores

in the range of 80 – 90%. Despite the high accuracy scores, the best ranking system under the classifier is very often in disagreement with human evaluations (marked as ✗ under the third sub-column of style of Table 1). A few works train regression-based models instead, using the training data of Pavlick and Tetreault (2016) that are human-annotated for formality on a 7-point scale—while, again, this resource is only available for EN.

Meaning Preservation Evaluation of this dimension is performed using a wider spectrum of approaches, as presented in the third column of Table 1. The most frequently used metric is reference-BLEU (r-BLEU), which is based on the n -gram precision of the system output compared to human rewrites of the desired formality. Other approaches include self-BLEU (s-BLEU), where the system output is compared to its input, measuring the semantic similarity between the system input and its output, or regression models (e.g., CNN, BERT) trained on data annotated for similarity-based tasks, such as the Semantic Textual Similarity task (STS) (Agirre et al., 2016).

Fluency Fluency is typically evaluated with model-based approaches (see fourth column of Table 1). Among those, the most frequent method is

that of computing perplexity (PPL) under a language model. The latter is either trained from scratch on the same corpus used to train the FoST models (i.e., GYAFC) using different underlying architectures (e.g., KenLM, LSTM), or employ large pre-trained language models (e.g., GPT). A few other works train models on EN data annotated for grammaticality (Heilman et al., 2014) or linguistic acceptability (Warstadt et al., 2019) instead.

Overall Systems’ overall quality (see fifth column of Table 2) is mostly evaluated using r-BLEU or by combining independently computed metrics into a single score (e.g., geometric mean - GM(.), harmonic mean - HM(.), F1(.)). Moreover, 6 out of 8 approaches that rely on combined scores do not include fluency scores in their overall evaluation.

English Focus Since most of the current work on FoST and ST is in EN, prior work relies heavily on EN resources for designing automatic evaluation methods. For instance, resources for training stylistic classifiers or regression models are not available for other languages. For the same reason, it is unclear whether model-based approaches for measuring meaning preservation and fluency can be ported to multilingual settings. Furthermore, reference-based evaluations (e.g., r-BLEU) require human rewrites that are only available for EN, BR-PT, IT, and FR. Finally, even though perplexity does not rely on annotated data, without standardizing the data language models are trained on, we cannot make meaningful cross-system comparisons.

2.4 Summary

Reviewing the literature shows the lack of standardized metrics for ST evaluation, which hampers comparisons across papers, the lack of agreement between human judgments and automatic metrics, which hampers system development, and the lack of portability to languages other than English which severely limits the impact of the work. These issues motivate the controlled multilingual evaluation of evaluation metrics in our paper.

3 Evaluating Evaluation Metrics

We evaluate evaluation metrics (described in §3.2) for multilingual FoST, in four languages for which human evaluation judgments (described in §3.1) on FoST system outputs are available.

3.1 Human Judgments

We use human judgments collected by prior work of Rao and Tetreault (2018) for EN and Briakou et al. (2021b) for BR-PT, FR, and IT. We include details on their annotation frameworks, the quality of human judges, and the evaluated systems below.

Human Annotations We briefly describe the annotation frameworks employed by Rao and Tetreault (2018) and Briakou et al. (2021b) to collect human judgments for each evaluation aspect: 1. **formality** ratings are collected—for each system output—on a 7-point discrete scale, ranging from -3 to $+3$, as per Lahiri (2015) (*Very informal, Informal, Somewhat Informal, Neutral, Somewhat Formal, Formal, Very Formal*); 2. **meaning preservation** judgments adopt the Semantic Textual Similarity annotation scheme of Agirre et al. (2016), where an informal input and its corresponding formal system output are rated on a scale from 1 to 6 based on their similarity (*Completely dissimilar, Not equivalent but on same topic, Not equivalent but share some details, Roughly equivalent, Mostly equivalent, Completely equivalent*); 3. **fluency** judgments are collected for each system output on a discrete scale of 1 to 5, as per Heilman et al. (2014) (*Other, Incomprehensible, Somewhat Comprehensible, Comprehensible, Perfect*); 4. **overall** judgments are collected following a *relative ranking* approach: all system outputs are ranked in order of their formality, taking into account both meaning preservation and fluency.

Human Annotators Both studies recruited workers from the Amazon Mechanical Turk platform after employing quality control methods to exclude poor quality workers (i.e., manual checks for EN, and qualification tests for BR-PT, FR, and IT). For all human evaluations and languages Briakou et al. (2021b) report at least moderate inter-annotator agreement.

Evaluated Systems The evaluated system outputs were sampled from 5 FoST models for each language, spanning a range of simple baselines to neural architectures (Rao and Tetreault, 2018; Briakou et al., 2021b). We also include detailed descriptions of them in Appendix C. For each evaluation dimension 500 outputs are evaluated for EN and 100 outputs per system for BR-PT, FR, and IT.

