

String Transduction with Target Language Models and Insertion Handling

Garrett Nicolai[†] and Saeed Najafi[‡] and Grzegorz Kondrak[‡]

[†]Department of Computer Science
Johns Hopkins University
gnicola2@jhu.edu

[‡]Department of Computing Science
University of Alberta
{snajafi, gkondrak}@ualberta.ca

Abstract

Many character-level tasks can be framed as sequence-to-sequence transduction, where the target is a word from a natural language. We show that leveraging target language models derived from unannotated target corpora, combined with a precise alignment of the training data, yields state-of-the-art results on cognate projection, inflection generation, and phoneme-to-grapheme conversion.

1 Introduction

Many natural language tasks, particularly those involving character-level operations, can be viewed as sequence-to-sequence transduction (Figure 1). Although these tasks are often addressed in isolation, they share a common objective — in each case, the output is a word in the target language.

The hypothesis that we investigate in this paper is that a single task- and language-independent system can achieve state-of-the-art results by leveraging unannotated target language corpora that contain thousands of valid target word types. We focus on low-data scenarios, which present a challenge to neural sequence-to-sequence models because sufficiently large parallel datasets are often difficult to obtain. To reinforce transduction models trained on modest-sized collections of source-target pairs, we leverage monolingual text corpora that are freely available for hundreds of languages.

Our approach is based on discriminative string transduction, where a learning algorithm assigns weights to features defined on aligned source and target pairs. At test time, an input sequence is converted into the highest-scoring output sequence. Advantages of discriminative transduction include an aptitude to derive effective models from small training sets, as well as the capability to incorporate diverse sets of features. Specifically, we build

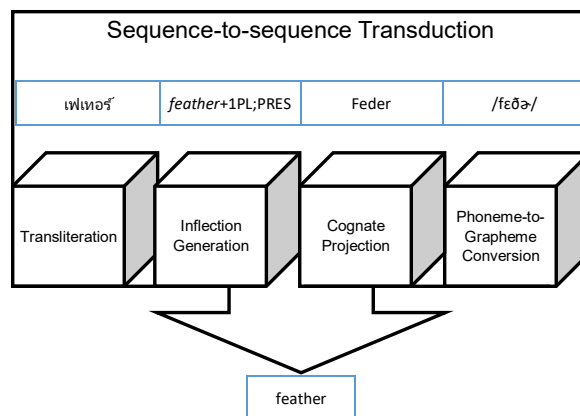


Figure 1: Illustration of four character-level sequence-to-sequence prediction tasks. In each case, the output is a word in the target language.

upon DIRECTL+ (Jiampojarn et al., 2010), a string transduction tool which was originally designed for grapheme-to-phoneme conversion.

We present a new system, DTLM, that combines discriminative transduction with character and word language models (LMs) derived from large unannotated corpora. Target language modeling is particularly important in low-data scenarios, where the limited transduction models often produce many ill-formed output candidates. We avoid the error propagation problem which is inherent in pipeline approaches by incorporating the LM feature sets directly into the transducer.

In addition, we bolster the quality of transduction by employing a novel alignment method, which we refer to as *precision alignment*. The idea is to allow null substrings (*nulls*) on the source side during the alignment of the training data, and then apply a separate aggregation algorithm to merge them nulls with adjacent non-empty substrings. This method yields precise many-to-many alignment links that lead to improved transduction accuracy.

The contributions of this paper include the following. (1) A novel method of incorporating strong target language models directly into discriminative transduction. (2) A novel approach to unsupervised alignment that is particularly beneficial in low-resource settings. (3) An extensive experimental comparison to previous models on multiple tasks and languages, which includes state-of-the-art results on inflection generation, cognate projection, and phoneme-to-grapheme generation. (4) Publicly available implementation of the proposed methods. (5) Three new datasets for cognate projection.

2 Baseline methods

In this section, we describe the baseline methods, including the alignment of the training data, the feature sets of DirecTL+ (henceforth DTL), and reranking as a way of incorporating corpus statistics.

2.1 Alignment

Before a transduction model can be derived from the training data, the pairs of source and target strings need to be aligned, in order to identify atomic substring transformations. The unsupervised M2M aligner (Jiampojarn et al., 2007) employs the Expectation-Maximization (EM) algorithm with the objective of maximizing the joint likelihood of its aligned source and target pairs. The alignment involves every source and target character. The pairs of aligned substrings may contain multiple characters on both the source and target sides, yielding *many-to-many* (M-M) alignment links.

DTL excludes insertions from its set of edit operations because they greatly increase the complexity of the generation process, to the point of making it computationally intractable (Barton, 1986). Therefore, the M2M aligner is forced to avoid nulls on the source side by incorporating them into many-to-many links during the alignment of the training data. Although many-to-many alignment models are more flexible than 1-1 models, they also generally require larger parallel datasets to produce correct alignments. In low-data scenarios, especially when the target strings tend to be longer than the source strings, this approach often yields sub-optimal alignments (e.g., the leftmost alignment in Figure 2).




