Dependency Treelet Translation: Syntactically Informed Phrasal SMT

Chris Quirk, Arul Menezes

Microsoft Research
One Microsoft Way
Redmond, WA 98052
{chrisq, arulm}@microsoft.com

Colin Cherry

University of Alberta
Edmonton, Alberta
Canada T6G 2E1
colinc@cs.ualberta.ca

Abstract

We describe a novel approach to statistical machine translation that combines syntactic information in the source language with recent advances in phrasal translation. This method requires a source-language dependency parser, target language word segmentation and an unsupervised word alignment component. We align a parallel corpus, project the source dependency parse onto the target sentence, extract dependency treelet translation pairs, and train a tree-based ordering model. We describe an efficient decoder and show that using these tree-based models in combination with conventional SMT models provides a promising approach that incorporates the power of phrasal SMT with the linguistic generality available in a parser.

1. Introduction

Over the past decade, we have witnessed a revolution in the field of machine translation (MT) toward statistical or corpus-based methods. Yet despite this success, statistical machine translation (SMT) has many hurdles to overcome. While it excels at translating domain-specific terminology and fixed phrases, grammatical generalizations are poorly captured and often mangled during translation (Thurmair, 04).

1.1. Limitations of string-based phrasal SMT

State-of-the-art phrasal SMT systems such as (Koehn et al., 03) and (Vogel et al., 03) model translations of *phrases* (here, strings of adjacent words, not syntactic constituents) rather than individual words. Arbitrary reordering of words is allowed within memorized phrases, but typically

only a small amount of phrase reordering is allowed, modeled in terms of offset positions at the string level. This reordering model is very limited in terms of linguistic generalizations. For instance, when translating English to Japanese, an ideal system would automatically learn large-scale typological differences: English SVO clauses generally become Japanese SOV clauses, English post-modifying prepositional phrases become Japanese pre-modifying postpositional phrases, etc. A phrasal SMT system may learn the internal reordering of specific common phrases, but it cannot generalize to unseen phrases that share the same linguistic structure.

In addition, these systems are limited to phrases contiguous in both source and target, and thus cannot learn the generalization that English *not* may translate as French *ne...pas* except in the context of specific intervening words.

1.2. Previous work on syntactic SMT¹

The hope in the SMT community has been that the incorporation of syntax would address these issues, but that promise has yet to be realized.

One simple means of incorporating syntax into SMT is by re-ranking the n -best list of a baseline SMT system using various syntactic models, but Och et al. (04) found very little positive impact with this approach. However, an n -best list of even 16,000 translations captures only a tiny fraction of the ordering possibilities of a 20 word sentence; re-ranking provides the syntactic model no opportunity to boost or prune large sections of that search space.

Inversion Transduction Grammars (Wu, 97), or ITGs, treat translation as a process of parallel parsing of the source and target language via a synchronized grammar. To make this process

¹ Note that since this paper does not address the word alignment problem directly, we do not discuss the large body of work on incorporating syntactic information into the word alignment process.

computationally efficient, however, some severe simplifying assumptions are made, such as using a single non-terminal label. This results in the model simply learning a very high level preference regarding how often nodes should switch order without any contextual information. Also these translation models are intrinsically word-based; phrasal combinations are not modeled directly, and results have not been competitive with the top phrasal SMT systems.

Along similar lines, Alshawi et al. (2000) treat translation as a process of simultaneous induction of source and target dependency trees using head-transduction; again, no separate parser is used.

Yamada and Knight (01) employ a parser in the target language to train probabilities on a set of operations that convert a target language tree to a source language string. This improves fluency slightly (Charniak et al., 03), but fails to significantly impact overall translation quality. This may be because the parser is applied to MT output, which is notoriously unlike native language, and no additional insight is gained via source language analysis.

Lin (04) translates dependency trees using paths. This is the first attempt to incorporate large phrasal SMT-style memorized patterns together with a separate source dependency parser and SMT models. However the phrases are limited to linear paths in the tree, the only SMT model used is a maximum likelihood channel model and there is no ordering model. Reported BLEU scores are far below the leading phrasal SMT systems.

