

English to Irish Machine Translation with Automatic Post-Editing

Meghan Dowling Teresa Lynn Yvette Graham John Judge

ADAPT Centre, Dublin City University, Dublin, Ireland

meghan.dowling@dcu.ie, tlynn@computing.dcu.ie, graham.yvette@gmail.com,
jjudge@computing.dcu.ie

RÉSUMÉ

Cet article présente l'adaptation d'un système de traduction automatique statistique, anglais→irlandais, à un nouveau domaine d'utilisation. Ce système nommé est actuellement utilisé par une équipe de traducteurs du gouvernement irlandais. Nous décrivons également le nouveau module de post-édition automatique qui a été développé pour améliorer le système actuel et faciliter le travail de post-édition des traducteurs.

ABSTRACT

English to Irish Machine Translation with Automatic Post-Editing

This paper reports on the continued development of a domain-tailored English→Irish Statistical Machine Translation system currently in use by an in-house translation team of an Irish government department. We describe the new automatic post-editing module that has been developed to enhance the current system and reduce the post-editing required of translators.

MOTS-CLÉS : Traduction automatique statistique, Post-édition automatique, langue morphologiquement riche, langue irlandaise..

KEYWORDS: statistical machine translation, automatic post-editing, morphologically rich language, Irish language.

1 Introduction

The Irish language holds the status of national and first official language in the Republic of Ireland. This status has led to a government requirement for all official documents and public services to be made accessible in both Irish and English, with the official status of English being a second official language. The demand for Irish-translated content exceeds the productivity capabilities of current translation services in Irish government departments.

In particular, the in-house translation team of the Department of Arts, Heritage and the Gaeltacht (DAHG), the government department responsible for Irish language affairs, has a significant workload and considerable amount of backlog of documents required to be translated into Irish, due to high demand from within their own and across other government departments.

In the past, translators relied solely on translation memory (TM) tools for translation into Irish. While TM goes some way towards speeding up manual translation and increasing productivity, its benefits are of course limited only to working with previously translated text that are similar to or fully

matching source language input text. When MT is not available and source language text has not been previously encountered, this results in translators translating text entirely from scratch. To this end, DAHG has provided funding for development of an English to Irish Statistical Machine Translation (SMT) system to bridge the wide gap between supply and demand of Irish language translations. The resulting system provides translators with the choice of an MT output translation in addition to any matching TM.

The specific requirement of the system was to achieve high-quality translation of domain specific data – that is, the system was required to produce high-quality translations specifically for public administration text. A feasibility study was carried out to determine the appropriate use-case, which amounted mostly to translation of documents, such as annual reports, staff announcements and public notices, for example. The feasibility study ensured the opportunity to ascertain the most appropriate data to train the required SMT system. As there was not a corpus of suitable quality within the required domain readily available, the priority of the project became the collection, cleaning and curation of parallel data. Dowling *et al.* (2015) provide a summary of this corpus development along with a report on preliminary translation scores for an English to Irish Phrase-based SMT system based on Moses (Koehn *et al.*, 2007), often referred to within the Irish-speaking community as the Tapadóir project.

This paper reports on recent enhancements to Tapadóir. In Section 3, we describe the development of an Automated Post-Editing (APE) module, which addresses morphological challenges encountered by the SMT system and results in modest BLEU score improvements. We then report on the evaluation of the APE module in Section 4. Finally, in Section 5, we show the success of the integration of this MT software into the translator’s work-flow by reporting on positive user-engagement with the newly introduced technology.

2 Related Work

There have been various approaches to addressing the problem of translation into morphologically-rich languages. For example, the approach taken by Avramidis & Koehn (2008) involves adding per-word linguistic information to the source language, while Virpioja *et al.* (2007) use unsupervised morphology learning. El Kholy & Habash (2012) report success in this area through the use of a discriminative lexicon model applied to the SMT system. The method suggested by Chahuneau *et al.* (2013) involves a two-tiered approach: building a discriminative model which can predict target-side inflections, and then using this model to generate additional translations which can be included in the standard translation model as “synthetic” phrases. More recently, the Dagstuhl seminar on Statistical Techniques for Translating to Morphologically Rich Languages (Fraser *et al.*, 2014), has brought together researchers from a number of NLP (natural language processing) disciplines to identify new techniques to translating into morphologically rich languages.

