

Enabling Medical Translation for Low-Resource Languages

Ahmad Musleh[†], Nadir Durrani[†], Irina Temnikova[†], Preslav Nakov[†],
Stephan Vogel[†], and Osama Alsaad[‡]

[†]Qatar Computing Research Institute, HBKU [‡]Texas A&M University in Qatar
{amusleh, ndurrani, itemnikova, pnakov, svogel}@qf.org.qa
osama.alsaad@qatar.tamu.edu

Abstract. We present research towards bridging the language gap between migrant workers in Qatar and medical staff. In particular, we present the first steps towards the development of a real-world Hindi-English machine translation system for doctor-patient communication. As this is a low-resource language pair, especially for speech and for the medical domain, our initial focus has been on gathering suitable training data from various sources. We applied a variety of methods ranging from fully automatic extraction from the Web to manual annotation of test data. Moreover, we developed a method for automatically augmenting the training data with synthetically generated variants, which yielded a very sizable improvement of more than 3 BLEU points absolute.

Keywords: Machine Translation, medical translation, doctor-patient communication, resource-poor languages, Hindi.

1 Introduction

In recent years, Qatar’s booming economy has resulted in rapid growth in the number of migrant workers needed for the growing number of infrastructure projects. These workers, who mainly come from the southern parts of Asia, usually know little or no English and do not know any Arabic either. This results in a communication barrier between them and the natives. More serious situation arises in the case of medical emergency. This causes serious problems as the public administration and services in Qatar mostly use Arabic and English. According to a 2012 report by the Weill Cornell Medical College (WCMC¹) in Qatar [15], almost 78% of the patients visiting the Hamad Medical Corporation (HMC, the main health-care provider in Qatar) did not speak Arabic or English. The study also pointed out that the five most spoken languages in Qatar in 2012 were (in that order) Nepali, Urdu, Hindi, English, and Arabic. The report also pointed out that even though HMC currently uses medical interpreters to overcome this problem, their number is not sufficient. This has urged the authorities to look into technology for alternatives.

¹ <http://qatar-weill.cornell.edu>

In this paper, we propose a solution to bridge this language gap. We present our preliminary effort towards developing a Statistical Machine Translation (SMT) system for doctor-patient communication in Qatar. The success of a data-driven system largely depends upon the availability of in-domain data.

This makes our task non-trivial as we are dealing with a low-resourced language pair and furthermore with the medical domain. We decided to focus on Hindi, one of the languages under question. Our decision was driven by the fact, that Hindi and Urdu are closely-related languages, often considered dialects of each other, and people from Nepal and other South-Asian countries working in Qatar typically understand Hindi. Moreover, we have access to more Hindi-English parallel data than for any other language pair involving the top-5 most spoken languages in Qatar.

Our focus in this paper is on data collection and data generation (Sections 3 and 4). We collected data from various sources (Section 3), including (i) Wikimedia (Wikipedia, Wiktionary, and OmegaWiki) parallel English-Hindi data, (ii) doctor-patient dialogues from YouTube videos, and movie subtitles, and (iii) parallel medical terms from BabelNet and MeSH. Moreover, we synthesized Hindi-English parallel data from Urdu-English data, by translating the Urdu part into Hindi. The approach is described in Section 4. Our results show improvement of up to +1.45 when using synthesized data, and up to +1.66, when concatenating the mined dictionaries on top of the synthesized data (Section 5).

Moreover, Section 2 provides an overview of related work on machine translation for doctor-patient dialogues and briefly discusses the Machine Translation (MT) approaches for low-resource languages; Section 5 presents the results of the manual evaluation, and Section 6 provides the conclusions and discusses directions for future work.

2 Related Work

Below we first describe some MT applications for doctor-patient communication. Then, we present more general research on MT for the medical domain.

2.1 Bi-directional Doctor↔Patient Communication Applications

We will first describe the pre-existing MT systems for doctor-patient communication, particularly the ones that required data collection for under-resourced languages. Several MT systems facilitating doctor-patient communication have been built in the past [4, 13, 6, 14, 20, 17]. Most of them are still prototypes, and only few have been fully deployed. Some of these systems work with under-resourced languages [4, 14, 20, 17]. Moreover, their solution relies on mapping the utterances to an interlingua, instead of using SMT.

MedSLT [4] is an interlingua-based speech-to-speech translation system. It covers a restricted set of domains, and covers English, French, Japanese, Spanish, Catalan, and Arabic. The system can translate doctor’s questions/statements to the patient, but not the responses by the patient back to the doctor.

