

SEVENTH FRAMEWORK PROGRAMME THEME ICT-2011.4.2(a) Language Technologies

ACCEPT Automated Community Content Editing PorTal

www.accept-project.eu

Starting date of the project: 1 January 2012

Overall duration of the project: 36 months

Definition of post-editing rules for English, French, German and Japanese

Workpackage n° 2	rkpackage n° 2 Name: Definition of Pre-Editing and Post-Editing Ru					
Deliverable n° 2.4	Name: Definition of Post-Editing Rules for English,					
	French, German and Japanese					
Due date: 30 June 2014	Submission date: 30 June 2014					
Dissemination level: PU						
Organisation name of lead contractor for this deliverable: Acrolinx						
Author(s): Pierrette Bouillon, Robert Grabowski						
Internal reviewer(s): Barry Haddow, Violeta Seretan						
Proofreading: Benjamin Gottesman, Victoria Porro, Manny Rayner, Violeta Seretan						

The research leading to these results has received funding from the European Community's Seventh Framework Programme (*FP7/2007-2013*) under *grant* agreement n° 288769.





Contents

1.	Obje	Objectives of the Deliverable					
2.	Scope of Post-Editing Rules						
3.	Colle	ection of Post-Editing Data					
4.	Ove	rview of Developed Rule Sets 4					
5.	Inte	ractive Rules for English					
5	5.1.	New Rules					
5	5.2.	Existing Rules					
5	5.3.	Evaluation Results					
6.	Auto	omatic Rules for French					
e	5.1.	Rules					
e	5.2.	Evaluation Results 10					
7.	Bilin	gual Rules for German					
7	7.1.	Motivation 11					
7	7.2.	Approach 11					
7	7.3.	Examples for Bilingual Rules12					
7	7.4.	Implementation					
7	7.5.	Evaluation Results					
8.	Engl	ish Pre-Editing Rules for Translation to Japanese14					
9.	Con	clusion					
Ref	erenc	es15					

Definition of Post-Editing Rules for English, French, German and Japanese

1. Objectives of the Deliverable

The main goals of this deliverable are to describe the set of post-editing rules that have been developed in *Task 2.2: Post-editing rules for MT* and to evaluate them individually. The task of evaluating their usefulness for post-editors is considered to fall under WP9; these rule sets are also being used in post-editing studies under WP7 and WP8. Task 2.2 will continue to the end of the project. We will therefore continue to improve the rules based on findings from the user studies.

As explained in D2.3, we have not developed Japanese post-editing rules, due to personnel changes within the project and consequent lack of expertise for developing Japanese rules; we have however evaluated the impact of English pre-editing rules on Japanese translations. The results are summarised at the end of this deliverable.

2. Scope of Post-Editing Rules

The goal of *Task 2.2: Post-editing rules for MT* is to create post-editing rules potentially capable of helping the community improve the quality of machine-translated user-generated content. These rules are designed to be used by the Acrolinx software.

Post-editing rules thus complement pre-editing rules, which improve the input text to reduce out-of-domain problems such as wrong or uncommon spelling. While pre-editing rules prepare the input to make it more similar to the SMT training corpus, post-editing rules aim at improving the quality of the MT output by directly addressing the shortcomings of SMT translations. Like pre-editing rules, post-editing rules can be interactive or automatic and can be applied at different stages of the post-editing process:

- before the post-editing process in an automatic way ("pre-postediting"), or
- during the process in an interactive way, or
- after the process, for checking human mistakes ("post-postediting").

3. Collection of Post-Editing Data

As explained in *Deliverable 2.3: Definition of post-editing rules for English, French, German and Japanese* and in the second periodic report, we started the task with the idea of automatically learning post-editing rules from manual post-editing actions performed by users. For this purpose, we machine-translated 2,000 forum sentences from English to French, and another 2,000 forum sentences from French to English using the Moses baseline systems developed in WP4. We then collected manual post-editing data from Amazon Mechanical Turk (AMT) workers for English, and from translators for French. Post-editors were given the task of manually correcting the English and French MT outputs. They were allowed to consult the source sentence while carrying out this task.

