

# Fine-grained Human Evaluation of Transformer and Recurrent Approaches to Neural Machine Translation for English-to-Chinese

**Yuying Ye**

Digital Humanities Programme  
University of Groningen  
The Netherlands  
y.ye.yuying@gmail.com

**Antonio Toral**

Center for Language and Cognition  
University of Groningen  
The Netherlands  
a.toral.ruiz@rug.nl

## Abstract

This research presents a fine-grained human evaluation to compare the Transformer and recurrent approaches to neural machine translation (MT), on the translation direction English-to-Chinese. To this end, we develop an error taxonomy compliant with the Multidimensional Quality Metrics (MQM) framework that is customised to the relevant phenomena of this translation direction. We then conduct an error annotation using this customised error taxonomy on the output of state-of-the-art recurrent- and Transformer-based MT systems on a subset of WMT2019’s news test set. The resulting annotation shows that, compared to the best recurrent system, the best Transformer system results in a 31% reduction of the total number of errors and it produced significantly less errors in 10 out of 22 error categories. We also note that two of the systems evaluated do not produce any error for a category that was relevant for this translation direction prior to the advent of NMT systems: Chinese classifiers.

## 1 Introduction

The field of machine translation (MT) has been revolutionised in the past few years by the emergence of a new approach: neural MT (NMT). NMT is a dynamic research area and we have witnessed two mainstream architectures already, the first of which is based on recurrent neural

networks (RNN) with attention (Bahdanau et al., 2014) while the second, referred to as Transformer, makes use of the self-attention mechanism in non-recurrent networks (Vaswani et al., 2017).

Several studies have analysed in depth, using both automatic and human evaluation methods, the resulting translations of NMT systems under the recurrent architecture and compared them to the translations of the previous mainstream approach to MT: statistical MT (Koehn et al., 2003), e.g. (Bentivogli et al., 2016; Castilho et al., 2017; Klubička et al., 2018; Popović, 2017; Shterionov et al., 2018). However, while the Transformer architecture has brought, at least when trained with sufficient data, considerable gains over the recurrent architecture (Vaswani et al., 2017), the research conducted to date that analyses the resulting translations of these two neural approaches is, to the best of our knowledge, limited to automatic approaches (Burlot et al., 2018; Lakew et al., 2018; Tang et al., 2018a; Tang et al., 2018b; Tran et al., 2018; Yang et al., 2019).

In this paper we conduct a detailed human analysis of the outputs produced by state-of-the-art recurrent and Transformer NMT systems. Namely, we manually annotate the errors found according to a detailed error taxonomy which is compliant with the hierarchical listing of issue types defined as part of the Multidimensional Quality Metrics (MQM) framework (Lommel et al., 2014). We carry out this analysis for the news domain in the English-to-Chinese translation direction. To this end, we define an error taxonomy that is relevant to the problematic linguistic phenomena of this translation direction. This taxonomy is then used to annotate errors produced by NMT systems that fall under the recurrent and Transformer architectures.

The main contributions of this paper can then be

summarised as follows:

1. We develop an MQM-compliant error taxonomy tailored to the English-to-Chinese translation direction.
2. We conduct, to the best of our knowledge, the first human fine-grained error analysis of Transformer-based versus recurrent NMT.

The rest of the paper is arranged in the following way. Section 2 presents a brief review of related work. Next, Section 3 outlines the recurrent- and Transformer-based NMT systems and the dataset used in our experiments. Subsequently, Section 4 presents the methodology for error annotation and the definition of the error taxonomy, followed by results and statistical analysis of the annotation. Finally, Section 5 gives a conclusion and suggestions for future work.

## 2 Related Work

This section provides an overview of related research on the two topics that correspond to our main contributions: human error analysis of MT outputs for the language pair English–Chinese (Section 2.1) and analyses of MT systems based on the recurrent and Transformer architectures (Section 2.2).

### 2.1 Human Error Analyses of MT for Chinese

One of the first taxonomies of MT errors, by Vilar et al. (2006), had a specific error typology for the Chinese-to-English translation direction, in accordance with the specific relevant phenomena of this language pair. Compared to their base taxonomy, a refined categorisation of word order was added to mark syntactic mistakes that appear in translations of questions, infinitives, declarative and subordinate sentences. In addition, the error type *Unknown words* was refined into four sub-types: *Person*, *Location*, *Organisation* and *Other proper names*.

