

Semantically-Informed Syntactic Machine Translation: A Tree-Grafting Approach

Kathryn Baker U.S. Dept. of Defense Fort Meade, MD 20755 klbake4@tycho.ncsc.mil	Michael Bloodgood JHU HLTCOE Baltimore, MD 21218 bloodgood@jhu.edu	Chris Callison-Burch JHU HLTCOE Baltimore, MD 21218 ccb@cs.jhu.edu	Bonnie J. Dorr University of Maryland College Park, MD 20742 bonnie@umiacs.umd.edu
Nathaniel W. Filardo JHU HLTCOE Baltimore, MD 21211 nwf@cs.jhu.edu	Lori Levin CMU Pittsburgh, PA 15213 lsl@cs.cmu.edu	Scott Miller BBN Cambridge, MA 02138 smiller@bbn.com	Christine Piatko JHU/APL Laurel, MD 20723 christine.piatko@jhuapl.edu

Abstract

We describe a unified and coherent syntactic framework for supporting a semantically-informed syntactic approach to statistical machine translation. Semantically enriched syntactic tags assigned to the target-language training texts improved translation quality. The resulting system significantly outperformed a linguistically naive baseline model (Hiero), and reached the highest scores yet reported on the NIST 2009 Urdu-English translation task. This finding supports the hypothesis (posed by many researchers in the MT community, e.g., in DARPA GALE) that both syntactic and semantic information are critical for improving translation quality—and further demonstrates that large gains can be achieved for low-resource languages with different word order than English.

1 Introduction

This paper describes a tree-grafting approach to incorporating named entities and modality into a unified and coherent syntactic framework, as a first step toward supporting Semantically-Informed Machine Translation (SIMT). The implementation of this approach was the result of a large effort undertaken in the summer of 2009. The most significant result of the SIMT effort was the integration of semantic knowledge into statistical machine translation in a unified and coherent syntactic framework. By augmenting hierarchical phrase-based translation rules with syntactic labels that were extracted from a parsed parallel corpus, and further augmenting the parse trees with semantic elements such as named-entity markers and modality (through a process we refer to as *grafting*), we produced a better

model for translating Urdu and English. The resulting system significantly outperformed the linguistically naive baseline Hiero model, and reached the highest scores yet reported on the NIST 2009 Urdu-English translation task.

We note that, while our largest gains were from syntactic enrichments to the model, smaller (but significant) gains were achieved by injecting semantic knowledge into the syntactic paradigm. Of course, entities and modalities are only a small piece of the much larger semantic space, but demonstrating success on these new, unexplored semantic aspects of language bodes well for (larger) improvements based on the incorporation of other semantic aspects (e.g., relations and temporal knowledge). Moreover, we believe this syntactic framework to be well suited for further exploration of the impact of many different types of semantics on MT quality. Indeed, it would not have been possible to initiate the current study without the foundational work that gave rise to a syntactic paradigm that could support these semantic enrichments. We believe this framework will be especially useful for exploring other languages with few resources and different word order than English.

The semantic units that we examined in this effort were named entities (such as people or organizations) and modalities (indications that a statement represents something that has taken place or is a belief or an intention). Other semantic units such as relations between entities and events, were not part of this effort, but we believe they could be similarly incorporated into the framework. We chose to examine semantic units that canonically exhibit two different syntactic types: nominal, in the case of named entities, and verbal, in the case of modality.

