TwittIrish: A Universal Dependencies Treebank of Tweets in Modern Irish

Anonymous ACL submission

Abstract

Modern Irish is a minority language lacking sufficient linguistic resources for the task of accurate automatic syntactic parsing of user-generated content. As with other languages, the linguistic style observed in Irish tweets differs, in terms of orthography, lexicon and syntax, to that of standard texts more commonly used in Natural Language Processing (NLP) for the development of language models and parsers. This paper reports on the development of TwittIrish, the first Irish Universal Dependencies Twitter Treebank. We describe our bootstrapping method, and report on preliminary parsing experiments.

1 Introduction

User-generated content (UGC) is a valuable resource for training syntactic parsers which can accurately process social media text. UGC is a domain with features different from those of both spoken language and standardised written language more traditionally used in NLP corpora. Given that the accuracy of syntactic parsing tools has been shown to decline when evaluated on noisy UGC data (Foster et al., 2011; Seddah et al., 2012) and that domain adaptation has been shown to improve parser performance in the case of dependency annotation of English tweets (Kong et al., 2014), the need for domain-specific resources is clear in order to reliably process this variety of data.

Universal Dependencies (UD) (Nivre et al., 2020) provides a cross-lingually consistent framework for dependency-based syntactic parsing. UGC, especially social media text, has recently become a popular focus within UD and NLP research more broadly (Silveira et al., 2014; Luotolahti et al., 2015; Albogamy and Ramsay, 2017; Wang et al., 2017; Zeldes, 2017; Bhat et al., 2018; Blodgett et al., 2018; Van Der Goot and van Noord, 2018; Cignarella et al., 2019; Seddah et al., 2020) and has encouraged active conversation around how best to represent it within this framework among the UD community.

This paper reports on the creation of the first Irish Twitter treebank, TwittIrish. Irish is a minority language mostly spoken in small communities in Ireland called 'Gaeltachtaí' (CSO, 2016) but social media sites, such as Twitter, provide a platform for Irish speakers to communicate electronically from any location. Users may reach a wide audience quickly, unconstrained by editors and publishers. These and other socio-linguistic factors contribute to the noncanonical language observed in this domain, which we analyse through the lens of orthographic, lexical and syntactic variation. In order to maintain optimum consistency with other UD treebanks, the annotation methodology employed in this research closely follows the general UD guidelines and the language-specific guidelines for Irish while aiming to incorporate the most up-to-date recommendations (Sanguinetti et al., 2020) for UGC in this evolving area of NLP.

We carry out preliminary parsing experiments with TwittIrish, investigating the following two questions: How effective is a parser trained on the Irish Universal Dependencies Treebank (Lynn and Foster, 2016), which contains only edited text and no UGC, when applied to tweets? And what difference do pre-trained contextualised word embeddings make? We observe a difference of approximately 23 LAS points between TwittIrish and the IUDT test set and find that the use of monolingual BERT embeddings (Barry et al., 2021) improves performance by over 10 LAS points.

The paper is structured as follows: Section 2 details the existing Irish NLP resources we use for our research, Section 3 outlines the development of the treebank, Section 4 describes the characteristics of UGC evident in Irish tweets, and Section 5 presents parsing experiments and error analysis.

2 Irish NLP Resources

We use the following resources:

Indigenous Tweets (IT)[1] This project compiles statistics on social media data of 185 minority and indigenous languages including Irish. All tweets in the TwittIrish treebank were sourced via IT.

Lynn Twitter Corpus (LTC)[2] **(Lynn et al., 2015)** A corpus of 1,493 lemmatised and POS-tagged Irish language tweets randomly sampled from 950k tweets by 8k users posted between 2006 and 2014, identified by IT. The LTC data also contains code-switching information (Lynn and Scannell, 2019).

Irish Universal Dependencies Treebank (IUDT)[3] **(Lynn and Foster, 2016)** A UD treebank consisting of 4910 sentences sampled from a balanced mixed-domain corpus for Irish.

gaBERT (Barry et al., 2021) A monolingual Irish BERT model, trained on approximately 7.9 million sentences, which outperforms Multilingual BERT (mBERT) (Devlin et al., 2019) and WikiBERT (Pyysalo et al., 2021) at the task of dependency parsing for Irish.

3 TwittIrish Development

We combined 700 POS-tagged tweets from the LTC with 166 tweets more recently crawled by IT in order to leverage previous linguistic annotations while also including newer tweets. This involved converting the LTC annotation scheme to that of the UD framework and then POS-tagging the new raw tweets. We provide further detail in Appendix A.

3.1 LTC Conversion

With regard to tokenisation, multiword expressions were automatically split into separate tokens following UD conventions. Only minor manual adjustments were required for lemmatisation to ensure alignment with the IUDT (to enable bootstrapping – see Section 3.3). Finally, the POS tagset used in the LTC was automatically converted to the UD tagset. Appendix A.2 describes this process.

[1] http://indigenoustweets.com/

[2] https://github.com/tlynn747/IrishTwitterPOS

[3] https://github.com/UniversalDependencies/UD_Irish-IDT

3.2 Preprocessing of Newly-crawled Tweets

Due to the lack of a tokeniser designed to deal specifically with UGC in Irish, we compared two tools for this task: UDPipe (Straka et al., 2016),[4] a language-agnostic trainable pipeline for tokenisation, tagging, lemmatisation and dependency parsing, and Tweettokenizer[5] from NLTK (Bird, 2006), a rule-based tokeniser designed for noisy UGC. The latter proved to be more effective for tokenising UGC phenomena such as emoticons, URLs and meta-language tags. Manual corrections were then applied in order to adhere to the Irish-specific tokenisation scheme within current UD guidelines. In order to establish the best system to use for automatic lemmatising and POS-tagging, two tools, Morfette (Chrupala et al., 2008) and UDPipe (Straka et al., 2016), were analysed with Morfette achieving higher scores on both tasks.

3.3 Syntactic Annotation

As a method shown to reduce manual annotation efforts in syntactic annotation (Judge et al., 2006; Seraji et al., 2012), we carry out a bootstrapping approach to dependency parsing as recommended by UD [6].

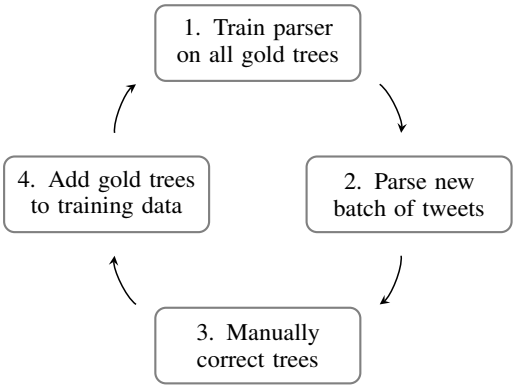

Figure 1: Bootstrapping approach to semi-automated syntax annotation.

