

# Olá, Bonjour, Salve!

## XFORMAL: A Benchmark for Multilingual Formality Style Transfer

**Eleftheria Briakou\***  
University of Maryland  
ebriakou@cs.umd.edu

**Di Lu**  
Dataminr, Inc.  
dlu@dataminr.com

**Ke Zhang**  
Dataminr, Inc.  
kzhang@dataminr.com

**Joel Tetreault**  
Dataminr, Inc.  
jtetreault@dataminr.com

### Abstract

We take the first step towards multilingual style transfer by creating and releasing XFORMAL, a benchmark of multiple formal reformulations of informal text in Brazilian Portuguese, French, and Italian. Results on XFORMAL suggest that state-of-the-art style transfer approaches perform close to simple baselines, indicating that style transfer is even more challenging when moving multilingual.<sup>1</sup>

### 1 Introduction

Style Transfer (ST) is the task of automatically transforming text in one style into another (for example, making an impolite request more polite). Most work in this growing field has focused primarily on style transfer within English, while *covering different languages* has received disproportional interest. Concretely, out of 35 ST papers we reviewed, all of them report results for ST within English text, while there is just a single work covering each of the following languages: Chinese, Russian, Latvian, Estonian, and French (Shang et al., 2019; Tikhonov et al., 2019; Korotkova et al., 2019; Niu et al., 2018). Notably, even though some efforts have been made towards multilingual ST, researchers are limited to providing system outputs as a means of evaluation, and progress is hampered by the scarcity of resources for most languages. At the same time, ST lies at the core of human communication: when humans produce language, they condition their choice of grammatical and lexical transformations to a target audience and a specific situation. Among the many possible stylistic variations, Heylighen et al. (1999) argue that “*a dimension similar to formality appears as the most important and universal feature distinguishing styles, registers or genres in different languages*”. Consider the informal excerpts and their formal reformulations in French (FR) and Brazilian Portuguese

#### BRAZILIAN-PORTUGUESE

saíam disso, força de vontade!!  
*get out of it, willpower!!*

Abandonem essa situação, tenham força de vontade.  
*Abandon this situation, have willpower!*

#### FRENCH

Il avait les yeux braqués ailleurs.  
*He had his eyes fixed elsewhere.*

Il ne prêtait pas attention à la situation.  
*He was not paying attention to the situation.*

#### ITALIAN

in bocca al lupo!  
*good luck!*

Ti rivolgo un sincero augurio!  
*I send you a sincere wish!*

Table 1: Informal-Formal pairs in XFORMAL.

(BR-PT) in Table 1. Both informal-formal pairs share the same content. However, the informal language conveys more information than is contained in the literal meaning of the words (Hovy, 1987). These examples relate to the notion of *deep formality* (Heylighen et al., 1999), where the ultimate goal is that of adding the context needed to disambiguate an expression. On the other hand, variations in formality might just reflect different situational and personal factors, as shown in the Italian (IT) example.

This work takes the first step towards a more language-inclusive direction for the field of ST by building the first corpus of style transfer for non-English languages. In particular, we make the following contributions: **1.** Building upon prior work on Formality Style Transfer (FoST) (Rao and Tetreault, 2018), we contribute an **evaluation dataset**, XFORMAL that consists of multiple formal rewrites of informal sentences in three Romance languages: Brazilian Portuguese (BR-PT), French (FR), and Italian (IT); **2.** Without assuming access to any gold-standard training data for the languages at hand, we **benchmark a myriad**

\*Work done as a Research Intern at Dataminr, Inc.

<sup>1</sup>Code and data: <https://github.com/Elbria/xformal-FoST>

**of leading ST baselines** through automatic and human evaluation methods. Our results show that FoST in non-English languages is particularly challenging as complex neural models perform on par with a simple rule-based system consisting of hand-crafted transformations. We make XFORMAL, our annotations protocols, and analysis code publicly available and hope that this study facilitates and encourages more research towards Multilingual ST.

## 2 Related Work

Controlling style aspects in generation tasks is studied in monolingual settings with an English-centric focus (intra-language) and cross-lingual settings together with Machine Translation (MT) (inter-language). Our work rests in intra-language ST with a multilingual focus, in contrast to prior work.

**ST datasets** that consist of parallel pairs in different styles include: GYAFC for formality (Rao and Tetreault, 2018), Yelp (Shen et al., 2017) and Amazon Product Reviews for sentiment (He and McAuley, 2016), political slant and gender controlled datasets (Prabhumoye et al., 2018), Expert Style Transfer (Cao et al., 2020), PASTEL for imitating personal (Kang et al., 2019), SIMILE for simile generation (Chakrabarty et al., 2020), and others.

**Intra-language ST** was first cast as generation task by Xu et al. (2012) and is addressed through methods that use either parallel data or unpaired corpora of different styles. **Parallel corpora designed for the task at hand** are used to train traditional encoder-decoder architectures (Rao and Tetreault, 2018), learn mappings between latent representation of different styles (Shang et al., 2019), or fine-tune pre-trained models (Wang et al., 2019). Other approaches use **parallel data from similar tasks** to facilitate transfer in the target style via domain adaptation (Li et al., 2019), multi-task learning (Niu et al., 2018; Niu and Carpuat, 2020), and zero-shot transfer (Korotkova et al., 2019) or create **pseudo-parallel data** via data augmentation techniques (Zhang et al., 2020; Krishna et al., 2020). Approaches that rely on **non-parallel data** include *disentanglement methods* based on the idea of learning style-agnostic latent representations (e.g., Shen et al. (2017); Hu et al. (2017)). However, they are recently criticized for resulting in poor content preservation (Xu et al., 2018; Jin et al., 2019; Luo et al., 2019; Subramanian et al., 2018) and alternatively, *translation-based models* are proposed

that use reconstruction and back-translation losses (e.g., Logeswaran et al. (2018); Prabhumoye et al. (2018)). Another line of work, focuses on *manipulation methods* that remove the style-specific attribute of text (e.g., Li et al. (2018); Xu et al. (2018)), while recent approaches use *reinforcement learning* (e.g., Wu et al. (2019); Gong et al. (2019)), probabilistic formulations (He et al., 2020), and masked language models (Malmi et al., 2020).

**Inter-language ST** is introduced by Mirkin and Meunier (2015) who proposed personalized MT for EN-French and EN-German. Subsequent MT works control for politeness (Sennrich et al., 2016a), voice (Yamagishi et al., 2016), personality traits (Rabinovich et al., 2017), user-provided terminology (Hasler et al., 2018), gender (Vanmassenhove et al., 2018), formality (Niu et al., 2017; Feely et al., 2019), morphological variations (Moryossef et al., 2019), complexity (Agrawal and Carpuat, 2019) and reading level (Marchisio et al., 2019).

## 3 XFORMAL Collection

We describe the process of collecting formal rewrites using data statements protocols (Bender and Friedman, 2018; Gebu et al., 2018).

**Curation rational** To collect XFORMAL, we firstly curate informal excerpts in multiple languages. To this end, we follow the procedures described in Rao and Tetreault (2018) (henceforth RT18) who create a corpus of informal-formal sentence-pairs in English (EN) entitled Grammarly’s Yahoo Answers Formality Corpus (GYAFC).

Concretely, we use the L6 Yahoo! Answers corpus that consists of questions and answers posted to the Yahoo! Answers platform.<sup>2</sup> The corpus contains a large number of informal text and allows control for different languages and different domains.<sup>3</sup> Similar to the collection of GYAFC, we extract all answers from the Family & Relationships (F&R) topic that correspond to the three languages of interest: *Família e Relacionamentos* (BR-PT), *Relazioni e famiglia* (IT), and *Amour et relations* (FR) (**Step 1**). We follow the same pre-processing steps as described in RT18 for consistency (**Step 2**). We filter out answers that: a) consist of questions; b) include URLs; c) have fewer

<sup>2</sup><https://webscope.sandbox.yahoo.com/catalog.php?datatype=1&did=11>

<sup>3</sup>More details are included under A.F.

