

Article

Human Evaluation of English–Irish Transformer-Based NMT

Séamus Lankford ^{1,2,*}, Haithem Afli ²  and Andy Way ¹ 

¹ ADAPT Centre, School of Computing, Dublin City University, D09 Y074 Dublin, Ireland; andy.way@adaptcentre.ie

² Department of Computer Sciences, Munster Technological University, T12 P928 Cork, Ireland; haithem.afli@mtu.ie

* Correspondence: seamus.lankford@adaptcentre.ie

Abstract: In this study, a human evaluation is carried out on how hyperparameter settings impact the quality of Transformer-based Neural Machine Translation (NMT) for the low-resourced English–Irish pair. SentencePiece models using both Byte Pair Encoding (BPE) and unigram approaches were appraised. Variations in model architectures included modifying the number of layers, evaluating the optimal number of heads for attention and testing various regularisation techniques. The greatest performance improvement was recorded for a Transformer-optimized model with a 16k BPE subword model. Compared with a baseline Recurrent Neural Network (RNN) model, a Transformer-optimized model demonstrated a BLEU score improvement of 7.8 points. When benchmarked against Google Translate, our translation engines demonstrated significant improvements. Furthermore, a quantitative fine-grained manual evaluation was conducted which compared the performance of machine translation systems. Using the Multidimensional Quality Metrics (MQM) error taxonomy, a human evaluation of the error types generated by an RNN-based system and a Transformer-based system was explored. Our findings show the best-performing Transformer system significantly reduces both accuracy and fluency errors when compared with an RNN-based model.

Keywords: human evaluation; MQM; neural machine translation; Irish; low-resource languages



Citation: Lankford, S.; Afli, H.; Way, A. Human Evaluation of English–Irish Transformer-Based NMT. *Information* **2022**, *13*, 309. <https://doi.org/10.3390/info13070309>

Academic Editor: Enrico Dentì

Received: 6 May 2022

Accepted: 21 June 2022

Published: 25 June 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

A new era of high-quality translations has been heralded with the advent of NMT. Given that large datasets are a prerequisite for high-quality NMT, these improvements are not always evident in the translation of low-resource languages. In the context of such languages, which suffer from a sparsity of data, alternative approaches must be adopted.

Developing applications and models to address the challenges of low-resource language technology is an important part of this research. This technology incorporates new methods, which reduce the impact that data scarcity has on the digital engagement of low-resource languages. One approach is to use a mechanism that helps NMT systems to learn from unlabeled data using dual-learning [1,2].

Out-of-the-box NMT systems, trained on English–Irish data, have been shown to achieve a lower translation quality compared with using a tailored SMT system [3]. It is in this context that further research is required in the development of NMT for low-resource languages, and the Irish language in particular.

Most research on the choice of subword models has focused on high-resource languages [4,5]. Translation, by its nature, requires an open vocabulary and the use of subword models aims to address the fixed-vocabulary problem associated with NMT. Rare and unknown words are encoded as sequences of subword units. By adapting the original BPE algorithm [6], the use of BPE submodels can improve translation performance [7,8]. In the context of developing models for English-to-Irish translation, there were no clear recommendations on the choice of subword model types. Character-based models were briefly explored due to their simplicity and reduced memory requirements. However, they were not considered suitable given that most single characters do not carry meaning in the

English and Irish languages. Therefore, one of the objectives of our research is to identify which type of subword model performs best in this low-resource scenario.

An important goal of this study is to extend our previous work [9] by providing a human evaluation (HE) and comparison of EN→GA machine translation (MT) on systems that use either a baseline RNN architecture or a subword-model optimized Transformer model.

This paper describes the context in which our research was conducted and provides a background of the types of available architecture in Section 2. A detailed overview of our approach is outlined in Section 3, where we provide details of the data and parameters used in our NMT systems. The empirical results, using both automatic metrics and a human evaluation, are presented in Section 4. Finally, our findings are discussed and possibilities for future work are described in Section 6.

2. Background

Native speakers of low-resource languages are often excluded from useful content since, more often than not, online content is not available to them in their language of choice. This digital divide experienced by second-language speakers has been well-documented in the research literature [10,11].

Research on MT in low-resource scenarios seeks to directly address this challenge of exclusion via pivot languages [12], and indirectly, via domain adaptation of models [13]. Consequently, research efforts focusing on NMT [14,15] have resulted in a state-of-the-art (SOA) performance being attained for multiple language pairs [16,17]. The Irish language is a primary example of a low-resource language that will benefit from this research. NMT involving Transformer model development will improve performance in specific domains of low-resource languages.

2.1. Hyperparameter Optimization

Hyperparameters are employed to customize machine learning models such as translation models. It has been shown that machine learning performance may be improved through hyperparameter optimization (HPO) rather than just using default settings [18]. The principal methods of HPO are Grid Search [19] and Random Search [20].

2.1.1. RNN

The tasks of natural language processing (NLP), speech recognition and MT are often performed by RNNs. This architecture enables previous outputs to be used as inputs while having hidden states. In the context of MT, such neural networks were ideal due to their ability to process inputs of any length. Furthermore, the model sizes do not necessarily increase with the input size. Commonly used variants of RNN include Bidirectional (BRNN) and Deep (DRNN) architectures. However, the problem of vanishing gradients coupled with the development of attention-based algorithms often leads to Transformer models performing better than RNNs.

2.1.2. Transformer

The greatest improvements have been demonstrated when either the RNN or the CNN architecture is abandoned completely and replaced with an attention mechanism creating a much simpler and faster architecture known as Transformer. Experiments in MT tasks show such models are better in quality due to greater parallelization while requiring significantly less time to train [21].

Transformer models use attention to focus on previously generated tokens. The approach allows for models to develop a long memory, which is particularly useful in the domain of language translation. Performance improvements to both RNN and CNN approaches may be achieved through the introduction of such attention layers in the translation architecture.

2.2. SentencePiece

Designed for NMT, SentencePiece is a language-independent subword tokenizer that provides an open-source C++ and a Python implementation for subword units. An attractive feature of the tokenizer is that SentencePiece directly trains subword models from raw sentences [22].

2.3. Human Evaluation

Human evaluation, within NLP and MT, is a topic of growing importance, which often has its own dedicated research track or workshop at major conferences [23]. This focus has resulted in many publications in the area of HE that relate to MT [24,25] and it has particularly benefited the evaluation of low-resource languages [26,27].

The best practice for the HE of MT has been published in the form of a series of recommendations [28]. As part of our research, we adopted these recommendations, which are in line with similar EN-GA HE studies at the ADAPT centre [29]. Specifically, these recommendations encourage the use of professional translators, evaluation at the document level and assessments of both fluency and accuracy. Original source texts were also used in the training and test data.

These recommendations have been complemented by a fine-grained human analysis, which uses both a Scalar Quality Metric (SQM) and the MQM.

3. Proposed Approach

Considerable performance improvements have been achieved using the HPO of RNN models in low-resource settings. One of the key research questions, evaluated as part of this study, is to identify the extent to which such optimization techniques may be applied to low-resource Transformer models. Evaluations included modifying the number of attention heads, changing the number of layers and experimenting with regularization techniques such as dropout and label smoothing. Most importantly, the choice of subword model type and the vocabulary size are evaluated. Furthermore, previous research focuses on using an automatic evaluation of performance, whereas we propose combining a HE approach with automatic metrics.

In order to test the effectiveness of our approach, optimization was carried out on an English–Irish parallel dataset: a general corpus of 52k lines from the Directorate General for Translation (DGT). With DGT, the test set used 1.3k lines and the development set comprised of 2.6k lines. All experiments involved concatenating source and target corpora to create a shared vocabulary and a shared SentencePiece subword model. The adopted approach is illustrated in Figure 1.