3.2 Evaluation Metrics

For the FOST evaluation aspects described below, we cover a broad spectrum of approaches that range from dedicated models for the tasks at hand to more lightweight methods relying on unsupervised approaches and automated metrics.

Formality We benchmark model-based approaches that fine-tune multilingual pre-trained language models (i.e., XLM-R, mBERT), where the task of formality detection is modeled either as a **binary classification** task (i.e., formal vs. informal), or as a **regression** task that predicts different formality levels on an ordinal scale.

Meaning Preservation We evaluate the BLEU score (Papineni et al., 2002) of the system output compared to the reference rewrite (r-BLEU) since it is the dominant metric in prior work. Prior reviews of meaning preservation metrics for paraphrase and sentiment ST tasks in EN (Yamshchikov et al., 2021) cover n -gram metrics and embedding-based approaches. We consider three additional metric classes to compare system outputs with inputs, as human annotators do:

1. **n -gram based metrics** include: s-BLEU (self-BLEU that compares system outputs with their inputs as opposed to references, i.e., r-BLEU), METEOR (Banerjee and Lavie, 2005) based on the harmonic mean of unigram precision and recall while accounting for synonym matches, and chrF (Popović, 2015) based on the character n -gram F-score;
2. **embedding-based methods** fall under the category of unsupervised evaluation approaches that rely on either *contextual* word representations extracted from pre-trained language models or *non-contextual pre-trained word embeddings* (e.g., word2vec (Mikolov et al., 2013); Glove (Pennington et al., 2014)). For the former, we use BERT-score (Zhang et al., 2020a) which computes the similarity between each output token and each reference token based on BERT contextual embeddings. For the latter, we experiment with two similarity metrics: the first is the cosine distance between the sentence-level feature representations of the compared texts extracted via averaging their word embeddings; the second is the *Word Mover’s Distance* (WMD) metric

of Kusner et al. (2015) that measures the dissimilarity between two texts as the minimum amount of distance that the embedded words of one text need to “travel” to reach the word embeddings of the other;

3. **semantic textual similarity (STS) models** constitute supervised methods that we model via fine-tuning multilingual pre-trained language models (i.e., XLM-R, mBERT) to predict a semantic similarity score for a pair of texts on an ordinal scale.

Fluency We experiment with **perplexity** (PPL) and **likelihood** (LL) scores based on probability scores of language models trained from scratch (e.g., KenLM (Heafield, 2011)), as well as **pseudo-likelihood scores** (PSEUDO-LL) extracted from pre-trained masked language models similarly to Salazar et al. (2020), by masking sentence tokens one by one.

4 Experiment Settings

Supervised Metrics For all supervised model-based approaches, we experiment with fine-tuning two multilingual pre-trained language models: 1. multilingual BERT, dubbed **mBERT** (Devlin et al., 2019)—a transformer-based model pre-trained with a masked language model objective on the concatenation of monolingual Wikipedia corpora from the 104 languages with the largest Wikipedias. 2. **XLM-R** (Conneau et al., 2020)—a transformer-based masked language model trained on 100 languages using monolingual Common-Crawl data. All models are based on the HuggingFace Transformers (Wolf et al., 2020)² library. We fine-tune with the Adam optimizer (Kingma and Ba, 2015), a batch size of 32, and a learning rate of $5e-5$ for 3 and 5 epochs for classification and regression tasks, respectively. We perform a grid search on held-out validation sets over learning rate with values: $2e-3$, $2e-4$, $2e-5$, and $5e-5$ and over number of epochs with values: 3, 5, and 8.

Cross-lingual Transfer For supervised model-based methods that rely on the availability of human-annotated instances to train dedicated models for specific tasks, we experiment with three standard cross-lingual transfer approaches (e.g., Hu et al. (2020)): 1. ZERO-SHOT trains a single model

²<https://github.com/huggingface/transformers>

DIMENSION	LANGUAGE CODE	DATASET	LINEAGE	LABELS	SIZE
Formality	EN	GYAFC Formality ratings	Rao and Tetreault (2018) Pavlick and Tetreault (2016)	informal vs. formal [-3, -2, -1, 0, 1, 2, 3]	105K 5K
Meaning	EN	STS	Cer et al. (2017)	[1, 2, 3, 4, 5]	5K
Fluency	EN, BR-PT, IT, FR	OpenSubtitles	Lison and Tiedemann (2016)	<i>none</i>	1M

Table 3: Details on training data used for model-based metrics across the three ST evaluation aspects.

on the EN training data and evaluates it on the original test data for each language; 2. TRANSLATE-TRAIN uses machine translation (MT) to obtain training data in each language through translating the EN training set—and trains independent systems for each language; 3. TRANSLATE-TEST trains a single model on the EN training data and evaluates it on the test data that are translated into EN using MT.

