Comparing different architectures of hybrid Machine Translation systems

Gregor Thurmair]
Linguatec GmbH, Munich
`g.thurmair@linguatec.de`

Abstract

The contribution discusses variants of architectures of hybrid MT systems. The three main types of architectures are: coupling of systems (serial or parallel), architecture adaptations (integrating novel components into SMT or RMT architectures, either by pre/post-editing, or by system core modifications), and genuine hybrid systems, combining components of different paradigms. The interest is to investigate which resources are required for which types of systems, and to which extent the proposals contribute to an overall increase in MT quality.

1   Introduction

In recent years, significant research in Machine Translation has been carried out, mainly in the area of data-driven MT (example-based and statistical (SMT)), as opposed to knowledge-driven approaches (rule-based (RMT), knowledge-based).

Recent evaluations (Callison-Burch et al. 2009) show that

- both types of systems reach comparable translation quality, but:
- the level of output acceptance (in terms of understandability) in the best language directions is at about 50%.

So the state of the art in MT is far from acceptability by human readers, which limits the success of the MT technology significantly.

Error analysis (Chen et al., 2007, Thurmair 2005) shows that the errors made by the different system types are complementary.

- RMT systems have weaknesses in lexical selection in transfer, and lack robustness in case of analysis failures sentences. However they translate more accurately by trying to represent every piece of the input.
- SMT systems are more robust and always produce output. They read more fluent, due to the use of Language Models, and are better in lexical selection. However, they have difficulties to cope with phenomena which require linguistic knowledge, like morphology, syntactic functions, and word order. Also, they lose adequacy due to missing or spurious translations (Vilar et al. 2006).

Systems which try to profit from the respective other approach, and avoid mistakes for which solutions already exist (albeit in another MT paradigm) must therefore be hybrid solutions, combining knowledge-driven and data-driven elements. The purpose of this paper is to discuss different architectures of hybrid systems which have been proposed recently.

The interest is to discuss the way how the hybrid systems overcome the restrictions of their respective paradigms, how they contribute to improve MT quality (in terms of fluency and adequacy), which requirements they have for language resources and performance, and how they can be adapted to new domains.

Chapter 2 will discuss systems which couple different systems (RMT and/or) SMT, either in a serial (PSE) or in a parallel way; the systems themselves are not modified. Chapter 3 considers systems which use either the SMT or the RMT paradigm as basic architecture, and extend it by either knowledge-driven or data-driven components. Chapter 4 presents approaches which have hybrid architectures, and combine components of RMT and SMT systems into novel architectures.

2 Coupling

Coupling means that two or more existing systems are used to produce improved MT output. Coupling can either be done in a serial way, the most researched approach being statistical post-editing (SPE) of a rule-based system. Or it can be done in a parallel way, whereby the best translation is selected/produced from the output of several systems.

2.1 Serial Coupling

Statistical post-editing of RMT (PSE)

Systems offering serial coupling nearly exclusively consist in modifying RMT output by means of a SMT post-editing component. The SMT component uses a ‚bilingual' training component (from RMT output to ‚good' output).

First systems using small domains (Simard et al. 2007, data from Canadian Job Bank) showed that even with relatively small training data of several thousand sentences, significant improvements of the MT output can be achieved, and a RMT+SPE component outperforms pure SMT systems in cases where only limited data are available.

Experiments have continued since then, on a broader data base, resulting in the following picture:

a. Combinations of RMT+SPE systems are highly competitive in MT quality (cf. Schwenk et al., 2009). The output tends to be grammatical, and the main effect of the combination is an increase in lexical selection quality (Dugast et al. 2007), one of the weak points of pure RMT systems.

b. However, care must be taken to avoid the introduction of errors by the SMT postprocessor. Such errors are: the syntactic structure of the output can be confused by the PSE component (Ehara 2007); accuracy drops as some parts of the translation are omitted, and special care needs to be taken to keep e.g. Named Entities in the output (Dugast et al. 2009).

c. To avoid such deteriorations, (Federmann et al. 2009) use a RMT system's syntactic structure, and only try local alternatives (using POS information). This helps for the lexical selection problem of the RMT systems, but less for the parse-failure problem.

PSE type systems require bilingual training data, to be able to align RMT output to good output.

