Chapter 7

Patent Translation

Isao Goto

Abstract The NTCIR patent translation task was the first task for the machine translation of patents that used large-scale patent parallel sentence pairs. In this chapter, we first present the history of machine translation; the contribution of evaluation workshops to machine translation research, and previous evaluation workshops; and the challenge of patent translation at the time of the first patent translation task at NTCIR. We then describe the innovations at NTCIR, including the sharing of research infrastructure, the progress of corpus-based machine translation technologies, and evaluation methods for patent translation. Finally, we outline the developments in machine translation technologies, including patent translation and remark on the future of patent translation.

7.1 Introduction

Research on machine translation began in the 1950s immediately after the birth of computers. The first machine translation technology was *Rule-Based Machine Translation (RBMT)*, which used manually built translation rules. RBMT was actively developed from the 1970s to the 1980s. In the late 1980s, research began on *Statistical Machine Translation (SMT)*, which is a learning-based machine translation technology based on corpus statistics, (Brown et al. 1993). However, there was little research on SMT for about 10 years. Then the situation changed. From the late 1990s to around 2000, that is, since high-performance computers began to be in widespread use, large parallel corpora became available, automatic evaluation methods, such as BLEU (Papineni et al. 2002), were developed, and research on SMT began to progress rapidly.

The progress of the research was facilitated by evaluation workshops. Evaluation workshops played a dual role in providing large datasets and making evaluations comparable using shared tasks. This made it possible to conduct experiments by

I. Goto (✉)
NHK, 1-10-11 Kinuta, Setagaya-ku Tokyo 157-8510, Japan
e-mail: goto.i-es@nhk.or.jp

sharing research infrastructure and to verify the effectiveness of methods by performing comparisons using the same data. Evaluation workshops made research more active, and research on machine translation progressed. The following is a list of major evaluation workshops on machine translation that were in existence by the mid-2000s:

- Defense Advanced Research Projects Agency (DARPA) Translingual Information Detection, Extraction, and Summarization (TIDES) project (2001 to 2005): The translation languages were Chinese to English and Arabic to English, and the target domain was news. This project was succeeded by the DARPA Global Autonomous Language Exploitation (GALE) project.
- International Workshop on Spoken Language Translation (IWSLT) (2004 to present (2019)¹): As of 2004 to 2007, speech translation of travel conversations was targeted. Several languages were included, including Japanese and English. The size of the training parallel corpus was 20,000 to 40,000 sentence pairs.
- Workshop on Statistical Machine Translation (WMT) (2006 to present (2019) see footnote 1): Machine translation between European languages is the target. As of 2006 to 2007, the proceedings of the European Parliament and news were the target domains.

As of 2007, research on SMT was in progress for several language pairs and fields. For the Japanese–English language pair, the domain covered in the evaluation workshops was travel conversations only. Because the sentence lengths were short and the topic was narrow, the shared task for travel conversation translation was technically easy. By contrast, there was no shared task for long sentence translation between Japanese and English, which is useful for advancing translation technology for long sentences between languages that differ significantly in word order.

As a domain that includes long sentence translation between Japanese and English, patent translation has substantial demand, such as translation for foreign applications and translation of patents in foreign languages to understand the content of existing patents. The machine translation of patents has been required by sectors that produce and use intellectual property in countries and many companies. Therefore, if machine translation performs well for patent translation, there will be a substantial impact on society.

In 2007, RBMT systems were on the market for the machine translation of patents between Japanese and English. Through years of research and development, RBMT systems have achieved translation quality at a level that is useful as a rough translation for manual post-editing.² However, there was a barrier to further improving the translation quality of RBMT. Simply increasing the number of translation rules did not improve translation quality. Manually adding translation rules so that the appropriate translation rules can be selected in accordance with the context from many

¹This chapter was written in 2019. Thus, this year does not indicate the final year.