Unsupervised Metrics For meaning preservation metrics, we use the open-sourced implementations of: Post (2018) for BLEU (Papineni et al., 2002); Banerjee and Lavie (2005) for METEOR; Popović (2015) for chrF.^{3,4,5} For BERT-score we use the implementation of Zhang et al. (2020a);⁶ non-contextualized embeddings-based approaches are based on fastText pre-trained embeddings.⁷ For fluency metrics, we use the implementation of Salazar et al. (2020) for computing pseudo-likelihood.⁸ PPL and LL scores are extracted from a 5-gram KenLM model (Heafield, 2011).⁹

Training Data Table 3 presents statistics on the training data used for supervised and unsupervised models across the 3 ST evaluation aspects. For datasets that are only available for EN, we use the already available *machine* translated resources for STS¹⁰ and formality datasets (Briakou et al., 2021b). The former employs the DeepL service (no information of translation quality is available) while the latter uses the AWS translation service¹¹ (with reported BLEU scores of 37.16 (BR-

PT), 33.79 (FR), and 32.67 (IT)).¹² The KenLM models for all the languages are trained on 1M randomly sampled sentences from the OpenSubtitles dataset (Lison and Tiedemann, 2016).

5 Experimental Results

We analyze the results of comparing the outputs from the several automatic metrics to their human-generated counterparts for formality style transfer (§5.1), meaning preservation (§5.2), fluency (§5.3) via conducting segment-level analysis—and then, turn into analyzing system-level rankings to evaluation overall task success (§5.4).

5.1 Formality Transfer Metrics

The field is divided on the best way to evaluate the style dimension – formality in our case. Practitioners use either a binary approach (is the new sentence formal or informal?) or a regression approach (how formal is the new sentence?). We discuss the first approach and its limitations in § 5.1.1, before moving to regression in § 5.1.2.

5.1.1 Evaluating Binary Classifiers

As discussed in §2, the vast majority of FoST works evaluate style transfer based on the accuracy of a binary classifier trained to predict whether human-written segments are formal or informal. Yet, as Table 1 indicates, this approach fails to identify the best system in this dimension 59% of the time. To better understand this issue, we evaluate these classifiers on human-written texts versus ST system outputs.

Human Written Texts Table 4 presents F1 scores when testing the binary formality classifiers on the task they are trained on: predicting whether human-written sentences from GYAFC and XFORMAL are formal or informal. First, the last column (i.e., $\delta(\text{XLM-R, mBERT})$) shows that XLM-R is a better model than mBERT for this task, across

¹²BLEU scores were computed on 5,000 randomly sampled data from OpenSubtitles.

³<https://github.com/mjpost/sacrebleu>

⁴<https://www.cs.cmu.edu/~alavie/METEOR/>

⁵<https://github.com/m-popovic/chrF>

⁶https://github.com/Tiiiger/bert_score

⁷<https://fasttext.cc>

⁸<https://github.com/awsmlabs/mlm-scoring>

⁹<https://github.com/kpu/kenlm>

¹⁰<https://github.com/PhilipMay/stsb-multi-mt>

¹¹<https://aws.amazon.com/translate>

METHOD	mBERT				XLM-R				$\delta(\text{XLM-R, mBERT})$			
	EN	BR-PT	FR	IT	EN	BR-PT	FR	IT	EN	BR-PT	FR	IT
ZERO-SHOT	90	87	84	88	90	90	86	89	0	+3	+2	+1
TRANSLATE-TRAIN		85 (-2)	82 (-2)	83 (-5)		87 (-3)	84 (-2)	84 (-5)		+2	+2	+1
TRANSLATE-TEST		82 (-5)	73 (-11)	79 (-11)		82 (-8)	74 (-12)	80 (-9)		0	+2	+1

Table 4: F1 scores of binary formality classifiers under different cross-lingual transfer settings. Numbers in parentheses indicate performance drops over ZERO-SHOT. ZERO-SHOT yields the highest scores across languages and pre-trained language models. XLM-R yields improvements over mBERT across most setting ($\delta(\text{XLM-R, mBERT})$).




Figure 1: Number of formal vs. informal predictions of binary classifiers for each formality bin. Binary classifiers are biased towards outputting formal predictions for BR-PT, FR, and IT, while their performance degrades when moving closer to more neutral formality levels.

METHOD	mBERT				XLM-R				$\delta(\text{XLM-R, mBERT})$			
	EN	BR-PT	FR	IT	EN	BR-PT	FR	IT	EN	BR-PT	FR	IT
ZERO-SHOT	66	64	51	51	67	72	59	61	+1	+8	+8	+10
TRANSLATE-TRAIN		61 (-3)	48 (-3)	49 (-2)		63 (-9)	52 (-7)	55 (-6)		+2	+4	+6
TRANSLATE-TEST		60 (-4)	39 (-12)	45 (-10)		61 (-11)	37 (-15)	51 (-10)		+1	-2	+6

Table 5: Spearman’s ρ correlation (%) of formality regression models. Numbers in parentheses indicate performance drops over ZERO-SHOT. ZERO-SHOT yields the highest scores across languages and pre-trained language models. XLM-R yields improvements over mBERT across most settings ($\delta(\text{XLM-R, mBERT})$).

languages, with the largest improvements in the ZERO-SHOT setting where XLM-R beats mBERT by +3, +2, +1 for BR-PT, FR, and IT respectively.

