

Improved Sentiment Detection via Label Transfer from Monolingual to Synthetic Code-Switched Text

Anonymous ACL submission

Abstract

Multilingual writers and speakers often alternate between two languages in a single discourse, a practice called “code-switching”. Existing sentiment detection methods are usually trained on sentiment-labeled monolingual text. Manually labeled code-switched text, especially involving minority languages, is extremely rare. Consequently, the best monolingual methods perform relatively poorly on code-switched text. We present an effective technique for synthesizing labeled code-switched text from labeled monolingual text, which is relatively readily available. The idea is to replace carefully selected subtrees of constituency parses of sentences in the resource-rich language with suitable token spans selected from automatic translations to the resource-poor language. By augmenting the scarce labeled code-switched text with plentiful synthetic labeled code-switched text, we achieve significant improvements in sentiment labeling accuracy (1.5%, 5.11%, 7.20%) for three different language pairs (English-Hindi, English-Spanish and English-Bengali). We also get significant gains for hate speech detection: 4% improvement using only synthetic and 6% if augmented with real text.

1 Introduction

Sentiment analysis on social media is critical for commerce and governance. Multilingual social media users often use code-switching, particularly to express emotion (Rudra et al., 2016). However, a basic requirement to train any sentiment analysis (SA) system is the availability of large sentiment-labeled corpora. These are extremely challenging to obtain (Chittaranjan et al., 2014; Vyas et al., 2014; Barman et al., 2014), requiring volunteers fluent in multiple languages.

We present **CSGen**, a system to synthesize unlimited sentiment-tagged code-switched text, without involving human labelers of code-switched text, or any linguistic theory or grammar for code-switching. These texts can then train state-of-the-art SA algorithms which, until now, primarily worked with monolingual text.

A common scenario in code-switching is that a resource-rich *source* language is mixed with a resource-poor *target* language. Given a sentiment-labeled source corpus, we first create a parallel corpus by translating to the target language, using a standard translator. Although existing neural machine translators (NMTs) can translate a complete source sentence to a target sentence with good quality, it is difficult to translate only designated source segments in isolation because of missing context and lack of coherent semantics.

Among our key contributions is a suite of approaches to automatic segment conversion. Broadly, given a source segment selected for code switching, we propose intuitive ways to select a corresponding segment from the target sentence, based on maximum similarity or minimum dissimilarity with the source segment, so that the segment blends naturally in the outer source context. Finally, the generated synthetic sentence is tagged with the same sentiment label as the source sentence. The source segment is carefully chosen based on an observation that apart from natural switching points dictated by syntax, there is a propensity to code-switch between highly opinionated segments.

Extensive experiments show that augmenting scarce natural labeled code-switched text with plentiful synthetic text associated with ‘borrowed’ source labels enriches the feature

space, enhances its coverage, and improves sentiment detection accuracy, compared to using only natural text. On four natural corpora having gold sentiment tags, we demonstrate that adding synthetic text can improve accuracy by 5.11% in English-Spanish, 7.20% in English-Bengali and (1.5%, 0.97%) in English-Hindi (Twitter, Facebook). The synthetic code-switch text, even when used by itself to train SA, performs almost as well as natural text in several cases. Hate speech is an extreme emotion expressed often on social media. On an English-Hindi gold-tagged hate speech benchmark, we achieve 6% absolute F1 improvement with data augmentation, partly because synthetic text mitigated label imbalance present in scarce real text.

2 Related Work

Recent SA systems are trained on labeled text (Sharma et al., 2015; Vilares et al., 2015; Joshi et al., 2016). For European and Indian code-switched sentiment analysis, several shared tasks have been initiated (Barman et al., 2014; Rosenthal et al., 2017; Patra et al., 2018; Sequiera et al., 2015; Solorio et al., 2014). Some of these involve human annotations on code-switched text. Vilares et al. (2015) have annotated the data set released for POS tagging by Solorio and Liu (2008). Joshi et al. (2016) had Hindi-English code-switched Facebook text manually annotated and developed a deep model for supervised prediction.

In a different direction, synthetic monolingual text has been created by Generative Adversarial Networks (GAN) (Kannan and Vinyals, 2017; Zhang et al., 2016, 2017; Maqsd, 2015), or Variational Auto Encoders (VAE) (Bowman et al., 2015). Some of these models can be used to generate sentiment-tagged synthetic text. However, they are not directly suitable for generating bilingual code-mixed text, due to the unavailability of sufficient volume of gold-tagged code-mixed text. Recently, Pratapa et al. (2018) used linguistic constraints arising from Equivalence Constraint Theory to design a code-switching grammar that guides text synthesis. Earlier, Bhat et al. (2016) presented similar ideas, but without empirical results. In contrast, CS-Gen uses a data-driven combination of word

alignment weights, similarity of word embeddings between source and target, and attention (Bahdanau et al., 2015).

3 Generation of code-switched text

CSGen takes a sentiment-labeled source sentence s and translates it into a target language sentence t . Then it generates text with language switches on particular constituent boundaries. This involves two sub-steps: select a segment in s (§3.1), and then select text from t that can replace it (§3.2–§3.3). This generation process is sketched in Algorithm 1.

