

Exploring Gap Filling as a Cheaper Alternative to Reading Comprehension Questionnaires when Evaluating Machine Translation for Gisting

Mikel L. Forcada

Dept. Lleng. i Sist. Inf.,
Universitat d'Alacant
E-03690 St. Vicent del Raspeig,
Spain
mlf@ua.es

**Carolina Scarton,
Lucia Specia**

Dept. of Comput. Sci.,
Univ. of Sheffield,
Sheffield S1 4DP, U.K
{l.specia, c.scarton}
@shef.ac.uk

**Barry Haddow,
Alexandra Birch**

School of Computing,
Univ. Edinburgh,
Edinburgh EH8 4AB, U.K
{bhaddow, a.birch}
@ed.ac.uk

Abstract

A popular application of machine translation (MT) is *gisting*: MT is consumed *as is* to make sense of text in a foreign language. Evaluation of the usefulness of MT for *gisting* is surprisingly uncommon. The classical method uses *reading comprehension questionnaires* (RCQ), in which informants are asked to answer professionally-written questions in their language about a foreign text that has been machine-translated into their language. Recently, *gap-filling* (GF), a form of *cloze* testing, has been proposed as a cheaper alternative to RCQ. In GF, certain words are removed from reference translations and readers are asked to fill the gaps left using the machine-translated text as a hint. This paper reports, for the first time, a comparative evaluation, using both RCQ and GF, of translations from multiple MT systems for the same foreign texts, and a systematic study on the effect of variables such as gap density, gap-selection strategies, and document context in GF. The main findings of the study are: (a) both RCQ and GF clearly identify MT to be useful; (b) global RCQ and GF rankings for the MT systems are mostly in agreement; (c) GF scores vary very widely across informants, making comparisons among MT systems hard, and (d) unlike RCQ, which is framed around documents, GF evaluation can be framed at the sentence level. These findings support the use of GF as a cheaper alternative to RCQ.

1 Introduction

1.1 Machine translation for gisting

Machine translation (MT) applications fall in two main groups: *assimilation* or *gisting*, and *dissemination*. Assimilation refers to the use of the raw MT output to make sense of foreign texts. Dissemination refers to the use of the MT output as a draft translation that can be *post-edited* into a

publishable translation. The needs of both groups of applications are quite different; for instance, an otherwise *perfect* Russian to English translation but with no articles (*some, a, the*), is likely to be fine for assimilation, but would need substantial post-editing for dissemination. State-of-the-art MT systems are however usually evaluated—even if manually—(and optimized) with respect to their ability to produce translations that resemble references, regardless of the intended application for the system.

Assimilation is by far the main use of MT in number of words translated. It is either explicitly invoked, for instance, by visiting webpages such as Google Translate, or integrated into browsers and social networks. Raw MT may sometimes be the only feasible option,¹ for instance when dealing with user-generated content or ephemeral material (such as product descriptions in e-commerce).

1.2 Evaluation of MT for gisting

A straightforward (but costly) way to evaluate MT for *gisting* *measures* the performance of target-language readers in a text-mediated task—for instance, a software installation task (Castilho et al., 2014)—by using raw MT and compares it with the performance reached using a professional translation of the text.

However, there may be scenarios without an obvious associated task: news, product and service reviews, or literature. On the other hand, even with a clear associated task, task completion evaluation is also quite expensive. It is therefore desirable to have alternative objective indicators which work as good surrogates for actual task-oriented

¹Twenty-five years ago, (Sager, 1993, p. 261) already hinted at MT-only scenarios: “there may, indeed, be no single situation in which either human or machine would be equally suitable.”

success.

Some authors have proposed eye-tracking (Doherty and O'Brien, 2009; Doherty et al., 2010; Stymne et al., 2012; Doherty and O'Brien, 2014; Castilho et al., 2014; Klerke et al., 2015; Castilho and O'Brien, 2016; Sajjad et al., 2016) as a measure of machine translation usefulness, but the technique is expensive and the evidence gathered is rather indirect and does not have a straightforward interpretation in terms of usefulness.

There are many methods in which informants are asked to *judge* the *quality* of machine-translated sentences, usually as regards their monolingual *fluency* (nativeness, grammaticality), their bilingual *adequacy* (how much of the information in the source sentence is present in the machine-translated sentence), or even monolingual adequacy (how much of the information in the reference sentence is present in the machine-translated sentence); informants may be asked either to *directly assess* MT outputs by giving values to these indicators in a predetermined scale or to *rank* a number of MT outputs for the same source sentence (sometimes being asked to consider aspects such as adequacy, fluency, or both). Direct assessments of adequacy and MT ranking are the official evaluation procedure for the most recent WMT translation shared task campaigns (Bojar et al., 2016, 2017). Other researchers use post-task questionnaires (Stymne et al., 2012; Doherty and O'Brien, 2014; Klerke et al., 2015; Castilho and O'Brien, 2016) to assess the perceived usefulness of MT output.