Second, ZERO-SHOT is surprisingly the best strategy to port EN models to other languages. TRANSLATE-TRAIN and TRANSLATE-TEST hurt F1 by 3 and 9 points on average compared to ZERO-SHOT, despite exploiting more resources in the form of machine translation systems and their training data. However, transfer accuracy is likely affected by regular translation errors (as suggested by larger F1 drops for languages with lower MT BLEU scores) and by formality-specific errors. Machine translation has been found to produce outputs that are more formal than its inputs (Briakou et al., 2021b), which yields noisy training signals for TRANSLATE-TRAIN and alters the formality of test samples for TRANSLATE-TEST.

System Outputs We now evaluate the best performing binary classifier (i.e., XLM-R in ZERO-

SHOT setting) on real system outputs—a setup in line with automatic evaluation frameworks. Figure 1 presents a breakdown of the number of formal vs. informal predictions of the classifiers binned by human-rated formality levels. Across languages, the performance of the classifier deteriorates as we move away from extreme formality ratings (i.e., very informal (-3) and very formal (+3)). This lack of sensitivity to different formality levels is problematic since system outputs across languages are concentrated around neutral formality values. In addition, when testing on BR-PT, FR, and IT (ZERO-SHOT settings), the classifier is more biased towards the formal class, which leads one to question its ability to correctly evaluate more formal outputs in multilingual settings. Taken together, these results suggest that validating the classifiers against human rewrites rather than system outputs is unrealistic and potentially misleading.

	METRIC	EN	IT	BR-PT	FR	
unsupervised	r-BLEU	0.31	0.00	0.05	0.12	
	s-BLEU	0.48	0.59	0.67	0.66	
	METEOR	0.54	0.58	0.64	0.60	
	chrF	0.59	0.67	0.70	0.75	
	WMD	0.51	0.52	0.57	0.55	
	Cosine	0.49	0.50	0.54	0.52	
	BERT-score	0.59	0.62	0.65	0.68	
TRANSLATE-TRAIN (STS)						
supervised	m-BERT	0.53	0.58	0.51	0.63	
	XLM-R	0.55	0.64	0.57	0.60	
	ZERO-SHOT (STS)					
	m-BERT	-	0.61	0.60	0.66	
XLM-R	-	0.67	0.68	0.65		

Table 6: Spearman’s ρ correlation of meaning preservation metrics with human judgments.

5.1.2 Regression Models

Table 5 presents Spearman’s ρ correlation of regression models’ predictions with human judgments. Again, XLM-R with ZERO-SHOT transfer yields the highest correlation across languages. More specifically, the trends across different transfer approaches and different pre-trained language models are similar to the ones observed on evaluation of binary classifiers: XLM-R outperforms mBERT for almost all settings, while ZERO-SHOT is the most successful transfer approach, followed by TRANSLATE-TRAIN, with TRANSLATE-TEST yielding the lowest correlations across languages. Interestingly, regression models highlight the differences between the generalization abilities of XLM-R and mBERT more clearly than the previous analysis on binary predictions: ZERO-SHOT transfer on XLM-R yields 8%, 8%, and 10% higher correlations than mBERT for BR-PT, FR, and IT—while both models yield similar correlations for EN.

5.2 Meaning Preservation Metrics

Table 6 presents Spearman’s ρ correlation of meaning preservation metrics with human judgments. chrF consistently yields the highest correlations across languages—this result is in line with prior observations on evaluating meaning preservation metrics for EN ST tasks (Yamshchikov et al., 2021) and is now confirmed in a multilingual setting. This trend might be explained by chrF’s ability to match spelling errors within words via character n -grams. XLM-R trained on STS with zero-shot trans-

METHOD	EN	IT	PT	FR
KenLM (LL)	0.33	0.27	0.43	0.39
KenLM (PPL)	0.40	0.35	0.45	0.41
mBERT (PSEUDO-LL)	0.42	0.28	0.43	0.41
XLM-R (PSEUDO-LL)	0.50	0.46	0.55	0.61

Table 7: Spearman’s ρ correlation of fluency metrics with human judgments.

fer is a close second to chrF, consistent with this model’s top-ranking behavior as a formality transfer metric. However, chrF outperforms the remaining more complex and expensive metrics, including BERT-score and mBERT models. In contrast to Yamshchikov et al. (2021), embedding-based methods (i.e., cosine, WMD) show no advantage over n -gram metrics, perhaps due to differences in word embedding quality across languages. Finally, it should be noted that r-BLEU is the worst performing metric across languages, and its correlation with human scores is particularly poor for languages other than English. This is remarkable because it has been used in 75% of automatic evaluations for FOST meaning preservation evaluation (as seen in Table 1). We, therefore, recommend discontinuing its use.