Source	Reference	pre-SIMT MT output
<p>ناگاؤں نے آسام میں آگ لگا دی</p> <p>بدھ کے روز مشتعل ناگا قبائلیوں نے منی پور کے دس سکولوں کو بھی نذر آتش کر دیا تھا۔</p> <p>پولیس کے مطابق سینکڑوں کی تعداد میں ناگالینڈ کے مسلح قبائلیوں نے آسام کے گلیکی اور سیسیسا گر کے تین گاؤں میں آگ لگا دی۔</p> <p>اس حملہ کے بعد بڑی تعداد میں مقامی باشندوں نے علاقوں کو خالی کر دیا ہے۔</p> <p>ناگالینڈ دعویٰ کرتا ہے کہ ریاست آسام اس کے بعض خطوں پر قابض ہے۔</p> <p>جبکہ ریاست آسام کا کہنا ہے کہ اس کے بعض علاقوں کو ناگالینڈ نے اپنے قبضے میں لے رکھا ہے۔</p> <p>ناگالینڈ کا ایک الگ ریاست کے طور پر قیام انیس ترستھ میں ہوا تھا جسے آسام کے ناگا قبائلیوں کی اکثریت والے اضلاع کو منقسم کر کے بنایا گیا تھا۔</p> <p>ناگا قبائل نے ریاست ناگالینڈ کے قیام کے لیے انیس سو چھپن میں مسلح جدوجہد کی شروعات کی تھی۔</p> <p>علیحدگی پسند تنظیم نیشنل سوشلسٹ کونسل آف ناگالینڈ کا عرصہ سے مطالبہ رہا ہے کہ 'گریٹر ناگالینڈ' کے قیام کے لیے آسام، منی پور اور اروناچل پردیش کے تمام ناگا علاقوں کا ناگالینڈ سے الحاق ہونا چاہیے۔</p> <p>ناگالینڈ کی حکومت دعویٰ کرتی ہے کہ اس کی ہزاروں کلومیٹر زمین آسام کے حصے میں ہے۔</p> <p>لیکن آسام کا الزام ہے کہ ناگالینڈ نے طاقت کے زور پر اس کے بہت بڑے خطے کو قبضے میں لے لیا ہے اور ایک مقام کو انتظامی امور کا نائب مرکز بھی بنا رکھا ہے جسے وہ نیوالینڈ کہتے ہیں۔</p>	<p>Nagas Set Fire in Assam</p> <p>On Wednesday, angry Naga tribesmen set 10 schools in Manipur on fire.</p> <p>According to police, hundreds of armed tribesmen of Nagaland set three villages of Gulleki and Sisagar in Assam.</p> <p>A large number of natives have vacated the area after this attack.</p> <p>Nagaland claims that Assam state is occupying some of its territory.</p> <p>While Assam state says that Nagaland has occupied some of its areas.</p> <p>Nagaland was established as a free state in 1963 which was created by dividing Assamese cities with Naga majority.</p> <p>Naga tribes started armed struggle for the creation of Nagaland state in 1956.</p> <p>The separatist Socialist Council of Nagaland has been claiming for a long time that for the creation of 'Greater Nagaland,' all the Naga areas of Assam, Manipur and Arunachal Pradesh should be joined with Nagaland.</p> <p>The Nagaland government claims that thousands of kilometers of its land lies in the Assamese part.</p> <p>But Assam accused Nagaland for occupying a very large part of its land by force and setting up a second centre of administrative affairs which they call Nevaland.</p>	<p>Has Imposed a Fire in Assam</p> <p>On Wednesday, the tribal mini pur enraged ten schools was also burnt.</p> <p>According to the police, the number of hundreds of armed tribesmen in the ratio of assam and three set the fire in the village.</p> <p>After this attack. Local residents in large numbers to the areas.</p> <p>Claim of assam. That this is the some regions.</p> <p>While of assam has said that this to some areas of his into custody</p> <p>A separate state of on 19 establishment of assam happened in Which the majority of the people of the districts was made.</p> <p>The state tribes for the establishment of the 19\$156 armed declare struggle of the beginning of.</p> <p>Separatist Council of National organization is the demand for a long time that 'greater', for the establishment of the assam, mini pur and all pradesh areas should be included with.</p> <p>The government of claim thousands of believes that the earth of assam.</p> <p>But has been accused of assam that the power of this on a large region has taken in the affairs and one Place, vice Center of which he is also</p>

Figure 1: An example of Urdu-English translation. Shown are an Urdu source document, a reference translation produced by a professional human translator, and machine translation output from a phrase-based model (Moses) without linguistic information, which is representative of state-of-the-art MT quality before the SIMT effort.

Named entities have been the focus of information extraction research since the Message Understanding Conferences of the 1980s (Grishman and Sundheim, 1996). Automatic taggers identify semantic types such as person, organization, location, date, facility, etc. In this research effort we tagged English documents using an HMM-based tagger derived from Identifinder (Bikel et al., 1999).

Modality is an extra-propositional component of meaning. In *John may go to NY*, the basic proposition is *John go to NY* and the word *may* indicates modality. Van der Auwera and Amman (2005) define core cases of modality: *John must go to NY* (epistemic necessity), *John might go to NY* (epistemic possibility), *John has to leave now* (deontic necessity) and *John may leave now* (deontic possibility). Many semanticists (Kratzer, 2009; von Stechow and Iatridou, 2009) define modality as quantifica-

tion over possible worlds. *John might go* means that there exist some possible worlds in which John goes. Another view of modality relates more to a speaker's attitude toward a proposition (Nirenburg and McShane, 2008; McShane et al., 2004). Modality resources built for this purpose have been described previously (Baker et al., 2010).

This paper will focus on a tree-grafting mechanism used to enrich the machine-translation output and on the resulting improvements to translation quality when the training process for the machine-translation systems included tagging of named entities and modality.

The next section provides the motivation behind the SIMT approach. Section 3 presents implementation details of the semantically-informed syntactic system. Section 4 describes the tree grafting algorithm. Section 5 provides the results of this work.

