

# Teaching and Assessing Empirical Approaches to Machine Translation

Andy Way & Nano Gough  
School of Computing,  
Dublin City University,  
Dublin, Ireland.

Email: {away, ngough}@computing.dcu.ie

## Abstract

Empirical methods in Natural Language Processing (NLP) and Machine Translation (MT) have become mainstream in the research field. Accordingly, it is important that the tools and techniques in these paradigms be taught to potential future researchers and developers in University courses. While many dedicated courses on Statistical NLP can be found, there are few, if any courses on Empirical Approaches to MT. This paper presents the development and assessment of one such course as taught to final year undergraduates taking a degree in NLP.

## 1 Introduction

It is relatively uncontroversial to state that empirical methods in Natural Language Processing (NLP) and Machine Translation (MT) have become mainstream in the research field. Accordingly, it is important that the tools and techniques in these paradigms be taught to potential future researchers and developers in University courses. Many degree courses nowadays contain specific modules on statistical NLP (as it relates to word-sense disambiguation, parsing, generation etc.). While more and more courses on MT address the statistical approaches which are currently in vogue, it is more a case that this is done in passing in a couple of classes, rather than devoting a whole module to empirical approaches to MT. Indeed, in a trawl of the Web, we could find no such specific courses on statistical paradigms in MT. Of course, that is not to say that none exist; merely that none could be found by us.

Naturally, courses on MT need to take into account the students' skills and demands. Kenny & Way (2001) describe how MT is taught in one institution to students of differing backgrounds. A distinction is made in that paper between *users* versus *developers*: while language students and translators can be expected to be able to use translation tools in their careers as translators, students of NLP with a

specialisation in MT might realistically be employed as designers and implementors of such tools in a programming or localisation environment.

This paper presents a course on empirical methods in MT taught to final year undergraduates taking a degree in NLP, focussing mostly on bitext alignment techniques and Example-Based MT (EBMT). These students have a strong background in programming, language skills and good competence levels in formal linguistics and NLP. Accordingly, the course is very practically oriented, and the students are expected by the end of the course to be in a position to develop a toy EBMT system. The presentation of the course in this paper is not intended to be prescriptive; rather, in reporting on the chosen methodology, it may serve as a basic model for others considering teaching such a module to similar students. In addition, by providing details of what worked and—more importantly, what did not—the hope is that others may benefit from lessons learnt.

The rest of the paper is organised as follows: in section 2, we discuss why teaching empirical approaches to MT is becoming ever more important, and provide some documentation of where such material is covered. Section 3 describes the content of the course taught by us. Section 4 describes the method of assessing the students to whom this course was delivered, and reports on the lessons

learned. Finally we conclude, and provide possible improvements and extensions to the course.

## 2 Teaching Empirical Approaches to MT

It is fair to say that empirical approaches to NLP and MT have matured to the extent that they may now be considered mainstream. Indeed, at any major conference on NLP/MT, nowadays one can expect to encounter more papers which favour an empirical approach than those which utilise rules and/or constraints.

If University students are to become familiar with such techniques and tools, dedicated modules have to be designed which address these topics. However, while there are many courses on statistical NLP already in existence, in a trawl of the Web, we could find no courses which solely address empirical approaches to MT. A similar parallel can be drawn with respect to textbooks: there are a number of volumes which specifically address statistical methods in NLP (e.g. Charniak, 1993; Manning & Schütze, 1999), but no textbooks are geared solely towards the concerns of empirical methods in MT. This hole in the market has recently been partially filled by (Carl & Way, 2003), which in turn may see an increase in the number of courses on EBMT, but there is still a need for a book on Statistical MT (SMT) itself (e.g. Brown *et al.*, 1990, 1992; Yamada & Knight, 2001; Soricut *et al.*, 2002).<sup>1</sup>

Nevertheless, while there would appear to be no courses (other than the one described in the next section) dedicated solely to the teaching of empirical methods in MT, SMT and EBMT have become so well established that any contemporary course on MT would be incomplete without at least equipping students with some superficial knowledge of these techniques. Some examples of courses which address these newer empirical approaches include, but are not limited to, the following:

- MSc. in MT, CALL and NLP at UMIST, UK;<sup>2</sup>

<sup>1</sup>While no standard University courses in SMT exist, Kevin Knight has given tutorials and short courses on SMT at various locations, including at the TMI in Kyoto in 1993, as well as at this MT Summit, of course! He also did a whole summer school at John Hopkins University in 1999, when he put together the *Statistical MT Workbook*, available at <http://www.clsp.jhu.edu/ws99/projects/mt/mt-workbook.htm>.

<sup>2</sup><http://www.ccl.umist.ac.uk/teaching/modules/3000/3003.htm>,

- MSc./Ph.D. Program in Language and Information Technologies at CMU, Pittsburgh.<sup>3</sup>

Others address the topic in modules on Empirical NLP, including:

- Programmes in Computer Science, ISI, CA;<sup>4</sup>
- Undergraduate Study in Computer Science at Brown University, Providence, RI;<sup>5</sup>
- Postgraduate programmes in Computer Science at UMIACS, MD.<sup>6</sup>

In addition, some of the newer textbooks on NLP/MT address these and related issues, e.g. Trujillo (1999, Chapter 8) goes into some detail on EBMT and SMT; Bowker (2002, Chapter 5) discusses the related area of Translation Memory (TM) systems; and to a lesser extent, Jurafsky & Martin (2001, Chapter 21) provide some detail on how empirical techniques can be used in MT. Melamed (2001) is geared specifically towards the exploitation of bitexts using empirical methods. However, until the advent of recent books dedicated to empirical methods in MT (e.g. Carl & Way, 2003), instructors in this area have had to rely on original papers and survey articles (e.g. Somers, 1999).

## 3 Course Content

Kenny & Way (2001) contrasts how courses on MT have to be tailored towards different sets of students with different backgrounds, even in the same institution. One of the authors of this paper also teaches a basic introduction to EBMT in two hours to a group of postgraduate students taking a degree in Translation Studies (TS). This is a small component of a module on Translation Technology. Given that TS students are more interested in TM tools, a superficial overview of EBMT, and especially the differences between TM and EBMT, suffices for this group of students. While both EBMT and TM require aligned corpora, for instance, the TS students

<http://www.ccl.umist.ac.uk/teaching/modules/5000/5092.html>

<sup>3</sup><http://www.lti.cs.cmu.edu/Courses/11-731-desc.html>

<sup>4</sup><http://www.isi.edu/natural-language/people/cs562-2003.htm>

<sup>5</sup><http://www.cs.brown.edu/courses/cs241/>

<sup>6</sup><http://benreilly.umiacs.umd.edu/~hwa/cmssc828-02/>

are far more likely to *use* built-in alignment tools such as Trados *WinAlign*, NLP students may be expected to *develop* their own alignment software.

The course presented here is geared specifically to a group of final year undergraduate students taking a degree in NLP. It would, therefore, be an inappropriate model for groups of students with differing backgrounds. The course in Empirical Approaches to MT taught by us consists of 3 hours a week lectures and a 2 hour practical session over a period of 8 weeks. The content of the course is as follows:

- Week 1:
  - Double lecture: Revision class on Perl.
  - Single lecture: Introduction to Course and Statistics-based MT.
- Week 2:
  - Lab: Perl exercises.
  - Double lecture: Corpus-based language and translation models. What corpora to use, and how to get them.
  - Single lecture: Alignment Methods: word-based, character-based etc.
- Week 3:
  - Lab: Sentential-level alignment using:
    - \* relative sentence position;
    - \* relative length of sentence.
  - Double lecture: Probabilistic Word (and Phrase) Models.
  - Single lecture: How to improve alignment with (a simple set of) heuristics (see Lab, Week 4).
- Week 4:
  - Lab: Add in cognates, paragraph markers, punctuation, HTML tags and other anchors to improve alignment.
  - Double lecture: EBMT. How it works, comparison with TM etc.
  - Single lecture: Marker Hypothesis as segmentation tool.