Li et al. (2009) carried out an error analysis for the Chinese-to-Korean translation direction with only three categories from the taxonomy of Vilar et al. (2006) (*Missing words*, *Wrong word order* and *Incorrect words*), and they replaced *Incorrect words* with two more specific categories: one for both wrong lexical choices and extra words and

another for wrong modality. The simplified taxonomy was used to check if their method of re-ordering verb phrases, prepositional phrases and modality-bearing words in the Chinese data resulted in an improved MT system.

Hsu (2014) adapted the classification scheme of Farrús et al. (2010) to conduct an error analysis for the Chinese-to-English translation direction. The error taxonomy of Farrús et al. (2010) was originally defined for Catalan→Spanish. Its first level corresponded to five types of errors, related to different linguistic levels: orthographic, morphological, lexical, semantic and syntactic.

Castilho et al. (2017) assessed the output of two MT systems (statistical and recurrent) on patents, also for the Chinese-to-English translation direction. For this, they used a custom error taxonomy consisting of the error types *Punctuation*, *Part of speech*, *Omission*, *Addition*, *Wrong terminology*, *Literal translation*, and *Word form*.

Hassan et al. (2018) analysed the output of a Transformer-based MT system, again for the Chinese-to-English translation direction, using a two-level taxonomy based on that by Vilar et al. (2006). The first level contains nine error types: *Missing words*, *Word repetition*, *Named entity*, *Word order*, *Incorrect words*, *Unknown words*, *Collocation*, *Factoid*, and *Ungrammatical*. Only the error type *Named entity* has a second level, with five subcategories: *Person*, *Location*, *Organisation*, *Event*, and *Other*.

As we can observe in these related works, fine-grained human evaluation for the English–Chinese language pair has been hitherto conducted, to the best of our knowledge, (i) only for the Chinese-to-English direction and (ii) with error taxonomies that were either developed prior to the advent of the MQM framework or that were designed ad-hoc and were not thoroughly motivated. The position of our paper in these regards is thus clearly novel: (i) our analysis is for the English-to-Chinese translation direction and (ii) we devise and use an error taxonomy that is compliant with the MQM framework.

### 2.2 Analyses of Recurrent versus Transformer MT Systems

Tang et al. (2018a) compared recurrent- and Transformer-based MT systems on a syntactic task that involves long-range dependencies (subject-verb agreement) and on a semantic task (word

sense disambiguation) The recurrent system outperformed Transformer on the syntactic task while Transformer was better than the recurrent system on the semantic task. The latter finding was corroborated by Tang et al. (2018b).

Tran et al. (2018) compared the recurrent and Transformer architectures with respect to their ability to model hierarchical structure in a monolingual setting, by means of two tasks: subject-verb agreement and logical inference. On both tasks, the recurrent system outperformed Transformer, slightly but consistently.

Burlot et al. (2018) confronted English→Czech Transformer- and recurrent-based MT systems submitted to WMT2018<sup>1</sup> on a test suite that addresses morphological competence, based on the error typology by Burlot and Yvon (2017). The recurrent system outperformed Transformer on cases that involve number, gender and tense, while both architectures performed similarly on agreement. It is worth noting that agreement here regards local agreement (e.g. an adjective immediately followed by a noun), while the aforementioned cases of agreement in which a recurrent system outperforms Transformer (Tang et al., 2018a; Tran et al., 2018) regard long-distance agreement.

Yang et al. (2019) assessed the ability of both architectures to learn word order. When trained on a specific task related to word order, word reordering detection, a recurrent system outperformed Transformer. However, when trained on a downstream task, MT, Transformer was able to learn better positional information.

Lakew et al. (2018) evaluated multilingual NMT systems under the Transformer and recurrent architectures in terms of their morphological, lexical, and word order errors. In both architectures lexical errors were found to be the most prominent ones, followed by morphological, and lastly come reordering errors. The authors compared the number of errors in bilingual, multilingual and zero-shot systems, both for recurrent and Transformer, and found multilingual and zero-shot systems to be more competitive with respect to bilingual models for Transformer than for recurrent.

### 3 Machine Translation Systems

This section reports on the MT systems and the dataset used in our experiments.

<sup>1</sup><http://www.statmt.org/wmt18/>

We have used output from systems that fall under the recurrent and Transformer architectures and were top-ranked at the news translation shared task at the Conference on Machine Translation (WMT). We chose the University of Edinburgh’s MT system (Sennrich et al., 2017) as our recurrent NMT system due to the fact that this system had the highest BLEU score (36.3) for the translation direction English→Chinese at WMT2017<sup>2</sup> and it was ranked first (tied with other two systems) in the human evaluation.