5.3 Fluency Metrics

Table 7 presents Spearman’s ρ correlation of various fluency metrics with human judgments. Pseudo-likelihood (PSEUDO-LL) scores obtained from XLM-R correlate with human fluency ratings best across languages. Their correlations are strong across languages, while other methods only yield weak (i.e., KenLM, mBERT) to moderate correlations (i.e., KenLM-PPL) for IT. We, therefore, recommend evaluating fluency using Pseudo-likelihood scores derived from XLM-R to help standardize fluency evaluation across languages.

5.4 System-level Rankings

Finally, we turn to predict the overall ranking of systems by focusing on how many correct pairwise system comparisons each metric gets correct. For each language, there are 5 systems, which means there are 10 pairwise comparisons, for a total of 40 given the 4 languages. We analyze corpus-level r-BLEU, commonly used for this dimension, along with leading metrics from the other dimensions: XLM-R formality regression models, chrF and XLM-R pseudo-likelihood. r-BLEU gets 30 out of 40 comparisons correct while the other metrics get

25, 22, and 19 respectively. This indicates that r-BLEU correlates with human judgments better at the corpus-level than at the sentence-level, as in machine translation evaluation (Mathur et al., 2020). We caution that these results are not definitive but rather suggestive of the best performing metric, given the ideal evaluation would be a larger number of systems with which to perform a rank correlation. The complete analysis for each language is in Appendix B.

6 Conclusions

Automatic (and human) evaluation processes are well-known problems for the field of Natural Language Generation (Howcroft et al., 2020; Clinciu et al., 2021) and the burgeoning subfield of ST is not immune. ST, in particular, has suffered from a lack of standardization of automatic metrics, a lack of agreement between human judgments and automatic metrics, as well as a blindspot to developing metrics for languages other than English. We address these issues by conducting the first controlled multilingual evaluation for leading ST metrics with a focus on formality, covering metrics for 3 evaluation dimensions and overall ranking for 4 languages. Given our findings, we recommend the formality style transfer community adopt the following best practices:

1. Formality XLM-R formality regression models in the ZERO-SHOT cross-lingual transfer setting yields the clear best metrics across all four languages as it correlates very well with human judgments. However, the commonly used binary classifiers do not generalize across languages (due to misleadingly over-predicting formal labels). We propose that the field use regression models instead since they are designed to capture a wide spectrum of formality rates.

2. Meaning Preservation We recommend using chrF as it exhibits strong correlations with human judgments for all four languages. We caution against using BLEU for this dimension, despite its overwhelming use in prior work as both its reference and self variants do not correlate as strongly as other more recent metrics.

3. Fluency XLM-R is again the best metric (in particular for French). However, it does not correlate well with human judgments as compared to the other two dimensions.

4. System-level Ranking chrF and XLM-R are the best metrics using a pairwise comparison evaluation. However, an ideal evaluation would be to have a large number of systems with which to draw reliable correlations.

5. Cross-lingual Transfer Our results support using ZERO-SHOT transfer instead of machine translation to port metrics from English to other languages for formality transfer tasks.

We view this work as a strong point of departure for future investigations of ST evaluation. Our work first calls for exploring how these evaluation metrics generalize to other *styles* and *languages*. Across the different ways of defining style evaluation (either automatic or human), prior work has mostly focused on the three main dimensions covered in our study. As a result, although our meta-evaluation on ST metrics focuses on formality as a case study, it can inform the evaluation of other style definitions (e.g., politeness, sentiment, gender, etc.). However, more empirical evidence is needed to test the applicability of the best performing metrics for evaluating style transfer beyond formality. Our work suggests that the top metrics based on XLM-R and chrF are robust across 4 Romance languages; yet, our conclusions and recommendations are currently limited to *this* set of languages. We hope that future work in multilingual style transfer will allow for testing their generalization to a broader spectrum of languages and style definitions. Furthermore, our study highlights that more research is needed on automatically ranking systems. For example, one could build a metric that combines metrics' outputs for the three dimensions, or one could develop a singular metric. In line with Briakou et al. (2021a), our study also calls for releasing more human evaluations and more system outputs to enable robust evaluation. Finally, there is still room for improvement in assessing how fluent a rewrite is. Our study provides a framework to address these questions systematically and calls for ST papers to standardize and release data to support larger-scale evaluations.

Acknowledgements

We thank Sudha Rao for providing references and materials of the GYAF dataset, Jordan Boyd-Graber, Pedro Rodriguez, the CLIP lab at UMD, and the EMNLP reviewers for their helpful and constructive comments.