²In fact, when the organizers of NTCIR-7 asked Japanese–English translators to produce multi-reference translations of the test sentences, the organizers found that an RBMT system was used for rough translation, and the reference translations had to be retranslated to avoid the bias of a specific machine translation (MT) system (Fujii et al. 2008).

candidates has been a serious challenge that requires craftsmanship. It was also a serious challenge to make sentences generated by combining translation rules into natural sentences as written by a person. Moreover, both the accumulated amount of bilingual patent data and computational power could be expected to increase over time. Thus, to overcome the barriers to RBMT and aim for translation quality at the level of human translation, corpus-based machine translation technology, which automatically acquires translation knowledge and sentence generation knowledge from patent data, was required. However, before 2007, there were few studies on corpus-based machine translation for the patent field.

7.2 Innovations at NTCIR

As explained in the previous section, in 2007, to advance long sentence translation technology between languages differing greatly in word order, it was appropriate timing for shared tasks of patent translation between Japanese and English. At that time, the NTCIR-7 organizers extracted over one million Japanese–English parallel sentence pairs from parallel patent applications and launched the shared task of patent translation. This led to research on corpus-based machine translation for long patent sentences between Japanese and English. Patent translation tasks were conducted four times, from NTCIR-7 to NTCIR-10, over six years (Fujii et al. 2008, 2010; Goto et al. 2011, 2013). In NTCIR-9, the Chinese–English patent translation task was added.

In the following, we present a summary of the comparison between SMT and RBMT for patent translation.

- From the evaluation results of NTCIR-7 in 2008, the translation quality of RBMT was higher than that of SMT for Japanese–English and English–Japanese translation.
- From the evaluation results of NTCIR-9 in 2011, the translation quality of SMT for English–Japanese caught up with that of RBMT. The translation quality of SMT for Chinese–English was higher than that of RBMT because the translation quality of RBMT was low.
- From the evaluation results of NTCIR-10 in 2013, SMT outperformed RBMT for English–Japanese translation. Although SMT could not catch up with RBMT for Japanese–English translation, the top SMT system for Japanese–English translation at NTCIR-10 improved compared with the top SMT system at NTCIR-9.

Thus, through four rounds of shared tasks over 6 years, the performance of SMT substantially improved for patent translation including long sentences for Japanese–English and English–Japanese, and Chinese–English. As a result, corpus-based machine translation could make it possible to overcome the challenges encountered by RBMT. This was the biggest innovation in the patent translation tasks.

In the following, the purpose of each patent translation task is described, and an overview of each of the four tasks, major findings, and innovations is provided. The goals of the patent machine translation tasks were as follows:

- to develop challenging and significant practical research into patent machine translation;
- to investigate the performance of state-of-the-art machine translation in terms of patent translations involving Japanese, English, and Chinese;
- to compare the effects of different methods of patent translation by applying them to the same test data;
- to explore practical MT performance in real scenarios for patent machine translation;
- to create publicly available parallel corpora of patent documents and human evaluation of MT results for patent information processing research;
- to drive machine translation research, which is an important technology for the cross-lingual access to information written in unfamiliar languages; and
- ultimately, to foster scientific cooperation.

7.2.1 Patent Translation Task at NTCIR-7 (2007–2008)

As described in Sect. 7.1, 2007 was a time when SMT technology was progressing. Because there was an open-source SMT tool called Moses (Koehn et al. 2007) at that time, it was easy to conduct experiments on SMT if a bilingual parallel corpus was available. SMT could translate short sentences, such as travel conversations, to some extent. By contrast, the translation quality of SMT was low for long sentences between language pairs with a largely different word order. Therefore, translating a patent document that included long sentences between Japanese and English, which largely differ in word order, was a serious challenge for SMT.