Although, as expected, the quantity of data was not sufficient to automatically learn and generalise post-editing rules, it proved valuable as a basis for manual development of rules. We identified critical patterns with the help of a tool that organises and groups the segments by the type of post-editing action performed.

4. Overview of Developed Rule Sets

We have developed French, English and German rules corresponding to the language pairs considered in ACCEPT by relying on the following data sets:

- French forum posts and NGO documents translated from English,
- English forum posts and NGO documents translated from French,
- German forum posts translated from English.

For all three target languages, we first prepared sets of interactive rules intended to be applied at the end of the actual post-editing phase to finalise the text. These rules are based on the pre-editing rules described in D2.2 and aim at correcting human mistakes (at the level of spelling, grammar and style). In the ACCEPT post-editing plug-in, these rules will be applied after the main post-editing process, i.e., after the post-editor has changed the translation and finished his or her post-editing project.

Very often, though, post-editing the MT output requires a different form of support. Since the MT output quality strongly depends on the language pair used, we investigated different types of additional rules:

- For English text translated from French, where the average translation quality is the highest among the language pairs of the project, we developed an additional set of rules that points the user to potential issues. This set can be applied separately in the portal, or in combination with the standard writing rules.
- For French text translated from English, the overall translation quality is lower, mainly due to typical issues that appear frequently and that repeatedly require the same post-editing patterns. Here, we developed a set of rules that can be applied automatically after the MT step, before the user is shown the text.
- For German text translated from English, the translation quality is the lowest. For example, the meaning of a sentence is often lost because words are completely missing in the MT output. For this translation pair, we developed a novel type of "bilingual rules" that take the English source text into account to verify the German translation.

In the remaining of this document, we describe the different rule sets and present evaluation results.

5. Interactive Rules for English

We developed a set of interactive rules for English, containing both newly-created rules and existing rules from the standard Acrolinx rule sets. The purpose of these rules is to identify translation problems that require fundamental changes to the output, possibly with the help of the source text. The focus of development was on reliably marking error patterns and explaining the issue to the user, for example signalling that the position of an adjective may be wrong. Some of the rules may also provide a correction.

5.1. New Rules

These rules are based on the corrections applied manually by post-editors in the AMT experiments. They differ from the standard Acrolinx writing rules in that they flag parts of the sentence that are likely to be gallicisms (that is, phrase structures or syntactic elements that are uncommon or wrong in English, but typical of French).

Examples:

Mark a wrongly-positioned modifier, where the adjective follows the noun:

The scan has allowed to detect *threats minor*. (from des menaces mineures)

Mark improper usage of an infinitive after "just", where simple past tense was expected:

I just buy new servers. (from *je viens d'acheter*)

Mark lexical errors caused by literal translation from French:

perhaps that (from peut-être que)

it *must be that* (from *il faut que*)

Mark certain noun cluster constructions with of:

Join the Support of Norton products. (from le Support des produits Norton)

I use the <u>version of Ghost 2003</u>. (from *la version 2003 de Ghost*)

Mark a pronoun that should be reflexive and in a different position, or removed:

You can <u>you</u> express openly. (from yous pouvez yous exprimer)

Mark questions that lack a main verb:

Did you your key? (from Tu as bien ta clé?)

5.2. Existing Rules

We also included some of the standard Acrolinx grammar rules in this set. These rules work exactly as if a text had been written by hand. As already explained, the aim is to point to global problems with the MT output beyond the specific error they are marking.

Examples:

Mark missing spaces after end-of sentence punctuation, since this usually means that Moses has not correctly determined the sentence boundary and left words untranslated:

This window is <u>ouvre.Donc</u> I have not access the history.

Mark sequences with duplicate parts of speech with the "avoid duplicates" rule, since this usually hints at missing words rather than too many words:

You can copy the details in the clipboard and paste the in a reply.

5.3. Evaluation Results

For the evaluation, we machine-translated 12,000 sentences (segments) from the French Norton forum to English using the Moses baseline system developed in WP4. An evaluator was given the task of judging for each flag whether it correctly marked the issue. The results can be found in Table 1. The rules flagged a relatively low number of segments (5.1%), but they had a very high precision (83% overall).

