MOSES CORE

2013

Annual Public Report

TAUS

Charles University, Prague

THE UNIVERSITY of EDINBURGH

CAPITA

FONDAZIONE BRUNO KESSLER

1 Overview

The diversity of languages in Europe makes translation vitally important to the economic, cultural and social lives of Europeans. Machine translation (MT) provides a way of fully or partially automating the translation process, and hence reducing the costs and enabling more text and speech to be translated.

Machine translation, however, is a complex field and presents many substantial barriers for entry to potential researchers, and users of the technologies. The principal aim of MosesCore is to reduce these barriers, making it easier to join and participate in the MT research community, and to become an MT user.

MosesCore achieves these aims by organising a variety of events targeted at users, developers and researchers of MT, and by promoting and coordinating the development and use of open-source MT tools, in particular the Moses toolkit.

In this report we will describe the events organised by MosesCore during 2013, as well as the main developments in Moses during this timeframe.

1.1 Key Facts

Project type	FP7 Coordination Action
Duration	February 1st 2012 - January 31st 2015
Financing	€1.2
Contact	Barry Haddow (info@mosescore.eu)

1.2 Partners

University of Edinburgh	United Kingdom
TAUS	Netherlands
Charles University, Prague	Czech Republic
Fondazione Bruno Kessler	Italy
Capita Translation and Interpreting (formerly Applied Language Solutions)	United Kingdom

1.3 Beneficiaries

Researchers have events in which they can showcase their research, compare their systems with others, and gather to implement new MT tech-

niques. They also have a state-of-the-art open-source platform to test out their ideas on.

Users and Developers have a stable and well supported open-source MT toolkit, and have forums to learn about new research developments in MT and share system building and deployment experience.

Everyone benefits from improved information exchange between developers, users and researchers.

2 Events

2.1 Machine Translation Marathon

In September 2013, the Machine Translation Marathon[1] took place in Prague, organised by Charles University. This week-long gathering of MT researchers and developers was the eighth such event to be organised, and the second by MosesCore. About 70 MT researchers, developers and users participated in the event.

At the MTM, participants spent the week working on a range of open-source MT projects, as well as hearing invited talks and tutorials from leading MT researchers, and paper presentations on new open-source MT tools. There was also a "summer school" of lectures and labs for those that are relatively new to the field. At this year's MT marathon we hosted a round table on Moses for commercial users[2], organised by TAUS.

The summer school was a fairly comprehensive overview of MT, starting with its history and covering all the current statistical paradigms and the integration of MT with translation tools. The course was given by MosesCore partners (FBK, Edinburgh and Charles University) , and invited lecturers from CMU, JHU, and the Universities of Sheffield and Zurich. As well as labs to accompany the course, we also held tutorials on the use of Charles University's syntactic MT tools, and the open-source rule-based translation system Apertium (run by Francis Tyers).

The invited talks in the MTM covered cutting edge research topics, as well as examples of practical usage of MT. Phil Blunsom (Oxford) gave a talk on deep learning (one of the hot topics in natural language processing for 2013)

[1] http://www.statmt.org/mtm13

[2] http://www.taus.net/events/moses-roundtable

and how his group is applying it to MT. There was also a talk from Alex Fraser (Munich) summarising the work of the 2012 John Hopkins workshop on domain adaptation in MT, and one from Kristina Toutanova (Microsoft Research) surveying methods to address translation of morphologically rich languages. Francis Tyers (now at the University of Tromsø) gave a talk on the usage of MT inside the EU, and Bonnie Dorr surveyed the US DARPA-funded research in MT.

For many of the Marathon's attendees though, the main business of the week was hacking the week-long projects. These are proposed on the Monday morning, and after a negotiation session, developers form into small groups to work on their projects. Projects may involve working on a new research idea, or creating a new implementation of a recently published idea, or improving an existing open-source tool. The projects are very varied in their scope, aims and results, but possibly their main impact is to bring members of the community together to discuss implementation in a very practical setting. This year the projects covered such topics as quality estimation, language modelling, translation of social media and improvements to syntax-based MT.