Corpus	BR-PT	FR	IT
L6 Yahoo!	230,302	225,476	101,894
+ <b>Step 1</b>	37,356	34,849	13,443
+ <b>Step 2</b>	14,448	14,856	4,095
+ <b>Step 3</b>	8,617	11,118	2,864

Table 2: Number of sentences in filtered versions of the L6 Yahoo! Corpus across curation steps and languages.

than five or more than 25 tokens; or d) constitute duplicates.<sup>4</sup> We automatically extract informal candidate sentences, as described in §5.3 (**Step 3**). Finally, we randomly sample 1,000 sentences from the pool of informal candidates for each language. Table 2 presents statistics of the curation steps.

**Procedures** We use the Amazon Mechanical Turk (MTurk) platform to collect formal rewrites for our informal sentences. For each language, we split the annotation into 20 batches of 50 Human Intelligence Tasks (HITS). In each HIT, Turkers are given an informal excerpt and asked to *generate its formal rewrite in the same language without changing its meaning*. We collect 4 rewrites per excerpt and release detailed instructions under A.G.

**Annotation Workflow & Quality Control** Our annotation protocol consists of multiple Quality Control (QC) steps to ensure the recruitment of high-quality annotators. As a first step, we use **location restrictions (QC1)** to limit the pool of workers to countries where native speakers are most likely to be found. Next, we run several small *pilot studies* (of 10 HITS) to recruit potential workers. To participate in the pilot study, Turkers have to pass a **qualification test (QC2)** consisting of multiple-choice questions that test workers’ understanding of formality (see A.L). The pilot study results are **reviewed by a native speaker (QC3)** of each language to exclude workers who performed consistently poorly. We find that the two main reasons for poor quality are: a) rewrites of minimum-level edits, or b) rewrites that change the input’s meaning. Table 3 presents the number of workers at each QC step. Only workers passing all quality control steps (last row of Table 3) contribute to the final task. Finally, we **post-process** the collected rewrites by a) removing instances consisting of normalization-based edits only and b) correcting minor spelling errors using an off-the-shelf tool.<sup>5</sup>

<sup>4</sup>We tokenize with nltk: <https://www.nltk.org/api/nltk.tokenize.html>

<sup>5</sup><https://languagetool.org/>

Step	Description	BR-PT	FR	IT
(QC1)	Location restriction	151	78	59
(QC2)	Qualification test	54	40	33
(QC3)	Rewrites review	<b>9</b>	<b>16</b>	<b>11</b>

Table 3: Number of Turkers **after** each QC step.

**Turkers’ demographics** We recruit Turkers from Brazil, France/Canada, and Italy for BR-PT, FR, and IT, respectively. Beyond their country of residence, no further information is available.

**Compensation** We compensate at a rate of \$0.10 per HIT with additional one-time bonuses that bumps them up to a target rate of over \$10/hour.

After this entire process, we have constructed a high-quality corpus of formality rewrites of 1,000 sentences for three languages. In the next section, we provide statistics and an analysis of XFORMAL.

## 4 XFORMAL Statistics & Analysis

**Types of formality edits** Following Pavlick and Tetreault (2016), we analyze the most frequent **edit** operations Turkers perform when formalizing the informal sentences. We conduct both an **automatic analysis** (details in A.I) of the whole set of rewrites, and a **human analysis** (details in A.H) of a random sample of 200 rewrites per language (we recruited a native speaker for each language). Table 4 presents both analyses’ results, where we also include the corresponding statistics for the English language (GYAFC). In general, we observe similar trends across languages: humans make edits covering both the "noisy-text" sense of formality (e.g., fixing punctuation, spelling errors, capitalization) and the more situational sense (paraphrase-based edits). Although cross-language trends are similar, we also observe differences: *deleting fillers* and word *completion* seems to be more prominent in the English rewrites than in other languages; *normalizing* abbreviations is a considerably frequent edit type for Brazilian Portuguese; *paraphrasing* is more frequent in the three non-English languages.

**Surface differences of informal-formal pairs** We quantify surface-level differences between the informal sentences and formal rewrites via computing their character-level Levenshtein distance (Figure 1) and their pairwise Lexical Difference (LeD) based on the percentages of **tokens** that are *not* found in both sentences (Table 5). Both analyses show that Italian rewrites have the most edits com-




Figure 1: Distribution of character-based Levenshtein distance between informal sentences and formal rewrites.




Figure 2: Number of informal sentences and formal rewrites binned according to formality score.

EDIT TYPES	BR-PT	FR	IT	EN
CAPITALIZATION	0.24	0.33	0.32	0.43
	0.43	0.56	0.60	0.46
PUNCTUATION	0.64	0.87	0.79	0.74
	0.76	0.64	0.66	0.40
SPELLING	0.36	0.38	0.32	0.29
	0.09	0.12	0.10	0.14
NORMALIZATION	0.17	0.01	0.04	0.07
	0.27	0.02	0.06	0.10
SPLIT SENTENCES	0.07	0.06	0.11	0.08
	0.07	0.07	0.07	0.04
PARAPHRASE	0.68	0.75	0.86	0.68
	0.74	0.75	0.58	0.47
UNCHANGED	0.00	0.00	0.00	0.01
LOWERCASE	0.13	0.06	0.07	0.13
REPETITION	0.01	0.01	0.02	0.09
DELETE FILLERS	0.02	0.05	0.11	0.26
COMPLETION	0.07	0.04	0.04	0.15
ADD CONTEXT	0.02	0.00	0.07	
CONTRACTIONS	0.02	0.01	0.01	
SLANG/IDIOMS	0.19	0.10	0.15	
POLITENESS	0.82	0.00	0.10	
RELATIVIZERS	0.01	0.00	0.00	

Table 4: Percentage categories of frequent edits calculated **automatically** (not highlighted) and **manually** (highlighted). Categories are not mutually exclusive.

pared to their corresponding informal sentences. French and Brazilian Portuguese follow, with English rewrites being closer to the informal inputs.

**Diversity of formal rewrites** Are Turkers making similar choices when formalizing text? Since

METRIC	BR-PT	FR	IT	EN
LeD	0.44	0.39	0.54	0.36
self-BLEU	0.43	0.35	0.21	0.51

Table 5: Surface differences **between** informal-formal pairs (LeD) and **within** formal rewrites (self-BLEU).

a large number of reformulations consist of paraphrase-based edits (more than 50%), we want to quantify the extent to which the formal rewrites of each sentence are **diverse**, in terms of their lexical choices. To that end, we quantify diversity via measuring self-BLEU (Zhu et al., 2018): considering one set of formal sentences as the hypothesis set and the others as references, we compute BLEU for each formal set and define the average BLEU score as a measure of the dataset’s diversity. Higher scores imply less diversity of the set. Results (last row of Table 4) show that XFORMAL consists of more diverse rewrites compared to GYAFC.

**Formality shift of rewrites** We analyze the formality distribution of the original informal sentences with their formal rewrites in GYAFC and XFORMAL, as predicted by formality mBERT models (§5.3). The distributions of formal rewrites are skewed towards positive values (Figure 2).

## 5 Multilingual FoST Experiments

We benchmark eight ST models on XFORMAL to serve as baseline scores for future research. We describe the models (§5.1), the experimental setting (§5.2), the human and automatic evaluation methods (§5.3 and §5.4), and results (§5.5).