Figure 2: Examples of different alignments in phoneme-to-letter conversion. The underscore denotes a null substring.

2.2 Features

DTL is a feature-rich, discriminative character transducer, which searches for a model-optimal sequence of character transformation operations for its input. The core of the engine is a dynamic programming algorithm capable of transducing many consecutive characters in a single operation, also known as a semi-Markov model. Using a structured version of the MIRA algorithm (McDonald et al., 2005), the training process assigns weights to each feature, in order to achieve maximum separation of the gold-standard output from all others in the search space.

DTL uses a number of feature templates to assess the quality of an operation: source context, target n -gram, and joint n -gram features. Context features conjoin the rule with indicators for all source character n -grams within a fixed window of where the rule is being applied. Target n -grams provide indicators on target character sequences, describing the shape of the target as it is being produced, and may also be conjoined with the source context features. Joint n -grams build indicators on rule sequences, combining source and target context, and memorizing frequently-used rule patterns. An additional copy feature generalizes the identity function from source to target, which is useful if there is an overlap between the input and output symbol sets.

2.3 Reranking

The target language modeling of DTL is limited to a set of binary n -gram features, which are based exclusively on the target sequences from the parallel training data. This shortcoming can be remedied by taking advantage of large unannotated corpora that contain thousands of examples of valid target words.

Nicolai et al. (2015) propose to leverage corpus statistics by reranking the n -best list of candidates generated by the transducer. They report consistent modest gains by applying an SVM-based

reranker, with features including a word unigram corpus presence indicator, a normalized character language model score, and the rank and normalized confidence score generated by DTL. However, such a pipeline approach suffers from error propagation, and is unable to produce output forms that are not already present in the n -best list. In addition, training a reranker requires a held-out set that substantially reduces the amount of training data in low-data scenarios.

3 Methods

In this section, we describe our novel extensions: precision alignment, character-level target language modeling, and corpus frequency. We make the new implementation publicly available.¹

3.1 Precision Alignment

We propose a novel alignment method that produces accurate many-to-many alignments in two stages. The first step consists of a standard 1-1 alignment, with nulls allowed on either side of the parallel training data. The second step removes the undesirable nulls on the source side by merging the corresponding 0-1 links with adjacent 1-1 links. This alignment approach is superior to the one described in Section 2.1, especially in low-data scenarios when there is not enough evidence for many-to-many links.²

Our precision alignment is essentially a 1-1 alignment with 1-M links added when necessary. In a low-resource setting, an aligner is often unable to distinguish valid M-M links from spurious ones, as both types will have minimal support in the training data. On the other hand, good 1-1 links are much more likely to have been observed. By limiting our first pass to 1-1 links, we ensure that only good 1-1 links are posited; otherwise, an insertion is predicted instead. On the second pass, the aligner only needs to choose between a small number of alternatives for merging the insertions, increasing the likelihood of a good alignment, and, subsequently, correct transduction.

Consider the example in Figure 2 where 5 source phonemes need to be aligned to 7 target letters. The baseline approach incorrectly links the

¹<http://github.com/GarrettNicolai/DTLM> and /M2MP

²The improvement in the alignment quality is relative to the performance of our transduction system, as we demonstrate in Section 4.6 — the alignments are not necessarily optimal from the linguistic point of view.

```

1: Algorithm:ForwardInsertionMerging
2: Input:  $(x^T, y^V)$ 
3: Output:  $(\alpha^{T+1, V+1})$ 
4:  $CI = 0, PI = 0$ 
5: for  $t=0; t \leq T$  do
6:   if  $t > 0$  and  $x_t == \_$  then
7:      $CI += 1$ 
8:   else
9:      $PI = CI$ 
10:     $CI = 0$ 
11:   for  $v=0; v \leq V$  do
12:     if  $t - CI == 0$  then
13:        $\alpha_{t,v} = 1$ 
14:       continue {insertions at the start of the word}
15:     if  $t > 0$  or  $v > 0$  then
16:        $\alpha_{t,v} = 0$ 
17:     if  $t > 0$  and  $v > 0$  then
18:       for  $k = 0; k \leq PI$  do
19:          $\alpha_{t,v} += \delta(x_{t-CI-k}^t, y_{v-CI-k}^v) * \alpha_{t-CI-k-1, v-CI-k-1}$ 

```

Figure 3: The forward step of M2M, modified to merge insertions to adjacent characters.

letter ‘a’ with the phoneme /w/ (the leftmost alignment in the diagram). Our first-pass 1-1 alignment (in the middle), correctly matches /ɔ/ to ‘a’, while ‘l’ is treated as an insertion. On the second pass, our algorithm merges the null with the preceding 1-1 link. By contrast, the second insertion, which involves /ə/, is merged with the substitution that follows it (the rightmost alignment).

