Moses & Treex Hybrid MT Systems Bestiary

Rudolf Rosa, Martin Popel, Ondřej Bojar, David Mareček, Ondřej Dušek
Charles University, Faculty of Mathematics and Physics,
Institute of Formal and Applied Linguistics,
Malostranské náměstí 25, 118 00 Prague, Czech Republic
{rosa,popel,bojar,marecek,odusek}@ufal.mff.cuni.cz

Abstract

Moses is a well-known representative of the phrase-based statistical machine translation systems family, which are known to be extremely poor in explicit linguistic knowledge, operating on flat language representations, consisting only of tokens and phrases. Treex, on the other hand, is a highly linguistically motivated NLP toolkit, operating on several layers of language representation, rich in linguistic annotations. Its main application is TectoMT, a hybrid machine translation system with deep syntax transfer. We review a large number of machine translation systems that have been built over the past years by combining Moses and Treex/TectoMT in various ways.

1 Introduction and Motivation

Phrase-based statistical machine translation (PB-SMT) systems, which have been the state-of-the-art approach to machine translation (MT) for many years, are known to contain very little explicit linguistic knowledge. While this characteristic has been at the core of their success, enabling fast development, training and tuning of the systems (as long as sufficient amounts of parallel data are available), it becomes a double-edged sword in many cases, e.g., when translating into a morphologically-rich language with frequent long-range dependencies, such as Czech.

It has been shown that many language phenomena hard to handle for the PB-SMT systems can be easily dealt with by linguistically motivated MT systems – although these systems often have other shortcomings, such as a tendency to translate very lexically, in a one-to-one fashion, due to lacking the (non-linguistic) phrase-based representation employed in PB-SMT systems.

This situation invites researchers to attempt to combine these conceptually different systems in a clever way so that their strengths combine and their shortcomings cancel out. In our paper, we review a set of such attempts, performed with Moses, a prominent representative of the PB-SMT systems, and Treex, a linguistically motivated NLP framework, featuring, among other, a full-fledged deep syntactic MT system, TectoMT.

As Treex and TectoMT have been primarily developed to process Czech language and to perform English-to-Czech translation, most of the existing system combination experiments have been performed on the English-to-Czech language pair.[1] Therefore, we limit ourselves to this setting in our work.

2 Individual Systems

2.1 Moses

Moses (Koehn et al., 2007) is a standard PB-SMT system. It features simple rule-based tokenization and true-casing scripts, which are sometimes language-specific, but the core of the decoder is purely statistical and oblivious of any linguistics. It relies on GIZA++ (Och and Ney, 2003) to compute word alignment of the training parallel corpus, used to extract lexicons and phrase tables that provide the knowledge of translation options to the decoder. A word-based language model is used to score possible translations, so that a fluent one can be produced as the output.

[1] A few combinations have been also applied to other translation pairs.

Proceedings of the 2nd Deep Machine Translation Workshop (DMTW 2016), pages 1–10,
Lisbon, Portugal, 21 October 2016.

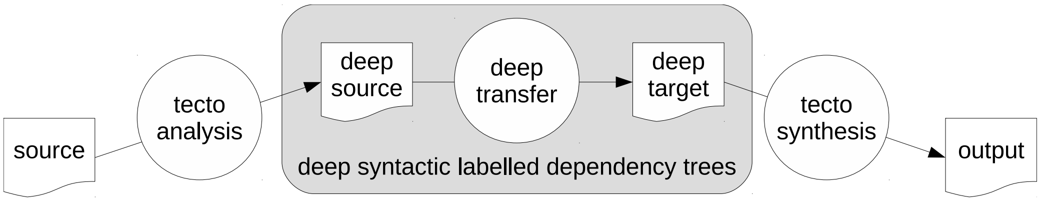

Figure 1: TectoMT

2.1.1 Factored Moses

In the more recent experiments that we report, the Moses system used is actually the Factored Moses of Bojar et al. (2012). It translates the source English text into a factored representation of Czech, where each word is represented by a tuple of a word form and a corresponding part-of-speech (PoS) tag. This enables Moses to use an additional language model which operates on PoS tags instead of word forms. This helps overcome data sparsity issues of the word-based language model and thus leads to a higher output quality, especially to its better grammaticality. Factored Moses is trained on parallel corpora pre-analyzed by Treex.

2.2 Treex

Treex[2] (Popel and Žabokrtský, 2010; Žabokrtský, 2011) is a linguistically motivated NLP framework. It consists of a large number of smaller components performing a specific NLP-task (blocks), both Treex-specific as well as Treex-wrapped external tools, which can be flexibly combined into processing pipelines. Sentences are represented by surface and deep syntactic dependency trees, richly annotated with numerous linguistic attributes, similarly to the Prague Dependency Treebank (Hajič, 1998).