3.1. Architecture Tuning

It is difficult and costly to tune systems using a conventional grid search approach given the long training times associated with NMT. Therefore, we adopted a random search approach in the HPO of our Transformer models.

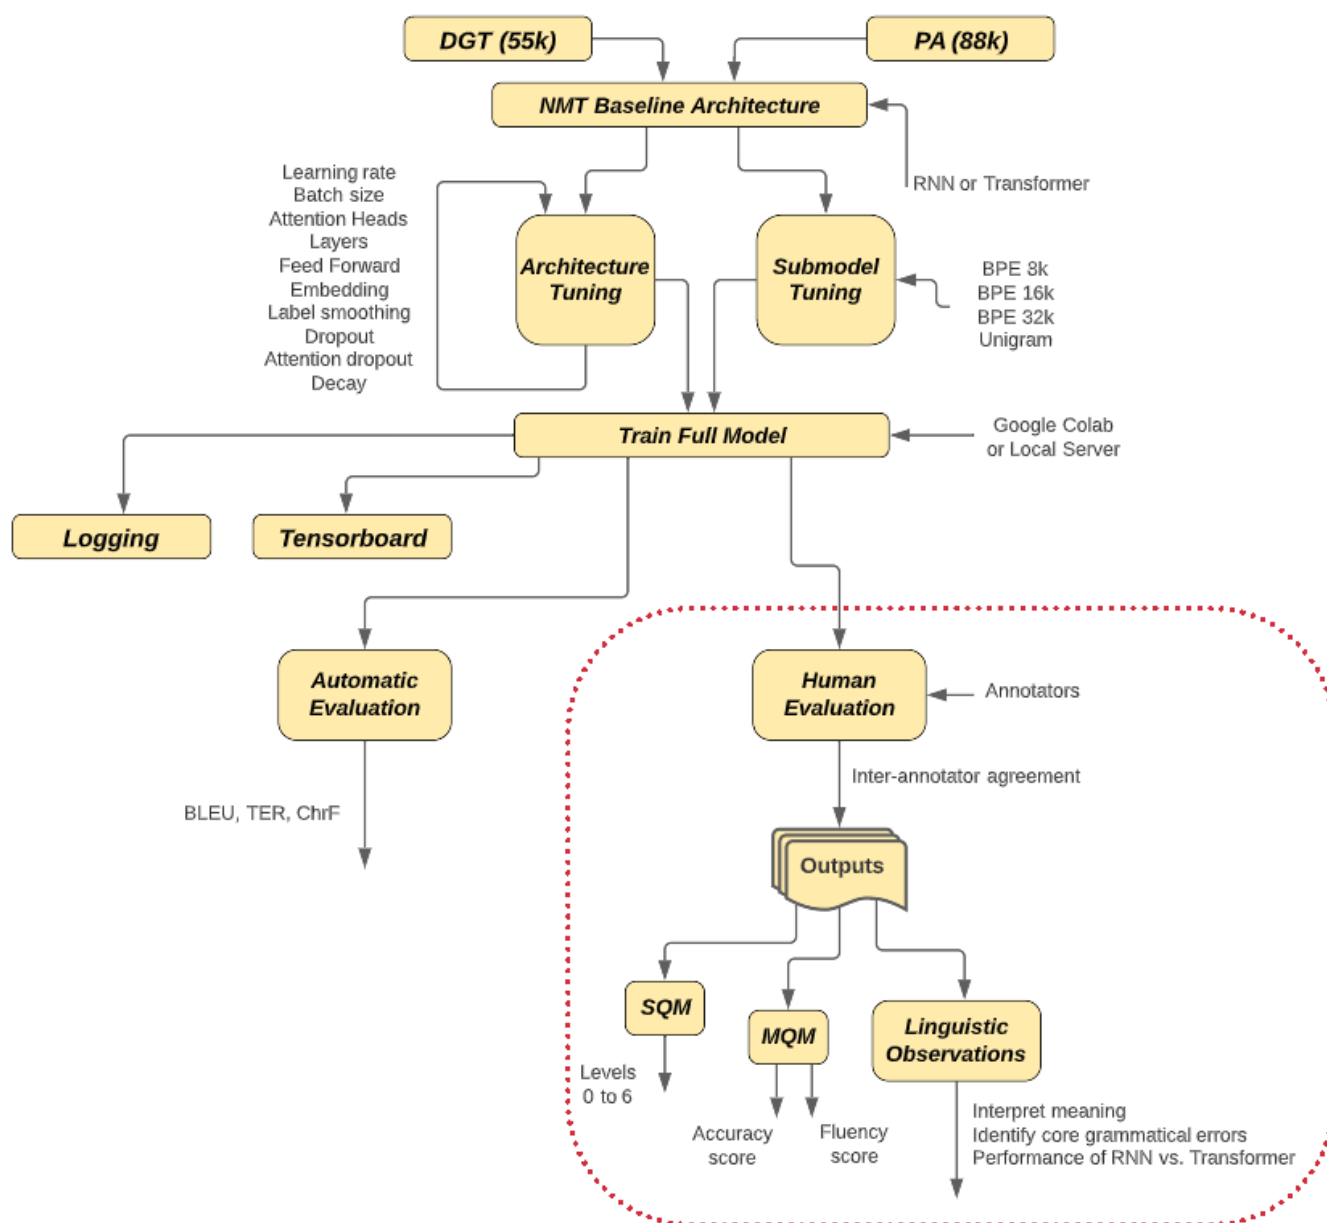
Using smaller and fewer layers with low-resource datasets has previously been shown to improve performance [30]. Furthermore, the use of shallow Transformer models has been demonstrated to improve the translation performance of low-resource NMT [31]. Guided by these findings, configurations were tested, which varied the number of neurons in each layer and modified the number of layers used in the Transformer architecture.

Varying degrees of dropout were applied to Transformer models to evaluate the impact of regularization. Configurations using smaller (0.1) and larger values (0.3) were applied to the output of each feed-forward layer.

3.2. Subword Models

Incorporating a word segmentation approach, such as BPE, is now standard practice when developing NMT models. Subword models are particularly beneficial for low-resource languages since rare words are often a problem. In the context of English-to-

Irish translation, there is no clear agreement as to what constitutes the best approach. Consequently, subword regularization techniques involving BPE and unigram models were evaluated as part of this study to determine the optimal parameters for maximizing translation performance. BPE models with varying vocabulary sizes of 4k, 8k, 16k and 32k were evaluated.



Fine-grained human evaluation

Figure 1. The proposed approach to evaluate the baseline architectures of RNN and Transformer models is illustrated above. Using a random search approach, the values outlined in Table 1 were tested to determine the optimal hyperparameters. Short cycles of 5k training steps were applied to test a range of values for each parameter. Once an optimal value was identified within the sampled range, it was locked in for tests on subsequent parameters. A fine-grained HE was conducted on the output from the DGT dataset and its results were compared with an automatic evaluation.

Table 1. Transformer HPO using a random search approach. The optimal hyperparameters are highlighted in bold. The best-performing model used two attention heads and was trained on a 55k DGT corpus.

Hyperparameter	Values
Learning rate	0.1, 0.01, 0.001, 2
Batch size	1024, 2048 , 4096, 8192
Attention heads	2 , 4, 8
Number of layers	5, 6
Feed-forward dimension	2048
Embedding dimension	128, 256 , 512
Label smoothing	0.1 , 0.3
Dropout	0.1, 0.3
Attention dropout	0.1
Average Decay	0, 0.0001

3.3. Human Evaluation of NMT

Morphological-rich languages, such as Irish, have a high degree of inflection and a free word order that gives rise to specific translation issues when translating from English. Grammatical categories such as gender or case inflections in nouns are often difficult to reliably generate in an Irish translation.