3.1 Sentiment-oriented source segment selection

In this step, our goal is to select a contiguous segment from the source sentence that could potentially be replaced by some segment in the target sentence. (Allowing non-contiguous target segments usually led to unnatural sentences.) Code switching tends to occur at constituent boundaries (Sankoff and Poplack, 1981), an observation that holds even for social media texts (Begum et al., 2016). Therefore, we apply a constituency parser to the source sentence. Specifically, we use the Stanford CoreNLP shift-reduce parser (Zhu et al., 2013) to generate a parse tree¹. Then we select segments under non-terminals, i.e., subtrees, having certain properties, chosen using heuristics informed by patterns observed in real code-switched text.

NP and VP. We allow as candidates all subtrees rooted at NP (noun phrase) and VP (verb phrase) nonterminals, which may cover multiple words. Translating single-word spans is more likely to result in ungrammatical output (Sankoff and Poplack, 1981).

SBAR. Bilingual writers often use a clause to provide a sentiment-neutral part and then switch to another language in another sentence-piece to express an opinion or vice-versa. An example is “Ramdhanu ended with tears *kintu sesh ta besh onho rokom etar*” (Ramdhanu ended with tears but the ending was quite different). Here the constituent “*but the ending was quite different*” comes under the subtree of SBAR.

¹<http://stanfordnlp.github.io/CoreNLP/>

Algorithm 1 CSGen overview.

```

1: Input: Sentiment-labeled source sentences  $S = \{(s_n, y_n)\}$ 
2: Output: Synthetic code-switched sentences  $C = \{(c, y)\}$ 
3:  $t_n \leftarrow \text{Translate}(s_n) \ \forall s_n \in S$  /* Make parallel corpus */
4:  $C \leftarrow \emptyset$ 
5: for each parallel sentence pair  $s, t$  do
6:   /* Collect word alignment signals */
7:    $a \leftarrow \text{AttentionScore}(s, t)$ ,  $g \leftarrow \text{GizaScore}(s, t)$ 
8:   /* Source segment selection */
9:    $P \leftarrow \text{SentimentOrientedSegmentSelection}(s)$ 
10:  for each segment  $p_s \in P$  to replace do
11:    /* Target segment selection */
12:     $q^1, q^2 \leftarrow \text{MaxSimTargetSeg}(s, p_s, t, a, g)$ 
13:     $q^3, q^4 \leftarrow \text{MinDissimTargetSeg}(s, p_s, t, 1-a, 1-g)$ 
14:    /* Code-switched text generation */
15:     $C_k \leftarrow \text{Project}(s, t, p_s, q^k)$  where  $k \in \{1, \dots, 4\}$ 
16:     $C \leftarrow C \cup \text{SelectBest}(\{C_k : k \in \{1, \dots, 4\}\})$ 
17:  end for
18: end for
19:  $C \leftarrow \text{Threshold}(C)$  /* Retain only best replacements */

```

Highly opinionated segments. We also include segments which have a strong opinion polarity, as detected by a (monolingual) sentiment analyzer (Gilbert, 2014). E.g., the tweet “*badiyan badiyan . . . kuch bhi ho humare teacher made us*” translates to “Great great . . . whatever we are, our teacher made us”.

An example sentence, its parse tree, and its candidate replacement segments are shown in Figure 1. In Algorithm 1, $p_s \in P$ denotes the set of candidate replacement subtrees, which correspond to segments. For each candidate segment, we generate a code-switched version of the source sentence, as described next.

3.2 Target segment selection

Given a source sentence s , corresponding target t , and one (contiguous) source segment $p_s = \{w_s^i \dots w_s^{i+x}\}$, the goal is to identify the best possible a contiguous target segment $q_t = \{w_t^j \dots w_t^{j+y}\}$ that could be used to replace p_s to create a realistic code-switched sentence. We adopt two approaches to achieve this goal: (a) selecting a target segment that has maximum similarity with p_s , and (b) selecting a target segment having minimum dissimilarity with p_s , for various definitions of similarity and dissimilarity. Below, we describe several methods that achieve this goal after describing several alignment scores which will be used in these methods. Overall, these

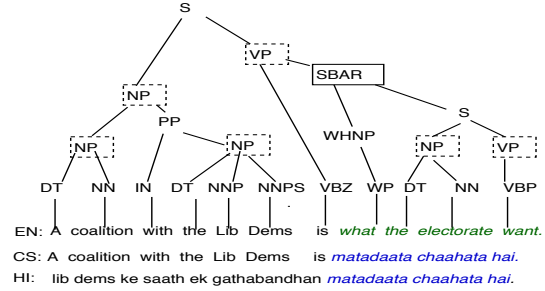


Figure 1: A phrase-structure tree for a sample synthesis. Dotted-boxes around constituents indicate that they are candidates for replacement on the source side (§3.1). EN: English source sentence, HI: Hindi target sentence, CS: code-switched sentence. The *italicized segment* is the target segment to replace the *source segment* under the non-terminal SBAR.

lead to target segments q_t^1, q_t^2, \dots shown in Algorithm 1, with t removed for clarity.

