
MOSES CORE

Deliverable D 4.4
Industrial Outreach Web Presence and e-Communications Report

Work Package:

Date (mm/yy):

Dissemination level:

Author:

WP 4: Industry Outreach

June 2014

Public

Yulia Korobova, Achim Ruopp

Table of Contents

1. Communications	3
1.1. Market research and reporting	3
1.2. Use cases	8
1.3. Publications	9
1.4. E-Bulletins	9
2. Moses tutorial.....	10
3. MosesCore website	11
4. Social media campaigns.....	13
5. General conclusions	14
6. Action points for 2014	15

Executive summary

To continue the workflow initiated in 2012, this year we have been providing the information about the Moses development, use case stories, market research, events and related communications. We continued to use the e-Bulletins, social media, and use cases on Slideshare, articles and Moses resources web repository to ensure sufficient exposure.

Participants

TAUS is a leader of Work package 4 “Industry Outreach” and is supported by UEDIN and ALS (now Capita Translation and Interpreting).

1. Communications

1.1. Market research and reporting

1.1.1 Spring 2013 Moses Users Survey

With the version 1.0 release of the O/S Moses toolkit in January 2013, Moses became easier to install and implement with the provision of binaries, pre-built models and installation packages for a range of operating systems. “TAUS Moses Survey: towards the excellence of the O/S MT solution” (April 2013) was aimed to determine priorities for the next release of Moses. Fifty-nine participants from various industries contributed to this research.

The findings and recommendations from this survey are reflected in the report [*“Are Moses users seeing common ground?”*](#) (August 2013)¹. In this report we answered questions like “How are other users using Moses?”, “What are user priorities?”, “What is in it for me?” and “How can I contribute?”, and we looked at the larger picture and attempt to answer questions like “What are the strengths and weaknesses of the Moses project?” and “What can industry users expect now and in the future?”.

This report has over 900 unique page views and 130 downloads.

We are presenting the major findings of the report here, as these provide important insights into the adoption of Moses by users and the requirements they have for Moses. The full report is included in the appendix 2. We have conducted similar surveys and provided reports in 2011 and 2012, which enables comparison, tracking of requirements over the years and also a larger sample base of surveyed parties.

In 2012 we already noticed clear adoption patterns of the Moses open source package:

Model 1	Model 2	Model 3	Model 4
Moses out-of-the-box	Moses with open source/free add-ons	Commercialized Moses	Hybrid MT
Download and build Moses and associated components yourself	Pre-built installation to make getting started easier Additional scripts to support	Web-hosted or self-hosted UI to build MT engines and	Moses combined with other MT technologies (e.g. RBMT engines)

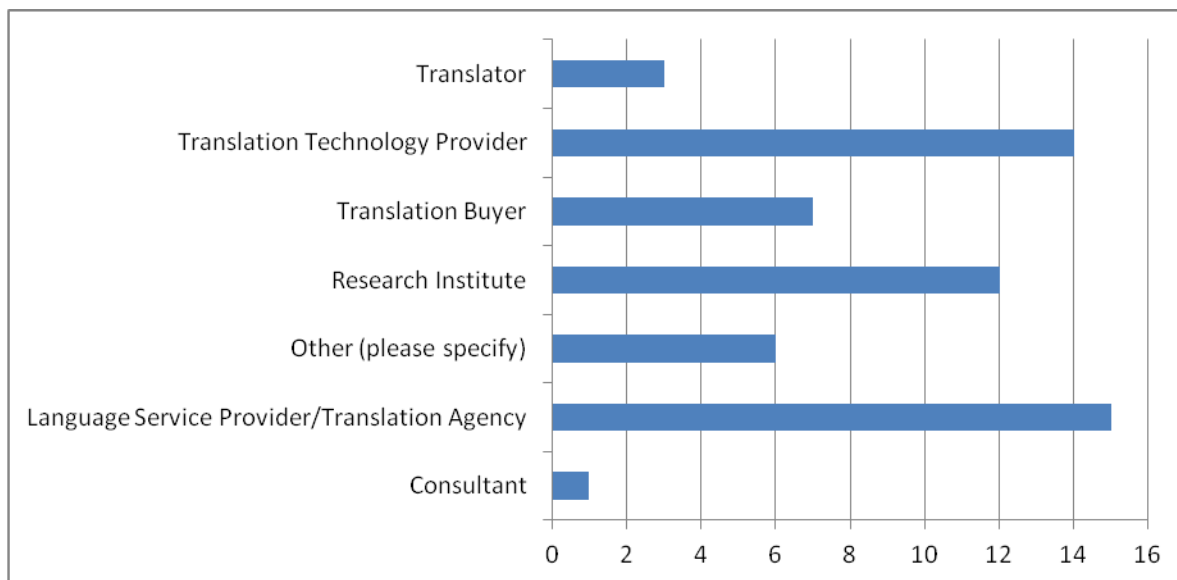
¹ <http://www.statmt.org/mosescore/index.php?n=Main.Publications>

	language industry needs	translate documents Various add-ons for data cleaning, post-editing, language support etc.	to provide high-quality MT output
Numerous translation buyers, research institutes, software developers, language service providers and a technically savvy translator or two Usage is hard to track because of the lack of download tracking on the Moses website and Github Internal Users Logrus Adobe European Commission WIPO SAP Crosslang EMC and many more (many not publicly known – see reason above)	DoMY Moses for Mere Mortals Moses for Localization Capita TI MosesCore packages (not open source)	Solutions: Tauyou Capita TI Smartmate Asia Online PangeaMT Let's MT Simple Shift Sovee DoMT KantanMT IPTranslator	Systran ProMT Safaba

Table 1: Moses adoption models

The demographics of adoption are the following according to our 2013 survey:

We received 58 survey responses this year in comparison with 43 responses in 2012 and 50 responses in 2011. Like in the years before we received survey responses from a broad cross-section of Moses users in industry and academia (see **Error! Reference source not found.**). About two thirds of the respondents have a technical role and the remaining third hold management and academic positions.



For the first time research institutes are not the largest group of respondents anymore, the largest group is now the group of language service providers/translation agencies. This shift indicates further adoption of the Moses toolkit by the language industry, which is encouraging. Language service providers as late adopters have downloaded the Moses toolkit more recently than the early adopters at translation buyers, technology providers and research institutes.

It is encouraging that more language service providers, which often do not have in depth technical or natural language processing skills, adopted Moses. However, the top requests for Moses have not changed much over the past 3 years:

2013 Rank	2012 Rank	2011 Rank	
1	1	3	Training and translation speed
2	4	1	Integrating Moses into existing workflow/system (e.g. TM integration)
3	3	2	Installing and using Moses
4	n/a	n/a	Terminology Management
5	2	4	Evaluation results (e.g. evaluating productivity)
6	5	5	Language-specific issues
7	n/a	n/a	Advanced features (e.g. tree-based translation)
8	7	7	Customer support
n/a	6	6	Easier to get the right human resources

Table 2: Ranking of requested Moses improvements

Despite improvements like the multi-threaded tokenizer and parallel training, performance remains the number one request. Requests that deal with the practical use of the core Moses engine in a language industry environment dominate the first third of the ranking. The middle third are requests that are surrounding, but not core to the Moses MT toolkit, like evaluation, terminology management and language support. At the bottom of the request ranking are advanced features and customer support – we suspect the former because advanced features are not yet on the radar of industry users and the latter because the support on the Moses support mailing list is so excellent and prompt.

Through the MosesCore project up to this point the efforts were focused on regular releases, improved code maintainability and, in continuation of the academic roots, integration of the latest research techniques into the code base. The project partner Capita TI released installation packages for Linux and Windows platforms, with still limited support and usability for non-technical users on Windows. Our survey results show that these efforts have not changed the prioritization of requested features among the user base.

On one hand this reflects the changing user basis of Moses with even more increased demands for the open source solution. The still largely academic developer base has not kept up with the industry-specific demands, which is understandable, given the different academic priorities. Industry developers have not stepped up to address these common issues in the open source code base.

Instead, an increasing number of MT system integrators use Moses as a basis supplemented with proprietary components to address the industry-specific gaps in the open source solution (Model 3 and 4). Commercial solutions are most often deployed as cloud-based SaaS (Software-as-a-Service) solutions that do not require the sharing of improvements with the larger community. Besides, the LGPL license under which Moses is licensed, even on distribution, does not require the open sourcing of supplementary components.