- Week 5:
  - Lab: Build word-alignment tool (i.e. probabilistic lexicon) using mutual information
  - Double lecture: Marker Hypothesis: advantages/disadvantages.
  - Single lecture: Problems for EBMT:
    - \* boundary definition;
    - \* boundary friction;
    - \* what examples to store etc.
- Week 6:
  - Lab: Build and Test EBMT system using 500 sentences of English, French and German data. Use sententially- and word-aligned databases.
  - Double lecture: Generalised Templates (cf. rules in rule-based methods in MT).
  - Single lecture: Generalised Templates.
- Week 7:
  - Lab: Improve model with generalised templates and word insertion.
  - Double lecture: Examples of EBMT systems.
  - Single lecture: Towards Hybrid Models.
- Week 8:
  - Lab: Finalise EBMT System.
  - Double lecture: Statistical MT.
  - Single lecture: Statistical MT.

The basic model followed here is that the material delivered in class during one week is put into practice in the labs during the following week. Essentially, the course is split into three chunks, namely Alignment (both sentential and word-level), EBMT and SMT.

In the introductory part of the course, students are made aware of the need for statistical language and translation models to be developed from large,

good quality, representative monolingual and bilingual corpora. By concocting toy corpora which do not fulfill these criteria, and asking students to calculate a number of unigram and bigram probabilities based on data contained in these corpora, it is quite easy to demonstrate that a number of undesirable effects follow when small, unrepresentative corpora are used. The advantages and disadvantages of bigram models are then presented to the students.

Despite the fact that the various mathematical techniques employed are, in principle at any rate, utilisable for any pair of languages, the fact that sentimentally aligned bilingual corpora exist only for a few language pairs renders these techniques somewhat less generally applicable. In order to try to overcome this problem, some consideration is given to using the Web as a corpus from which usable bitexts might be extracted (cf. Grefenstette, 1999; Resnik & Smith, 2003).

Some of the major algorithms for aligning bilingual corpora are then presented (Brown *et al.*, 1991; Gale & Church, 1993; Kay & Röscheisen, 1993). These are interestingly different, in that the method of Brown *et al.* uses a length-based metric which counts words, that of Gale & Church uses a character-based model, while Kay & Röscheisen require the use of a bilingual dictionary (a ‘Word Alignment Table’) constructed automatically from the bitext.

As for sub-sentential alignments, students are shown how to estimate co-occurrence using Mutual Information. With particular respect to EBMT, other methods of segmentation are also presented, especially Marker-Based segmentation (e.g. Veale & Way, 1997; Way & Gough, 2003). The ‘Marker Hypothesis’ (Green, 1979) is a universal psycholinguistic constraint which states that natural languages are ‘marked’ for complex syntactic structure at surface form by a closed set of specific lexemes and morphemes. Marker-Based MT constructs sets of marker words (e.g. determiners, conjunctions, quantifiers etc.) for source and target languages and segments a sentimentally aligned corpus into marker-headed chunks.  $\langle \text{source}, \text{target} \rangle$  sub-sentential aligned pairs can then be automatically constructed if the number of marker chunks for a source string is equal to the number of marker chunks for the target equivalent, and the chunks are

headed by marker words of the same category.

Students are also made aware of the need for extracted sub-sentential alignments to be made more general, in order to improve coverage and robustness. Some of the techniques presented include:

- Extracting transfer rules from examples (e.g. Furuse & Iida, 1992);
- Generalising by syntactic category (e.g. Kaji *et al.*, 1992);
- Generalising by semantic features (e.g. Matsumoto & Kitamura, 1995);
- Generalising Parse Trees (e.g. Way, 2003);
- Generalising Strings (e.g. Cicekli & Güvenir, 2003; McTait, 2003);
- Generalising using Placeables (e.g. Brown, 1999).