Rule name	Flags (in 12k segments)	Evaluated flags	Correct flags	Wrong flags	Unsure	Precision
New rules						
missing imperfect	47		36	8	3	77%
noun cluster with "of"	35	(same as total flags)	25	10	0	71%
lexical errors	14		14	0	0	100%
wrong pronoun	38		30	8	0	79%
modifier position	73		55	13	5	75%
questions without verb	10		10	0	0	100%
Existing rules	Existing rules					
wrong verb form	184		149	11	24	81%
wrong sequence of words	153	(same as	118	3	32	77%
missing space	14	total	12	0	2	86%
missing word	30	iiags)	21	2	7	70%
avoid duplicates	16		16	0	0	100%
TOTAL	614	614	486	55	73	83%

Table 1: Evaluation results for interactive English post-editing rules

6. Automatic Rules for French

6.1. Rules

For post-editing French, we developed 26 automatic monolingual rules, which can be grouped in two main classes, according to the phenomena they treat (see Table 2):

- Rules that target MT-specific errors (e.g., "séquence_incorrecte"),
- Traditional grammar and style rules (e.g., "accord_sujet_verbe").

In contrast to English, all the rules have unique suggestions and can be applied automatically. They are meant to help post-editors speed up their work by reducing the number of edits they have to perform, and to improve their working experience by enhancing text readability.

Rule name	Description of the rule	Flags (in 10k segments)	Evaluated Flags (max 50)	Correct flags	Wrong flags	Precision	
	MT-specific errors						
séquence_incorr	ecte						
	Incorrect sequence of words	306	50	47	3	94%	
Je suis en espérant qu'ils vont résoudre le problème.							
	J' <mark>espère</mark> qu'ils vont résoudre	le problème.					
nom_de_nom	•						
	Incorrect noun sequence	43	43	35	8	81%	
	Le installation problème a ét	é résolu.					
	Le problème d'installation a	été résolu.					
forme_verbale_i	ncorrecte						
	Wrong verb form	255	50	49	1	98%	
	Il n' <mark>a</mark> pas <mark>faire</mark> ça.						
	Il n' <mark>a</mark> pas <mark>fait</mark> ça.						
mots_doublés							
	Doubled words	32	32	30	2	94%	
	J'ai fait <mark>le un</mark> erreur.						
	J'ai fait <mark>un</mark> erreur.						
nom_sigle							
	Reorder nouns and abbreviations	79	50	43	7	86%	
	J'ai reçu un DLL message d'erre	ur.					
	J'ai reçu un <mark>message d'erreur D</mark>	DLL.					
ne_pas_élider	•						
	Incorrect elision	37	37	36	1	97%	
	Le "DiskPart" <mark>d</mark> 'XP ne permet p	as.					
	Le "DiskPart" <mark>de</mark> XP ne permet	pas.					
reordonner_phra	se_nominale						
	Incorrect order of noun phrase components	199	50	37	13	74%	
	Le Norton technicien m'a conse	eillé de le faire 1 f	ois/semaine.				
	Le technicien Norton m'a conse	eillé de le faire 1 f	ois/semaine.				