Figure 3 demonstrates how we modify the forward step to merge insertions with adjacent substitutions; similar modifications are made for the backward step, expectation, and decoder. The input consists of a source string \mathbf{x} of length T , and a target string \mathbf{y} of length V . Both \mathbf{x} and \mathbf{y} may contain underscores, which represent nulls from the first alignment pass. The α score represents the sum of the likelihoods of all paths that have been traversed through source character t and target character v . In a 1-1 alignment, all α scores accumulate along the diagonal, while in a many-to-many alignment, other cells of the α matrix may be filled. Our precision alignment is a compromise between these two methods: we consider adjacent characters, but force the α score to accumulate on the diagonal. By allowing insertions and deletions in the first pass, we force \mathbf{x} and \mathbf{y} to be of equal length. We then perform a 1-1 alignment, expanding the alignment size only when the source character is a null.

We supplement the forward algorithm of M2M with two counters: PI is the number of adjacent insertions immediately to the left of the current character, while CI is the number of insertions that

have been encountered since the last substitution. The loop at line 18 executes the α score accumulation, where δ is the likelihood of a specific sequence alignment, effectively merging insertions with adjacent substitutions. An extended example that illustrates the operation of the algorithm is included in the Appendix.

3.2 Character-level language model

In order to incorporate a stronger character language model into DTL, we propose an additional set of features that directly reflect the probability of the generated subsequences. We train a character-level language model on a list of types extracted from a raw corpus in the target language, applying Witten-Bell smoothing and back-off for unseen n -grams. During the generation process, the transducer incrementally constructs target sequences character-by-character. The normalized log-likelihood score of the current output sequence is computed according to the character language model.

For consistency with other sets of features, we convert these real-valued scores into binary indicators by means of binning. Development experiments led to the creation of bins that represent a normal distribution around the mean likelihood of words. Features fire in a cumulative manner, and a final feature fires only if no bin threshold is met. For example, if a sequence has a log-likelihood of -0.85 , the feature for -0.9 fires, as does the one for -0.975 , and -1.05 , etc.

3.3 Corpus frequency counts

We also extend DTL with a feature set that can be described as a unigram word-level language model. The objective is to bias the model towards generating output sequences that correspond to words observed in a large corpus. Since an output sequence can only be matched against a word list after the generation process is complete, we propose to estimate the final frequency count for each prefix considered during the generation process. Following [Cherry and Suzuki \(2009\)](#) we use a prefix trie to store partial words for reference in the generation phase. We modify their solution by also storing the count of each prefix, calculated as the sum of all of the words in which the prefix occurs.

As with our language model features, unigram features are binned. A unigram feature fires if the

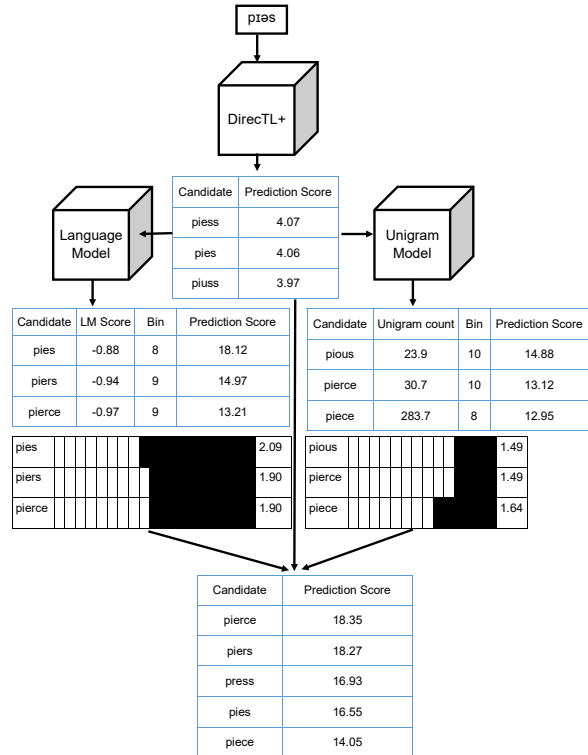


Figure 4: An example of transduction with target language models. Black cells represent firing features.

count of the generated sequence surpasses the bin threshold, in a cumulative manner.

We found that the quality of the target unigram set can be greatly improved by language-based corpus pruning. Although unannotated corpora are more readily available than parallel ones, they are often noisier. Specifically, crowd-sourced corpora such as Wikipedia contain many English words that can unduly influence our unigram features. In order to mitigate this problem, we preprocess our corpora by removing all word unigrams that have a higher probability in an English corpus than in a target-language corpus.