MSR-MT (Menezes & Richardson, 01) parses both source and target languages to obtain a logical form (LF), and translates source LFs using memorized aligned LF patterns to produce a target LF. It utilizes a separate sentence realization component (Ringger et al., 04) to turn this into a target sentence. As such, it does not use a target language model during decoding, relying instead on MLE channel probabilities and heuristics such as pattern size. Recently Aue et al. (04) incorporated an LF-based language model (LM) into the system for a small quality boost. A key disadvantage of this approach and related work (Ding & Palmer, 02) is that it requires a parser in *both* languages, which severely limits the language pairs that can be addressed.

2. Dependency Treelet Translation

In this paper we propose a novel dependency tree-based approach to phrasal SMT which uses tree-based ‘phrases’ and a tree-based ordering model in combination with conventional SMT models to produce state-of-the-art translations.

Our system employs a source-language dependency parser, a target language word segmentation component, and an unsupervised word alignment component to learn treelet translations from a parallel sentence-aligned corpus. We begin by parsing the source text to obtain dependency trees and word-segmenting the target side, then applying an off-the-shelf word alignment component to the bitext.

The word alignments are used to project the source dependency parses onto the target sentences. From this aligned parallel dependency corpus we extract a treelet translation model incorporating source and target treelet pairs, where a *treelet* is defined to be an arbitrary connected subgraph of the dependency tree. A unique feature is that we allow treelets with a wildcard root, effectively allowing mappings for siblings in the dependency tree. This allows us to model important phenomena, such as *not ... → ne...pas*. We also train a variety of statistical models on this aligned dependency tree corpus, including a channel model and an order model.

To translate an input sentence, we parse the sentence, producing a dependency tree for that sentence. We then employ a decoder to find a combination and ordering of treelet translation pairs that cover the source tree and are optimal according to a set of models that are combined in a log-linear framework as in (Och, 03).

This approach offers the following advantages over string-based SMT systems: Instead of limiting learned phrases to contiguous word sequences, we allow translation by all possible phrases that form connected subgraphs (treelets) in the source and target dependency trees. This is a powerful extension: the vast majority of surface-contiguous phrases are also treelets of the tree; in addition, we gain discontinuous phrases, including combinations such as verb-object, article-noun, adjective-noun etc. regardless of the number of intervening words.

Figure 1. An example dependency tree.

Another major advantage is the ability to employ more powerful models for reordering source language constituents. These models can incorporate information from the source analysis. For example, we may model directly the probability that the translation of an object of a preposition in English should precede the corresponding postposition in Japanese, or the probability that a pre-modifying adjective in English translates into a post-modifier in French.

2.1. Parsing and alignment

We require a source language dependency parser that produces unlabeled, ordered dependency trees and annotates each source word with a part-of-speech (POS). An example dependency tree is shown in Figure 1. The arrows indicate the head annotation, and the POS for each candidate is listed underneath. For the target language we only require word segmentation.

To obtain word alignments we currently use GIZA++ (Och & Ney, 03). We follow the common practice of deriving many-to-many alignments by running the IBM models in both directions and combining the results heuristically. Our heuristics differ in that they constrain many-to-one alignments to be contiguous in the source dependency tree. A detailed description of these heuristics can be found in Quirk et al. (04).

2.2. Projecting dependency trees

Given a word aligned sentence pair and a source dependency tree, we use the alignment to project the source structure onto the target sentence. One-to-one alignments project directly to create a target tree isomorphic to the source. Many-to-one alignments project similarly; since the ‘many’ source nodes are connected in the tree, they act as if condensed into a single node. In the case of one-to-many alignments we project the source node to the rightmost² of the ‘many’ target words, and make the rest of the target words dependent on it.

² If the target language is Japanese, leftmost may be more appropriate.

(a) Word alignment.

(b) Dependencies after initial projection.

(c) Dependencies after reattachment step.

Figure 2. Projection of dependencies.

Unaligned target words³ are attached into the dependency structure as follows: assume there is an unaligned word t_j in position j . Let $i < j$ and $k > j$ be the target positions closest to j such that t_i depends on t_k or vice versa: attach t_j to the lower of t_i or t_k . If all the nodes to the left (or right) of position j are unaligned, attach t_j to the left-most (or right-most) word that is aligned.