One of the goals of this research is to explore how an NMT system handles these issues compared with an RNN approach. Existing research suggests NMT systems should improve these linguistic aspects. NMT, with its use of subword models, implicitly addresses the problem in an unsupervised manner, without understanding the actual formal rules of grammatical categories.

Previous HE studies that evaluate English–Irish MT performance have focused on the differences between an SMT and an NMT approach [3]. In the context of our research, HE was conducted on purely NMT methods, which included RNN and Transformer approaches. Furthermore, our study is differentiated by using both SQM and MQM as our HE metrics.

It is clear from our earlier experimental findings, based solely on automatic evaluation metrics, that a Transformer approach leads to significant improvements compared to traditional RNN systems. However, as with most automatic scoring methods, these simply provide an overall score for each system but do not indicate the exact nature of the linguistic problems that may be encountered in translation. Therefore, it can be said that automatic evaluation does not address the question of the linguistic or grammatical quality of the target output. Nuances, such as how gender or cases are handled, are not covered by this approach.

To achieve a deeper understanding of the linguistic errors created by our RNN and Transformer systems, a fine-grained HE was conducted. The outputs from these systems were systematically analyzed and compared in a manual error analysis. This approach captures the nature of the translation errors for each of the evaluated systems. The output from this study forms the basis of future work, which will help to improve the translation quality of our models. The annotation framework, the overall annotation process and inter-annotator agreement are discussed below, and broadly follow the approach adopted by other fine-grained HE studies [32].

3.3.1. Scalar Quality Metrics

SQM [33] adapts the WMT shared-task settings to collect segment-level scalar ratings with a document context. SQM uses a scale from 0 to 6 for translation quality assessment. This is a modification of the WMT approach [34], which uses a range from 0 to 100.

With this evaluation approach, annotators must select a rating from 0 through 6 when presented with the source and target sentences. The SQM quality levels for 0, 2, 4 and 6 are outlined in Table 2. Annotators may also choose intermediate levels of 1, 3 and 5 in cases where the translations do not exactly match the core SQM levels.

Table 2. SQM levels explained [33].

SQM Level	Details of Quality
6	Perfect Meaning and Grammar: The meaning of the translation is completely consistent with the source and the surrounding context (if applicable). The grammar is also correct.
4	Most Meaning Preserved and Few Grammar Mistakes: The translation retains most of the meaning of the source. This may contain some grammar mistakes or minor contextual inconsistencies.
2	Some Meaning Preserved: The translation preserves some of the meaning of the source but misses significant parts. The narrative is hard to follow due to fundamental errors. Grammar may be poor.
0	Nonsense/No meaning preserved: Nearly all information is lost between the translation and source. Grammar is irrelevant.

3.3.2. Multidimensional Quality Metrics

As part of QTLaunchpad project (<https://www.qt21.eu/>, accessed on 5 May 2022) the MQM framework (<https://www.qt21.eu/mqm-definition/definition-2015-12-30.html>, accessed on 5 May 2022) was developed to provide a framework of how manual evaluation could be performed via a detailed error analysis. A single metric for all uses is not imposed. Instead, a comprehensive catalogue of quality issue types, with standardized names and definitions, is provided. This catalogue may be customized for specific tasks. In addition to forming a reliable methodology for quality assessment, it also allows for us to specify which error tags were relevant to our task.

To adapt the generic MQM framework for our context, we followed the official guidelines for scientific research [35]. The details of our customization of MQM are discussed below.

A large variety of tags, on several annotation layers, are proposed within the original MQM guidelines. However, this full MQM tagset is too detailed for a specific annotation task. Therefore, when evaluating our MT output, the smaller default set of evaluation categories, specified in the core tagset, were used. These standard top-level categories of accuracy and fluency, which are proposed by the MQM guidelines, are illustrated in Figure 2. A special non-translation error was used to tag an entire sentence, which was too badly translated to allow for the identification of individual errors.

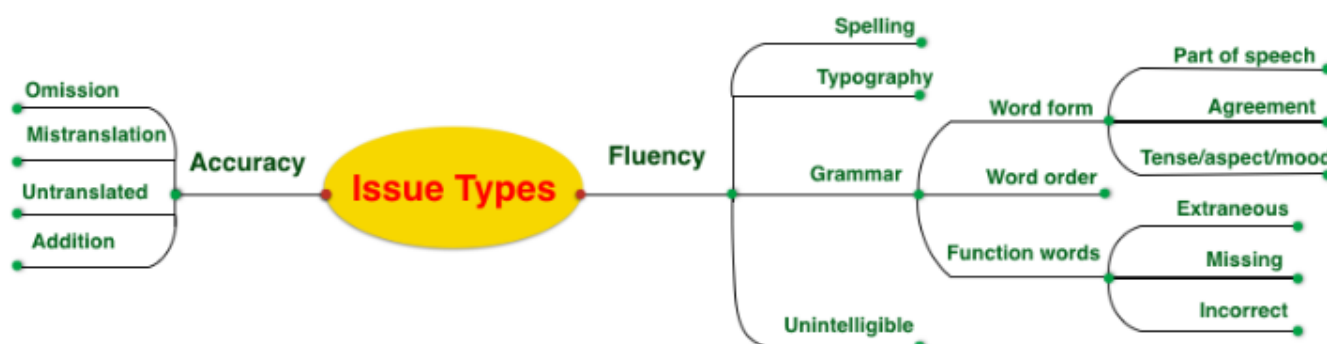


Figure 2. The core set of error categories proposed by the MQM guidelines.

Error severities are specified as either major or minor errors and are assigned independently of category. These correspond to actual translation/grammatical errors or smaller imperfections, respectively. The recommended default weights [35] were used, which allocate a weight of 1 to minor errors whereas major errors are assigned a weight of 10.

Furthermore, the non-translation category was allocated a weight of 25, an approach which is line with the best practice established in previous studies [33].

The annotators were instructed to identify all errors within each sentence of the translated output for both systems. The error categories used by the annotators are outlined in Table 3.

Table 3. Description of error categories within the core MQM framework [33].

Category	Sub-Category	Description
Non-translation		Impossible to reliably characterize the 5 most severe errors.
Accuracy	Addition	Translation includes information not present in the source.
	Omission	Translation is missing content from the source.
	Mistranslation	Translation does not accurately represent the source.
	Untranslated text	Source text has been left untranslated.
Fluency	Punctuation	Incorrect punctuation
	Spelling	Incorrect spelling or capitalization.
	Grammar	Problems with grammar, other than orthography.
	Register	Wrong grammatical register (e.g., inappropriately informal pronouns).
	Inconsistency	Internal inconsistency (not related to terminology).
	Character encoding	Characters are garbled due to incorrect encoding.

3.3.3. Annotation Setup

Annotations were carried using the simpler SQM approach and a more detailed, fine-grained MQM approach. The hierarchical taxonomy of our MQM implementation is illustrated in Figure 2, whereas the SQM categories are summarized in Table 2.

Two annotators with similar backgrounds, were used for the annotation of outputs from an RNN system and a Transformer system. Both annotators are native speakers of Irish and neither had prior experience with MQM. Prior to annotation, they were thoroughly familiarized with the process and the official MQM annotation guidelines. These guidelines offer detailed instructions for annotation within the MQM framework.