The bootstrapping process is illustrated in Figure 1. After converting the LTC and new tweets to the CoNLL-U format, we manually annotated a small set of 166 tweets and began the bootstrapping cycle.[7] (Step 1) A parsing model[8] was then

[4] Trained on IUDT v2.8 with no pre-trained embeddings.

[5] https://www.nltk.org/api/nltk.tokenize.html

[6] https://universaldependencies.org/how_to_start.html

[7] All manual annotation and correction was performed by one linguist annotator.

[8] Biaffine Parser (Dozat and Manning, 2017) with mBERT (Devlin et al., 2019) embeddings.

trained on the IUDT in combination with the newly annotated tweets. (Step 2) The parsing model was used to automatically annotate the next batch of 100 tweets. (Step 3) These tweets were then manually corrected. (Step 4) The corrected tweets were then added to the training data. Steps 1 to 4 were repeated until all 866 tweets were fully parsed. This dataset represents the TwittIrish test set in the UD version 2.8 release.

4 Annotating Irish UGC

This section describes the linguistic features that can create challenges when parsing Irish social media text. We provide Irish language examples and discussion around the factors that influence these phenomena.

4.1 Orthographic Variation

Orthographic variation refers to deviation from the conventional spelling system of the language and is observed at the token level. Therefore, it can affect the lemmatisation of a token in an NLP pipeline, potentially affecting other downstream areas of annotation. In the TwittIrish dataset, 2.5% of tokens contained some orthographic variation. Table 1 exemplifies some frequently-occurring phenomena in Irish tweets which deviate from standard orthography.

Diacritic variation Diacritic marks are often omitted or incorrectly added to Twitter text. The acute accent or *síneadh fada* is used in Irish to indicate a long vowel and is necessary to disambiguate between certain word pairs. Example 1 shows the most probable intended word *léacht* 'lecture' rendered as *leacht* 'liquid'.

(1) ***Leacht*** *faoi stair Príosún Dún Dealgain*
'Lecture about the history of Dundalk Prison'

Abbreviation Predictable shorthand forms can occur in standard Irish texts e.g. *lch* as an abbreviated form of *leathanach* 'page'. While more unconventional, and thus less predictable, abbreviations are observed in Irish tweets, as per Example 2 in which the word *seachtain* 'week' is shortened to *seacht* 'seven'. Abbreviations are common in tweets than standard text as the character limit and real-time, up-to-date nature of the platform encourages the user to be efficient with time and space.

(2) *Bím de ghnáth ach sa bhaile an* ***tseacht*** *seo*
'I usually am but home this week'

Lengthening This refers to the elongation of a token by repeating one or more characters. This can be thought of as an encoding of sociophonetic information (Tatman, 2015) and has been shown to be strongly linked to sentiment. Despite incentives to save time and space while tweeting, users often elongate certain words for expressive purposes (Brody and Diakopoulos, 2011). Example 3 demonstrates the lengthening of the word *buí* 'yellow'.

(3) *tá siad go léir* ***buuuuuuí***
'They are all yellooooooow'

Case Variation Nonstandard use of upper- and lowercase text is another method of encoding sociophonetic information by focusing attention or emotion on a particular word or phrase. Heath (2021) discusses the association between the use of all-caps and perceived shouting as in Example 4.

(4) *Níl todhchaí na Gaeilge sa Ghaeltacht, ach in aon áit* ***AR DOMHAIN***
'The future of Irish is not in the Gaeltacht but anywhere ON EARTH'

Punctuation Variation Punctuation is used creatively in UGC to format or emphasise strings of text. However, due to the lack of standardisation, occurrences of unconventional punctuation can make text difficult to parse for both human and machine, as in Example 5 which shows a phrase from an Irish tweet appended by two punctuation characters '-)'. It is unclear whether this should be interpreted as some form of punctuation, creative formatting or a smiley e.g. ':-)'.

(5) *sin a dhóthain****-)***
'That's enough-)'

Other Spelling Variation These are mostly slight variations very close to the intended word and may occur due to typographical error. Typos are very common in UGC due to lack of editing or proof-reading and may occur via insertion, deletion, substitution or transposition of characters.

(6) tus **staith** 6 de Imeall
'start of season 6 of Imeall'

Example 6 shows *sraith* (season) rendered as **staith*. Due to their phonetic dissimilarity and the fact that 't' and 'r' are adjacent on the QWERTY keyboard layout, it is reasonable to infer that the substitution was unintentional. Less commonly, disguise or censorship of words or phrases may occur to encrypt profanity or taboo language of some variety.

Phenomenon	Example	Standard form	Gloss
Diacritic Variation	***nior*** *fhoghlaim* ***tu***	*níor fhoghlaim tú*	'you did not learn'
Abbreviation	***fhoir*** *rugbaí na* ***hÉir***	*fhoireann rugbaí na hÉireann*	'Irish rugby team'
Lengthening	***obairrrr***	*obair*	'work'
Case Variation	*ceolchoirm* ***DEN SCOTH***	*ceolchoirm den scoth*	'excellent concert'
Punctuation Variation	***folúntas***	*folúntas*	'vacancy'
Other Spelling Variation	*O 'Bama*	*Obama*	'Obama'

Table 1: Examples of orthographic variation in Irish tweets.

4.2 Lexical Variation

Just 38.32% of the set of unique lemmata that make up the vocabulary of the TwittIrish treebank occur in the IUDT training data. Table 2 shows examples of lexical variation in Irish tweets.

Dialectal Vocabulary Irish has three major dialects; Connaught, Munster and Ulster. Distinctive features of these dialects in the form of lexical variation are evident in spoken language and informal text such as tweets. Example 7 shows the use of *domh*, the Ulster variant of *dom* meaning 'to me'.

(7) *Ba chóir* ***domh*** *rá!*
'I should say!'

Initialism Multiword phrases are frequently represented by the initial letter of each of their constituent tokens. Example 8 shows *GRMA* 'Thank you' used to represent its expanded form *Go raibh maith agat*.

(8) *Scaip an scéal!* ***GRMA****!*
'Spread the word! Thank you!'

Pictogram Emojis, emoticons etc. can be added to text to emulate gesture (Gawne and McCulloch, 2019) or they may play a syntactic role in a phrase, replacing a word as in Example 9, in which the symbol, ❤, acts as the object of a verb. Pictograms tend not to have one-to-one correspondence with natural language words.

(9) *Conas a deireann tú* ❤*?*
'How do you say ❤'

Truncation Due to the current limit of 280 characters per tweet, the end of the text may be truncated, sometimes mid-sentence or mid-word. Example 10 features an utterance in which the end has been unnaturally attenuated.