Converser [6] is a commercial doctor-patient speech-to-speech bidirectional MT system for the English \leftrightarrow Spanish language pair. It has been deployed in several US hospitals and has the following features: users can correct the automatic speech recognition (ASR) and MT outputs, back-translation (re-translation of the translation) to the user is made to allow this. The system maps concepts to a lexical database specially created from various sources, and also allows “translation shortcuts” (i.e., translation memory of previous translations that do not need verification).

Jibbig [13] is a travel and medical speech-to-speech MT system, deployed as an iPhone mobile app. Jibbig allows English-to-Spanish and Spanish-to-English medical speech-to-speech translation.

S-MINDS [14] is a two-way doctors-patient MT system, which uses an in-house “interpretation” software. It matches the ASR utterances to a finite set of concepts in the specific semantic domain and then paraphrases them. In case the utterance cannot be matched, the system uses an SMT engine.

Accultran [20] is an automatic translation system prototype, which features back translation to the doctor and yes/no or multiple-choice questions (MCQs) to the patient. It allows the doctor to confirm the translation to the patient, and has a cross cultural adviser. It flags sensitive utterances that are difficult to translate to the patient. The system maps the utterances to SNOMED-CT or Clinical Document Architecture (CDA-2) standards, which are used as an interlingua.

IBM MASTOR [17] is a speech-to-speech MT system for two language pairs (English-Mandarin and English-dialectal Arabic), which relies on ASR, SMT, and Speech Synthesis components. It works both on laptops and PDAs.

English-Portuguese SLT [37] is an English-Portuguese speech-to-speech MT system, composed of an ASR, MT (relying on HMM) and speech synthesis. It works as an online service and as a mobile application.

None of the above systems handles the top-5 languages of interest to Qatar.

2.2 Uni-directional Doctor \leftrightarrow Patient Communication Applications

Besides the above-described MT systems, there are a number of mobile or web applications, which are based on pre-translated phrases. The phrases are pre-recorded by professionals or native speakers and can be played to the patient. Most of these applications work only in the doctor-to-patient direction. The most popular ones are UniversalDoctor², MediBabble³, Canopy⁴, MedSpeak, Mavro Emergency Medical Spanish⁵, and DuoChart⁶.

Unfortunately, these applications do not allow free, unseen, or spontaneous translations, and do not cover the language pairs of interest for Qatar. Moreover, some of them (e.g., UniversalDoctor) require a paid subscription.

² <http://www.universaldocor.com>

³ <http://medibabble.com>

⁴ <http://www.canopyapps.com>

⁵ <http://mavroinc.com/medical.html>

⁶ <http://duochart.com>

2.3 General Research in Medical Machine Translation

A number of systems have been developed and participated in the WMT'14 Medical translation task. It is a Cross-Lingual Information Retrieval (CLIR) task divided into two sub-tasks: (i) translation of user search queries, and (ii) translation of summaries of retrieved documents:

A system described in [12], part of the Khresmoi project⁷, uses the phrase-based Moses and standard methods for domain adaptation. [25] also uses the phrase-based Moses system and achieved the highest BLEU score for the English-German intrinsic query translation evaluation. Another system [26] combined web-crawled in-domain monolingual data and a bilingual lexicon in order to complement the limited in-domain parallel corpora. A third one [34] proposed a terminology translation system for the query translation subtask and used 6 different methods for terminology extraction. A fourth system [35] used a combination of n-gram based NCODE and phrase-based Moses to the subtask of sentence translation. The system of [40] applied a combination of domain adaptation techniques on the medical summary sentence translation task and achieved the first and the second best BLEU scores. Then, the system of [44] used the Moses phrase-based system and worked on the medical summary WMT'14 task and experimented with translation models, re-ordering models, operation sequence models, and language models, as well as with data selection. A study on quality analysis of machine translation systems in medical domain was carried in [27]. Most of this work focused on European language pairs and did not cover languages of interest to us, nor did it involve low-resource languages in general.

3 Data Collection

As the main problem of low resource languages is data collection [24, 30], we have adopted a variety of approaches, in order to collect as much parallel English-Hindi data as possible.