Many users of the cloud-based solutions do not even know that they are using Moses in their MT solutions. Our surveys cannot capture these users. Often the Moses-community does not know either, as disclosure of use of the open source code base is not required.

This trend to Moses-based cloud MT solutions has solidified over the last three years along with larger industry trends to cloud-based solutions. These cloud-solutions are now, in TAUS' assessment, de-facto the largest and most rapidly growing deployment of Moses in the foreseeable future. The direct user base of open source Moses, in our assessment, will remain small, consisting of academia, MT system integrators, large internal users and a small enthusiast user base of individuals and SMB. Such a development of the commercial vs. academic open source MT solutions, paralleling a similar development in the speech recognition market in earlier years, was predicted by Daniel Marcu in a presentation at the TAUS User Conference 2011 – an excerpt of his slides is included in the appendix. Daniel Marcu is Research Project Leader at the Information Sciences Institute of the University of Southern California and Chief Science Officer at SDL PLC, the largest language service provider. Last year Daniel Marcu was invited as independent reviewer of the META-NET project.

1.1.2 Fall 2013 Moses Audience Survey

In Q4 2013 we carried out another Moses Users survey to understand the current level of Moses adoption and identify the influence of the MosesCore project on the MT community. Seventy-one participant from various industries participated in this questionnaire (29 January 2014).

Initial results show that the demographics of survey respondents and event participants are heavily skewed towards academia and larger industry participants (buyers, translation providers and technology providers). We also notice gradual growth of language services providers.

Participants in Moses Survey, 2012-2013




Chart 1 Moses Survey participation, per type of activity, 2012-2013

It is still a challenge to involve the larger translation community in Moses research and implementation due to the still required technical knowledge and a strong bias against MT in some quarters of the translator community. This is for example a quote from the International Association of Professional Translators and Interpreters: “Free or commercially available machine translation (MT) software cannot be expected to render with accuracy the meaning of a foreign-language text, and thus MT systems cannot guarantee that their output is suitable for any purpose other than obtaining the gist of a text. The result can be a very poor translation, leading to a negative image for the client and potentially loss of business. [...] machine translation does not offer the consistently high quality and accountability that a professional translator does, and information obtained by means of machine translation cannot be considered reliable. Therefore, users of publicly and commercially available machine-translation software should be aware of the dangers involved, such as compromised quality and confidentiality, and making any decisions or taking any actions based on machine-translated information shall be at the users' own risk.”²

Chart 2 shows that the number of users who have downloaded and keep using the Moses decoder in the last two years is growing.

² <https://www.iapti.org/speaks-out/art8-iaptis-recommendation-on-machine-translation.html>




Chart 2 Moses Survey: Use of Moses decoder, 2012-2013

According to the results presented in Chart 3, a level of adaptation of Moses gradually shifted from piloting and implementation that dominated in 2012- beginning 2013 to production. On this stage Moses technical support and businesswise recommendations from MosesCore are essential and should be addressed during events and news resources.

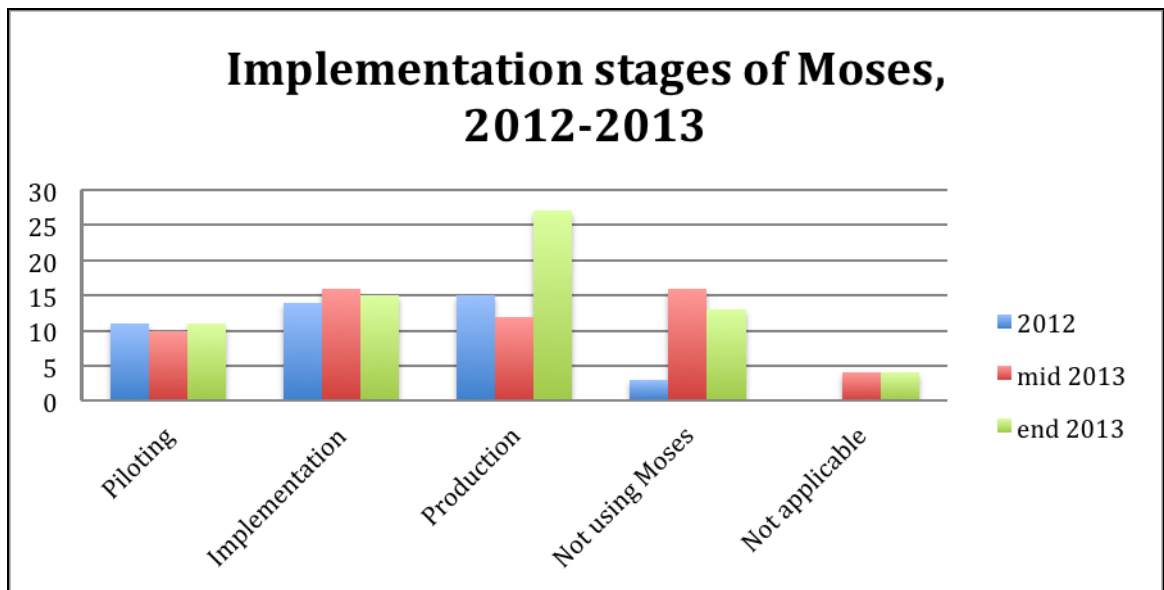


Chart 3 Moses Survey: Implementation stages, 2012-2013

Other survey outcomes will be analysed and shared with the Moses community at the end of Q1 2014.

1.2. Use cases

In 2013 we organized three complimentary MT Showcases, two of which were held as a pre-conference day of Localization World (see report on D4.5: Report on second year’s industry outreach events). All the presentations are publicly available via the MosesCore project website.

- MT Showcases in Singapore, 10 April 2013
- MT Showcases in London, UK, 12 June 2013
- MT Showcases in Santa Clara, USA, 9 October 2013

The twenty use case presentations are accessible on Slideshare, MosesCore project website and the TAUS MT repository. These presentations had been viewed a total of 7,117 times, making an average of 356 views per talk (16 January 2013).

Use cases from 2013 together with the presentations from the previous year can be found via [this link](#).³

1.3. Publications

Two articles published in 2013 cover the outcomes of the MT Showcases.

[“Does Moses Have a Future?”](#) by Rahzeb Choudhury

During the MT Showcase in Singapore TAUS presented an overview of the Moses timeline and highlighted some of the myths around the toolkit. This article discusses a few questions raised about the factors affecting the currently active role of proprietary MT suppliers in comparison to the previously quickly developing open source MT offerings.

[“Moses MT - Collaboration beats Fragmentation”](#) by Achim Ruopp

This article covers the outcomes of MT Marathon in Prague and the MT Showcase in Santa Clara discussing the ideas on cooperation for better custodianship of Moses for commercial users with Moses users from the industry.

These publications had a total of 923 page visits (582 unique page views) with an average time on page of two minutes. These articles had a high response rate from the industry peers. All the comments are published on the TAUS website.

1.4. E-Bulletins



Image 1 MosesCore E-Bulletin

This year TAUS sent 5 MosesCore e-bulletins⁴ to an opt-in distribution list. This list was created at the beginning of the project and is regularly updated with contact details of the MT Showcases participants, users of Moses toolkit and tutorial, as well as readers of Moses publications and e-bulletin subscribers.

January/February: [Moses SMT releases v.1, Help needed, MT Showcases 2013](#)

This bulletin was dedicated to the Moses SMT release.

E-campaign: 650 recipients

Looking in to the future with Moses - new events and publications.

E-campaign: 910 recipients

May/June: [Week Long Hackathon, Moses in The Mix, Showcase in Santa Clara](#)

³ <http://www.statmt.org/mosescore/index.php?n=Main.Videos>

⁴ <http://www.statmt.org/mosescore/index.php?n=Main.Newsletters>

This bulletin focused on results of the MT Showcase London and the preparations for MT Marathon and the MT Showcase in Santa Clara.

E-campaign: 915 recipients

July/August: [Free TAUS Report, Complimentary Events in Prague and Santa Clara](#)

The key message of this bulletin is the introduction of the recently published report on “Moses users seeking common ground”

E-campaign: 544 recipients

November/December: [Moses: Collaboration beats Fragmentation](#)

This newsletter is about the Moses User survey and the results reached in 2013.