In 2007, the organizers constructed a Japanese–English parallel patent dataset that consisted of approximately 1.8 million parallel sentence pairs and launched the shared tasks of Japanese–English and English–Japanese patent translation. This was the first time that more than one million parallel sentence pairs in Japanese and English became widely available for research. The task organizers extracted the Japanese–English parallel patent sentence pairs from Japanese–English bilingual patent families. A patent family is a set of patents taken in more than one country to protect a single invention. The extraction of parallel sentence pairs was conducted by applying an automatic sentence alignment method (Utiyama and Isahara 2007) to approximately 85,000 patent families from 10 years of Japanese patents published by the Japan Patent Office (JPO) and 10 years of English patents published by the United States Patent and Trademark Office.

In the NTCIR-7 patent translation task, human evaluation was performed. For Japanese–English translation, human evaluation was performed for a total of 15 system outputs that consisted of the 14 system outputs submitted by the participating teams and a system output of the SMT tool Moses used by the organizers. The results showed that the automatic evaluation BLEU-4 score of SMT was higher than that of RBMT; however, in the human evaluation, the results indicated that the actual translation quality of RBMT was better than that of SMT. For English–Japanese

translation, human evaluation was performed for some representative systems, and the results showed that the trend of the comparison between SMT and RBMT was similar to that of Japanese–English translation.

Additionally, the organizers compared the effect when English–Japanese machine translation was used for cross-lingual patent retrieval (CLPR) as an extrinsic evaluation. They used a standard retrieval method for CLPR. Because the standard retrieval method did not use the order of words in queries and documents, the order of words did not affect the retrieval results. The CLPR results were highly correlated with the BLEU score, and SMT was better than RBMT; that is, the results showed that SMT was more effective than RBMT in terms of translation word selection.

7.2.2 Patent Translation Task at NTCIR-8 (2009–2010)

The Japanese–English and English–Japanese patent translation tasks continued. The organizers expanded the size of the bilingual corpus by extracting parallel sentence pairs from 15 years of patent families, and provided the task participants with a Japanese–English parallel corpus that consisted of approximately 3.2 million sentence pairs. In the tasks, no purely RBMT system was included in the evaluation and no human evaluation was performed. Therefore, SMT and RBMT could not be compared.

The system with the highest BLEU score for Japanese–English translation first translated Japanese sentences into English using RBMT, and then post-edited the translation results using SMT (Ehara 2010). The results showed that the word reordering performance of SMT had not caught up with that of RBMT. Additionally, the shared task of the automatic evaluation of machine translation was also conducted using the human evaluation results of NTCIR-7. The task evaluated automatic evaluation methods based on the human evaluation results.

7.2.3 Patent Translation Task at NTCIR-9 (2010–2011)

The organizers³ added a Chinese–English patent translation task in addition to the Japanese–English and English–Japanese patent translation tasks. Chinese–English translation is a globally required language pair and is popular in the machine translation research community. For the Japanese–English and English–Japanese translation tasks, the training dataset was the same as that of NTCIR-8, that is, approximately 3.2 million sentence pairs, and the test dataset was newly produced. For the Chinese–English translation task, the organizers provided the task participants with a training dataset that consisted of one million parallel sentence pairs of Chinese–English bilingual patents. The organizers produced translation results using com-

³The organizers of the patent translation task at NTCIR-9 changed from the organizers at NTCIR-8.

mercial RBMT systems to compare SMT and RBMT. They also performed human evaluation. Twenty-one teams around the world participated in the patent translation tasks. The introduction of the Chinese–English translation task led to the participation of top international teams, such as BBN (Ma and Matsoukas 2011), IBM Watson Research (Lee et al. 2011), and RWTH Aachen University (Feng et al. 2011).

The findings obtained from the evaluation results were as follows: For English–Japanese translation, the top SMT system achieved a translation quality equal to or better than that of the top RBMT system. For the first time in patent translation from English to Japanese, the top SMT system had caught up with the top RBMT system. The top SMT system improved substantially in translation quality by improving word reordering performance using a pre-ordering method (Sudoh et al. 2011). It became clear that separating word reordering from the decoding process could obtain a large effect in a simple manner. For Chinese–English translation, the translation quality of SMT was higher than that of RBMT because the performance of the Chinese–English RBMT systems was low.