The target dependency tree created in this process may not read off in the same order as the target string, since our alignments do not enforce phrasal cohesion. For instance, consider the projection of the parse in Figure 1 using the word alignment in Figure 2a. Our algorithm produces the dependency tree in Figure 2b. If we read off the leaves in a left-to-right in-order traversal, we do not get the original input string: *de démarrage* appears in the wrong place.

A second reattachment pass corrects this situation. For each node in the wrong order, we reattach it to the lowest of its ancestors such that it is in the correct place relative to its siblings and parent. In Figure 2c, reattaching *démarrage* to *et* suffices to produce the correct order.

³ Source unaligned nodes do not present a problem, with the exception that if the root is unaligned, the projection process produces a forest of target trees anchored by a dummy root.

2.3. Extracting treelet translation pairs

From the aligned pairs of dependency trees we extract all pairs of aligned source and target treelets along with word-level alignment linkages, up to a configurable maximum size. We also keep treelet counts for maximum likelihood estimation.

2.4. Order model

Phrasal SMT systems often use a model to score the ordering of a set of phrases. One approach is to penalize any deviation from monotone decoding; another is to estimate the probability that a source phrase in position i translates to a target phrase in position j (Koehn et al., 03).

We attempt to improve on these approaches by incorporating syntactic information. Our model assigns a probability to the order of a target tree given a source tree. Under the assumption that constituents generally move as a whole, we predict the probability of each given ordering of modifiers independently. That is, we make the following simplifying assumption (where c is a function returning the set of nodes modifying t):

$$P(\text{order}(T) | S, T) = \prod_{t \in T} P(\text{order}(c(t)) | S, T)$$

Furthermore, we assume that the position of each child can be modeled independently in terms of a head-relative position:

$$P(\text{order}(c(t)) | S, T) = \prod_{m \in c(t)} P(\text{pos}(m, t) | S, T)$$

Figure 3a demonstrates an aligned dependency tree pair annotated with head-relative positions; Figure 3b presents the same information in an alternate tree-like representation.

We currently use a small set of features reflecting very local information in the dependency tree to model $P(\text{pos}(m, t) | S, T)$:

- The lexical items of the head and modifier.
- The lexical items of the source nodes aligned to the head and modifier.
- The part-of-speech ("cat") of the source nodes aligned to the head and modifier.
- The head-relative position of the source node aligned to the source modifier.⁴

As an example, consider the children of *propriété* in Figure 3. The head-relative positions

(a) Head annotation representation

(b) Branching structure representation.

Figure 3. Aligned dependency tree pair, annotated with head-relative positions

of its modifiers *la* and *Cancel* are -1 and +1, respectively. Thus we try to predict as follows:

$$P(\text{pos}(m_1) = -1 | \text{lex}(m_1) = "la", \text{lex}(h) = "propriété", \text{lex}(\text{src}(m_1)) = "the", \text{lex}(\text{src}(h)) = "property", \text{cat}(\text{src}(m_1)) = \text{Determiner}, \text{cat}(\text{src}(h)) = \text{Noun}, \text{position}(\text{src}(m_1)) = -2) \cdot$$

$$P(\text{pos}(m_2) = +1 | \text{lex}(m_2) = "Cancel", \text{lex}(h) = "propriété", \text{lex}(\text{src}(m_2)) = "Cancel", \text{lex}(\text{src}(h)) = "property", \text{cat}(\text{src}(m_2)) = \text{Noun}, \text{cat}(\text{src}(h)) = \text{Noun}, \text{position}(\text{src}(m_2)) = -1)$$

The training corpus acts as a supervised training set: we extract a training feature vector from each of the target language nodes in the aligned dependency tree pairs. Together these feature vectors are used to train a decision tree (Chickering, 02). The distribution at each leaf of the DT can be used to assign a probability to each possible target language position. A more detailed description is available in (Quirk et al., 04).

2.5. Other models

Channel Models: We incorporate two distinct channel models, a maximum likelihood estimate (MLE) model and a model computed using Model-1 word-to-word alignment probabilities as in (Vogel et al., 03). The MLE model effectively captures non-literal phrasal translations such as idioms, but suffers from data sparsity. The word-

⁴ One can also include features of siblings to produce a Markov ordering model. However, we found that this had little impact in practice.

to-word model does not typically suffer from data sparsity, but prefers more literal translations.