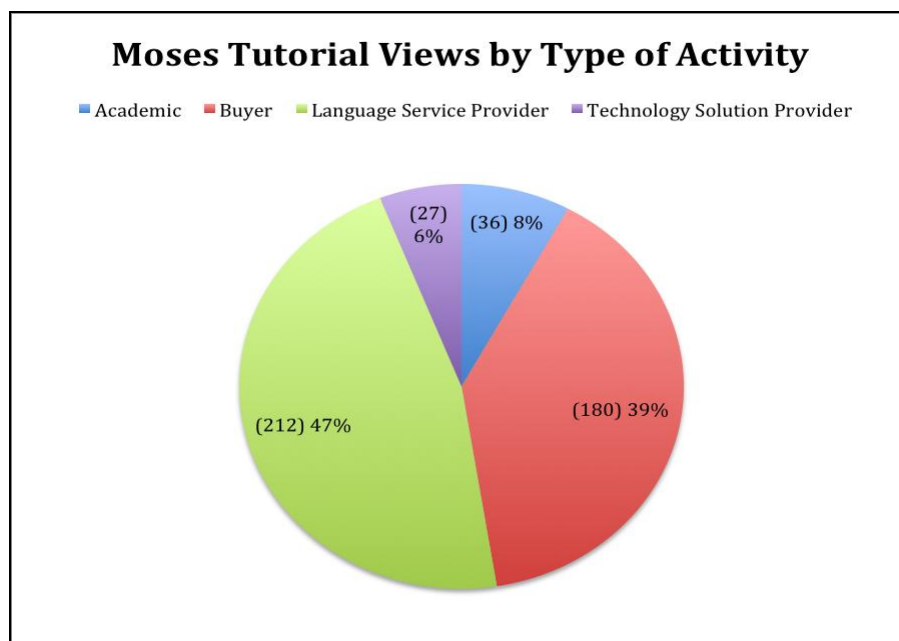
E-campaign: 803 recipients

An average open rate for the Moses e-Bulletins in 2013 was 35% of the total of recipients, which is higher than a common rate in technology mailings but a bit lower than in the previous year of the MosesCore project. The lower number of recipients and a drop in the open rate has to do with the change of our CRM (smaller but more up-to date database, new profile) in summer 2013. The lessons learned from these transitions would be taken into consideration for our future mailings.

On average the click through rates on the Moses mailings ranged between 25-35%, which is similar to 2012 data.⁵

2. Moses tutorial

In 2013 we have managed to grow the list of Moses tutorial users up to 1,070 (in combination with other TAUS services), 453 representatives from 98 countries are registered for the tutorial exclusively. Similar to the year before a relatively small number (13) of the 453 registered tutorial users attended MT Showcase events in 2013. Low effort is still one of the benefits of this free online resource in comparison to the real life event.



⁵ As soon as it is impossible to be 100% sure in the open rate, we do not take it as a hard and fast number. It is used here as general guide, and as a way of measuring the trends on our email campaigns.

Chart 2 Moses tutorial views by type of activity




Chart 3 Familiarity with Moses tutorial among the Moses users, 2013

Results of the survey show that impact of MT tutorial is less than we would like to see. Only 41% of Moses users in the industry know about tutorial. As most of the users in the Moses community are rather advanced in employing Moses technology and do not have a need in tutorial anymore, we can conclude that the reach of 41% of Moses users is a positive result of our communication campaigns. Although, the increase of this number is not on our priority action list for 2014, we will definitely continue to work on usability of this resource, which is essential for the new users of Moses.

3. MosesCore website

This year we have continued to use the MosesCore website to communicate about Moses developments.

Web presence: 1,309 visits (531 unique visits)

More than 2,266 page views.



Image 2 MosesCore home page

“MosesCore Project Website: Audience Overview” chat shows that majority of the visitors this year are new. By means of cross linking with TAUS Labs and using e-campaigns we have managed to maintain awareness about the project, as well as to reach and attract different audience.




Chart 4 MosesCore website: audience overview

The chart below highlights the global user base of MosesCore resources. The figures for the Netherlands exclude traffic arising from TAUS, who are based in Amsterdam.




Chart 5 MosesCore website: demographics

Next image shows an overview of the most popular pages on the MosesCore website. Results show that there is a demand for information about the events. We will take this information as one of the action points for the next year and will make an effort in crafting Moses events.

Page Title	Pageviews	% Pageviews
1. MosesCore Main/Home Page	1,288	38.27%
2. MosesCore Main/Events	425	12.63%
3. MosesCore Main/Moses Releases	276	8.20%
4. MosesCore Main/Publications	272	8.08%
5. MosesCore Main/Partners	190	5.64%
6. MosesCore Main/Tutorials	182	5.41%
7. MosesCore Main/Videos	158	4.69%
8. MosesCore Main/News	152	4.52%
9. MosesCore Main/Deliverables	95	2.82%
10. MosesCore Main/Other Projects	76	2.26%

Image 3 MosesCore website: page analytics, 2013

In 2013 we also continued to use and extend the MosesCore resources on the TAUS Labs web site. Image 4 shows that Moses tutorial was the fourth most visited section on TAUS Labs in 2013. 915 web users visited the MosesCore section.

Page	Pageviews	% Pageviews
1. labs.taus.net/	4,035	15.34%
4. labs.taus.net/mt/moses-tutorial-guest	1,305	4.96%
8. labs.taus.net/mt/mosescore	915	3.48%

Image 4 MosesCore sections on the TAUS Labs: page analytics, 2013

4. Social media campaigns

In 2013 we kept using various Social media channels to promote Moses content. Every publication and event had a cross-channel promotion and referencing to #MosesCore on Twitter.

In addition to the existing channels this year we have initiated the Moses Google + Group (<https://plus.google.com/108294242360153392968/posts>) This group is still under development and did not get much influence on the total communication flow though.

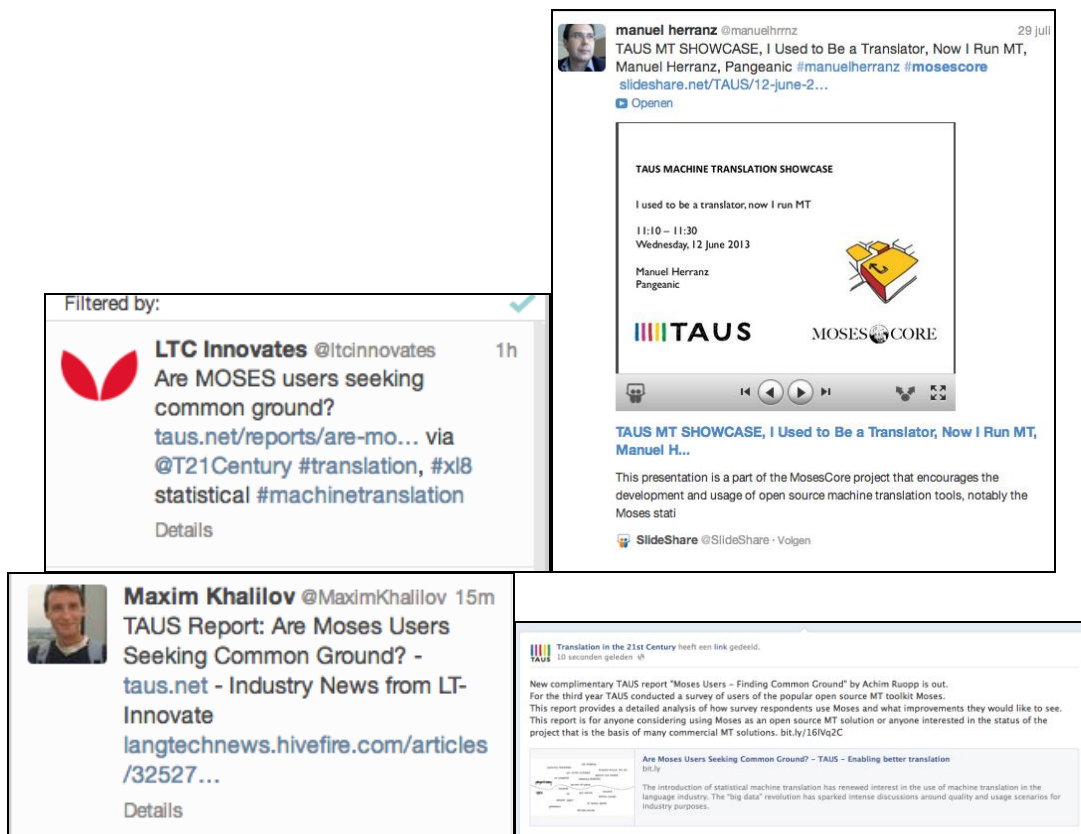


Image 5 MosesCore on Social media, 2013

5. General conclusions

In 2013 we have succeeded in communication activities targeted at language service providers. The use cases published on the project website, as well as on TAUS Labs attracted a good number of web visitors who were exploring Moses and the related knowledgebase.