2.2 Parallel Coupling

This coupling employs several MT systems in parallel, and uses some mechanism to select or produce the best output from the result set. Two main paradigms are followed in these approaches:

The first approach identifies the best translations from a list of n-best translations (Hildebrand/Vogel 2008). They search for the best n-grams in all output hypotheses available, and then select the best hypothesis from the candidate list. They report an improvement of 2-3 BLEU compared to the best single system, as the resulting text can integrate sentences from different MT system outputs.

The second approach does not work on whole sentences but on smaller segments (phrases, words). It uses confusion networks, and *generates* an output sentence on the basis of the available MT outputs. In its first variant, a skeleton is selected as a basis, and for each position of the skeleton the best translation alternative is identified and composed to the overall output sentence. Skeletons can be selected on sentence level (cf. Rosti et al. 2007) but also on phrase level (Heafield et al. 2009); the choice of the best skeleton is critical as it determines the structure / word order of the target sentence. Although most MT output stems from SMT systems, RMT output seems to add interesting hypotheses (Leusch et al. 2009), and is sometimes used itself as the skeleton (Chen et al. 2009).

To overcome the risk of skeleton selection, techniques have been applied to build confusion network such as to let *every* hypothesis be the skeleton, and calculate the overall best solution; this has been done in the context of consensus translation (Matusov et al., 2006).

The results of parallel coupling seem to improve BLEU by 2-3 points; however, the 2009 MT workshop results seem to indicate that system combinations can perform as well as the best individual systems but not significantly better (Callison-Burch et al. 2009).

On top, a parallel system approach seems to be difficult to be used in practical applications; mainly for reasons of computational resources, and availability of MT systems. In praxi, at most two systems would be able to run in parallel; and the reduced number of output candidates would lead to a loss in efficiency for the decision process.

3 Architecture Extensions

While coupling means that the architecture of the participating systems is not changed, by extension we mean that the system architecture basically follows the RMT or SMT paradigm but is modified by including resources of the respective other approach. Modifications can occur as pre-editing (i.e. the system data are pre-processed), or core modification (e.g. phrase tables are extended, dictionaries are enlarged etc. by the respective other approach).

3.1 RMT Extensions

Approaches to improve rule-based systems with data-driven procedures focus on two problems:

- **Pre-editing** is tried, both on the dictionary side, by running Term-Extraction tools, and on the grammar side, by automatically extracting grammar rules from corpora.
- Modification of the **system core** is attempted, both by adding probability information to the analysis / parsing process, and by manipulating the transfer selection process.

3.1.1 Pre-Editing

Pre-Editing refers to the preparation of the language resources for RMT. Main language resources, dictionaries and grammar rules, can be set up using data-driven technology.

Learning of dictionary entries

Pre-Editing in rule-based systems means to apply data-driven techniques for terminology extraction from corpora, either on a monolingual basis (to find missing entries in the system's dictionaries), or from bilingual corpora, to find translation candidates, and to load them into the system dictionary.

Such approaches are already in use in RMT systems. The challenges are:

- recognition of multiword terms: Most of the semantically meaningful words are multiword terms (like ‚*nuclear power plant*‘), having an internal linguistic structure.
- linguistic annotation of the recognised terms. Terms must be brought into correct citation form (i.e. lemmatised), and annotated with (POS etc.) Approaches are described in (Dugast et al. 2009, Eisele et al., 2008).

Results reported show that the MT quality improves moderately, depending on the amount of reductions of the out-of-vocabulary words, which in turn depends on the size and coverage of the already existing dictionary. The approach helps to fill dictionary gaps, and to adapt to new domains.

However, in MT systems with already large dictionaries[1] the problem of lexical selection aggravates, as the amount of translations between which to select increases. This problem turns out to be much more difficult to solve than the problem of dictionary gaps.

Learning of rules in RMT

Research on learning grammar rules by data-driven techniques does not seem to have improved MT output quality significantly. The challenge for learning grammar rules seems to be that very many rule candidates are identified, even for small corpora, and that it is difficult to select the low-frequent ‚good‘ rules from noise produced by the extraction technique. Existing RMT already use large grammars covering a lot of specific linguistic phenomena. As with dictionaries, the main problem is less that some structures are not covered but much more that the grammar rules interact and lead to problems of combinatorics and unexpected side-effects which require massive pruning and often led to parse failures.