2.2.1 TectoMT

The main application of Treex is TectoMT[3] (Žabokrtský et al., 2008; Dušek et al., 2015), a linguistically motivated hybrid machine translation system. Its pipieline consists of three main steps: analysis of each source sentence up to t-layer (a deep syntactic representation of the sentence in a labelled dependency t-tree), transfer of the source t-tree to the target t-tree (i.e., the translation per se), and generation of the target sentence from the target t-tree (see Figure 1).

The transfer is performed by copying the t-tree structure and *grammatemes*[4] (attributes describing grammatical meaning) from source, and predicting target lemmas and *formemes*[5] (deep morphosyntactic attributes (Dušek et al., 2012)) using a set of machine-learned translation models. In the current transfer implementation, TectoMT translates t-tree nodes one-to-one; however, as function words are abstracted from, a one-to-one correspondence between t-trees in different languages is present in most cases.

3 System Combinations

This section contains description and evaluation of several system combination setups. We list a number of combinations of Moses and Treex/TectoMT that we are aware of, both successful and unsuccessful.

Results of automatic evaluation of the setups, as reported in available literature,[6] are provided in Table 1. We report absolute *differences* in BLEU scores[7] versus the base systems, rather than the absolute scores themselves – the setups were evaluated on many different test sets, and it is well known that BLEU scores are not directly comparable across datasets. Still, for each of the references in Table 1, we also list the absolute scores of the base system(s) in Table 2. We round up the scores to one decimal digit.

[2] `http://ufal.mff.cuni.cz/treex`

[3] `http://ufal.mff.cuni.cz/tectomt`

[4] `https://ufal.mff.cuni.cz/pdt2.0/doc/manuals/en/t-layer/html/ch05s05.html`

[5] `https://ufal.mff.cuni.cz/pcedt2.0/en/formemes.html`

[6] Except for "Moses + TectoMT post-editing" (Section 3.4), which we ran and evaluated ourselves.

[7] The scores are either case-sensitive or case-insensitive BLEU scores, depending on what was reported in the referenced paper. We do not include information on statistical significance of the score differences, as most of the authors did not report that. We kindly ask the interested reader to refer directly to the referenced papers or to their authors for any further details.

	Setup	Δ BLEU versus base Moses	 TectoMT	Reference
§ 3.1	TectoMoses: TectoMT with Moses transfer		−2.2	Popel (2015)
§ 3.2	PhraseFix: TectoMT + Moses post-editing		+2.7	Bojar et al. (2013a)
			+3.2	Galuščáková et al. (2013)
§ 3.3	Moses + Moses post-editing, simple	−0.1		Rosa (2013)
	Moses + Moses post-editing, TwoStep	−0.1		Bojar and Kos (2010)
§ 3.4	Google Translate + TectoMT post-editing	*−0.9		Majliš (2009)
	Moses + TectoMT post-editing	−2.4	+2.4	Section 3.4 & Bojar et al. (2016)
§ 3.5	Moses + Depfix post-editing	+0.1		Mareček et al. (2011)
		+0.1		Rosa et al. (2012)
		+0.4		Rosa (2013)
§ 3.6	Joshua + Treex pre-processing	**+0.5		Zeman (2010)
	Moses + Treex pre-/post-processing	+0.4		Rosa et al. (2016)
§ 3.7	Two-headed Chimera: Moses + TectoMT		+4.7	Bojar et al. (2013a)
		+0.6	+5.4	Bojar et al. (2013b)
			+5.5	Bojar et al. (2014)
		+1.1	+5.3	Bojar et al. (2015)
		+1.6		Bojar and Tamchyna (2015)
		+1.3	+6.1	Bojar et al. (2016)
§ 3.8	Chimera: Moses + TectoMT + Depfix		+5.0	Bojar et al. (2013a)
		+0.5	+5.3	Bojar et al. (2013b)
			+5.7	Bojar et al. (2014)
		+1.2	+5.4	Bojar et al. (2015)
		+1.5	+6.3	Bojar et al. (2016)
		+1.1		Tamchyna et al. (2016)

Table 1: System combinations. Difference in BLEU versus the Moses and/or TectoMT base system; * versus Google Translate, ** versus Joshua.