3.2.1 Word alignment signals

Signals based on word alignment methods are part of the recipe in choosing the best possible q_t given the sentence pair and p_s .

GIZA score. The standard machine translation word alignment tool Giza++ (Och and Ney, 2003) uses IBM statistical word alignment models 1–5 (Fernández, 2008; Schoenemann, 2010; Brown et al., 1993; Riley and Gildea, 2012). This tool incorporates principled probabilistic formulations of the IBM models and gives a correspondence score $G[w_t^i, w_s^j]$ between target and source words for a given sentence pair. This word-pair score is used as a signal to find the best q_t .

NMT attention score. Given an attention-guided trained sequence-to-sequence neural machine translation (NMT) model (Bahdanau et al., 2015; Luong et al., 2015) and sentence pair s, t , we use the attention score matrix $A[w_t^i, w_s^j]$ as an alignment signal.

Inverse document frequency (rarity). The inverse document frequency (IDF) of a word in a corpus signifies its importance in the sentence (van Rijsbergen, 1979). We define $\mathcal{I}(w) = \sigma(a \text{ IDF}(w) - b)$ as a shifted, squashed IDF that normalizes the raw corpus-level score. Parameters a and b are empirically determined. This IDF-based signal is optionally incorporated while choosing q_t .

3.2.2 Target segment with maximum similarity

Given word-pair scores derived from either Giza++ or NMT attention described in §3.2.1, we formulate two methods for identifying target segments. First, we identify the best target segment given Giza++ scores, $G[\cdot, \cdot]$, as follows:

$$\hat{q}_t^1 \leftarrow \operatorname{argmax}_{q_t} \prod_{w_t \in q_t} \sum_{w_s \in p_s} G[w_t, w_s] \quad (1)$$

For each word in q_t , we compute the total attention score concentrated in p_s and then multiply them as if they are independent.

Second, we use the attention score learned by the NMT system of Luong et al. (2017) (a bidirectional LSTM model with attention). Essentially, given the attention score $A[\cdot, \cdot]$ between target and source words, we intend to select the target segment q_t whose maximum attention is concentrated in the given p_s .

Initial exploration of the above method revealed that the attention of a target word may spread out over several related but less appropriate source words, and accrue better overall similarity than a single more appropriate word. Here IDF can come to the rescue, the intuition being that words w_t^i and w_s^j with very different IDFs are less likely to align, because (barring polysemy and synonymy) rare (common) words in one language tend to translate to rare (common) words in another. This intuition is embodied in the improved formulation:

$$\hat{q}_t^2 \leftarrow \operatorname{argmax}_{q_t} \prod_{w_t \in q_t} \mathcal{I}(w_t) \sum_{w_s \in p_s} \mathcal{I}(w_s) A[w_t, w_s] \quad (2)$$

Informally, if a source segment contains many rare words, the target segment should also have a similar number of rare words from the target domain, and vice-versa.

3.2.3 Target segment with minimum dissimilarity

We examine an alternative method for identifying target segments that leverage the *Earth Mover’s Distance (EMD)* (Vaserštejn, 1969). Kusner et al. (2015) extended EMD to the *Word Mover Distance* to measure the dissimilarity between documents by ‘transporting’ word vectors from one document to the vectors of the other. In the same spirit, we define a dissimilarity measure between p_s and candidate target segments using EMD. We present

here EMD as a minimization over fractional transportation matrix $\mathbf{F} \in \mathbb{R}^{|q_t| \times |p_s|}$ as below:

$$\operatorname{EMD}(q_t, p_s) = \min_{\mathbf{F}} \sum_{i=1}^{|q_t|} \sum_{j=1}^{|p_s|} F_{i,j} d_{i,j} \quad (3)$$

where $\sum_i F_{i,j} = \frac{1}{|q_t|}$ and $\sum_j F_{i,j} = \frac{1}{|p_s|}$ and $d_{i,j}$ is a distance metric between a target and a source word pair, given suitable representations. Finally, we choose the target segment which is least dissimilar to a given source segment defined by the EMD. We compute $d_{i,j}$ in two ways, described below.

Attention-based distance. Here the distance between the embeddings is defined as:

$$d_{i,j}^A = 1 - A[w_t^i, w_s^j] \quad (4)$$

Giza-based distance. Similarly we can compute the distance using Giza score as:

$$d_{i,j}^G = 1 - G[w_t^i, w_s^j] \quad (5)$$

Given the two types of distances in Eq. (4)–(5) and the definition of EMD in Eq. (3) we can formulate two methods for identifying target segments:

$$\hat{q}_t^k \leftarrow \operatorname{argmin}_{q_t} \min_{\mathbf{F}} \sum_{i=1}^{|q_t|} \sum_{j=1}^{|p_s|} F_{i,j} d_{i,j}^k \quad (6)$$

where $k \in \{3, 4\}$ and $d_{i,j}^3 \equiv d_{i,j}^A$ and $d_{i,j}^4 \equiv d_{i,j}^G$.

We can also use Euclidean distance as $d_{i,j}$. However, this method requires multilingual word embeddings for every word to calculate the distance. The volume of labeled source text we can use is usually smaller than the vocabulary size, making it difficult to learn reliable word embeddings. Also, if these corpora contains informal social media text like the ones described in §4.1, then publicly available pretrained word embeddings exclude a significant percentage of them.