3.1.2 RMT core system modifications

Modifications of the system core of RMT systems have been tried in several respects. The option to use probabilistic information in parsing has already been implemented in several existing RMT systems.

Transfer

Current hybrid approaches focus more on translation selection in the transfer phase, which is one of the weaknesses of RMT systems, esp. if dictionaries grow. Traditional approaches to RMT transfer selection rely on two techniques:

- Assignment of subject area codes to translations; if a text belongs to a given subject area (which can be automatically detected, cf. Thurmair 2006), the respective translation is activated. However, even in specific domains,

---

[1] The Linguatec ‚Personal Translator‘ has close to 1 mio entries in its German-English dictionary.

general readings of the terms in question are also found, so that this method is not reliable.

- Tests and actions on certain contextual / structural properties (like: presence of direct object, certain prepositions, passive voice etc.), which trigger a specific translation. However, often such conditions cannot be reliably stated for lexical selection, esp. if the number of alternative translation grows; in addition, such tests rely on correct parses of the input sentence which cannot be guaranteed.

Therefore, additional and robust means for lexical selection need to be developed.

An obvious means is to use the more frequently used translation of a given term as default. But this technique is not sensitive to the specific context in which a term must be translated, and mostly returns the default.

A second option is to use contextual disambiguation in the lexical selection process. Relevant clusters of (source language) contextual terms for a given candidate translation are built at training time from a corpus; at runtime these contexts are matched against the context of the text to be translated, and the best translation is selected. This technique, (cf. Thurmair 2006), requires broadening the analysis scope of the system (from sentence-based to paragraph-based contexts); it achieves very good disambiguation results for the terms it was built for. Improvements of accuracy are also reported by (Kim et al., 2002); they use a smaller contextual window and follow a (Probabilistic) Latent Semantic Analysis approach.

As a result, core modifications in RMT can improve the transfer selection process significantly; however they are less successful in case of robustness / parse failures.

3.2 SMT Extensions

Like RMT systems, SMT systems have also been extended to improve translation quality. Again,

- Pre-editing is tried to prepare the data; the most important steps are morphology, POS-information / syntactic information, and word reordering.
- System core modifications are tried as well, by adding RMT information to the phrase tables, and by using factored translation.

3.2.1 Pre-editing

Morphology: Morphology has been researched rather extensively, mainly in languages with rich morphological schemas. Lemmatisation and POS tagging was used both on the source side (e.g. de Gispert et al., 2006) and on the target side (Vandeghinste et al. 2006); the aim is to reduce data sparseness using lemma-based language models instead of textform-based ones. It seems to improve results for smaller corpora. Also, it seems that both textform and lemma based analysis should be done, as surface information has also shown to be beneficial (Koehn/Hoang 2007). Factored translation (cf. below) is able to work on both levels simultaneously.

Another research area in morphology is compounding (of English) / decompounding (of German words), to parallelise alignment (Stymne et al. 2008, Popović et al. 2006). Moreover, in languages with agglutinative behaviour, like Turkish (Hakkani-Tür et al. 2004), Hungarian or Arabic (Habash 2007), preprocessing is required to split complex word strings (including pronouns, case markers etc.) into meaningful parts to be able to align them.

Syntax: Syntactic preprocessing id tried e.g. in (Hannemann et al. 2009); the idea is to parse source and target side of a corpus, and only let syntactically well-formed phrases enter the phrase table. Both corpora are parsed, matching subtrees (mainly on NP level) are identified and aligned in the phrase table. The parsed phrases are still a minority on the phrase table but can help improving the MT output, in particular for local reorderings.

Reordering: Reordering is a major challenge for SMT systems, not just because languages have different word and constituent order (SVO vs. SOV etc.) but also because the constituent order is meaning-bearing (e.g. case marking in English). While standard phrase-based models can handle local reorderings (e.g. noun-adjective position) to some extent, longer distance reordering requires different means.

Proposals have been made to extend the input word sequence into a lattice containing different reorderings of the input words (based e.g. on POS information). Distortion rules can be set up manually or automatically, for contiguous and discontinuous POS sequences (Niehues/Kolss 2009), by matching them on source and target side of the training corpus. The input lattice contains the re-

spective distorted strings, with weights on the probability of the distortion.