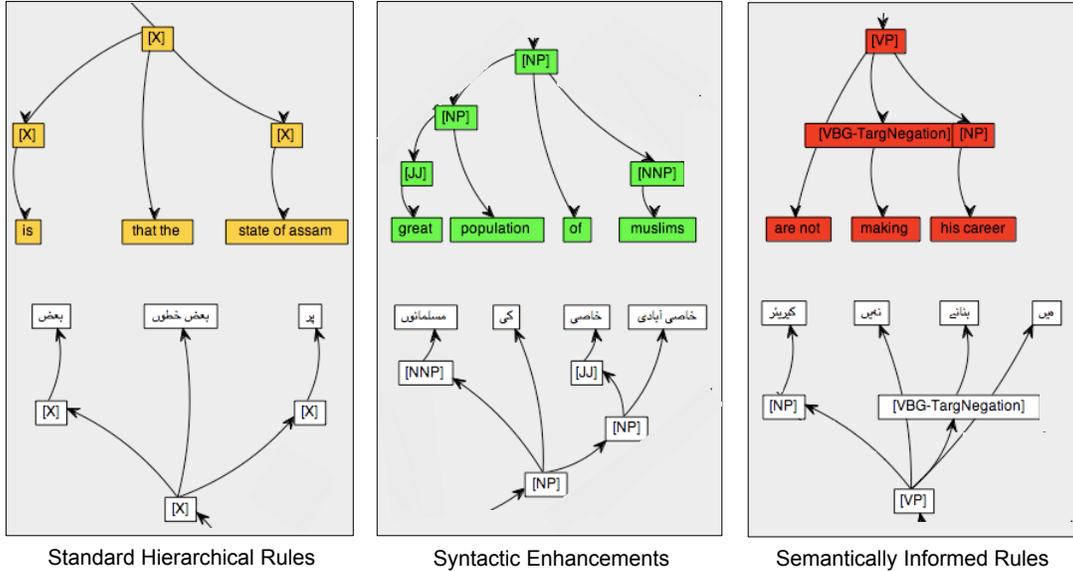


Figure 2: The evolution of a semantically informed approach to our synchronous context free grammars (SCFGs). At the start of summer the decoder used translation rules with a single generic non-terminal symbol, later syntactic categories were used, and by the end of the summer the translation rules included semantic elements such as named entities and modalities.

Following this, Section 6 examines work that is related to our approach. Finally, Section 7 presents conclusions and future work.

2 Motivation

The aim of the SIMT effort was to provide a generalized framework for representing structured semantic information, such as named entities and modality, and to investigate whether incorporating this sort of information into machine translation (MT) systems could produce better translations. The SIMT effort differs from other efforts in MT, most notably the DARPA Global Autonomous Language Exploitation (GALE) initiative, in at least two ways:

1. The SIMT effort worked on translation for a low-density language, with a minimal amount of bilingual training data. In GALE, hundreds of millions of words worth of bilingual texts are used to train statistical translation models. In the SIMT effort, only 1.7 million words of Urdu-English texts were available. Table 1 provides the data set sizes used in our experiments.
2. The SIMT effort showed significant improvements from incorporating syntax and semantics into machine translation, whereas syntactic translation models have not shown dramatic improvements in GALE’s Arabic-

English translation task. The improvements for Urdu translation described here are probably due to the fact that it is a low-resource, verb-final language and so requires generalization beyond phrase-based or hierarchical phrase-based models.

These differences created novel research directions for our effort, and resulted in promising findings that suggest that both syntactic and semantic information are critical for improving translation quality.

It is informative to look at an example translation to understand the challenges of translating important semantic entities when working with a low-resource language pair. Figure 1 shows an example taken from the 2008 NIST Urdu-English translation task, and illustrates the translation quality of a state-of-the-art Urdu-English system (prior to the SIMT effort). The small amount of training data for this language pair (see Table 1) results in significantly degraded translation quality compared, e.g., to an Arabic-English system that has more than 100 times the amount of training data.

The machine translation output in Figure 1 was produced using Moses (Koehn et al., 2007), a state-of-the-art phrase-based machine translation system that by default does not incorporate any linguistic information (e.g., syntax or morphology or translit-

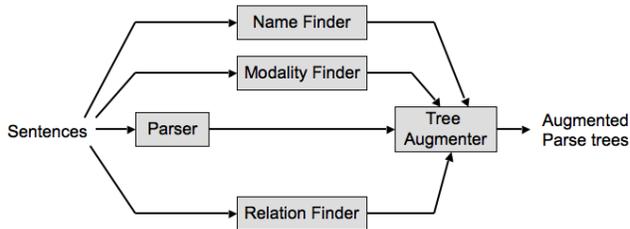


Figure 3: Workflow for producing semantically-grafted parse trees. The English side of the parallel corpus is automatically parsed, and also tagged with modality and named-entity markers. These tags are then grafted onto the syntactic parse trees. The relation finder was designed for additional tagging but was not implemented in the current work. (Future work will test relations as another component of meaning that may contribute toward improved MT output.)

eration knowledge). As a result, words that were not directly observed in the bilingual training data were untranslatable. Names, in particular, are problematic. For example, the lack of translation for *Nagaland* and *Nagas* induces multiple omissions throughout the translated text. This is because out of vocabulary words are deleted from the Moses output.