The organizers created and applied a new human evaluation criterion, that is, “Acceptability,” in addition to “Adequacy,” which is a conventional human evaluation criterion. The criteria for each grade of Adequacy were ambiguous, and the actual ratings were compared mainly on a relative basis to distinguish between the systems to be evaluated. Therefore, the translation quality was not necessarily the same for the same grade. For example, grade 3 when only low-level systems were evaluated and grade 3 when only high-level systems were evaluated would be different translation qualities. Thus, it was not possible to know the actual quality using such relatively scored grades. By contrast, Acceptability was defined as an objective and clearer standard, with the aim of making the quality of the same grade constant. The Acceptability results showed that the percentage of translated sentences that could convey all the meanings of the source sentences was 60% for the top systems for both Japanese–English and English–Japanese translation, and the percentage was 80% for the top system for Chinese–English translation.

7.2.4 Patent Translation Task at NTCIR-10 (2012–2013)

The Japanese–English, English–Japanese, and Chinese–English patent translation tasks were continued at NTCIR-10. The training dataset was the same as that at NTCIR-9 and the test dataset was newly produced. Twenty-one teams participated in the tasks.

The findings obtained from the evaluation results were as follows: For English–Japanese translation, the top SMT system (Sudoh et al. 2013) outperformed the RBMT systems in terms of translation quality. For Japanese–English translation, RBMT was still better than SMT; however, the translation quality of the top SMT system had improved from NTCIR-9 (Sudoh et al. 2013). For Chinese–English trans-

lation, the top system used neural networks in a language model to improve performance (Huang et al. 2013), and the effectiveness of neural networks for machine translation was thus demonstrated.

If the test data was simply selected from the automatically extracted parallel corpus, biases, such as lengths or included expressions, may result. To reduce biases, the organizers selected test sentences using two methods. For one method, the organizers first calculated the distribution of sentence lengths in monolingual patent documents in the source language, and divided the cumulative length distribution into quartiles (25% each). Each quartile was called a sentence length class. Next, they classified the automatically aligned sentences in the source language into four classes according to their sentence lengths and extracted the same number of sentences from each class as test sentences. For the other method, the organizers randomly selected test sentences from all the description sentences in the source language patents for bilingual patents. Translators translated the test sentences to produce their reference translations. The data produced by the second method was used for the human evaluation.

At NTCIR-9, the top systems performed well for sentence-level evaluations. Therefore, the NTCIR-10 organizers wanted to see how useful the top systems were for practical scenarios. Patent examination was one of the practical scenarios. The organizers performed Patent Examination Evaluation (PEE), which measures the usefulness of MT systems for patent examinations. PEE is described as follows: PEE assumes that the patent is examined in English. When a patent application in English is filed, an examiner examines existing patents and rejects the patent application if almost identical technology is described in an existing patent. If a patent application is rejected by referencing an existing patent, the examiner writes the final decision document (*Shinketsu*), which describes the facts about the existing patent on which the rejection is based. Assuming that the referenced patents were written in a foreign language, the organizers extracted the part that described the facts from the referenced patents and used the extracted sentences as test data. The test data in foreign languages (Japanese/Chinese) were translated into English using machine translation, and the translation results were evaluated according to whether the facts that were used to reject patent applications could be recognized from the translation result. PEE was performed by two experienced patent examiners. For Japanese–English translation, for the best system, all facts were recognized in 66% of referenced patents, and at least half of the facts were recognized in 100% of referenced patents. For Chinese–English translation, for the best system, all facts were recognized in 20% of referenced patents, and at least half of the facts were recognized in 88% of referenced patents. PEE achieved the evaluation of usefulness in one representative practical scenario of patent machine translation. The PEE results and translations can be used as standards of usefulness in patent examination. Specifically, by comparing new translation results for the PEE test data with the PEE evaluated translations at NTCIR-10, their usefulness in patent examination for other systems can be assessed roughly.