## 5.1 Models

**Simple baselines** We define three baselines: 1. **COPY** Motivated by Pang and Gimpel (2019) who notice that untransferred sentences with no alterations have the highest BLEU score by a large margin for ST tasks, we use this simple baseline as a lower bound; 2. **RULE-BASED** Based on the quantitative analysis of §4 and similarly to RT18, we develop a rule-based approach that performs a set of predefined edits-based operations defined by handcrafted rules. Example transformations include fix casing, remove repeated punctuation, handcraft a list of contraction expansions—a detailed description is found at A.C; 3. **ROUND-TRIP MT** Inspired by Zhang et al. (2020) who identify useful training pairs from the paraphrases generated by round-trip translations of millions of sentences, we devise a simpler baseline that starts from a text in language  $x$ , pivots to EN and then backtranslates to  $x$ , using the AWS translation service.<sup>6</sup>

### NMT-based models with synthetic parallel data

We follow the TRANSLATE TRAIN (Conneau et al., 2018; Artetxe et al., 2020) approach to collect data in multilingual settings: we obtain *pseudo-parallel* corpora in each language via machine translating an EN resource of informal-formal pairs (§5.2).<sup>7</sup> Then, starting with TRANSLATE TRAIN we benchmark the following NMT-based models: 1. **TRANSLATE TRAIN TAG** extends a leading EN FoST approach (Niu et al., 2018) and trains a unified model that handles either formality direction via attaching a source tag that denotes the desired target formality; 2. **MULTI-TASK TAG-STYLE** Niu et al. (2018) augments the previous approach with bilingual data that is automatically identified as formal (§5.3). The models are then trained in a multi-task fashion; 3. **BACKTRANSLATE** augments the TRANSLATE TRAIN data with back-translated sentences of automatically detected informal text (Sennrich et al., 2016b), using 1. as the base model. We exclude backtranslated pairs consisting of copies. The output of the RULE-BASED system is given as input to each model at inference time. For all three models, we run each system with 4 random seeds, and combine them in a linear ensemble for decoding.

**Unsupervised approaches** We benchmark two unsupervised methods that are used for EN ST: 1. **UNSUPERVISED NEURAL MACHINE TRANS-**

<sup>6</sup><https://aws.amazon.com/translate/>

<sup>7</sup>Details on the AWS performance are found in A.B.

**LATION (UNMT)** (Subramanian et al., 2018) defines a pseudo-supervised setting and combines denoising auto-encoding and back-translation losses; 2. **DEEP LATENT SEQUENCE MODEL (DLSM)** (He et al., 2020) defines a probabilistic generative story that treats two unpaired corpora of separate styles as a partially observed parallel corpus and learns a mapping between them, using variational inference.

## 5.2 Experimental setting

**Training data** For TRANSLATE TRAIN TAG we use GYAFC, a large set of 110K EN informal-formal parallel sentence-pairs obtained through crowdsourcing. Additionally, we augment the translated resource with OpenSubtitles (Lison and Tiedemann, 2016) bilingual data used for training MT models.<sup>8</sup> Given that bilingual sentence-pairs can be noisy, we perform a filtering step to extract noisy bitexts using the Bicleaner toolkit (Sánchez-Cartagena et al.).<sup>9</sup> Furthermore, we apply the same filtering steps as in §3 (Curation rational). Finally, each of the remaining sentences is assigned a formality score (§5.3), resulting in two pools of informal and formal text. Training instances are then randomly sampled from those pools: formal parallel pairs are used for MULTI-TASK TAG-STYLE; informal target side sentences are backtranslated for BACKTRANSLATE; both informal and formal target-side texts are independently sampled from the two pools for training unsupervised models. Finally, for unsupervised FoST in FR, we additionally experiment with in-domain data from the L26 French Yahoo! Answer Corpus that consists of 1.7M FR questions.<sup>10, 11</sup> Table 6 includes statistics on training sizes.<sup>12</sup>

METHODS	GYAFC	OpenSubs.	L6 Yahoo!
TRANSL. TRAIN TAG	110K	–	–
MULTI-TASK TAG-STYLE	110K	2M	–
BACKTRANSLATE	110K	2M	–
UNMT/DLSM	–	110K	–
UNMT/DLSM (IN-DOMAIN)	–	–	2M

Table 6: Number of training instances for each model.

<sup>8</sup>Data are available at: <http://opus.nlpl.eu/>.

<sup>9</sup>We use the publicly available pretrained Bicleaner models: <https://github.com/bitextor/bicleaner>, and discard sentences with a score lower than 0.5.

<sup>10</sup><https://webscope.sandbox.yahoo.com/catalog.php?datatype=l&did=74>

<sup>11</sup>Split into 6.2M/6.6M formal/informal sentences.

<sup>12</sup>Bilingual data statistics are in A.K.

**Preprocessing** We preprocess data consistently across languages using MOSES (Koehn et al., 2007). Our pipeline consists of three steps: a) normalization; b) tokenization; c) true-casing. For NMT-based approaches, we also learn joint source-target BPE with 32K operations (Sennrich et al., 2016b).

**Model Implementations** For NMT-based and unsupervised models we use the open-sourced implementations of Niu et al. (2018) and He et al. (2020), respectively.<sup>13,14</sup> We include more details on model architectures in A.D.

### 5.3 Automatic Evaluation

Recent work on ST evaluation highlights the **lack of standard evaluation practices** (Yamshchikov et al., 2020; Pang, 2019; Pang and Gimpel, 2019; Mir et al., 2019). We follow the most frequent evaluation metrics used in EN tasks and measure the quality of the system’s outputs with respect to four dimensions, while we leave an extensive evaluation of automatic metrics for future work.

**Meaning Preservation** We compute self-BLEU (Papineni et al., 2002) which compares system outputs with the informal sentences.

**Formality** We average the style transfer score of transferred sentences computed by a formality regression model. We fine-tune mBERT (Devlin et al., 2019) pre-trained language models on the machine-translated answers genre from Pavlick and Tetreault (2016) that consists of about 4K human-annotated sentences rated on a 7-point formality scale. To acquire an annotated corpus in the languages of interest, we follow the TRANSLATE TRAIN transfer approach: we propagate the original EN training data’s human ratings to their corresponding translations, assuming that translation preserves formality.<sup>15</sup> To evaluate the multilingual formality regression models’ performance, we crowdsourced human judgments of 5 Turkers for 200 sentences per language. We report Spearman correlations of 64 (BR-PT), 70 (IT), 71 (FR), and 81 (EN).

**Fluency** We compute the logarithm of each sentence’s probability—computed by a 5-gram Kneser-Ney language model (Kneser and Ney, 1995)—and normalize it by the sequence length. We train each

language model on 2M random sample of the non-English side of OpenSubtitles formal data.

**Overall** We compute multi-BLEU (Post, 2018) via comparing with multiple formal rewrites on XFORMAL. Freitag et al. (2020) shows that correlation with human judgments improves when considering multiple references for MT evaluation.

### 5.4 Human evaluation

Given that automatic evaluation of ST lacks standard evaluation practices—even in cases when EN is considered—we turn to human evaluation to reliably assess our baselines following the protocols of RT18. We sample a subset of 100 sentences from XFORMAL per language, evaluate outputs of 5 systems, and collect 5 judgments per instance. We open the task to all workers passing QC2 in Table 3. We include inter-annotator agreement results in A.E.

**Formality** We collect formality ratings for the original informal reference, the formal human rewrite, and the formal system outputs on a 7-point discrete scale of -3 to 3, following Lahiri (2015) (*Very informal* → *Informal* → *Somewhat Informal* → *Neutral* → *Somewhat Formal* → *Formal* → *Very Formal*).