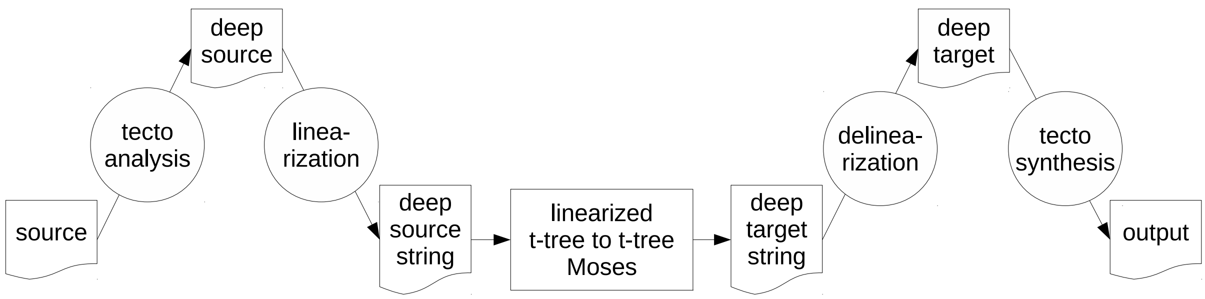

Figure 2: TectoMoses: TectoMT with Moses Transfer

While most of the setups have been properly described and evaluated in a peer-reviewed publication, others, especially some of the unsuccessful ones, were never properly published; in such cases, the descriptions and results are based on semi-official materials provided by the authors of the experiments.

3.1 TectoMoses: TectoMT with Moses Transfer

In the TectoMoses experiment of Popel (2013), which is depicted in Figure 2, the transfer step of TectoMT is substituted with Moses. This means that after the analysis to t-layer, each source-language t-tree is linearized into a sequence of lemmas and formemes (either as two factors, or interleaved). This linearized sequence is translated by Moses (trained on such data) into a target-language sequence of lemmas and formemes. Afterwards, dependencies are projected (using Moses alignment) from the source t-tree to the target sequence to reconstruct the target t-tree. Grammatemes and other attributes are also projected along the alignment. Finally, target-language synthesis is performed (as usual in TectoMT).

TectoMT's main transfer is isomorphic, which means translating one t-node to one t-node and keeping the dependency structure of the t-tree unchanged. This is much more powerful that surface word-to-word translation because t-nodes can represent e.g. complex verb forms ("*have been done*" is translated as "*bylo uděláno*"). However, there are still many cases which cannot be translated isomorphically on the t-layer. One of the advantages of TectoMoses is that it allows non-isomorphic transfer on t-layer, e.g.

System	BLEU	Reference
TectoMT	14.2	Bojar et al. (2013a)
	14.7	Bojar et al. (2013b)
	14.7	Galuščáková et al. (2013)
	15.4	Bojar et al. (2014)
	13.9	Bojar et al. (2015)
	12.4	Popel (2015)
	14.7	Bojar et al. (2016)
Moses	14.2	Bojar and Kos (2010)
	16.0	Mareček et al. (2011)
	15.4	Rosa et al. (2012)
	16.4	Rosa (2013)
	19.5	Bojar et al. (2013b)
	17.6	Bojar et al. (2015)
	22.6	Bojar and Tamchyna (2015)
	19.5	Bojar et al. (2016)
	19.1	Tamchyna et al. (2016)
	23.3	Rosa et al. (2016)
Google Translate	5.3	Majliš (2009)
Joshua	8.6	Zeman (2010)

Table 2: Base systems.

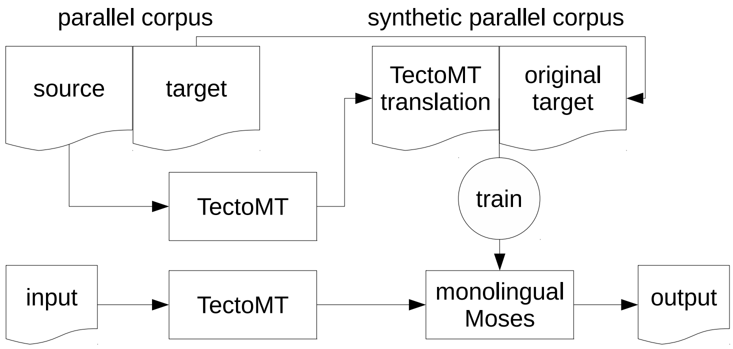

Figure 3: PhraseFix

translating one t-node with two or more t-nodes or deleting some t-nodes.[8] It also uses MERT tuning and it should scale with more training data. In the experiments with two factors (Popel, 2013), two language models were used: one for lemmas and one for formemes. Unfortunately, the TectoMoses experiment brought negative results, presumably due to additional noise introduced by the added transformations.