Direct assessment, ranking or post-task questionnaire evaluation methods are clearly subjective and require informants to make “in vitro” *judgements* about the *quality* of MT outputs, without considering their *usefulness* for a specific “in vivo”, real-world application.

1.3 Reading comprehension questionnaires

Reading comprehension questionnaires (RCQ), as used in the assessment of foreign-language learning, are the standard approach to evaluate MT for gisting that measures reader performance in response to MT. Readers answer questions using either a machine-translated or a professionally-translated version of the source text and their performance on the tests (i.e. to what extent they answer questions correctly) using the two sets of texts is then compared. RCQ are however quite

costly: a human translation is needed for a control group and questions need to be professionally written and often manually marked.

RCQ has a long history as an MT evaluation method. Tomita et al. (1993), Fuji (1999), and Fuji et al. (2001) evaluate the *informativeness* or *usefulness* of English–Japanese MT by using standardized English-as-a-foreign-language RCQs (TOEFL, TOEIC) which have been machine translated into Japanese and they are sometimes capable of distinguishing MT systems. Jones et al. (2005b), Jones et al. (2005a), Jones et al. (2007), and Jones et al. (2009) use the structure of standardized language proficiency tests (Defence Language Proficiency Test, Interagency Language Roundtable) to evaluate the readability of Arabic–English MT texts. MT’ed documents are found to be harder to understand than professional translations, and that they may be assigned an intermediate level of English proficiency. Berka et al. (2011) collected a set of English short paragraphs in various domains, created yes/no questions in Czech about them, and machine translated the English paragraphs into Czech with different MT systems. They found that outputs produced by different MT systems lead to different accuracy in the annotators’ answers. Weiss and Ahrenberg (2012) evaluate comprehension of Polish–English translations using RCQ tests and found that a text with more MT errors have less correct answers than a text with fewer MT errors. Finally, Stymne et al. (2012) use RCQ to validate eye-tracking as a tool for MT error analysis for English–Swedish. Interestingly, for one of their systems, the number of correct answers in the RCQ tests were higher than for the human translation. However, test takers were more confident in answering questions about the human translations than about the MT outputs.

In this paper we explore RCQ as a measure of MT quality by using the CREG-mt-eval corpus (Scarton and Specia, 2016). In contrast to previous work, this paper presents an evaluation of MT quality based on open questions that have different levels of difficulty (as presented in Section 2) for a considerable amount of documents (36 in contrast to only 2 analysed by Weiss and Ahrenberg (2012)).

1.4 An alternative: evaluation via gap-filling

An alternative approach to RCQs, *gap filling* (GF), has been recently proposed (Trosterud and Unhammer, 2012; O’Regan and Forcada, 2013; Ageeva et al., 2015; Jordan-Núñez et al., 2017) based on another typical way of measuring reading comprehension: *cloze* (or *closure*) testing (Taylor, 1953). Instead of a question, readers get an incomplete sentence with one or more words replaced by gaps, and are asked to fill the gaps. Indeed, GF may be seen as equivalent to the answering of simple reading comprehension questions: for instance, a question like *Who was the president of the Green Party in 2011?* would be equivalent to the sentence with one gap *In 2011, _____ was the president of the Green Party.*

GF tasks are prepared by automatically punching gaps in *reference* sentences taken from a professional translation of the source text. Informants are given the machine-translated sentence as a “hint” for the gap-filling task; therefore, we may view GF as a way of automatically generating questions to evaluate the MT output. The evaluation measure is the proportion of gaps that can be successfully filled using MT as a hint. This can be compared with the success rate in the case where no hint (MT) is provided, to give an estimate of the usefulness of MT output.

Note that *cloze* testing evaluation of machine translation was attempted decades ago in a completely different *readability* setting: gaps were then punched in *machine-translated output* and informants tried to complete them without any further hint (Crook and Bishop, 1965; Sinaiko and Klare, 1972). This work was reviewed and extended later by Somers and Wild (2000). But filling gaps in machine-translated output may be unnecessarily challenging and therefore make evaluation less adequate: for instance, informants would sometimes have to fill gaps in disfluent or ungrammatical text, which is much harder than filling them in a fluent, professionally translated reference, or, even in fluent output, a crucial content word that has been removed may be very hard to guess unless the surrounding text is very redundant. Moreover, the GF method described here has an easier interpretation in terms of its analogy to RCQ.