Both annotators have been very involved in the education sector for decades. One of the annotators has edited numerous English-language and Irish-language books during her career as a university lecturer. The second annotator has a PhD in Irish-language place names. In addition, he has written numerous books in both English and Irish. Given their experience and strong language backgrounds, they were well-equipped to handle the task at hand.

Using a test set of 20 randomly selected sentences, the annotators were presented with the English source text, an Irish reference translation and the two unannotated system outputs: one generated using an RNN model and the other created using a Transformer model. Potential bias was removed by using blind annotation such that annotators did not know which model the translation output came from. The annotators worked independently of each other but were occasionally in contact to discuss the process and how to approach difficult sentences.

Translations from the RNN and the Transformer system were annotated by both annotators, meaning that each system translated the same 20 sentences and each annotator annotated the resulting 40 translated sentences (20 source sentences for 2 MT systems), producing a total of 80 annotated sentences. The annotated dataset is publicly available on GitHub (<https://github.com/seamusl/isfeidirlinn>, accessed on 5 May 2022)

Once the annotation data were extracted, each annotator analyzed the output to determine the performance of each system for each error category.

3.3.4. Inter-Annotator Agreement

Low inter-annotator agreement (IAA) scores is a common problem experienced when using manual MT evaluation approaches such as MQM [36,37]. To determine the validity of

the findings of our research, it is important to check the level of agreement between our annotators [38].

Cohen's kappa (k) [39] was used to determine inter-annotator agreement. Agreement was calculated based on the annotations of each individual system, with agreement being observed at the sentence level. With this approach, the differences in agreement across systems was explored and we also gained a high-level view of overall agreement between the annotators. Furthermore, Cohen's kappa was calculated separately for every error type and the findings are outlined in Table 4.

Table 4. Inter-annotator agreement using Cohen values

Error Type	RNN	NMT
Non-translation	1.0	1.0
Accuracy	1.0	1.0
Addition	1.0	1.0
Omission	1.0	1.0
Mistranslation	−0.071	1.0
Untranslated text	0.0	1.0
Fluency		
Punctuation	0.651	1.0
Spelling	0.0	0.0
Grammar	0.867	0.895
Register	1.0	1.0
Inconsistency	1.0	1.0
Character Encoding	1.0	1.0

4. Empirical Evaluation

4.1. Experimental Setup

4.1.1. Datasets

The performance of the Transformer and RNN approaches is evaluated on a publicly available English-to-Irish parallel dataset from the Directorate General for Translation (DGT) (<https://ec.europa.eu/info/departments/translation>, accessed on 5 May 2022). The Joint Research Centre of the DGT has made all its translation memory (i.e. sentences and their professionally produced translations) available, which covers the official European Union languages [40]. Included in the training data are parallel texts from the Digital Corpus of the European Parliament (DCEP) and the DGT. Crawled data, from sites of a similar domain, are also incorporated. This dataset is broadly categorised as generic and is publicly available.

4.1.2. Infrastructure

Model development was conducted using local workstations, each of which was built with an AMD Ryzen 7 2700X processor, 16 GB memory, a 256 SSD and an NVIDIA GeForce GTX 1080 Ti.

In addition, a Google Colab Pro subscription enabled rapid prototype development and created zero-emission models. The available computing power of the Google Cloud was much higher than our local infrastructure and provided servers with 16 GB graphic cards (NVIDIA Tesla P100 PCIe) and up to 27 GB of memory [41]. Larger Transformer models were built on local infrastructure, since long builds timed out on Colab due to Google restrictions. The Pytorch implementation of OpenNMT 2.0, an open-source toolkit for NMT [42], was used to train all MT models.

4.1.3. Metrics

The performance of all models was evaluated using the automated metrics of BLEU [43], TER [44] and ChrF [45]. Case-insensitive BLEU scores are reported at the corpus level.

4.2. Automatic Evaluation Results

4.2.1. Performance of Subword Models

The impact that choice of subword model has on translation is highlighted in Tables 5 and 6. Incorporating any subword model type led to improvements in model accuracy when training both RNN and Transformer architectures.

Table 5. RNN performance on DGT dataset of 52k lines. There were zero carbon emissions in building these models, since smaller RNN models were trained on Google Colab servers, which are carbon-neutral.

Architecture	BLEU ↑	TER ↓	ChrF3 ↑	Steps	Runtime (h)	kgCO ₂
dgt-rnn-base	52.7	0.42	0.71	75k	4.47	0
dgt-rnn-bpe8k	54.6	0.40	0.73	85k	5.07	0
dgt-rnn-bpe16k	55.6	0.39	0.74	100k	5.58	0
dgt-rnn-bpe32k	55.3	0.39	0.74	95k	4.67	0
dgt-rnn-unigram	55.6	0.39	0.74	105k	5.07	0

Table 6. Transformer performance on 52k DGT dataset. The highest performing model uses 2 attention heads. All other models use 8 attention heads. Transformer models were long-running builds, which had to be carried out on local servers.

Architecture	BLEU ↑	TER ↓	ChrF3 ↑	Steps	Runtime (h)	kgCO ₂
dgt-trans-base	53.4	0.41	0.72	55k	14.43	0.81
dgt-trans-bpe8k	59.5	0.34	0.77	200k	24.48	1.38
dgt-trans-bpe16k	60.5	0.33	0.78	180k	26.90	1.52
dgt-trans-bpe32k	59.3	0.35	0.77	100k	18.03	1.02
dgt-trans-unigram	59.3	0.35	0.77	125k	21.95	1.24

A baseline RNN model, illustrated in Table 5, achieved a BLEU score of 52.7, whereas the highest-performing BPE variant, using a 16k vocab, recorded an improvement of nearly three points, with a score of 55.6.

In the context of Transformer architectures, highlighted in Table 6, the use of subword models delivers significant performance improvements. The performance gains for Transformer models are much higher compared with the improvements recorded by the RNN models. A baseline Transformer model achieves a BLEU score of 53.4, whereas a Transformer model, with a 16k BPE submodel, has a score of 60.5, representing a BLEU score improvement of 13% at 7.1 BLEU points.

For translating into a morphologically rich language, such as Irish, the ChrF metric has proven successful in showing a strong correlation with human translation [46]. In the context of our experiments, this worked well in highlighting the performance differences between RNN and Transformer architectures.

4.2.2. Transformer Performance Compared with RNN

The performance of RNN models is contrasted with the Transformer approach in Figures 3 and 4. Transformer models, as anticipated, outperformed all their RNN counterparts. It is interesting to note the impact of choosing the optimal vocabulary size for BPE submodels. Choosing a BPE vocabulary of 16k words yields the highest performance.

Furthermore, the TER scores highlighted in Figure 4 reinforce the findings that using 16k BPE submodels on Transformer architectures leads to a better translation performance. The TER score for the 16k BPE Transformer model is significantly better (0.33) when compared with the baseline performance (0.41).