3.2 PhraseFix: TectoMT with Moses Post-editing

The PhraseFix system of Galuščáková et al. (2013) is based on the work of Simard et al. (2007), who introduced the idea of automatically post-editing a first-stage MT system by a second-stage MT system, trained to "translate" the output of the first-stage system into a reference translation. This has been shown to be particularly beneficial for conceptually different MT systems. In PhraseFix, the source English side of the CzEng parallel corpus of Bojar and Žabokrtský (2009) is translated by TectoMT into Czech, and Moses is then trained in a monolingual setting to translate the TectoMT-Czech into reference-Czech, i.e., the target side of CzEng (see Figure 3). Evaluation shows that this approach works well in principle, significantly improving the quality of the output as compared to the base TectoMT system. However, it does not surpass the translation quality provided by a standard standalone bilingual Moses.

3.3 Moses with Moses Post-editing

In case one does not have two different systems to combine, the simple approach of Oflazer and El-Kahlout (2007) can always be tried, who were the first to report translation quality improvements by

[8] PhraseFix (Section 3.2) also allows non-isomorphic translation, but only as post-processing. All other Moses-based systems (including Chimera, Section 3.7 & Section 3.8) allow non-isomorphic translations, but their transfer is on the t-layer.

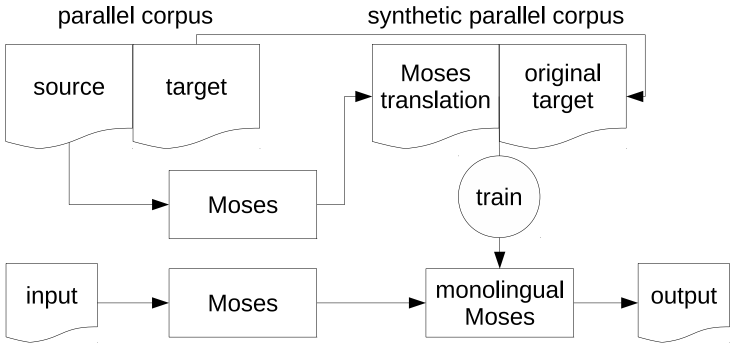

Figure 4: Simple post-editing of Moses by Moses

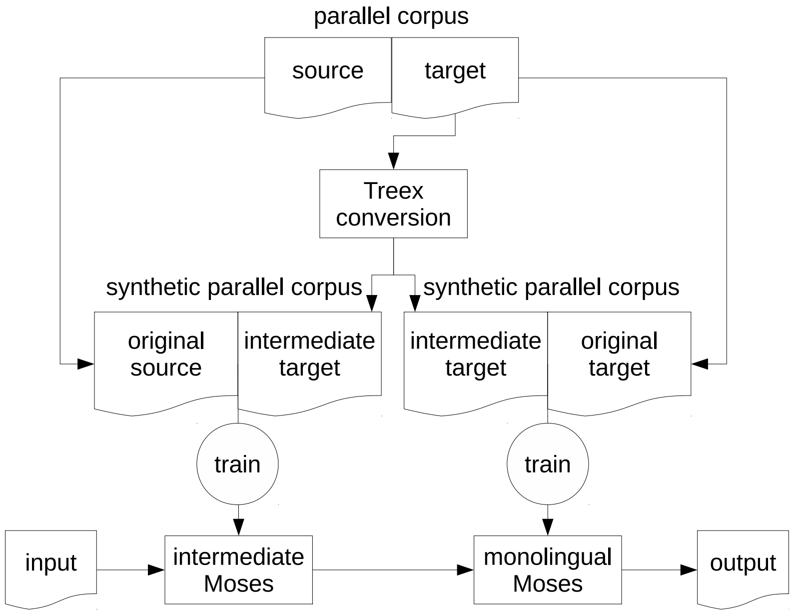

Figure 5: TwoStep Moses translation

training Moses to post-edit its own output. The setup, shown in Figure 4, is generally identical to that described in Section 3.2, except for using a standard bilingual Moses as the first-stage system, and then again Moses, this time in a monolingual setting, as the second-stage system. This setup was implemented and evaluated for English-to-Czech translation in (Rosa, 2013, section 7.4.1), but no improvements were found; based on a review of previous papers reporting positive results, the authors noted that this approach is probably only useful in cases where the available parallel training corpus is very small.