(10) *Súil agam go bheas sé mar sin don…*
'I hope it will be like that for the…'

Transliteration The practice of transliteration, in which a word in one language is written using the writing system of another, is common within the language pair of Irish and English. In the TwittIrish treebank, the English language phrase 'fair play' occurs twice while variations 'fair plé', as shown in Example 11 and 'féar plé' occur once each.

(11) ***Fair plé*** *daoibh* 💗'
'Fair play to you 💗'

Insertional Code-switching As with transliteration, code-switching occurs as a result of the high levels of contact in the Irish-English language pair. 66.74% of tokens in the TwittIrish treebank are in Irish, 4.85% of tokens are in English and the remainder are classified as neither, or indeed both in the case of intraword code-switching in which the morphologies of two languages are combined in a single word. Lynn and Scannell (2019) note the propensity of Irish speakers to conjugate an English language verb with the Irish gerund suffix *áil*.

(12) *Eachtra i ndiaidh* ***Happenáil*** *i nGaoth Dobhair*
'An event after happening in Gweedore'

Example 12 of intraword code-switching in which an Irish utterance uses the English verb root 'happen' instead of the Irish equivalent *tarlaigh*.

Single word code-switches and the use of loan words[9] are a distinctive feature of informal Irish. 74.71% of the tweets in the TwittIrish treebank were considered to be entirely in Irish, the remaining 25.29% of tweets being considered bi- or multilingual. Example 13 shows a section of an Irish tweet utilizing the English word 'Dubs' a nickname for 'Dubliners'.

(13) *Roimh na* ***Dubs***
'Before the Dubs'

[9] Hickey (2009) and Stenson (2011) comment on the fuzzy boundary between single word switches and loan words in Irish. The TwittIrish corpus does not provide sufficient evidence to distinguish meaningfully between these terms.

Phenomenon	Example	Standard form	Gloss
Dialectal Vocabulary	**anso**	*anseo*	'here'
Initialism	**BÁC**	*Baile Átha Cliath*	'Dublin'
Pictogram	***<3*** *mór*	*Grá mór*	'Lots of love'
Truncation	*thart fa' 53* **nó...**	*thart fa' 53 nóiméad*	'over 53 mi...(minutes)'
Transliteration	***raight***	*ceart*	'right'
Insertional Code-switching	*sa* **town** *amárach*	*sa bhaile amárach*	'in town tomorrow'
Other Nonstandard Lexical Forms	**TochaltÓr**	*Tochaltóir óir*	'Gold-digger'

Table 2: Examples of lexical variation in Irish tweets.

Other Nonstandard Lexical forms Other unfamiliar terms may occur in the form of hypercorrection and neologsims. Hypercorrection occurs when an autocorrection system is either not activated or available in a user's language of choice. As a result, their attempts to type a word are corrected to a word with a similar spelling in another language.

(14) *Mhúscail mé i mo leaba féin ar maidin i ndiaidh* ***concise*** *mór*
'I woke up in my own bed after a big fortnight'

Example 14 shows the Irish word *coicíse* rendered as 'concise' probably due to automatic English spelling correction software. It is often difficult to distinguish between hypercorrection, neologisms, typos or other spelling variations.

(15) *tá an teanga ag fáil bháis* ***agua*** *níl ach uaireanta*
'the language is dying and there are only hours'

Example 15 shows *agus* (and) rendered as *agua* which may have occurred due to hypercorrection as 'agua' (water) is a frequent token in other languages such as Portuguese and Spanish. However it could also be a simple typo.

4.3 Syntactic Variation

Grammatical phenomena observed in Irish tweets are described in this section. As these idiosyncrasies occur at the phrasal rather than token level, their effect is observed on the structure of the parse trees. Table 3 exemplifies syntactic variation in Irish tweets.

Contraction Much like abbreviation at the token level, contraction is defined here as the fusion of several tokens for the purpose of brevity sometimes mimicking spoken pronunciation. Figure 2 shows the phrase *tá a fhios agam* (lit. its knowledge is at me) reduced to *tá's agam*.

Over-splitting The inclusion of extra white space within tokens is often observed in Irish tweets e.g. *Nílim* ***ró chinnte***. The prefix *ró-* ('too') is conventionally fused with the adjective it precedes in standardised text and so such tokens are annotated with the `goeswith` label as shown in Figure 3.

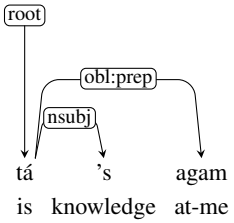

Figure 2: Contraction 'I know'.

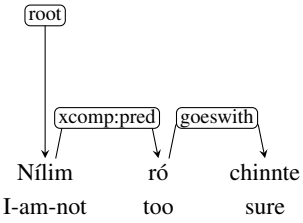

Figure 3: Over-splitting 'I am not too sure'.

Alternational Code-switching This refers to code-switching which alters the structure of the syntax tree, due to the differing word orders of the languages involved, thus complicating the task of dependency parsing. For example, in Irish the adjectival modifier usually *follows* the noun it modifies whereas the inverse is true for English. Figure 4 shows how the English adjective 'hippy-dippy' is incorporated in an Irish language utterance.

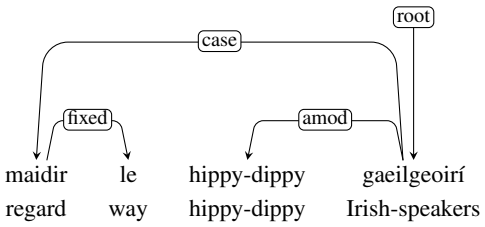

Figure 4: Alternational code-switching 'as for hippy-dippy Irish speakers'.

Dialectal Grammar Figures 5 and 6 show semantically equivalent statements rendered using the synthetic and analytic verb forms respectively. The synthetic form is more common in the Munster dialect of Irish.

Ellipsis It is also common for tweets to consist of incomplete sentences. Example 16 shows such a

Phenomenon	Example	Standard form	Gloss
Contraction	*go* ***dtí'n***	*go dtí an*	'until the'
Over-splitting	***ana shuimiúil***	*an-suimiúil*	'very interesting'
Alternational Code-switching	*Tá an* ***tweet machine*** *ró-tapa*	*Tá inneall na tvuíte ró-tapa*	'The tweet machine is too fast'
Dialectal Grammar	*Ní* ***fhacthas***	*ní fhaca mé*	'I did not see'
Ellipsis	***jab iontach déanta aige***	*tá jab iontach déanta aige*	'he has done a wonderful job'
Meta-language Tags	***#sonas***	*sonas*	'happiness'
Non-sentential Segmentation	***haha:) tá súil agam go raibh sé ann***	*ha ha! Tá súil agam go raibh sé ann.*	'haha:) I hope he was there.'
Other Grammatical Variation	***ce ata an athair?***	*cé hé an t-athair?*	'who is the father?'