At the early stage of the MosesCore project we have mainly focused on Model 3 and 4 of Moses adaptation (See also D4.1 MosesCore Industrial Outreach Plan, 2012-2015). In 2014 we will pay special attention to the internal users engaged in Model 1 (Moses out-of-the-box) and Model 2 (Moses with open source/free add-ons), which are often individuals or smaller companies. Definitions of the Models can be found in section 1.1.1 above.

New release of Moses v2.1 introduced a number of useful features particularly interesting for the industrial users:

1. Transliteration Phrase-Table
2. Placeholder
3. Ondisk phrase-table for phrase-based model
4. Updated Windows GUI

Moses is now available in the form of Amazon EC2 images.

Among others, communication campaigns for 2014 will be focused on emphasizing these benefits of Moses direct adopters (Model 1 and Model 2).

6. Action points for 2014

In order to address some of the registration concerns of the reviewers we are planning to make more resources available on the MosesCore website and popular 3rd party publication sites:

- We will make all MosesCore related TAUS reports downloadable directly from MosesCore page without registration
- We will make the Machine Translation and Moses tutorial videos visible/searchable on YouTube
- We will continue to make the slides from the MosesCore events that TAUS organizes (for a detailed list see D4.5) available on Slideshare
 - We will add lead generation to the slides to encourage leaving contact information voluntarily in order to be able to follow up with interested parties for in-depth feedback requested by the consortium
 - We will add the slides from Tilde for the October 2013 MT Showcase

To increase discoverability of resources we add and maintain links on the MosesCore site to related projects and the numerous Moses-based industry solutions

(<http://statmt.org/mosescore/index.php?n=Main.OtherProjects>).

To improve the tracking of the impact of the different outreach activities we are planning to improve the user tracking across different online properties by improving user tracking on statmt.org and track users across the statmt.org and taus.net domains using Google Analytics. This also includes adding robots.txt rules to statmt.org to improve the analyzability of web server access logs.

Supplementary TAUS will continue actively involve social media for the MosesCore project. Optimization of social media channels stays a challenging task. As Facebook and LinkedIn groups had a private character and did not have much weight in our communication campaigns, we have decided focus our effort of Twitter and Google+ only in 2014. The private character of the LinkedIn group is due to a technical limitation we cannot fix. We intend to encourage users to join the Google+ group and close the LinkedIn group.

These are the communication KPIs set for 2014:

- Attract existing Moses community to the MosesCore website. At the moment the average number of visits is 159 a month. In 2014 we plan to increase it to 200 by means of cross-references with Social media channels and TAUS Labs, as well as by improving the CEO.
- Increase public interaction and improve community management on Twitter and Google+. Through mutual efforts of the consortium members in generating the relevant and up-to-date content about MosesCore and Open MT, we strive to grow our presence on Twitter (@MosesSMT) to 500, and 100 followers on Google+ till the end of 2014. TO support these plan TAUS will approach its members with a special request to follow and engage in Moses activities on social media.

Reaching these KPIs requires the effort of all MosesCore consortium members, not just TAUS. All partners should encourage their contacts to follow the Twitter account and the Google+ group. The partners should also post news regularly to these social media accounts. In order to ensure regular posting TAUS will monitor the moses-support mailing list for important news and encourage partners to post to social media if they have not done so yet. We would also like to get permission to post on the Twitter account. We are hoping for the support of the partners to meet the outreach goals.

An editorial calendar will be put in place in June 2014. Moses related news, publications and presentations provided by consortium members, as well as news collected by means of Google Alerts, feedly.com and mt-list will be used to strengthen the MosesCore web presence and position consortium as

a thought-leader maintaining and supporting the knowledgebase in relation to Moses and open MT. Following table provides the high-level weekly plan for our communication effort on Moses.

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Moses News and events	Check Moses tutorial	Weekly selection of MT-list update	Look back at the MT showcase results: focus on presentations	Moses News	Moses articles and reports

Table 2: Initial draft of the editorial calendar Q3-4, 2014

With the market developments we observed in the Moses User Survey reports in past years and in order to capture more qualitative feedback requested by the MosesCore project leader, Barry Haddow, we plan to interview internal users of Moses (mainly at large companies like Oracle and large organizations like the European Commission) and Moses integrators (for a list see section 1.1.1). We also plan to interview more individuals or SMB users of Moses, although they are hard to come by.

The interview responses and additional original market research by TAUS will be compiled into a comprehensive Moses Market Report which will take the place of the smaller scoped deliverable D4.8.

We currently have 1070 registered users for the Machine Translation and Moses Tutorial, which are 100 more than the current number of subscribers to the moses-support mailing list (970). Given that the awareness of the tutorial among existing Moses users is low (41%), most likely due to the introductory nature of the tutorial and the expertise of many users, we conclude that we already introduced many novices to Moses and MT in general through the tutorial. We will add communication to the Moses user base and the larger industry community advertising the tutorial. We will later survey the existing tutorial users and newly recruited users on the tutorial and feedback on Moses in general and feedback for novice users in particular. As mentioned we also plan to make the tutorial videos available publicly on YouTube, although it will be harder to capture feedback from YouTube users as this does not require registration and users can voluntarily leave comments.

Our communication activities will keep their informative character. In 2014 we will carry on e-bulletins focused on the Moses use case stories, reports and research results.

Appendix 1 Daniel Marcu Slides from TAUS User Conference 2011

Letting 1000 MT Systems Bloom

Daniel Marcu, CTO
SDL Language Technologies



A bit of history

Speech Recognition

- Early 1990s
 - A few startups
 - Open source toolkits – HTK
- Late 1990s – Early 2000s
 - Dragon Naturally Speaking
 - Open source toolkits – HTK
 - Lots of startups using HTK & variants
 - Lots of experimentation
- 2000s – Today
 - Product, IP, and talent consolidation
 - Open Source toolkits for R&D

A bit of history

Speech Recognition

- Early 1990s
 - A few startups
 - Open source toolkits – HTK
- Late 1990s – Early 2000s
 - Dragon Naturally Speaking
 - Open source toolkits – HTK
 - Lots of startups using HTK & variants
 - Lots of experimentation
- 2000s – Today
 - Product, IP, and talent consolidation:
 - Open Source toolkits for R&D

Machine Translation

- Early 2000s
 - A few startups
- 2005 – Today
 - A few products
 - Open source toolkits – Moses
 - Lots of startups using Moses & variants
 - Lots of experimentation
- Today – Future
 - ?

A bit of history

Speech Recognition

- Early 1990s
 - A few startups
 - Open source toolkits – HTK
- Late 1990s – Early 2000s
 - Dragon Naturally Speaking
 - Open source toolkits – HTK
 - Lots of startups using HTK & variants
 - Lots of experimentation
- 2000s – Today
 - Product, IP, and talent consolidation
 - Open Source toolkits

- Product characteristics
 - Works out of the box --- to a certain extent
 - Product can be adapted to user needs



State of affairs today

Speech Recognition

- Product characteristics
 - Works out of the box to a certain extent
 - Product can be adapted to user needs



Machine Translation

- Product characteristics
 - Works out of the box --- to a certain extent
 - Product can be adapted to user/business needs by highly qualified people

State of affairs today/tomorrow

Speech Recognition

- Product characteristics
 - Works out of the box to a certain extent
 - Product can be adapted to user needs



Machine Translation

- Product characteristics
 - Works out of the box to a certain extent
 - Product can be adapted to user/business needs by **highly qualified people**

Appendix 2 Report “Moses Users: Finding Common Ground”



Enabling better translation

TAUS RESEARCH REPORT

Moses Users - Finding Common Ground



July 2013

MOSES  CORE

TAUS - Enabling better translation

COPYRIGHT © TAUS 2013

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system of any nature, or transmitted or made available in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of TAUS. TAUS will pursue copyright infringements.