These methods are embedded in a basic outline of the EBMT process. Comparisons are made with TM, some experience of which the students have had before. The major EBMT problems of boundary definition and boundary friction are presented: boundary definition is concerned with the problem that retrieved fragments may not be well-formed constituents, so that syntactic well-formedness may not be ensured in the generation of the target string, while the main problem with boundary friction is that context may not be taken into account in the retrieval process, so that ill-formed output is constructed despite the fact that the derivation of that output was correct according to the knowledge in the system’s databases.

In addition, the storage of examples is discussed, issues pertaining to segmentation are put forward, and the matching and recombination stages of EBMT are explained. We then present the IBM Models 1 and 2 of SMT.<sup>7</sup> Finally, given that these students have previously taken a module on rule-based MT, we discuss possible hybrid models which combine elements of both paradigms as a more effective solution to the problems of translation.

<sup>7</sup>Note that some students choose to do their final year projects in this area, cf. <http://www.redbrick.dcu.ie/~grover/project/index.html> for a particularly good example.

## 4 Assessment

The course was designed to be an ‘assessment only’ module (i.e. no end of module exam), for two reasons. Firstly, being a second semester course, any final examinations would be scheduled during time which students would otherwise be spending on implementing their final year project (which constitutes 33% of their overall degree classification). Secondly, students at our University are classified only on their final year marks, so if students were to fail either the exam or assessment component of any final year module, they would not be eligible for an Honours degree. It was felt, therefore, that by having the module evaluated purely by continuous assessment, these issues could be best avoided.

There were two assignments:

1. a labtest on building an aligner (week 5);
2. a group presentation/demonstration on building an EBMT system (week 8).

The labtest was a 3-hour assessment, in which the students were individually asked to develop a number of programs in Perl, namely:<sup>8</sup>

- to calculate the average sentence length of the  $\langle source, target \rangle$  sentences provided in terms of both words and characters;
- to calculate the ratio of  $\langle source, target \rangle$  words and characters per sentence;
- to write a length-based sentence aligner, in terms of both words and characters;
- to compare the alignment results against a ‘gold standard’ provided;
- to segment the ‘gold standard’ reference solution according to the marker hypothesis;
- to propose sub-sentential alignments using the marker hypothesis.

In addition, there were three discussion questions on aspects of the course.

<sup>8</sup>For a set of freely available such tools, see Dan Melamed’s page at <http://www.cs.nyu.edu/~melamed/software.html>.

Note that all of these programs had been tackled in the lab sessions during the course. We considered three hours to be a reasonable time limit given that one of the authors was able to write programs to perform the various tasks in one hour. Nevertheless, we found that we had overestimated quite considerably the amount we thought the students were capable of in the time available: none of them completed all questions, and in general, too many students spent far too long on programs for which very few marks were awarded (as indicated on the question paper).

While the students’ answers were marked benevolently (for instance, where students provided pseudocode instead of actual Perl code, full marks were given if the pseudocode was a complete solution to the problem at hand), using the original schema, over half of the class had failed, with the top mark being just 57%. However, a compromise was developed whereby the marks were divided by 0.7 in order to give a truer indication of each student’s performance (top mark 81%, lowest 19%, average 53%).

The second assessment was a group project, where the students were divided into groups of three and were asked to develop an EBMT system, based on the sentence-, phrasal- and word-level alignments written in preparation for the first assignment. Marks were awarded both for system design and functionality, and for documentation. No one segmentation method was preferred over any other; indeed, some groups used the marker-based approach, others used a bigram approach, etc. The groups presented their systems to these authors, who found their efforts to be extremely good (highest mark 90%, lowest 57%).<sup>9</sup> Finally, in order to derive the final mark for the module, the first assignment was weighted 0.35, with the group assignment weighted 0.65. All students passed the module (highest mark 87%, lowest 44%, average 65%).