Al alternative approach is proposed e.g. in (Bangalore et al. 2007); they do not use position at all, and try a global alignment in a kind of sentence-based bag-of-words strategy. In decoding, they create all possible permutations allowed by the Language Model (in a given window). However, apart from practical problems (window size), as all source language information is missing, results are not too promising; in addition, multiple occurrences of words in the target (‚the') need to be handled.

(Birch et al. 2009) even claim that reordering problems determine the selection of the translation models: Long term reordering is better handled by hierarchical models (Hiero) while for short and medium reorderings, phrase-based models show better results. This remains to be researched.

3.2.2 SMT core system modifications

Three approaches can be found to incorporate RMT resources into an SMT architecture: Extension of the Phrase Table, rule-based control of the Language-Model-based generation, and factored translation.

Importing RMT resources into the phrase table

It was proposed (e.g. by Eisele et al. 2008) to run RMT systems in addition to SMT systems, and enrich the SMT phrase tables by terms and phrases produced by RMT systems. This approach makes use of the knowledge coded in the bilingual dictionaries of the RMT systems.

Results show that the coverage of the system can indeed be increased, esp. in cases of texts from different domains; however, as the SMT decoder runs last, the effect is that the output can be less grammatical than the one of the original RMT.

The proposal reacts on the data sparseness problem of the SMT training; it does not react on the output grammaticality problem.

Improving decoding using target grammars

First rather dramatic improvements had been reported by (Charniak et al. 2003) where the number of grammatical translations was increased in tests by 45%. Other results were less encouraging (e.g. Och et al., 2003) but this may have been due to the selection of an problematic evaluation metric. In recent times, using syntax in decoding is a major topic of research.

Several proposals exist how to learn grammar and transfer rules from bilingual corpora. (Lavoie et al. 2002, Hannemann et al. 2008) identify structural contexts for translation selection from bitexts. Melamed 2004 adapts parsing to allow for multiple input strings (multitrees). Hierarchical translation (Chiang 2007) uses synchronous context-free grammars in decoding: Different grammar and parsing alternatives are given e.g. in (Zollmann /Venugopal 2006, Galley et al. 2004). An open-source toolset for target langauge parsing, Joshua, has recently been presented (Li et al., 2009).

Including syntax into the decoding process, esp. in the context of hierarchical translation, is a promising approach to boost the grammaticality of the MT output.

Factored Translation

While using structural information for decoding attracts increasing interest, Factored Translation (cf. Koehn/Hoang 2007) aims at enriching systems 'bottom-up', by providing more information at word level. It treats words not just as simple textforms but as vectors of features, such features being the lemma, the POS, morphology, and others. The approach decomposes phrase translations into a sequence of mapping steps, with translation steps operating on phrase level, and generation steps on word level. Models are combined in a log-linear fashion.

Several papers (e.g. Stymne et al. 2008) show that phenomena like NP-agreement and compounding can be handled efficiently within a factored translation framework.

As a result, treating words as feature bundles in factored translation, and using structural information for both source-to-target mapping as well as target decoding, allows significant quality improvements for systems combining these factors. They would use both knowledge-driven (dictionaries and grammars) and data-driven (phrase tables, language models) information. However, they rely on the availability of (possibly even linguistically pre-processed) bilingual corpora. This fact may reduce their applicability.

4 Genuine hybrid architectures

Genuine hybrid architectures do not just use add-ons to their system architecture but combine whole system components of the respective approaches into novel systems. They use three basic components: identification of source language ‚chunks‘ (words, phrases or equivalents thereof), transformation of such chunks into the target language by means of a bilingual resource, and generation of a target language sentence.

Several proposals have been made how such systems could look like.

4.1 Rule-based analysis, bilingual dictionary, target language model

Such an approach has been investigated in the METIS projects (Vandeghinste et al. 2006). Analysis is done using available NL tools (lemmatisers, taggers, chunkers); transfer is based on existing dictionaries (consisting basically of lemma and POS in source and target language, including single and multiword terms), and generation uses a language model (based on a tokenized and tagged English corpus (BCN)).

To ensure that the LM based generation produces grammatical sentences, several approaches have been investigated for different languages; e.g.