This paper systematically builds upon previous work on GF to obtain experimental evidence that gap-filling is a viable, lower-cost alternative to

RCQ evaluation. Its main **contributions** are:

- While Trosterud and Unhammer (2012), O’Regan and Forcada (2013), and Ageeva et al. (2015) used GF just to demonstrate the usefulness of a single rule-based MT system for each language pair studied, this paper, like Jordan et al.’s (2017), performs a comparison of several MT systems for the same language pair.
- Previous work (Trosterud and Unhammer, 2012; O’Regan and Forcada, 2013; Ageeva et al., 2015; Jordan-Núñez et al., 2017) simply assumes the validity of GF as an evaluation method for MT gisting, in some cases arguing about its equivalence to RCQ. Ours is the first work to actually compare GF and RCQ evaluation of the same MT systems.
- Previous work used sentences (Trosterud and Unhammer, 2012; O’Regan and Forcada, 2013; Ageeva et al., 2015) or short excerpts of text (Jordan-Núñez et al., 2017), but did not study the influence of a larger, document-level machine-translated context around the target sentence, as it is done here.
- This paper explores for the first time a gap-positioning strategy based on an approximate computation of gap entropy, and compares it to random placing of gaps.

The paper is organized as follows: section 2 describes the design and implementation of both evaluation methods, RCQ and GF; then section 3 reports and discusses the results obtained; and, finally, concluding remarks (section 4) close the paper.

2 Methodology

2.1 Data and informants

We use an extended version of CREG-mt-eval (Scarton and Specia, 2016), a version of the expert-built CREG reading comprehension corpus (Ott et al., 2012) for 2nd-language learners of German. CREG was originally created to build and evaluate systems that automatically correct answers to open questions. CREG-mt-eval contains 108 source (German) documents with different domains, including literature, news, job adverts, and others (on average 372 words and 33 sentences per document). The original documents

were machine-translated in December 2015 into English using four systems: an in-house baseline² statistical phrase-based Moses (Koehn et al., 2007) system trained on WMT 2015 data (Bojar et al., 2015), Google Translate,³ Bing⁴ and Systran.⁵ CREG-mt-eval also contains professional translations of a subset of 36 documents (90–1500 words) as a control group to check whether the questions are adequate for the task. All questions from the CREG original questionnaires (in German) were professionally translated to English. On average, there are 8.8 questions per document.

The questions in CREG-mt-eval are classified (Meurers et al., 2011) as: *literal*, when they can be answered directly from the text and refer to explicit knowledge, such as names, dates (79% of the total number of questions); *reorganization*, also based on literal text understanding, but requiring the combination of information from different parts of the text (12% of the total number of questions); and *inference*, which involve combining literal information with world knowledge (9% of the total number of questions).

Following Scarton and Specia (2016), test takers (informants) for both GF and RCQ were fluent English-speaking volunteers, staff and students at the University of Sheffield, who were paid (with a 10 GBP online gift certificate) to complete the task.

2.2 Reading comprehension questionnaire task

For the version of CREG-mt-eval used herein, thirty informants were given a set of six documents each and answered three to five questions per document, using only the English document (either machine- or human-translated) provided. Therefore, for each of the 36 original documents, questions were answered using each machine translation system or the human translation. Each document was only evaluated by one informant. The original German document was not given. The guidelines were similar to those used in other reading comprehension tests: test takers were asked to answer the questions based on the

²<http://www.statmt.org/moses/?n=moses.baseline>

³<http://translate.google.co.uk/>, presumably a statistical system at that time.

⁴<https://www.bing.com/translator/>, also presumably a statistical system at that time.

⁵<http://www.systransoft.com/>, presumably a hybrid rule-based / statistical system at that time.

document provided. They were also advised to read the questions first and then look for the information required on the text in order to speed up the task. Questions in CREG-mt-eval were marked as proposed by Ott et al. (2012): *correct answer* (1 mark), if the answer is correct and complete; *extra concept* (0.75 marks), when incorrect additional concepts are added; *missing concept* (0.5 marks), when important concepts are missing; *blend* (0.25 marks) when there are both extra and missing concepts; and *incorrect* (0 marks), when the answer is incorrect or missing.

Given the marks and the type of question, RCQ overall scores (f) are calculated as:

$$f = \alpha \cdot \frac{1}{N_l} \sum_{k=1}^{N_l} l_k + \beta \cdot \frac{1}{N_r} \sum_{k=1}^{N_r} r_k + \gamma \cdot \frac{1}{N_i} \sum_{k=1}^{N_i} i_k,$$

where N_l , N_r and N_i are the number of literal, reorganization and inference questions, respectively, l_k , r_k and i_k are real values between 0 and 1, according to the mark of question k , and α , β and γ are weights for the different types of questions.

We experiment with three different types of scores: *simple* (same weight for all question types: $\alpha = \beta = \gamma = 1.0$), i.e. marks are averaged giving all questions the same importance; *weighted*, i.e. marks are averaged using different weights for different types of question ($\alpha = 1$, $\beta = 2$ and $\gamma = 3$),⁶ and *literal*, where only marks for literal questions are used to compute the average quality score ($\alpha = 1$, $\beta = \gamma = 0$). The last score is interesting because literal questions are the most similar to gap-filling problems and correspond to almost 80% of the corpus and they should be easier to answer than other types. Therefore, problems in answering a literal question may be a sign of a bad quality translation.