Figure 1: The MT Marathon in Prague

2.2 Industrial Outreach Events

To continue the work started in 2012, this year our activities were aimed at outlining the latest market developments and placing Moses adoption within

the context of other MT adoptions and offerings. A number of dedicated events, web and e-campaigns helped us to reach the goals and engage the target audience.

2.2.1 MT Showcases

In light of our key learning points from the previous events, in 2013 we have repositioned the outreach workshops as MT Showcases, where a big share of the content is focused on Moses and/or open source MT.

We organised the following TAUS Machine Translation Showcases:

- October 10, 2013. Hosted by LocWorld, Santa Clara, CA (USA)
  Use cases: KantanMT, Microsoft, Welocalize, Tilde and TAUS
- June 12, 2013, London (UK)
  Use cases: University of Edinburgh, Safaba, LexWorks, Pangeanic and TAUS
- April 10, 2013. Hosted by Localization World Singapore
  Use cases: Institute for Infocomm, Google, Asia Online, Precision Translation Tools, Hunnect and TAUS

All the uses cases presented at the events can be found on the Moses project website[3]

In 2013 about 100 people participated in the TAUS MT Showcases. Twenty-one presentations are publicly available on Slideshare. These use cases have been viewed a total of 5,801 times at 4 November 2013. On average each presentation has been viewed 276 times.

2.2.2 Moses Round Table

TAUS Moses Round Table, September 11, 2013, Prague (Czech Republic)
An industry Round Table, a half-day event at the MT Marathon in Prague, September 2013 was initiated to continue higher-level industry discussions started in Moses reports in 2011[4] and 2012[5].

The round table had broad participation from industry participants, government, representatives of the Moses core team and researchers interested in the topic.

[3] http://www.statmt.org/mosescore/index.php?n=Main.Videos
[4] https://www.taus.net/reports/moses-users-experiences-and-future-requirements
[5] https://www.taus.net/reports/moses-users-changing-priorities

2.3 Workshop on Machine Translation

The Workshop in Machine Translation (WMT) and its associated collection of shared tasks is one of the highlights of the MT research calendar. The shared tasks allow researchers from around the world to compare their techniques against other groups, using standard benchmarks, and to report their results at the workshop. The data used for the shared tasks is made available for researchers in the future to use for the testing and validation of new techniques as they are developed. MosesCore provides the significant resources required to create the reference translations for the new test sets for the task, as well as providing overall coordination for the tasks and workshop.

The shared task campaign ran from February to June 2013. This year there were three different tasks:

Translation Participants had to translate common test sets using their MT systems. The language pairs for this year were English to and from French, Spanish, Russian, Czech and German. The Russian reference translations were donated by Yandex.

Metrics This task was concerned with measuring the quality of machine translation output as compared to the reference.

Quality Estimation This task consisted of several subtasks all of which were concerned with measuring the quality of an MT system in the absence of a reference. The task was supported by QTLaunchPad[6].

The translation task of 2013 was the first to be run by MosesCore, and, aware of some criticism of the quality of the reference translations provided in previous years, we instituted some extra quality control procedures. We provided specific instructions to the translators about how to produce translations suitable for MT evaluation, and we recruited independent translators to quality check a sample of the translations produced.

In all we had 24 teams competing in the translation task and 8 in the metrics task, drawn from North America, Europe, the Middle East and Asia. Across the 10 language pairs in the translation task, there were between 9 and 17 submissions per language pair. Once we had assembled all submissions, we obtained human evaluations of them, both from the participating academics, and crowd-sourced via Amazon Mechanical Turk (AMT). The

[6] http://www.qt21.eu/launchpad/

crowd-sourced judgements were funded by MosesCore and Microsoft, and we used the researchers' judgements to perform quality control on the crowd-sourced judgements. The judgements were collected using the Appraise[7] system to ask judges to rank groups of 5 different MT outputs, given source and human-produced reference. This year we managed to collect significantly more judgements that last year (a 9-fold increase) enabling a more fine-grained ranking of the systems.