Automatic Post-Editing (APE) of MT aims to improve MT output quality in order to reduce post-editing effort required of professional translators (Knight & Chander, 1994). The most widely applied method of APE for MT currently in use is statistical phrase-based post-editing, proposed by Simard *et al.* (2007), where the APE uses the MT output and its corresponding human post-edited data as a parallel corpus. Béchara *et al.* (2011) propose a significant variant that includes the source information along with the MT output on the source side of the parallel corpus. Chatterjee *et al.* (2015) compare these two approaches for English to Spanish MT, the approach of Simard *et al.* (2007) achieving

lower TER scores. Pal *et al.* (2015) apply hybrid word alignment techniques, while Wisniewski *et al.* (2015) take a rule-based approach in addition to Statistical APE. In this paper, we apply a simple rule-based approach to APE for English to Irish MT.

3 Automated Post Editing for Irish

Usability and user experience are extremely important factors in the Tapadóir project. As the primary aim of Tapadóir is to improve the speed and productivity of translators, it is crucial to produce a tool that does not hinder the user in any way. As part of our translator-developer feedback loop, translators reported some repetitive errors in the MT output that were causing frustration. On closer examination, most of the errors were grammatical problems arising from Irish language morphology that Tapadóir was not yet equipped to deal with. In comparison to English, Irish has a richer morphology, such as inflected prepositions and the initial consonant mutations, and causing challenges for SMT due to data sparsity. This problem is compounded in the case of lesser-resourced languages where there are low instances of various inflected forms in the training data.

This gap in knowledge could be bridged through a number of methods such as increasing the volume of training data (where the system becomes familiar with various inflected forms of a word), factored models (where the system uses part-of-speech and lemma information to improve its knowledge) or through the introduction of post-processing module that could address simple grammatical issues on a word level basis.

To this end, we designed an Automated Post Editing (APE) module that could address trivial spelling issues or contraction issues that challenged the SMT system. By automatically post-editing these errors, translators can dedicate more time to more important issues such as language style. The addition of APE is intended to improve the translator user-experience and avoid any negative impact of repetitive grammatical or orthographic errors, thus creating a more enjoyable user experience.

3.1 Designing the APE module

To develop the APE module, our translator-developer feedback loop enabled us to acquire information on frequently occurring errors, and occurrences of mistranslations. On inspection, translations contained a high number of errors related to Irish language prepositions, eclipsis, lenition and contractions. This motivated the development of a set of manually written rules to correct regularly occurring errors in Irish MT output. Rule sets were developed for individual prepositions and contractions and are triggered by the presence of lexical items in MT output. The APE module is split into two parts: one part which deals solely with orthographic rules, and another which addresses errors caused by grammatical case. In total there are 167 hand-written rules, which have been divided into 55 rule groups (according to preposition and error type).

3.1.1 MT Errors related to orthographic rules in Irish

16 of the most common Irish simple prepositions can be inflected to mark pronominal objects (Christian-Brothers, 1962), (Christian-Brothers, 1960), known as prepositional pronouns or pronominal prepositions. For example, it is ungrammatical in Irish for a pronominal object to occur separated

from the preposition (Ó Múrchú, 2013). Such occurrences on occasion arise in the translation output, however, possibly due to a specific phrase being unseen by the MT system and subsequently translating the phrase on the individual word level. An example of an APE rule now implemented in the systems produces correctly inflected forms of these prepositions when the system incorrectly generates word for word translations (see examples 1 and 2).

Examples of rules:

(1) le mé* → liom
'with me'

(2) ag sinn* → againn
'with us'

Irish includes orthographical rules that aid pronunciation and reduce ambiguity from sentences, such as the rule driven by the pronunciation of neighbouring vowels. For example, if a word ending in a vowel is followed by a vowel-initial word, morphophonemic rewrite rules are applied to change the spelling to aid pronunciation (Ó Siadhail, 1989). Examples 3 and 4 show eclipsis and h-prefixing respectively being applied to prevent vowel elision.