## 5 Conclusion

Empirical approaches to NLP and MT have reached a reasonable stage of maturity. It is important, therefore, that the tools and techniques underpinning these fields be taught to University students, who are

<sup>9</sup>For a good example, consult <http://www.computing.dcu.ie/~sfoy-cl4/ebmt.html>.

likely to form a pool from which future researchers and developers in these areas are to be found. While dedicated courses on statistical NLP have made their way into many University curricula, we were unable to find any courses on empirical methods in MT in a trawl of the Web. Similarly, while a number of textbooks have appeared on statistical NLP, the first such book on one of the flavours of empirical MT, namely EBMT, has only recently appeared.

This paper presents the development and assessment of one such course as taught to final year undergraduates taking a degree in NLP. It focusses mainly on Alignment, EBMT and SMT. It was designed so that student performance was evaluated purely in terms of continuous assessment. We commented on the problems that arose in scheduling a laboratory test of the students' understanding of alignment, both at the sentential and sub-sentential levels. Nonetheless, when asked to develop an EBMT system in small groups, the students rose to the challenge and performed well on this task.

As for improvements/changes to the course, the students may be assessed on a more ongoing basis in weekly labs rather than in one laboratory examination. In addition, the material in (Carl & Way, 2003) may be used in study classes, with students presenting this material to the class on a weekly basis. Finally, we may choose to focus more on the building of SMT systems, using the excellent, freely available Egypt toolkit.<sup>10</sup>

In sum, despite some teething problems, it can be said with some confidence that the module was successful. The developers and teachers of this course have certainly learned from the experience, and it is hoped that others who are considering the development of similar courses may find some value in the sharing of our experiences. That said, nothing in this paper is intended to be prescriptive: if the course structure and methods of assessment are of use to others, then fine, but if some other model is chosen, then we too would hope to benefit from the development of similar or related material.

## References

[1] Lynne Bowker. 2002. *Computer-Aided Translation Tech-*

<sup>10</sup><http://www.clsp.jhu.edu/ws99/projects/mt/>

- nology*, University of Ottawa Press, Ottawa, Canada.
- [2] Peter Brown, John Cocke, Stephen Della Pietra, Fred Jelinek, Vincent Della Pietra, John Lafferty, Robert Mercer and Paul Roossin. 1990. A Statistical Approach to Machine Translation, *Computational Linguistics* **16**:79–85.
- [3] Peter Brown, Jennifer Lai and Robert Mercer. 1991. Aligning sentences in parallel corpora. In *Proceedings of 29th Annual Meeting of the Association for Computational Linguistics*, University of California, Berkeley, pp.169–176.
- [4] Peter Brown, Stephen Della Pietra, Vincent Della Pietra, John Lafferty and Robert Mercer. 1992. Analysis, Statistical Transfer, and Synthesis in Machine Translation. In *4th International Conference on Theoretical and Methodological Issues in Machine Translation*, Montreal, Canada, pp.83–100.
- [5] Ralf Brown. 1999. Adding Linguistic Knowledge to a Lexical Example-based Translation System. In *Proceedings of the 8th International Conference on Theoretical and Methodological Issues in Machine Translation (TMI 99)*, Chester, England, pp.22–32.
- [6] Michael Carl and Andy Way, eds. 2003. *Recent Advances in Example-Based Machine Translation*, Kluwer Academic Publishers, Dordrecht, The Netherlands.
- [7] Eugene Charniak. 1993. *Statistical Language Learning*, MIT Press, Cambridge, MA.
- [8] Ilyas Cicekli and H.Altay Güvenir. 2003. Learning Translation Templates from Bilingual Translation Examples. In *Recent Advances in Example-Based Machine Translation*, M. Carl & A. Way, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, pp.255–286.
- [9] Osamu Furuse and Hitoshi Iida. 1992. An Example-Based Method for Transfer-Driven Machine Translation. In *Fourth International Conference on Theoretical and Methodological Issues in Machine Translation: Empiricist vs. Rationalist Methods in MT (TMI-92)*, Montréal, Canada, pp.139–150.
- [10] William A. Gale and Ken W. Church. 1993. A Program for Aligning Sentences in Bilingual Corpora. *Computational Linguistics* **19**(1):75–102.
- [11] Thomas R.G. Green. 1979. The Necessity of Syntax Markers. Two experiments with artificial languages. *Journal of Verbal Learning and Behavior* **18**:481–496.
- [12] Greg Grefenstette. 1999. The World Wide Web as a Resource for Example-Based Machine Translation Tasks. In *Proceedings of the ASLIB Conference on Translating and the Computer* **21**, London, [pages not numbered].
- [13] Daniel Jurafsky and James H. Martin. 2000. *Speech and Language Processing*, Prentice Hall, Upper Saddle River, NJ.
- [14] Hiroyuki Kaji, Yuuko Kida and Yuji Morimoto. 1992. Learning Translation Templates from Bilingual Text. In *COLING-90, Papers Presented to the 13th International Conference on Computational Linguistics*, Helsinki, Finland, pp.672–678.
- [15] Martin Kay and Martin Röscheisen. 1993. Text-translation alignment. *Computational Linguistics* **19**(1):121–142.