In spite of careful preparation and editing, this publication may contain errors and imperfections. Authors, editors, and TAUS do not accept responsibility for the consequences that may result thereof.

Published by TAUS BV, De Rijp, The Netherlands.
For further information, please email info@taus.net

Relevance

For a third year TAUS conducted a survey of users of the popular open source machine translation toolkit Moses.

This report provides a detailed analysis of how survey respondents use Moses and what improvements they would like to see. With the data gathered over the past three years we look at the larger picture and analyze where the strengths and weaknesses of the project lie and what users can expect in the future.

This report is for anyone considering using Moses as an open source MT solution or anyone interested in the status of the project that is the basis of many commercial MT solutions.

Author

Achim Ruopp

Reviewers

Rahzeb Choudhury and Maxim Khalilov



This is a MosesCore project report supported by the European Commission Grant Number 288487 under the 7th Framework Programme.

For more information about the MosesCore initiative, please check <http://www.statmt.org/mosescore/index.php?n=Main.HomePage>

Contents

1. Introduction	6
2. Demographic Composition	7
3. Survey Results	8
3.1. Comparison of 2011 - 2013 Rankings	8
3.2. Detailed Analysis of Results	9
3.2.1. Training and Translation Speed	9
3.2.2. Integrating Moses into Existing Workflows/Systems	10
3.2.3. Installing and Using Moses	10
3.2.4. Terminology Management	11
3.2.5. Evaluation	11
3.2.6. Language-Specific Issues	12
3.2.7. Advanced Features for Moses Commercialization	13
3.2.8. Customer Support	13
3.2.9. Moses Strengths and Weaknesses versus other MT Solutions	14
3.3. Implementation of Previous Recommendations	15
3.3.1. 2011: Four of Six Recommendations Implemented	15
3.3.2. 2012: Five of Eight Recommendations Implemented	16
4. Finding Common Ground	18

1. Introduction

We are now a few years into the adoption of statistical machine translation in the language industry. Judging from discussions on social networks, if we assume a Gartner hype cycle for the adoption of statistical machine translation, we have now moved from the “*Peak of Inflated Expectations*” to the “*Trough of Disillusionment*” (see Figure 1). This is also reflected in the increased focus on industry quality initiatives and admittedly also in the TAUS article “*Does Moses have a Future?*”.

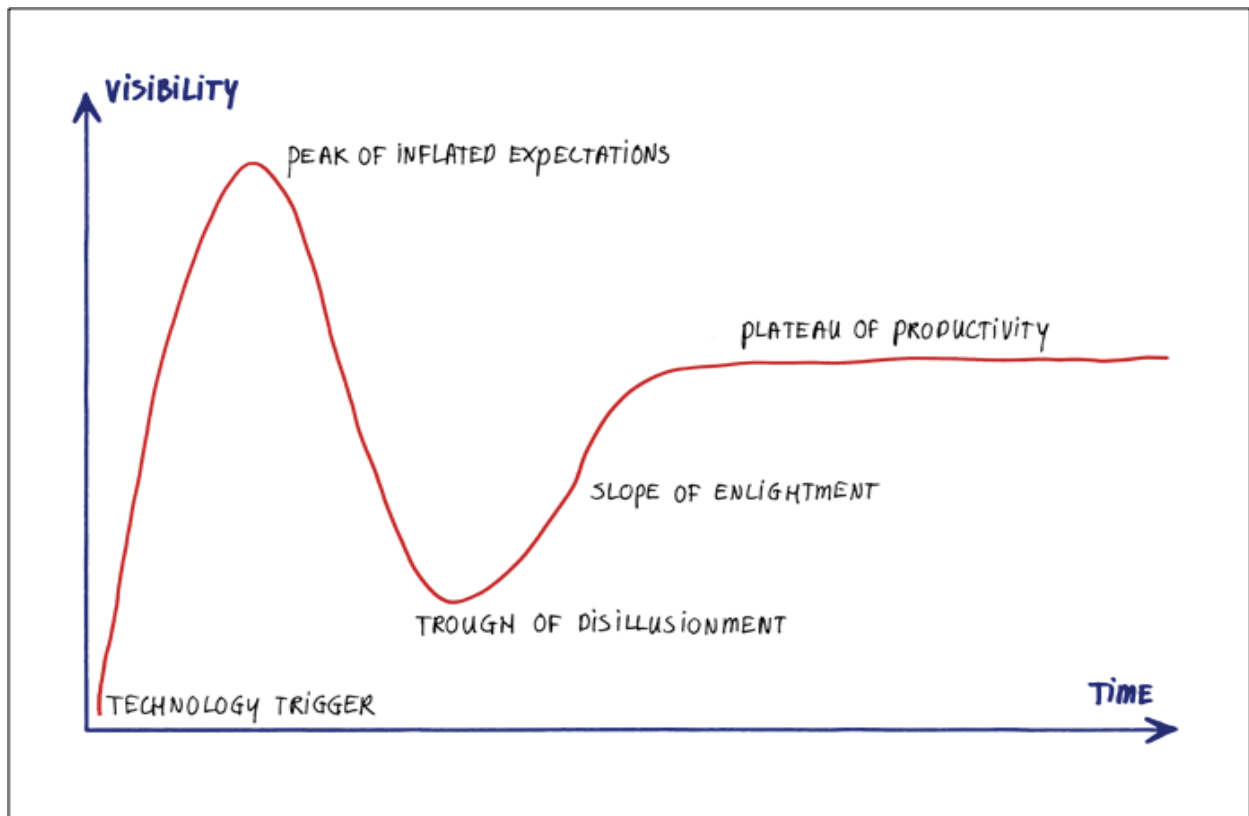


Figure 1: Gartner Hype Cycle Diagram ©Jeremy Kemp CC-BY-SA 3.0

As in 2011 and 2012 we conducted a survey of Moses statistical machine translation toolkit users. The results cut through the hype cycle and show that the adoption of this toolkit in academia and industry has continued uninterrupted and has moved into a more mature stage.

We also look back at the recommendations we made in the 2011 and 2012 reports based on the improvements most requested by users. A lot of progress has already been made on some recommendations, less on others. We analyze where the strengths of the publicly funded Moses project lie and try to adjust our new recommendations to focus resources on areas most beneficial to academia and industry.

2. Demographic Composition

We received 58 survey responses this year in comparison with 43 responses in 2012 and 50 responses in 2011. Like in the years before we received survey responses from a broad cross-section of Moses users in industry and academia (see Figure 2). About two-thirds of the respondents have a technical role and the remaining one-third hold management and academic positions.

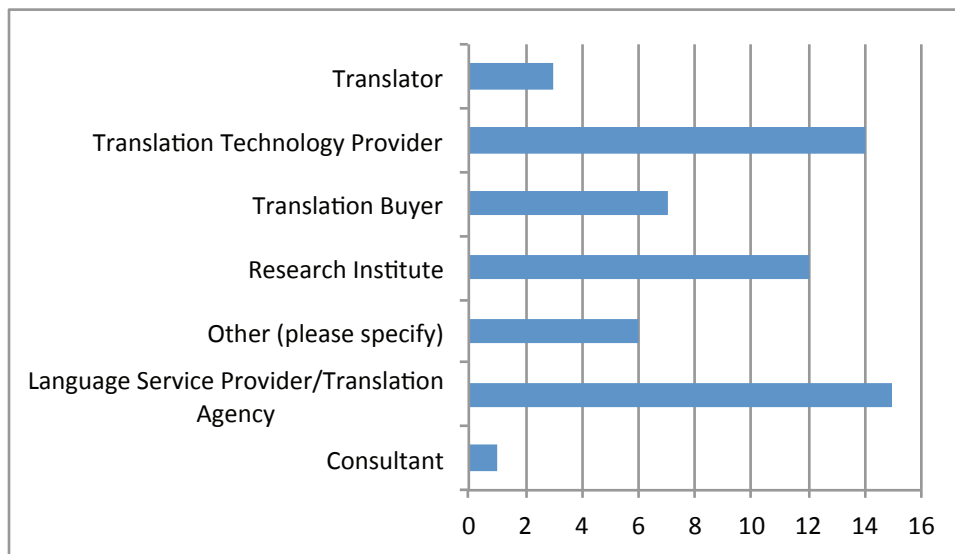


Figure 2: Survey Respondents

For the first time research institutes are not the largest group of respondents, the largest group is now the group of language service providers/translation agencies. This shift indicates further adoption of the Moses toolkit by the language industry, which is encouraging. Language service providers as late adopters have downloaded the Moses toolkit more recently than the early adopters at translation buyers, technology providers and research institutes.