(3) Eclipsis

(i + vowel) → (in + vowel)
i Éirinn → in Éirinn
'in Ireland'

(4) h-prefix

(le + vowel) → (le + h+vowel)
le úll → le húll
'with an apple'

3.1.2 MT Errors with Grammatical Case in Irish

The second type of error the APE module is designed to correct arise due to the system's occasional incorrect choice of grammatical case. Modern Irish includes three main grammatical cases: nominative, genitive and vocative. In Irish, nouns are marked with case through various morphological changes such as lenition (e.g. *an buidéal* 'the bottle' → *dath an bhuidéil* 'colour of the bottle'), eclipsis (e.g. *na fir* 'the men' → *foirgneamh na bhfear* 'the men's building'), and slenderisation or broadening of consonants (e.g. *an dochtúir* 'the doctor' → *ainm an dochtúra* 'the doctor's name'). The nominative form is sometimes regarded as the 'common case' (Christian-Brothers (1962), Christian-Brothers (1960)) as it also replaces the dative and accusative cases. While the dative case is not expressly marked in Modern Irish, definite nouns that are objects of prepositions still undergo an inflection process. This morphological change may also vary depending on dialect.

The Irish language has three main dialects – the Ulster dialect, Connacht dialect and Munster dialect. Inflection of definite prepositional objects (in the form of initial mutation) is realised through either

lenition (Ulster dialect) or eclipsis (Connacht and Munster dialects) (Ó Siadhail, 1989). From a spelling standards perspective, the translators in the DAHG follow the standard orthography for Irish (An Caighdeán Oifigiúil (Rannóg an Aistriucháin, 1962)), which means they should be consistent within a document, given their chosen type of initial mutation. This means that, while MT output of a lenited form of prepositional object may in fact be grammatically correct, it often requires correction to ensure consistency. Through observation of the data at hand, we chose to consistently use eclipsis as the default for the APE. If the translator wishes to instead apply lenition in a given document, they have the option to then post-edit the text manually.

In some instances, the nominal prepositional object is directly translated as a unigram (i.e. without taking into context the other elements of the prepositional phrase such as preposition and determiner) resulting in the use of an incorrectly inflected form. This is likely to be the result of the MT system backing off to translate on a unigram basis due to data sparsity in the training data. Example 5¹ shows the editing step required in such cases. Our APE module, removes the need for this correction and ensures consistency by applying rewrite rules to capture the mapping between the two dialectal forms.

- (5) **MT output:** *leis an phróiseas pleanála teanga*
Post-APE output: *leis an bpróiseas pleanála teanga*
‘with the language planning policy’

In example 6, we show two rewrite rules, which inflect definite nouns following the prepositions *as* ‘from’ and *ar* ‘on’ to conform to the official standard spelling.

- (6) **(PREP + DEF. ART + NOUN) → (PREP + DEF. ART + eclipsed NOUN)**
as an baile → *as an mbaile*
‘from the town’

ar an geata → *ar an ngeata*
‘on the gate’

Rule precedence The order in which the APE rules are applied are important. We apply the orthographic rules described in Section 3.1.1 ahead of the grammatical case rules described in Section 3.1.2. Example 2 shows the steps (1 & 2) of the APE module working together on the phrase *faoin gcathaoir* ‘under the chair’.

- (7) **(vowel-final-PREP + DEF.ART + NOUN) → (contracted-PREP/DEF.ART + eclipsed NOUN)**
1. **Contraction**
faoi an cathaoir → *faoin cathaoir*
 2. **Eclipsis**
faoin cathaoir → *faoin gcathaoir*
‘under the chair’

The combination of vowels in ‘*faoi*’ and ‘*an*’ contract to form ‘*faoin*’ (see example 7.1). The presence of *faoin* before an eclipsable consonant in turn triggers an initial mutation (‘*gcathaoir*’ instead

¹Taken from actual system output.

of ‘*cathaoir*’ in example 7.2). Rule precedence is clearly important here so that the orthography component of the APE module is run before the case component, resulting in the output of the first set of rules triggering the need for the second set of rules.