Consider an example of how our new features benefit a transduction model, shown in Figure 4. Note that although we portray the extensions as part of a pipeline, their scores are incorporated jointly with DTL’s other features. The top- n list produced by the baseline DTL for the input phoneme sequence */pɪəs/* fails to include the correct output *piēce*. However, after the new language model features are added, the correct form makes its way to the top predictions. The new features combine with the original features of DTL, so that the high unigram count of *piēce* is not sufficient to make it the top prediction on the right

side of the diagram. Only when both sets of new features are incorporated does the system manage to produce the correct form, as seen at the bottom of the diagram.

4 Experiments

In this section, we present the results of our experiments on four different character-level sequence-to-sequence tasks: transliteration, inflection generation, cognate projection, and phoneme-to-grapheme conversion (P2G). In order to demonstrate the generality of our approach, the experiments involve multiple systems and datasets, in both low-data and high-data scenarios.

Where low-data resources do not already exist, we simulate a low-data environment by sampling an existing larger training set. Low-data training sets consist of 100 training examples, 1000 development examples, and 1000 held-out examples, except for cognate projection, where we limit the development set to 100 training examples, and the held-out set to the remaining examples. An output is considered correct if it exactly matches any of the targets in the reference data.

4.1 Systems

We evaluate DTLM, our new system, against two strong baselines and two competitive tools. Parameter tuning was performed on the same development sets for all systems.

We compare against two baselines. The first is the standard DTL, as described in Section 2.2. The second follows the methodology of Nicolai et al. (2015), augmenting DTL with a reranker (DTL+RR), as described in Section 2.3. Both baselines use the default 2-2 alignment with deletions produced by the M2M aligner. We train the reranker using 10-fold cross validation on the training set, using the reranking method of Joachims (2002). Due to the complexity of its setup on large datasets, we omit DTL+RR in such scenarios. Except where noted otherwise, we train 4-gram character language models using the CMU toolkit³ with Witten-Bell smoothing on the UniMorph corpora of inflected word forms.⁴ Word counts are determined from the first one million lines of the corresponding Wikipedia dumps.

We also compare against Sequitur (SEQ), a generative string transduction tool based on joint

source and target n -grams (Bisani and Ney, 2008), and a character-level neural model (RNN). The neural model uses the encoder-decoder architecture typically used for NMT (Sutskever et al., 2014). The encoder is a bi-directional RNN applied to randomly initialized character embeddings; we employ a soft-attention mechanism to learn an aligner within the model. The RNN is trained for a fixed random seed using the Adam optimizer, embeddings of 128 dimensions, and hidden units of size 256. We use a beam of size 10 to generate the final predictions. We experimented with the alternative neural approach of Makarov et al. (2017), but found that it only outperforms our RNN when the source and target sides are largely composed of the same set of symbols; therefore, we only use it for inflection generation.

4.2 Transliteration

Transliteration is the task of converting a word from a source to a target script on the basis of the word’s pronunciation.

Our low-resource data consists of three back-transliteration pairs from the 2018 NEWS Shared Task: Hebrew to English (HeEn), Thai to English (ThEn), and Persian to English (PeEn). These languages were chosen because they represent back-transliteration into English. Since the original forms were originally English, they are much more likely to appear in an English corpus than if the words originated in the source language. We report the results on the task’s 1000-instance development sets.

Since transliteration is mostly used for named entities, our language model and unigram counts are obtained from a corpus of named entities. We query DBPedia⁵ for a list of proper names, discarding names that contain non-English characters. The resulting list of 1M names is used to train the character language model and inform the word unigram features.

The results in Table 1 show that our proposed extensions have a dramatic impact on low-resource transliteration. In particular, the seamless incorporation of the target language model not only simplifies the model but also greatly improves the results with respect to the reranking approach. On the other hand, the RNN struggles to learn an adequate model with only 100 training examples.

³<http://www.speech.cs.cmu.edu/SLM>

⁴unimorph.org

⁵<https://wiki.dbpedia.org>

System	HeEn	ThEn	PeEn
DTL	13.2	1.1	8.7
DTL+RR	19.0	2.7	13.6
DTLM	36.7	9.6	26.1
RNN	5.4	1.3	2.6
SEQ	7.8	4.4	8.5

Table 1: Word-level accuracy on transliteration (in %) with 100 training instances.

System	HeEn	ThEn	PeEn
DTL	21.9	37.0	23.6
DTLM	38.7	48.0	36.8
RNN	25.8	43.8	26.7
SEQ	25.5	44.9	31.2

Table 2: Word-level accuracy on transliteration (in %) with complete training sets.

We also evaluate a larger-data scenario. Using the same three languages, we replace the 100 instance training sets with the official training sets from the 2018 shared task, which contain 9,447, 27,273, and 15,677 examples for HeEn, TnEn, and PeEn, respectively. The language model and frequency lists are the same as for the low-resource experiments. Table 2 shows that DTLM outperforms the other systems by a large margin thanks to its ability to leverage a target word list. Additional results are reported by Najafi et al. (2018b).