**Fluency** We collect fluency ratings for the original informal reference, the formal human rewrite, and the formal system outputs on a discrete scale of 1 to 5, following Heilman et al. (2014) (*Other* → *Incomprehensible* → *Somewhat Comprehensible* → *Comprehensible* → *Perfect*).

**Meaning Preservation** We adopt the annotation scheme of Semantic Textual Similarity (Agirre et al., 2016): given the informal reference and formal human rewrite or the formal system outputs, Turkers rate the two sentences’ similarity on a 1 to 6 scale (*Completely dissimilar* → *Not equivalent but on same topic* → *Not equivalent but share some details* → *Roughly equivalent* → *Mostly equivalent* → *Completely equivalent*).

**Overall** We collect overall judgments of the system outputs using *relative ranking*: given the informal reference and a formal human rewrite, workers are asked to rank system outputs in the order of their overall formality, taking into account both fluency and meaning preservation. An overall score is then computed for each model via averaging results across annotating instances.

<sup>13</sup><https://github.com/xingniu/multitask-ft-fsmt>

<sup>14</sup><https://github.com/cindyxinyiwang/deep-latent-sequence-model>

<sup>15</sup>See A.A for discussion on this assumption.

Method	self-BLEU (↑)			MBERT SCORE (↑)			PERPLEXITY (↓)			multi-BLEU (↑)		
	BR-PT	IT	FR	BR-PT	IT	FR	BR-PT	IT	FR	BR-PT	IT	FR
<b>Automatic</b>												
GOLD STANDARD	23.8	16.2	19.2	-0.48	-0.39	-0.42	5.32	5.84	4.70	46.2	40.2	44.3
COPY				-1.67	-1.94	-2.04	8.77	8.58	8.36	53.1	43.1	46.1
<b>RULE-BASED</b>	<b>79.8</b>	<b>83.4</b>	<b>85.4</b>	-1.33	-1.41	-1.59	7.21	7.24	7.07	53.1	43.1	46.1
ROUND-TRIP MT	33.0*	39.7*	32.5*	-1.10*	-1.27*	-1.20*	6.27*	6.84*	6.03*	43.0*	34.5*	41.4*
TRANSLATE TRAIN TAG	60.9*	57.0*	51.8*	-0.96*	-0.90*	-0.94*	5.89*	6.03*	5.36*	52.4*	42.6	44.9*
MULTI-TASK TAG-STYLE	57.6*	54.6*	48.8*	<b>-0.91*</b>	<b>-0.91*</b>	<b>-0.86*</b>	<b>5.85*</b>	<b>5.83*</b>	<b>5.19*</b>	<b>55.0*</b>	42.5	47.9*
BACKTRANSLATE	75.7*	68.8*	71.9*	-1.12*	-1.07*	-1.19*	6.39*	6.28*	5.93*	54.6*	42.9	<b>48.3*</b>
UNMT	15.3*	15.1*	19.7*	-1.74*	-1.84*	-1.91*	6.42*	6.40*	6.61*	14.8*	11.5*	17.0*
DLSM	20.4*	19.6*	21.3*	-1.84*	-1.61*	-1.90*	6.97*	7.45*	6.97*	18.4*	12.4*	17.4*
UNMT (IN-DOMAIN)			52.8*			-1.71*			5.69*			35.0*
DLSM (IN-DOMAIN)			71.2*			-1.66			6.78*			42.1*
<b>Human</b>												
GOLD STANDARD	4.88	5.18	5.08	+0.66	+1.46	+1.14	4.54	4.79	4.51			
COPY				-1.18	-1.06	-0.75	3.91	4.09	3.88			
<b>RULE-BASED</b>	<b>5.70</b>	<b>5.96</b>	<b>5.95</b>	-0.51	-0.36	-0.36	4.06	4.24	3.94	2.81	3.23	2.89
ROUND-TRIP MT	4.80*	5.03*	4.50*	-0.17*	-0.34	+0.25*	4.07	4.03*	3.97	2.89	2.54*	2.94
TRANSLATE TRAIN TAG	4.97*	5.18*	4.56*	-0.26*	+0.07*	+0.10*	4.15	4.33	3.84	3.03*	3.02*	2.83
MULTI TASK TAG-STYLE	5.07*	5.18*	4.81*	-0.27	+0.01*	+0.05*	4.15	4.22	3.88	<b>3.24*</b>	2.93*	<b>3.19*</b>
BACKTRANSLATE	5.54*	5.72*	5.59*	-0.49	-0.16	-0.13	<b>4.19*</b>	<b>4.37</b>	<b>4.00</b>	3.01*	<b>3.25</b>	3.13*

Table 7: Automatic and human evaluation results for multilingual FOST. \* denotes statistical significance differences compared to RULE-BASED ( $p < 0.05$ ). NMT-based results are ensembles of 4 systems; unsupervised models are average results across 4 reruns. Darker colors denote higher rankings, bold numbers denote best systems.

## 5.5 Results

Table 7 shows automatic results for all models across the four dimensions as well as human ratings for selected top models.

**NMT-based model evaluation** Concretely, the RULE-BASED baselines are significantly ( $p < 0.05$ ) the best performing models in terms of **meaning preservation** across languages. This result is intuitive as the RULE-BASED models act at the surface level and are unlikely to change the informal sentence’s meaning. The BACKTRANSLATE ensemble systems are the second-best performing models in terms of meaning preservation, while the ROUND-TRIP MT outputs diverge semantically from the informal sentences the most. Those results are consistent across languages and human/automatic evaluations. On the other hand, when we compare systems in terms of their **formality**, we observe the opposite pattern: the RULE-BASED and BACKTRANSLATE outputs are the most informal compared to the other ensemble NMT-based approaches across languages. Interestingly, the ROUND-TRIP MT outputs exhibit the largest formality shift for BR-PT and FR as measured by human evaluation. The trade-off between meaning preservation and formality among models was also observed in EN (RT18). Moreover, when we move to **fluency**, we notice similar results across systems. Specifically, human evaluation assigns almost all models an average score of  $> 4$ , denoting that system outputs are comprehensible on average, with small differences

between systems not being statistically significant. Notably, perplexity tells a different story: all system outputs are significantly better compared to the RULE-BASED systems across configurations and languages. This result denotes that perplexity might not be a reliable metric to measure fluency in this setting, as noticed in Mir et al. (2019) and Krishna et al. (2020). When it comes to the **overall ranking** of systems, we observe that the NMT-based ensembles are better than the RULE-BASED baselines for BR-PT and FR, yet by a small margin as denoted by both multi-BLEU and human evaluation. However, the corresponding results for IT denote that there is no clear win, and the NMT-based ensembles still fail to surpass the naive RULE-BASED models, yet by a small margin. Finally, all ensembles outperform the trivial COPY baseline. Table 8 presents examples of system outputs. As a side note, we followed the recommendation of Tikhonov and Yamshchikov (2018) to show the performance of ST models of individual runs and visualize trade-offs between metrics better. Unlike their work which found that reruns of the same model showed wide performance discrepancies, we found that most of our NMT-based models did not vary in performance on XFORMAL. The results can be visualized in A.M.