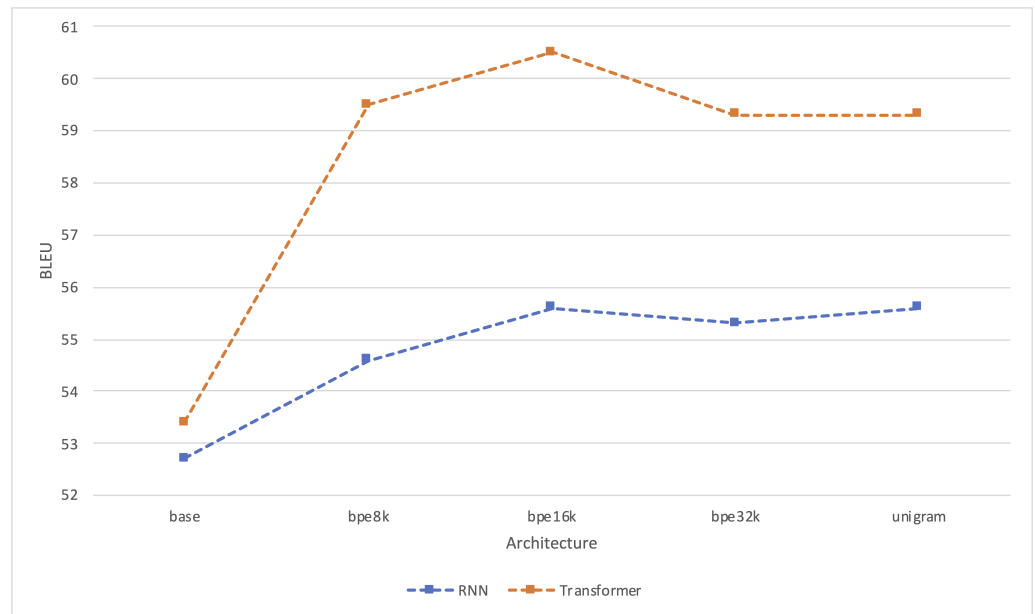


Figure 3. BLEU performance for all model architectures is compared. The use of a BPE subword model improved translation performance in all cases. The best-performing model was built using a 16k BPE subword model on a Transformer architecture.




Figure 4. TER performance for all model architectures. The highest-performing model uses a 16k BPE subword model on a Transformer architecture. In all instances, incorporating a subword model improves TER.

4.3. Human Evaluation Results

The aggregate total of errors found by annotators for each system is highlighted in Table 7. Looking at the aggregate data alone, it is evident that both annotators have judged that the RNN system contains more errors, and that the NMT system contains less errors.

Table 7. Total errors found by each annotator using the MQM metric.

System	Annotator 1		Annotator 2	
	RNN	Transformer	RNN	Transformer
Total Errors	41	23	43	23

While such a high-level view is instructive in determining which system is better, it lacks the granularity required to pinpoint the linguistic aspects of how these translations can be improved. To achieve a deeper insight, a fine-grained analysis of the error types was conducted, the results of which are displayed in Table 8. Categorized by error type, the sum of error tags by each annotator for each system is outlined.

Table 8. Transformer and RNN approach is compared using concatenated annotation data across both annotators. In all MQM error categories, the Transformer architecture performs better, apart from a tie in the omission category.

Error Type	RNN	NMT
	Error	Error
Non-translation	0	0
Accuracy		
Addition	10	4
Omission	12	12
Mistranslation	26	14
Untranslated text	4	1
Fluency		
Punctuation	5	4
Spelling	1	0
Grammar	20	11
Register	2	0
Inconsistency	2	0
Character Encoding	0	0
Total errors	82	46

5. Environmental Impact

The environmental impact of all aspects of computing has received increased research interest in recent times. Much of this effort has concentrated on NMT's carbon footprint [47,48]. To assess the environmental impact of our NMT models, we tracked energy consumption during their development.

Prototype model development was carried out using Google Colab which is a carbon neutral platform [49]. However, longer running Transformer experiments were conducted on local servers using 324 g CO₂ per kWh (<https://www.seai.ie/publications/Energy-in-Ireland-2020.pdf>, accessed on 5 May 2022) [50]. The net result was just under 10 kg CO₂, created for a full run of model development. Models developed during this study will be reused for ensemble experiments in future work.

The environmental costs of our model development were tracked to serve as a benchmark for future work. Awareness of such costs will impose a discipline on our work, such that we opt for carbon-neutral cloud providers. In cases where models are developed on local infrastructure, this will encourage the use of more efficient GPUs and the utilization of techniques that result in faster builds.

6. Discussion

Validation accuracy and model perplexity (PPL) in developing the baseline and optimal Transformer models are illustrated in Figures 5 and 6. Training a Transformer model with a 16k BPE subword model boosted the validation accuracy by over 8% compared to its baseline.




Figure 5. Transformer baseline.




Figure 6. Transformer 16k BPE subword model.

Rapid convergence was observed while training the baseline model such that little accuracy improvement occurs after 20k steps. Including a subword model led to slower converging models, with only marginal gains recorded after 60k steps. Examining Figures 5 and 6, PPL achieves a lower global minimum when the Transformer approach is used with a 16k BPE submodel. The PPL global minimum (2.7) is over 50% lower than the corresponding PPL for the Transformer base model (5.5). This finding illustrates that choosing an optimal subword model delivers significant performance gains.

Translation engine performance, at the corpus level, was benchmarked against Google Translate's (<https://translate.google.com/>, accessed on 5 May 2022) English-to-Irish translation service, which is freely available on the internet. Four random samples were selected from the English source test file and are presented in Table 9. Translation of these samples was carried out on the optimal Transformer model using Google Translate. Case-insensitive, sentence-level BLEU scores were recorded and are presented in Table 10. It must be acknowledged that this comparison is not entirely valid given that Google does not have access to our training data, nor do we have unlimited access to the Google cloud infrastructure. Nonetheless, the results are encouraging and indicate a good performance by our translation models on the DGT dataset.

Table 9. Random samples of human reference translations taken from the test dataset.

Source Language (English)	Reference Human Translation (Irish)
A clear harmonised procedure, including the necessary criteria for disease-free status, should be established for that purpose.	Ba cheart nós imeachta comhchuibhithe soiléir, lena n-áirítear na critéir is gá do stádas saor ó ghalair, a bhunú chun na críche sin.
the mark is applied anew, as appropriate.	déanfar an mharcáil arís, mar is iomchuí.
If the court decides that a review is justified on any of the grounds set out in paragraph 1, the judgment given in the European Small Claims Procedure shall be null and void.	Má chinneann an chúirt go bhfuil bonn cirt le hathbhreithniú de bharr aon cheann de na forais a leagtar amach i mír 1, beidh an breithiúnas a tugadh sa Nós Imeachta Eorpach um Éilimh Bheaga ar neamhní go hiomlán.
households where pet animals are kept;	teaghlaigh ina gcoimeádtar peataí;

Table 10. Transformer model compared with Google Translate using random samples from the DGT corpus. Full evaluation of Google Translate's engines on the DGT test set, with 1.3k lines, generated a BLEU score of 46.3 and a TER score of 0.44. Comparative scores on the test set using our Transformer model, with 2 attention heads and 16k BPE submodel realised 60.5 for BLEU and 0.33 for TER.

Transformer (16k BPE)	BLEU ↑	Google Translate	BLEU ↑
Ba cheart nós imeachta soiléir comhchuibhithe, lena n-áirítear na critéir is gá maidir le stádas saor ó ghalair, a bhunú chun na críche sin.	61.6	Ba cheart nós imeachta comhchuibhithe soiléir, lena n-áirítear na critéir riachtanacha maidir le stádas saor ó ghalair, a bhunú chun na críche sin.	70.2
go gcuirtear an marc i bhfeidhme, de réir mar is iomchuí.	21.4	cuirtear an marc i bhfeidhm as an nua, de réir mar is cuí.	6.6
Má chinneann an chúirt go bhfuil bonn cirt le hathbhreithniú ar aon cheann de na forais a leagtar amach i mír 1, beidh an breithiúnas a thugtar sa Nós Imeachta Eorpach um Éilimh Bheaga ar neamhní.	77.3	Má chinneann an chúirt go bhfuil údar le hathbhreithniú ar aon cheann de na forais atá leagtha amach i mír 1, beidh an breithiúnas a thugtar sa Nós Imeachta Eorpach um Éilimh Bheaga ar neamhní	59.1
teaghlaigh ina gcoimeádtar peataí;	100	teaghlaigh ina gcoinnítear peataí;	30.2

The optimal parameters selected in this discovery process are identified in bold in Table 1. A higher initial learning rate of 2 coupled with an average decay of 0.0001 led to longer training times but more accurate models. Despite setting an early stopping parameter, many of the Transformer builds continued for the full cycle of 200k steps over periods of 20+ hours.