As for the Transformer-based MT system used in our research, we have taken the PATECH submission to WMT2019.<sup>3</sup> We conducted our experiments before the human evaluation of WMT2019 was available, and therefore we chose the PATECH’s system based on the automatic evaluation of WMT2019, in which this system was the best performing one.<sup>4</sup> However, PATECH’s system was not included in the human evaluation of WMT2019. Therefore we carried out an additional annotation on the top-performing system from that human evaluation: the Transformer system developed by Kingsoft AI Lab (Guo et al., 2019), hereafter referred to as KSAI.

Before our human error analysis, we would like to compare the recurrent and Transformer MT systems in terms of an automatic evaluation metric. This is not possible from their outputs since they correspond to two different test sets (`newstest2017` and `newstest2019`). In order to be able to compare them, we asked the developer of the recurrent system to provide us with the output from their system for `newstest2019`. As shown in Table 1, the use of the Transformer architecture leads to a considerable improvement compared to the recurrent system (on average 31.4% relative in terms of BLEU). While the gap between the two architectures is large based on BLEU, this is an overall metric and therefore does not provide any insight into which aspects of the translation have improved with Transformer with respect to the recurrent system. To gain further insight we conduct a fine-grained human error analysis in the following section.

<sup>2</sup><http://www.statmt.org/wmt17/>

<sup>3</sup><http://matrix.statmt.org/systems/show/4243>

<sup>4</sup>[http://matrix.statmt.org/matrix/systems\\_list/1908](http://matrix.statmt.org/matrix/systems_list/1908)

RNN	Transformer (PATECH)	Transformer (KSAI)
33.1	44.6	42.4

**Table 1:** Automatic evaluation (BLEU scores) of the 3 MT systems on the WMT 2019 news test set.

## 4 Error Annotation

This section details the annotation setup (Section 4.1), explains how we defined our MQM-compliant error taxonomy adapted to the relevant characteristics of translating from English into Chinese and the challenges faced by NMT systems in this translation direction (Section 4.2) and presents the results of the annotation, as well as analysis and discussion thereof (Section 4.3).

### 4.1 Annotation Setup

We use `translate5`,<sup>5</sup> an open-source web-based tool, as the annotation environment. `translate5` was installed on a cloud server, so that it could be accessed remotely by annotators. The source text and reference translation are provided next to the NMT translations.

The annotation was performed by two annotators who are native Chinese speakers with fluent English. They both had an academic background and experience in translation. Prior to annotation, they were fully informed on the annotation environment and were provided with annotation instructions, comprising MQM’s usage guidelines and decision tree (Burchardt and Lommel, 2014).

The dataset used in our experiments is the test set from WMT2019 (`newstest2019`) for English→Chinese. This test set is chosen due to the fact that we have outputs for the RNN- and Transformer-based MT systems (see Section 3), and also because it is a commonly-used benchmark in the MT community. In our error annotation we use two subsets of this test set.

- A calibration set, made of the first 40 sentences from the testset. This refers to a small sample of annotation data that annotators work on before the actual annotation task takes place. Its purpose is twofold: (i) we use it to find out which error types occur in the translations and therefore use it to guide the refinement of the error taxonomy in a data-driven way; (ii) we also use it to identify disagreements between the annotators.

<sup>5</sup><http://www.translate5.net>

- An evaluation set, made up of 100 sentences from the test set. In order to have intersentential context, these sentences are taken from six documents (five full documents and the first sentences of the sixth document up to 100 sentences are reached). Using this evaluation set led then to the annotation of 500 sentences (100 distinct sentences times two MT systems (RNN and PATECH) times two annotators, plus the annotation of the 100 sentences for a third system (KSAI) by one annotator).

The annotators annotated the calibration set with our custom error taxonomy (see Figure 2), after which they discussed difficult cases and reached agreement on how to annotate them. Then they annotated the translations of the evaluation set. Once annotators started working on the evaluation set, they were not allowed to discuss problems in annotation any more.

### 4.2 Error Taxonomy


We decided to develop our error taxonomy based on the MQM framework developed at the QT-LaunchPad project (Lommel et al., 2014), after reviewing different translation quality evaluation frameworks. MQM stands out with its extensive standardised issue types<sup>6</sup> which are provided with clear definitions and explanations. In addition, a thorough guideline and decision tree<sup>7</sup> are available to assist annotators. Furthermore, this framework allows the building of customised error taxonomies.

Following the method of Klubička et al. (2018), our customisation process started with the sample MQM-compliant hierarchy for diagnostic MT evaluation (Figure 1) as the initial stage of our error taxonomy. The sample MQM tagset went through the preliminary selection of issue types to be used for fine-grained MT evaluation.