**Unsupervised model evaluation** We also benchmark the unsupervised models but focus solely on automatic metrics since they lag behind their supervised counterparts. As shown in Table 7, when

(BR-PT)	<b>INFORMAL</b>	n preciso pedir pois sei q ela vai vir atras!!
	<b>XFORMAL REWRITE</b>	Não é preciso pedir, pois sei que ela virá em busca!
	RULE-BASED	não preciso pedir pois sei <b>que</b> ela vai vir atras!
	ROUND-TRIP MT	Não preciso <b>perguntar porque</b> sei <b>que</b> ela <b>virá atrás de mim!</b>
	TRANSLATE TRAIN TAG	<b>Eu não</b> preciso pedir <b>que eu saiba que</b> ela vai vir atras!
	MULTI-TASK TAG-STYLE	<b>Eu não</b> preciso pedir <b>que eu saiba que</b> ela vai vir atras!
TRANSLATE TRAIN BACK.	Não preciso pedir <b>porque</b> sei <b>que</b> ela vai vir atras!	
(FR)	<b>INFORMAL</b>	drôle heinnnnnnnnn s étais ma femme de ménage!
	<b>XFORMAL REWRITE</b>	Le plus amusant est le fait qu'il s'agissait de mon ancienne employée de maison.
	RULE-BASED	Drôle <b>hein</b> s étais ma femme de ménage!
	ROUND-TRIP MT	<b>C'était</b> drôle ma <b>bonne!</b>
	TRANSLATE TRAIN TAG	<b>C'est</b> drôle, <b>mais j'</b> étais ma femme de ménage.
	MULTI-TASK TAG-STYLE	<b>C'était</b> drôle, <b>c'était</b> ma femme de ménage!
TRANSLATE TRAIN BACK.	<b>C'était</b> drôle, <b>mais j'</b> étais ma femme de ménage!	
(IT)	<b>INFORMAL</b>	un po'di raffreddore ma tutto ok!!!
	<b>XFORMAL REWRITE</b>	Sono affetta da un leggero raffreddore ma per il resto va tutto bene.
	RULE-BASED	<b>Un</b> po'di raffreddore ma tutto ok!
	ROUND-TRIP MT	<b>Un</b> po' <b>freddo</b> ma <b>va bene!</b>
	TRANSLATE TRAIN TAG	<b>Un</b> po'di raffreddore, ma tutto <b>va bene.</b>
	MULTI-TASK TAG-STYLE	<b>Un</b> po'di raffreddore, ma <b>va bene.</b>
TRANSLATE TRAIN BACK.	<b>Un</b> po'di raffreddore, ma tutto <b>va bene.</b>	

Table 8: Example system outputs on random sentences of XFORMAL.




Figure 3: Heatmaps of LeD scores showing the lexical difference between pairs of systems.

using out-of-domain data (e.g., OpenSubtitles) for training, the models perform worse than their NMT counterparts across all three languages. The difference is most stark when considering self-BLEU and multi-BLEU scores. However, given access to large *in-domain* corpora (e.g., L26 Yahoo! French Answers) the gap between the two model classes closes with DSLM achieving a multi-BLEU score of 42.1 compared to 48.3 for the best performing NMT model BACKTRANSLATE. This shows the promise of unsupervised methods, assuming a large amount of in-domain data, on multilingual ST tasks.

**Lexical differences of system outputs** Finally, in Figure 3 we analyze the diversity of outputs by leveraging LeD scores resulting from pair-wise comparisons of different NMT systems. A larger LeD score denotes a larger difference between the lexical choices of the two systems under comparison. First, we observe that the ROUND-TRIP MT outputs have the smallest lexical overlap with the informal input sentences. However, when this observation is examined together with human evaluation results, we conclude that the large number of lexical edits happens at the cost of diverging se-



manically from the input sentences. Moreover, we observe that the average lexical differences within NMT-based systems are small. This indicates that different systems perform similar edit operations that do not deviate a lot from the input sentence in terms of their lexical choices. This is unfortunate given that multilingual FoST requires systems to perform more phrase-based operations, as shown in the analysis in §4.

**Evaluation Metric** While evaluating evaluation metrics is not a goal of this work (though the data can be used for that purpose), we observe that the top models identified by the automatic metrics generally align with the top models identified by humans. While promising, further work is required to confirm if the automatic measures really do correlate with human judgments.

## 6 Conclusions & Future Directions

This work extends the task of formality style transfer to a multilingual setting. Specifically, we contribute XFORMAL, an evaluation testbed consisting of informal sentences and multiple formal rewrites spanning three languages: BR-PT, FR, and IT. As in Rao and Tetreault (2018) Turkers can be effective in creating high quality ST corpora. In contrast to the aforementioned EN corpus, we find that the rewrites in XFORMAL tend to be more diverse, making it a more challenging task.

Additionally, inspired by work on cross-lingual transfer and EN FoST, we benchmark several methods and perform automatic and human evaluations on their outputs. We found that NMT-based ensembles are the best performing models for FR and BR-PT—a result consistent with EN—however, they perform comparably to a naive RULE-BASED baseline for IT. To further facilitate reproducibility of our evaluations and corpus creation processes, as well as drive future work, we will release our scripts, rule-based baselines, source data, and annotation templates, on top of the release of XFORMAL.

Our results open several avenues for future work in terms of benchmarking and evaluating FoST in a more language inclusive direction. Notably, current supervised **and** unsupervised approaches for EN FoST rely on parallel in-domain data—with the latter treating the parallel set as two unpaired corpora—that are not available in most languages. We suggest that **benchmarking** FoST models in multilingual settings will help understand their ability to generalize and lead to safer conclusions when

comparing approaches. At the same time, multilingual FoST calls for more language-inclusive consideration for **automatic evaluation** metrics. Model-based approaches have been recently proposed for evaluating different aspects of ST. However, most of them rely heavily on English resources or pre-trained models. How those methods can be extended to multilingual settings and how we evaluate their performance remain open questions.

## Acknowledgements

We thank Sudha Rao for providing references and materials of the GYAFC dataset, Chris Callison-Burch and Courtney Napoles for discussions on MTurk annotations, Svebor Karaman for helping with data collection, our colleagues at Dataminr, and the NAACL reviewers for their helpful and constructive comments.

## 7 Ethical Considerations

Finally, we address ethical considerations for **dataset papers** given that our work proposes a new corpus XFORMAL. We reply to the relevant questions posed in the NAACL 2021 Ethics FAQ.<sup>16</sup>

### 7.1 Dataset Rights

The underlying data for our dataset as well as training our FoST models and formality classifiers are from Yahoo! Answers L6 dataset. We were granted written permission by Yahoo (now Verizon) to make the resulting dataset public for academic use.

### 7.2 Dataset Collection Process

Turkers are paid over 10 USD an hour. We targeted a rate higher than the US national minimum wage of 7.50 USD given discussions with other researchers who use crowdsourcing. We include more information on collection procedures in §3.

### 7.3 IRB Approval

This question is not applicable for our work.

### 7.4 Dataset Characteristics

We follow Bender and Friedman (2018) and Gebru et al. (2018) and report characteristics in §3 and §4.

<sup>16</sup><https://2021.naacl.org/ethics/faq/>

## 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.
- Sweta Agrawal and Marine Carpuat. 2019. [Controlling text complexity in neural machine translation](#). 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 1549–1564, Hong Kong, China. Association for Computational Linguistics.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- 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.
- Tuhin Chakrabarty, Smaranda Muresan, and Nanyun Peng. 2020. [Generating similes effortlessly like a pro: A style transfer approach for simile generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6455–6469, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. 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.
- Weston Feely, Eva Hasler, and Adrià de Gispert. 2019. [Controlling Japanese honorifics in English-to-Japanese neural machine translation](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 45–53, Hong Kong, China. Association for Computational Linguistics.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be guilty but references are not innocent](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 61–71, Online. Association for Computational Linguistics.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daumé, and Kate Crawford. 2018. [Datasheets for datasets](#). *ArXiv*, abs/1803.09010.
- 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.
- Eva Hasler, Adrià de Gispert, Gonzalo Iglesias, and Bill Byrne. 2018. [Neural machine translation decoding with terminology constraints](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 506–512, New Orleans, Louisiana. Association for Computational Linguistics.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. [A probabilistic formulation of unsupervised text style transfer](#). In *Proceedings of International Conference on Learning Representations*.
- Ruining He and Julian McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, page 507–517, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- 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.