Training Transformer models with a reduced number of attention heads led to a marginal improvement in translation accuracy with a smaller corpus. Our best-performing model achieved a BLEU score of 60.5 and a TER score of 0.33 with 2 heads and a 16k BPE submodel. By comparison, using 8 heads with the same architecture and dataset yielded 60.3 for BLEU and 0.34 in terms of TER.

Transformer models developed, using state-of-the-art techniques, were evaluated as part of the LoResMT2021 Shared Task [51]. Models developed using our approach, as outlined above, were entered into the competition, and the highest-performing EN-GA system was submitted by our team (ADAPT) [52].

6.1. Inter-Annotator Reliability

In Cohen's original article [39], the interpretation of specific k scores is clearly outlined. There is no agreement with values ≤ 0 , none to slight agreement when scores are in the range of 0.01–0.20, fair agreement is represented by 0.21–0.40, 0.41–0.60 is moderate agreement, 0.61–0.80 is substantial agreement, and 0.81–1.00 is almost perfect agreement.

The literature [53] recommends a minimum of 80% agreement for good inter-annotator agreement. As illustrated in Table 4, there is almost perfect agreement between the annotators when evaluating output from the NMT models. In the case of the RNN outputs, there is disagreement in the mistranslation category but agreement in all other categories. Given these scores, we have a high degree of confidence in our human evaluation of both the RNN and NMT outputs.

6.2. Performance of Is Féidir Linn Models Relative to Google

Using standard Transformer parameters, such as a batch size of 2048 and setting the number of encoder/decoder layers to 6, were observed to perform well. Increasing the regularization dropout to 0.3 and reducing hidden neurons to 256 improved translation performance. Consequently, these values were selected when building all Transformer models.

6.3. Linguistic Observations

A linguistic analysis of the outputs from the Transformer-optimized model is illustrated in Table 11. The English language source sentences and their Irish language translations are presented. The sentences have been selected from fine-grained human evaluation, since they highlight some of the key error types that are encountered. The analysis focuses on the shortcomings of our model outputs, which fall into the following categories: interpretative meaning, core grammatical errors and commonly used irregular verbs. Finally, using the HE metrics of SQM and MQM, the performance of an RNN approach is contrasted with that of the Transformer approach.

6.3.1. Interpreting Meaning

The generic Irish verb “déan” (to do or to make) is used to express more precise concepts such as “to conduct”, “to put into effect” or “to carry out”. Both the RNN and Transformer systems make use of “déan” in a generic way, but they fail to capture the refinement of concept expressed in each of these meanings. An example of this problem is illustrated in GA-1 in Table 11. In this context, a more natural and intuitive translation to capture the expression “to conduct” would be to substitute “a dhéanamh” with “a sheoladh”.

A similar lack of refinement from both systems is also found with the usage of other words. For example, “cuid” (part) is used to translate “operative part” in GA-2. However, a more precise interpretation would be the usage of “gné”, leading to the correct translation “gné oibríochtúil” i.e., “operative part”.

Another example where the translation models failed to correctly interpret the true sense of an English source word into a corresponding Irish translation can be seen in GA-3. The Irish verb “Mainnigh” meaning “to default” would not be used in the context of the source text in EN-3. Using the Irish verb “teip”, meaning “to fail”, is the correct translation of the idea “fails to meet the performance requirements”: “má theipeann an t-oibreoir na ceanglais feidhmíochta a chomhlíonadh.” This error was observed in both the RNN and Transformer model outputs.

Table 11. Linguistic analysis of system outputs. Sources of errors are flagged in blue and in red.

Type	Sentence
EN-1	The lead supervisory authority may request at any time other supervisory authorities concerned to provide mutual assistance pursuant to Article 61 and may conduct joint operations pursuant to Article 62, in particular for carrying out investigations or for monitoring the implementation of a measure concerning a controller or processor established in another Member State.
GA-1	Féadfaidh an príomhúdarás maoirseachta iarraidh, tráth ar bith, ar bith eile lena mbaineann cúnamh frithpháirteach a chur ar fáil de bhun Airteagal 61 agus féadfaidh sé oibríochtaí comhpháirteacha a dhéanamh de bhun Airteagal 62, go háirithe maidir le himscrúduithe a dhéanamh nó maidir le faireachán a dhéanamh ar chur chun feidhme beart i ndáil le rialaitheoir nó próiseálaí atá bunaithe i mBallstát eile.
EN-2	The Office shall mention the judgment in the Register and shall take the necessary measures to comply with its operative part .
GA-2	Luafaidh an Oifig an breithiúnas sa Chlár agus glacfaidh sí na bearta is gá chun cloí lena chuid oibríochtúil.
EN-3	The competent authority may at any time wholly or partially suspend or terminate the contract awarded under this provision if the operator fails to meet the performance requirements.
GA-3	Féadfaidh an t-údarás inniúil an conradh a dámhadh faoin bhforáil seo a chur ar fionraí nó a fhoirceannadh go hiomlán nó go páirteach má mhainníonn an t-oibreoir na ceanglais feidhmíochta a chomhlíonadh.
EN-4	This Directive shall enter into force on the day following that of its publication in the Official Journal of the European Union.
GA-4	Tiocfaidh an Treoir seo i bhfeidhm an lá tar éis lá a fhoilsithe in Iris Oifigiúil An Aontais Eorpaigh.
EN-5	Such special measures are interim in nature, and shall not be subject to the conditions set out in Article 7(1) and (2).
GA-5	Tá bearta speisialta den sórt sin eatramhach, agus ní bheidh said faoi réir na gcoinníollacha a leagtar amach in Airteagal 7(1) agus (2) iad.

6.3.2. Core Grammatical Errors

Grammatical mistakes in the form of the misuse of lenitions (e.g., GA-4), incorrect pronouns (e.g., GA-5) and register errors (e.g., GA-5) were observed in both translation architectures. However, as is evident from both the automatic and MQM evaluations, there were far fewer errors with the Transformer model. Evidence of this can be seen in Table 11. In the case of GA-4, the RNN model included the lenition in “a foilsithe”, whereas the Transformer model correctly removed “h”. The correct use of the feminine noun “treoir” requires the removal of “h” in “foilsithe”.

The misuse of pronouns was observed in the RNN translation model and, to a lesser degree, in the Transformer model. In the case of GA-5, the RNN’s incorrect use of the pronoun “ní bheidh siad” (they will not) is illustrated, whereas the Transformer approach used the correct form “ní bheidh sé” (he will not).

Within the same sentence, GA-5, there is also evidence of a register error. In the English source text EN-5, the use of “shall not be subject to” expresses a stipulation. This is not registered in the Irish translation of “ní bheidh said”, which simply, and less forcefully, means “they will not”. This incorrect use of register was observed with both the RNN and the Transformer approaches. A more formal and closer interpretation of the English source would be the use of the imperative mode: “ná bídis” (let it not be).