We use modality and named-entity tags as higher-order symbols inside the translation rules used by the translation models. Generic symbols in translation rules (e.g., the non-terminal symbol “X”) were replaced with structured information at multiple levels of abstraction, using a tree-grafting approach, as described in more detail in the following sections. Figure 2 illustrates the evolution of the translation rules that we used, first replacing “X” with grammatical categories and then with semantic categories.

set	lines	Urdu		English	
		tokens	types	tokens	types
training	202k	1.7M	56k	1.7M	51k
dev	981	21k	4k	19k	4k
devtest	883	22k	4k	19-20k	4k
test	1792	42k	6k	38-41k	5k

Table 1: The size of the various data sets used for the experiments in this paper including the training, development (dev), incremental test set (devtest) and blind test set (test). The dev/devtest was a split of the NIST08 Urdu-English test set, and the blind test set was NIST09.

3 Tree-Grafting to refine translation grammars with semantic categories

We use synchronous context free grammars (SCFGs) as the underlying formalism for our statistical models of translation. SCFGs provide a convenient and theoretically grounded way of incorporating linguistic information into statistical models of translation, by specifying grammar rules with syntactic non-terminals in the source and target languages. We refine the set of non-terminal symbols so that they not only include syntactic categories, but also semantic categories.

Chiang (2005) re-popularized the use of SCFGs for machine translation, with the introduction of his hierarchical phrase-based machine translation system, Hiero. Hiero uses grammars with a single non-terminal symbol “X” rather than using linguistically informed non-terminal symbols. When moving to linguistic grammars, we use the Syntax Augmented Machine Translation (SAMT) developed by Venugopal et al. (2007). In SAMT the “X” symbols in translation grammars are replaced with nonterminal categories derived from parse trees that label the English side of the Urdu-English parallel corpus.¹ We refine the syntactic categories by combining them with semantic categories. This progression is illustrated in Figure 2.

We extracted SCFG grammar rules containing named entities and modality using an extraction procedure that requires parse trees for one side of the parallel corpus. While it is assumed that these trees are labeled and bracketed in a syntactically motivated fashion, the framework places no specific requirement on the label inventory. We take advantage of this characteristic by providing the rule extraction algorithm with augmented parse trees containing syntactic labels that have named entities and modalities grafted onto them so that they additionally express semantic information.

Our strategy for producing semantically-grafted parse trees involves three steps:

1. The English sentences in the parallel training data are parsed with a syntactic parser. In our work, we used the lexicalized probabilistic

¹For non-constituent phrases, composite CCG-style categories are used (Steedman, 1999).

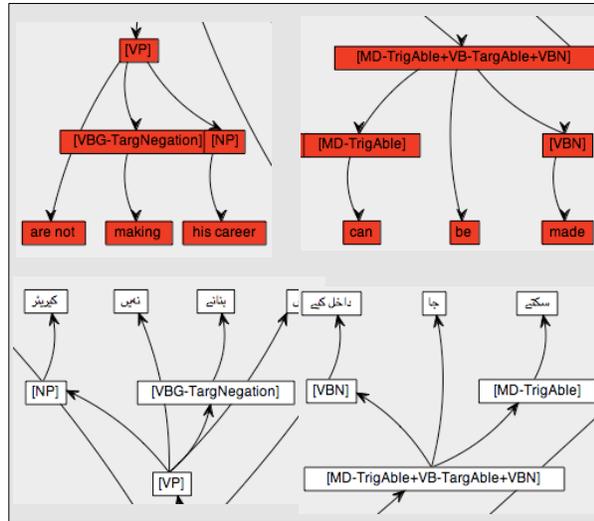
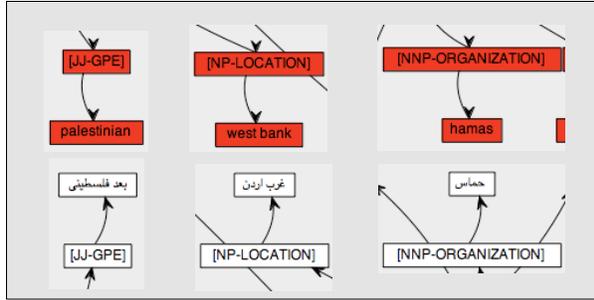


Figure 5: Example translation rules with named entity tags and modalities combined with syntactic categories.

parent node or nodes in the corresponding syntactic parse tree that dominate the word sequence covered by the tag. The following tests are then applied:

- If the semantic and syntactic units correspond exactly, graft the name of the semantic tag onto the highest corresponding syntactic constituent in the tree. For example, in Figure 4, the NNP “Lebanon” receives a GPE (geo-political entity) tag at the NP constituent level.
- For the case of named entities: If the semantic tag corresponds to words that are adjacent daughters in a syntactic constituent, but less than the full constituent, insert an NP node dominating those words into the parse tree, as a daughter of the original syntactic constituent. The name of the semantic tag is grafted onto the new NP node. This is a case of rule splitting.
- If a syntactic constituent selected for grafting has already been labeled with a semantic tag,

Named Entity	Example
AGE	50 years old
DATE	September 26, 2009
FACILITY	Southwestern Medical Center
GPE (Geo-political entity)	New York
GPE-ite	Australian
LOCATION	West Sea
MONEY	15,000 pounds
OCCUPATION	governor
ORGANIZATION	United Nations
ORGANIZATION-ite	marines
PERCENT	3.1 percent
PERSON	Tony Blair
TIME	2030 GMT

Table 2: Named entity tags

Require	NOTPermit
Permit	NOTRequire
Succeed	NOTSucceed
SucceedNegation	NOTSucceedNegation
Effort	NOTEffort
EffortNegation	NOTEffortNegation
Intend	NOTIntend
IntendNegation	NOTIntendNegation
Able	NOTAble
AbleNegation	NOTAbleNegation
Want	NOTWant
Belief	NOTBelief
Firm_Belief	NOTFirm_Belief
Negation	

Table 3: Modality tags with their negated versions

overlay that tag.