Given a set of treelet translation pairs that cover a given input dependency tree and produce a target dependency tree, we model the probability of source given target as the product of the individual treelet translation probabilities: we assume a uniform probability distribution over the decompositions of a tree into treelets.

Target Model: Given an ordered target language dependency tree, it is trivial to read off the surface string. We evaluate this string using a trigram model with modified Kneser-Ney smoothing.

Miscellaneous Feature Functions: The log-linear framework allows us to incorporate other feature functions as ‘models’ in the translation process. For instance, using fewer, larger treelet translation pairs often provides better translations, since they capture more context and allow fewer possibilities for search and model error. Therefore we add a feature function that counts the number of phrases used. We also add a feature that counts the number of target words; this acts as an insertion/deletion bonus/penalty.

3. Decoding

The challenge of tree-based decoding is that the traditional left-to-right decoding approach of string-based systems is inapplicable. Additional challenges are posed by the need to handle treelets—perhaps discontinuous or overlapping—and a combinatorially explosive ordering space.

Our decoding approach is influenced by ITG (Wu, 97) with several important extensions. First, we employ treelet translation pairs instead of single word translations. Second, instead of modeling rearrangements as either preserving source order or swapping source order, we allow the dependents of a node to be ordered in any arbitrary manner and use the order model described in section 2.4 to estimate probabilities. Finally, we use a log-linear framework for model combination that allows any amount of other information to be modeled.

We will initially approach the decoding problem as a bottom up, exhaustive search. We define the set of all possible treelet translation pairs of the subtree rooted at each input node in the following manner: A treelet translation pair x is said to *match* the input dependency tree S iff

(a) Example input dependency tree.

(b) Example treelet translation pair.

Figure 4. Example decoder structures.

there is some connected subgraph S' that is identical to the source side of x . We say that x *covers* all the nodes in S' and is *rooted* at source node s , where s is the root of matched subgraph S' .

We first find all treelet translation pairs that match the input dependency tree. Each matched pair is placed on a list associated with the input node where the match is rooted. Moving bottom-up through the input dependency tree, we compute a list of *candidate translations* for the input subtree rooted at each node s , as follows:

Consider in turn each treelet translation pair x rooted at s . The treelet pair x may cover only a portion of the input subtree rooted at s . Find all descendants s' of s that are *not* covered by x , but whose parent s'' is covered by x . At each such node s'' look at all interleavings of the children of s'' specified by x , if any, with each translation t' from the candidate translation list⁵ of each child s' . Each such interleaving is scored using the models previously described and added to the candidate translation list for that input node. The resultant translation is the best scoring candidate for the root input node.

As an example, see the example dependency tree in Figure 4a and treelet translation pair in 4b. This treelet translation pair covers all the nodes in 4a except the subtrees rooted at *software* and *is*.

⁵ Computed by the previous application of this procedure to s' during the bottom-up traversal.

We first compute (and cache) the candidate translation lists for the subtrees rooted at *software* and *is*, then construct full translation candidates by attaching those subtree translations to *installés* in all possible ways. The order of *sur* relative to *installés* is fixed; it remains to place the translated subtrees for *the software* and *is*. Note that if c is the count of children specified in the mapping and r is the count of subtrees translated via recursive calls, then there are $(c+r+1)!/(c+1)!$ orderings. Thus $(1+2+1)!/(1+1)! = 12$ candidate translations are produced for each combination of translations of *the software* and *is*.

3.1. Optimality-preserving optimizations

Dynamic Programming

Converting this exhaustive search to dynamic programming relies on the observation that scoring a translation candidate at a node depends on the following information from its descendants: the order model requires features from the root of a translated subtree, and the target language model is affected by the first and last two words in each subtree. Therefore, we need to keep the best scoring translation candidate for a given subtree for each combination of (head, leading bigram, trailing bigram), which is, in the worst case, $O(V^5)$, where V is the vocabulary size. The dynamic programming approach therefore does not allow for great savings in practice because a trigram target language model forces consideration of context external to each subtree.