3.3 Projecting target segments

Given a source sentence s with designated segment p_s to replace, and target sentence t , we have by now identified four possible target segments \hat{q}_t^k where $k \in \{1, \dots, 4\}$ as described in §3.2.2–§3.2.3. We now *project* the target segment to the source sentence, meaning, (a) replace the source segment with the target segment and (b) transliterate the replacement using the Google Transliteration API to the source script². This creates four possible

²<http://www.google.com/transliterate?langpair=hi|en&text=<text>>

synthetic code-switched sentences for each instance of (s, p_s, t) . Finally, we transfer the labels of the original monolingual corpora to the generated synthetic text corpora.

3.4 Best candidate via reverse translation

From these four code-switched sentences c_1, \dots, c_4 , we wish to retain the one that retains most of the syntactic structures of the source sentence. Each code-switched sentence c_k has an associated score as defined in §3.2. We use two empirically tuned thresholds: a lower cut-off for the similarity score of c_1, c_2 and an upper cut-off for the dissimilarity score of c_3, c_4 , to improve the quality of candidates retained. These scores are not normalized and cannot be compared across different methods. Therefore, we perform a reverse translation of each candidate back to the source language using the Google translation API to obtain \tilde{s} . We retain the candidate whose retranslated version \tilde{s} has the highest BLEU score (Papineni et al., 2002) wrt s . In case of a tie, we select the candidate with maximum word overlap with s .

3.5 Thresholding and stratified sampling

In addition to retaining only the best among code-switched candidates $c_{1,\dots,4}$, we discard the winner if its BLEU score is below a tuned threshold. Further, we sample source sentences such that the surviving populations of sentiment labels of the code-switched sentences match the populations in the low-resource evaluation corpus. Another tuned system parameter is the amount of synthetic text to generate to supplement the gold text.

We do not depend on any domain coherence between the source corpus used to synthesize text and the gold ‘payload’ corpus — this is the more realistic situation. Our expectation, therefore, is that adding some amount of synthetic text should improve sentiment prediction, but excessive amounts of off-domain synthetic text may hurt it. In our experiments we grid search the synthetic:gold ratio between 1/4 and 2 using 3-fold cross validation.

4 Experiments

We demonstrate the effectiveness of augmenting gold code-switched text with synthetic code-switched text. We also measure the usefulness of synthetic text without gold text. In this section, we will first describe the data sets used to generate the synthetic text and then the resource-poor labeled code-switched text used for evaluation. Next, we will present the method used for sentiment detection, baseline performance, and finally our performance, along with a detailed comparative analysis.

4.1 Source corpora for text synthesis

We use publicly available monolingual sentiment-tagged (positive, negative or neutral) gold corpora in the source language.

ACL. Dong et al. (2014) released about 6000 manually labeled English tweets.

Election. Wang et al. (2017) published about 5000 human-labeled English tweets.

Mukherjee. This data set contains about 8000 human-labeled English tweets (Mukherjee and Bhattacharyya, 2012; Mukherjee et al., 2012).

Semeval shared task. This provides about 10000 human-labeled English tweets (Rosenthal et al., 2017).

Union. This is the union of above mentioned different data sets.

Hate speech. We collected 15K tagged English tweets from (Founta et al., 2018) which consists of 4.7K *abusive*, 1.7K *hateful* and 4K *normal* tweets.

We picked Spanish, which is homologous to English, and Hindi and Bengali, which are comparatively dissimilar to English, for our experiments. We translated these monolingual tweets to Spanish, Hindi and Bengali using Google Translation API³ and used as parallel corpus to train attention-based NMT models and statistical MT model (GIZA) to learn the word alignment signals as described in §3.2.1.

4.2 Preliminary qualitative analysis

Analysis of texts synthesized by various mechanisms proposed in §3.2 shows that similarity based methods contribute 82–85% of the best

³<https://translation.googleapis.com>

candidates and the rest come from dissimilarity based methods. Similarity-based methods using NMT attention and Giza perform well because the segments selected for replacement often constitute nouns and entity mentions, which have a very strong alignment in the corresponding target segment. NMT attention and Giza-EMD perform well when segments contract or expand in translation.

4.3 Low-resource evaluation corpora

Each data set below was divided into 70% training, 10% validation and 20% testing folds. The training fold was (or was not) augmented with synthetic labeled text to train sentiment classifiers, which were then applied on the test fold judge the quality of synthesis.

ES-EN (English-Spanish). We collected 2883 labeled tweets specified by [Vilares et al. \(2015\)](#).

HI-EN,FB (Hindi-English, Facebook). [Joshi et al. \(2016\)](#) released around 4000 labeled code-switched sentences from the Facebook timeline of Narendra Modi (Indian Prime Minister) and Salman Khan (Bollywood actor).

HI-EN,TW (Hindi-English, Twitter). This is a shared task from ICON 2017 ([Patra et al., 2018](#)) with 15575 instances.

BN-EN (Bengali-English). This is another shared task from ICON 2017 ([Patra et al., 2018](#)) with 2499 instances.