We annotated the calibration set with the sample MQM-compliant hierarchy to find out what types of errors occur in the outputs of our MT systems. Based on the results of the calibration set, we defined the complete tagset (shown in Figure 2). In the following subsections we provide detailed information concerning each of the modifications made to the error taxonomy.

<sup>6</sup><http://www.qt21.eu/mqm-definition/issues-list-2015-12-30.html>

<sup>7</sup><http://www.qt21.eu/downloads/MQM-usage-guidelines.pdf>



**Figure 1:** The sample MQM-compliant error hierarchy for diagnostic MT evaluation. The italicised issue types are not included in the standard MQM issue types.



**Figure 2:** The MQM-compliant error taxonomy for the translation direction English→Chinese. All the changes are marked by boxes with grey dotted lines and the issue types that are not included in the MQM issue types are italicised.

#### 4.2.1 Word Form & Spelling

Given that Chinese is an analytic language without inflection and its writing system is logographic, the issue types *Word form* and *Spelling* are of no interest to our research agenda.

#### 4.2.2 Classifier

We add one of the distinctive features of Chinese part-of-speech, the usage of classifiers, which have been researched thoroughly in Chinese linguistics (Jin, 2018) and Chinese language processing (Huang et al., 2017). In short, classifiers are special linguistic units located behind a number, demonstrative or certain quantifiers. These classifiers do not have a counterpart in English, which might give rise to translation problems. Examples of classifiers are shown in Table 2. How MT systems handle such a specific linguistic phenomenon is of interest to us.

Pronoun	Classifier	Noun
每( <i>mei</i> ) Every	<u>个</u> ( <i>ge</i> )	角落( <i>jiaoluo</i> ) corner
一( <i>yi</i> ) One	<u>架</u> ( <i>jia</i> )	飞机( <i>feiji</i> ) plane

**Table 2:** Examples of classifiers in Chinese. The classifiers are underscored.

#### 4.2.3 Typography

We extend the issue type *Typography* into two specific subtypes, based on the result of the calibration set. Though an unpaired quote or a misuse of punctuation is less likely to damage the comprehension of the content critically than other errors, as stated in Vilar et al. (2006), the Chinese→English error annotation conducted by Hsu (2014) shows that punctuation accounts for 10% of the errors. Such a high amount of punctuation mistakes could be a nuisance in the MT output. Incorrect usage of *Typography* could negatively influence the reception of a translation, since the reader might consider such an error as a sign of lack of professionalism, and therefore react by distrusting the content.

#### 4.2.4 Mistranslation

Preliminarily, we observe that *Mistranslation* is a major issue in the calibration set and that related translation errors *Overly-literal* and *Entity* appear frequently. We have thus decided to specify them as sub-types of *Mistranslation*. Vilar et al. (2006) also included entity errors in their error typology for the Chinese→English language pair and further divided them into specific sub-types. As their result showed, this issue type only amounted to a small percentage of the errors. Therefore, the is-

sue type *Entity* is not further specified in our taxonomy.

#### 4.2.5 Function word

*Function word* is extended to one extra layer under *Extraneous* with the intention of covering westernised Chinese expressions that were observed in the calibration set. Westernised Chinese refers to a cross-lingual phenomenon of imposing English grammar on Chinese, which is manifested in many problematic forms, abuse of function words especially (Tse, 2001). The relations between sentence parts, tenses and aspects are often shown through word order, particles or context in Chinese, due to its lack of inflection. Specifying the types of extraneous function words into three common types, *Preposition*, *Adverb* and *Particle* could be useful to discuss whether there is difference among these word classes.

The two other sub-types of *Function word* (*Incorrect* and *Missing*) are not specified in conformity with the initial examination of the data. Not only might adding the extra layer for both sub-error types not prove practical, but it is also not advised by the MQM guidelines to have the error taxonomy so big that it could challenge annotators’ memory limit (Burchardt and Lommel, 2014).

### 4.3 Results and Discussion

#### 4.3.1 Inter-annotator Agreement

Inter-annotator agreement (IAA) was calculated with Cohen’s Kappa ( $\kappa$ ) (Cohen, 1960) on the annotations of the calibration and evaluation sets for the RNN and PATECH’s Transformer systems (Table 3). It is worth noting that the IAA values of the evaluation set improve considerably upon those of the calibration set ( $\kappa = 0.44$  versus 0.27). It shows that the discussion of annotation disagreements can contribute to improving the level of agreement notably.

IAA	RNN	Transformer (PATECH)	Both
Calibration set	0.31	0.22	0.27
Evaluation set	0.45	0.43	0.44

**Table 3:** Total and average inter-annotator agreement (Cohen’s  $\kappa$  values) for the MQM calibration set and evaluation set.