- F. Heylighen, J. Dewaele, and L. Apostel. 1999. Formality of language: definition, measurement and behavioral determinants.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. 2018. [The sockeye neural machine translation toolkit at AMTA 2018](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 200–207, Boston, MA. Association for Machine Translation in the Americas.
- Eduard Hovy. 1987. [Generating natural language under pragmatic constraint](#). *Journal of Pragmatics*, 11(6):689–719.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596, International Convention Centre, Sydney, Australia. PMLR.
- Zhijing Jin, Di Jin, Jonas Mueller, Nicholas Matthews, and Enrico Santus. 2019. [Imat: Unsupervised text attribute transfer via iterative matching and translation](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 3095–3107. Association for Computational Linguistics.
- Dongyeop Kang, Varun Gangal, and Eduard Hovy. 2019. [\(male, bachelor\) and \(female, ph.d\) have different connotations: Parallely annotated stylistic language dataset with multiple personas](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong.
- Reinhard Kneser and Hermann Ney. 1995. [Improved backing-off for m-gram language modeling](#). In *ICASSP*, pages 181–184. IEEE Computer Society.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Elizaveta Korotkova, Agnes Luhtaru, Maksym Del, Krista Liin, Daiga Dekšne, 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.
- 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.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. [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.
- 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).
- Lajanugen Logeswaran, Honglak Lee, and Samy Bengio. 2018. [Content preserving text generation with attribute controls](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 5103–5113. Curran Associates, Inc.
- 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, IJCAI 2019*.
- Eric Malmi, Aliaksei Severyn, and Sascha Rothe. 2020. [Unsupervised text style transfer with padded masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8671–8680, Online. Association for Computational Linguistics.
- Kelly Marchisio, Jialiang Guo, Cheng-I Lai, and Philipp Koehn. 2019. [Controlling the reading level of machine translation output](#). In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 193–203, Dublin, Ireland. European Association for Machine Translation.

- 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.
- Shachar Mirkin and Jean-Luc Meunier. 2015. Personalized machine translation: Predicting translational preferences. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Amit Moryossef, Roei Aharoni, and Yoav Goldberg. 2019. [Filling gender & number gaps in neural machine translation with black-box context injection](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 49–54, Florence, Italy. Association for Computational Linguistics.
- Xing Niu and Marine Carpuat. 2020. [Controlling neural machine translation formality with synthetic supervision](#). In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 8568–8575.
- Xing Niu, Marianna Martindale, and Marine Carpuat. 2017. [A study of style in machine translation: Controlling the formality of machine translation output](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2814–2819, Copenhagen, Denmark. 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.
- 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 on Association for Computational Linguistics*, ACL '02, page 311–318, 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.
- 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.
- 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.
- Ella Rabinovich, Raj Nath Patel, Shachar Mirkin, Lucia Specia, and Shuly Wintner. 2017. [Personalized machine translation: Preserving original author traits](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1074–1084, Valencia, Spain. 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.
- Víctor M. Sánchez-Cartagena, Marta Bañón, Sergio Ortiz-Rojas, and Gema Ramírez-Sánchez. Prompt’s submission to wmt 2018 parallel corpus filtering shared task. In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, Brussels, Belgium. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [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.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- 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.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6830–6841. Curran Associates, Inc.
- Sandeep Subramanian, Guillaume Lample, Eric Michael Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2018. [Multiple-attribute text style transfer](#). *CoRR*, abs/1811.00552.
- Aleksey Tikhonov and Ivan Yamshchikov. 2018. [Sounds wilde. phonetically extended embeddings for author-stylized poetry generation](#). In *Proceedings of the Fifteenth Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 117–124, Brussels, Belgium. Association for Computational Linguistics.
- Aleksey 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.
- Eva Vanmassenhove, Christian Hardmeier, and Andy Way. 2018. [Getting gender right in neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3003–3008, Brussels, Belgium. 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.
- Chen Wu, Xuancheng Ren, Fuli Luo, and Xu Sun. 2019. [A hierarchical reinforced sequence operation method for unsupervised text style transfer](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4873–4883.
- Jingjing Xu, Xu Sun, Qi Zeng, Xuancheng Ren, Xiaodong Zhang, Houfeng Wang, and Wenjie Li. 2018. [Unpaired sentiment-to-sentiment translation: A cycled reinforcement learning approach](#). In *ACL*.
- 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.
- Hayahide Yamagishi, Shin Kanouchi, Takayuki Sato, and Mamoru Komachi. 2016. [Controlling the voice of a sentence in Japanese-to-English neural machine translation](#). In *Proceedings of the 3rd Workshop on Asian Translation (WAT2016)*, pages 203–210, Osaka, Japan. The COLING 2016 Organizing Committee.
- Ivan P. Yamshchikov, Viacheslav Shibaev, Nikolay Khlebnikov, and Alexey Tikhonov. 2020. [Style-transfer and paraphrase: Looking for a sensible semantic similarity metric](#). *ArXiv*, abs/2004.05001.
- Yi Zhang, Tao Ge, and Xu SUN. 2020. [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.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Txygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research; Development in Information Retrieval, SIGIR ’18*, page 1097–1100, New York, NY, USA. Association for Computing Machinery.

## A Does translation preserve formality?

We examine the extent to which machine translation—through the AWS service—affects the formality level of an input sentence: starting from a set of English sentences we have formality judgments for (ORIGINAL-**EN**), we perform a round-trip translation via pivoting through an auxiliary language (PIVOT-**X**). We then compare the formality prediction scores of the English Formality regression model for the two versions of the English input. In terms of Spearman correlation, the model’s performance drops by 7.5 points on average when tested on round-trip translations. To better understand what causes this drop in performance, we present a per formality bin analysis in Table 9. On average we observe that translation preserves the formality level of formal sentences considerably well, while at the same time it tends to shift the formality level of informal sentences towards formal values—by a margin smaller than 1 point—most of the times. To account for the formalization effect of translation, we draw the line between formal and informal sentences at the value of  $-1$  for scores predicted by multilingual regression models. This decision is based on the following intuition: if the formality shift of machine translated informal sentences is around  $+1$  value, the propagation of English formality labels imposes a negative shift of formal sentences in the model’s predictions.

## B Amazon Web Service details

We compute the performance of the AWS system on 2.5K randomly sentences from OpenSubtitles, as a sanity check of translation performance. We report BLEU of 37.16 (BR-PT), 33.79 (FR), and 32.67 (IT).

## C Rule-based baselines

We develop a set of rules to automatically make an informal sentence more formal via performing surface-level edits similar to the EN rule-based system of Rao and Tetreault (2018). The set of extracted rules are shared across languages with the only difference being the list of abbreviations:

**Normalize punctuation** We remove punctuation symbols that are repeated  $\geq 2$  times.

**Character repetition** We trailed characters repeated  $\geq 3$  times (e.g., *ciaooo*  $\rightarrow$  *ciao* (IT)).

**Normalize casing** Several sentences might consist of words that are written in upper case. We lower case all characters apart from the first characters of the first word.

**Normalize abbreviations** Informal text might contain slang words that are abbreviated. We hand-craft a list of expansions for each language. For example, we replace *kra*  $\rightarrow$  *cara* (BR-PT), *ta*  $\rightarrow$  *ti amo* (IT), and *bjr*  $\rightarrow$  *bonjour* (FR), using publicly available resources.<sup>17, 18, 19</sup> The resulting lists sizes are: 64 (BR-PT), 21 (IT), and 48 (FR).