Duplicate elimination

To eliminate unnecessary ordering operations, we first check that a given set of words has not been previously ordered by the decoder. We use an order-independent hash table where two trees are considered equal if they have the same tree structure and lexical choices after sorting each child list into a canonical order. A simpler alternate approach would be to compare bags-of-words. However since our possible orderings are bound by the induced tree structure, we might overzealously prune a candidate with a different tree structure that allows a better target order.

3.2. Lossy optimizations

The following optimizations do not preserve optimality, but work well in practice.

N-best lists

Instead of keeping the full list of translation candidates for a given input node, we keep a top-scoring subset of the candidates. While the decoder is no longer guaranteed to find the optimal translation, in practice the quality impact is minimal with a list size ≥ 10 (see Table 5.6).

Variable-sized n-best lists: A further speedup can be obtained by noting that the number of translations using a given treelet pair is exponential in the number of subtrees of the input not covered by that pair. To limit this explosion we vary the size of the n -best list on any recursive call in inverse proportion to the number of subtrees uncovered by the current treelet. This has the intuitive appeal of allowing a more thorough exploration of large treelet translation pairs (that are likely to result in better translations) than of smaller, less promising pairs.

Pruning treelet translation pairs

Channel model scores and treelet size are powerful predictors of translation quality. Heuristically pruning low scoring treelet translation pairs before the search starts allows the decoder to focus on combinations and orderings of high quality treelet pairs.

- Only keep those treelet translation pairs with an MLE probability above a threshold t .
- Given a set of treelet translation pairs with identical sources, keep those with an MLE probability within a ratio r of the best pair.
- At each input node, keep only the top k treelet translation pairs rooted at that node, as ranked first by size, then by MLE channel model score, then by Model 1 score. The impact of this optimization is explored in Table 5.6.

Greedy ordering

The complexity of the ordering step at each node grows with the factorial of the number of children to be ordered. This can be tamed by noting that given a fixed pre- and post-modifier count, our order model is capable of evaluating a single ordering decision independently from other ordering decisions.

One version of the decoder takes advantage of this to severely limit the number of ordering possibilities considered. Instead of considering all interleavings, it considers each potential modifier position in turn, greedily picking the most

		English	French
Training	Sentences	570,562	
	Words	7,327,251	8,415,882
	Vocabulary	72,440	80,758
	Singletons	38,037	39,496
Test	Sentences	10,000	
	Words	133,402	153,701

Table 4.1 Data characteristics

probable child for that slot, moving on to the next slot, picking the most probable among the remaining children for that slot and so on.

The complexity of greedy ordering is linear, but at the cost of a noticeable drop in BLEU score (see Table 5.4). Under default settings our system tries to decode a sentence with exhaustive ordering until a specified timeout, at which point it falls back to greedy ordering.

4. Experiments

We evaluated the translation quality of the system using the BLEU metric (Papineni et al., 02) under a variety of configurations. We compared against two radically different types of systems to demonstrate the competitiveness of this approach:

- Pharaoh: A leading phrasal SMT decoder (Koehn et al., 03).
- The MSR-MT system described in Section 1, an EBMT/hybrid MT system.

4.1. Data

We used a parallel English-French corpus containing 1.5 million sentences of Microsoft technical data (e.g., support articles, product documentation). We selected a cleaner subset of this data by eliminating sentences with XML or HTML tags as well as very long (>160 characters) and very short (<40 characters) sentences. We held out 2,000 sentences for development testing and parameter tuning, 10,000 sentences for testing, and 250 sentences for lambda training. We ran experiments on subsets of the training data ranging from 1,000 to 300,000 sentences. Table 4.1 presents details about this dataset.

4.2. Training

We parsed the source (English) side of the corpus using NLPWIN, a broad-coverage rule-based parser developed at Microsoft Research able to

produce syntactic analyses at varying levels of depth (Heidorn, 02). For the purposes of these experiments we used a dependency tree output with part-of-speech tags and unstemmed surface words.

For word alignment, we used GIZA++, following a standard training regimen of five iterations of Model 1, five iterations of the HMM Model, and five iterations of Model 4, in both directions.