HI-EN, Hate speech. [Bohra et al. \(2018\)](#) published 4000 manually-labeled code-switched Hindi-English tweets: 1500 exhibiting *hate speech* and 2500 *normal*. We also found a significant number of abusive tweets marked *hate speech*. For uniformity, we merged hate speech tweets and abusive tweets.

4.4 Sentiment classifier

We adopt the sub-word-LSTM system of [Joshi et al. \(2016\)](#). We prefer this over feature-based methods because (a) feature extraction for code-switched text is very difficult, and varies widely across language pairs, and (b) the vocabulary is large and informal, with many tokens outside standard (full-) word embedding vocabularies.

Loss functions. If the sentiment labels

$\{-1, 0, +1\}$ are regarded as categorical, cross-entropy loss is standard. However, prediction errors between the extreme polarities $\{-1, +1\}$ need to be penalized more than errors between $\{-1, 0\}$ or $\{0, +1\}$. Hence, we use ordinal cross-entropy loss ([Niu et al., 2016](#)), introducing a weight factor proportional to the order of intended penalty multiplied with the cross entropy loss. On the test fold, we report 0/1 accuracy and per class micro-averaged F1 score.

Baseline and prior art. Our baseline scenario is a self-contained train-dev-test split of the gold corpus. The primary prior art is the work of [Pratapa et al. \(2018\)](#).

Feature space coherence. Our source corpora are quite unrelated to the gold corpora. Table 1 shows that the average Euclidean distance between feature space of gold training and testing texts is much lower than that between gold and synthetic texts. While this may be inescapable in a low-resource situation, the gold baseline does not pay for such decoherence, which can lead to misleading conclusions.

	HI,EN,TW	HI,EN,FB	ES,EN	BN,EN
Mukherjee	2.47 (2.33)	3.82 (2.26)	1.64 (1.64)	5.18 (2.50)
ACL	2.21 (2.13)	3.72 (2.24)	2.09 (1.73)	4.11 (2.67)
Semeval	2.23 (2.11)	4.04 (2.26)	1.69 (1.67)	3.63 (2.59)
Election	2.40 (2.12)	6.27 (2.67)	1.58 (1.49)	5.23 (2.43)
Union	2.55 (2.15)	3.80(2.65)	1.65 (1.53)	5.48 (2.56)
Gold	2.05	1.87	1.64	1.83

Table 1: Average pairwise Euclidean distance between training data and test data features. Rows correspond to standalone (respectively, augmented) text for training. Gray: reference distance of gold test from gold train. Red: largest distance observed.

Training regimes. Absence of coherent tagged gold text may lead to substantial performance loss. Hence, along with demonstrating the usefulness of augmenting natural with synthetic text, we also measure the efficacy of synthetic text on its own. We train the SA classifier with three labeled corpora: (a) limited gold code-switched text, (b) gold code-switched text augmented with synthetic text and (c) only synthetic text. Then we evaluate the resulting models on labeled gold code-switched test fold.

Train \ Test	HI_EN (TW)	HI_EN (FB)	ES_EN	BN_EN	HI_EN (TW)	HI_EN (FB)	ES_EN	BN_EN
	Categorical cross entropy training loss				Ordinal cross entropy training loss			
Mukherjee	53.76 (54.69)	64.82 (65.85)	47.06 (46.36)	49.79 (57.40)	53.79 (53.53)	59.66 (64.57)	44.85 (43.07)	51.00 (49.40)
ACL	51.80 (52.59)	62.59 (65.33)	44.80 (48.84)	52.59 (59.81)	52.68 (53.76)	62.72 (64.22)	45.69 (50.31)	50.08 (51.60)
Semeval	52.99 (54.19)	65.33 (65.46)	47.36 (44.69)	53.40 (59.99)	52.99 (53.83)	63.26 (64.66)	47.36 (44.64)	53.14 (57.59)
Election	52.84 (54.59)	65.59 (66.80)	43.07 (43.07)	57.99 (59.00)	52.89 (54.64)	64.88 (65.26)	45.21 (45.80)	53.40 (55.60)
Union	53.28 (54.64)	65.50 (65.99)	44.24 (45.32)	57.23 (59.89)	53.65 (53.69)	64.30 (67.65)	44.04 (46.00)	53.40 (57.40)
MSR	54.50	65.58	48.14	59.79	53.69	62.80	47.50	52.8
Gold	52.26	65.37	42.6	55.19	52.34	64.29	45.20	50.39

Table 2: Accuracy (%) on 20% test data after training with augmented and only gold text. Rows correspond to sources of augmentation. In most cells we show (A) no thresholding or stratification and (B) with thresholding and stratification (within brackets). Gray: reference accuracy with only gold training. Blue: A or B or MSR outperforms gold. Green: B performs best. Row ‘MSR’ uses text synthesized by Pratapa et al. (2018).