## D NMT architecture details

We used the NMT implementations of Niu et al. (2018) that are publicly available: <https://github.com/xingniu/multitask-ft-fsmt>. The NMT models are implemented as bi-directional LSTMs on Sockeye (Hieber et al., 2018), using the same configurations across languages to establish fair comparisons. We use single LSTMs consisting of a single of size 512, multilayer perceptron attention with a layer size of 512, and word representations of size 512. We apply layer normalization and tie the source and target embeddings as well as the output layer’s weight matrix. We add dropout with probability 0.2 (for the embeddings and LSTM cells in both the encoder and the decoder). For training, we use the Adam optimizer with a batch size of 64 sentences and checkpoint the model every 1000 updates. Training stops after 8 checkpoints without improvement of validation perplexity. For decoding, we use a beam size of 5.

## E Inter-annotator agreement on human evaluation

To quantify inter-annotator agreement for the tasks of *formality*, *meaning preservation*, and *fluency* we measure the correlation of their ordinal ratings using inter-class correlation (ICC) and their categorical agreement using a variations of Cohen’s  $\kappa$  coefficient. For the latter, given that we collect human evaluation judgments through crowd-sourcing, we follow the simulation framework of Pavlick and Tetreault (2016) to quantify agreement. Concretely, we simulate two annotators (*Annotator 1*,

<sup>17</sup><https://braziliangringo.com/brazilianinternetslang/>

<sup>18</sup><https://www.dummies.com/languages/italian/texting-and-chatting-in-italian/>

<sup>19</sup><https://frenchtogether.com/french-texting-slang/>

	Informal bins			Formal bins		
	[-3,-2]	[-2,-1]	[-1,0]	[0,1]	[1,2]	[2,3]
GOLD-STANDARD	-2.35 ± 0.31	-1.49 ± 0.36	-0.57 ± 0.33	0.46 ± 0.32	1.33 ± 0.34	2.27 ± 0.25
ORIGINAL-EN	-1.81 ± 0.59	-1.18 ± 0.73	-0.53 ± 0.71	0.30 ± 0.82	1.17 ± 0.68	1.50 ± 0.45
PIVOT-IT	-1.58 ± 0.67	-0.95 ± 0.74	-0.35 ± 0.68	0.34 ± 0.86	1.12 ± 0.73	1.41 ± 0.54
Δ	<b>+0.23 ± 0.45*</b>	<b>+0.23 ± 0.54*</b>	<b>+0.18 ± 0.53*</b>	+0.04 ± 0.39	<b>-0.05 ± 0.23</b>	<b>-0.09 ± 0.24</b>
PIVOT-FR	-1.44 ± 0.77	-0.93 ± 0.75	-0.37 ± 0.72	0.38 ± 0.80	1.12 ± 0.82	1.50 ± 0.47
Δ	<b>+0.37 ± 0.54*</b>	<b>+0.25 ± 0.50*</b>	<b>+0.15 ± 0.48</b>	<b>+0.08 ± 0.42</b>	<b>-0.05 ± 0.49</b>	-0.00 ± 0.17
PIVOT-PT	-1.46 ± 0.71	-0.87 ± 0.80	-0.31 ± 0.70	0.34 ± 0.83	1.14 ± 0.71	1.47 ± 0.47
Δ	<b>+0.35 ± 0.49*</b>	<b>+0.31 ± 0.57*</b>	<b>+0.21 ± 0.55*</b>	+0.03 ± 0.39	-0.03 ± 0.31	-0.03 ± 0.14

Table 9: Average formality scores of human annotations (GOLD-STANDARD) and model’s predictions on the original and round-trip translated PT16 test set grouped in 6 bins of varying formality. Δ gives the average formality shift of the English sentences resulting from round-trip translation. Scores in blue and red indicate that mean is above and below zero, respectively. \* denotes statistical significant formality shifts with  $p < 0.05$ . Translation preserves formality of formal sentences while informal sentences exhibit a shift towards formal values.

Annotator 2) via randomly choosing one annotator’s judgment for a given instance as the rating of Annotator 1 and taking the mean rating of the rest judgments as the rating of Annotator 2. We then compute Cohen’s  $\kappa$  for these two simulated annotators. We repeat this process 1,000 times, and report the median and standard deviation of results. For measuring agreement of the overall ranking evaluation task we use the same simulated framework and report results of Kendall’s  $\tau$ . Table 10 presents Inter-annotator agreement results on human evaluation across evaluation aspects and languages.

## F Yahoo! L6 language statistics

Table 11 presents the total number of questions included in L6 Yahoo! Corpus, for each language. Although almost 90% of questions are in English, the corpus contains a non-negligible number of informal sentences in Spanish, French, Portuguese, Italian, and German, with a long tail of few questions for other languages.

## G Instructions for XFORMAL

**Summary of the task** Given an **informal sentence** in French (or Portuguese/Italian), generate its **formal rewrite**, *without changing its meaning*.

**Detailed Instructions** Given an informal sentence, provide us with its informal rewrite. The informal rewrite should only change the formality attribute of the original sentence and preserve its meaning. Each sentence should be treated independently while rewrites should only rely on the information available in the sentences. There is no need to guess what additional information might

be available in the documents the sentences come from.

**Examples** Following we include examples of good and bad rewrites given to Turkers:

INFORMAL *Wow, I am very dumb in my observation skills.....*

GOOD FORMAL *I do not have good observation skills.*

REASONING *Formality is properly transferred and meaning is preserved.*

BAD FORMAL *Wow, I am very foolish in my observation skills..*

REASONING *Formality is not properly transferred.*

BAD FORMAL *I am very unintelligent and I don’t have good observation skills.*

REASONING *Meaning has changed.*

## H Qualitative analysis of XFORMAL

Following, we include the instructions given to **native speakers** for the qualitative analysis of XFORMAL, as described in §4.

**Background context** In this task you will be asked to judge the quality of formal rewrites of a set of informal sentences. To give more background context, your work would serve as a quality check over annotations obtained from the Amazon Mechanical Turk (AMT) platform. AMT workers were given an informal sentence, e.g., “I’d say it is punk though”, and then asked to provide us with its formal rewrite while maintaining the original

TASK	BR-PT		FR		IT	
	Weighted $\kappa$	ICC	Weighted $\kappa$	ICC	Weighted $\kappa$	ICC
Formality	$0.57 \pm 0.02$	$0.72 \pm 0.01$	$0.56 \pm 0.02$	$0.71 \pm 0.01$	$0.67 \pm 0.02$	$0.79 \pm 0.02$
Fluency	$0.43 \pm 0.02$	$0.62 \pm 0.02$	$0.58 \pm 0.02$	$0.73 \pm 0.01$	$0.46 \pm 0.02$	$0.64 \pm 0.02$
Meaning Preservation	$0.53 \pm 0.02$	$0.70 \pm 0.02$	$0.71 \pm 0.02$	$0.83 \pm 0.01$	$0.55 \pm 0.03$	$0.71 \pm 0.02$
Kendall's $\tau$						
Overall	$0.41 \pm 0.03$		$0.44 \pm 0.03$		$0.41 \pm 0.03$	

Table 10: Inter-annotator agreement for human evaluation results.