Table 3: Examples of syntactic variation in Irish tweets.

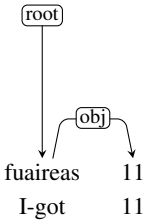

Figure 5: Synthetic verb form 'I got 11'.

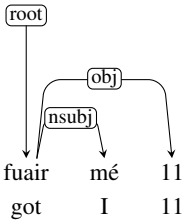

Figure 6: Analytic verb form 'I got 11'.

sentence fragment lacking a main verb. The probable inferred full phrase is *tá báisteach anseo* 'rain is here'.

(16) *báisteach anseo*
'rain here'

Meta-language Tags Hashtags are used in tweets to render the topic searchable and at-mentions are used to address or refer to another user. Both can play syntactic roles in the sentence as shown in Figure 7.

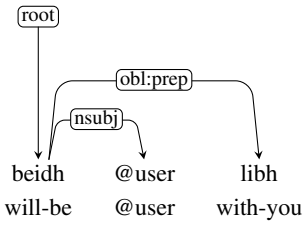

Figure 7: Syntactic meta-language tag '@user will be with you'.

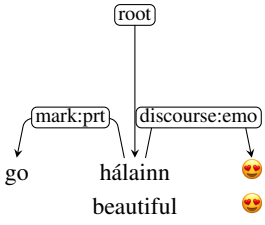

Figure 8: Non-sentential tweet using emoji as punctuation 'beautiful 😍'.

Non-sentential Structure In tweets, the sentence is not an appropriate unit of segmentation as frequently non-standard punctuation, or none at all, is used. Figure 8 exemplifies a tweet utilising an emoji instead of punctuation.

Other Grammatical Variation Grammatical variation can also occur via unintentional deviation from conventional spelling or grammar by an L2 Irish speaker. Example 17 shows a grammatically incorrect phrase roughly translating to 'I have to *going'. In such cases, though the annotator may be able to infer the intended phrase *Caithfidh mé dul* 'I have to go', no corrections are made by the annotator to the surface form, however this information can be represented in the annotation via the label `CorrectForm` as described in (Sanguinetti et al., 2020).

(17) *Caithfidh mé* ***ag*** *dul*
'I have to *going'

Additionally, Irish tweets contain extremely unconventional constructions. This can occur in the form of unnatural phrases that have been machine translated or generated by bots.

(18) *Conas a Faigh tonna de Morgáiste*
'*How to get a tonne of mortgage'

Example 18 shows an ungrammatical construction that appears to have been translated automatically word by word. A more natural construction might be *conas tonna morgáiste a fháil* 'How to get a tonne of mortgage'. Some examples of this variety are easy to identify from surrounding context such as links to websites with similar content however, tweets may consist of text alone making it difficult to infer whether the author is human or machine.

5 Parsing Experiments

We compare the performance of two widely used neural dependency parsers on the TwittIrish test set, and examine the effect of using pre-trained contextualised word embeddings from a monolingual Irish BERT model (gaBERT). We report parsing performance by UPOS and dependency label and carry out a manual error analysis. Further information is detailed in Appendix B.

5.1 Parser Comparison

We experiment with two neural dependency parsing architectures: UDPipe (Straka et al., 2016), an NLP

System	LAS IUDT	LAS TwittIrish
UDPipe v1	70.58	47.33
AllenNLP	71.56	48.73
AllenNLP + gaBERT	**84.25**	**59.34**

Table 4: Comparison of parsing systems UDPipe v1, Biaffine Dependency Parser (Dozat and Manning, 2017): Biaffine dependency parser with BiLSTM encoder, and AllenNLP + gaBERT: Biaffine dependency parser where BiLSTM is replaced with pretrained BERT model of Barry et al. (2021). All were trained on IUDT version 2.8 and tested on the IUDT and TwittIrish test sets.

pipeline which includes a transition-based non-projective parser, and AllenNLP (Gardner et al., 2018), a biaffine dependency parser with a BiLSTM encoder (Dozat and Manning, 2017). Both systems are trained on IUDT version 2.8 [10] and tested on the IUDT and TwittIrish test sets for comparison. Gold standard tokenization is provided to the models which then predict UPOS tags and dependency relations and labels. As the TwittIrish test set is the only gold annotated treebank of Irish UGC, no UGC is used as training or development data in these experiments. We opt to preserve it as a test set so that our results and results of future research in this area will be comparable.

In order to leverage the substantial advances in accuracy achieved in dependency parsing by the use of pre-trained contexualized word representations (Che et al., 2018; Kondratyuk and Straka, 2019; Kulmizev et al., 2019), we use AllenNLP with token representations obtained from the last hidden layer of the gaBERT model (Barry et al., 2021) and passed to the biaffine parser.

Table 4 shows that, when tested on the IUDT version 2.8 test set, UDPipe achieves 70.58% labelled attachment score (LAS). In comparison, UDPipe achieves a much lower LAS score of 47.33 on the TwittIrish test set. Similarly to UDPipe, AllenNLP achieves 71.56 LAS on the IUDT test set with a similar decrease of 22.83 points on the TwittIrish test set. The highest accuracy of 84.25 LAS is achieved by gaBERT with a difference of 24.91 points when tested on the TwittIrish test set. The lower accuracy obtained by parsers on the TwittIrish test is unsurprising given the linguistic differences between the training and test sets. The 10+ LAS improvement provided by the gaBERT embeddings is seen in both test sets.

[10]Models were trained with and without XPOS and feature annotation. The results shown here are without XPOS and features. The addition of XPOS and features constituted a difference of approximately +/-1 LAS.

5.2 Analysis

Analysis was carried out on the AllenNLP parser with gaBERT embeddings using Dependable (Choi et al., 2015).