4.5 Sentiment detection accuracy

Table 2 shows the benefits of augmenting natural with synthetic text. Test accuracy increases further (shown in brackets) if thresholding and stratified sampling are used. Gains for HI_EN,TW, HI_EN,FB, ES_EN and BN_EN are 1.5% (2.43%), 0.23% (1.43%), 4.76% (6.24%), and 2.8% (4.8%) respectively. Categorical cross-entropy loss was used here. Similar improvements in accuracy of are observed after training with ordinal loss function. Our conclusion is that *careful augmentation with synthetic data can lead to useful gain in accuracy*. Moreover, by selecting synthetic text which is syntactically more natural, even larger gains can be achieved. Notably, the distance between training and test features (Table 1) is negatively correlated with accuracy gain (Pearson correlation coefficient of -0.48).

Comparison with Pratapa et al. (2018).

They depend on finding correspondences between constituency parses of the source and target sentences. However, the common case is that a constituency parser is unavailable or ineffective for the target language, particularly for informal social media. They are thus restricted to synthesizing text from only a subset of monolingual data. Training SA with natural text augmented with their synthesized text leads to poorer accuracy, albeit by a small amount, than using CSGen. The performance is worse for target languages that are more resource-poor.

Ordinal vs. categorical loss. Table 2 shows that ordinal loss helps when the neutral label dominates. However, neither is a clear winner and the gains are small. Therefore, we use categorical loss henceforth.

Choice of monolingual corpus. Across all monolingual corpora, *Election* performs consistently well. Best test performance on HI_EN,TW was obtained by synthesizing from the *Mukherjee* corpus. Text synthesized from *Election* provides the best results for HI_EN,FB for both setups. The performance of *Union* is also good but not the best. This is because although a larger and diverse amount of data is available which ensures its quality, the Euclidean distance between test data and some individual corpora is still large.

	Categorical Cross Entropy training			Ordinal Cross Entropy training		
	Pos	Neu	Neg	Pos	Neu	Neg
	HI_EN,TW					
CSGen	0.52	0.62	0.38	0.55	0.63	0.34
Gold	0.48	0.63	0.24	0.50	0.62	0.35
	HI_EN,FB					
CSGen	0.59	0.73	0.56	0.62	0.71	0.55
Gold	0.60	0.74	0.54	0.60	0.71	0.44
	ES_EN					
CSGen	0.38	0.53	0.37	0.48	0.50	0.42
Gold	0.47	0.44	0.41	0.40	0.53	0.43
	BN_EN					
CSGen	0.63	0.49	0.58	0.55	0.47	0.59
Gold	0.55	0.51	0.65	0.37	0.49	0.61

Table 3: F1 score for each class prediction. Blue: CSGen is better than Gold.

4.6 Sentiment detection F1 score

Beyond 0/1 accuracy, Table 3 shows F1 score gains. *Election* yields consistently good results. We have reported the F1 score gain for different sentiment classes only for *Election* in Table 3 for brevity. Augmenting synthetic data with gold data yields better F1 score than training only with gold tagged data. Also, it is interesting to observe that there is a sharp drop of F1 score for HI_EN,FB and BN_EN data sets for Gold data while training with or-

dinal cross entropy function across all the sentiment labels. As described in §4.5, this is due to non-discriminative features. However, mixing them with synthetic data helps in achieving better results.

Train \ Test	HI_EN (TW)	HI_EN (FB)	ES_EN	BN_EN
Mukherjee	46.22	48.98	39.76	44.42
ACL	40.33	49.96	38.40	47.81
Semeval	45.80	48.38	39.18	45.99
Election	47.22	48.78	31.20	42.44
Union	43.50	49.80	41.9	41.2
Gold	52.26	65.37	42.6	55.19

Table 4: Percent accuracy on 20% test data after training on only synthetic and only gold text. Each row corresponds to a source. Red: Accuracy achieved with only gold training. Blue: The closest accuracy achieved to best.

4.7 Performance of standalone synthetic data

The accuracy of using only synthetic data as training is reported in Table 4. We can see that for EN_HI,TW and EN_ES the synthetic data is very close to the gold data performance (lagging by 5.04% and 2.84%). However, it performed poorly for HI_EN,FB and BN_EN dataset. This is because there is heavy mismatch between the synthetic text set generated and the test data distribution (Table 1) in these two dataset. The Pearson rank correlation coefficient between the distance (between test and training set) measures and relative accuracy gain is highly negative, -0.66 .

To further establish the importance of domain coherence, we report on an experiment performed with HI_EN,FB gold dataset. This dataset has texts corresponding to two different entities namely *Narendra Modi* and *Salman Khan*. Training SA with natural text corresponding to one entity and testing on the rest leads to a steep accuracy drop from 65.37% to 52.32%.

Failure category	Example sentence	Gold	Pred.
Different polarity keywords	• twin brothers lost in fair reunited in adulthood amidst dramatic circumstances ei themer movie akhon ar viewers der attract kore na	0	+1
	• kon se duniya hai muslim bhai jo kafir hai usko acha oor great bol ray hai	-1	+1
Ambiguous meaning	• elizabq ellen quiere entrevistat julianna margulies clooney says she is hard one to crack	-1	0
	• hum kam se kam fight ker haaray lekin tum loog zillat ki maut maaray gaye	+1	-1

Table 5: Sample sentiment prediction failures. Red, green: Negative, positive polarity words.