References

- Eneko Agirre, Carmen Banea, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, Rada Mihalcea, German Rigau, and Janyce Wiebe. 2016. [SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 497–511, San Diego, California. Association for Computational Linguistics.
- 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.
- Eleftheria Briakou, Sweta Agrawal, Ke Zhang, Joel Tetreault, and Marine Carpuat. 2021a. [A review of human evaluation for style transfer](#). In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 58–67, Online. 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.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Yu Cheng, Zhe Gan, Yizhe Zhang, Oussama Elachqar, Dianqi Li, and Jingjing Liu. 2020. [Contextual text style transfer](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2915–2924, Online. Association for Computational Linguistics.
- Miruna-Adriana Clinciu, Arash Eshghi, and Helen Hastie. 2021. [A study of automatic metrics for the evaluation of natural language explanations](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2376–2387, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, 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.
- Hongyu Gong, Suma Bhat, Lingfei Wu, JinJun Xiong, and Wen-mei Hwu. 2019. [Reinforcement learning based text style transfer without parallel training corpus](#). 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 3168–3180, Minneapolis, Minnesota. Association for Computational Linguistics.
- Navita Goyal, Balaji Vasan Srinivasan, Anandhavelu Natarajan, and Abhilasha Sancheti. 2020. [Multi-dimensional style transfer for partially annotated data using language models as discriminators](#). *CoRR*, abs/2010.11578.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. [A probabilistic formulation of unsupervised text style transfer](#). In *International Conference on Learning Representations*.
- Kenneth Heafield. 2011. [KenLM: Faster and smaller language model queries](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197, Edinburgh, Scotland. Association for Computational Linguistics.
- Michael Heilman, Aoife Cahill, Nitin Madnani, Melissa Lopez, Matthew Mulholland, and Joel Tetreault. 2014. [Predicting grammaticality on an ordinal scale](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 174–180, Baltimore, Maryland. Association for Computational Linguistics.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson.

2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2020. [Deep learning for text style transfer: A survey](#). *CoRR*, abs/2011.00416.
- Tomoyuki Kajiwara. 2019. [Negative lexically constrained decoding for paraphrase generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6047–6052, Florence, Italy. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Elizaveta Korotkova, Agnes Luhtaru, Maksym Del, Krista Liin, Daiga Deksnė, and Mark Fishel. 2019. [Grammatical error correction and style transfer via zero-shot monolingual translation](#). *CoRR*, abs/1903.11283.
- 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 International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 957–966, Lille, France. PMLR.
- Shibamouli Lahiri. 2015. [Squinky! A corpus of sentence-level formality, informativeness, and implicature](#). *CoRR*, abs/1506.02306.
- Dianqi Li, Yizhe Zhang, Zhe Gan, Yu Cheng, Chris Brockett, Bill Dolan, and Ming-Ting Sun. 2019. [Domain adaptive 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 3304–3313, Hong Kong, China. Association for Computational Linguistics.
- Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael Lyu, and Irwin King. 2020. [Unsupervised text generation by learning from search](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 10820–10831. Curran Associates, Inc.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Yixin Liu, Graham Neubig, and John Wieting. 2021. [On learning text style transfer with direct rewards](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4262–4273, Online. Association for Computational Linguistics.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Xu Sun, and Zhifang Sui. 2019. [A dual reinforcement learning framework for unsupervised text style transfer](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5116–5122. International Joint Conferences on Artificial Intelligence Organization.
- 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.
- 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.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- 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.
- Richard Yuanzhe Pang. 2019. [Towards actual \(not operational\) textual style transfer auto-evaluation](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 444–445, Hong Kong, China. Association for Computational Linguistics.
- Richard Yuanzhe Pang and Kevin Gimpel. 2019. [Unsupervised evaluation metrics and learning criteria for non-parallel textual transfer](#). In *Proceedings of*

- the 3rd Workshop on Neural Generation and Translation, pages 138–147, Hong Kong. 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.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Sudha Rao and Joel Tetreault. 2018. [Dear sir or madam, may I introduce the GYAF 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.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Abhilasha Sancheti, Kundan Krishna, Balaji Vasanth Srinivasan, and Anandhavelu Natarajan. 2020. Reinforced rewards framework for text style transfer. In *Advances in Information Retrieval*, pages 545–560, Cham. Springer International Publishing.
- Mingyue Shang, Piji Li, Zhenxin Fu, Lidong Bing, Dongyan Zhao, Shuming Shi, and Rui Yan. 2019. [Semi-supervised text style transfer: Cross projection in latent space](#). 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 4937–4946, 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.
- Yunli Wang, Yu Wu, Lili Mou, Zhoujun Li, and Wenhao Chao. 2019. [Harnessing pre-trained neural networks with rules for formality 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 3573–3578, Hong Kong, China. Association for Computational Linguistics.
- Yunli Wang, Yu Wu, Lili Mou, Zhoujun Li, and Wenhao Chao. 2020. [Formality style transfer with shared latent space](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2236–2249, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. 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.
- Yu Wu, Yunli Wang, and Shujie Liu. 2020. [A dataset for low-resource stylized sequence-to-sequence generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9290–9297.
- Ruochen Xu, Tao Ge, and Furu Wei. 2019. [Formality style transfer with hybrid textual annotations](#). *CoRR*, abs/1903.06353.
- Ivan P. Yamshchikov, Viacheslav Shibaev, Nikolay Khlebnikov, and Alexey Tikhonov. 2021. [Style-transfer and paraphrase: Looking for a sensible semantic similarity metric](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):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. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Yi Zhang, Tao Ge, and Xu Sun. 2020b. [Parallel data augmentation for formality style transfer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3221–3228, Online. Association for Computational Linguistics.
- 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 Cross-metric Correlation Analysis