7.3 Developments After NTCIR-10

The evaluation workshop on Asian translation (WAT) for machine translation was launched in 2014. WAT targets machine translation between language pairs that include Asian languages. The activities of WAT have promoted the construction and sharing of research infrastructure for machine translation involving Asian languages. WAT features an open innovation platform. The test data and reference translations have been published with the training data, and the use of the same test data every year facilitates comparisons. In the following, we describe the activities and findings of WAT.

In the first workshop (WAT 2014) (Nakazawa et al. 2014), the organizers set the shared tasks of scientific paper translation between Japanese and English, and between Japanese and Chinese. An SMT system using syntactic structures achieved the highest performance.

In the second workshop (WAT 2015) (Nakazawa et al. 2015), in addition to the scientific paper translation tasks, Chinese–Japanese and Korean–Japanese patent translation tasks were included. The size of the training dataset for each patent translation task was one million sentence pairs. The results showed that the translation quality of the top SMT system was higher than that of the RBMT systems for patent translation for Chinese–Japanese and Korean–Japanese. For scientific paper translation, a reranking method using *Neural Machine Translation (NMT)* achieved the highest translation quality. The effectiveness of the scoring by NMT was thus demonstrated.

In the third workshop (WAT 2016) (Nakazawa et al. 2016), Japanese–English and English–Japanese patent translation tasks were added. The size of the training dataset for each patent translation task was one million sentence pairs. For Japanese–English patent translation, the results confirmed that the translation quality of NMT and SMT outperformed the translation quality of RBMT. This was the first time that a corpus-based machine translation system yielded Japanese–English patent translation results comparable with those of RBMT systems. The translation quality of NMT evaluated by humans was higher than that of SMT for Japanese–English patent translation. For Japanese–English and English–Japanese scientific paper translation, pure NMT systems, not SMT reranking, achieved the best performance. In the field of machine translation, where large-scale parallel data was available, the mainstream technology for machine translation was changed from SMT to NMT. For English–Japanese patent translation, NMT achieved a translation quality close to that of the top SMT.

In the fourth workshop (WAT 2017) (Nakazawa et al. 2017), news translation tasks between Japanese and English and recipe translation tasks between Japanese and English were added. In Japanese–English patent translation, the results showed that 86% of translated sentences conveyed all the meanings of the source sentences for the top NMT system, which was trained using ten million parallel sentence pairs in addition to the shared task data of one million parallel sentence pairs. By contrast, for Japanese–English news translation, 5% of translated sentences conveyed all the meanings for the top NMT system. This percentage is substantially lower than that of the top system for Japanese–English patent translation. The small size of the

training data was one of the reasons. An essential reason was that the quality of the parallel translation of news was lower than that of patents. The reason for the low quality of parallel translation of news compared with that of patents is as follows: In patent applications, because the content in Japanese is translated literally to make an English version of the patent to file as a patent family, the translation quality at the sentence level is high. By contrast, news translation is not only translation but news writing. In news writing, writers select the content in consideration of the difference between readers of news in the source language and readers of news in the target language, and writers edit articles to change the structure to that of an English news structure. Thus, even if the sentences are aligned in same-topic bilingual news articles in Japanese and English, the parallel translation quality at the sentence level is lower than that of patents. It was shown that the translation of news with low-quality parallel data was a challenge for machine translation. Additionally, in the Chinese–Japanese patent translation task, 62% of translated sentences conveyed all the meanings of the source sentences. The performance improved from 29% in the previous year. Chinese–Japanese patent translation is in high demand in Japan.

In the fifth workshop (WAT 2018) (Nakazawa et al. 2018), the translation tasks between Myanmar and English, and between seven Indic languages and English were added. For Japanese–English scientific paper translation, the percentage of translated sentences that conveyed all the meanings of the source sentences improved from 34% in WAT 2017 to 61% in WAT 2018.