4.8 Error analysis

We found two dominant error modes where synthetic augmentation confuses the system. Table 5 shows a few examples. The first error mode can be triggered by the presence of words of different polarities, one polarity more common than the other, and the gold label being the minority polarity. The second error mode is prevalent when the emotion is weak or mixed. Either there is no strong opinion, or there are two agents, one regarded positively and the other negatively.

4.9 Hate speech detection results

Table 6 shows hate speech detection results. Training with only synthetic text after thresholding and stratified sampling outperforms training with only gold-tagged text by 4% F1, and using both gold and synthetic text gives a F1 boost of 6% beyond using gold alone. Remarkably, synthetic text alone outperforms gold text, because gold text has high class imbalance, leading to poorer prediction. Because we can create arbitrary amounts of synthetic text, we can balance the labels to achieve better prediction.

	Prec	Recall	F-score
Only synthetic	0.58 (0.63)	0.60 (0.63)	0.51 (0.52)
Synthetic +Gold	0.59 (0.60)	0.63 (0.63)	0.53 (0.54)
Gold	0.40	0.62	0.48

Table 6: Hate speech results (3-fold cross val.). In most cells we show performance without thresholding and stratification (within bracket with thresholding and stratification). Green: Best performance in each column.

5 Conclusion

Code-mixing is an important and rapidly evolving mechanism of expression among multilingual populations on social media. Monolingual sentiment analysis techniques perform poorly on code-mixed text, partly because code-mixed text often involves resource-poor languages. Starting from sentiment-labeled text in resource-rich source languages, we propose an effective method to synthesize labeled code-mixed text without designing switching grammars. Augmenting scarce natural text with synthetic text improves sentiment detection accuracy.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.
- Utsab Barman, Amitava Das, Joachim Wagner, and Jennifer Foster. 2014. Code mixing: A challenge for language identification in the language of social media. In *Proceedings of the first workshop on computational approaches to code switching*, pages 13–23.
- Rafiya Begum, Kalika Bali, Monojit Choudhury, Koustav Rudra, and Niloy Ganguly. 2016. Functions of code-switching in Tweets: An annotation scheme and some initial experiments. *Proceedings of LREC*.
- Gayatri Bhat, Monojit Choudhury, and Kalika Bali. 2016. Grammatical constraints on intra-sentential code-switching: From theories to working models. *arXiv preprint arXiv:1612.04538*.
- Aditya Bohra, Deepanshu Vijay, Vinay Singh, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018. A dataset of hindi-english code-mixed social media text for hate speech detection. In *Proceedings of the Second Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, pages 36–41.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, and Dai. 2015. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*.
- Peter F Brown, Vincent J Della Pietra, Stephen A Della Pietra, and Robert L Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263–311.
- Gokul Chittaranjan, Yogarshi Vyas, Kalika Bali, and Monojit Choudhury. 2014. Word-level language identification using CRF: Code-switching shared task report of MSR India system. In *Proceedings of The First Workshop on Computational Approaches to Code Switching*, pages 73–79.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. Adaptive recursive neural network for target-dependent twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 49–54.
- Pablo Malvar Fernández. 2008. *Improving Word-to-word Alignments Using Morphological Information*. Ph.D. thesis, San Diego State University.
- Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. *arXiv preprint arXiv:1802.00393*.
- CJ Hutto Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*. Available at (20/04/16) <http://comp.social.gatech.edu/papers/icwsm14.vader.hutto.pdf>.
- Aditya Joshi, Ameya Prabhu, Manish Shrivastava, and Vasudeva Varma. 2016. Towards subword level compositions for sentiment analysis of hindi-english code mixed text. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2482–2491.
- Anjali Kannan and Oriol Vinyals. 2017. Adversarial evaluation of dialogue models. *arXiv preprint arXiv:1701.08198*.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International Conference on Machine Learning*, pages 957–966.
- Minh-Thang Luong, Eugene Brevdo, and Rui Zhao. 2017. Neural machine translation (seq2seq) tutorial. <https://github.com/tensorflow/nmt>.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of EMNLP*.
- Umar Maqsood. 2015. Synthetic text generation for sentiment analysis. In *Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.
- Subhabrata Mukherjee and Pushpak Bhattacharyya. 2012. Sentiment analysis in twitter with lightweight discourse analysis. *Proceedings of COLING 2012*, pages 1847–1864.
- Subhabrata Mukherjee, Akshat Malu, Balamurali AR, and Pushpak Bhattacharyya. 2012. Twisent: a multistage system for analyzing sentiment in twitter. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 2531–2534. ACM.
- Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. 2016. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4920–4928.