We then projected the dependency trees and used the aligned dependency tree pairs to extract treelet translation pairs and train the order model as described above. The target language model was trained using only the French side of the corpus; additional data may improve its performance. Finally we trained lambdas via Maximum BLEU (Och, 03) on 250 held-out sentences with a single reference translation, and tuned the decoder optimization parameters (n -best list size, timeouts etc) on the development test set.

Pharaoh

The same GIZA++ alignments as above were used in the Pharaoh decoder. We used the heuristic combination described in (Och & Ney, 03) and extracted phrasal translation pairs from this combined alignment as described in (Koehn et al., 03). Except for the order model (Pharaoh uses its own ordering approach), the same models were used: MLE channel model, Model 1 channel model, target language model, phrase count, and word count. Lambdas were trained in the same manner (Och, 03).

MSR-MT

MSR-MT used its own word alignment approach as described in (Menezes & Richardson, 01) on the same training data. MSR-MT does not use lambdas or a target language model.

5. Results

We present BLEU scores on an unseen 10,000 sentence test set using a single reference translation for each sentence. Speed numbers are the end-to-end translation speed in sentences per minute. All results are based on a training set size of 100,000 sentences and a phrase size of 4, except Table 5.2 which varies the phrase size and Table 5.3 which varies the training set size.

Results for our system and the comparison systems are presented in Table 5.1. Pharaoh monotone refers to Pharaoh with phrase reordering disabled. The difference between Pharaoh and the Treelet system is significant at the 99% confidence level under a two-tailed paired t-test.

	BLEU Score	Sents/min
Pharaoh monotone	37.06	4286
Pharaoh	38.83	162
MSR-MT	35.26	453
Treelet	40.66	10.1

Table 5.1 System comparisons

Table 5.2 compares Pharaoh and the Treelet system at different phrase sizes. While all the differences are statistically significant at the 99% confidence level, the wide gap at smaller phrase sizes is particularly striking. We infer that whereas Pharaoh depends heavily on long phrases to encapsulate reordering, our dependency tree-based ordering model enables credible performance even with single-word ‘phrases’. We conjecture that in a language pair with large-scale ordering differences, such as English-Japanese, even long phrases are unlikely to capture the necessary reorderings, whereas our tree-based ordering model may prove more robust.

Max. size	Treelet BLEU	Pharaoh BLEU
1	37.50	23.18
2	39.84	32.07
3	40.36	37.09
4 (default)	40.66	38.83
5	40.71	39.41
6	40.74	39.72

Table 5.2 Effect of maximum treelet/phrase size

Table 5.3 compares the same systems at different training corpus sizes. All of the differences are statistically significant at the 99% confidence level. Noting that the gap widens at smaller corpus sizes, we suggest that our tree-based approach is more suitable than string-based phrasal SMT when translating from English into languages or domains with limited parallel data.

	1k	3k	10k	30k	100k	300k
Pharaoh	17.20	22.51	27.70	33.73	38.83	42.75
Treelet	18.70	25.39	30.96	35.81	40.66	44.32

Table 5.3 Effect of training set size on treelet translation and comparison system

We also ran experiments varying different system parameters. Table 5.4 explores different ordering strategies, Table 5.5 looks at the impact of discontinuous phrases and Table 5.6 looks at the impact of decoder optimizations such as treelet pruning and n -best list size.

Ordering strategy	BLEU	Sents/min
No order model (monotone)	35.35	39.7
Greedy ordering	38.85	13.1
Exhaustive (default)	40.66	10.1

Table 5.4 Effect of ordering strategies

	BLEU Score	Sents/min
Contiguous only	40.08	11.0
Allow discontinuous	40.66	10.1

Table 5.5 Effect of allowing treelets that correspond to discontinuous phrases

	BLEU Score	Sents/min
Pruning treelets		
Keep top 1	28.58	144.9
... top 3	39.10	21.2
... top 5	40.29	14.6
... top 10 (default)	40.66	10.1
... top 20	40.70	3.5
Keep all	40.29	3.2
N-best list size		
1-best	37.28	175.4
5-best	39.96	79.4
10-best	40.42	23.3
20-best (default)	40.66	10.1
50-best	39.39	3.7

Table 5.6 Effect of optimizations

6. Discussion

We presented a novel approach to syntactically-informed statistical machine translation that leverages a parsed dependency tree representation of the source language via a tree-based ordering model and treelet phrase extraction. We showed that it significantly outperforms a leading phrasal SMT system over a wide range of training set sizes and phrase sizes.