As with any language, there are exceptions to these rules. For example, in some instances, the combination of both rules can produce non-grammatical character strings (e.g. *ngC*, *mbhF*). Therefore, a small number of ‘clean-up’ rules were introduced to prevent the module introducing such errors. See Example 8 for a list of these rules.

- (8)
1. *ngc* → *gc*
 2. *ngC* → *gC*
 3. *mbp* → *bp*
 4. *mbP* → *bP*
 5. *mbhf* → *bhf*
 6. *mbhF* → *bhF*

Currently this post-editing module alters 13% of sentences on average, with 4% of these sentences having both sets of APE rules applied.

4 Evaluation

In this section, we describe experiments carried out to evaluate the addition of our APE module. We summarise the training data used to train and test the MT system. We then highlight the BLEU score changes following the introduction of the APE module. In addition, we discuss our observation that improvements introduced by the APE from a post-editing perspective may not always be reflected in an increase in BLEU scores.

4.1 Experiment Set-up

Training Data Our training data comprises mainly data received from the DAHG. The Tapadóir project represents a specific use case for professional translators working in the Department of Arts, Heritage and Gaeltacht (DAHG). As the system is tailored to their specific translation demands, it is important that the MT output is of a certain domain and register. The type of text generally translated by this team comprises of annual reports, staff notices, public announcements, and so on. To achieve accurate domain-specific translation, we have worked closely with the translation team to ensure that we can retrain the system at regular intervals on text they have translate in the interim. This text is provided to us in the form of translation memory (TMX) files. Such a data format is easily fed back into the MT system as it is well-structured, aligned, and does not require much cleaning or pre-processing. This data set is the most crucial component of the training corpus as it helps to tune the system to the text genre of the DAHG use case. Currently the Tapadóir training set benefits from 42,500 sentence-pairs of DAHG data.

To add to the domain-specific data, we also make use of two additional translation memories, (Digital Corpus of the European Parliament)² and DGT-TM³ (Directorate General for Translation, Translation Memories). Together they provide us with 29,000 sentence-pairs of good quality data of a similar domain.

While parallel data from the DAHG, DCEP and DGT is extremely beneficial to the Tapadóir project, it also requires some support from general-domain data. To achieve this, we used the ILSP web-crawler (Papavassiliou *et al.*, 2013)⁴ to gather parallel English-Irish data from websites. Websites containing public reference material were crawled in order to ensure (i) a high level of quality and (ii) close alignment to our domain as possible. Currently 10,000 sentence-pairs of this parallel data crawl are included in the training set.

In addition to this, we made use of some previously publicly available datasets: Corpas Comhthreomhar Gaeilge-Béarla (CCGB), a bilingual corpus crawled from the web⁵ and ‘Paradocs’, a parallel English-Irish corpus of legal texts⁶. While this data did not reflect our domain accurately enough, it was, however, useful in the language model. CCGB and Paradocs contain 6,000 and 89,000 sentence-pairs respectively.

Test data A random sample of 1,500 sentence pairs received from DAHG were held out from the training set to form the test set. The test set is therefore domain-specific, and representative of the type of texts the system will be used to translate (letters, reports, press releases, etc.).

4.2 APE Results

In Table 1, we present BLEU scores for various data combinations before and after the APE module has been included in the Tapadóir pipeline evaluated on our held-out test set. The results show a modest increase in BLEU across the board when the APE module is applied to correction of errors. The maximum increase in BLEU scores occurs when the system is trained on the translation memory and crawled data combined of +0.1 BLEU. Although the increase is small, we believe the impact on translation quality to be more substantial than is apparent from the BLEU scores alone, as approximately 200 of the 1500 test set translations are changed by the APE. Therefore, a small-scale human evaluation of the sentences was carried out for translations of the best-performing model to investigate the precision of our rule application.

Sentence-Level Analysis To further analyse the performance of the APE, we conducted a sentence-level BLEU analysis, which brought to light several instances where the inclusion of the APE module triggered a decrease in BLEU, even though the sentence was in fact improved from a post-editing perspective. In order to understand this conflict, the nuances of Irish grammar need to be understood first.

For example, where the translation from English included some words in French, and lenition was applied to the French words in the sentence. In Irish, however, foreign words should not be lenited.