- 900 Franz Josef Och and Hermann Ney. 2003. A
901 systematic comparison of various statistical
902 alignment models. *Computational Linguistics*,
903 29(1):19–51. 950
- 904 Kishore Papineni, Salim Roukos, Todd Ward, and
905 Wei-Jing Zhu. 2002. BLEU: a method for au-
906 tomatic evaluation of machine translation. In
907 *Proceedings of ACL*. 951
- 908 Braja Gopal Patra, Dipankar Das, and Amitava
909 Das. 2018. Sentiment analysis of code-mixed
910 indian languages: An overview of sail_code-
911 mixed shared task@ icon-2017. *arXiv preprint*
912 *arXiv:1803.06745*. 952
- 913 Adithya Pratapa, Gayatri Bhat, Monojit Choud-
914 hury, Sunayana Sitaram, Sandipan Dandapat,
915 and Kalika Bali. 2018. Language modeling for
916 code-mixing: The role of linguistic theory based
917 synthetic data. In *Proceedings of ACL*. 953
- 918 C J van Rijsbergen. 1979. *Information Re-*
919 *trieval*. Butterworths, London. Online at
920 <http://www.dcs.gla.ac.uk/Keith/Preface.html>. 954
- 921 Darcey Riley and Daniel Gildea. 2012. Improv-
922 ing the IBM alignment models using variational
923 Bayes. In *Proceedings of ACL*. 955
- 924 Sara Rosenthal, Noura Farra, and Preslav Nakov.
925 2017. SemEval-2017 task 4: Sentiment analysis
926 in Twitter. In *Proceedings of the 11th Inter-*
927 *national Workshop on Semantic Evaluation, Sem-*
928 *Eval '17, Vancouver, Canada*. Association for
929 Computational Linguistics. 956
- 930 Koustav Rudra, Shruti Rijhwani, Rafiya Begum,
931 Kalika Bali, Monojit Choudhury, and Niloy
932 Ganguly. 2016. Understanding language prefer-
933 ence for expression of opinion and sentiment:
934 What do hindi-english speakers do on twitter?
935 In *Proceedings of the 2016 Conference on Em-*
936 *pirical Methods in Natural Language Processing*,
937 pages 1131–1141. 957
- 938 David Sankoff and Shana Poplack. 1981. A for-
939 mal grammar for code-switching. *Research on*
940 *Language & Social Interaction*, 14(1):3–45. 958
- 941 Thomas Schoenemann. 2010. Computing optimal
942 alignments for the IBM-3 translation model. In
943 *Proceedings of CoNLL*. 959
- 944 Royal Sequiera, Monojit Choudhury, and Kalika
945 Bali. 2015. POS tagging of Hindi-English code
946 mixed text from social media: Some machine
947 learning experiments. In *Proceedings of Inter-*
948 *national Conference on NLP*. 960
- 949 Shashank Sharma, PYKL Srinivas, and
950 Rakesh Chandra Balabantaray. 2015. Text
951 normalization of code mix and sentiment
952 analysis. In *Advances in Computing, Com-*
953 *munications and Informatics (ICACCI), 2015*
954 *International Conference on*, pages 1468–1473.
955 IEEE. 961
- 956 Tamar Solorio, Elizabeth Blair, Suraj Mahar-
957 jan, Steven Bethard, Mona Diab, Mahmoud
958 Ghoneim, Abdelati Hawwari, Fahad AlGhamdi,
959 Julia Hirschberg, Alison Chang, et al. 2014.
960 Overview for the first shared task on language
961 identification in code-switched data. In *Proceed-*
962 *ings of the First Workshop on Computational*
963 *Approaches to Code Switching*, pages 62–72. 962
- 964 Tamar Solorio and Yang Liu. 2008. Part-
965 of-speech tagging for English-Spanish code-
966 switched text. In *Proceedings of EMNLP*. 963
- 967 Leonid Nisonovich Vaseršteĭn. 1969. Markov pro-
968 cesses over denumerable products of spaces de-
969 scribing large systems of automata. *Problems of*
970 *Information Transmission*, 5(3):47–52. 964
- 971 David Vilares, Miguel A Alonso, and Carlos
972 Gómez-Rodríguez. 2015. Sentiment analysis on
973 monolingual, multilingual and code-switching
974 twitter corpora. In *Proceedings of the 6th Work-*
975 *shop on Computational Approaches to Subjectiv-*
976 *ity, Sentiment and Social Media Analysis*, pages
977 2–8. 965
- 978 Yogarshi Vyas, Spandana Gella, Jatin Sharma,
979 Kalika Bali, and Monojit Choudhury. 2014.
980 POS tagging of English-Hindi code-mixed social
981 media content. In *Proceedings of EMNLP*. 966
- 982 Bo Wang, Maria Liakata, Arkaitz Zubiaga, and
983 Rob Procter. 2017. Tdparse: Multi-target-
984 specific sentiment recognition on twitter. In
985 *Proceedings of the 15th Conference of the Eu-*
986 *ropean Chapter of the Association for Computa-*
987 *tional Linguistics: Volume 1, Long Papers*, vol-
988 *ume 1*, pages 483–493. 967
- 989 Yizhe Zhang, Zhe Gan, and Lawrence Carin. 2016.
990 Generating text via adversarial training. In
991 *NIPS workshop on Adversarial Training*. 968
- 992 Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ri-
993 cardo Henao, Dinghan Shen, and Lawrence
994 Carin. 2017. Adversarial feature match-
995 ing for text generation. *arXiv preprint*
996 *arXiv:1706.03850*. 969
- 997 Muhua Zhu, Yue Zhang, Wenliang Chen, Min
998 Zhang, and Jingbo Zhu. 2013. Fast and accurate
999 shift-reduce constituent parsing. In *Proceedings*
1000 *of ACL*. 970