- in Greek-to-English (Tambouratsis et al., 2005, Markantonatou et al., 2006), a pattern matcher is applied to search for the best matching patterns containing the respective lexical head, and number + POS of modifiers; the selected pattern is then analysed recursively down the structure (sentence level – chunk level – unit level) for the best matching sub-patterns. The best patterns undergoes language.model-based target search.
- in German-to-English, a ‘structural transfer’ type of mapping component is implemented to prepare good LM-based search

Other languages explored in METIS implement other solutions, like bag-of-words.

Evaluation of the technology show that results are similar to basic SMT systems but worse than a complete (rule-based) system like SYSTRAN in all language combinations (Vandeghinste et al. 2008). However this is not surprising comparing the effort invested in the two systems. However, it needs to be seen if the proposed architecture has the potential to produce superior MT quality once the effort is increased.

4.2 Data driven analysis and generation, bilingual dictionary

Instead of rule-based analysis, an alternative data-driven approach has been proposed by (Carbonell et al. 2006). The required resources are: a (full-form) bilingual dictionary, and a n-gram indexed target language corpus. In analysis, an n-gram window is moved over the sentence, and all words in the window are translated using the bilingual dictionary; based on these translations, the target language corpus is searched for the closest n-gram (ideally containing all words of the source, and no additional ones). The result is a lattice of n-gram translations. Of these, the segments with the strongest left and right overlaps, and the highest density of terms, are selected by the decoder.

While this approach also circumvents the problem of the availability of bilingual resources and uses a dictionary as main translation resource, it does not attempt any ‘phrase’ analysis of the input (while METIS uses phrases produced by linguistic chunkers), and any knowledge-based analysis or generation resource. It needs to be seen how grammaticality of the output can be ensured (e.g. proper morphology, word order problems), and how accuracy can be produced, as the technique seems to ignore out-of-dictionary words (like proper names) and to insert spurious translations in the target language n-grams.

5 Domain Adaptation

A special issue to be considered is domain adaptation. All kinds of MT systems must cope with the fact that they will be used not only in the domain for which they had been developed but also for other domains. While RMT systems support adaptation by dictionary import and coding, which in turn can be based on domain corpus collections, the situation is less obvious for SMT-based systems, and a significant drop of quality (up to 10 BLEU) had been observed.

The most promising approach for SMT systems seems to be to use large out-of-domain training data (e.g. Europarl), and with a small in-domain training set, build different resources for both kinds of data. While the phrase tables of the out-of-domain data moderately improve the in-domain

ones (by closing gaps in the translations), the most efficient approach seems to be to run a target Language Model trained only with in-domain data (Koehn/Schroeder 2007).

Experiments have also been made for integrating customer terminology (a bilingual list of terms) into an SMT system (Itagaki/Aikawa 2008). Artificial contexts are created to identify how the phrase tables would translate a given source language term; in a post-processing phase, the phrase table translations of the terms are replaced by the target expressions of the term list. While this seems to be a significant and error-prone effort, options to manipulate the phrase tables directly (a shown above) could be more promising.

6 Conclusion

The selection of the ‚best' architecture for a practical MT system depends on three basic factors:

- the intended use case, e.g. the translation domain(s). A single-domain application with enough bilingual training data is the exception rather than the rule. SMT approaches to low-resource languages are presented e.g. in (Nießen/Ney 2004)
- the translation quality which can be achieved, both for the domain in which the system was trained, and other domains in which the systems are supposed to be used
- the availability of resources and data, both on monolingual and bilingual level.

For the determination of the MT quality, most of the presented systems claim to outperform some baseline system; however, the results are difficult to compare, also due to the fact that the used metrics often are not adequate as some of them do not treat different system types equally (Dugast et al. 2009). In the present context, where several types of systems need to be compared, this is a drawback. Recently, several approaches for sentence-based metrics have been proposed (an overview is given in Callison-Burch et al. 2009); however there is no consolidated picture, and different metrics seem to perform best for different language directions.

Much more relevant, from a practical point of view, is the availability of resources. For many language directions and many domains, sufficient amounts of bilingual data still do not exist, or cannot be accessed. In this case, architectures which rely on monolingual data and use bilingual dictionaries would have to be preferred. So the selection of the best alternative would depend on quality criteria, and on the availability of (training) data.

However, whichever approach is taken, there is still a long way to go before machine translation systems reach acceptable quality.