6.3.3. Commonly-Used Irregular Verbs

One of the main inadequacies observed in both the RNN and Transformer systems is a lack of refinement of verbal usage, particularly when using the verbs “déan” (to do or to make) and “bí” (to be). As in many languages, the fact that these are possibly the two most universally used verbs in Irish further exacerbates the problem. An illustration of this problem can be seen in the output GA-1, which highlights the incorrect usage of “déan”. In a similar fashion, GA-5 demonstrates how the system misinterprets the usage of the verb “bí”, e.g., “ní bheidh said”.

6.3.4. Performance of RNN Approach Relative to Transformer Approach

There is a strong correlation between automatic and human evaluation of the translation systems that we developed. The automatic BLEU scores are contrasted with the HE scores for both the RNN and Transformer models in Table 12.

Table 12. Transformer approach compared to the RNN approach across all metrics for the DGT dataset. The results from our HE, using SQM and MQM metrics, validate the BLEU automatic evaluation results.

Approach	BLEU ↑	SQM ↑	MQM ↑
Transformer	60.5	4.53	77.92
RNN	52.7	3.30	43.05

6.4. Limitations of the Study

Certain aspects of this study could be further developed, given more time and resources. Although there is high inter-annotator agreement, it would help to have more annotators. In addition, the human evaluation of a greater number of lines, coupled with a more detailed MQM taxonomy, may provide greater insight into the MT outputs. This would help in uncovering other aspects, such as how gender is handled by the MT models.

7. Conclusions and Future Work

With this research, we have presented the first HE study that compares the output of EN-GA RNN systems with that of Transformer-based EN-GA systems. Automatic metrics were shown to differentiate the systems and highlighted that Transformer models are superior to RNN models. In our paper, we demonstrated that a random search approach to HPO enabled the development of high-performing translation models. We have shown there is a high level of correlation between an HE and an automatic approach. Both the automatic metrics and our HE demonstrated that the Transformer-based system is the most accurate.

The importance of selecting hyperparameters when training low-resource Transformer models was also demonstrated. By increasing dropout and reducing the number of hidden-layer neurons, our models performed significantly better than Google Translate and our baseline models.

We have demonstrated that choosing the correct subword models is an important performance driver for low-resource MT. Within the context of low-resource English-to-Irish translations, we achieved optimal performance on a 55k generic corpus when a Transformer architecture with a 16k BPE subword model was used. Improvements in the performance of our optimized Transformer models was observed across all key indicators, namely, PPL

was achieved at a lower global minimum, with a lower post-editing effort and a higher translation accuracy.

As part of future work, steps can be taken to deal with the inadequacies highlighted in our linguistic analysis. The issue of misusing common irregular verbs could be addressed by fine-tuning our models with a dataset specifically tailored for that purpose. In a similar fashion, fine-tuning after the careful selection of training data would also reduce the register errors encountered in our linguistic analysis. As it is difficult to train systems for all eventualities, using post-editing tools would be the best approach to correcting core grammatical errors involving pronouns, lenitions and lemmatization.

Author Contributions: Writing—original draft, S.L.; Writing—review & editing, H.A. and A.W. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by ADAPT, which is funded under the SFI Research Centres Programme (Grant 13/RC/2016) and is co-funded by the European Regional Development Fund. This research was also funded by the Munster Technological University.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The data presented in this study are openly available at <https://github.com/seamusl/isfeidirlinn> (accessed on 5 May 2022).

Acknowledgments: We would like to thank the annotators, Éamon Lankford and Máirín Lankford for their meticulous work in annotating the system outputs.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Glossary

Irish terms referenced and used in this manuscript:

Déan	To do or to make
Bí	To be
Ná bídis	Let it not be
Ní bheidh siad	They will not
Ní bheidh sé	He will not

References

1. He, D.; Xia, Y.; Qin, T.; Wang, L.; Yu, N.; Liu, T.Y.; Ma, W.Y. Dual learning for machine translation. In Proceedings of the Advances in Neural Information Processing Systems 29 (NIPS 2016), Barcelona, Spain, 5–10 December 2016; Volume 29.
2. Ahmadnia, B.; Dorr, B.J. Augmenting neural machine translation through round-trip training approach. *Open Comput. Sci.* **2019**, *9*, 268–278. [[CrossRef](#)]
3. Dowling, M.; Lynn, T.; Poncelas, A.; Way, A. SMT versus NMT: Preliminary comparisons for Irish. In Proceedings of the AMTA 2018 Workshop on Technologies for MT of Low Resource Languages (LoResMT 2018), Boston, MA, USA, 21 March 2018; pp. 12–20.
4. Ding, S.; Renduchintala, A.; Duh, K. A call for prudent choice of subword merge operations in neural machine translation. *arXiv* **2019**, arXiv:1905.10453.
5. Gowda, T.; May, J. Finding the optimal vocabulary size for neural machine translation. *arXiv* **2020**, arXiv:2004.02334.
6. Gage, P. A new algorithm for data compression. *C Users J.* **1994**, *12*, 23–38.
7. Sennrich, R.; Haddow, B.; Birch, A. Neural machine translation of rare words with subword units. *arXiv* **2015**, arXiv:1508.07909.
8. Kudo, T. Subword regularization: Improving neural network translation models with multiple subword candidates. *arXiv* **2018**, arXiv:1804.10959.
9. Lankford, S.; Alfi, H.; Way, A. Transformers for Low-Resource Languages: Is Féidir Linn! In Proceedings of the 18th Biennial Machine Translation Summit (Volume 1: Research Track), Virtual, 16–20 August 2021; pp. 48–60.
10. MacFarlane, A.; Glynn, L.G.; Mosinkie, P.I.; Murphy, A.W. Responses to language barriers in consultations with refugees and asylum seekers: A telephone survey of Irish general practitioners. *BMC Fam. Pract.* **2008**, *9*, 1–6. [[CrossRef](#)]
11. Alam, K.; Imran, S. The digital divide and social inclusion among refugee migrants. *Inf. Technol. People* **2015**, *28*, 344–365. [[CrossRef](#)]