A more elaborate attempt in the same direction was presented as the TwoStep setup of Bojar and Kos (2010), this time bringing in Treex as well. TwoStep uses a first-stage Moses to translate from English into *intermediate Czech*, where each word is represented by a tuple of its lemma and a label marking several morphological features (such as detailed PoS, morphological number, grade, and negation). The second-stage Moses then translates from intermediate Czech into Czech (see Figure 5). The conversion of Czech into intermediate Czech is performed by a Treex pipeline described by Bojar and Žabokrtský (2009), with the main component being the Morče tagger of Spoustová et al. (2007). Unfortunately, this complex setup has not been found to have any benefit either.

3.4 Moses with TectoMT Post-editing

This setup uses Moses with transfer-less TectoMT as a post-editing tool (see Figure 6). A transfer-less TectoMT performs a tecto-analysis of the input, and then immediately proceeds with the tecto-synthesis

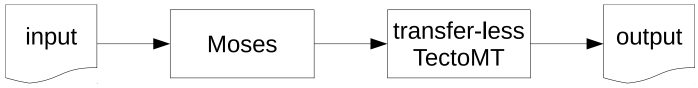

Figure 6: Moses with TectoMT post-editing

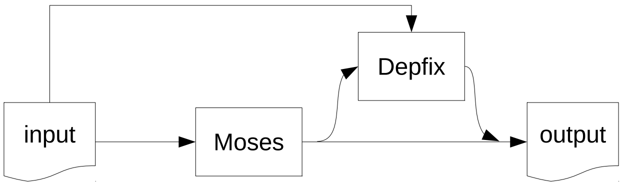

Figure 7: Moses with Depfix post-editing

of the output, completely omitting the interlingual transfer step.

Theoretically, analysis and a subsequent synthesis of a correct sentence should lead to the output being identical to the input (except for real synonymy). The motivation of Moses with transfer-less TectoMT post-editing was that incorrect sentences should be fixed this way, especially with respect to grammatical agreement. However, even the first assumption of identical output for correct sentences is not always true in practice, as some of the Treex blocks are not 100% accurate. Unfortunately, the assumption about fixing incorrect sentences also did not stand the practical test, mainly because the incorrect sentence on input tends to confuse the analysis pipeline and often leads to a largely incorrect analysis being produced (even if we disregard the fact that it is hard to define a correct analysis of an incorrect sentence).

We have been unable to find any work evaluating this particular setup, apart from the project of Majliš (2009), who applied TectoMT post-editing to Google Translate.[9] Therefore, we rerun the experiment ourselves, using current TectoMT[10] to post-edit the output of Moses obtained from the website of the WMT 2016 translation task (Bojar et al., 2016),[11] confirming the negative result reported by Majliš.

3.5 Moses with Depfix Post-editing

Similarly to the previous setup, Moses is complemented by a post-editing system implemented in Treex; this time, the system is Depfix (see Figure 7). Depfix (Mareček et al., 2011; Rosa, 2014) consists of several dozens rule-based post-editing Treex blocks. It focuses mainly on enforcing grammatical correctness, e.g., marking the subject and object by inflectional endings based on analysis of the source sentence, or inflecting adjectives to morphologically agree with their head nouns in gender, number, and case. However, contrary to the TectoMT post-editing (Section 3.4), it only modifies the erroneous parts of the output, thus avoiding generating too much noise; its second strength is the availability of the source analysis to the post-editing blocks, which enables them to make better-informed decisions regarding the intended meaning of the target sentence. This leads to a small but consistent improvement in BLEU.

3.6 Moses with Treex Pre- and Post-processing

Here, Treex is used in a more aggressive way, modifying the input and/or output to account for phenomena that the PB-SMT system is known not to be able to handle well (see Figure 8).

[9] `https://translate.google.com/`

[10] The 13th September 2016 version of the Treex repository, `https://github.com/ufal/treex/`

[11] `cu-plain` Moses output downloaded from `http://matrix.statmt.org/systems/show/2807`, test set downloaded from `http://matrix.statmt.org/test_sets/list`.