As can be seen in the increasing number of pre-packaged solutions based on Moses in the market, the interest in Moses by translation technology providers is still very strong and it is unlikely that Moses as the core technology in these solutions will be replaced any time soon, if ever. Examples include Asia Online, PangeaMT, Tauyou, Safaba Translation Solutions, LingoMT and others.

In addition to having a broad reach across different types of institutions, Moses is also used in many different countries. While the bulk of respondents still come from Europe, Moses also is used extensively in the Middle East, Asia and the Americas (as anybody monitoring the Moses support mailing list can attest to).

3. Survey Results

3.1. Comparison of 2011 - 2013 Rankings

Before diving into the details of the individual survey questions we first look at Moses area for improvement, as ranked by importance by the participants over the past three years:

Rankings			
2013	2012	2011	
1	1	3	Training and translation speed
2	4	1	Integrating Moses into existing workflow/system (e.g. TM integration)
3	3	2	Installing and using Moses
4	n/a	n/a	Terminology Management
5	2	4	Evaluation results (e.g. Evaluating productivity)
6	5	5	Language-specific issues
7	n/a	n/a	Advanced features (e.g. tree-based translation)
8	7	7	Customer support
n/a	6	6	Easier to get the right human resources

Table 1: Ranking of Requested Moses Improvements

Despite improvements like the multi-threaded tokenizer and parallel training, performance remains the number one request. Requests that deal with the practical use of the core Moses engine in a language industry environment dominate the first third of the ranking. The middle third are requests that are surrounding, but not core to the Moses MT toolkit, like evaluation, terminology management and language support. At the bottom of the request ranking are advanced features and customer support – we suspect the former because advanced features are not yet on the radar of industry users and the latter because the support on the Moses support mailing list is so excellent and prompt.

3.2. Detailed Analysis of Results

3.2.1. Training and Translation Speed

3.2.1.1. Are you satisfied with the speed of Moses training and translation processes? What are your comments and suggestions?

Most users, being fairly experienced with Moses and statistical MT in general by now, are aware that training SMT systems and translating with them requires a considerable amount of computing resources. However, management and customer requests for quick turn-around translations and re-training of MT engines with new data, leads users to push the performance envelope further and some of them implement custom solutions to increase speed.

The core Moses team recently made several improvements to speed up MT system training with the introduction of a multi-threaded tokenizer and parallel training. It is evident from survey comments - *“The training workflow is still based on some outdated assumptions that preclude it from optimally utilizing available computing resources”, “multi-core parallelization should be available to the entire tool chain”* – is that speed improvements need to be integrated into the training tool chain and/or users need to be made more aware of improvements.

There are also calls to improve the decoder speed in multi-threaded environments even more (“we notice that running Moses decoder with 10 threads consumes 10 cores, but provides only very limited speed-up”). Moses is already very fast in comparison to other academic MT decoders. See the presentation *“Moses Past, Present and Future”* by Hieu Hoang at the MosesCore/TAUS Machine Translation Showcase in London. The question is whether such performance optimization work should be performed by academia or if the development should be funded by industry, which has a greater need for speedy online translations.

Recommendation: Integrate recent training speed improvements into the training tool chain and document recommendations how to best use the training speed improvements. Further performance optimization for multi-threaded decoding.

3.2.2. Integrating Moses into Existing Workflows/Systems

3.2.2.1. Which improvements related to integration of Moses into the existing workflow would you like to see in the future releases of Moses?

Like in previous years the desire for a RESTful HTTP API and translation of different file formats with related tag handling are mentioned in the responses (*“XLIFF handling, better tag handling/placement”, “RESTful API”*). With the broadening of the user base there are also some interesting new requests for Moses integration: *“Easier integration to ASR systems”, “perfect integration with TM and QA tools”* and *“Integration in the [redacted] TMS is still shoddy, although this is also the responsibility of the TMS development team itself”*.

TAUS - Enabling better translation

There are also some calls for improved portability of trained Moses models: *“Moses translation models should be portable to another location/system with Moses”* and *“some tokenizer that would be compatible with tokenization required for parsers so the users would not have to convert from one tokenization to another all the time”*.

It is clear that with broader adoption, which is welcome for the continued vitality of the Moses project, workflow integration requests will become even more varied than before. It is impossible for the Moses core team to address each one of them individually.

Recommendation: Our recommendation is that the Moses core team completes comprehensive, stable and well documented APIs to the decoder and data produced by it. This information can be used by developers to implement varied integration scenarios.

The Okapi/M4Loc project already provides some file format support including tagging and awareness about this should be raised. The language industry might also fund the implementation of web APIs modeled after the popular Google Translate/Microsoft Translator/TAUS API specifications.

3.2.3. Installing and Using Moses

3.2.3.1. What are your suggestions or comments on the installation procedure of the Moses Decoder and/or related components?

3.2.3.2. Is lack of compatibility with Microsoft Windows operating systems a barrier to adoption?

Installation experiences reported by users range from *“Nothing specific... it was quite easy for me to have it setup.”* and *“No problems”* to *“installation procedure?”* *“Is there one listed?”* and *“very complex to understand and to implement”*. These extreme differences can be attributed to very different experience levels with software installation on Unix and easier installation on some platforms (e.g. current main-stream Linux distributions with few pre-installed packages). Survey participants commented that the release of Moses v1.0 made the installation much easier. However, installation of some required 3rd party components still creates problems. Some of this should be solved by the installable packages to be released by the MosesCore partner Capita TI.

There isn't unanimous agreement that the Moses team should provide installable packages – one provided commenter writes *“Let commercial vendors create packages. [...] Simply benchmarking versions per year is enough.”*

This year we asked a new question whether the lack of compatibility with Microsoft Windows is a barrier to adoption. Quantitatively the respondents were evenly split on this question while in the comments nobody mentioned this as a real barrier to adoption. Though clearly, non-users maybe a better response group for this question.

Recommendation: Occasional stable releases of Moses as installable packages across different platforms are greatly desired by the Moses user community. Following the example of similar open source projects, we think the Moses team should continue the releases it started. The main issue is still dependent components like language modeling toolkits and alignment tools. Beyond coordination it does not make sense to ask the Moses team to release components actively developed and released by 3rd parties. However, it may be a good idea for the team to take on the maintenance and release of components abandoned by their original developers.

3.2.4. Terminology Management

3.2.4.1. Would you like to see an integrated model for terminology management in the future releases of Moses module? What would the requirements be?

Since mid last year TAUS is participating in the FP7 TaaS project who's goals are to address the need for instant access to the most up-to-date terms, user participation in the acquisition and sharing of multilingual terminological data, and efficient solutions for terminology resources reuse (for more information see <http://www.taas-project.eu/>).

One of the use cases of the project is to use the terminological resources from the TaaS project in SMT systems like Moses. We therefore added a question this year asking survey participants whether they would like to see a terminology management module added to Moses and what the requirements for such a terminology module would be.

All respondents expressed the need for terminology management (*"Moses should be able to identify terminologies and use the recommended translation", "run-time terminology injection feature w/o retraining"*) and many respondents are aware of the XML input option for Moses, allowing basic terminology injection. Respondents would like to see this feature better explained and documented as there are concerns about the feature affecting translation output quality, specifically fluency. Some respondents asked for more advanced features like terminology (named entity) recognition in the source and terminology support beyond word lists, e.g. using semantical information.

Recommendation: Better documentation of the XML input feature, ensuring that the XML input feature minimally impacts translation quality of the overall sentence and possibly the option to provide more linguistic information to the decoder, so that it can make the right terminology choice. As terminology recognition, a.k.a. named entity recognition is its own separate research field in computational linguistics with many different toolkits available, the integration of named entity recognition tools is not recommended. However, it should be ensured that input marked up with named entities can be easily used as input for the Moses decoder.