12. Liu, C.H.; Silva, C.C.; Wang, L.; Way, A. Pivot machine translation using chinese as pivot language. In Proceedings of the China Workshop on Machine Translation, Wuyishan, China, 25–26 October 2018; Springer: Berlin/Heidelberg, Germany, 2018; pp. 74–85.
13. Ghifary, M.; Kleijn, W.B.; Zhang, M.; Balduzzi, D.; Li, W. Deep reconstruction-classification networks for unsupervised domain adaptation. In Proceedings of the European Conference on Computer Vision, Amsterdam, The Netherlands, 8–16 October 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 597–613.
14. Bahdanau, D.; Cho, K.; Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv* **2014**, arXiv:1409.0473.
15. Cho, K.; Van Merriënboer, B.; Bahdanau, D.; Bengio, Y. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv* **2014**, arXiv:1409.1259.
16. Bojar, O.; Chatterjee, R.; Federmann, C.; Graham, Y.; Haddow, B.; Huang, S.; Huck, M.; Koehn, P.; Liu, Q.; Logacheva, V.; et al. Findings of the 2017 Conference on Machine Translation (WMT17). In Proceedings of the Second Conference on Machine Translation, Copenhagen, Denmark, 7–11 September 2017; Association for Computational Linguistics: Copenhagen, Denmark, 2017; pp. 169–214. [[CrossRef](#)]
17. Bojar, O.; Federmann, C.; Fishel, M.; Graham, Y.; Haddow, B.; Koehn, P.; Monz, C. Findings of the 2018 Conference on Machine Translation (WMT18). In Proceedings of the Third Conference on Machine Translation: Shared Task Papers, Belgium, Brussels, 31 October–1 November 2018; Association for Computational Linguistics: Belgium, Brussels, 2018; pp. 272–303. [[CrossRef](#)]
18. Sanders, S.; Giraud-Carrier, C. Informing the use of hyperparameter optimization through metalearning. In Proceedings of the 2017 IEEE International Conference on Data Mining (ICDM), New Orleans, LA, USA, 18–21 November 2017; IEEE: New York, NY, USA, 2017; pp. 1051–1056.
19. Montgomery, D.C. *Design and Analysis of Experiments*; John Wiley & Sons: Hoboken, NJ, USA, 2017.
20. Bergstra, J.; Bengio, Y. Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **2012**, *13*, 281–305.
21. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. *arXiv* **2017**, arXiv:1706.03762.
22. Kudo, T.; Richardson, J. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv* **2018**, arXiv:1808.06226.
23. Belz, A.; Agarwal, S.; Graham, Y.; Reiter, E.; Shimorina, A. Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval). In Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval), Kiev, Ukraine, 19–20 April 2021; Association for Computational Linguistics (ACL): Stroudsburg, PA, USA, 2021.
24. Toral, A.; Castilho, S.; Hu, K.; Way, A. Attaining the unattainable? reassessing claims of human parity in neural machine translation. *arXiv* **2018**, arXiv:1808.10432.
25. Castilho, S.; Moorkens, J.; Gaspari, F.; Calixto, I.; Tinsley, J.; Way, A. Is neural machine translation the new state of the art? *Prague Bull. Math. Linguist.* **2017**, *108*, 109. [[CrossRef](#)]
26. Bayón, M.D.C.; Sánchez-Gijón, P. Evaluating machine translation in a low-resource language combination: Spanish-Galician. In Proceedings of the Machine Translation Summit XVII: Translator, Project and User Tracks, Dublin, Ireland, 19–23 August 2019; pp. 30–35.
27. Imankulova, A.; Dabre, R.; Fujita, A.; Imamura, K. Exploiting out-of-domain parallel data through multilingual transfer learning for low-resource neural machine translation. *arXiv* **2019**, arXiv:1907.03060.
28. Läubli, S.; Castilho, S.; Neubig, G.; Sennrich, R.; Shen, Q.; Toral, A. A set of recommendations for assessing human-machine parity in language translation. *J. Artif. Intell. Res.* **2020**, *67*, 653–672. [[CrossRef](#)]
29. Dowling, M.; Castilho, S.; Moorkens, J.; Lynn, T.; Way, A. A human evaluation of English-Irish statistical and neural machine translation. In Proceedings of the 22nd Annual Conference of the European Association for Machine Translation, Lisbon, Portugal, 3–5 November 2020; pp. 431–440.
30. Araabi, A.; Monz, C. Optimizing Transformer for Low-Resource Neural Machine Translation. *arXiv* **2020**, arXiv:2011.02266.
31. Van Biljon, E.; Pretorius, A.; Kreutzer, J. On optimal transformer depth for low-resource language translation. *arXiv* **2020**, arXiv:2004.04418.
32. Klubička, F.; Toral, A.; Sánchez-Cartagena, V.M. Quantitative fine-grained human evaluation of machine translation systems: A case study on English to Croatian. *Mach. Transl.* **2018**, *32*, 195–215. [[CrossRef](#)]
33. Freitag, M.; Foster, G.; Grangier, D.; Ratnakar, V.; Tan, Q.; Macherey, W. Experts, errors, and context: A large-scale study of human evaluation for machine translation. *Trans. Assoc. Comput. Linguist.* **2021**, *9*, 1460–1474. [[CrossRef](#)]
34. Ma, Q.; Graham, Y.; Wang, S.; Liu, Q. Blend: A novel combined MT metric based on direct assessment—CASICT-DCU submission to WMT17 metrics task. In Proceedings of the Second Conference on Machine Translation, Copenhagen, Denmark, 7–11 September 2017; pp. 598–603.
35. Lommel, A. Metrics for translation quality assessment: A case for standardising error typologies. In *Translation Quality Assessment*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 109–127.
36. Lommel, A.; Burchardt, A.; Popović, M.; Harris, K.; Avramidis, E.; Uszkoreit, H. Using a new analytic measure for the annotation and analysis of MT errors on real data. In Proceedings of the 17th Annual Conference of the European Association for Machine Translation, Dubrovnik, Croatia, 16–18 June 2014; pp. 165–172.
37. Callison-Burch, C.; Fordyce, C.S.; Koehn, P.; Monz, C.; Schroeder, J. (Meta-) evaluation of machine translation. In Proceedings of the Second Workshop on Statistical Machine Translation, Prague, Czech Republic, 23 June 2007; pp. 136–158.

38. Artstein, R. Inter-annotator agreement. In *Handbook of Linguistic Annotation*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 297–313.
39. Cohen, J. A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* **1960**, *20*, 37–46. [[CrossRef](#)]
40. Steinberger, R.; Eisele, A.; Klocek, S.; Pilos, S.; Schlüter, P. DGT-TM: A freely available translation memory in 22 languages. *arXiv* **2013**, arXiv:1309.5226.
41. Bisong, E. Google colabatory. In *Building Machine Learning and Deep Learning Models on Google Cloud Platform*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 59–64.
42. Klein, G.; Kim, Y.; Deng, Y.; Senellart, J.; Rush, A.M. Openmt: Open-source toolkit for neural machine translation. *arXiv* **2017**, arXiv:1701.02810.
43. Papineni, K.; Roukos, S.; Ward, T.; Zhu, W.J. Bleu: A method for automatic evaluation of machine translation. In Proceedings of the 40th annual meeting of the Association for Computational Linguistics, Philadelphia, PA, USA, 7–12 July 2002; pp. 311–318.
44. Snover, M.; Dorr, B.; Schwartz, R.; Micciulla, L.; Makhoul, J. A study of translation edit rate with targeted human annotation. In Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers, Cambridge, MA, USA, 8–12 August 2006; Citeseer: Forest Grove, OR, USA, 2006; Volume 200.
45. Popović, M. chrF: character n-gram F-score for automatic MT evaluation. In Proceedings of the Tenth Workshop on Statistical Machine Translation, Lisboa, Portugal, 17–18 September 2015; pp. 392–395.
46. Stanojević, M.; Kamran, A.; Koehn, P.; Bojar, O. Results of the WMT15 metrics shared task. In Proceedings of the Tenth Workshop on Statistical Machine Translation, Lisboa, Portugal, 17–18 September 2015; pp. 256–273.
47. Jooste, W.; Haque, R.; Way, A. Knowledge Distillation: A Method for Making Neural Machine Translation More Efficient. *Information* **2022**, *13*, 88. [[CrossRef](#)]
48. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual, 3–10 March 2021; pp. 610–623.
49. Lacoste, A.; Luccioni, A.; Schmidt, V.; Dandres, T. Quantifying the carbon emissions of machine learning. *arXiv* **2019**, arXiv:1910.09700.
50. SEAI. *Sustainable Energy in Ireland*; SEAI: Dublin, Ireland, 2020.
51. Ojha, A.K.; Liu, C.H.; Kann, K.; Ortega, J.; Shatam, S.; Fransen, T. Findings of the LoResMT 2021 Shared Task on COVID and Sign Language for Low-resource Languages. *arXiv* **2021**, arXiv:2108.06598.
52. Lankford, S.; Afli, H.; Way, A. Machine Translation in the Covid domain: An English-Irish case study for LoResMT 2021. In Proceedings of the 4th Workshop on Technologies for MT of Low Resource Languages (LoResMT2021), Virtual, 16 August 2021; pp. 144–150.
53. McHugh, M.L. Interrater reliability: The kappa statistic. *Biochem. Medica* **2012**, *22*, 276–282. [[CrossRef](#)]