- If the words covered by the semantic tag fall across two syntactic constituents, do nothing. This is a case of crossing brackets.

Our tree-grafting procedure was simplified to accept a single semantic tag per syntactic tree node as the final result. The algorithm keeps the last tag seen as the tag of precedence. In practice, we established a precedence ordering for modality tags over named entity tags by grafting named entity tags first and modalities second. Our intuition was that, in case of a tie, finer-grained verbal categories would be more helpful to parsing than finer-grained nominal categories.² In case a word was tagged both as

²In testing we found that grafting named entities first and modalities last yielded a slightly higher Bleu score than the reverse order.

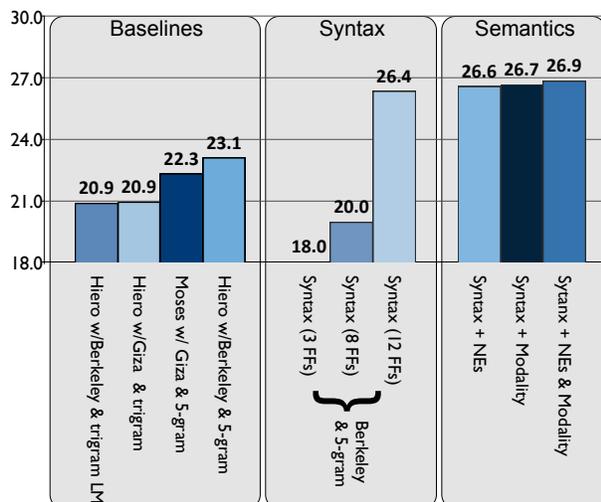


Figure 6: Results for a range of experiments conducted during the SIMT effort. Results show scores for baseline systems, which here include a phrase-based model (Moses) and a hierarchical phrase-based model (Hiero), neither of which make use of syntactic information. These also show the substantial improvements when syntax is introduced, along with different numbers of feature functions (FFs), and further improvements from semantic elements. The scores are lowercased Bleu calculated on the held-out devtest set.

a modality target and a modality trigger, we gave precedence to the target tag. This is because, while modality targets vary, modality triggers are generally identifiable with lexical items. Finally, we used a simplified specificity ordering of modality tags, borrowing from an approach described in (Baker et al., 2010), to ensure precedence of more specific tags over more general ones. Table 3 lists the modality types from highest (Require modality) to lowest (Negation modality) precedence.³

5 Results

Figure 6 gives the results for a number of experiments conducted during the SIMT effort.⁴ The ex-

³Future work could include exploring additional methods of resolving tag conflicts or combining tag types on single nodes, e.g. by inserting multiple intermediate nodes (effectively using unary rewrite rules) or by stringing tag names together.

⁴These experiments were conducted on the devtest set, containing 883 Urdu sentences (21,623 Urdu words) and four reference translations per sentence. The Bleu score for these experiments is measured on uncased output, which in general should be higher, but the devtest effectively had only three reference translations. This explains why the scores are lower than the

periments are broken into three groups: baselines, syntax, and semantics. To contextualize our results we experimented with a number of different baselines that were composed from two different approaches to statistical machine translation—phrase-based and hierarchical phrase-based SMT—along with different combinations of language model sizes and word aligners. Our best performing baseline was a Hiero model with a 5-gram language model and word alignments produced using the Berkeley aligner. The Bleu score for this baseline on the development set was 23.1 Bleu points.

After experimenting with syntactically motivated grammar rules, we conducted three experiments on the effects of incorporating semantic elements (e.g., named entities and modality markers) into the translation grammars. In our devtest set our taggers tagged on average 3.5 named entities (NEs) per sentence and 0.35 modalities per sentence. These were included by grafting NEs and modality markers onto the parse trees. Individually, each of these made modest improvements over the syntactically-informed system alone. Grafting named entities onto the parse trees improved the Bleu score by 0.2 points. Modalities improved it by 0.3 points. Doing both simultaneously had an additive effect and resulted in a 0.5 Bleu score improvement over syntax alone. This improvement was the largest improvement that we got from anything other than the move from linguistically naive models to syntactically-informed models.

Figure 7 shows example output from the final SIMT system. Notice that even in the title of the article, the SIMT system produces much more coherent English output than that of the linguistically naive system. The figure also shows improvements due to transliteration, which are described in Irvine et al. (2010). The scores reported in Figure 6 do not include transliteration improvements.

6 Related Work

This section describes related work in monolingual techniques for augmenting parsing, where parsing is applied to one language in the parallel text.