4.3 Inflection generation

Inflection generation is the task of producing an inflected word-form, given a citation form and a set of morphological features. For example, given the Spanish infinitive *liberar*, with the tag $V; IND; FUT; 2; SG$, the word-form *liberarás* should be produced.

In recent years, inflection generation has attracted much interest (Dreyer and Eisner, 2011; Durrett and DeNero, 2013; Nicolai et al., 2015; Ahlberg et al., 2015). Aharoni and Goldberg (2017) propose an RNN augmented with hard attention and explicit alignments for inflection, but have difficulty consistently improving upon the results of DTL, even on larger datasets. Furthermore, their system cannot be applied to tasks where the source and target are different languages, due to shared embeddings between the encoder and decoder. Ruzsics and Samardzic (2017) incorporate a language model into the decoder of

System	Average
DTL	40.7
DTLM	49.0
CLUZH	40.9

Table 3: Word-level accuracy (in %) on inflection generation with 100 training instances.

a canonical segmentation system. Our model differs in that we learn the influence of the language model during training, in conjunction with DTL’s other features. Deutsch et al. (2018) place a hard constraint on the decoder, so that it only produces observed derivational forms. We instead implement a soft constraint, encouraging candidates that look like real words, but allowing the model to generalize to unseen forms.

Our inflection data comes from the 2017 CoNLL–SIGMORPHON Shared Task on Reinflection (Cotterell et al., 2017). We use the datasets from the low-resource setting of the inflection generation sub-task, in which the training sets are composed of 100 source lemmas with inflection tags and the corresponding inflected forms. We supplement the training data with 100 synthetic “copy” instances that simply transform the target string into itself. This modification, which is known to help in transduction tasks where the source and target are nearly identical, is used for the inflection generation experiments only. Along with the training sets from the shared task, we also use the task’s development and test sets, each consisting of 1000 instances.

Since Sequitur is ill-suited for this type of transduction, we instead train a model using the method of the CLUZH team (Makarov et al., 2017), a state-of-the-art neural system that was the winner of the 2017 shared task. Their primary submission in the shared task was an ensemble of 20 individual systems. For our experiments, we selected their best individual system, as reported in their system paper. For each language, we train models with 3 separate seeds, and select the model that achieves the highest accuracy on the development set.

Table 3 shows that DTLM improves upon CLUZH by a significant margin. The Appendix contains the detailed results for individual languages. DTLM outperforms CLUZH on 46 of the 52 languages. Even for languages with large morphological inventories, such as Basque and Polish,

with the sparse corpora that such inventories supply, we see notable gains over DTL. We also see large gains for languages such as Northern Sami and Navajo that have relatively small Wikipedias (fewer than 10,000 articles).

DTLM was also evaluated as a non-standard submission in the low-data track of the 2018 Shared Task on Universal Morphological Inflection (Cotterell et al., 2018). The results reported by Najafi et al. (2018a) confirm that DTLM substantially outperforms DTL on average. Furthermore, a linear combination of DTLM and a neural system achieved the highest accuracy among all submissions on 34 out of 103 tested languages.

4.4 Cognate projection

Cognate projection, also referred to as cognate production, is the task of predicting the spelling of a hypothetical cognate in another language. For example, given the English word *difficulty*, the Spanish word *dificultad* should be produced. Previously proposed cognate projection systems have been based on SVM taggers (Mulloni, 2007), character-level SMT models (Beinborn et al., 2013), and sequence labeling combined with a maximum-entropy reranker (Ciobanu, 2016).

In this section, we evaluate DTLM in both low- and high-resource settings. Our low-resource data consists of three diverse language pairs. The first set corresponds to a mother-daughter historical relationship between reconstructed Vulgar Latin and Italian (VL-IT) and contains 601 word pairs manually compiled from the textbook of Boyd-Bowman (1980). English and German (EN-DE), close siblings from the Germanic family, are represented by 1013 pairs extracted from Wiktionary. From the same source, we also obtain 438 Slavic word pairs from Russian and Polish (RU-PL), which are written in different scripts (Cyrillic vs. Latin). We make the new datasets publicly available.⁶

The results are shown in Table 4. Of the systems that have no recourse to corpus statistics, the RNN appears crippled by the small training size, while SEQ is competitive with DTL, especially on the difficult EN-DE dataset. On the other hand, the other two systems obtain substantial improvements over DTL, but the gains obtained by DTLM are 2-3 times greater than those obtained by DTL+RR. This demonstrates the advantage of incorporating the language model features directly

System	EN-DE	RU-PL	VL-IT
DTL	4.3	23.5	39.2
DTL+RR	7.1	32.8	43.6
DTLM	17.7	43.9	52.5
RNN	2.2	1.7	15.7
SEQ	9.2	22.3	36.9

Table 4: Word-level accuracy (in %) on cognate projection with 100 training instances.