References

Bangalore,S., Haffner,P., Kanthak,St., 2007: Statistical Machine Translation through Global Lexical Selection and Sentence Reconstruction. Proc. ACL, Prague

Birch,A., Blunsom,Ph., Osborne, M., 2009: A Quantitative Analysis of Reordering Phenomena. Proc 4th Workshop on SMT, Athens

Callison-Burch,Chr., Koehn,Ph., Monz,Ch., Schroeder,J., 2009: Findings of the 2009 Workshop on Statistical Machine Translation. Proc 4th Workshop on SMT, Athens

Carbonell,J., Klein,St., Miller,D., Steinbaum,M., Grassiany,T., Frey,J., 2006: Context-based Machine Translation. Proc. 7th ACL in the Americas, Cambridge, Ms.

Charniak,E., Knight,K., Yamada,K., 2003: Syntax-based Language Models for Statistical Machine Translation. Proc. MT Summit IX, New Orleans

Chen,Yu, Jellinghaus, M., Eisele,A., Yi Zhang, Hunsicker, S., Theison, S., Federmann, Ch., Uszkoreit, H., 2009: Combining Multi-Engine Translations with Moses. Proc 4th Workshop on SMT, Athens

Chiang, D., 2007: Hierarchical Phrase-Based Translation. in: Computational Linguistics, 33,2

de Gispert,A., Gupta,D., Popović,M., Lambert,P., Mariño,J., Federico,M., Ney,H., Banchs,R.,2006: Improving Statistical Word Alignments with Morpho-syntactic Transformations. Proc. 5th FinTAL, Turku

Dugast,L., Senellart,J., Koehn,Ph., 2007: Statistical Post-Editing on SYSTRAN's Rule-Based Translation System. Proc. SMT-WS 2 WTS, Prague

Dugast,L., Senellart, J., Koehn,Ph., 2009: Statistical Post Editing and Dictionary Extraction: Systran/Edinburgh Submissions for ACL-WMT2009

Ehara, T., 2007: Rule Based Machine Translation Combined with Statistical Post Editor for Japanese to English Patent Translation. Proc. MT Summit, Copenhagen, Workshop on Patent Translation

Eisele,A., Federmann,Chr., Uszkoreit,H., Saint-Amand,H., Kay,M., Jellinghaus,M., Hunsicker,S., Herrmann,T., Yu Chen, 2008: Hybrid Machine Translation Architectures within and beyond the EuroMatrix project. Proc. 12th EAMT, Hamburg

Federmann, Chr., Theison, S., Eisele,A., Uszkoreit, H., Yu Chen, Jellinghaus, M., Hunsicker, S., 2009: Translation Combination using Factored Word Substitution. Proc 4th Workshop on SMT, Athens

Galley,M., Graehl,J., Knight,K., Marcu,D., Deneefe,St., Wang;Wie, Thayer,I., 2006: Scalable Inference and Training of Context-Rich Syntactic Translation Models. Proc. COLING Sydney.

Habash, N., 2007: Syntactic Preprocessing for Statistical Machine Translation. Proc. MT-Summit Copenhagen

Hakkani-Tür,D.Z., Oflazer,K., Tür,G., 2002: Statistical Morphological Disambiguation for Agglutinative Languages. in: Compouters and the Humanities 36,4

Hanneman,Gr., Ambati,V., Clark,J.H., Parlikar, A., Lavie, A., 2009: An Improved Statistical Transfer System for French-English Machine Translation. Proc 4th Workshop on SMT, Athens

Hannemann,Gr., Huber,E., Agarwal,A., Ambati,A., Parlikar,A., Peterson,E., Lavie,A., 2008: Statistical Transfer Systems for French-English and German-English Machine Translation. Proc. 3rd Workshop on SMT, Columbus.

Heafield,K., Hanneman,Gr., Lavie,A.2009: Machine Translation System Combination with Flexible Word Ordering. Proc 4th Workshop on SMT, Athens

Hildebrand,A.S., Vogel,St., 2008: Combination of Machine Translation Systems via Hypothesis Selection from Combined N-Best Lists. Proc. AMTA, Hawaii

Hildebrand,A.S., Vogel,St., 2009: CMU System Combination for WMT'09. Proc 4th Workshop on SMT, Athens

Itagaki,M., Aikawa,T., 2008: Post-MT Term Swapper: Supplementing a statistical machine translation system with a user dictionary. Proc. LREC Marrakech