3.1 Wiki Dumps

We downloaded, extracted and mined all language links from Wikipedia,⁸ Wiktionary,⁹ and OmegaWiki¹⁰ in order to provide a one-to-one word mapping from English into Hindi. We then extracted page links and language links from Wikipedia and Wiktionary. Moreover, we used OmegaWiki to provide a bilingual word dictionary containing the word, its synonyms, its translation, and its lexical, terminological and ontological forms. We extracted the data using two OmegaWiki sources: bilingual dictionaries and an SQL database dump. Table 1 shows the number of Hindi words we collected from all three sources.

⁷ <http://www.khresmoi.eu>

⁸ <http://www.wikipedia.org>

⁹ <http://www.wiktionary.org>

¹⁰ <http://www.omegawiki.org>

Source	Word Pairs
Wikipedia	40,764
Wiktionary	10,352
OmegaWiki	3,476

Table 1. Number of word translation pairs collected from Wikimedia sources.

3.2 Doctor-Patient YouTube Videos and Movie Subtitles

We used a mixed approach to extract doctor-patient dialogues from medical YouTube subtitles. As using the YouTube-embedded automatic subtitling is inefficient, the dialogues of the videos were first extracted by manually typing the audio found in the videos. However, as this process was very time consuming, we started a screenshot session in order to collect all the visual representations of the subtitles. Next, the subtitles were extracted using Tesseract,¹¹ an open source Optical Character Recognition (OCR) reader provided by Google, on the screenshots captured. The subtitles were then manually corrected, translated into Hindi using Google Translate, and post-edited by a Hindi native speaker. This resulted in a parallel corpus of medical dialogues with 11,000 Hindi words (1,200 sentences). These sentences were later used for tuning and testing our MT system. Additionally, we used a web crawler to extract a small number of non-medical parallel English-Hindi movie subtitles (nine movies) from Open Subtitles.¹²

3.3 BabelNet and MeSH

We extracted medical terms from BabelNet [32] using their API. As Medical Subject Headings (MeSH¹³) represents the largest source of Medical terms, we downloaded their dumps and extracted the 198,958 MeSH terms which we overlapped with the previously mined results of Wiki Dumps.

4 Data Synthesis

Hindi and Urdu are closely-related languages that share grammatical structure and largely overlap in vocabulary. This provides strong motivation to transform an Urdu-English parallel data into Hindi-English by translating the Urdu part into Hindi. We made use of the Urdu-English segment of the Indic multi-parallel corpus [36], which contains about 87,000 sentence pairs. The Hindi-English segment of this corpus is a subset of the parallel data that was made available for the WMT’14 translation task, but its English side is completely disjoint from the English side of the Urdu-English segment.

¹¹ <https://github.com/tesseract-ocr>

¹² <http://www.opensubtitles.com>

¹³ <http://www.ncbi.nlm.nih.gov/mesh>

Initially, we trained an Urdu-to-Hindi SMT system using the tiny EMILLE¹⁴ corpus [1]. However, we found this system to be useless for translating the Urdu part of the Indic data due to domain mismatch and the high proportion of Out-of-Vocabulary (OOV) words (approximately 310,000 tokens). Thus, in order to reduce data sparseness, we synthesized additional phrase tables using interpolation and transliteration.

4.1 Interpolation

We built two phrase translation tables $p(\bar{u}_i|\bar{e}_i)$ and $p(\bar{e}_i|\bar{h}_i)$, from Urdu-English (Indic corpus) and Hindi-English (HindEnCorp [3]) bitexts. Given the phrase table for Urdu-English $p(\bar{u}_i|\bar{e}_i)$ and the phrase table for English-Hindi $p(\bar{e}_i|\bar{h}_i)$, we induced an Urdu-Hindi phrase table $p(\bar{u}_i|\bar{h}_i)$ using the model [39, 43]:

$$p(\bar{u}_i|\bar{h}_i) = \sum_{\bar{e}_i} p(\bar{u}_i|\bar{e}_i)p(\bar{e}_i|\bar{h}_i) \quad (1)$$

The number of entries in the baseline Urdu-to-Hindi phrase table were approximately 254,000. Using interpolation, we were able to build a phrase table containing roughly 10M phrases. This reduced the number of OOV tokens from 310K to approximately 50,000.