LAS by Number of Tokens per Sentence/Tweet The mean sentence length of the IUDT is 23.5 tokens, whereas the mean tweet length in TwittIrish is 17.8. Figure 9 shows that, when tested

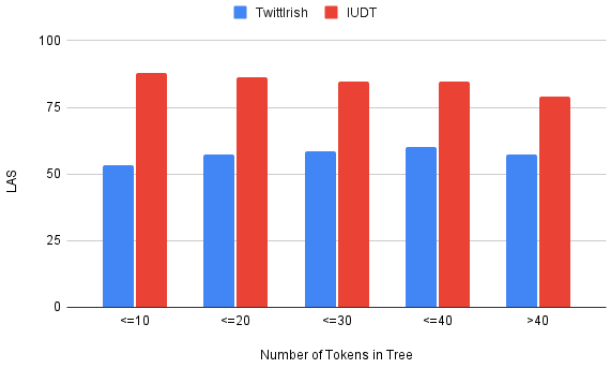

Figure 9: LAS by number of tokens per tweet achieved by AllenNLP Parser with gaBERT embeddings on the IUDT and TwittIrish test sets.

on the IUDT, parsing accuracy decreases as the length of the sentence increases. The highest accuracy of 87.92 LAS is associated with sentences of 10 tokens or fewer and the lowest accuracy is observed in sentences of 40 tokens or more. This is an unsurprising trend as a higher number of tokens increases the probability of longer dependency distances and more complex constructions within a sentence. While the range of scores is smaller and trend less pronounced, the opposite effect is observed when the same parser is tested on TwittIrish, whereby LAS tends to increase as the length of the tweet increases. The highest LAS is associated with tweets of 31 to 40 tokens in length and the lowest accuracy is associated with tweets of 10 tokens or less. This trend is also observed when gaBERT representations are not used, suggesting that, in this case, deep contextualised word embeddings do not cause this effect as observed in (Kulmizev et al., 2019). From manual inspection of the data, we observe that the genre-specific phenomena which challenge the parser such as ellipsis, meta-language tags and URLs, occur in higher proportions in shorter tweets therefore causing this

trend.

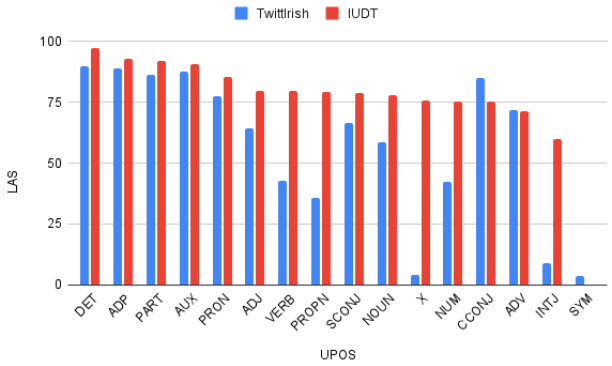

Figure 10: LAS by UPOS tag achieved by AllenNLP Parser with gaBERT embeddings on the IUDT and TwittIrish test sets.

LAS by UPOS Figure 10 shows LAS associated with each UPOS tag when tested on the IUDT and TwittIrish. LAS is higher when tested on the IUDT for all UPOS tags except `CCONJ`, `ADV` and `SYM` and in these cases the difference is small (<10 LAS). The most notable differences are `X` (71.6 LAS), `INTJ` (51.3 LAS), `PROPN` (43.5 LAS). These differences are due to 1) the divergent genres of the treebanks e.g. in the TwittIrish treebank the UPOS tag `X` is used for all non-syntactic hashtags, and `PROPN` is used for all at-mentions, neither of which occur in the IUDT and 2) differing annotation conventions e.g. in the IUDT, the tag `X` is used mostly for foreign-language tokens. Whereas, in TwittIrish, due to the high proportion of English language tokens, non-Irish words are annotated with their true UPOS tag where the language is known to the annotator.

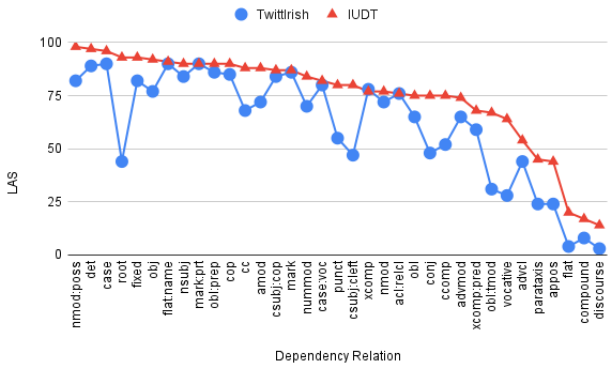

Figure 11: LAS achieved by AllenNLP Parser with gaBERT embeddings on the IUDT and TwittIrish test sets by dependency relation.

LAS by Dependency Relation Figure 11 shows parsing accuracy broken down by dependency relation. The parser obtains higher scores on the IUDT for all dependency relations except `xcomp` for which it is just one point higher when tested on TwittIrish. The largest differences between the accuracy of the two test sets are associated with the labels `root`, `vocative`, `obl:tmod`, `csubj:cleft`, `conj`, and `punct`.

Error Analysis In order to assess the effect of the UGC phenomena present in Irish tweets, we analyse the most and least accurate parses. 7 tweets (76 tokens) were parsed with LAS between 0 and 5. There were 15 occurrences of emojis which were most commonly incorrectly labelled `punct`. The 10 occurrences of code-switching were most commonly attached via `flat:foreign`. The 9 (2 syntactic) occurrences of usernames were most commonly labelled as `root`. There were 5 occurrences of ellipsis in the form of verb omission obfuscating the task of root selection. The 3 hashtags were most commonly mislabelled as `nmod` as were the 3 URLs. 1 occurrence of spelling variation was observed in the form of diacritic omission wherein the word *ár* 'our' was rendered as *ar* 'on' causing the parser to misinterpret the dependency label of `nmod:poss` as `case`. 7 tweets (89 tokens) were parsed with an accuracy between 95 and 100 LAS. All of these were grammatical, well-formed sentences. There were 3 usernames and 1 hashtag all of which were syntactically integrated and so they were parsed correctly. There was one occurrence of insertional single-word code-switching which was accurately parsed. There were 2 occurrences of spelling variation, both in the form of diacritic omission but, as these do not resemble any other words, they were parsed correctly.

6 Conclusion

Presented in this paper is the novel resource, TwittIrish, the first Universal Dependencies treebank for Irish UGC. Analysis of this linguistic genre and anonymised examples of Irish tweets are presented. This research facilitates the development of NLP tools such as dependency parsers for Irish by providing a test set on which future Irish language technology can be tested. Future work will involve both further annotation and exploration of semi-supervised techniques.

References

Fahad Albogamy and Allan Ramsay. 2017. Universal dependencies for Arabic tweets. In *RANLP*, pages 46–51.

James Barry, Joachim Wagner, Lauren Cassidy, Alan Cowap, Teresa Lynn, Abigail Walsh, Mícheál J Ó Meachair, and Jennifer Foster. 2021. gabert–an irish language model. *arXiv preprint arXiv:2107.12930*.

Irshad Bhat, Riyaz A. Bhat, Manish Shrivastava, and Dipti Sharma. 2018. Universal Dependency parsing for Hindi-English code-switching. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 987–998, New Orleans, Louisiana. Association for Computational Linguistics.