System	EN-ES	EN-DE	EN-RU
BZG-13	45.7	17.2	8.3
DTL	30.3	24.3	13.3
DTLM	56.8	33.5	45.9
RNN	34.3	20.5	15.0

Table 5: Word-level accuracy (in %) on large-scale cognate projection.

into the training phase over simple reranking.

Our high-resource data comes from a previous study of Beinborn et al. (2013). The datasets were created by applying romanization scripts and string similarity filters to translation pairs extracted from Bing. We find that the datasets are noisy, consisting mostly of lexical loans from Latin, Greek, and English, and include many compound words that share only a single morpheme (i.e., *informatics* and *informationswissenschaft*). In order to alleviate the noise, we found it beneficial to disregard all training pairs that could not be aligned by M2M under the default 2-2 link setting.

Another problem in the data is the overlap between the training and test sets, which ranges from 40% in EN-ES to 94% in EN-EL. Since we believe it would be inappropriate to report results on contaminated sets, we decided to ignore all test instances that occur in the training data. (Unfortunately, this makes some of the test sets too small for a meaningful evaluation.) The resulting dataset sizes are included in the Appendix. Along with the datasets, Beinborn et al. (2013) provide the predictions made by their system. We re-calculate the accuracy of their predictions (BZG-13), discarding any forms that were present in the training set, and compare against DTL and DTLM, as well as the RNN.

Table 5 shows striking gains. While DTL and the RNN generally improve upon BZG-13, DTLM vastly outstrips either alternative. On EN-RU,

⁶<http://github.com/GarrettNicolai/CognateData>

DTLM correctly produces nearly half of potential cognates, 3 times more than any of the other systems. We conclude that our results constitute a new state of the art for cognate projection.

4.5 Phoneme-to-grapheme conversion

Phoneme-to-grapheme (P2G) conversion is the task of predicting the spelling of a word from a sequence of phonemes that represent its pronunciation (Rentzepopoulos and Kokkinakis, 1996). For example, a P2G system should convert [t r ae n z d ah k sh ah n] into *transduction*. Unlike the opposite task of grapheme-to-phoneme (G2P) conversion, large target corpora are widely available. To the best of our knowledge, the state of the art on P2G is the joint n -gram approach of Bisani and Ney (2008), who report improvements on the results of Galescu and Allen (2002) on the NetTalk and CMUDict datasets.

Our low-resource dataset consists of three languages: English (EN), Dutch (NL), and German (DE), extracted from the CELEX lexical database (Baayen et al., 1995).

Table 6 shows that our modifications yield substantial gains for all three languages, with consistent error reductions of 15-20% over the reranking approach. Despite only training on 100 words, the system is able to convert phonetic transcriptions into completely correct spellings for a large fraction of words, even in English, which is notorious for its idiosyncratic orthography. Once again, the RNN is hampered by the small training size.

We also evaluate DTLM in a large-data scenario. We attempt to replicate the P2G experiments reported by (Bisani and Ney, 2008). The data comes from three lexicons on which we conduct 10-fold cross validation: English NetTalk (Sejnowski and Rosenberg, 1993), French Brulex (Mousty and Radeau, 1990), and English CMUDict (Weide, 2005). These corpora contain 20,008, 24,726, and 113,438 words, respectively, in both orthographic and phonetic notations. We note that CMUDict differs from the other two lexicons in that it is much larger, and contains predominantly names, as well as alternative pronunciations. When the training data is that abundant, there is less to be gained from improving the alignment or the target language models, as they are already adequate in the baseline approach.

Table 7 shows the comparison of the results. The P2G results obtained by Sequitur in our exper-

System	EN	NL	DE
DTL	13.9	30.6	33.5
DTL+RR	25.3	32.6	51.5
DTLM	39.6	43.7	60.5
RNN	2.7	6.6	14.1
SEQ	15.9	30.5	28.6

Table 6: Word-level accuracy (in %) on phoneme-to-grapheme conversion with 100 training instances.

	NetTalk	Brulex	CMU
DTL	61.0	68.0	48.3
DTLM	75.2	76.8	49.0
RNN	55.8	67.1	48.0
SEQ	62.7	71.5	48.6

Table 7: Word-level accuracy (in %) on phoneme-to-grapheme conversion with large training sets.

iments are slightly lower than those reported in the original paper, which is attributable to differences in data splits, tuned hyper-parameters, and/or the presence of stress markers in the data. Sequitur still outperforms the baseline DTL, but DTLM substantially outperforms both Sequitur and the RNN on both the NetTalk and Brulex datasets, with smaller gains on the much larger CMUDict. We conclude that our results advance the state of the art on phoneme-to-grapheme conversion.