Traditiona	l grammar and	style rules			
répétez_sujet					
Repeat the subject of the phrase	31	31	24	7	77%
Suis intéressée par un feedback					
Je suis intéressée par un feedbac	:k.				
négation_incomplète					
Incomplete negation	195	50	44	6	88%
je vais le relire pour être sûr de	<mark>pas avoir</mark> de pro	oblèmes!			
je vais le relire pour être sûr de	ne pas avoir de	problèmes!			
accord_sujet_verbe					
Error in subject-verb agreement	580	50	42	8	84%
Tout d'abord <mark>je</mark> vous <mark>fait</mark> part de	ma configuratio	on:			
Tout d'abord <mark>je</mark> vous <mark>fais</mark> part de	ma configurati	on:			
homophones_divers					
Confusion of homophones such as "diagnostic" et "diagnostique"	9	9	9	0	100%
Il ne <mark>diagnostic</mark> aucune erreur.					
Il ne diagnostique aucune erreur	r.				
espaces_autour_ponctuation					
Incorrect spaces before or after punctuation	1997	50	49	1	98%
Comme l'a stipulé Misstigry , il s	s'agit de service	es.			
Comme l'a stipulé Misstigry, il s	'agit de services	s.			
ajouter_virgule					
Missing comma	215	50	50	0	100%
je ne comprends pas ce qui s'est	<mark>passé</mark> mais j'ai	trois comptes No	orton.		
je ne comprends pas ce qui s'est	<mark>passé,</mark> mais j'ai	i trois comptes N	orton.		
ajouter_tiret			_		
Missing hyphenation	98	50	44	6	88%
c'est à dire ?					
c'est-à-dire ?					
accord_phrase_nominale			_		
Agreement error in noun phrase	347	49	45	4	92%
Bonjour, quel version de norton	360 avez-vous ?	0			
Bonjour, quelle version de norto	n 360 avez-vous	5 ?			
évitez_les_anglicismes					
Avoid anglicisms	193	50	49	1	98%
Hi, merci pour les informations.					
Bonjour, merci pour les informat	tions.				
abréviation_incorrecte					
Use of an incorrect abbreviation	76	50	46	4	92%
J'ai un <mark>2e</mark> PC installé de la même	e façon et sur ce	lui-ci, je la trouv	e sans problèi	me.	
J'ai un <mark>2ème</mark> PC installé de la mê	me façon et sur	r celui-ci, je la tro	ouve sans pro	blème.	

ajoutez_virgule_a	après_PP					
	Add a comma after a prepositional phrase	18	18	18	0	100%
	Sous l'onglet Périphériques vous p	oouvez sélectio	nner le PC que vo	ous utilisez.		
	Sous l'onglet Périphériques, vous	pouvez sélectio	onner le PC que v	ous utilisez.		
ajoutez_un_blan	2					
	Missing space	62	50	49	1	98%
	J'ai essayé d'uploader un fichier d	le <mark>4ko</mark> seulemer	nt, et le problèm	e est le mêm	е.	
	J'ai essayé d'uploader un fichier d	le <mark>4 ko</mark> seuleme	nt, et le problèm	ne est le mêm	ie.	
espace_en_trop						
	Delete extra blanks	0	0	0	0	0%
	Le programme <u>est_bl</u> oqué.					
	Le programme <u>est bl</u> oqué.					
utilisez_impérati						
	Use the imperative form	98	50	47	3	94%
	Regardes en bas si tu vois "Boot					
	Regarde en bas si tu vois "Boot					
ponctuation_inco	prrecte					
	Incorrect punctuation or doubled ponctuation	162	49	44	5	90%
	Blocage des <mark>appels</mark> Pas de messa	ges.				
	Blocage des <mark>appels</mark> . Pas de messa	iges.				
utilisez_subjonct	f					
	Use the subjunctive form	203	50	49	1	98%
	Bien que je ne <mark>comprends</mark> pas ce o	que tu as fait q	uand tu écris :			
	Bien que je ne <mark>comprenne</mark> pas ce	que tu as fait q	uand tu écris :			
évitez_les_questi	ons_directes					
	Use inversion instead of direct questions	92	50	47	3	94%
	Tu as lu le tuto sur le forum ?					
	As-tu lu le tuto sur le forum ?					
élidez_ce_mot						
	Missing or incorrect elision	143	50	46	4	92%
	Est-ce <mark>que il</mark> s'agit ici de Norton O	nline Backup o	u du backup de I	Norton 360 ?		
	Est-ce <mark>qu'il</mark> s'agit ici de Norton On	iline Backup ou	du backup de N	orton 360 ?		
terme_incorrect						
	Wrong term	181	50	48	2	96%
	J'ai récemment installé Nortons s	ur mon bureau.				
	J'ai récemment installé Norton su	r mon bureau.				
TOTAL		5651	1118	1027	91	92%

Table 2: Evaluation results for automatic and monolingual French post-editing rules

6.2. Evaluation Results

For the evaluation, we translated 10,000 unseen sentences (segments) from the English Norton forum to French and applied the rules individually. In contrast to English, at most 50 instances of each rule were included in the evaluation (see Table 2). Two evaluators were asked to judge, for each flag, whether it correctly marked and corrected the issue it was meant to flag. Judging was carried out in a monolingual setting, by looking only at what was corrected in the target (for example, for agreement we checked if the highlighted issues were corrected in French, independently of the source text). Results were discussed between evaluators and consolidated on this basis. The results can be found in Table 2.