²<https://ec.europa.eu/jrc/en/language-technologies/dcep>

³<https://ec.europa.eu/jrc/en/language-technologies/dgt-translation-memory>

⁴Maligna Jassem & Lipski (2008) was used to align segments

⁵<http://borel.slu.edu/corpas/index.html>

⁶<http://gaois.ie/crp/en/>

System Training Data	No APE (BLEU)	APE (BLEU)
TM	42.21	42.28
TM + Crawled	42.24	42.33
TM + Paradocs	42.91	42.96
TM + Paradocs + Crawled	42.79	42.83
TM + (Paradocs)	43.11	43.19
TM + Crawled + (Paradocs)	43.13	43.18
TM + (Crawled)	42.89	42.99

Table 1: BLEU evaluations for the Tapadóir system trained on various combinations of the data available, with and without the APE module. Brackets indicate that the data was used to train the language model, but not the translation model.

For example, *sa* ‘in the’ normally triggers lenition on words beginning with *b, c, d, f, g, m, p, s, t*. However, this rule cannot apply to non-Irish words (e.g. *sa Chôte d’Azur**). This type of incorrect use of lenition results in an error output in the APE.

An additional example occurs when the APE module is applied to the phrase given in Example 9, there is a decrease in BLEU from 25.93 to 25.68, yet the overall grammaticality of the sentence has been improved ⁷. In this example, the reference translation for the phrase ‘with my department’s officials’ is *le mo chuid oifigh* ‘with my own officials’ (*chuid* does not trigger a h-prefix on *oifigh*). However, the MT output is actually more exact than the reference translation: *le oifigh* mo Roinne* ‘with my department’s officials’, although it does still contain a grammatical error *oifigh**. This machine translation, while matching the orthography of the reference translation (thus contributing to a higher BLEU score), is missing a h-prefix that should be triggered by the preposition *le* ‘with’. The APE accurately corrects this error, resulting in an accurate and grammatical translation of the source text and removing the need for post-editing. However, the application of the APE rule lowers the BLEU score because of the increased edit distance from the reference translation. This is a clear example of how the BLEU metric can miss grammatical improvements in translation output. These differing analyses of automated translation are therefore worth considering in the case of MT evaluation.

- (9) *Source*: the Minister said : “I recently met with my department’s **officials**..”
Irish reference: dúirt an tAire: “bhí cruinniú agam le déanaí le mo chuid **oifigh**”
Before APE: dúirt an tAire: “chas mé le déanaí le **oifigh** mo Roinne..”
After APE: dúirt an tAire: “chas mé le déanaí le **hoifigh** mo Roinne..”

BLEU decrease: 25.93 to 25.68

- (10) *Source*: submissions received about the public advisory **process**...
Irish reference: aighneachtaí a fuarthas mar chuid den **bpróiseas** comhairliúcháin phoiblí...
Before APE: aighneachtaí a fuarthas **faoi an próiseas** comhairliúcháin phoiblí...
After APE: aighneachtaí a fuarthas **faoin bpróiseas** comhairliúcháin phoiblí...

BLEU increase: 35.43 to 38.60

⁷The words changed as a result of the APE module are highlighted in bold.

Example 10⁸, taken from MT output, shows the importance of rule precedence (see also example 7). The contraction of *faoi an* to *faoin* is carried out by the first set of rules in the APE module. The presence of the word '*faoin*' then triggers an eclipsis, mutating '*próiseas*' to '*bpróiseas*'. Had the rules been applied in reverse order, this eclipsis would not have been triggered. The sentence-level BLEU score for this translation is increased from 35.43 to 38.60. Similar to example 9, the reference translation and automated output differ somewhat in their translation of 'about the public advisory process' (*mar chuid den bpróiseas* vs *faoin bpróiseas*). Yet, in contrast to Example 9, both of these possible translations of 'about' trigger initial mutation of *próiseas*, and thus the APE results in an increase in the BLEU score.