Figure 1 shows an example of the questionnaires presented to the test takers. In this example, the first, second and last questions are inference questions, whilst the third and fourth questions are literal questions.

2.3 Gap filling task

Twenty different kinds of configurations were used in problems posed to informants. Sixteen configurations used the four MT systems to generate hints, in two modalities (showing the full

⁶These values reflect the expected relative difficulty of questions: inference harder than reorganization, and reorganization harder than literal.

Questionnaire 4

*Required

at the end of the street surrounded by single-family houses you achieve your new home .
from the door step in the top of this wonderful accommodation unit , the owners very
carefully and with love for detail removiert .
on the ground floor there is a bathroom with daylight , guest WC , the kitchen and a very
spacious living / dining room with a fascinating view of the well-kept garden .
down the stairs are another bathroom with shower - not yet used - as well as two rooms ,
both with a view to the beautiful garden and the adjacent forests .
the rooms in this area are in high quality laminate .
if you like to stay at the same time , only a family environment , this apartment for you right
.

KU-38.1: For whom is this apartment ideal?

Your answer

KU-38.3: Is the apartment in a new building or an old building?

Your answer

KU-38.4: Name two rooms on the ground floor.

Your answer

KU-38.5: Where is the apartment?.

Your answer

KU-38.8: How many rooms are in the basement?

Your answer

Figure 1: A screenshot of a RCQ questionnaire.

machine-translated document, or just the problem sentence) and with two different gap densities (10% or 20%). We added 4 additional configurations with no hint, using the same two gap densities, and with two different gap-selection strategies (statistical language model entropy and random).

The gap entropy at position k of sentence w_1^N is given by,

$$H(k, w_1^N) = - \sum_{x \in V} p(x|w_1^N, k) \log_2 p(x|w_1^N, k),$$

with V the target vocabulary (including the unknown word UNK), and with

$$p(x|w_1^N, k) = \frac{p(w_1^{k-1} x w_{k+1}^N)}{\sum_{x' \in V} p(w_1^{k-1} x' w_{k+1}^N)},$$

estimated using a 3-gram language model trained trained using KenLM (Heafield, 2011) on the English NewsCommentary version 8 corpus.⁷ Gaps

⁷<http://www.statmt.org/wmt13>

are punched in order of decreasing entropy, disallowing gaps at stop-words or punctuation, and ensuring that two gaps are never consecutive or separated only by stop-words or punctuation.

To select important sentences for the test, for each of the reference documents, the best single-sentence summary was selected as the problem sentence using GenSim.⁸

Each of 60 informants was given exactly one problem per document. Problem configurations were assigned such that each informant tackled at least one problem in each configuration, and each document was evaluated 3 times in each configuration. The mean time per problem was about 1 minute.

To create the user interface for the task we modified⁹ Ageeva et al.'s (2015) version of an older

⁸<https://rare-technologies.com/text-summarization-with-gensim/>; the percentage of text to be kept in the summary is reduced until it contains a single sentence.

⁹<https://github.com/mlforcada/Appraise>

version (2014) of Federmann’s (2012) Appraise.¹⁰ Each problem was presented in Appraise in a single screen, divided in three sections. The top of each screen reminded informants about the objective of the task. Immediately below, a machine-translated *Hint text* is provided for those 16 configurations that have one. The sentence in the hint text corresponding to the problem sentence is highlighted when a complete document is provided. At the bottom of the screen, the *Problem sentence* containing the gaps to be filled is provided. Figure 2 shows a screenshot of the interface, where a whole machine-translated document is shown as a hint, with the key sentence highlighted. The score for each problem and configuration is simply the ratio of correctly filled gaps.

3 Results

Table 1 shows, for each system, the averaged informant performance (see Appendix A for details) for the GF and RCQ quality scores explained previously; BLEU and NIST scores are also given as a reference. In view that score distributions are actually very far from normality, the usual significance tests (such as Welch’s *t*-test) are not applicable; therefore, statistical significances of differences between RCQ and GF scores will be reported throughout using the distribution-agnostic Kolmogorov–Smirnov test.¹¹ Note that previous work in RCQ did not provide statistical significance when comparing different hinting conditions, and that only Jordan et al. (2017) provided that information for GF.

3.1 Reading comprehension questionnaire scores

According to all three variations of RCQ scores, and contrary to BLEU and NIST, Systran appears to be better than the homebrew Moses. The RCQ scores for the professionally translated documents (‘Human’ row on the table) are higher than those for the best MT system, which shows that the questions are answerable from the texts and that informants did follow the guidelines as expected.

We also report the statistical significance of score differences and find (a) the only statistically significant difference at $\alpha < 0.05$ between MT systems for any score type is between Google

and the homebrew Moses; (b) all three scores of Bing, Google and Systran are statistically indistinguishable among them; (c) some (but not all) scores obtained with the professional translation are not statistically different from those obtained with Google, Bing or Systran MT output; and (d) all three scores obtained with the professional translation are statistically distinguishable from those with Moses output.