As shown in Table 3, the difference of IAA scores between Transformer and RNN is slight in our evaluation set. The average IAA value

(0.44), corresponds to moderate agreement, according to Cohen (1960). When interpreting these results, it should be taken into account that IAA scores are known to be low in human evaluation of MT. For example, Callison-Burch et al. (2007) observed fair agreements for fluency and accuracy for eight language pairs, and, though the MQM framework is rigorously defined and supported by clear guidelines, in the experiments by Lommel and Burchardt (2014) MQM led to relatively low IAA, due to span-level difference, ambiguous categorisation and differences of opinion. Klubička et al. (2018) reported a moderate agreement on English–Croatian, higher than that by Lommel and Burchardt (2014), probably because the agreement was calculated on errors annotated for each sentence, thus not taking the spans of the annotations into account. Our own IAA results do not differ greatly with aforementioned research.

	RNN	Transformer (PATECH)	Both
Accuracy	0.60	<b>0.61</b>	<b>0.61</b>
Mistranslation	0.50	0.52	0.51
Entity	-0.03	0.39	0.18
Overly-literal	0.24	0.21	0.23
Omission	0.52	<b>0.67</b>	0.60
Addition	0.37	0.00	0.19
Untranslated	<b>0.73</b>	<b>0.71</b>	<b>0.72</b>
Fluency	0.01	0.07	0.04
Grammar	0.36	0.24	0.30
Function word	0.17	-0.01	0.08
Extraneous	0.32	-0.01	0.16
Preposition	<b>0.65</b>	-0.01	0.32
Adverb	0.00	N/A	N/A
Particle	-0.02	-0.03	-0.03
Incorrect	-0.02	-0.01	-0.02
Missing	0.32	0.00	0.16
Word order	0.45	0.29	0.37
Classifier	N/A	N/A	N/A
Unintelligible	0.20	-0.02	0.09
Typography	0.22	0.28	0.25
Punctuation	0.21	0.29	0.25
Unpaired-mark	N/A	N/A	N/A

**Table 4:** Inter-annotator agreement (Cohen’s  $\kappa$  values) on the evaluation set for the RNN and PATECH’s Transformer systems and their average. Substantial scores (0.61–0.80) are shown in bold. N/A is given to the error categories that were never used, since no data points could be used to calculate the IAA score.

In addition to overall IAA, Cohen’s ( $\kappa$ ) was also calculated for each issue type in the evaluation set

individually (Table 4). For both systems, the IAA scores for *Accuracy* and its sub-types are considerably higher than those under *Fluency*. It is an expected result taken into account that accuracy errors are more straightforward and less open to interpretation. The  $\kappa$  values are relatively consistent between Transformer and RNN, except a striking plunge in agreement scores for Transformer in some categories (*Function word* and its subtypes, *Word order* and *Unintelligible*) and the opposite, a considerably lower agreement for RNN, for *Entity*.

The source of these disagreements can be traced back to the annotation output. For example, in the case of *Unintelligible*, the evaluators annotated different sentences with this error category. As for *Entity*, it is worth mentioning that disagreement arose over this category in the annotation of the calibration set. It seems that, despite the discussion, the understanding of entity was still not shared by the two annotators. It is also possible that due to the improved translation quality of Transformer, mistakes such as *Function word* are more subtle and harder to detect.

#### 4.3.2 Annotated Errors

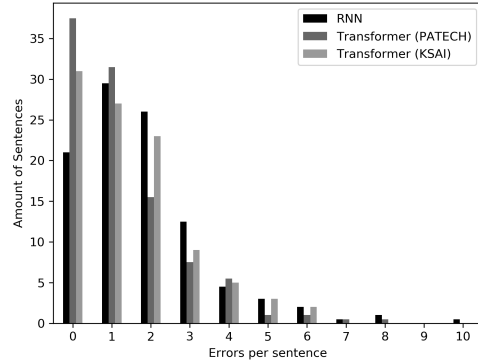
Table 5 presents the overall number of annotated error tags in the output of each system by each annotator. One can clearly observe that both annotators have annotated relatively less errors in Transformer’s output (PATECH) than in RNN’s; the error reduction is of 35% in the case of annotator 1 and of 27% in the case of the second annotator. The Transformer system from KSAI only reduces the number of errors by 12.5%, compared to the RNN system.