We have outlined research trends in machine translation, including patent translation from the activities of WAT. In the following, we describe other events. Google Translate changed from SMT to NMT in 2016. The change to NMT improved the translation quality, and people recognized the effectiveness of NMT. As a global trend, artificial intelligence (AI) technologies using deep learning have attracted attention since 2012. NMT is an AI technology. NMT's translation quality first caught up with SMT's translation quality in 2014, and NMT's translation quality has improved each year. There were very rapid advances in translation quality in the four years from 2015 to 2018.

Finally, we discuss the future of patent translation. Patent translation is an area in which large-scale high-quality parallel corpora are available. For example, a parallel corpus exists that contains over 100 million sentences.⁴ Although machine translation is not perfect, the translation quality of NMT will become close to translators for sentences without low-frequency words or new words as a result of training using a parallel corpus with the scale of 100 million sentence pairs. Because patent claims in Japanese have special styles, special pre-processing is necessary. The translation of sentences in claim sections is expected to be of high quality in the future. However, the translation of low-frequency words and new words is a problem that is difficult to solve using a corpus-based mechanism alone, and another approach will be necessary. Methods that use subword units, such as byte pair encoding (Sennrich et al. 2016), alleviate this problem. However, the translation of low-frequency words whose elements are not compositional and low-frequency subwords is still a problem.

⁴ALAGIN JPO corpus <https://alaginrc.nict.go.jp/resources/jpo-info/jpo-outline.html>.

There have been some studies on using automatically discovered bilingual words, and such techniques might be applied to NMT. Although machine translation may make errors, machine translation can do many things. Machine translation can be used for new translation needs that take advantage of its low cost and high speed. The patent offices of several countries, such as JPO, have already incorporated machine translation into their work. Machine translation has also been used in commercial services that provide foreign language patents in their customers' preferred language. Machine translation of patents will be used in society as an indispensable tool to overcome the language barrier in intellectual property.

Acknowledgements The author would like to thank Prof. Mikio Yamamoto for his support in writing this manuscript, particularly the chapters on NTCIR-7 and NTCIR-8. The author also thanks Prof. Doug Oard for his helpful comments. The author thanks Maxine Garcia, PhD, from Edanz Group for editing a draft of this manuscript.

References

- Brown PE, Della Pietra SA, Della Pietra VJ, Mercer RL (1993) The mathematics of statistical machine translation: parameter estimation. *Comput Linguist* 19(2):263–311
- Ehara T (2010) Machine translation for patent documents combining rule-based translation and statistical post-editing. In: *Proceedings of the NTCIR-8 workshop*, pp 384–386
- Feng M, Schmidt C, Wuebker J, Peitz S, Freitag M, Ney H (2011) The RWTH Aachen system for NTCIR-9 PatentMT. In: *Proceedings of the NTCIR-9 workshop*, pp 600–605
- Fujii A, Utiyama M, Yamamoto M, Utsuro T (2008) Overview of the patent translation task at the NTCIR-7 workshop. In: *Proceedings of the NTCIR-7 workshop*, pp 389–400
- Fujii A, Utiyama M, Yamamoto M, Utsuro T, Ehara T, Echizen-ya H, Shimohata S (2010) Overview of the patent translation task at the NTCIR-8 workshop. In: *Proceedings of the NTCIR-8 workshop*, pp 371–376
- Goto I, Lu B, Chow KP, Sumita E, Tsou BK (2011) Overview of the patent machine translation task at the NTCIR-9 workshop. In: *Proceedings of the NTCIR-9 workshop*, pp 559–578
- Goto I, Chow KP, Lu B, Sumita E, Tsou BK (2013) Overview of the patent machine translation task at the NTCIR-10 workshop. In: *Proceedings of the NTCIR-10 workshop*, pp 260–286
- Huang Z, Devlin J, Matsoukas S (2013) BBN's systems for the Chinese-English sub-task of the NTCIR-10 PatentMT. In: *Proceedings of the NTCIR-10 workshop*, pp 287–293
- Koehn P, Hoang H, Birch A, Callison-Burch C, Federico M, Bertoldi N, Cowan B, Shen W, Moran C, Zens R, Dyer C, Bojar O, Constantin A, Herbst E (2007) Moses: Open source toolkit for statistical machine translation. In: *Proceedings of the 45th annual meeting of the association for computational linguistics companion volume proceedings of the demo and poster sessions, association for computational linguistics, Prague, Czech Republic*, pp 177–180
- Lee YS, Xiang B, Zhao B, Franz M, Roukos S, Al-Onaizan Y (2011) IBM Chinese-to-English PatentMT system for NTCIR-9. In: *Proceedings of the NTCIR-9 workshop*, pp 606–613
- Ma J, Matsoukas S (2011) BBN's systems for the Chinese-English sub-task of the NTCIR-9 PatentMT evaluation. In: *Proceedings of the NTCIR-9 workshop*, pp 579–584
- Nakazawa T, Mino H, Goto I, Kurohashi S, Sumita E (2014) Overview of the 1st workshop on Asian translation. In: *Proceedings of the 1st workshop on Asian translation (WAT2014), workshop on Asian translation, Tokyo, Japan*, pp 1–19. <https://www.aclweb.org/anthology/W14-7001>