The workshop took place in August 2013, collocated with the ACL (Association for Computational Linguistics) conference in Sofia, and once again was the largest workshop at ACL. It included research paper presentations, poster sessions on all the shared tasks, and an invited talk from Andreas Eisele of the EU's DGT, talking about how MT is used in the European Commission.

We are already planning for next year's shared task and workshop, with a data release for the translation task planned for December 2013, and the workshop to take place in June 2014 in Baltimore, again with ACL. For 2014, we have the following changes to the shared task:

- Spanish is to be replaced by Hindi in the translation task, providing us with a low resource language.
- A new task on translation in the medical domain, supported by Khresmoi.[8]

3 The Moses Toolkit

3.1 Background

Moses[9] is an open-source toolkit for building statistical machine translation systems. It provides tools to train such systems from parallel data, and a decoder to translate sentences using models trained with the toolkit. The two main statistical MT paradigms (phrase-based and hierarchical/syntactic) are both implemented in Moses, and its comprehensive coverage of current technologies, together with its liberal LGPL license have made it popular with both academic and commercial users.

[7] https://github.com/cfedermann/Appraise
[8] http://www.khresmoi.eu/
[9] http://www.statmt.org/moses

The MosesCore project aims to retain Moses' place as (arguably) the most popular open-source SMT toolkit by continuing to incorporate new research, whilst improving stability and support. It has funded the appointment of a "Moses Coordinator" (Hieu Hoang) to oversee Moses development.

3.2 Releases

In January 2013, thanks to the support of the MosesCore project, we were able to ship a version 1.0 release of Moses[10]. The release process included a major increase in the unit and system test coverage of Moses, as well as 2 weeks of full end-to-end testing. The release was accompanied by binaries for Linux, Mac and Windows (cygwin), as well as translation models for 10 language pairs.

Project partners Capita have also created Moses installers for Windows 7 and 8, which can install native Windows versions of the decoders, as well as the pre- and post-processing tools. These are available for the Moses packages website.[11]

3.3 Current Development

After the release in January, Hieu began a major Moses refactoring effort aimed at making the decoder smaller, and more maintainable and extensible. The refactoring has included an overhaul of the feature function interface, as well as extending lattice decoding to work with syntactic and hierarchical models. Lattice decoding makes it possible to provide many possible input sentences to the decoder, for example from the output of automatic speech recognition (ASR). The refactoring also included the deletion of unused and duplicate code, and a general effort to improve code quality.

The refactoring is due to be complete and tested in time for the 2014 release of Moses. Hieu gave a talk on the refactoring at the Machine Translation Marathon, as well as spending time at the University of Stuttgart helping them update their Moses-related code to use the new interfaces.

Apart from the refactoring, contributors have continued to add new features and bug fixes to Moses, with over 4000 commits to the master branch in the last year. Some highlights from 2013 include:

[10] http://www.statmt.org/moses/?n=Moses.Releases
[11] http://www.statmt.org/moses/?n=Moses.Packages

- **Multi-model phrase table**, from Rico Sennrich, providing a mechanism to dynamically adapt translation systems.
- **Support for placeholders** enabling dates, numbers and other entities to be protected during decoding and translated separately. This was implemented by Achim Ruopp and Hieu.
- **Integration of the Operation Sequence Model** developed by Nadir Durrani in his PhD work at the University of Stuttgart. This addition to phrase-based Moses has been shown to give consistent performance increases in the University of Edinburgh's already very strong submissions to the WMT shared task.
- **Integration of the Neural Network Language Model** from Ashish Vaswani of ISI.
- **Support for Pipeline Creation Language (PCL)**, a domain specific language[12] designed for developing pipelines of reusable components, developed by Capita as part of MosesCore. Moses now includes definitions of the training pipeline in PCL.
- **Integration of Cache-based Language and Translation Models**, from Nicola Bertoldi of FBK.