4.6 Ablation

Table 8 shows the results of disabling individual components in the low-resource setting of the P2G task. The top row reproduces the full DTLM system results reported in Table 6. The remaining three show the results without the character-level LM, word unigram, and precision alignment, respectively. We observe that the accuracy substantially decreases in almost every case, which demonstrates the contribution of all three components.

In a separate experiment on the English G2P dataset, we investigate the impact of the alignment quality by applying several different alignment approaches to the training sets. When M2M aligner uses unconstrained alignment, it favors long alignments that are too sparse to learn a transduction model, achieving less than 1% accuracy. Kubo et al. (2011)’s MPALIGNER, which employs a length penalty to discourage such overlong substring matches, improves moderately, achiev-

System	EN	NL	DE
DTLM	39.6	43.7	60.5
-Language model	37.8	38.2	48.5
-Freq	22.0	37.1	56.7
-Precision	34.9	43.7	46.7

Table 8: Ablation test on P2G data with 100 training instances.

ing 27.5% accuracy, while constraining M2M to 2-2 improves further, to 34.9%. The accuracy increases to 39.6% when the precision alignment is employed. We conclude that in the low-resource setting, our proposed precision alignment is preferable to both MPALIGNER and the standard M2M alignment.

4.7 Error Analysis

The following error examples from three different tasks demonstrate the advantages of incorporating the character-level LM, word frequency, and precision alignment, respectively. For the purpose of insightful analysis, we selected test instances for which DTLM produces markedly better output than DTL.

In inflection generation, the second person plural form of *knechten* is correctly predicted as *knechtetet*, instead of *knechttet*. In this case, our character language model derived from a large text corpus rightly assigns very low probability to the unlikely 4-gram sequence `chtet`, which violates German phonotactic constraints.

In the phoneme-to-grapheme conversion task, [tɪləmətri] is transduced to *telemetry*, instead of *tilemetry*. In English, a reduced vowel phoneme such as [ɪ] may correspond to any vowel letter. In this example, DTLM is able to successfully leverage the occurrence of the correct word-form in a raw corpus.

In cognate projection, the actual cognate of *Kenyan* is *kenianisch*, rather than *kenyisch*. This prediction can be traced to the alignment of the adjectival suffix *-an* to *-anisch* in the training data. The match, which involves a target substring of considerable length, is achieved through a merger of multiple insertion operations.

The errors made by DTLM fall into a few different categories. Occasionally, DTLM produces an incorrect form that is more frequent in the corpus. For example, DTLM incorrectly guesses a subjunctive form of the verb *versetzen* to be the high-

frequency *versetzt*, rather than the unseen *versetztet*. More important, the transducer is incapable of generalizing beyond source-target pairs seen in training. For example, consider the doubling of consonants in English orthography (e.g. *betting*). Unlike the RNN, DTLM incorrectly predicts the present participle of *rug* as **ruging*, because there is no instance of the doubling of ‘g’ in the training data.

5 Conclusion

We have presented DTLM: a powerful language- and task-independent transduction system that can leverage raw target corpora. DTLM is particularly effective in low-resource settings, but is also successful when larger training sets are available. The results of our experiments on four varied transduction tasks show large gains over alternative approaches.

Acknowledgments

The first author was supported in part by the Defense Advanced Research Projects Agency’s (DARPA) Low Resource Emergent Incidents (LORELEI) program, under contract No. HR0011-15-C0113. Any opinions and conclusions expressed in this material are those of the authors, and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA).

The second and third authors were supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

We thank the members of the University of Alberta teams who collaborated with us in the context of the 2018 shared tasks on transliteration and morphological reinflection: Bradley Hauer, Rashed Rubby Riyadh, and Leyuan Yu.

References

- Roe Aharoni and Yoav Goldberg. 2017. [Morphological inflection generation with hard monotonic attention](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2004–2015, Vancouver, Canada. Association for Computational Linguistics.
- Malin Ahlberg, Markus Forsberg, and Mans Hulden. 2015. [Paradigm classification in supervised learning of morphology](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Lan-*