4.2 Transliteration

As Urdu and Hindi are written in different scripts (Arabic and Devanagiri, respectively), we added a transliteration component to our Urdu-to-Hindi system. While it can be used to translate all 50,000 OOV words, previous research has shown that transliteration is useful for more than just translating OOV words when translating closely related language pairs [29, 31, 38, 41, 42]. Following [10], we transliterate all Urdu words to Hindi and hypothesize n-best transliterations, along with regular translations. The idea is to generate novel Hindi translations that may be absent from the regular and interpolated phrase table, but for which there is evidence in the language model. Moreover, the overlapping evidence in the translation and transliteration phrase tables improves the overall system.

We learn an unsupervised transliteration model [11] from the word-alignments of Urdu-Hindi parallel data. We were able to extract around 2,800 transliteration pairs. To learn a richer transliteration model, we additionally fed the interpolated phrase table, as described above, to the transliteration miner. We were able to mine about 21,000 additional transliteration pairs and to build an Urdu-Hindi character-based model from it. In order to fully capitalize on the large overlap in Hindi-Urdu vocabulary, we transliterated each word in the Urdu test data to Hindi and we produced a phrase table with 100-best transliterations. We then used the two synthesized (triangulated and transliterated) phrase tables along with the baseline Urdu-to-Hindi phrase table in a log-linear model.

¹⁴ EMILLE contains about 12,000 sentences of comparable data in Hindi and Urdu. We were able to align about 7,000 sentences to build an Urdu-to-Hindi system.

Table 2 shows development results from training an Urdu-to-Hindi SMT system. By adding interpolated phrase tables and transliteration, we obtain a very sizable gain of +3.35 over the baseline Urdu→Hindi system. Using our best Urdu-to-Hindi system ($\mathbf{B}_{u,h}\mathbf{T}_g\mathbf{T}_r$), we translated the Urdu part of the multi-Indic corpus to form a Hindi-English bi-text. This yielded a synthesized bi-text of $\approx 87,000$ Hindi-English sentence pairs. Detailed analysis can be found in [8].

System	PT	Tune	Test	System	Tune	Test	System	Tune	Test
$\mathbf{B}_{u,h}$	254K	34.18	34.79	$\mathbf{B}_{u,h}\mathbf{T}_g$	37.65	37.58	$\mathbf{B}_{u,h}\mathbf{T}_g\mathbf{T}_r$	38.0	37.99
\mathbf{T}_g	10M	15.60	15.34	$\mathbf{B}_{u,h}\mathbf{T}_r$	34.77	35.76			
\mathbf{T}_r		9.54	9.93	$\mathbf{T}_g\mathbf{T}_r$	17.63	18.11			

Table 2. Evaluating triangulated and transliterated phrase tables for Urdu-to-Hindi SMT. Notation: $\mathbf{B}_{u,h}$ = Baseline Phrase Table, \mathbf{T}_g = Triangulated Phrase Table, and \mathbf{T}_r = Transliteration Phrase Table.

5 Experiments and Evaluation

5.1 Machine Translation

5.1.1 Baseline Data We trained the Hindi-English systems using Hindi-English parallel data [3] composed by compiling several sources including the Hindi-English segment of the Indic parallel corpus. It contains 287,202 parallel sentences, and 4,296,007 Hindi and 4,026,860 English tokens. We used 635 sentences (6,111 tokens) for tuning and 636 sentences (5,231) for testing, collected from doctor-patient communication dialogues in YouTube videos. The sentences were translated into Hindi by a human translator. We trained interpolated language models using all the English and Hindi monolingual data made available for the WMT’14 translation task: 287.3M English and 43.4M Hindi tokens.

5.1.2 Baseline System We trained a phrase-based system using Moses [22] with the following settings: a maximum sentence length of 80, GDFSA symmetrization of GIZA++ alignments [33], an interpolated Kneser-Ney smoothed 5-gram language model with KenLM [19] used at runtime, 100-best translation options, MBR decoding [23], Cube Pruning [21] using a stack size of 1,000 during tuning and 5,000 during testing. We tuned with the k -best batch MIRA [5]. We additionally used msd-bidirectional-fe lexicalized reordering, a 5-gram OSM [7], class-based models [9]¹⁵ sparse lexical and domain features [18], a distortion limit of 6, and the no-reordering-over-punctuation heuristic. We used an unsupervised transliteration model [11] to transliterate the OOV words. These are state-of-the-art settings, as used in [2].

Pair	\mathbf{B}_0	+Syn	Δ	+PT	Δ
hi-en	21.28	22.67	+1.39	22.1	+0.82
en-hi	22.52	23.97	+1.45	23.28	+0.76

Table 3. Results using synthesized (Syn) Hindi-English parallel data. Notation used: \mathbf{B}_0 = System without synthesized data, +PT = System using synthesized data as an additional phrase table.