4 Communications

4.1 In the Media

Moses was featured in an article in the 2013 edition of Edinburgh University's "Infinite Magazine[13]". The article mentions MosesCore and includes quotes from Jie Jiang of Capita, and a photo of the University of Edinburgh MosesCore team.

You can keep up with Moses on twitter (@MosesSMT), LinkedIn, Facebook and Google+.

[12] https://github.com/ianj-als/pcl.git
[13] http://issuu.com/edinburghresearchandinnovation

4.2 Industrial Outreach Newsletter and Website

To continue the work started in 2012, we kept the project website and the dedicated section[14] on the TAUS website up to date.

Bi-monthly newsletters have been sent to an opt-in list. These can be found via at http://www.statmt.org/mosescore/index.php?n=Main.Newsletters

All e-campaigns have been supported by social media activity across the major platforms.

In the second half of 2013, we launched MosesCore Google+ page. This page is under development and aims to reflect project and Moses related news in general.

4.3 Industrial Outreach Publications

In order to determine priorities for the next release of Moses (the first release of the open source Moses toolkit was issued in January 2013) we invited industry peers to make an input that can ensure the growing success of the Moses SMT toolkit.

We received 58 survey responses this year in comparison with 43 responses in 2012 and 50 responses in 2011. About two-thirds of the respondents have a technical role and the remaining one-third hold management and academic positions. Detailed analysis of the survey results outlined in the TAUS report Are Moses users seeking common ground?[15].

Does Moses have a future?[16] (April 2013) summarises the presentations and discussions from the TAUS MT Showcase in Singapore and raises a few questions around factors affecting the likelihood of our prediction that open source machine translation would have less market impact going forward, compared to the last few years.

In the late 2013 we will publish an article presenting a summary of the fall 2013 MT Events (the Machine Translation Marathon and the TAUS MT Showcase at Localization World Santa Clara).

[14] https://labs.taus.net/mt/about-mosescore
[15] https://www.taus.net/reports/are-moses-users-seeking-common-ground
[16] https://www.taus.net/articles/does-moses-have-a-future

4.4 Tutorials

Hieu Hoang and Philipp Koehn gave a tutorial "Open Source Statistical Machine Translation" at MT Summit in September 2013.

The TAUS Moses tutorial provides guidelines on how to get up and running with Moses and is publicly available through this link - http://www.statmt.org/mosescore/index.php?n=Main.Tutorials

This year we improved the overall responsiveness of the tutorial site and increased the number of registered users to 396 in 79 countries by November 2013.

In December 2013 we plan to follow up to gather indicators on Moses adoption resulting from tutorial usage in the last year.

Bibliography

[1] Nicola Bertoldi. Dynamic Models in Moses for Online Adaptation. *(To appear in) Prague Bulletin of Mathematical Linguistics*, 2014.

[2] Ondřej Bojar, Christian Buck, Chris Callison-Burch, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. Findings of the 2013 Workshop on Statistical Machine Translation. In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pages 1–44, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

[3] Hieu Hoang, Kenneth Heafield, Barry Haddow, Matt Post, Eva Hasler, Phil Williams, Chris Dyer, and Philipp Koehn. Updating the Feature Function Framework in the Moses Decoder. *(To appear in) Prague Bulletin of Mathematical Linguistics*, 2014.

[4] Ian Johnson. Open Machine Translation Core: An Open API for Machine Translation Systems. *Prague Bulletin of Mathematical Linguistics*, 100:91–100, 2013.

[5] Ian Johnson. Pipeline Creation Language for Machine Translation. *(To appear in) Prague Bulletin of Mathematical Linguistics*, 2014.