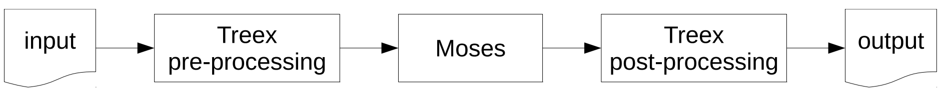

Figure 8: Moses with TectoMT pre- and post-processing

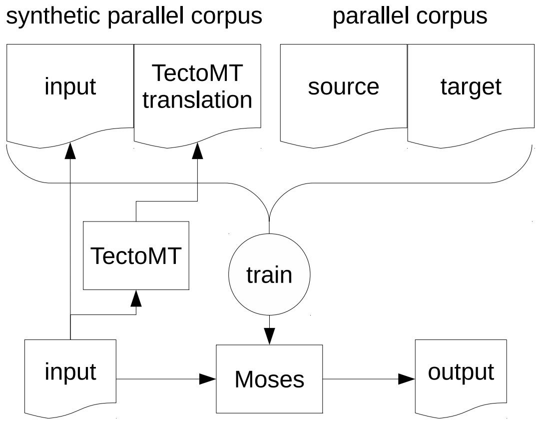

Figure 9: Two-headed Chimera

Zeman (2010) used several pre-processing steps to make the source English text more similar to Czech, such as removing articles, marking subjects by artificial suffixes ("/Sb"), and reordering auxiliary verbs to neighbor their main verbs. Of course, the SMT system was also trained on texts preprocessed in that way; in these experiments, the Joshua PB-SMT system (Li et al., 2009) was used instead of Moses. This approach may seem too aggressive, prone to making the input noisier as well as being potentially lossy. However, the author showed that with careful selection and tuning of the pre-processing steps, a significant improvement of translation quality can be achieved; moreover, this was also confirmed on English-to-Hindi translation.

Rosa et al. (2016) successfully apply Treex pre-processing and post-processing to Moses, but this time with the main objective being an adaptation of Moses trained on general-domain data to a specific domain (namely the domain of Information Technology). The authors use Treex to perform *forced translation* of identified named entities, using a named entity recognizer and a bilingual gazetteer, as well as *forced non-translation* of special structures (URLs, e-mail addresses, computer commands and filenames); Moses XML annotation is used to preserve the forcedly translated items.[12] Apart from domain adaptation, simpler general Treex pre- and post-processing steps were also successfully used, such as projection of letter case in identical words from source to target.

3.7 Two-headed Chimera: Moses with Additional TectoMT Phrase-table

The Two-headed Chimera or AddToTrain (Bojar et al., 2013b; Bojar and Tamchyna, 2015)is a combination of full TectoMT with full Moses (see Figure 9). First, the input is translated by TectoMT. TectoMT translations are then joined with the input to create a small synthetic parallel corpus, from which a secondary phrase table is extracted. This is then used together with the primary phrase table, extracted from the large training data, to train Moses. Finally, the input is translated by the resulting Moses system.

This setup enables Moses to use parts of the TectoMT translation that it considers good, while still having the base large phrase table at its disposal. This has been shown to have a positive effect, e.g., in choosing the correct inflection of a word when the language model encounters an unknown context, or in generating a translation for a word that constitutes an out-of-vocabulary item for Moses (as TectoMT can abstract from word forms to lemmas and beyond, which Moses cannot).

3.8 Chimera: Moses with Additional TectoMT Phrase-table and Depfix Post-editing

The Three-headed Chimera, or simply Chimera (Bojar et al., 2013b; Tamchyna et al., 2016), is a combination of TectoMT and Moses, as in Section 3.7, complemented by a final post-editing step performed by Depfix, as in Section 3.5 (see Figure 10). It has been repeatedly confirmed as the best system by both automatic and manual evaluations, not only among the ones reported in this paper, but also in general,

[12] `http://www.statmt.org/moses/?n=Advanced.Hybrid`

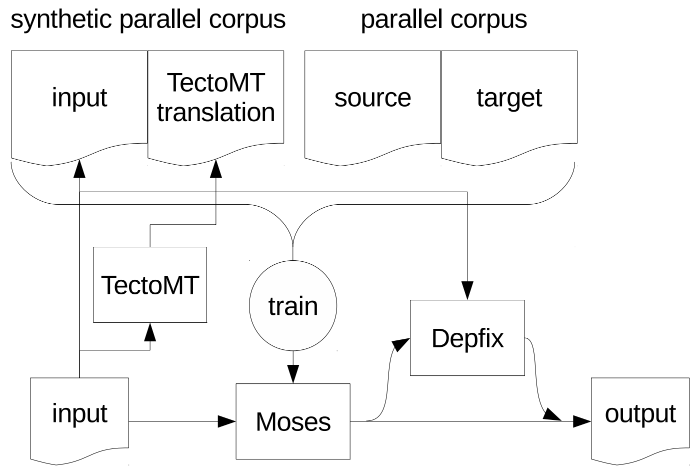

Figure 10: Three-headed Chimera

being the winner of the WMT English-to-Czech translation task in the years 2013, 2014 and 2015 (Bojar et al., 2013a; Bojar et al., 2014; Bojar et al., 2015).