5 Integration into the User Workbench

The use of technology in translation work-flow has changed considerably over the past two decades. Computer Assisted Translation (CAT) tools such as translation memories (TM) have been widely embraced by the translation community as they help to eliminate repetitive errors and increase consistency in terminology use (García (2006), Heyn (1998)). In more recent years, there has been a drive towards the integration of MT systems into the translator's work-flow. In general, MT does not aim to replace TM, but instead complement it.

When integrating a SMT system into an existing translation work-flow, it is important to consider translator experience or preconceptions of MT as it is widely acknowledged that there is still some resistance amongst the translation community towards using MT (Lingo *et al.*, 2013).

Fortunately, the in-house translation team were open to trying new types of technology and as a result, integration of MT into translators' daily work-flow has been practically seamless. Figure 1 is a screen-shot of the typical DAHG translator's view within SDL Trados Studio 15⁹. Within the workspace, translators are given a choice to post-edit output which has been found in the translation memory or generated by the Tapadóir MT system. The lower section shows the current segment being translated. The upper section (lines 1 and 2) show the sentence translation options for the current segment as presented by the TM (line 1, indicated by a 71% fuzzy match) and the MT system (line 2, indicated by AT (Automated Translation)).

Figure 2 shows the number of words translated by MT as part of the DAHG translators' work-flow during the period April-August 2015. The steady rise from month to month¹⁰ indicates that the translators have responded positively to the inclusion of MT, and are embracing it in as part of their day-to-day workload.¹¹

6 Conclusion

The Irish government Department of Arts Heritage and the Gaeltacht (DAHG) have funded the development of the Tapadóir SMT project to assist their in-house translators meet the growing demand for English to Irish translation.

⁸The sentence was shortened for clarity in this example.

⁹<http://www.translationzone.com/products/trados-studio/>

¹⁰The dip in activity in July is a result of the Irish parliament summer break period.

¹¹The total number of translated words per month is unfortunately unavailable to us at present.