3.2 Gap-filling

Gap placement strategy: Filling of gaps in the absence of a hint was done in two configurations: one where gaps were punched at random, and one where gaps were punched where LM entropy was maximum. Entropy appears to make gap filling more difficult in the absence of hints (19.6% vs. 25.8% success rate) The value of $p_{KS} = 0.081$, above the customary $\alpha = 0.05$ significance threshold, would however tentatively support our use of entropy-selected gaps in all situations where MT was used as a hint.

Comparing MT systems: Taking all MT systems together, one can see that the success rate (58%) is, as expected, 3 times larger than that obtained without MT using the entropy-driven gap placing strategy (19%) and this difference is statistically significant. The homebrew Moses system is the least helpful (55.9%), and Bing the most helpful (62.6%), but the only statistically significant difference is between these two ($p_{KS} = 0.005$) and between Bing and Systran ($p_{KS} = 0.044$). Even with 432 problems solved for each system, MT systems were hard to distinguish by success rate (Jordan et al. (2017) report clearer differences between systems, but the paper does not clarify whether they are running the same problems through all MT systems to ensure the independence of their comparisons).

Figure 3 shows box-and-whisker plots of the distribution of performance across all 60 informants for each MT system. The large overlap observed among the four MT systems illustrates how hard it is to simply average gap-filling scores to evaluate them.

Even if annotators are quite different, each one of them may still be consistent in the relative scores they give to different MT systems. Plotting the average score each informant gives to each MT system against their average score for all systems after removing four clearly outlying in-

¹⁰<https://github.com/cfedermann/Appraise>

¹¹https://en.wikipedia.org/wiki/Kolmogorov-Smirnov_test

Instructions: Fill each one of the gaps in the "problem sentence" at the bottom with the most fitting **single word**, using only information from the *hint text* (if there is one).

Hint text: (you might need to scroll to find some highlighted text)

The Federal Republic of Germany after 1945 experienced a huge economic boom, which was the economic basis for a stable democracy.

In the German Democratic Republic the socialist one-party dictatorship of the SED and the socialist planned economy have been introduced at the same time.

Until 1989, the GDR had therefore great economic problems.

The consequences had a major impact on life in the GDR.

Problem sentence: At the same time in the German Democratic Republic , the socialist one-party dictatorship of the SED and

state-planned were introduced .

✓ Submit

Figure 2: A screenshot of the gap-filling evaluation interface, showing a whole machine-translated document as a hint (with the key sentence highlighted).

	BLEU	NIST	RCQ scores			GF scores		
			Simple	Weighted	Literal	Overall	10%	20%
Google	0.306	4.66	0.753	0.748	0.776	0.592	0.565	0.619
Bing	0.281	4.40	0.709	0.695	0.734	0.618	0.595	0.640
Homebrew	0.241	4.51	0.594	0.577	0.608	0.550	0.547	0.553
Systran	0.203	3.05	0.680	0.670	0.701	0.569	0.544	0.595
MT Average			0.684	0.673	0.705	0.582	0.563	0.602
Human	1.000	10.0	0.813	0.810	0.872			
No hint (random)						0.258	0.302	0.213
No hint (entropy)						0.193	0.191	0.195
No hint (average)						0.225	0.247	0.204

Table 1: A comparison of BLEU and NIST scores, RCQ marks in the three possible weightings, and GF success rates at different densities.

formants, Pearson correlations are only moderate (ranging between 0.47 and 0.73), and the slopes a_{system} of line fits of the form $\text{score}(\text{system}) = a_{\text{system}} \text{score}(\text{all})$ show the same ranking as average scores: $a_{\text{homebrew}} = 0.95$, $a_{\text{Systran}} = 0.97$, $a_{\text{Google}} = 1.00$, $a_{\text{Bing}} = 1.06$, but are very close to each other and their confidence intervals overlap substantially.

Effect of context: In half of the configurations with MT hints, a single machine-translated sentence was shown; in the other half, the whole machine-translated document was shown as a hint. The results indicate that extended context, instead of helping, seems to make the task slightly more difficult (58.3% vs. 59.5% success rate), but differences are not statistically significant; therefore, GF scores in Table 1 are average scores obtained with and without context. This supports evaluation

through simpler GF tasks based on single-sentence hints.

Effect of gap density: Gaps were punched with two different densities, 10% and 20%, to check if a higher gap density would make the problem harder. Contrary to intuition, the task becomes easier when gap density is higher, and the result is statistically significant ($p_{\text{KS}} < 0.001$). This unexpected result is however easily explained as follows: problems with 20% gap density contain all of the high-entropy gaps present in 10% problems, plus additional lower-entropy gaps, which are easier to fill successfully, and therefore, the average success rate rises. In the no-hint situation, however, as shown in Table 1, higher densities would seem to make the problem harder, perhaps because the only information available to fill the gaps comes from the problem sentence itself,




Figure 3: Box-and-whisker plots of the distribution of informant performance for each MT system.

and higher gap densities substantially reduce the number of available content words in the sentence. However, the differences are not statistically significant.