4 Conclusion

We reviewed a range of existing methods of combining the linguistically poor Moses phrase-based machine translation system with linguistically rich systems implemented in the Treex NLP framework, most notably the TectoMT system, including their automatic evaluation via BLEU as reported in the literature. Some of the methods have been shown to achieve significant improvements in the translation quality, as measured by BLEU as well as by human evaluation. The most successful are the Chimera methods, which constituted the state-of-the-art in English-to-Czech machine translation in several WMT translation shared tasks.

On the other hand, many other methods have not brought any significant improvement, or have even lead to a deterioration of the translation quality. However, we believe these methods to be worth considering as well, as they bring more insight into the problematics of hybrid translation. Moreover, some of them might be further modified or combined in future, and may eventually become useful, possibly for a different language pair or for a specific domain.

Acknowledgements

This research was supported by the grants FP7-ICT-2013-10-610516 (QTLeap), GAUK 1572314, GAUK 2058214, and SVV 260 333. This work has been using language resources and tools developed, stored and distributed by the LINDAT/CLARIN project of the Ministry of Education, Youth and Sports of the Czech Republic (project LM2015071).

References

Ondřej Bojar and Kamil Kos. 2010. 2010 failures in English-Czech phrase-based MT. In *Proceedings of the Joint Fifth WMT and MetricsMATR*, pages 60–66, Uppsala, Sweden. Uppsala Universitet, ACL.

Ondřej Bojar and Aleš Tamchyna. 2015. CUNI in WMT15: Chimera strikes again. In *Proceedings of the 10th WMT*, pages 79–83, Stroudsburg, PA, USA. ACL.

Ondřej Bojar and Zdeněk Žabokrtský. 2009. CzEng 0.9, building a large Czech-English automatic parallel treebank. *PBML*, (92):63–83.

Ondřej Bojar, Bushra Jawaid, and Amir Kamran. 2012. Probes in a taxonomy of factored phrase-based models. In *Proceedings of the Seventh WMT*, pages 253–260, Montréal, Canada. ACL.

Ondřej Bojar, Christian Buck, Chris Callison-Burch, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2013a. Findings of the 2013 WMT. In *Proceedings of the Eight WMT*, pages 1–44, Sofija, Bulgaria. Bălgarska akademija na naukite, ACL.

Ondřej Bojar, Rudolf Rosa, and Aleš Tamchyna. 2013b. Chimera – three heads for English-to-Czech translation. In *Proceedings of the Eight WMT*, pages 92–98, Sofija, Bulgaria. Bălgarska akademija na naukite, ACL.

Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. Findings of the 2014 WMT. In *Proceedings of the Ninth WMT*, pages 12–58, Baltimore, MD, USA. ACL.

Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. Findings of the 2015 WMT. In *Proceedings of the 10th WMT*, pages 1–46, Stroudsburg, PA, USA. ACL.

Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Névéol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. Findings of the 2016 conference on machine translation (WMT16). In Ondřej Bojar and et al ., editors, *Proceedings of the First Conference on Machine Translation (WMT). Volume 2: Shared Task Papers*, volume 2, pages 131–198, Stroudsburg, PA, USA. ACL.

Ondřej Dušek, Luís Gomes, Michal Novák, Martin Popel, and Rudolf Rosa. 2015. New language pairs in TectoMT. In *Proceedings of the 10th WMT*, pages 98–104, Stroudsburg, PA, USA. ACL.

Ondřej Dušek, Zdeněk Žabokrtský, Martin Popel, Martin Majliš, Michal Novák, and David Mareček. 2012. Formemes in English-Czech deep syntactic MT. In *Proceedings of the Seventh WMT*, page 267–274.

Petra Galuščáková, Martin Popel, and Ondřej Bojar. 2013. PhraseFix: Statistical post-editing of tectoMT. In *Proceedings of the Eight WMT*, pages 141–147, Sofija, Bulgaria. Bălgarska akademija na naukite, ACL.

Jan Hajič. 1998. Building a Syntactically Annotated Corpus: The Prague Dependency Treebank. In Eva Hajičová, editor, *Issues of Valency and Meaning. Studies in Honor of Jarmila Panevová*, pages 12–19. Prague Karolinum, Charles University Press.

Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *ACL 2007, Proceedings of the 45th Annual Meeting of the ACL Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic, June. ACL.

Zhifei Li, Chris Callison-Burch, Sanjeev Khudanpur, and Wren Thornton. 2009. Decoding in Joshua: Open source, parsing-based machine translation. *PBML*, 91:47–56.