Correlations across meaning preservation metrics Figure 3 presents a cross-metric correlation-based analysis of the different approaches for measuring meaning preservation. We observe consistent trends across languages: methods that are similar in nature correlate well with each other. Concretely, across settings, n -gram based methods (i.e., BLEU, METEOR, and chrF) yield 0.8 – 0.9 correlation scores. The latter also holds when looking at correlations within the group of embedding-based methods (cosine and WMD) and and group of STS approaches for EN, FR, and IT, while for BR-PT we observe that the correlation between XLM-R and mBERT based approaches is smaller (0.7 vs. 0.8 for other languages). Finally, n -gram approaches correlate better with STS methods (with correlations in the range of 0.7 – 0.8) across languages, while the lowest correlations (0.5–0.6) are observed between embedding-based methods (i.e., cosine, WMD) and each of the rest metrics.

Correlations within and across formality-fluency metrics Figure 4 presents results of cross-metric correlations for the studied approaches that capture formality transfer and fluency. For formality, each of the translate-based approaches (i.e., TRANSLATE-TRAIN and TRANSLATE-TEST) yields high correlations (0.8 – 0.9) between models that fine-tune XLM-R vs. mBERT, while their correlations decrease (0.7) for IT and BR-PT in the zero-shot setting. Finally, pseudo-perplexity metrics extracted from XLM-R—that consists the best correlated metric with human judgments for fluency—yield positive correlations with all formality metrics.

B System-level Analysis

Table 8 presents the number of correct system-level pair-wise comparisons of automatic metrics based on human judgments. For STS, chrF, F.REG*, F.CLASS*, and PSEUDO-LKL*, system-level scores are extracted via averaging sentence-level scores. For s-BLEU and r-BLEU the system scores are extracted at the corpus-level. The total number of pair-wise comparisons for each language is 10 (given access to 5 systems). Among the meaning preservation metrics (i.e., STS, s-BLEU, and chrF), chrF yields the highest number of correct comparisons (i.e., 37 out of 40 for all languages). The formality regression models (i.e., F.REG*) result in correct

rankings more frequently than the formality classifiers (i.e., F.CLASS*) yielding 35 out of 40 correct comparisons. Reference-BLEU (i.e., r-BLEU) is compared with overall ranking judemnts. It ranks 8 out of 10 systems correctly for EN, FR, and BR-PT and only 6 for IT. Finally, perplexity (i.e., PPL) results in the fewest correct rankings at system-level (i.e., 22 out of 40), despite correlating well with human judgments at the segment-level.

Additionally, in Figure 2 we visualize the differences between relative rankings induced by human judgments and the best segment-level correlated metrics for each dimension, averaged per system.

C Evaluated Systems Details

For each of BR-PT, IT, and FR, outputs are sampled from:

1. Rule-based systems consisting of hand-crafted transformations (e.g., fixing casing, normalizing punctuation, expanding contractions, etc.);
2. Round-trip translation models that pivot to EN and backtranslate to the original language;
3. Bi-directional neural machine translation (MT) models that employ side constraints to perform style transfer for both directions of formality (i.e., informal \leftrightarrow formal)—trained on (machine) translated informal-formal pairs of an English parallel corpus (i.e., GYAFC);
4. Bi-directional NMT models that augment the training data of 3. via backtranslation of informal sentences;
5. A multi-task variant of 3. that augments the training data with parallel-sentences from bilingual resources (i.e., OpenSubtitles) and learns to translate jointly between and across languages.

For EN, the outputs were sampled from:

1. A rule-based system of similar transformations to ones for BR-PT, FR, and IT;
2. A phrase-based machine translation model trained on informal-formal pairs of GYAFC;
3. An NMT model trained on GYAFC to perform style transfer uni-directionally;

Language	MEANING			FORMALITY		FLUENCY		OVERALL			
	STS	s-BLEU	r-BLEU	chrF	F.REG*	F.CLASS*	PPL	r-BLEU	F.REG*	chrF	pseudo-LL
EN	6	7	8	7	9	9	7	8	6	6	3
BR-PT	10	9	7	9	7	7	5	8	8	4	5
FR	9	7	10	10	10	9	5	8	6	3	7
IT	10	10	8	10	9	8	5	6	5	9	4
ALL	35	33	35	36	35	33	22	30	25	22	19

Table 8: Number of correct system level pair-wise comparisons between 5 systems for each language.