5.1.3 Results Table 3 shows the results when adding the synthesized Hindi-English bi-text on top of the baseline system (+Syn). The synthesized data was simply concatenated to the baseline data to train the system. We also tried building phrase tables (+PT) separately from the baseline data and from the synthesized one and used as a separate features in the log-linear model as done in [28, 29, 41, 42]. We found that concatenating synthetic data with the baseline data directly was superior to training a separate phrase-table from it. We obtained improvements of up to +1.45 by adding synthetic data.

Table 4 shows the results from adding the mined dictionaries to the baseline system. The baseline system (\mathbf{B}_0) used in this case is the best system in Table 3. Again, we simply concatenated the dictionaries with the baseline data and we gained improvements of up to +1.66 BLEU points absolute. Cumulatively, by using dictionaries and synthesized phrase-tables, we were able to obtain statistically significant improvements of more than 3 BLEU points.

Pair	\mathbf{B}_0	+Dict	Δ
hi-en	22.67	23.29	+0.62
en-hi	23.97	25.63	+1.66

Table 4. Evaluating the effect of Dictionaries. \mathbf{B}_0 = System without dictionaries.

5.2 Manual Evaluation

In addition to the above evaluation, we ran a small manual evaluation experiment, using the Appraise platform [16]. The two sections below describe the results of the Hindi-to-English (Section 5.2.1) and the English-to-Hindi (Section 5.2.2) evaluations.

5.2.1 Hindi-to-English The evaluation was conducted by 3 monolingual English speakers, using 321 randomly selected sentences, divided into three batches of evaluation. Similar to the setup at evaluation campaigns such as WMT, the evaluators were shown the translations and references.

¹⁵ We used mkcls to cluster the data into 50 clusters.

The evaluators were asked to assign one of the following three categories to each translation: (a) *helpful in this situation*, (b) *misleading*, and (c) *doubtful that people will understand it*. As shown in Table 5, over 37% of the cases were classified as *helpful* (good translations), 39% as *doubtful* (mediocre), and 24% as *misleading* (really bad translations). Annotators did not always agree, e.g., Judge 1 and Judge 2 were more lenient than Judge 3.

Response	Judge 1	Judge 2	Judge 3	Total	Percentage
Helpful	129	129	96	354	37%
Doubtful	114	145	118	377	39%
Misleading	78	47	107	232	24%

Table 5. Manual sentence evaluation for Hindi-to-English translation.

5.2.2 English-to-Hindi In order to check the output of the English-to-Hindi system, we asked a bilingual judge to evaluate 328 sentences. She was asked to classify the sentences in the same categories as for the Hindi-English evaluation. Table 6 shows the results; we can see that 55.8% of the sentences were found *helpful in this situation*. This is hardly because English-Hindi system was any better, but more likely because the human evaluator was lenient. Unfortunately, we could not find a second Hindi speaker to evaluate our translations, and thus we could not calculate inter-annotator agreement.

Response	Number	Percentage
Helpful	183	55.8%
Doubtful	111	33.8%
Misleading	34	10.4%

Table 6. Manual sentence evaluation for English-to-Hindi translation.

5.2.3 Analysis In order to understand the problems with the Hindi output, we conducted an error analysis on 100 sentences classifying the errors into the following categories:

- missing/untranslated words;
- wrongly translated words;
- word order problems;
- other error types, e.g., extra words.

Table 7 shows the results. We can see that most of the problems are associated with *word order problems* (84%) or *wrongly translated words* (74%).

Category	Percentage
Missing/Untranslated words	45%
Wrongly translated words	74%
Word order problems	84%
Other types of errors	13%

Table 7. Error analysis results for our English-to-Hindi translation.

6 Conclusions

We presented our preliminary efforts towards building a Hindi↔English SMT system for facilitating doctor-patient communication. We improved our baseline system using two approaches, namely (i) additional data collection, and (ii) automatic data synthesis. We mined useful dictionaries from Wikipedia in order to improve the coverage of our system. We made use of the relatedness between Hindi and Urdu to generate synthetic Hindi-English bi-texts by automatically translating 87,000 Urdu sentences into Hindi. Both our data collection and our synthesis approach worked well and have shown significant improvements over the baseline system, yielding a total improvement of +3.11 BLEU points absolute for English-to-Hindi and +2.07 for Hindi-to-English. We also carried out human evaluation for the best system. In the error analysis of the Hindi outputs, we found that most errors were due to ordering of the words in the output, or to wrong lexical choice.