- Nakazawa T, Mino H, Goto I, Neubig G, Kurohashi S, Sumita E (2015) Overview of the 2nd workshop on Asian translation. In: Proceedings of the 2nd workshop on Asian translation (WAT2015), workshop on Asian translation, Kyoto, Japan, pp 1–28. <https://www.aclweb.org/anthology/W15-5001>
- Nakazawa T, Ding C, Mino H, Goto I, Neubig G, Kurohashi S (2016) Overview of the 3rd workshop on Asian translation. In: Proceedings of the 3rd workshop on Asian translation (WAT2016), the COLING 2016 organizing committee, Osaka, Japan, pp 1–46. <https://www.aclweb.org/anthology/W16-4601>
- Nakazawa T, Higashiyama S, Ding C, Mino H, Goto I, Kazawa H, Oda Y, Neubig G, Kurohashi S (2017) Overview of the 4th workshop on Asian translation. In: Proceedings of the 4th workshop on Asian translation (WAT2017), Asian federation of natural language processing, Taipei, Taiwan, pp 1–54. <https://www.aclweb.org/anthology/W17-5701>
- Nakazawa T, Sudoh K, Higashiyama S, Ding C, Dabre R, Mino H, Goto I, Pa WP, Kunchukuttan A, Kurohashi S (2018) Overview of the 5th workshop on Asian translation. In: Proceedings of the 32nd Pacific Asia conference on language, information and computation: 5th workshop on Asian translation: 5th workshop on Asian translation, association for computational linguistics, Hong Kong. <https://www.aclweb.org/anthology/Y18-3001>
- Papineni K, Roukos S, Ward T, Zhu WJ (2002) Bleu: a method for automatic evaluation of machine translation. In: Proceedings of 40th annual meeting of the association for computational linguistics, pp 311–318
- Sennrich R, Haddow B, Birch A (2016) Neural machine translation of rare words with subword units. In: Proceedings of the 54th annual meeting of the association for computational linguistics (Vol 1: Long Papers), association for computational linguistics, Berlin, Germany, pp 1715–1725. <https://doi.org/10.18653/v1/P16-1162>, <https://www.aclweb.org/anthology/P16-1162>
- Sudoh K, Duh K, Tsukada H, Nagata M, Wu X, Matsuzaki T, Tsujii J (2011) NTT-UT statistical machine translation in NTCIR-9 PatentMT. In: Proceedings of the NTCIR-9 workshop, pp 585–592
- Sudoh K, Suzuki J, Tsukada H, Nagata M, Hoshino S, Miyao Y (2013) NTT-NII statistical machine translation for NTCIR-10 PatentMT. In: Proceedings of the NTCIR-10 workshop, pp 294–300
- Utiyama M, Isahara H (2007) A Japanese-English patent parallel corpus. In: Proceedings of the machine translation summit XI, pp 475–482

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.