The rules flag 56.5% of the segments and their overall precision is around 92%. The less precise rules are "reordonner_phrase_nominale" (74%), "répétez_sujet" (77%) and "nom_de_nom" (81%). In these rules, the noise is often caused by wrong tagging, as illustrated in Figure 1, where *desinstallé* was considered as an adjective. However, it must be noted that in these cases, the rules do not necessarily degrade the translations, but merely change an incorrect sequence into another incorrect sequence (as in the example in Figure 1). As mentioned above, the usefulness of the rules for post-editors will be evaluated in WP9, where we will look at the impact of applying the set of rules as a whole. This will improve performance: in this specific example, another rule will first add the auxiliary verb (*"Lorsqu'il a désinstallé programme, il sauvegarde les fichiers"*), which is expected to lead to a better tagging for next rules.

Rule: "reordonner_phrase_nominale"

Target: Lorsqu'il désinstallé programme, il sauvegarde les fichiersAutomatically post-edited target: Lorsqu'il programme désinstallé , il sauvegarde les fichiers

Figure 1: Example of an incorrect sequence being converted into another equally incorrect sequence

7. Bilingual Rules for German

7.1. Motivation

In previous work in WP2, we found fundamental quality problems in the SMT output specific to the language pair English-German, which were related to the distance between the two languages, in syntactical terms. Two concrete patterns are the following:

- The SMT system fails to generate a German long-distance dependency not present in the English source, such as between an auxiliary and a main verb, or a separable verb and its particle.
- The SMT system translates certain grammatical constructions and idioms literally into German, instead of using the appropriate translation.

Issues like these are among the most challenging to correct, as the post-editor needs to understand the source text in order to find out which part of the meaning has changed. This gives reason to hope that providing support in these cases may lead to big gains in the effectiveness of post-editing. At the same time, a standard Acrolinx rule only has access to the German target sentence, which makes it hard to reliably detect missing or incorrect translations, and impossible to provide suggestions.

7.2. Approach

To address these issues, we developed so-called "bilingual" rules that look for certain patterns on the English side, and then verify whether the expected translation of the pattern also appears in the German translation. For example, there is a bilingual rule that looks for modal verbs followed by an infinitive verb in English, and in parallel for the corresponding structure in the German translation (also a modal verb followed by infinitive verb). Since the infinitive verb usually appears at the end of the German sentence, this poses a serious challenge for automatically-learned alignment in the MT training phase, and often leads to MT output where the infinitive verb is missing.

The bilingual rule can detect this situation. When the "modal verb + infinitive" construction is present in the source but not in the target sentence, we can point the author to the specific problem as follows:

- mark the part of the target sentence that is aligned with the phrase where the source construction appeared;
- tell the user that a German "modal verb + infinitive" construction was expected at that point;
- provide the English modal verb + infinitive as found in the source.

Additionally, we have also examined bilingual rules that specifically search for German translations that are known to be wrong. For example, the English present progressive tense should generally be translated to the German simple present. We developed a bilingual rule that checks whether an English progressive was translated literally to German; if the corresponding "Germanised" progressive tense is present, the rule indicates a translation error.

While this mechanism may still require the intervention of a bilingual post-editor, it helps to quickly identify and fix the issue, without having to understand the source text deeply. Also, in contrast to many Quality Estimation (QE) techniques, we can express the issue in relatively precise linguistic terms, rather than in statistical features which may be of little relevance to the post-editor.