Our tree-grafting approach is related to a technique used for tree augmentation in (Miller et al., 2000), where parse-tree nodes are augmented with scores on the NIST 2009 test set.

semantic categories. Miller et al. augment tree nodes with named entities and relations, while we used named entities and modalities. The parser is subsequently retrained for both semantic and syntactic processing. The semantic annotations were done manually by students following a set of guidelines and then merged with the syntactic trees automatically. In our work we tagged our corpus with entities and modalities automatically and then grafted them onto the syntactic trees automatically, for the purpose of training a statistical machine translation system. An added benefit of the extracted translation rules is that they are capable of producing semantically-tagged Urdu parses, despite that the training data were processed by only an English parser and tagger.

Related work in syntax-based MT includes (Huang and Knight, 2006), where a series of syntax rules are applied to a source language string to produce a target language phrase structure tree. The Penn English Treebank (Marcus et al., 1993) is used as the source for the syntactic labels and syntax trees are relabeled to improve translation quality. In this work, node-internal and node-external information is used to relabel nodes, similar to earlier work where structural context was used to relabel nodes in the parsing domain (Klein and Manning, 2003). Klein and Manning’s methods include lexicalizing determiners and percent markers, making more fine-grained VP categories, and marking the properties of sister nodes on nodes. All of these labels are derivable from the trees themselves and not from an auxiliary source.

In the parsing domain, the work of (Petrov and Klein, 2007) is related to the current work. Petrov and Klein use a technique of rule splitting and rule merging in order to refine parse trees during machine learning. Hierarchical splitting leads to the creation of learned categories that have linguistic relevance, such as a breakdown of a determiner category into two subcategories of determiners by number, i.e., *this* and *that* group together as do *some* and *these*. We use rule splitting in cases where a semantic category is inserted as a node in a parse tree, after the English side of the corpus has been parsed by a statistical parser (as described in section 4).

7 Conclusions and Future Work

We have described a technique for translation that shows particular promise for low-resource languages. We have integrated linguistic knowledge into statistical machine translation in a unified and coherent framework. We demonstrated that augmenting hierarchical phrase-based translation rules with semantic labels (through “grafting”) resulted in a 0.5 Bleu score improvement over syntax alone.

Although our largest gains were from syntactic enrichments to the Hiero model, demonstrating success on the integration of new semantic aspects of language bodes well for future improvements based on the incorporation of other semantic aspects, e.g., relations and temporal knowledge, into the translation rules, would further improve the translations. The syntactic framework is unique in its ability to support the exploration of the impact of many different types of semantics on MT quality.

Our findings indicate that the use of syntactic and semantic information radically improves translation quality for low-resource languages with different word order than English. Urdu has SOV (subject, object, verb) word order compared to English SVO (subject, verb, object). Thus, our observed improvements are likely to be transferable to languages like Korean and Farsi, as well as a host of other low-resource languages with different word order.

The work presented here represents the first small steps toward a full integration of MT and semantics. Efforts underway in DARPA’s GALE program have already demonstrated the potential for combining MT and semantics (termed *distillation*) to answer the information needs of monolingual speakers using multilingual sources. In previous work, however, semantic processing proceeded largely independently of the MT system, operating only on the translated output. Our approach is significantly different in that it combines syntax, semantics, and MT into a single model, offering the potential advantages of joint modeling and joint decision-making. It would be interesting to explore whether the integration of MT with syntax and semantics can be extended to provide a single-model solution for tasks such as cross-language information extraction and question answering, and to evaluate our integrated approach, e.g., using GALE distillation metrics.

Acknowledgments

We thank Aaron Phillips for help converting the output of the entity tagger for ingest by the tree-grafting program. We also thank Basis Technology Corporation for their generous contribution of software components to this work. This work is supported, in part, by the Johns Hopkins Human Language Technology Center of Excellence, by the National Science Foundation under grant IIS-0713448, and by BBN Technologies under GALE DARPA/IPTO Contract No. HR0011-06-C-0022. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsor.