In future work, we plan to collect more data for Hindi, but also to synthesize Urdu data. We further plan to develop a system for Nepali-English. Finally, we would like to add Automatic Speech Recognition (ASR) and Speech Synthesis components in order to build a fully-functional speech-to-speech system, which we would test and gradually deploy for use in real-world scenarios.

Acknowledgments

The authors would like to thank Naila Khalisha and Manisha Bansal for their contributions towards the project.

References

1. Paul Baker, Andrew Hardie, Tony McEnery, Hamish Cunningham, and Robert J. Gaizauskas. EMILLE, A 67-Million Word Corpus of Indic Languages: Data Collection, Mark-up and Harmonisation. In *Proceedings of the Third International Language Resources and Evaluation Conference, LREC'02*, Las Palmas, Canary Islands, Spain, 2002.
2. Alexandra Birch, Matthias Huck, Nadir Durrani, Nikolay Bogoychev, and Philipp Koehn. Edinburgh SLT and MT system description for the IWSLT 2014 evaluation. In *Proceedings of the 11th International Workshop on Spoken Language Translation, IWSLT '14*, Lake Tahoe, CA, USA, 2014.

3. Ondřej Bojar, Vojtěch Diatka, Pavel Rychlý, Pavel Straňák, Aleš Tamchyna, and Dan Zeman. Hindi-English and Hindi-only Corpus for Machine Translation. In *Proceedings of the Ninth International Language Resources and Evaluation Conference, LREC'14*, pages 3550–3555, Reykjavik, Iceland, 2014.
4. Pierrette Bouillon, Glenn Flores, Maria Georgescu, Ismahene Sonia Halimi Mallem, Beth Ann Hockey, Hitoshi Isahara, Kyoko Kanzaki, Yukie Nakao, Emmanuel Rayner, Marianne Elina Santaholma, Marianne Starlander, and Nikos Tsourakis. Many-to-many multilingual medical speech translation on a PDA. In *Proceedings of the Eighth Conference of the Association for Machine Translation in the Americas, AMTA'08*, pages 314–323, Waikiki, Hawaii, USA, 2008.
5. Colin Cherry and George Foster. Batch Tuning Strategies for Statistical Machine Translation. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, HLT-NAACL'12*, pages 427–436, Montréal, Canada, 2012.
6. Mike Dillinger and Mark Seligman. Converser: Highly Interactive Speech-to-Speech Translation for Healthcare. In *Proceedings of the COLING-ACL 2006 Workshop on Medical Speech Translation*, pages 36–39, Sydney, Australia, 2006.
7. Nadir Durrani, Alexander Fraser, Helmut Schmid, Hieu Hoang, and Philipp Koehn. Can Markov Models Over Minimal Translation Units Help Phrase-Based SMT? In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, Sofia, Bulgaria, 2013. Association for Computational Linguistics.
8. Nadir Durrani and Philipp Koehn. Improving Machine Translation via Triangulation and Transliteration. In *Proceedings of the 17th Annual Conference of the European Association for Machine Translation, EAMT'14*, pages 71–78, Dubrovnik, Croatia, 2014.
9. Nadir Durrani, Philipp Koehn, Helmut Schmid, and Alexander Fraser. Investigating the Usefulness of Generalized Word Representations in SMT. In *Proceedings of the 25th Annual Conference on Computational Linguistics, COLING'14*, pages 421–432, Dublin, Ireland, 2014.
10. Nadir Durrani, Hassan Sajjad, Alexander Fraser, and Helmut Schmid. Hindi-to-Urdu Machine Translation through Transliteration. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL'14*, pages 465–474, Uppsala, Sweden, 2010.
11. Nadir Durrani, Hassan Sajjad, Hieu Hoang, and Philipp Koehn. Integrating an Unsupervised Transliteration Model into Statistical Machine Translation. In *Proceedings of the 15th Conference of the European Chapter of the ACL, EACL'14*, pages 148–153, Gothenburg, Sweden, 2014.
12. Ondřej Dušek, Jan Hajic, Jaroslava Hlaváčová, Michal Novák, Pavel Pecina, Rudolf Rosa, Aleš Tamchyna, Zdenka Urešová, and Daniel Zeman. Machine Translation of Medical Texts in the Khresmoi Project. In *Proceedings of the 52nd Annual Meeting of the Association of Computational Linguistics, ACL'14*, pages 221–228, Baltimore, Maryland, USA, 2014.
13. Matthias Eck, Ian Lane, Ying Zhang, and Alex Waibel. Jibbig: Speech-to-Speech Translation on Mobile Devices. In *Proceedings of IEEE Spoken Language Technology Workshop, SLT2010*, pages 165 – 166, Berkeley, California, USA, 2010.
14. Farzad Ehsani, Jim Kimzey, Demitrios Master, Karen Sudre, and Hunil Park. Speech to Speech Translation for Medical Triage in Korean. In *Proceedings of the COLING-ACL 2006 Workshop on Medical Speech Translation*, pages 13–19, New York City, New York, USA, 2006.