7.3. Examples for Bilingual Rules

Search for by followed by a gerund in the English source sentence, and expect *indem* followed by a verb or *durch* followed by a noun in the German translation. The following translation is marked as correct:

Open the software <u>by clicking</u> on the icon. Öffnen Sie die Software, <u>indem</u> Sie auf das Symbol <u>klicken</u>.

Search for a verb in past or perfect tense in English, and expect a verb in past or perfect tense in German. The following translation is marked as incorrect, as no such verb is found in the translation:

I <u>came in</u> this morning and found one of my cats sleeping on the CPU. Ich heute Morgen und einer meiner Katzen schlafenden auf der CPU.

Search for a verb in present progressive tense in English, and look for a (wrong) literal translation of the present progressive form in German. The following translation is marked as incorrect:

I <u>am working</u> on this problem. Ich <u>bin</u> an diesem Problem <u>arbeiten</u>.

7.4. Implementation

We have implemented the concept of a "bilingual" rule as a pair of two monolingual Acrolinx rules, one for English and one for German, which flag the English source and expected German target construction, respectively. All of these rule pairs are packaged as two parallel Acrolinx rule sets (one for English and one for German), where corresponding rules have the same name.

We then developed a tool that runs two Acrolinx checks in parallel on the source and target texts, and finds correlations between flags for corresponding rules. The tool also uses the SMT phrase alignment information to search only for flags in the sentence parts that are aligned with each other. While Moses can already output this alignment, we needed to extend the Moses server wrapper (specifically the translation broker service and its Translate API) to correctly convey the alignment information to the bilingual checking tool.

For the time being, the mechanism is not integrated into the Acrolinx software, but rather runs as a prototype offline tool that makes use of an Acrolinx server for checking the source and the target side. For that reason, "bilingual rules" cannot straightforwardly be integrated into the ACCEPT post-editing plug-in. The integration is planned for M31-36 of the project.

7.5. Evaluation Results

In view of the two main types of issues presented above, we have developed two sets of bilingual rules. The rules of the first set check whether an expected translation is missing. The rules of the second set check whether a translation in the target sentence is known to be wrong.

The hypothesis for the evaluation of the first type of rules (that mark missing expected translations) is that if a bilingual rule matches on the source side but not on the target side, then the construction it is looking for is mistranslated; if it matches on the source and the target, then the construction is translated correctly. For the rules that mark incorrect translations, the hypothesis is inverted: the translation is considered bad if and only if the flag appears in the source and in the target.

To verify these hypotheses, we extracted 200 English forum posts (about 1,600 sentences) from the Symantec corpus, machine-translated them to German, and ran the tool on the source and translated texts with the developed rule sets. An evaluator was given the task of judging for each indicated translation problem whether there was indeed a problem with the translation of the marked pattern. Moreover, for each indicated "good" translation, the evaluator verified whether there was indeed no problem with the translation, or whether it was a false negative.

The results are presented in Table 3 for the rules that look for missing good translations, and in Table 4 for the rules that look for wrong translations. As can be seen from the tables, the overall results for the rules are very good, which makes us confident that our approach can indeed be used to improve the post-editing performance for the English>German translation pair.

	Matches on source but not		Matches on source and		
	on target (expecting		target (expecting good		
	translation problems)		translation	translation)	
				— 1.11	
		Iranslation		Iranslation	
	Instances	indeed wrong	Instances	indeed correct	
"by" + progressive verb	20	18 (90%)	30	27 (90%)	
imperative	19	17 (89%)	31	29 (93%)	
modal verb + infinitive	25	19 (76%)	25	20 (80%)	
modal verb + infinitive + past participle	30	28 (93%)	20	19 (95%)	
l + past tense	25	24 (96%)	25	22 (88%)	
verb + that	25	20 (80%)	25	19 (76%)	
noun + that	25	25 (100%)	25	21 (84%)	
subordinate clause with verbs	25	21 (84%)	25	19 (76%)	
progressive tense	25	24 (96%)	25	24 (96%)	

Table 3: Evaluation results for the rules that mark missing expected translations

	Rule matches target (expect problems)	on source and ting translation	Rule matches on source but not on target (expecting good translation)		
	Instances	Translation indeed wrong	Instances	Translation indeed correct	
progressive tense	25	24 (96%)	25	24 (96%)	

Table 4: Evaluation results for the rules that mark wrong translations

8. English Pre-Editing Rules for Translation to Japanese

For the English-Japanese language pair, we investigated how to improve translation quality when performing statistical machine translation of English support forum posts into Japanese. For this purpose, we explored two approaches to domain adaptation: combining in-domain monolingual data with close-domain bilingual corpora for MT training; and developing pre-editing rules to reformulate phrases that are difficult for machine translation to Japanese. All details of this work can be found in a bachelor's thesis [1] and a conference publication [2]; here, we briefly summarise the results.