- guage Technologies, pages 1024–1029. Association for Computational Linguistics.
- Harald R. Baayen, Richard Piepenbrock, and Leon Gulikers. 1995. *The CELEX Lexical Database. Release 2 (CD-ROM)*. Linguistic Data Consortium, University of Pennsylvania, Philadelphia, Pennsylvania.
- G. Edward Barton. 1986. [Computational complexity in two-level morphology](#). In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pages 53–59, New York, New York, USA. Association for Computational Linguistics.
- Lisa Beinborn, Torsten Zesch, and Iryna Gurevych. 2013. [Cognate production using character-based machine translation](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 883–891, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Maximilian Bisani and Hermann Ney. 2008. Joint-sequence models for grapheme-to-phoneme conversion. *Speech communication*, 50(5):434–451.
- Peter Boyd-Bowman. 1980. *From Latin to Romance in sound charts*. Georgetown University Press.
- Colin Cherry and Hisami Suzuki. 2009. [Discriminative substring decoding for transliteration](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1066–1075. Association for Computational Linguistics.
- Alina Maria Ciobanu. 2016. [Sequence labeling for cognate production](#). In *Knowledge-Based and Intelligent Information and Engineering Systems: Proceedings of the 20th International Conference KES-2016*, pages 1391–1399.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D. McCarthy, Katharina Kann, Sebastian Mielke, Garrett Nicolai, Miikka Silfverberg, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, Brussels, Belgium. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2017. [CoNLL–SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages](#). In *Proceedings of the CoNLL–SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 1–30, Vancouver. Association for Computational Linguistics.
- Daniel Deutsch, John Hewitt, and Dan Roth. 2018. [A distributional and orthographic aggregation model for english derivational morphology](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1938–1947. Association for Computational Linguistics.
- Markus Dreyer and Jason Eisner. 2011. Discovering morphological paradigms from plain text using a Dirichlet process mixture model. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 616–627. Association for Computational Linguistics.
- Greg Durrett and John DeNero. 2013. Supervised learning of complete morphological paradigms. In *HLT-NAACL*, pages 1185–1195.
- Lucian Galescu and James F. Allen. 2002. Pronunciation of proper names with a joint n-gram model for bi-directional grapheme-to-phoneme conversion. In *Seventh International Conference on Spoken Language Processing*.
- Sittichai Jiampojarn, Colin Cherry, and Grzegorz Kondrak. 2010. Integrating joint n-gram features into a discriminative training network. In *NAACL-HLT*, pages 697–700, Los Angeles, CA. Association for Computational Linguistics.
- Sittichai Jiampojarn, Grzegorz Kondrak, and Tarek Sherif. 2007. [Applying many-to-many alignments and hidden markov models to letter-to-phoneme conversion](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 372–379. Association for Computational Linguistics.
- Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 133–142. ACM.
- Keigo Kubo, Hiromichi Kawanami, Hiroshi Saruwatari, and Kiyohiro Shikano. 2011. Unconstrained many-to-many alignment for automatic pronunciation annotation. In *Proceedings of the Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*.
- Peter Makarov, Tatiana Ruzsics, and Simon Clematide. 2017. Align and copy: UZH at SIGMORPHON 2017 shared task for morphological reinflection. *arXiv preprint arXiv:1707.01355*.
- Ryan McDonald, Koby Crammer, and Fernando Pereira. 2005. Online large-margin training of dependency parsers. In *ACL*.
- Philippe Mousty and Monique Radeau. 1990. Brulex. Une base de données lexicales informatisée pour le français écrit et parlé. *L'année Psychologique*, 90(4):551–566.

- Andrea Mulloni. 2007. [Automatic prediction of cognate orthography using support vector machines](#). In *Proceedings of the ACL 2007 Student Research Workshop*, pages 25–30, Prague, Czech Republic. Association for Computational Linguistics.
- Saeed Najafi, Bradley Hauer, Rashed Rubby Riyadh, Leyuan Yu, and Grzegorz Kondrak. 2018a. [Combining neural and non-neural methods for low-resource morphological reinflection](#). In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, Brussels, Belgium. Association for Computational Linguistics.
- Saeed Najafi, Bradley Hauer, Rashed Rubby Riyadh, Leyuan Yu, and Grzegorz Kondrak. 2018b. [Comparison of assorted models for transliteration](#). In *Proceedings of the Seventh Named Entities Workshop*, pages 84–88. Association for Computational Linguistics.
- Garrett Nicolai, Colin Cherry, and Grzegorz Kondrak. 2015. [Inflection generation as discriminative string transduction](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 922–931. Association for Computational Linguistics.
- Panagiotis A. Rentzepopoulos and George K. Kokkinakis. 1996. [Efficient multilingual phoneme-to-grapheme conversion based on HMM](#). *Computational Linguistics*, 22(3):351–375.
- Tatyana Ruzsics and Tanja Samardzic. 2017. [Neural sequence-to-sequence learning of internal word structure](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 184–194, Vancouver, Canada. Association for Computational Linguistics.
- TJ Sejnowski and CR Rosenberg. 1993. NETtalk corpus. [URL < ftp://svrftp.eng.cam.ac.uk/pub/comp.speech/dictionaries](ftp://svrftp.eng.cam.ac.uk/pub/comp.speech/dictionaries).
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in neural information processing systems*, pages 3104–3112.
- Robert Weide. 2005. The Carnegie Mellon pronouncing dictionary [cmudict.0.6].