3.2.5. Evaluation

3.2.5.1. Do you have any suggestions or comments on the translation quality evaluation functionality built into Moses?

3.2.5.2. Would you like to see a productivity testing system integrated into one of the next releases of Moses?

Asked about translation quality evaluation functionality built into Moses, many survey respondents suggested an expansion of the metrics that can be used to tune Moses MT systems. Participants list a whole range of metrics, including metrics that model productivity in MT+ post-editing usage scenarios.

Some participants want to see additional evaluation functionality and productivity testing systems integrated into Moses (*"It would be interesting to have some integrated tools for translation visualization and manual evaluation"*, *"Should mimic a production post-editing environment and evaluate on the segment level."*), while many don't see the need for such tight integration (*"Testing productivity is different from MT and highly dependent on the user's situation. I am afraid that trying to develop such tools would distract Moses developers too much from the core issues of improving MT quality."*, *"There are other ways to do this, and it seems unnecessary."*).

Recommendation: Integrate tuning metrics into Moses that allow optimizing systems for the MT+post-editing usage scenario.

3.2.6. Language-Specific Issues

3.2.6.1. Would you like to see any language-specific features included into the next releases of Moses? Please explain which and why.

Due to the history of the Moses project the built-in language support in Moses is focused on a relatively small set of European languages. For languages outside of this set Moses users need to find and install tools to process them.

Our survey participants would like to see tools for more languages included: *"It would be nice to have tokenizers or word segmentors readily available"*, *"If I can dream big: integration of a segmentor for unsegmented languages (i.e. Chinese, other Asian languages, etc.)"*.

Some users develop their own tools, but run into Unicode bugs as Moses hasn't been extensively used or tested with some scripts: *"[.]Sinhala and Tamil [...] when we tokenize it removes the zero-width joiner also. Then our characters are breaking and then those characters are consider as another word."*

Recommendation: Test and improve Unicode support in the language-independent core. Provide documentation on which language tools have been tested and are well established, along with recommendations on how to use them. Users should be more active in reporting Unicode issues to the mailing list and in providing language-specific non-breaking prefix files for tokenization of more languages.

3.2.7. Advanced Features for Moses Commercialization

3.2.7.1. Which are the advanced technologies that in your opinion are most urgent for further Moses commercialization?

When we asked which advanced technologies participants consider to be most urgent for further Moses commercialization, we received a broad cross-section of requests already discussed in the sections above. This illustrates the broadening commercial user base and the numerous scenarios in which Moses is used. Industrial users have very practical requests when using this still maturing open source project. We should not forget, however, that researchers across the world also develop cutting-edge technologies that can lead to significant improvements in MT quality and productivity for the industry. Some of them are already included in Moses and many more are about to be integrated into the code base – this is one of the major strengths of Moses.

Recommendation: The Moses core team should not only respond to requests from industry, but also start a conversation explaining how newly developed methods and technologies can help the industry to address critical MT issues. For example by demonstrating quality improvements for certain languages and giving recommendations how to deal with morphologically rich languages – too often conversations still happen in distinct academic and industry silos.

3.2.8. Customer Support

3.2.8.1. Do you have any suggestions for improvement of Moses customer support?

Improvement of customer support is ranked at the bottom of priorities of respondents and most of the written responses we received praised the support provided on the Moses support mailing list: *“I find that the Moses support mailing list is very good.”*, *“I’m very impressed by the prompt answers I get on the mailing list.”*, *“the mailing list works perfectly”*.

We received very few requests for professional support or even faster support response times. Some respondents would like to see improved documentation, which as mentioned in some of the sections above, could enable industrial users to adopt some of the advanced features that are already there.

Recommendation: The Moses core team should continue to participate as actively in the Moses support mailing list as it does today; documentation for some industry-relevant features could be expanded to enable industry users to address their problems (the Moses core team already started to raise awareness of some features the *“Moses Past, Present and Future”* at the MosesCore/TAUS Machine Translation Showcase in London in mid-June).

3.2.9. Moses strengths and weaknesses versus other MT solutions

3.2.9.1. What are the major strengths and weaknesses of Moses in contrast to other MT solutions?

Survey respondents universally like the availability of Moses as a free open source project and the community that supports it. The open source nature enables the customizability and flexibility that many users desire. Comments like “best performance” and “stable, reliable, with many support tools” show that Moses has reached a stage of maturity where users feel very comfortable about building MT systems with the toolkit and use it for their MT needs, be they academic or commercial, on a daily basis.

When it comes to commercial use however, there is still a gap between what users expect from a MT toolkit and what is provided, even an open source one. Comments here repeat many of the requests that we already mentioned in the sections above: *“language-specific and domain-specific support”, “complicated to set up and use. [...] no GUI for daily use”, “Inability to handle XLIFF and tags, awkward file handling; error-prone and awkward training steps”*. Surprisingly there are also complaints about the statistical nature of the system and large corpus requirements – which are issues that are inherent to all statistical MT systems.

To project whether this gap can be closed in the context of the current project organization we take a look back at the recommendations we made in the 2011 and 2012 Moses user survey reports. We examine which recommendations were implemented and which ones have not been addressed. This allows us to not only determine the strengths and weaknesses of the Moses toolkit, but also of the project organization developing and supporting it

By analyzing where the strengths of the publicly funded Moses project lie, we can give more informed recommendations on where to focus its efforts - first to use the project’s strengths, but also to focus available resources on areas that are most beneficial to the academia/industry partnership.

This report complements the article *“Does Moses have a future?”*, April 2013, by considering related questions: What has the Moses project provided to help industrial users/integrators to be successful in the marketplace? What can the project team work on in the future to best support the industry, without duplicating what industry integrators readily supply? This can also serve as a clarification for the industry on what they can and cannot expect from the public open source project.

TAUS - Enabling better translation

3.3. Implementation of the 2011 - 2012 Moses User Survey Report Recommendations

3.3.1. 2011: Four of Six Recommendations Implemented

The Moses team and the community implemented four out of six recommendations from our 2011 Moses User Survey Report – demonstrating the responsiveness of the community and the competitiveness of the open source approach in comparison to proprietary MT software.

Recommendation	Description	Full Implementation	Status
A more modern API	Exposing decoder information through a programming language API and a modern web API (see also section 3.2.2.)	No	The Moses team integrated many new technologies since 2011 and kept up exposing the additional data, however, there hasn't been much progress on a modern web API
Defined releases and regression tests	Implementation of a release management process and tested releases for various platforms; more regression/unit tests	Yes	Moses release management successfully released Moses v1.0; more tests were added to the code base
Improved parallelization and restart	The parallelization and restart was desired for the training process	Yes	Parallelization for all steps in the training process was recently enabled
Incremental training and guidance in documentation	Many users want to feed data obtained in post-editing back into the training process	Yes	Incremental training is implemented, but considered experimental
Expose decoding information to help post-editors	Expose data that enables post-editors to choose better translations	No	While word-alignments are now included by default and other data is available, there is no guidance how to use that data for post-editing
Better information resource	Better documentation, recommendations for language tools and collaborative documentation features	Yes	The documentation has been updated, and the TAUS Technology Directories cover many tools; documentation published in Wiki (without commenting)

TAUS - Enabling better translation

3.3.2. 2012: Five of Eight Recommendations Implemented

The Moses team and the community implemented five out of eight recommendations from our 2012 Moses User Survey Report. Two of them – Linux installation packages and the online tutorial - were realized through targeted funding from the EU MosesCore project that started in February 2012.