15. Maha Elnashar, Huda Abdelrahim, and Michael D Fetters. Cultural Competence Springs Up in the Desert: The Story of the Center for Cultural Competence in Health Care at Weill Cornell Medical College in Qatar. *Academic Medicine*, 87(6):759–766, 2012.
16. Christian Federmann. Appraise: An Open-source Toolkit for Manual Evaluation of MT Output. *The Prague Bulletin of Mathematical Linguistics*, 98:25–35, 2012.
17. Yuqing Gao, Liang Gu, Bowen Zhou, Ruhi Sarikaya, Mohamed Afify, Hong-Kwang Kuo, Wei-Zhong Zhu, Yonggang Deng, Charles Prosser, Wei Zhang, et al. IBM MASTOR System: Multilingual Automatic Speech-to-Speech Translator. In *Proceedings of the COLING-ACL 2006 Workshop on Medical Speech Translation*, pages 53–56, Sydney, Australia, 2006.
18. Eva Hasler, Barry Haddow, and Philipp Koehn. Sparse Lexicalised features and Topic Adaptation for SMT. In *Proceedings of the Seventh International Workshop on Spoken Language Translation*, IWSLT’12, pages 268–275, Hong Kong, China, 2012.
19. Kenneth Heafield. KenLM: Faster and Smaller Language Model Queries. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, WMT’11, pages 187–197, Edinburgh, Scotland, United Kingdom, 2011.
20. Daniel T. Heinze, Alexander Turchin, and V. Jagannathan. Automated Interpretation of Clinical Encounters with Cultural Cues and Electronic Health Record Generation. In *Proceedings of the COLING-ACL 2006 Workshop on Medical Speech Translation*, pages 20–27, Sydney, Australia, 2006.
21. Liang Huang and David Chiang. Forest Rescoring: Faster Decoding with Integrated Language Models. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, ACL’07, pages 144–151, Prague, Czech Republic, 2007.
22. Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. Moses: Open Source Toolkit for Statistical Machine Translation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, ACL’07, pages 177–180, Prague, Czech Republic, 2007.
23. Shankar Kumar and William J. Byrne. Minimum Bayes-Risk Decoding for Statistical Machine Translation. In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, HLT-NAACL’04, pages 169–176, Boston, Massachusetts, USA, 2004.
24. William D. Lewis, Robert Munro, and Stephan Vogel. Crisis MT: Developing a Cookbook for MT in Crisis Situations. In *Proceedings of the EMNLP’11 Sixth Workshop on Statistical Machine Translation*, WMT’11, pages 501–511, Edinburgh, Scotland, United Kingdom, 2011.
25. Jianri Li, Se-Jong Kim, Hwidong Na, and Jong-Hyeok Lee. Postechs System Description for Medical Text Translation Task. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, WMT’14, pages 229–233, Baltimore, Maryland, USA, 2014.
26. Yi Lu, Longyue Wang, Derek F Wong, Lidia S. Chao, Yiming Wang, and Francisco Oliveira. Domain Adaptation for Medical Text Translation Using Web Resources. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, WMT’14, pages 233–238, Baltimore, Maryland, USA, 2014.
27. Marta R. Costa-Jussa Mireia Farrus, Jordi Serrano Pons. Machine Translation in Medicine. A Quality Analysis of Statistical Machine Translation in the Medical Domain. In *Proceedings of the 1st Virtual International Conference on Advanced Research in Scientific Areas*, ARSA’12, pages 1995–1998, 2012.