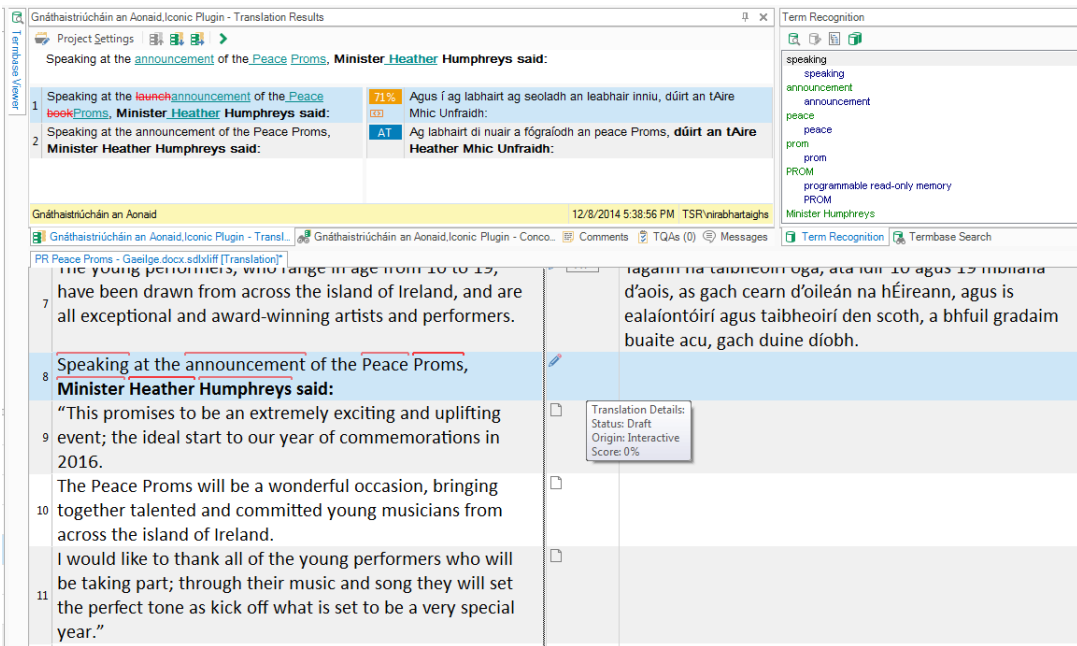


Figure 1: Integration of Tapadóir into SDL Trados Studio 15

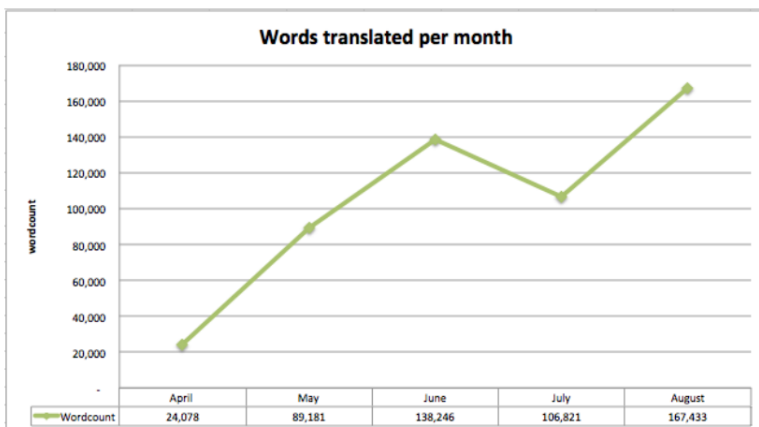


Figure 2: Words translated by Tapadóir in DAHG translation workflow

While we have evaluated the system using traditional MT evaluation metrics such as BLEU (Papineni *et al.*, 2002) in earlier work (Dowling *et al.*, 2015), we show here that we are also focusing on improving the post-editing user-experience as much as possible. We have described in this paper how, through analysis of examples of MT output inaccuracies (provided by DAHG translators) there is still plenty of room for improvement and we plan to embark on further development and improvement of

the system.

We identified grammatical output errors that could easily be addressed by the introduction of an APE module. We also summarised the various nuances of Irish orthography and how to produce the rewrite rules to seamlessly include them in a post-processing step, thus reducing the need for translators to consistently correct simple mistakes.

Thus far the addition of this APE prototype has shown promising results. Therefore the expansion of this module is a natural next step. Future work will also include the adaption of resources such as rules contained in Irish language grammar-checkers (Scannell, 2008) to the domain-specific translation required by the Tapadóir project, as well as the application of factored models (Koehn & Hoang, 2007) to improve translation with respect to Irish language morphology. We also hope to adapt factor templates originally developed for deep-syntax transfer rules (Graham & van Genabith, 2010; Graham & van Genabith, 2008) to factored phrase-based models.

Acknowledgements

This work was part-funded by the Department of Arts, Heritage and the Gaeltacht (DAHG), and by the European Union Horizon 2020 research and innovation programme under grant agreement 645452 (QT21), and is also supported by the ADAPT Centre for Digital Content Technology, which is funded under the SFI Research Centres Programme (Grant 13/RC/2016) and is co-funded by the European Regional Development Fund. We would also like to thank the three anonymous reviewers for their useful comments.

References

- AVRAMIDIS E. & KOEHN P. (2008). Enriching morphologically poor languages for statistical machine translation. In *Proceedings of the Annual Conference of the Association for Computational Linguistics*, p. 763–770.
- BÉCHARA H., MA Y. & VAN GENABITH J. (2011). Statistical post-editing for a statistical mt system. In *Proceedings of Machine Translation Summit Conference*, volume 13, p. 308–315.
- CHAHUNEAU V., SCHLINGER E., SMITH N. A. & DYER C. (2013). Translating into morphologically rich languages with synthetic phrases. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, p. 1677–1687: Association for Computational Linguistics.
- CHATTERJEE R., TURCHI M. & NEGRI M. (2015). The fbk participation in the wmt15 automatic post-editing shared task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, p. 210–215, Lisbon, Portugal: Association for Computational Linguistics.
- CHRISTIAN-BROTHERS (1960). *Graiméar Gaeilge na mBráithre Críostaí*. Dublin: M.H. Mac an Ghoill agus a Mhac, Tta.
- CHRISTIAN-BROTHERS (1962). *New Irish Grammar*. Dublin: C J Fallon.