System	RNN	Transformer (PATECH)	Transformer (KSAI)
Annotator 1	168	109	147
Annotator 2	193	141	

**Table 5:** Total amounts of error per annotator and system, as annotated in MQM.

To delve deeper into the error distribution, we plot a histogram to show how many errors appear in each sentence and how many of these sentences are there in the output from each system. The mean of both annotators’ annotations for the first two systems are used, amounting to 100 sentences per system. The histogram is shown in Figure 3. It can be observed that more than 35 sentences in the Transformer (PATECH) output are not annotated with any error while only slightly

over 20 sentences in the RNN are marked as errorless. The two systems have similar amount of sentences with one mistake, while PATECH’s output contains considerably less sentences than RNN with more than one error.



**Figure 3:** Error distribution per system. For RNN and Transformer (PATECH), the average of annotation data from both annotators has been used.

We can also see notable differences between the two Transformer systems in Figure 3. Fewer sentences in the KSAI’s output are annotated without error, while considerably more sentences are tagged with two errors in this output than in PATECH’s system.

While comparing the systems in terms of their total number of errors gives us a clear indication of their relative performance, we note that a fairer comparison should take into account their outputs’ lengths. To that end, we make use of the normalisation approach proposed by Klubička et al. (2018): tokens annotated with errors are counted for each system’s output and they are then used to compute each system’s error ratio, which equals to the total number of erroneous tokens (Chinese characters) divided by the total number of tokens in the system’s output. This error ratio can serve then as a general score for each system. We also apply the same normalisation procedure to each issue type. Statistical significance for the total amount of errors and each issue type is computed with a pairwise chi-squared ( $\chi^2$ ) test (Plackett, 1983), following its application to normalised MQM errors introduced by Klubička et al. (2018).

Table 6 shows the error ratios (both overall and for each issue type) for each system, together with an indication of whether there are significant differences between each pair of systems. In terms of total error ratio, compared to RNN, the error reduc-

tion by PATECH amounts to 34% relative (11.85% versus 17.93%) and is significant ( $p < 0.001$ ). No significant difference is observed between the two Transformer-based systems.

For nearly half of the error types, the decrease in error ratio for Transformer (PATECH), compared to RNN, is statistically significant. For example, the number of tokens with *Fluency* errors decreased by 45% (6.45% verse 3.56%,  $p < 0.001$ ). The reduction is particularly notable for its child category *Unintelligible*, for which the number of erroneous tokens decreased by 55% (2.1% verse 0.93%,  $p < 0.001$ ). This Transformer-based system also managed to generate significantly less extraneous *Function words*, gaining a decrease of 47% (0.51% verse 0.27%,  $p < 0.05$ ). In addition, Transformer manages to produce significantly less extraneous *Prepositions*, (0.2% verse 0.04%,  $p < 0.05$ ). Though it also produces less *Overly-literal* translations (1.20% verse 0.86%) and no extraneous *Adverb* (0.06% verse 0%), these differences are not significant.

Conversely, this Transformer-based system underperforms on *Punctuation* (0.2% verse 0.37%), although the difference is not significant. By tracing this back to the annotation, we can observe that Transformer (PATECH) produces several cases of missing, wrong or redundant punctuation marks. For example, in one instance an English period (.) was used instead of a Chinese full stop (。). This Transformer system also had issues with adding guillemets (《》) around newspaper names and putting commas after adverbials, which are required in Chinese grammar.

Between the two Transformer systems, we can see that except for the category *Entity* and *Untranslated*, the two Transformer systems do not produce statistically significant different amount of errors. It proves that there are few significant discrepancies between these two systems.

Finally, we note that the error category *Unpaired-mark*, has not been used by any of the annotators for any of the three MT systems and the category *Classifier* has only been used to annotate 6 tokens (0.16%) in the third system’s output. While these categories were relevant in MT in the past (see Section 2), our results seem to indicate that they can be considered to have been solved by NMT.

	RNN	Transformer (PATECH)	Transformer (KSAI)
Accuracy	<b>11.48</b>	8.29**	7.41
Mistranslation	<b>7.49</b>	4.50**	4.39
Entity	0.24	0.23	<b>0.59*</b>
Overly-literal	1.20	0.86	0.51
Omission	<b>0.61</b>	0.33**	0.35
Addition	0.23	0.19	0.22
Untranslated	3.16	<b>3.27</b>	2.45*
Fluency	<b>6.45</b>	3.56**	3.02
Grammar	<b>3.08</b>	1.83**	2.24
Function word	<b>0.51</b>	0.27**	0.40
Extraneous	<b>0.35</b>	0.12**	0.30
Preposition	<b>0.20</b>	0.04**	0.13
Adverb	0.06	0	0.05
Particle	0.07	0.08	0.08
Incorrect	0.06	0.08	0
Missing	0.10	0.07	0.11
Word order	<b>2.32</b>	1.41**	1.46
Classifier	0	0	0.16
Unintelligible	<b>2.10</b>	0.93**	0
Typography	0.20	0.37	0.59
Punctuation	0.20	0.37	0.59
Unpaired-mark	0	0	0
Total error ratio	<b>17.93</b>	11.85**	10.40

**Table 6:** Error ratio (%) for each error type and overall. The annotations on RNN and Transformer (PATECH) from both annotators are concatenated. \* indicates  $p$ -value  $< 0.05$  and \*\*  $p$ -value  $< 0.001$ , when a system is compared to the system adjacent to its left side. Numbers shown in bold indicate that the system has significantly more erroneous tokens in the pair comparison.