28. Preslav Nakov. Improving English-Spanish Statistical Machine Translation: Experiments in Domain Adaptation, Sentence Paraphrasing, Tokenization, and Re-casing. In *Proceedings of the Third Workshop on Statistical Machine Translation*, WMT'08, pages 147–150, Columbus, Ohio, USA, 2008.
29. Preslav Nakov and Hwee Tou Ng. Improved Statistical Machine Translation for Resource-poor Languages Using Related Resource-rich Languages. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 3*, EMNLP '09, pages 1358–1367, Singapore, 2009.
30. Preslav Nakov and Hwee Tou Ng. Improving statistical machine translation for a resource-poor language using related resource-rich languages. *Journal of Artificial Intelligence Research (JAIR)*, pages 179–222, 2012.
31. Preslav Nakov and Jörg Tiedemann. Combining Word-level and Character-level Models for Machine Translation Between Closely-related Languages. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers - Volume 2*, ACL'12, pages 301–305, Jeju Island, Korea, 2012.
32. Roberto Navigli and Simone Paolo Ponzetto. BabelNet: The Automatic Construction, Evaluation and Application of a Wide-Coverage Multilingual Semantic Network. *Artificial Intelligence*, 193:217–250, 2012.
33. Franz J. Och and Hermann Ney. A Systematic Comparison of Various Statistical Alignment Models. In *In Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, ACL'03, pages 19–51, Sapporo, Japan, 2003.
34. Tsuyoshi Okita, Ali Hosseinzadeh Vahid, Andy Way, and Qun Liu. The DCU Terminology Translation System for the Medical Query Subtask at WMT14. In *Proceedings of the ACL 2014 Workshop on Statistical Machine Translation*, WMT'14, pages 239–245, Baltimore, Maryland, USA, 2014.
35. Nicolas Pécheux, Li Gong, Quoc Khanh Do, Benjamin Marie, Yulia Ivanishcheva, Alexandre Allauzen, Thomas Lavergne, Jan Niehues, Aurélien Max, and Francois Yvon. LIMSIS@ WMT14 Medical Translation Task. In *Proceedings of the ACL 2014 Workshop on Statistical Machine Translation*, WMT'14, pages 246–253, Baltimore, Maryland, USA, 2014.
36. Matt Post, Chris Callison-Burch, and Miles Osborne. Constructing Parallel Corpora for Six Indian Languages via Crowdsourcing. In *Proceedings of the Seventh Workshop on Statistical Machine Translation*, WMT'12, pages 401–409, Montréal, Canada, 2012.
37. João António Santos Gomes Rodrigues. *Speech-to-Speech Translation to Support Medical Interviews*. PhD thesis, Universidade de Lisboa, Portugal, 2013.
38. Jörg Tiedemann and Preslav Nakov. Analyzing the Use of Character-Level Translation with Sparse and Noisy Datasets. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, RANLP'13, pages 676–684, Hissar, Bulgaria, 2013.
39. Masao Utiyama and Hitoshi Isahara. A Comparison of Pivot Methods for Phrase-Based Statistical Machine Translation. In *Proceedings of the 2007 Meeting of the North American Chapter of the Association for Computational Linguistics*, NAACL'07, pages 484–491, Rochester, New York, USA, 2007.
40. Longyue Wang, Yi Lu, Derek F Wong, Lidia S Chao, Yiming Wang, and Francisco Oliveira. Combining Domain Adaptation Approaches for Medical Text Translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, WMT'14, pages 254–259, Baltimore, Maryland, USA, 2014.
41. Pidong Wang, Preslav Nakov, and Hwee Tou Ng. Source Language Adaptation for Resource-Poor Machine Translation. In *Proceedings of the 2012 Joint Conference*

- on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, EMNLP-CoNLL'12, pages 286–296, Jeju Island, Korea, 2012.
42. Pidong Wang, Preslav Nakov, and Hwee Tou Ng. Source Language Adaptation Approaches for Resource-Poor Machine Translation. *Computational Linguistics*, pages 1–44, 2016.
 43. Hua Wu and Haifeng Wang. Pivot Language Approach for Phrase-Based Statistical Machine Translation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, ACL'07, pages 856–863, Prague, Czech Republic, 2007.
 44. Jian Zhang, Xiaofeng Wu, Iacer Calixto, Ali Hosseinzadeh Vahid, Xiaojun Zhang, Andy Way, and Qun Liu. Experiments in Medical Translation Shared Task at WMT 2014. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, WMT'14, pages 260–265, Baltimore, Maryland, USA, 2014.