References

- Johan Van Der Auwera and Andreas Ammann. 2005. Overlap between situational and epistemic modal marking. In Martin Haspelmath, Matthew S. Dryer, David Gil, and Bernard Comrie, editors, *World Atlas of Language Structures*, chapter 76, pages 310–313. Oxford University Press.
- Kathryn Baker, Michael Bloodgood, Bonnie J. Dorr, Nathaniel W. Filardo, Lori Levin, and Christine Pitatko. 2010. A modality lexicon and its use in automatic tagging. In *7th International Conference on Language Resources and Evaluation (LREC)*, Malta, May. Language Resources and Evaluation Conference.
- Daniel M. Bikel, Richard Schwartz, and Ralph M. Weischedel. 1999. An algorithm that learns what’s in a name. *Mach. Learn.*, 34(1-3):211–231.
- David Chiang. 2005. A hierarchical phrase-based model for statistical machine translation. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL-2005)*, Ann Arbor, Michigan.
- Ralph Grishman and Beth Sundheim. 1996. Message Understanding Conference-6: A brief history. In *COLING*, pages 466–471.
- Bryant Huang and Kevin Knight. 2006. Relabeling syntax trees to improve syntax-based machine translation quality. In *HLT-NAACL*, New York.
- Ann Irvine, Mike Kayser, Zhifei Li, Wren Thornton, and Chris Callison-Burch. 2010. Integrating output from specialized modules in machine translation: Transliteration in Joshua. *The Prague Bulletin of Mathematical Linguistics*, 93:107–116.
- Dan Klein and Christopher D. Manning. 2003. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 423–430.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the ACL-2007 Demo and Poster Sessions*.
- Angelika Kratzer. 2009. Plenary address at the annual meeting of the Linguistic Society of America.
- Zhifei Li, Chris Callison-Burch, Chris Dyer, Sanjeev Khudanpur, Lane Schwartz, Wren Thornton, Jonathan Weese, and Omar Zaidan. 2009. Joshua: An open source toolkit for parsing-based machine translation. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 135–139, Athens, Greece, March. Association for Computational Linguistics.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: the Penn Treebank. *Computational Linguistics*, 19(2):313–330.
- Marjorie McShane, Sergei Nirenburg, and Ron Zacharsky. 2004. Mood and modality: Out of the theory and into the fray. *Natural Language Engineering*, 19(1):57–89.
- Scott Miller, Heidi J. Fox, Lance A. Ramshaw, and Ralph M. Weischedel. 2000. A novel use of statistical parsing to extract information from text. In *Proceedings of Applied Natural Language Processing and the North American Association for Computational Linguistics*.
- Nirenburg and McShane. 2008. The formulation of modalities (speaker attitude) in OntoSem.
- Slav Petrov and Dan Klein. 2007. Learning and inference for hierarchically split PCFGs. In *AAAI 2007 (Nectar Track)*.
- Alexander E. Richman and Patrick Schone. 2008. Mining wiki resources for multilingual named entity recognition. In *Proceedings of ACL-08: HLT*, pages 1–9, Columbus, Ohio, June. Association for Computational Linguistics.
- Mark Steedman. 1999. Alternating quantifier scope in CCG. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL)*, College Park, Maryland.
- Ashish Venugopal and Andreas Zollmann. 2009. Grammar based statistical MT on Hadoop: An end-to-end toolkit for large scale PSCFG based MT. *Prague Bulletin of Mathematical Linguistics*, 91.
- Ashish Venugopal, Andreas Zollmann, and Stephan Vogel. 2007. An efficient two-pass approach to synchronous-CFG driven statistical MT. In *Proceedings of the Human Language Technology Conference of the North American chapter of the Association for Computational Linguistics (HLT/NAACL-2007)*, Rochester, New York.
- Kai von Fintel and Sabine Iatridou. 2009. Morphology, syntax, and semantics of modals. Lecture notes for 2009 LSA Institute class.