Recommendation	Description	Full Implementation	Status
Increase training speed and incremental training	Code optimization to improve training speed and incremental training for MT+post-editing scenario	Yes	Training steps parallelized and incremental training released (still considered experimental)
Cluster Simplification	Tutorial for cluster setups and simplification of cluster installation	No	No progress – most efforts focused on multi-threading
Evaluation and Productivity Measure Guidance	Guidance on academic quality metrics and development of post-editing productivity metrics	Yes (external)	External initiative TAUS Dynamic Quality Framework
Linux Installation Packages	Installation Packages for the most popular Linux distributions including related components	Yes	Packages for Redhat- and Debian-based Linux distributions were released through MosesCore; DoMY CE installation packages
Flexible integration options	Wide variety of integration scenarios requires flexible/ customizable solution (see section 3.2.2)	No	Integration needs for various scenarios unclear; lack of industry guidance/ participation
Translation of different file formats	Parsing of file formats and translation with Moses while maintaining inline markup	Yes	Implemented with M4Loc/ Okapi solution
Improved language support	Improve support for under-resourced, morphologically rich and non-European languages (see section 3.2.6)	No	While limited tools listing is available, guidance for tool use in combination with Moses is not available
Training materials	Training materials for users without a computer science/ computational linguistics background	Yes	Online TAUS Machine Translation and Moses Tutorial

TAUS - Enabling better translation

As a side note, we marked some recommendations as completed, which were implemented, but are still considered experimental, for example iterative training, parallel training and preservation of inline markup. Industry users often requested these features and we believe that only active use and testing by all users, and reporting of issues back to the Moses team can bring these features to production quality. Otherwise these features might linger and degrade.



4. Finding Common Ground

We now look for patterns in the recommendations that the Moses team and other contributors implemented.

We observed that implemented recommendations fall into two broad categories: in the first category are recommendations which lie at the core of the Moses toolkit, e.g. in the decoder, the integrated language model and the training process. This is the expertise of the Moses core team and the core contributors and they are very responsive to implementing such recommendations. The implementation of such recommendations provides benefits to virtually all users of the Moses toolkit.

The second category consists of recommendations that are publicly funded, partially publicly funded or funded by industry groups like TAUS. The implemented recommendations here are specifically funded goals of the respective initiatives. This includes the online tutorial, the installation packages, the evaluation/productivity measures and tools for the translation of different file formats.

What is missing from this picture, in our opinion, are significant contributions from industry users, specifically long-time Moses users in the translation buyer and the translation technology provider category. The two exceptions are contributions by Adobe (Adobe Moses Tools) and Precision Translation Tools (DoMY CE).

There is a lot of activity in the market of Moses-backed MT solutions and perhaps most vendors vie for an edge in the market and see solutions they have developed for some of the open issues as a competitive advantage. We believe however, that many of the issues we describe in the recommendations don't provide a competitive edge and that the wheel is invented many times over.

At the same time there seem to be oversized expectations on the part of smaller language service providers and translation agencies new to machine translation who expect Moses to be a quickly installable tool with pre-trained engines to instantly meet their machine translation needs. As any open source solution Moses has a relatively steep learning curve and requires an investment of time and resources to customize it for the intended use.

With public funding and the contributions of the academic community Moses has been set on a path to mature into a stable and long-lasting open source project for a broad user base in industry. Academic developers have contributed a tremendous amount of work towards this goal and have probably contributed all they can at this point for industry users.

TAUS - Enabling better translation

In a time of shrinking public budgets it is now up industry to decide whether to continue on this path or to leave Moses to a few expert vendors on which they will become dependent upon. We believe that a balance is possible and the Moses project can continue in small steps on the path to open source project maturity with ample opportunities for commercial vendors.

What this means practically is that the Moses project should refrain from attempting to fulfill oversized industry expectations that can and should be provided by commercial vendors. For example this means providing APIs instead of a graphical user interface, exposing information useful for post-editing in a data structure or API rather than developing a post-editing environment and maybe only providing installation packages for the core Moses binaries without 3rd party components.

Sustainability of the features and release process over the longer term is more important than developing a complete industry solution. Complete industry solutions, including convenient setup, ease of use, customization and related support can be best supplied by commercial vendors. Open source projects like Linux and PostgreSQL, successful over not just years, but decades, demonstrate that such a model can work well, even when public funding dries up.

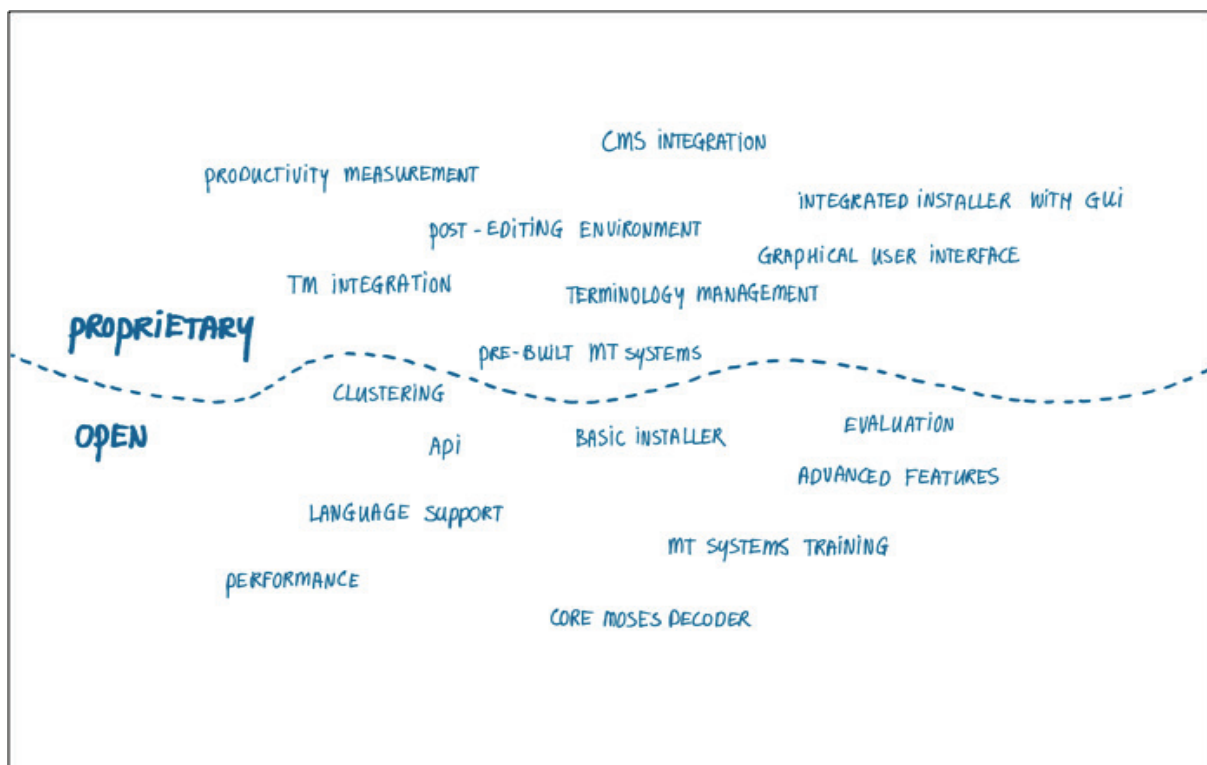


Figure 3: Open/Proprietary

TAUS - Enabling better translation

The above diagram suggests a separation between features that are best addressed by the open source project and features that should be provided by commercial vendors. However, the industry should not understand this line as an invitation to disengage from the open source project. Rather industry and academia should continuously discuss where this line should be drawn at any point in time. Having such a line provides planning certainty and a roadmap. It avoids duplication of effort and allows the most efficient use of available resources. If the language industry wants to see its needs met, it needs to become more actively involved in the process, be it through code contributions and funding, or the participation in an industry advisory board.

We noticed that, while the Moses support mailing list is great for in-depth technical questions, there is not much coverage of higher-level industry requests that we discussed in this report and the 2011/2012 reports. TAUS will initiate this discussion with a half-day industry roundtable at the MT Marathon in Prague in September 2013. We also believe it would be useful to provide a forum for continued industry consultation following the meeting for knowledge sharing and coordinated activity.

We are hopeful that following a balanced path will help the industry/academia partnership to, in the language of the Gartner Hype Cycle, inch up the *"Slope of Enlightenment"* and reach the *"Plateau of Productivity"* for the adoption of (open source) statistical machine translation.

TAUS - Enabling better translation

TAUS is an innovation think tank and platform for industry-shared services, resources and research for the translation sector globally.

We envision translation as a standard feature, a ubiquitous service. Like the internet, electricity and water, translation is one of the basic needs of human civilization.

Our mission is to increase the size and significance of the translation industry to help the world communicate better.

TAUS supports entrepreneurs and principals in the translation industry to share and define new strategies through a comprehensive range of events, publications and knowledge tools.