## 5 Conclusion

This paper presented a fine-grained manual evaluation for English→Chinese on the two mainstream architectures of NMT: RNN and Transformer. The evaluation was approached in the form of a human error annotation based on a customised MQM error taxonomy.

The error taxonomy was developed from the MQM core taxonomy for MT evaluation. Chinese linguistic features and issues emerged in the calibration set were taken into account by including customised error types, such as *Extraneous function word*, *Classifier* and *Typography*. The error type *Extraneous function word* underpins investigating westernised Chinese phenomena of extraneous function words by specifying it into three word



classes: *Preposition*, *Adverb* and *Particle*.

From our analysis, it is clear that Transformer-based systems generate significantly more accurate, fluent and comprehensible translation with less westernised Chinese expressions. However, Transformer systems do not handle typography as well as RNN. We also note that none of the MT systems did produce any errors related to unpaired-marks and only one system produced errors related to classifiers, which were very unfrequent (0.16% of the tokens). We can conclude that Transformer systems produce an overall better translation compared to RNN when translating from English to Chinese, which corroborates findings of prior studies on other language pairs. A limitation worth mentioning is that our annotation was conducted by only two annotators on a limited amount of data.

Our taxonomy could be of use for further error analysis on Chinese MT quality. Future research could include a larger annotation sample to investigate if punctuation is a common issue in NMT systems based on Transformer and to verify that NMT is able to produce correct classifiers. Also, as Transformer still shows a major problem in mistranslation, the error taxonomy can be extended with more specific categories to explore this issue in more detail.

The annotations for the three MT systems and the code used for the analysis thereof are publicly available.<sup>8</sup>

## Acknowledgements

We would like to thank Rico Sennrich, for translating the test set used in this paper with a system he had co-developed for a previous edition of the WMT news translation shared task, and Filip Klubička, for providing us with the code to perform the statistical analysis of MQM output.

## References

- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of International Conference on Learning Representations 2015*, pages 1–15, San Diego, CA, USA, September.
- Bentivogli, Luisa, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. Neural versus Phrase-Based Machine Translation Quality: a Case Study.

In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 257–267, Austin, Texas. Association for Computational Linguistics.

- Burchardt, Aljoscha and Arle Lommel, 2014. *Practical Guidelines for the Use of MQM in Scientific Research on Translation Quality*.
- Burlot, Franck and François Yvon. 2017. Evaluating the morphological competence of machine translation systems. In *Proceedings of the Second Conference on Machine Translation*, pages 43–55, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Burlot, Franck, Yves Scherrer, Vinit Ravishankar, Ondřej Bojar, Stig-Arne Grönroos, Maarit Koponen, Tommi Nieminen, and François Yvon. 2018. The WMT’18 morphoEval test suites for English-Czech, English-German, English-Finnish and Turkish-English. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 546–560, Belgium, Brussels, October. Association for Computational Linguistics.
- Callison-Burch, Chris, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. (Meta-) evaluation of machine translation. In *Proceedings of the Second Workshop on Statistical Machine Translation - StatMT ’07*, pages 136–158, Prague, Czech Republic. Association for Computational Linguistics.
- Castilho, Sheila, Joss Moorkens, Federico Gaspari, Iacer Calixto, John Tinsley, and Andy Way. 2017. Is Neural Machine Translation the New State of the Art? *The Prague Bulletin of Mathematical Linguistics*, 108(1):109–120.
- Cohen, Jacob. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Farrús, Mireia, Marta Costa-jussa, Jose Bernardo Mariño Acebal, and José Fonollosa. 2010. Linguistic-based evaluation criteria to identify statistical machine translation errors. In *Proceedings of the 13th Annual Conference of the EAMT*, pages 52–57, Barcelona, Spain, 01.
- Guo, Xinze, Chang Liu, Xiaolong Li, Yiran Wang, Guoliang Li, Feng Wang, Zhitao Xu, Liuyi Yang, Li Ma, and Changliang Li. 2019. Kingsoft’s neural machine translation system for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 196–202, Florence, Italy, August. Association for Computational Linguistics.
- Hassan, Hany, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide,