Gap density and MT evaluation: When comparing MT systems using only the 10% gap density problems, no differences are found to be statistically significant. This means that for very hard gaps, systems would appear to behave similarly. When selecting a value of 20% for the gap density (some easier gaps are included), Bing and Google do appear to be significantly better than the homebrew Moses.

Inter-annotator agreement: As 3 different informants filled the gaps for exactly the same set of problems and configurations, with 20 such sets available, we studied the pairwise Pearson correlation r of their GF success in each of the 36 problems.¹² All values of r were found to be positive, averaging around 0.58, a sign of rather good inter-annotator agreement. After removing two outlying informants ($r < 0.1$), results did not appreciably change.

Allowing for synonyms: The GF success scores reported thus far have been computed by giving credit only to *exact* matches. We have studied giving credit to *synonyms* observed in informant work, namely to those appearing at least twice (in the work of all informants) that, according to one

¹²The usual Fleiss' kappa statistic cannot be applied here because the labels are not nominal or taken from a discrete set, but rather numerical success rates.

of the authors, preserved the meaning of the problem sentence, or were trivial spelling or case variations. A total of 124 frequent valid substitutions were considered. As expected, GF success rates (see table 2) increase considerably, for example, from 22.7% to 32.2% for no hint, or from 58.9% to 75.5% for all systems averaged. The relative ranking of MT systems is maintained; the statistical significance of the homebrew Moses results versus Bing results is maintained, and two additional statistically significant differences appear: Google vs. homebrew Moses and Systran vs. homebrew Moses. The statistical significance of the effect of gap density disappears when allowing for synonyms. This indicates that it would be beneficial to assign credit to *synonyms* if the necessary language resources are available or if further analysis of actual GF results is feasible.

3.3 Correlation between GF and RCQ

One of our main goals was to explore whether GF would be able to reproduce the results of the established method in the field, RCQ. Table 1 shows reasonable agreement between RCQ and GF scores: both give the homebrew Moses system the worst score, and commercial statistical systems (Bing and Google) get the best scores. Also, as commonly found for subjective *judgements* (for example, Callison-Burch et al. (2006)), BLEU and NIST penalize the rule-based Systran system with respect to the statistical homebrew system, while *measurements* of human performance do not, but the differences observed are however not statistically significant.

System	GF scores with synonyms			GF scores without synonyms		
	Overall	10%	20%	Overall	10%	20%
Google	0.757	0.711	0.776	0.592	0.565	0.619
Bing	0.795	0.785	0.804	0.618	0.595	0.640
Homebrew	0.704	0.711	0.697	0.550	0.547	0.553
Systran	0.765	0.750	0.781	0.569	0.544	0.595
MT Average	0.755	0.746	0.765	0.582	0.563	0.602
No hint (random)	0.339	0.379	0.299	0.258	0.302	0.213
No hint (entropy)	0.306	0.322	0.290	0.193	0.191	0.195
No hint (average)	0.322	0.350	0.294	0.225	0.247	0.204

Table 2: Effect in success rates of allowing for synonyms in GF

On the other hand, GF and RCQ scores assigned to specific (document, MT system) pairs show low correlation. This may be due to the scarcity of RCQ data (only one data point per document–MT system pair, as compared to of 12 data points for GF), or to the fact that, while RCQ takes the whole document into account, GF only looks at a specific sentence. In addition, the RCQ tests and the sentence selected for GF for a given document may not directly correspond, i.e. the information required from the document to answer the RCQ tests may differ from the information required to fill the gaps in a given sentence. This happens because the comprehension questions may target different parts of the text and do not require the sentence selected by our GF approach. A natural follow up of this work is to use sentences for GF directly related to the RCQ tests.

4 Concluding remarks

We have compared two methods for the evaluation of MT in gisting applications: the well-established method using reading comprehension questionnaires and an alternative method: gap filling. While RCQ require the manual preparation of questionnaires for each document, and grading of answers to open questions, GF is cheaper, as it only needs reference translations for one or a few sentences in each document and both questions and scores can be obtained automatically. GF is fast and easily crowdsourcable.

In GF, without a hint, we found that entropy-selected gaps appear to be harder than random gaps. We therefore recommend using entropy-selected gaps to discourage guesswork and incentivize annotators to rely on the MT hints. Providing the whole machine-translated document as a hint does not seem to help as compared with pro-

viding only the machine-translated version of the problem sentence. This would suggest the possibility of framing GF evaluation around single sentences.

RCQ scores obtained using a machine-translated text range between 70% and 95% of the scores obtained using a professionally-translated text. In GF, the presence of a machine-translated text clearly improves performance (by about 3 times). Both results are a clear indication of the usefulness of raw MT in gisting applications.