The ACCEPT baseline system had been trained on bilingual Symantec product manuals and monolingual Japanese forum data, with a BLEU score of 20.37 on the test set. With the addition of a colloquial language corpus, as well as a large out-of-domain corpus that broadened the base vocabulary of the system, we achieved a considerable improvement of the BLEU score to 22.10 on the same test set.

For this improved system, we developed new English pre-editing rules. These rules are meant to be automatically applied immediately before the machine translation. Most of the rules aim to make the English input more like the English training data, while the others modify the used language slightly towards customary Japanese expressions. The rules can be grouped into four categories:

- reformulating slang and punctuation
 examples: "cuz" → "because", "guess so" → "I guess so", "&" → "and", "AFAIK" → "as far as
 I know", deletion of "LOL"
- splitting long sentences
 Sentences are split at "then", "therefore", "so" and "but"
- shortening phrases & removing ambiguities
 examples: "going to" → "will", "not" → "un-", "have to" → "must", "not...any" → "no..."
- modifying English towards Japanese add "please" to imperatives, delete discourse marker "well".

We performed a comparative evaluation similar to the ones for English and French in Task 2.1. We extracted 100 original English segments, such that for each segment, exactly one rule matched. We then generated the 100 corresponding automatically pre-edited segments, and translated the segment pairs automatically. For each segment pair, three Japanese native-speakers were asked to judge whether the translation of the original or the translation of the pre-edited segment had a

better quality, resulting in 300 individual ratings. We also computed the majority judgement for each segment pair. Table 5 shows the results.

	All individual ratings	Majority judgements
Original better	92 (31%)	25 (25%)
Pre-edited better	127 (42%)	46 (46%)
Similar quality	51 (17%)	8 (8%)
Identical translation	30 (10%)	10 (10%)
Disagreement	-	11 (11%)
Total	300 (100%)	100 (100%)

Table 5: Evaluation results for English pre-editing rules for Japanese

The results show that the pre-editing rules significantly improve the quality of the MT output (p < 0.02), and that there was relatively little disagreement among the judges. These positive results agree with those derived from automatic metrics: the BLEU score for the improved MT system on the automatically pre-edited test set further increased to 22.69. This makes us confident that pre-editing is a useful approach to improve the MT output quality for this difficult language pair.

9. Conclusion

We have developed stable sets of post-editing rules for the three language pairs of the project (French>English, English-French, and English>German). The rule sets are provided on the hosted Acrolinx servers for the ACCEPT project, and will be integrated into different stages of the post-editing workflow, depending on their use case. The French automatic rules will be applied on the MT output text, such that the manual post-editing tasks are performed on an automatically post-processed MT output. The interactive English rules are presented to the user during the main post-editing stage, whereas the general rules sets mentioned in the beginning are applied afterwards. An exception is the German bilingual rules, which are not yet integrated in the portal, as mentioned above.

The rule sets will be the subject of further experiments and evaluations. We will measure if the rules help the post-editing process, and adapt and improve them based on our findings.

References

[1] Torsten Jachmann

Exploring Domain Adaptation in Machine Translation for English to Japanese Bachelor's thesis, Saarland University, Saarbrücken, 2013.

[2] Torsten Jachmann, Robert Grabowski and Mayo Kudo

Machine-Translating English Forum Posts to Japanese: On Pre-editing Rules as Part of Domain Adaptation

In Proceedings of the 20th Annual Meeting of the Association for Natural Language Processing (ANLP), Sapporo, Japan, 2014.