Martin Majliš. 2009. Google Translate + TectoMT. `http://www.slideshare.net/martin.majlis/google-translate-tectomt`, May. Oral talk.

David Mareček, Rudolf Rosa, Petra Galuščáková, and Ondřej Bojar. 2011. Two-step translation with grammatical post-processing. In Chris Callison-Burch, Philipp Koehn, Christof Monz, and Omar Zaidan, editors, *Proceedings of the Sixth WMT*, pages 426–432, Edinburgh, UK. University of Edinburgh, ACL.

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

Kemail Oflazer and Ilknur Durgar El-Kahlout. 2007. Exploring different representational units in English-to-Turkish statistical machine translation. In *Proceedings of the Second WMT*, pages 25–32. ACL.

Martin Popel and Zdeněk Žabokrtský. 2010. TectoMT: Modular NLP framework. In Hrafn Loftsson, Eirikur Rögnvaldsson, and Sigrun Helgadottir, editors, *LNAI, Proceedings of the 7th International Conference on Advances in NLP (IceTAL 2010)*, volume 6233 of *LNCS*, pages 293–304, Berlin / Heidelberg. Iceland Centre for Language Technology (ICLT), Springer.

Martin Popel. 2013. Machine translation zoo: Tree-to-tree transfer and discriminative learning. `https://ufal.mff.cuni.cz/~popel/papers/2013_05_06_zoo.pdf`, May. Oral talk.

Martin Popel. 2015. Machine translation and discriminative models: Tree-to-tree transfer and discriminative learning. `http://ufal.mff.cuni.cz/~popel/papers/2015_03_23_discriminative_tectomt.pdf`, March. Oral talk.

Rudolf Rosa, David Mareček, and Ondřej Dušek. 2012. DEPFIX: A system for automatic correction of Czech MT outputs. In *Proceedings of the Seventh WMT*, pages 362–368, Montréal, Canada. ACL.

Rudolf Rosa, Roman Sudarikov, Michal Novák, Martin Popel, and Ondřej Bojar. 2016. Dictionary-based domain adaptation of MT systems without retraining. In Ondřej Bojar and et al ., editors, *Proceedings of the First Conference on Machine Translation (WMT). Volume 2: Shared Task Papers*, volume 2, pages 449–455, Stroudsburg, PA, USA. ACL.

Rudolf Rosa. 2013. Automatic post-editing of phrase-based machine translation outputs. Master's thesis, Charles University in Prague, Faculty of Mathematics and Physics, Praha, Czechia.

Rudolf Rosa. 2014. Depfix, a tool for automatic rule-based post-editing of SMT. *PBML*, 102:47–56.

Michel Simard, Cyril Goutte, and Pierre Isabelle. 2007. Statistical phrase-based post-editing. In *HLT 2007: The Conference of the North American Chapter of the ACL; Proceedings of the Main Conference*, pages 508–515, Rochester, New York, April. ACL.

Drahomíra Spoustová, Jan Hajič, Jan Votrubec, Pavel Krbec, and Pavel Květoň. 2007. The best of two worlds: Cooperation of statistical and rule-based taggers for Czech. In *Proceedings of the Workshop on Balto-Slavonic NLP 2007*, pages 67–74, Praha, Czechia. Univerzita Karlova v Praze, ACL.

Aleš Tamchyna, Roman Sudarikov, Ondřej Bojar, and Alexander Fraser. 2016. CUNI-LMU submissions in WMT2016: Chimera constrained and beaten. In Ondřej Bojar and et al ., editors, *Proceedings of the First Conference on Machine Translation (WMT). Volume 2: Shared Task Papers*, volume 2, pages 385–390, Stroudsburg, PA, USA. ACL, ACL.

Zdeněk Žabokrtský. 2011. Treex – an open-source framework for natural language processing. In Markéta Lopatková, editor, *Information Technologies – Applications and Theory*, volume 788, pages 7–14, Košice, Slovakia. Univerzita Pavla Jozefa Šafárika v Košiciach.

Zdeněk Žabokrtský, Jan Ptáček, and Petr Pajas. 2008. TectoMT: highly modular MT system with tectogrammatics used as transfer layer. In *Proceedings of the Third WMT*, pages 167–170. ACL.

Daniel Zeman. 2010. Using TectoMT as a preprocessing tool for phrase-based statistical machine translation. In Petr Sojka, Aleš Horák, Ivan Kopeček, and Karel Pala, editors, *Text, Speech and Dialogue*, volume 6231 of *LNCS*, pages 216–223, Berlin / Heidelberg. Masarykova univerzita, Springer.