Both RCQ and GF rank a low-quality homebrew Moses system worst, but differ as regards the best MT system, although differences are not always statistically significant. It would seem as if informants *make do* with any MT system regardless of small differences in quality. The discriminative power of RCQ and GF evaluations is, however, quite low; this may be due to the scarcity of data; if one expects that the collection of larger amounts of human evaluation data (like the crowdsourced direct assessment (judgement) results described by Bojar et al. (2016)) would increase the discriminative power of the evaluation method, this would be much more feasible using GF, than the more costly RCQ.

Acknowledgements: Work supported by the Spanish government through project EFFORTUNE (TIN2015-69632-R) and through grant PRX16/00043 for MLF, and by the European Commission through project Health in my Language (H2020-ICT-2014-1, 644402). CS is supported by the EC project SIMPATICO (H2020-EURO-6-2015, grant number 692819). We would like to thank Dr Ramon Ziai (University of Tübingen) for making the CREG corpus available and answering our questions.

References

- Ekaterina Ageeva, Francis M Tyers, Mikel L Forcada, and Juan Antonio Pérez-Ortiz. 2015. Evaluating machine translation for assimilation via a gap-filling task. In *Proceedings of EAMT*, pages 137–144.
- Jan Berka, Martin Černý, and Ondřej Bojar. 2011. Quiz-based evaluation of machine translation. *The Prague Bulletin of Mathematical Linguistics*, 95:77–86.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, volume 2, pages 131–198, Berlin, Germany.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. Findings of the 2017 conference on machine translation. In *Proceedings of the Second Conference on Machine Translation: Volume 2, Shared Task Papers*, volume 2, pages 169–214, Copenhagen, Denmark.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. Findings of the 2015 workshop on statistical machine translation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 1–46, Lisbon, Portugal. Association for Computational Linguistics.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. Re-evaluating the role of BLEU in machine translation research. In *EACL*, volume 6, pages 249–256.
- Sheila Castilho and Sharon O’Brien. 2016. Evaluating the Impact of Light Post-Editing on Usability. In *The Tenth International Conference on Language Resources and Evaluation*, pages 310–316, Portorož, Slovenia.
- Sheila Castilho, Sharon O’Brien, Fabio Alves, and Morgan O’Brien. 2014. Does post-editing increase usability? a study with Brazilian Portuguese as target language. In *Proceedings of the 17th Annual conference of the European Association for Machine Translation, EAMT 2014*, pages 183–190. European Association for Machine Translation.
- M Crook and H Bishop. 1965. Evaluation of machine translation, final report. Technical report, Institute for Psychological Research, Tufts University, Medford, MA.
- Stephen Doherty and Sharon O’Brien. 2009. Can MT output be evaluated through eye tracking? In *The 12th Machine Translation Summit*, pages 214–221, Ottawa, Canada.
- Stephen Doherty and Sharon O’Brien. 2014. Assessing the Usability of Raw Machine Translated Output: A User-Centred Study using Eye Tracking. *International Journal of Human-Computer Interaction*, 30(1):40–51.
- Stephen Doherty, Sharon O’Brien, and Michael Carl. 2010. Eye tracking as an automatic MT evaluation technique. *Machine Translation*, 24:1–13.
- Christian Federmann. 2012. Appraise: An open-source toolkit for manual evaluation of machine translation output. *The Prague Bulletin of Mathematical Linguistics*, 98:25–35.
- M. Fuji, N. Hatanaka, E. Ito, S. Kamei, H. Kumai, T. Sukehiro, T. Yoshimi, and H. Isahara. 2001. Evaluation Method for Determining Groups of Users Who Find MT “Useful”. In *The Eighth Machine Translation Summit*, pages 103–108, Santiago de Compostela, Spain.
- Masaru Fuji. 1999. Evaluation experiment for reading comprehension of machine translation outputs. In *The Seventh Machine Translation Summit*, pages 285–289, Singapore, Singapore.
- Kenneth Heafield. 2011. KenLM: Faster and smaller language model queries. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197. Association for Computational Linguistics.
- Douglas Jones, Edward Gibson, Wade Shen, Neil Granoin, Martha Herzog, Douglas Reynolds, and Clifford Weinstein. 2005a. Measuring human readability of machine generated text: three case studies in speech recognition and machine translation. In *Proceedings.(ICASSP’05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, volume 5, pages v:1009–v:1012. IEEE.
- Douglas Jones, Martha Herzog, Hussny Ibrahim, Arvind Jairam, Wade Shen, Edward Gibson, and Michael Emonts. 2007. ILR-based MT comprehension test with multi-level questions. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 77–80. Association for Computational Linguistics.
- Douglas Jones, Wade Shen, and Martha Herzog. 2009. Machine translation for government applications. *Lincoln Laboratory Journal*, 18(1).
- Douglas A. Jones, Edward Gibson, Wade Shen, Neil Granoin, Martha Herzog, Douglas Reynolds, and Clifford Weinstein. 2005b. Measuring Translation Quality by Testing English Speakers with a New Defense Language Proficiency Test for Arabic. In *The*