LANGUAGE	CODE	# QUESTIONS
English	EN	3,895,407
Spanish	ES	258,086
French	FR	125,393
Portuguese	PT	105,813
Italian	IT	55,027
German	DE	43,149
Persian	FA	30
Catalan	CA	27
Dutch	NL	13
Danish	DA	10
Arabic	AR	8
Romania	RO	8
Norwegian	NO	8
Swedish	SV	5
Estonian	ET	5
Finnish	FI	4
Malay	MS	4
Turkish	TR	4
Slovak	SK	3
Latvian	LV	3
Albania	SQ	3
Croatian	HR	3
Tagalog	TL	3
Czech	CS	3
Slovenian	SL	3
Vietnamese	VI	2
Icelandic	IS	2
Hungarian	HU	2
Polish	PL	2

Table 11: Number of question per language included in L6 Yahoo! Answers.

content and being grammatical/fluent, e.g., “However, I do believe it to be punk”. In our work, workers were presented with sentences in French, Italian and Brazilian Portuguese. For our analysis we want to know a) whether the quality of the collected annotations is good (Task 1), b) what are the types of edits workers performed when formalizing the input sentence (Task 2). Both tasks consist of the same 200 informal-formal sentences-pairs and could be performed either in parallel (e.g., judging a single informal-formal sentence-pair both in terms of quality and types of edit at the same time; Task 1 and 2 in Google sheet), or individually (Task 1, Task 2 in Google sheet). More information and

examples for the two tasks are included below.

**Task 1** Given an informal sentence you are asked to assess the quality of its formal rewrite. For each sentence pair type excellent, acceptable, or poor under the rate column. Read the instructions below before performing the task!

**What constitutes a good formal rewrite?**

- The style of the rewrite should be formal
- The overall meaning of the informal sentence should be preserved
- The rewrite should be fluent

**How should I interpret the provided options?**

Below we include detailed instructions on how to interpret the provided options.

- **Excellent** the rewrite is formal, fluent and the original meaning is maintained. There is very little that could be done to make a better rewrite.
- **Acceptable** the rewrite is generally an improvement upon the original but there are some minor issues (such as the rewrite contains a typo or missed transforming some informal parts of the sentence, for example).
- **Poor** the rewrite offers a marginal or no improvement over the original. There are many aspects that could be improved.

**Task 2** In this task the goal is to characterize the types of edits workers made while formalizing the informal input sentence. For each informal-formal sentence pair you should check each of the provided boxes. Note that: a) Multiple types of edits might hold at the same type (e.g., in the informal—formal “Lexus cars are awesome!”—“Lexus cars



are very nice.” you should check both the paraphrase and the punctuation boxes.); b) At the end of the Google sheet there is a column named ‘Other’ you are welcome to write down in plain text any additional type that you observed and it is not covered by the existing classes. The provided classes are: capitalization, punctuation, paraphrase, delete fillers, completion, add context, contractions, spelling, normalization, slang/idioms, politeness, split sentences, and relativizers.

## I Quantitative analysis of XFORMAL

Following, we include more details on the qualitative analysis procedures.

**Capitalization** A rewrite performs a capitalization edit if it contains tokens that appear in capital letters in the informal text but in lowercase letters in the formal text.

**Punctuation** A rewrite contains a punctuation edits if any of the punctuation of the informal-formal texts differs.

**Spelling** We identify spelling errors based on the character-based Levenshtein distance between informal-formal tokens.

**Normalization** We identify normalization edits based on a hand-crafted list of abbreviations for each language.

**Split sentences** We split sentences using the NLTK toolkit.

**Paraphrase** A formal rewrite is considered to contain a paraphrase edit if the token level Levenshtein distance between the informal-formal text is greater than 3 tokens.

**Lowercase** A rewrite performs a capitalization edit if it contains tokens that appear in lower case in the informal text but in capital letters in the formal text.

**Repetition** We identify repetition tokens using regular expressions (a token that appears more than 3 times in a row is considered a repetition.)

## J XFORMAL: Data Quality

We ask a **native speaker** to judge the quality of a sample of 200 rewrites for each language via choosing one of the following three options: “excellent”, “acceptable”, and “poor”. The details of this analysis are included in A.H (Task 1). Results indicate

that on average less than 10% of the rewrites were identified as of poor quality across languages while more than 40% as of excellent quality. The small number of rewrites identified as “poor” consists mostly of edits where humans add context not appearing in the original sentence. We choose not to exclude any of the rewrites from the dataset as we provide multiple reformulations of each informal instance.

## K OpenSubtitles data

OpeSubtitles	Preprocessing	BR-PT	IT	FR
All	✗	61.3M	35.2M	41.8M
Informal	✓	29.8M	16.8M	19.0M
	✓	25.7M	14.7M	16.3M

Table 12: OpenSubtitles statistics in three languages.

## L Qualification tests

Table 13 presents questions and answers used in QC2 of our annotation protocol. Turkers have to score 80 and above to participate in the task. To compute an average score for each test, we assume that an answer is incorrect if it deviates more than 1 point from the gold-standard scores given in the second column of Table 13.

## M Trade-off plots

**Trade-off plots** Figure 4 presents trade-off plots between multi-BLEU vs. formality score and fluency, across different reruns as proposed by Tikhonov and Yamshchikov (2018). First, we observe that models exhibit small variations in terms of formality and fluency scores across different reruns, and larger variations across BLEU for most cases. Notably, the single seed BACKTRANSLATE systems are the most consistent across reruns for all metrics and languages. Furthermore, in almost all cases, models trained on 2M data, perform better than the naive COPY baseline, across metrics. However, single seed models fail to consistently outperform the RULE-BASED baselines in almost all cases, with the exception of BACKTRANSLATE for BR-PT and FR which report an improvement of about 1 BLEU score.

## N Compute time & Infrastructure

All experiments for benchmarking both NMT-based and unsupervised approaches use Amazon EC2 P3 instances: <https://aws.amazon.com/>




Figure 4: BLEU vs. STYLE SCORE vs. PERPLEXITY trade-off plots across 4 reruns.

---

**FRENCH**

---

Heureux ceux qui ont la chance de pouvoir faire la sieste devant Derrick !	-2
je kiffe trop ce film il déchire carément!!	-3
Moi je les données à des vieilles personnes qui ne comprennent pas le fonctionnement d'un lecteur DVD!	-1
Le producteur de ce morceau, c'est Scott Torch.	2
Il s'agit bien du groupe America et le titre "A horse with no name".	2
La France ne gagnera jamais si elle continue avec des chansons qui ont l'air de dater des années 70.	0
Le but n'est pas de se faire du profit dessus!	-1
Il existe surtout sur les chaînes du câble et du satellite, TL9, TPS foot, par exemple.	2

---

**ITALIAN**

---

TI do questo sito qui puoi trovare qualunque cosa su Tru Calling buona lettura!	-1
Parla proprio di amore in chat.	-2
ma parla cm mangi!!!!!!!!!!!!	-3
io ho un programma si kiama evil lyrics.	-1
Se non ci si dà una svegliata "dove andremo a finire?"	-1
Alla fine del Festival ci saranno due vincitori, uno per categoria.	2
è quello che stò passando ora.	-1
Il montaggio risulta quindi ridotto al minimo.	2

---

**BRAZILIAN-PORTUGUESE**

---

Me parece que você usa BASTANTE essa palavra e seus derivados, né?!!	-2
Mas vem cá, você não tem nada melhor pra fazer do que escutar RBD ao contrário?!	-2
Então na minha opinião é babaquice daquele que tem preconceitos desse tipo.	-1
Vi a propaganda na locadora, na minha cidade.	0
Na minha opinião, o mínimo existencial subentende a palavra oportunidade.	1
A única diferença dessa mulher para um assassino comum, é que ela nem chega a conhecer sua vítima.	1
Dentre as hipóteses menos consideradas por historiadores sérios está o famoso mito celta.	1
Dado um ambiente qualquer, uma sala ou um quarto, por exemplo, as ondas sonoras "respondem" de formas diferentes.	-1

---

Table 13: Qualification test (questions and answers) for QC2

[ec2/instance-types/p3/](#), on Tesla V100 GPUs. Concretely, NMT-based experiments are run on 4 GPUS (~ 1-2 hours), while unsupervised models are run on single GPUS, with average run time spanning from ~ 5 hours (for out-of-domain data—100K), to ~ 72 hours (for in domain data—2M).