pre-SIMT	SIMT	Reference
<p>'first nuclear experiment in 1990 was'</p> <p>Thomas red Unilever National Laboratory of the United States in ویپن designer, are already working on the book of Los ایلبوس National Laboratory ڈیسی, former director of the technical Laboratory انڈیجینس written with the cooperation of سٹیلمن.</p> <p>This book 'nuclear express: political history and the expansion of bomb' has been written, and the two writers have also claimed that the country has made nuclear bomb is he or any other country's nuclear secrets to چرای or that of any other nuclear power cooperation is achieved.</p> <p>Thomas Reid said in a news عواس interview that in 1990 in the era of Benazir Bhutto China had the experience of Pakistan's first nuclear bomb.</p> <p>Thomas red said that on the basis of many reasons he was sure that China had the experience of Pakistan's first nuclear bomb.</p> <p>reasons in the bomb design and the China scientists mentioned During the conversation with Information.</p> <p>He further said that this was the reason that only two weeks in Pakistan in 1998 and within three days in response to India's nuclear experience to nuclear experiment was able to.</p> <p>Thomas red reminded that in 61 in Russia has suddenly nuclear experience and was in response to the United States were to experience began 17 days in despite the fact that the United States had the bomb from a long period.</p> <p>He further said that the nuclear bomb in 1998 that Pakistan may experience of what was he was made from very carefully and confidence was to meet on the Pakistani scientists.</p> <p>Thomas was red when this question that China has provided the nuclear technology to Pakistan, In response, He said as China and India was joint enemy of Pakistan.</p> <p>He said that it is also present proof of that Dr. Abdul Qadeer Khan after the Chinese design. apart from this, The کی obtained documents in Libya were is also confirmed it from them.</p> <p>To another question whether the joint nuclear tests is common, He said in سروے in the US Open in the desert نواڈا servants for Britain's nuclear experiment.</p> <p>He said that we are guesses also believed that Israel should also provide access to the results of the this experience.</p> <p>Thomas Reid and this question was also is China's program of nuclear proliferation.</p> <p>On this Thomas Reid said that in 1991 to set up in China ایرفران city of Iran nuclear installation instructions for the raw material, and ڈی is provided.</p> <p>Besides this, China has provided of enrichment plant in Iran خوج also.</p> <p>He said that China, North Korea, Iran, Syria, Pakistan, through Egypt, Libya and Yemen is to provide nuclear technology.</p>	<p>The First Nuclear Test Was in 1990.</p> <p>Thomas red of the United States, the National Laboratory in designer are already working on the book of Los Alamos National Laboratory, former director of the technical intelligence, with the cooperation of Diana steelman wrote.</p> <p>This book under the title of the spread of nuclear expressway: the political history of the bomb and this has been written and the two writers have claimed that the country also has made nuclear bomb or any other country, Korea nuclear secrets, or any of the other nuclear power cooperation.</p> <p>Thomas Reid said in an interview to the news that Benazir Bhutto in 1990 in the era of China The first nuclear bomb test.</p> <p>Thomas Reid said that on the basis of several reasons to believe that China has the first nuclear bomb test.</p> <p>reasons of this bomb design and China information obtained during the conversation with the scientists.</p> <p>He further said that this was the reason that Pakistan only two weeks in 1998 and was able to nuclear experiment in response to India's nuclear experience within three days.</p> <p>Thomas Reid had reminded that Russia has sudden nuclear experiment in 61 and in response to the United States to experiment in 17 days. despite the fact that the United States had quite a long period from the bomb.</p> <p>He further said that the experience of the atom bomb Pakistan in May 1998, he was to make very carefully and confidence on Pakistani scientists.</p> <p>Thomas red when was this question that China has provided to Pakistan nuclear technology, he said in response to China, Pakistan and India as a common enemy.</p> <p>He said that there is evidence that Dr. Abdul Qadeer Khan has also used the Chinese design after did the documents from Libya in were also confirmed to them.</p> <p>On a question whether the joint nuclear tests is common, he said in Hanoi in the United States In the Nevada desert open detainees for Britain's nuclear experiment.</p> <p>He said that it also that Israel should also provide access to the results of this experience.</p> <p>Thomas red with this question was whether China's program of nuclear proliferation.</p> <p>Thomas red said that in the Iranian city of China with 491 in Spain, the raw material for nuclear installation instructions and provides.</p> <p>In addition to this, China has also provided the enrichment plant in Cairo.</p> <p>He said that China would provide nuclear technology to North Korea, Iran, Syria, Pakistan, Egypt, Libya and Yemen.</p>	<p>"First Nuclear Experiment conducted in 1990"</p> <p>Thomas Reed, who has worked as a weapon designer in Lawrence Livermore National Laboratory of America, has written this book in collaboration with Danny Stillman, the former Director of Technical Intelligence of Los Alamos National Laboratory.</p> <p>This book has been written with the title 'Nuclear Express: A Political History of the Bomb And its Proliferation,' and in this both the authors have claimed that any country that has made an atomic bomb has either stolen the nuclear secrets of another country or has had cooperation with some other nuclear power.</p> <p>Thomas Reed said in an interview to US News that in 1990, in the era of Benazir Bhutto, China had conducted the experiment of Pakistan's first nuclear bomb.</p> <p>Thomas Reed said that he is convinced on the basis of several reasons that China has conducted the experiment of Pakistan's first nuclear bomb.</p> <p>Those reasons include the design of the bomb and information obtained while talking to the scientists of China.</p> <p>He further said that this was the reason why in 1998, Pakistan was able to conduct a nuclear experiment just in two weeks and three days in response to India's nuclear experiment.</p> <p>Thomas Reed also reminded that in 1961 Russia suddenly carried out a nuclear experiment and it took 17 days for America to do the experiment in response to this, although America already had this bomb for awhile.</p> <p>He further said that the atom bomb, whose experiment was done in 1998 by Pakistan, was developed with extreme care and Pakistani scientists had full confidence in it.</p> <p>When Thomas Reed was asked if China had provided the nuclear technology to Pakistan, he replied that India was a common enemy of China and Pakistan.</p> <p>He said that the proof to this also exists in that Dr. Abdul Qadeer Khan used the Chinese design, and, apart from this, the documents retrieved from Libya afterwards also proved this.</p> <p>To another question as to whether it is usual to carry out nuclear experiments with others, he said that in 1990 America openly conducted a nuclear experiment for Britain in the desert of Nevada.</p> <p>He said that we may also presume that Israel, too, was given access to the results of this experiment.</p> <p>Thomas Reed was also asked whether China's nuclear proliferation program is active.</p> <p>On this, Thomas Reid said that since 1991, China has been providing raw material, instructions, and designs for the nuclear structure situated in Isfahan, a city in Iran.</p> <p>Besides this, China has also provided an enrichment plant to Iran in Karaj.</p> <p>He said that China has been providing nuclear technology to Iran, Syria, Pakistan, Egypt, Libya, and Yemen through North Korea.</p>

Figure 7: An example of the improvements to Urdu-English translation before and after the SIMT effort. Output is from the baseline Hiero model, which does not use linguistic information, and from the final model, which incorporates syntactic and semantic information.