- International Conference on Intelligence Analysis*, McLean, VA.
- Kenneth Jordan-Núñez, Mikel L. Forcada, and Esteve Clua. 2017. Usefulness of MT output for comprehension — an analysis from the point of view of linguistic intercomprehension. In *Proceedings of MT Summit XVI*, volume 1. Research Track, pages 241–253.
- Sigrud Klerke, Sheila Castilho, Maria Barrett, and Anders Sjøgaard. 2015. Reading metrics for estimating task efficiency with MT output. In *The Sixth Workshop on Cognitive Aspects of Computational Language Learning*, pages 6–13, Lisbon, Portugal.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *The Annual Meeting of the Association for Computational Linguistics, demonstration session*, Prague, Czech Republic.
- Ramon Ziai Meurers, Niels Ott, and Janina Kopp. 2011. Evaluating Answers to Reading Comprehension Questions in Context: Results for German and the Role of Information Structure. In *TextInfer 2011 Workshop on Textual Entailment*, pages 1–9, Edinburgh, UK.
- Jim O’Regan and Mikel L. Forcada. 2013. Peeking through the language barrier: the development of a free/open-source gisting system for Basque to English based on apertium.org. *Procesamiento del Lenguaje Natural*, pages 15–22.
- Niels Ott, Ramon Ziai, and Detmar Meurers. 2012. Creation and analysis of a reading comprehension exercise corpus: Towards evaluating meaning in context. In T. Schmidt and K. Worner, editors, *Multilingual Corpora and Multilingual Corpus Analysis*, Hamburg Studies on Multilingualism (Book 14), pages 47–69. John Benjamins Publishing Company, Amsterdam, The Netherlands.
- Juan C. Sager. 1993. *Language engineering and translation: consequences of automation*. Benjamins, Amsterdam.
- Hassan Sajjad, Francisco Guzman, Nadir Durrani, Houda Bouamor, Ahmed Abdelali, Irina Teminkova, and Stephan Vogel. 2016. Eyes Don’t Lie: Predicting Machine Translation Quality Using Eye Movement. In *The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1082–1088, San Diego, CA.
- Carolina Scarton and Lucia Specia. 2016. A reading comprehension corpus for machine translation evaluation. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- H. Wallace Sinaiko and George R. Klare. 1972. Further experiments in language translation. *International Journal of Applied Linguistics*, 15:1–29.
- Harold Somers and Elizabeth Wild. 2000. Evaluating machine translation: the cloze procedure revisited. In *Translating and the Computer 22: Proceedings of the Twenty-second International Conference on Translating and the Computer*.
- Sara Stymne, Henrik Danielsson, Sofia Bremin, Hongzhan Hu, Johanna Karlsson, Anna Prytz Liljkull, and Martin Wester. 2012. Eye Tracking as a Tool for Machine Translation Error Analysis. In *The 8th International Conference on Language Resources and Evaluation*, pages 1121–1126, Istanbul, Turkey.
- Wilson L Taylor. 1953. “Cloze procedure”: a new tool for measuring readability. *Journalism Bulletin*, 30(4):415–433.
- Masaru Tomita, Shirai Masako, Junya Tsutsumi, Miki Matsumura, and Yuki Yoshikawa. 1993. Evaluation of MT systems by TOEFL. In *The Fifth International Conference on Theoretical and Methodological Issues in Machine Translation*, pages 252–265, Kyoto, Japan.
- Trond Trosterud and Kevin Brubeck Unhammer. 2012. Evaluating North Sámi to Norwegian assimilation RBMT. In *Proceedings of the Third International Workshop on Free/Open-Source Rule-Based Machine Translation (FreeRBMT 2012)*.
- Sandra Weiss and Lars Ahrenberg. 2012. Error profiling for evaluation of machine-translated text: a Polish–English case study. In *LREC*, pages 1764–1770.

A Supplemental material

Raw gap-filling results for 2159 problems,¹³ 60 informants, 36 documents, and 20 configurations, are available for download at the following address: <http://www.dlsi.ua.es/~mlf/wmt2018/raw-gap-filling-results.csv>.

Raw reading comprehension test results for 36 documents, four different MT systems (Google, Bing, Moses and Systran) and one human reference are available, totalling 180 documents. Each document was assessed by one test taker. The markings for questions available in each document and the final document scores

¹³Should have been 2160 = 36 × 60, but data for one specific document, informant and configuration, was lost due to a bug in the Appraise system.

used in this paper (namely simple, weighted or literal) are available for download at: <http://www.dlsi.ua.es/~mlf/wmt2018/raw-reading-comprehension-results.csv>.