Constituents vs. dependencies: Most attempts at

syntactic SMT have relied on a constituency analysis rather than dependency analysis. While this is a natural starting point due to its well-understood nature and commonly available tools, we feel that this is not the most effective representation for syntax in MT. Dependency analysis, in contrast to constituency analysis, tends to bring semantically related elements together (e.g., verbs become adjacent to all their arguments) and is better suited to lexicalized models, such as the ones presented in this paper.

7. Future work

The most important contribution of our system is a linguistically motivated ordering approach based on the source dependency tree, yet this paper only explores one possible model. Different model structures, machine learning techniques, and target feature representations all have the potential for significant improvements.

Currently we only consider the top parse of an input sentence. One means of considering alternate possibilities is to build a packed forest of dependency trees and use this in decoding translations of each input sentence.

As noted above, our approach shows particular promise for language pairs such as English-Japanese that exhibit large-scale reordering and have proven difficult for string-based approaches. Further experimentation with such language pairs is necessary to confirm this. Our experience has been that the quality of GIZA++ alignments for such language pairs is inadequate. Following up on ideas introduced by (Cherry & Lin, 03) we plan to explore ways to leverage the dependency tree to improve alignment quality.

References

Alshaw, Hiyon, Srinivas Bangalore, and Shona Douglas. Learning dependency translation models as collections of finite-state head transducers. *Computational Linguistics*, 26(1):45–60, 2000.

Aue, Anthony, Arul Menezes, Robert C. Moore, Chris Quirk, and Eric Ringger. Statistical machine translation using labeled semantic dependency graphs. *TMI* 2004.

Charniak, Eugene, Kevin Knight, and Kenji Yamada. Syntax-based language models for statistical machine translation. *MT Summit* 2003.

Cherry, Colin and Dekang Lin. A probability model to improve word alignment. *ACL* 2003.

Chickering, David Maxwell. The WinMine Toolkit. Microsoft Research Technical Report: MSR-TR-2002-103.

Ding, Yuan and Martha Palmer. Automatic learning of parallel dependency treelet pairs. *IJCNLP* 2004.

Heidorn, George. (2000). “Intelligent writing assistance”. In Dale et al. *Handbook of Natural Language Processing*, Marcel Dekker.

Koehn, Philipp, Franz Josef Och, and Daniel Marcu. Statistical phrase based translation. *NAACL* 2003.

Lin, Dekang. A path-based transfer model for machine translation. *COLING* 2004.

Menezes, Arul and Stephen D. Richardson. A best-first alignment algorithm for automatic extraction of transfer mappings from bilingual corpora. *DDMT Workshop, ACL* 2001.

Och, Franz Josef and Hermann Ney. A systematic comparison of various statistical alignment models, *Computational Linguistics*, 29(1):19-51, 2003.

Och, Franz Josef. Minimum error rate training in statistical machine translation. *ACL* 2003.

Och, Franz Josef, et al. A smorgasbord of features for statistical machine translation. *HLT/NAACL* 2004.

Papineni, Kishore, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. *ACL* 2002.

Quirk, Chris, Arul Menezes, and Colin Cherry. Dependency Tree Translation. Microsoft Research Technical Report: MSR-TR-2004-113.

Ringger, Eric, et al. Linguistically informed statistical models of constituent structure for ordering in sentence realization. *COLING* 2004.

Thurmair, Gregor. Comparing rule-based and statistical MT output. Workshop on the amazing utility of parallel and comparable corpora, *LREC*, 2004.

Vogel, Stephan, Ying Zhang, Fei Huang, Alicia Tribble, Ashish Venugopal, Bing Zhao, and Alex Waibel. The CMU statistical machine translation system. *MT Summit* 2003.

Wu, Dekai. Stochastic inversion transduction grammars and bilingual parsing of parallel corpora. *Computational Linguistics*, 23(3):377–403, 1997.

Yamada, Kenji and Kevin Knight. A syntax-based statistical translation model. *ACL*, 2001.