<sup>8</sup><https://github.com/yy-ye/mqm-analysis>

- Xu Tan, Fei Tian, Lijun Wu, and Ming Zhou. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv:1803.05567 [cs]*, pages 1–25, 03.
- Hsu, John D. 2014. Error Classification of Machine Translation A Corpus-based Study on Chinese-English Patent Translation. *Translation Studies Quarterly*, 18:121–136.
- Huang, Chu-Ren, Shu-Kai Hsieh, Keh-Jiann Chen, Shu-Kai Hsieh, and Keh-Jiann Chen. 2017. *Mandarin Chinese Words and Parts of Speech : A Corpus-based Study*. Routledge, 7.
- Jin, Jing. 2018. *Partition and Quantity : Numeral Classifiers, Measurement, and Partitive Constructions in Mandarin Chinese*. Routledge, 06.
- Klubička, Filip, Antonio Toral, and Víctor M. Sánchez-Cartagena. 2018. Quantitative Fine-Grained Human Evaluation of Machine Translation Systems: a Case Study on English to Croatian. *arXiv e-prints*, 32(3):195–215.
- Koehn, Philipp, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1*, NAACL '03, pages 48–54, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Lakew, Surafel Melaku, Mauro Cettolo, and Marcello Federico. 2018. A comparison of transformer and recurrent neural networks on multilingual neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 641–652, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Li, Jin-Ji, Jungi Kim, Dong-Il Kim, and Jong-Hyeok Lee. 2009. Chinese Syntactic Reordering for Adequate Generation of Korean Verbal Phrases in Chinese-to-Korean SMT. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 190–196. Association for Computational Linguistics.
- Lommel, Arle and Aljoscha Burchardt. 2014. Assessing inter-annotator agreement for translation error annotation. In *MTE: Workshop on Automatic and Manual Metrics for Operational Translation Evaluation*, Reykjavik, Iceland.
- Lommel, Arle, Hans Uszkoreit, and Aljoscha Burchardt. 2014. Multidimensional quality metrics (mqm): A framework for declaring and describing translation quality metrics. *Tradumàtica: tecnologies de la traducció*.
- Plackett, R. L. 1983. Karl pearson and the chi-squared test. *International Statistical Review / Revue Internationale de Statistique*, 51(1):59–72.
- Popović, Maja. 2017. Comparing language related issues for nmt and pbmt between german and english. *The Prague Bulletin of Mathematical Linguistics*, 108(1):209–220.
- Sennrich, Rico, Alexandra Birch, Anna Currey, Ulrich Germann, Barry Haddow, Kenneth Heafield, Antonio Valerio Miceli Barone, and Philip Williams. 2017. The University of Edinburgh’s Neural MT Systems for WMT17. In *Proceedings of the Second Conference on Machine Translation*, pages 389–399, Copenhagen, Denmark, 09. Association for Computational Linguistics.
- Shterionov, Dimitar, Riccardo Superbo, Pat Nagle, Laura Casanellas, Tony O’ Dowd, and Andy Way. 2018. Human versus automatic quality evaluation of nmt and pbsmt. *Machine Translation*, 32(3):217–235.
- Tang, Gongbo, Mathias Müller, Annette Rios, and Rico Sennrich. 2018a. Why self-attention? a targeted evaluation of neural machine translation architectures. *arXiv preprint arXiv:1808.08946*.
- Tang, Gongbo, Rico Sennrich, and Joakim Nivre. 2018b. An analysis of attention mechanisms: The case of word sense disambiguation in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 26–35, Brussels, Belgium, October. Association for Computational Linguistics.
- Tran, Ke, Arianna Bisazza, and Christof Monz. 2018. The importance of being recurrent for modeling hierarchical structure. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4731–4736, Brussels, Belgium. Association for Computational Linguistics.
- Tse, Yiu Kay 谢耀基. 2001. ”hanyu yufan ouhua zongshu”汉语语法欧化综述[a review on the westernised chineses grammar]. *语文研究*, 1:17–22.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. [abs/1706.03762](https://arxiv.org/abs/1706.03762).
- Vilar, David, Jia Xu, Luis Fernando D’Haro, and Hermann Ney. 2006. Error Analysis of Statistical Machine Translation Output. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, pages 697–702.
- Yang, Baosong, Longyue Wang, Derek F. Wong, Lidia S. Chao, and Zhaopeng Tu. 2019. Assessing the ability of self-attention networks to learn word order. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3635–3644, Florence, Italy, July. Association for Computational Linguistics.