DOWLING M., CASSIDY L., MAGUIRE E., LYNN T., SRIVASTAVA A. & JUDGE J. (2015). Tapadóir: Developing a statistical machine translation engine and associated resources for irish. In *The 4th LRL Workshop: Language Technologies in support of Less-Resourced Languages*.

EL KHOLY A. & HABASH N. (2012). Orthographic and morphological processing for english–arabic statistical machine translation. *Machine Translation*, **26**(1-2), 25–45.

FRASER A. M., KNIGHT K., KOEHN P., SCHMID H. & USZKOREIT H. (2014). Statistical Techniques for Translating to Morphologically Rich Languages (Dagstuhl Seminar 14061). *Dagstuhl Reports*, **4**(2), 1–16.

GARCÍA I. (2006). Translators on translation memories: a blessing or a curse? *Translation Technology and its Teaching*, p. 97–105.

GRAHAM Y. & VAN GENABITH J. (2008). Packed rules for automatic transfer-rule induction.

GRAHAM Y. & VAN GENABITH J. (2010). Factor Templates for Factored Machine Translation Models. In *Proceedings of the seventh International Workshop on Spoken Language Translation*, p. 275–282.

HEYN M. (1998). Translation memories: Insights and prospects. In L. BOWKER, M. CRONIN, D. KENNY & J. PEARSON, Eds., *Unity in Diversity. Current Trends in Translation Studies*, p. 123–136: St Jerome Publishing.

JASSEM K. & LIPSKI J. (2008). A new tool for the bilingual text aligning at the sentence level. *Intelligent Information Systems*, p. 279–286.

KNIGHT K. & CHANDER I. (1994). Automated post-editing of documents. In *Proceedings of 12th National Conference on Artificial Intelligence*, p. 779–784.

KOEHN P. & HOANG H. (2007). Factored translation models. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, p. 868–876.

KOEHN P., HOANG H., BIRCH A., CALLISON-BURCH C., FEDERICO M., BERTOLDI N., COWAN B., SHEN W., MORAN C., ZENS R., DYER C., BOJAR O., CONSTANTIN A. & HERBST E. (2007). Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions, ACL '07*, p. 177–180, Stroudsburg, PA, USA: Association for Computational Linguistics.

LINGO A. W., GREATER G. & UK M. (2013). Traditional and emerging use-cases for machine translation.

Ó MÚRCHÚ P. (2013). A grammar of modern irish: An annotated guide to graiméar gaeilge na mbráithre críostaí.

Ó SIADHAIL M. (1989). *Modern Irish: Grammatical structure and dialectal variation*. Cambridge University Press.

PAL S., VELA M., NASKAR S. K. & VAN GENABITH J. (2015). Usaar-sape: An english–spanish statistical automatic post-editing system. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, p. 216–221, Lisbon, Portugal: Association for Computational Linguistics.

PAPAVASSILIOU V., PROKOPIDIS P. & THURMAIR G. (2013). A modular open-source focused crawler for mining monolingual and bilingual corpora from the web. In *Proceedings of the Sixth Workshop on Building and Using Comparable Corpora*, p. 43–51, Sofia, Bulgaria: Association for Computational Linguistics.

PAPINENI K., ROUKOS S., WARD T. & ZHU W.-J. (2002). Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, p. 311–318: Association for Computational Linguistics.

RANNÓG AN AISTRIUCHÁIN (1962). *Gramadach na Gaeilge agus Litriú na Gaeilge: An Caighdeán Oifigiúil*. Oifig an tSoláthair.

SCANNELL K. P. (2008). An gramadóir: A grammar-checking framework for the celtic languages and its applications. In *14th annual NAACL conference*.

SIMARD M., GOUTTE C. & ISABELLE P. (2007). Statistical phrase-based post-editing. p. 508–515.

VIRPIOJA S., VÄYRYNEN J. J., CREUTZ M. & SADENIEMI M. (2007). Morphology-aware statistical machine translation based on morphs induced in an unsupervised manner. In *Proceedings of Machine Translation Summit XI*, p. 491–498.

WISNIEWSKI G., PÉCHEUX N. & YVON F. (2015). Why predicting post-edition is so hard? failure analysis of limsi submission to the ape shared task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, p. 222–227, Lisbon, Portugal: Association for Computational Linguistics.