Figure 2: Difference in relative ranking between human judgments and automatic metrics across systems (i.e., represented by different markers) for different evaluation dimensions. STS, s-BLEU and chrF are compared with meaning rankings, r-BLEU (reference-BLEU) with overall, XLM-R classifiers (*F.CLASS) and regression (*F.REF) models with formality, and XLM-R pseudo-perplexity (*PPL) with fluency.

4. A variant of 3. that incorporates a copy-enriched mechanism that enables direct copying of words from input;
5. A variant of 4. trained on additional back-translated data of target style sentences using 2.

In general, neural models performed best for all languages according to overall human judgments, while the simpler baselines perform closer to the more advanced neural models for BR-PT, FR, and IT. For each evaluation dimension 500 outputs are evaluated for EN and 100 outputs per system for BR-PT, FR, and IT.

D Meaning Preservation Metrics (reference-based)

Table 9 presents supplemental results on meaning preservation metrics for reference-based settings.

METRIC	EN	IT	PT	FR
r-BLEU	0.306	0.004	0.047	0.122
METEOR	0.279	-0.005	0.061	0.124
chrF	0.319	0.065	0.044	0.174
WMD	0.316	0.039	0.098	0.198
Cosine	0.218	0.027	0.048	0.161
BERT-score	0.359	-0.023	0.054	0.112
mBERT (TRANSLATE-TRAIN)	0.369	0.077	0.167	0.165
mBERT (ZERO-SHOT)	-	0.124	0.197	0.179
XLM-R (TRANSLATE-TRAIN)	0.385	0.183	0.136	0.259
XLM-R (ZERO-SHOT)	-	0.179	0.153	0.258

Table 9: Spearman’s ρ correlation of meaning preservation metrics for reference-based meaning.

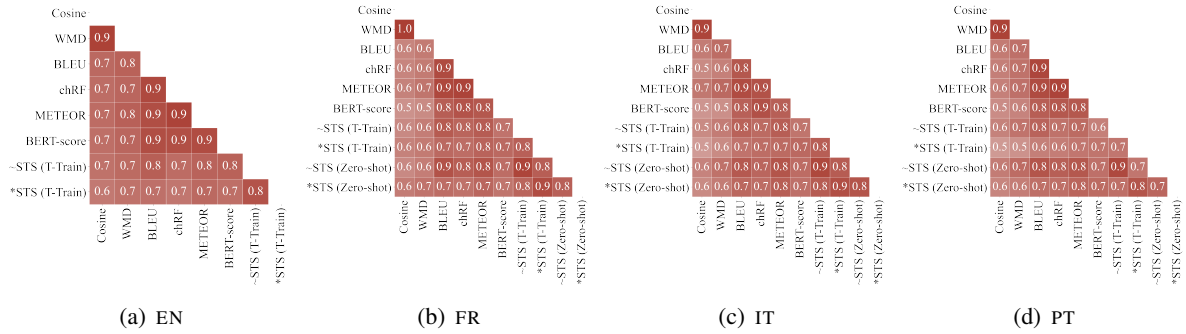


Figure 3: Cross-metric correlations for meaning preservation automatic evaluation metrics. Mode-based metrics marked with * use XLM-R while markers ~ use mBERT as the base pre-trained language model.

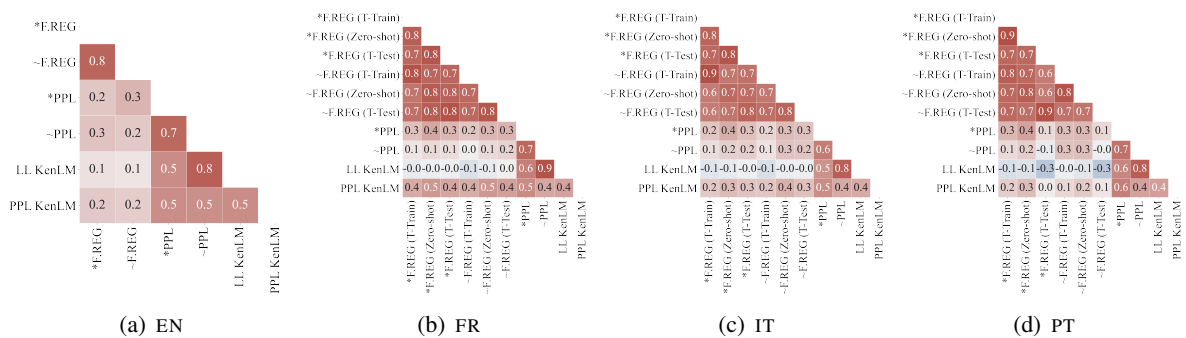


Figure 4: Cross-metric correlations for fluency and formality automatic evaluation metrics. Mode-based metrics marked with * use XLM-R while markers ~ use mBERT as the base pre-trained language model. F.REG refers to formality regression models, PPL to perplexity, and LL to likelihood.