Steven Bird. 2006. Nltk: the natural language toolkit. In *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, pages 69–72.

Su Lin Blodgett, Johnny Wei, and Brendan O'Connor. 2018. Twitter Universal Dependency parsing for African-American and mainstream American English. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1415–1425, Melbourne, Australia. Association for Computational Linguistics.

Samuel Brody and Nicholas Diakopoulos. 2011. Cooooooooooooooollllllllllll!!!!!!!!!!!!!! using word lengthening to detect sentiment in microblogs. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 562–570, Edinburgh, Scotland, UK. Association for Computational Linguistics.

Wanxiang Che, Yijia Liu, Yuxuan Wang, Bo Zheng, and Ting Liu. 2018. Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 55–64, Brussels, Belgium. Association for Computational Linguistics.

Jinho D. Choi, Joel Tetreault, and Amanda Stent. 2015. It depends: Dependency parser comparison using a web-based evaluation tool. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 387–396, Beijing, China. Association for Computational Linguistics.

Grzegorz Chrupala, Georgiana Dinu, and Josef van Genabith. 2008. Learning morphology with Morfette. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).

Alessandra Teresa Cignarella, Cristina Bosco, and Paolo Rosso. 2019. Presenting TWITTIRO-UD: An Italian Twitter treebank in universal dependencies. In *Proceedings of the Fifth International Conference on Dependency Linguistics (Depling, SyntaxFest 2019)*, pages 190–197.

CSO. 2016. Census of Population 2016 – Profile 10 Education, Skills and the Irish Language. Publisher: Central Statistics Office.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Jennifer Foster, Ozlem Cetinoglu, Joachim Wagner, Joseph Le Roux, Stephen Hogan, Joakim Nivre, Deirdre Hogan, and Josef Van Genabith. 2011. # hardtoparse: Pos tagging and parsing the twitterverse. In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*.

Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.

Lauren Gawne and Gretchen McCulloch. 2019. Emoji as digital gestures. *Language@Internet*, 17(2).

Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-speech tagging for Twitter: Annotation, features, and experiments. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 42–47, Portland, Oregon, USA. Association for Computational Linguistics.

Maria Heath. 2021. NO NEED TO YELL: A Prosodic Analysis of Writing in All Caps. *University of Pennsylvania Working Papers in Linguistics*, 27(1).

Tina Hickey. 2009. Code-switching and borrowing in Irish. *Journal of Sociolinguistics*, 13(5):670–688.

John Judge, Aoife Cahill, and Josef Van Genabith. 2006. Questionbank: Creating a corpus of parse-annotated questions. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 497–504.

Dan Kondratyuk and Milan Straka. 2019. 75 languages, 1 model: Parsing Universal Dependencies universally. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2779–2795, Hong Kong, China. Association for Computational Linguistics.

Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archna Bhatia, Chris Dyer, and Noah A. Smith. 2014. A Dependency Parser for Tweets. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1001–1012, Doha, Qatar. Association for Computational Linguistics.

Artur Kulmizev, Miryam de Lhoneux, Johannes Gontrum, Elena Fano, and Joakim Nivre. 2019. Deep contextualized word embeddings in transition-based and graph-based dependency parsing - a tale of two parsers revisited. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2755–2768, Hong Kong, China. Association for Computational Linguistics.

Juhani Luotolahti, Jenna Kanerva, Veronika Laippala, Sampo Pyysalo, and Filip Ginter. 2015. Towards Universal Web Parsebanks. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 211–220. Uppsala University, Uppsala, Sweden.

Teresa Lynn and Jennifer Foster. 2016. Universal Dependencies for Irish. In *In Proceedings of the 2nd Celtic Language Technology Workshop*, Paris, France.

Teresa Lynn and Kevin Scannell. 2019. Code-switching in Irish tweets: A preliminary analysis. In *In Proceedings of the 3rd Celtic Language Technology Workshop*.

Teresa Lynn, Kevin Scannell, and Eimear Maguire. 2015. Minority Language Twitter: Part-of-Speech Tagging and Analysis of Irish Tweets. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 1–8, Beijing, China. Association for Computational Linguistics.

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. Universal Dependencies v2: An evergrowing multilingual treebank collection. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.

Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. A universal part-of-speech tagset. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2089–2096, Istanbul, Turkey. European Language Resources Association (ELRA).

Sampo Pyysalo, Jenna Kanerva, Antti Virtanen, and Filip Ginter. 2021. WikiBERT models: Deep transfer learning for many languages. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 1–10, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.

Manuela Sanguinetti, Cristina Bosco, Lauren Cassidy, Özlem Çetinoğlu, Alessandra Teresa Cignarella, Teresa Lynn, Ines Rehbein, Josef Ruppenhofer, Djamé Seddah, and Amir Zeldes. 2020. Treebanking user-generated content: A proposal for a unified representation in Universal Dependencies. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5240–5250, Marseille, France. European Language Resources Association.

Djamé Seddah, Farah Essaidi, Amal Fethi, Matthieu Futeral, Benjamin Muller, Pedro Javier Ortiz Suárez, Benoît Sagot, and Abhishek Srivastava. 2020. Building a user-generated content North-African Arabizi treebank: Tackling hell. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1139–1150, Online. Association for Computational Linguistics.

Djamé Seddah, Benoit Sagot, Marie Candito, Virginie Mouilleron, and Vanessa Combet. 2012. The french social media bank: a treebank of noisy user generated content. In *COLING 2012-24th International Conference on Computational Linguistics*.

Mojgan Seraji, Beáta Megyesi, and Joakim Nivre. 2012. Bootstrapping a persian dependency treebank. *Linguistic Issues in Language Technology*, 7(18).

Natalia Silveira, Timothy Dozat, Marie Catherine De Marneffe, Samuel R. Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A Gold Standard Dependency Corpus for English. In *Proceedings of the 9th International Conference on Language Resources and Evaluation, LREC 2014*, pages 2897–2904. ELRA.

Nancy Stenson. 2011. *21. Code-switching vs. borrowing in Modern Irish*, pages 559–580. Max Niemeyer Verlag.

Nancy Stenson. 2019. *Modern Irish: A Comprehensive Grammar*. Routledge Comprehensive Grammars. Taylor & Francis.

Milan Straka, Jan Hajic, and Jana Straková. 2016. UDPipe: trainable pipeline for processing CoNLL-U

files performing tokenization, morphological analysis, pos tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 4290–4297.

Rachael Tatman. 2015. #go awn: Sociophonetic Variation in Variant Spellings on Twitter. *Working Papers of the Linguistics Circle*, 25(2):97–108. Number: 2.