- [16] Dorothy Kenny and Andy Way. 2001. Teaching Machine Translation & Translation Technology: A Contrastive Study. In *Proceedings of the Workshop on Teaching Machine Translation, MT Summit VIII*, M. Forcada, J-A. Pérez-Ortiz & D. Lewis, eds., Santiago de Compostela, Spain, pp.13–17.
- [17] Chris Manning and Hinrich Schütze. 1999. *Foundations of Statistical Natural Language Processing*, MIT Press, Cambridge, MA.
- [18] Yuki Matsumoto and Mihoko Kitamura. 1995. Acquisition of Translation Rules from Parallel Corpora. In *Recent Advances in Natural Language Processing: Selected Papers from the Conference*, R. Mitkov and N. Nicolov, eds., John Benjamins, Amsterdam, pp.405–416.
- [19] Kevin McTait. 2003. Translation Patterns, Linguistic Knowledge and Complexity in EBMT. In *Recent Advances in Example-Based Machine Translation*, M. Carl & A. Way, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, pp.307–338.
- [20] Dan Melamed. 2001. *Empirical Methods for Exploiting Parallel Texts*, MIT Press, Cambridge, MA.
- [21] Philip Resnik and Noah Smith. 2003. The Web as a Parallel Corpus. *Computational Linguistics* **29**(3) (to appear).
- [22] Harold Somers. 1999. Review Article: Example-based Machine Translation. *Machine Translation* **14**(2):113–157 (revised and updated in M. Carl & A. Way, eds. 2003. *Recent Advances in Example-Based Machine Translation*, Kluwer Academic Publishers, Dordrecht, The Netherlands, pp.3–57).
- [23] Radu Soricut, Kevin Knight and Daniel Marcu. 2002. Using a Large Monolingual Corpus to Improve Translation Accuracy. In *Machine Translation: From Research to Real Users, 5th Conference of the Association for Machine Translation in the Americas (AMTA-2002)*, S. Richardson, ed., LNAI 2499, Springer Verlag, Berlin/Heidelberg, Germany, pp.155–164.
- [24] Arturo Trujillo. 1999. *Translation Engines: Techniques for Machine Translation*, Springer, London.
- [25] Tony Veale and Andy Way. 1997. *Gaijin*: A Bootstrapping, Template-Driven Approach to Example-Based Machine Translation. In *International Conference, Recent Advances in Natural Language Processing*, Tzigov Chark, Bulgaria, pp.239–244.
- [26] Andy Way. 2003. Translating with Examples: The LFG-DOT Models of Translation. In *Recent Advances in Example-Based Machine Translation*, M. Carl & A. Way, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, pp.443–472.
- [27] Andy Way and Nano Gough. 2003. *wEBMT*: Developing and Validating an Example-Based Machine Translation System using the World Wide Web. *Computational Linguistics* **29**(3) (to appear).
- [28] Kenji Yamada and Kevin Knight. 2001. A Syntax-Based Statistical Translation Model. In *39th Annual Meeting and 10th Conference of the European Chapter of the Association for Computational Linguistics*, Toulouse, France, pp.523–530.