Kim, Yu-Seop, Chang, Jeong-Ho, Zhang, Byoung-Tak, 2002: A Comparative Evaluation of Data-driven Models in Translation Selection of Machine Translation. Proc. COLING Taipei

Koehn,Ph., Hoang,H., 2007: Factored Translation Models. Proc. ACL, Prague

Koehn,Ph., Schroeder,J., 2007: Experiments in Domain Adaptation for Statistical Machine Translation. Proc. 2nd Workshop on SMT, Prague

Lavoie,B., White,M., Korelsky,T., 2002: Learning Domain-Specific Transfer Rules: An Experiment with Korean to English Translation. Proc. COLING, Tapiei

Leusch,Gr., Matusov,E., Ney,H.: The RWTH System Combination System for WMT 2009. Proc 4th Workshop on SMT, Athens

Li,Zhifei, Callison-Burch, Chr., Dyer,Chr., Khudanpur,S., Schwartz,L., Thornton,W., Weese, J., Zaidan, O.,2009: Joshua: An Open Source Toolkit for Parsing-Based Machine Translation. Proc 4th Workshop on SMT, Athens

Markantonatou,St, Sofianopoulos,S., Spilioti,V., Tambouratzis,G., Vassiliou,M., Yannoutsou,O., 2006: Using patterns for Machine Translation (MT). Proc EAMT Oslo

Matusov,E., Ueffing,N., Ney,H., 2006: Computing Consensus Translation from Multiple Machine Translation Systems Using Enhanced Hypotheses Alignment. Proc EACL, Trento

Melamed,I.D., 2004: Statistical Machine Translation by Parsing. Proc. ACL, Barcelona

Niehues,J., Kolss,M., 2009: A POS-Based Model for Long-Range Reorderings in SMT. Proc 4th Workshop on SMT, Athens

Nießen,S., Ney,H., 2004: Statistical Machine Translation with Scarce resources Using Morpho-syntactic Information. in: Comnputational Linguistics 30,2

Och,F., Gildea,D.,Khudanpur,S., et al., 2003: Syntax for Statistical Machine Translation, Final Report. John Hopkins Summer Workshop. www.clsp.jhu.edu/ws03/groups/translate

Popović,M., Stein,D., Ney,H., 2006: Statistical Machine Translation of German Compound Words. Proc. FinTAL, Turku

Rosti,A., Ayan,N., Xiang,B., Matsoukas,S., Schwartz,R., Dorr,B., 2007: Combining Outputs from Multiple Machine Translation Systems. Proc NAACL, Rochester

Simard, M., Ueffing, N., Isabelle, P., Kuhn, R., 2007: Rule-based Translation With Statistical Phrase-based Post-editing. ACL 2007 Second Workshop on Statistical Machine Translation. Prague

Stymne,S., Holmqvist,M., Ahrenberg,L., 2008: Effects of Morphological Analysis in Translation between German and English. Proc. 3rd Workshop on SMT, Columbus.

Tambouratzis,G., Sofianopoulos,S., Spilioti,V., Vassiliou,M., Yannoutsou,O., Markantonatou,St.,Ioannou,N., 2005: Monolingual Corpus-based MT using Chunks. Proc. MT Summit X, Phuket

Thurmair,Gr., 2005: Hybrid Architectures for Machine Translation Systems. in: Language Resources and Evaluation, 39,1

Thurmair,Gr., 2006: Using Corpus Information to Improve MT Quality. in: Proc. Workshop LR4Trans-III, LREC Genoa

Vandeghinste,V., Dirix,P., Schuurman,I., Markantonatou,St., Sofianopoulos,S., Vassiliou,M., Yannoutsou,O., Badia,T., Melero,M., Boleda,G., Carl,M., Schmidt,P., 2008: Evaluation of a Machine Translation System for Low Resource languages: METIS-II. Proc. LREC Marrakech

Vandeghinste,V., Schuurman,I., Carl,M., Markantonatou,St., Badia,T., 2006: METIS-II: Machine Translation for Low Resource Languages. Proc LREC Genoa

Vilar,D., Xu,Jia, D'Haro,F., Ney,H., 2006: Error Analysis of Statistical Machine Translation Output. Proc. LREC 2006, Genoa

Zollmann,A., Venugopal,A., 2006: Syntax Augmented Machine Translation via Chart Parsing. Proc. Workshop on SMT, New York.