Rob Van Der Goot and Gertjan van Noord. 2018. Modeling Input Uncertainty in Neural Network Dependency Parsing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4984–4991.

Hongmin Wang, Yue Zhang, Guang Yong Leonard Chan, Jie Yang, and Hai Leong Chieu. 2017. Universal Dependencies Parsing for Colloquial Singaporean English. In *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, volume 1, pages 1732–1744.

Amir Zeldes. 2017. The GUM Corpus: Creating Multilayer Resources in the Classroom. *Language Resources and Evaluation*, 51(3):581–612.

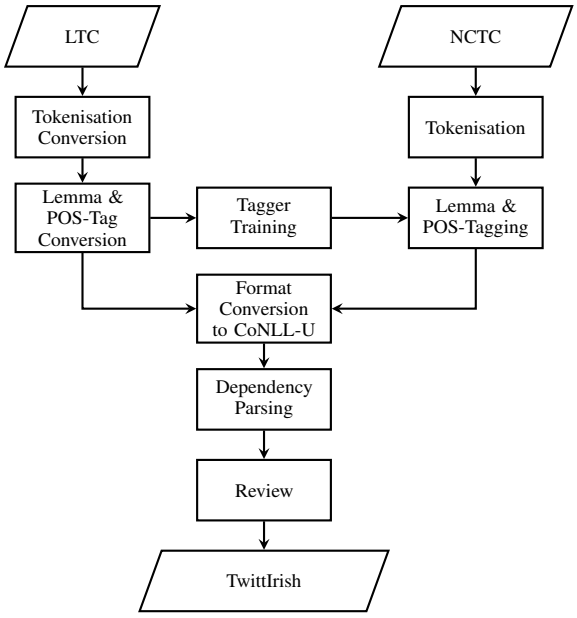

Figure 12: Diagram of the TwittIrish Creation Process. Where NCTC refers to Newly-crawled tweets.

A TwittIrish Development

A.1 LTC Tokenisation Conversion

The most notable difference in the tokenisation approach of LTC as compared to that of UD, was in the treatment of multi-word expressions (MWEs). In LTC, the individual tokens of MWEs are fused with an underscore whereas words with spaces are not allowed in UD [11]. Several minor differences were also observed between the two tokenisation schemes such as whether or not certain symbols, abbreviations or punctuation marks should be attached to the token they follow or considered as a separate token. e.g. 5%, ama..., 1-0, 10pm. UD tends to favour the approach of separating such combinations[12] therefore we resolved to manually separate such occurrences in the TwittIrish tokenisation scheme.

A.2 LTC POS-tag Conversion

Table 5 shows the mapping of LPOS to UPOS. LPOS tags were automatically converted to the corresponding UPOS tag where a one-to-one or many-to-one mapping existed. In the case of one-to-many relationships, automatic identification and manual correction was performed. [13]

[11] https://universaldependencies.org/v2/mwe.html

[12] however not all treebanks apply this consistently.

[13] Both the Gimpel et al. (2011) and UD tagsets derived from the Google Universal POS tagset (Petrov et al., 2012)

LTC POS	UPOS
N, VN	NOUN *
^, @	PROPN *
O	PRON
V	VERB, AUX †
A	ADJ
R	ADV
D	DET
P	ADP
T	PART
,	PUNCT
&	CCONJ, SCONJ †
$	NUM
!	INTJ
U, ~, E	SYM *
#, #MWE	X *
EN	any †
G	any †

Table 5: POS tag Mapping
* Many-to-one relation
† One-to-many relation

Surface	LPOS	UPOS
@user	@	PROPN
#cutie	#	X
ca	R	ADV
bhfuil	V	VERB
an	D	DET
ghra	N	NOUN
you	EN	PRON
ask	EN	VERB
@user #cutie ca bhfuil an ghra you ask [14]		
'@user #cutie where is the love you ask'		

Table 6: Example Irish tweet with LTC and corresponding universal POS tags.

Table 6 demonstrates the mapping of a sample tweet from one scheme to the other. As all English language tokens were annotated with a single tag 'EN' in the LPOS scheme, these tags were converted to the appropriate UPOS tag in the TwittIrish treebank.

Table 7 shows that, using the LPOS tagset, all verbs are tagged `V`. According to UD, the Irish copula (e.g. *is*, *ní*) is tagged as `AUX` distinguishing it from other verbs (e.g. *tá*, *níl*) which are tagged `VERB`.

A.3 Preprocessing of Newly-crawled Tweets

Table 8 shows that hashtags and emoticons were not correctly handled by the UDPipe tokenizer trained on the IUDT. Despite being trained on Irish data, due to differences in domain, Twitter-specific features such as meta-language tags are not present in its training data.

Surface	LPOS	UPOS
Ní	V	AUX
duine	N	NOUN
cáilúil	A	ADJ
é	O	PRON
ach	&	CCONJ
táim	V	VERB
bródúil	A	ADJ
#Grá	#	X
Ní duine cáiliúil é ach táim bródúil #Grá		
'He is not a celebrity but I'm proud #Love '		

Table 7: Example Irish tweet with LTC and corresponding universal POS tags.

UDPipe (IUDT)	NLTK Tweettokenizer
Dé	Dé
Céadaoin	Céadaoin
# Midweek	#Midweek
# Beagnachann	#Beagnachann
:)	:)
:)	:)
Dé Céadaoin #Midweek #Beagnachann :) :)	
'Wednesday #Midweek #Almostthere :) :)'	

Table 8: Example Irish tweet with UDPipe and NLTK tokenization

A.4 Lemma and POS-Tagging of Newly Crawled Tweets

A.5 Conversion to CoNLL-U Format

Table 9 shows that the Morfette format is a subset of the CoNLL-U format used by UDPipe. The LTC and CTC were thus converted automatically from the 3-column Morfette format, consisting of the token, lemma and POS-tag to the 10-column CoNLL-U format. CoNLL-U enables additional token-level annotation i.e. a token id, language-specific part-of-speech tags (XPOS), morphological features, the head of the current word, the dependency relation, an enhanced dependency graph in the form of a list of head-deprel pairs and any other miscellaneous annotation.[15] CoNLL-U also requires a sentence ID and the original raw text to be included preceding the annotation. Further, in the miscellaneous column, the label 'SpaceAfter=No' encodes infor-

[15] In order to make optimum use of the time spent by the annotator, language-specific part-of-speech tags, morphological features and enhanced dependency annotation were not included in this version of the TwittIrish dataset. These elements can be automatically added in later versions of the treebank.

CoNLL-U	Morfette			CoNLL-U					
ID	**FORM**	**LEMMA**	**UPOS**	**XPOS**	**FEATS**	**HEAD**	**DEPREL**	**DEPS**	**MISC**
1	Cuirfidh	cuir	VERB	_	_	0	root	_	_
2	mé	mé	PRON	_	_	1	nsubj	_	_
3	DM	DM	NOUN	_	_	1	obj	_	_
4	chuici	chuig	ADP	_	_	1	obl:prep	_	_
‘Cuirfidh mé DM chuici’ ‘I will send her a DM’									

Table 9: Example conversion of Irish tweet from Morfette to CoNLL-U format

mation about which tokens have a space after them in the original text for detokenisation purposes enabling automatic conversion from raw text to tree and vice versa.

A.6 Review

In order to assess the accuracy of the dependency annotation, a subset of the annotated data, consisting of 46 trees (773 tokens), was reviewed for errors by another Irish speaker trained in linguistic annotation. The task of the reviewer was to flag possible errors in the form of a token with an incorrect head and/or label. 46 possible errors were identified by the reviewer. The possible errors were then discussed by a team of two expert annotators to confirm whether the possible errors were true errors. 32 possible errors were confirmed as true errors. The overall accuracy of the treebank annotation can be estimated as 95.86% by dividing the number of correctly annotated tokens by the total number of tokens in the review.

16 tokens (2.07% of all tokens in the review) had an incorrect label and correct head. Figure 13 exemplifies one such correction. *Go* is a common particle in Irish, which can precede an adjective to create an adverb. When used for this function it is roughly equivalent to the suffix ‘-ly’ in English. e.g. *Ainnis* (‘miserable’), *go hainnis* (‘miserably’). For this reason, a parser is likely to annotate this construction as `advmod`. However, these constructions also appear as the predicate of the substantive Irish verb *bí*(‘to be’) and in this case they should be considered as `xcomp:pred`.[16]

12 tokens (1.55% of all tokens in the review) had an incorrect head and correct label. The most common error (5 instances) was incorrect punctuation attachment. Only 4 tokens (0.52%) were identified as having both incorrect head and label. Figure 14 shows the phrase *maith sibh* (‘good on you’) incorrectly annotated with *sibh* as the `root` and *maith* as its `adjectival modifier`. It was identified in the review that *maith* should be considered the adjective predicate of an elided copula (Stenson, 2019).

The full phrase is thought to be *is maith sibh* and the corrected annotation is shown in Figure 15.

B Parsing Experiments

B.1 LAS by UPOS

LAS	TwittIrish High	TwittIrish Low
IUDT High	DET, ADP, PART, AUX, PRON, SCONJ	VERB, PROPN, PUNCT, X, INTJ
IUDT Low	ADJ, CCONJ, ADV	NOUN, NUM, SYM

Table 10: Confusion matrix of LAS by UPOS tag acheived by AllenNLP Parser with gaBERT embeddings on the IUDT and TwittIrish test sets

Table 10 shows which UPOS tags are associated with higher or lower than average LAS in both test sets. High accuracy is correlated with tokens which occur frequently and have low variation.[17]

UPOS tags `DET`, `ADP`, `PART`, `AUX`, `PRON` and `SCONJ` are associated with higher than average LAS in both the TwittIrish and IUDT test sets. In the IUDT, a high proportion, 8.87%, of tokens have the UPOS tag `DET`. As is common with function words, `DET` comprises of a closed set of lemmata and thus has the low variation of 0.21%.

The tags `ADJ`, `CCONJ`, and `ADV` are associated with higher than average LAS in the TwittIrish tet set but lower than average LAS in the IUDT. This might be because these tags are more likely to be involved in more complex, ambiguous or long-distance attachments.

The tags `VERB`, `PROPN`, `PUNCT`, `X`, and `INTJ` are

[16] https://universaldependencies.org/ga/dep/xcomp-pred.html

[17] Variation is calculated by dividing the number of occurrences by then number of unique lemmata

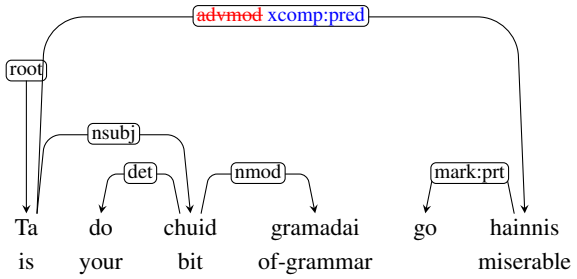

'Your grammar is awful'

Figure 13: Reviewed tweet with corrected label.

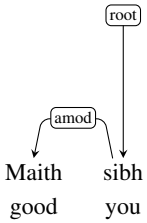

'Good on you'

Figure 14: Example of tweet with incorrect head and label.

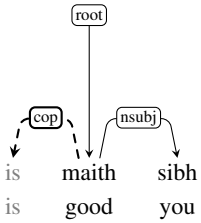

'Good on you'

Figure 15: Reviewed tweet with corrected head and label.

LAS	TwittIrish High	TwittIrish Low
IUDT High	nmod:poss, det, case, fixed, obj, flat:name, nsubj, mark:prt, obl:prep, cop, cc, amod, csubj:cop, mark, nummod, case:voc	root, csubj:cleft, punct
IUDT Low	xcomp:pred, advmod, obl, acl:relcl, nmod, xcomp	discourse, compound, flat, appos, parataxis, advcl, vocative, obl:tmod, ccomp, conj

Table 11: Confusion matrix of LAS by dependency label acheived by AllenNLP Parser with gaBERT embeddings on the IUDT and TwittIrish test sets

associated with higher than average LAS in the IUDT tet set but lower than average LAS in TwittIrish. In the case of VERB and PUNCT, this can be attributed to the non-sentential nature of tweets. UPOS tags NOUN, NUM and SYM are associated with lower than average LAS in both the TwittIrish and IUDT test sets. In the IUDT, a low proportion, 0.02%, of tokens have the UPOS tag SYM. The variation is high 83.33%.

B.2 LAS by Dependency Relation

Table 11 shows that high accuracy is associated with dependency relations nmod:poss, det, case, fixed, obj, flat:name, nsubj, mark:prt, obl:prep, cop, cc, amod, csubj:cop, mark, nummod, case:voc in both the IUDT and TwittIrish. root, csubj:cleft and, punct are associated with higher than average LAS in both the TwittIrish and IUDT test sets. xcomp:pred, advmod, obl, acl:relcl, nmod, and xcomp are associated with higher than average LAS in the TwittIrish tet set but lower than average LAS in the IUDT. are associated with higher than average LAS in the IUDT tet set but lower than average LAS in TwittIrish. discourse, compound, flat, appos, parataxis, advcl, vocative, obl:tmod, ccomp, and conj are associated with lower than average LAS in both the TwittIrish and IUDT test sets.